Pensions, household saving, and welfare:
A dynamic analysis of crowd out

Dᴀᴠɪᴅ M. Bʟᴀᴜ
Department of Economics, The Ohio State University and IZA

This paper specifies a life cycle model of saving and employment and uses it to analyze crowd out of private household saving by public and private pensions. Some parameters of the model are estimated and others are calibrated to match life cycle employment and asset profiles, and Social Security claiming decisions. Simulation results indicate that defined benefit (DB) and defined contribution (DC) pensions on average crowd out household wealth by $0.09 and $0.37 per dollar of pension wealth, respectively, while crowd out by Social Security is $0.56. The magnitude of crowd out is sensitive to model specification, with more restrictive versions of the model (e.g., no employment decision, no bequest motive, no uncertainty) generally resulting in larger simulated crowd out. A welfare analysis implies that DB pensions and Social Security are not valued by households. The longevity insurance provided by such plans is offset by a high degree of impatience and, for Social Security, low benefits relative to taxes paid. A typical DC pension is valued at about one quarter of its expected present discounted value.

Kᴇʏᴡᴏʀᴅs. Retirement, pensions, savings.

JEL ᴄʟᴀssɪғɪᴄᴀᴛɪᴏɴ. J26.

1. Iɴᴛʀᴏᴅᴜᴄᴛɪᴏɴ

The effect of pensions on household saving is an issue of longstanding interest to economists and policymakers. Financial imbalances in public pension systems have led to substantial benefit cuts around the world, and rapid population aging makes further cuts virtually inevitable. In the United States, there has been a major shift away from employer-sponsored defined benefit (DB) pension plans toward defined contribution (DC) plans.[1] The retirement and saving incentives of these types of pensions are very

[1]DB pension plans provide employees with an annuity determined by age, length of service, and earnings history at the pension-providing firm. DB plans typically provide a strong financial incentive to remain

different. Understanding how household saving behavior is affected by cuts and structural changes in public pensions and the changing characteristics of private pensions is crucial for optimal pension design and welfare analysis.

The life cycle model of intertemporal decision-making implies that households respond to the implicit or explicit savings accumulated in their public and private pension plans by saving less in other forms. Thus pensions displace or "crowd out" household saving. A large literature uses regression analysis to estimate the magnitude of such crowd out. The studies in this literature vary along many dimensions: time period, country, type of pension analyzed, age range of households, type of data used, estimation method, and source of identification. These studies do share one common feature, however: the empirical specification is implicitly or explicitly based on a very simple stylized version of the life cycle model.

The assumptions include fixed retirement and pension claiming ages, no borrowing constraint, no bequest motive, and little or no uncertainty. The model predicts one-for-one crowd out: an additional dollar of pension wealth (the present discounted value of future benefits) causes a 1 dollar increase in consumption expenditure, spread over the remaining lifetime. The increased consumption is financed by saving 1 dollar less (Gale (1998)). The logic is straightforward: consumption is the only good in the model. There is nothing else, such as leisure, bequests, or self-insurance against risks, on which to spend the additional pension wealth. The absence of a borrowing constraint allows households to smooth consumption regardless of the timing of pension receipt. Under the assumptions of this model, one can compute a measure of pension wealth and use it as an explanatory variable in a linear regression model of household saving. Alessie, Kapteyn, and Klijn (1997), Gale (1998), and Attanasio and Rohwedder (2003) demonstrate this (Attanasio and Weber (2010), review the literature on life cycle models of savings). It has long been understood that intuition about crowd out may not hold in more realistic settings (e.g., Feldstein (1974), Gale (1998)). However, it is difficult to know how to interpret parameter estimates based a very restrictive model when the assumptions of the model do not hold.

In this paper, I specify a richer version of the life cycle model in which several of the key restrictions of the stylized model are relaxed. The model incorporates employment and pension claiming decisions, a liquidity constraint, a bequest motive, several sources of uncertainty, realistic tax treatment of pensions, and institutional constraints on pension claiming. Some of the parameters of the model are estimated using data from the Health and Retirement Study (HRS) and the Survey of Income and Program Participation (SIPP), and others are calibrated to match average life cycle employment and asset profiles and Social Security (SS) claiming decisions. The effects of public and private pensions on household saving are analyzed by solving the model numerically and simulating behavior under alternative pension scenarios. The results are used to measure

---

with the employer until reaching a benchmark age and/or years of service, and relatively little incentive to remain thereafter. DC plans specify the fraction of the employee's pretax earnings contributed to the pension account by the employee, and the rate at which the employer matches the employee contribution. The balance in the account is allocated by the employee among the investment options available in the plan. The returns and capital gains accrue to the account tax free. The funds in the account become available to the employee upon retirement from the firm.

the magnitude of crowd out by comparing simulated wealth profiles with and without pensions, all else equal.

There is no single correct way to measure crowd out, and different approaches correspond to different conceptual experiments. My approach addresses the broad question of how providing an individual with a pension with specific characteristics affects saving behavior compared to the counterfactual of no pension, in a specific institutional environment. The model is solved for the sequence of optimal saving and employment decisions and the wealth trajectory for an individual with no pension. The model is re-solved for the same individual with a pension with specified characteristics—type (DB, DC, SS), employer match rate, DB benefit formula, SS benefit formula, and so forth—allowing optimal employment and job switching behavior, and optimal claiming behavior in the case of a DC pension and SS. Crowd out is measured as the age profile of the difference in the individual's asset holdings in the two scenarios, scaled by an age-specific measure of pension wealth also derived from the model solution. In contrast, the regression approach measures the marginal effect of an additional dollar of pension wealth, holding constant earnings, employment and claiming behavior, pension contributions, and payroll taxes. There are a number of reasons why the regression measure of crowd out could differ from the measure I compute, including the fact that I simulate for a given individual while regression estimates use representative samples.[2] I do not claim that my approach is better. Rather, it provides an opportunity to explore the sensitivity of crowd out to modeling assumptions.

To compare the magnitude of crowd out implied by the model to the estimate one would derive from the usual empirical approach, I use the simulated data to estimate regression models of household wealth like those found in the literature. Comparing regression estimates on the simulated data to simulated crowd out, I can determine the sensitivity of the estimates to the strong assumptions required to rationalize the usual empirical specification.

The results indicate that crowd out of household saving by typical DB and DC pensions is −$0.09 and −$0.37, respectively, per dollar of pension wealth. Crowd out by SS is −$0.56 per dollar of SS wealth.[3] The SS crowd out measure assumes no private pension coverage. With private pension coverage, SS crowd out is smaller: −$0.25. The magnitude of simulated crowd out by SS is sensitive to model specification, moving toward −1 as the restrictive assumptions of the stylized life cycle model are imposed. Regression estimates of crowd out using the simulated data are −$0.38 for DB pensions, −$0.64 for DC pensions, and −$0.47 for SS. The regression estimates of crowd out are also sensitive to specification.

To measure the value of pensions to households, I compute the compensating variation (CV): the amount by which the initial assets of a household without a pension must

---

[2]The model is too computation-intensive to solve for a large number of individuals. This also limits the possibility of structurally estimating the model, as discussed below.

[3]Crowd out varies by age. These results measure crowd out at the last age at which no more than half of the simulated observations have claimed the pension or SS benefit. Using other reasonable benchmark ages has little impact on the results.

be increased so as to equate the household's optimized lifetime utility with and without a pension. The welfare calculations indicate that, conditional on the existence of SS, DB pensions have essentially no value to households, with a CV of −$1.7K relative to pension wealth at age 25 of $6.6K, thus reducing welfare by −$0.26 per dollar of pension wealth. DC pensions are an asset to households, but are valued at only $0.24 per dollar of DC pension wealth. For a household without an employer-provided pension, SS is also a liability: it is worth −$0.07 per dollar of Social Security wealth. The absence of value in DB and SS is a result of three factors: (1) a high calibrated discount factor, (2) the highly illiquid nature of DB and SS pension wealth, and (3) in the case of SS, the low return on taxes paid.

This paper contributes to two literatures. The first is the literature on structural estimation or calibration of models of saving behavior in the presence of pensions or SS. The most closely related paper is Samwick (2003). The motivation for his analysis is quite similar to mine: understanding crowd out behavior in the framework of a model of optimal decision-making. His analysis focuses on the choice of DC contribution rate and the effects on crowd out of alternative values of DC parameters such as the employer match rate, contribution limit, and early withdrawal penalty. His model is more stylized than mine in several respects (e.g., simplified DB and SS plans, no employment and SS claiming decisions), but has some features that I omit, such as persistent shocks. Other simulation studies of the effects of pensions on saving have accounted for some of the features modeled here, but in a more limited way, and none has incorporated all of them.[4] A unique contribution of this paper is the systematic analysis of the sensitivity of crowd out to model assumptions concerning employment, pension claiming, uncertainty, and bequests.

The second relevant literature is on empirical estimation of crowd out. The most common approach in this literature is to estimate the impact on saving of an additional dollar of pension wealth, using public pension reforms to provide identification. This requires computing a measure of pension wealth. This is typically done by applying the rules of the pension plan to compute the annuity to which an individual will be entitled, assuming specific retirement and claiming ages and a given path of earnings. The annuity is converted to a wealth measure by a standard present discounted value calculation. Implicit in the latter is the assumption of a perfect capital market and little or no uncertainty (typically only mortality risk). The wealth measure is used as an explanatory variable in a cross section or cohort regression analysis of saving or wealth. Examples include Attanasio and Brugiavini (2003), Attanasio and Rohwedder (2003), Chetty, Friedman, Leth-Petersen, Nielsen, and Olsen (2014), Hurd, Michaud, and Rohwedder (2012), and Kapteyn, Alessie, and Lusardi (2005). Empirical estimates from these and other papers are discussed in Section 4. My approach differs in that it does not use pension reform to identify crowd out, but uses the simulated data to estimate regressions

---

[4]Related papers include Scholz, Seshadri, and Khitatrakun (2006), Engen, Gale, and Uccello (1999), and Laibson, Repetto, and Tobacman (1998). These papers do not allow for choice of retirement age, but they do incorporate earnings and/or medical expenditure uncertainty, persistent shocks, and in some cases a liquidity constraint. French (2005) and van der Klaauw and Wolpin (2008) structurally estimate models of saving and retirement behavior incorporating Social Security, and, in the case of French, a stylized pension plan.

like those in the literature to explore the sensitivity of results to modeling assumptions and specification choices.

The following section of the paper describes the life cycle model used in the analysis and illustrates the implications of the restrictions implied in the more stylized version. Section 3 describes the data, parameter calibrations, and initial conditions. The simulation results are discussed and extensively explored in Section 4. Section 5 describes an application of the model to simulating the effects of SS reform, and conclusions are offered in Section 6. Additional material is available in supplementary files on the journal website, http://qeconomics.org/supp/349/supplement.pdf and http://qeconomics.org/supp/349/code_and_data.zip.

2. A life cycle model of saving, employment, and pension claiming

2.1 *Model description*

The model developed here characterizes the behavior of a married man from entry to the labor force at age 25 to the end of the life cycle.[5] The individual makes a categorical employment choice j_t and a continuous consumption choice c_t at each age t, where consumption is defined as net of out-of-pocket medical expenditure. The employment choice set is (0) nonemployment, (1) a job with a new employer, and (2) the job with the period $t-1$ employer (age and period are used interchangeably). The latter alternative is available only if 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 provide pension coverage. Allowing job switching is important because pensions are usually employer-specific, and it is often possible to leave the pension-providing employer, collect a pension benefit, and work for another employer.[6] The length of a period is 1 year. The last age to which the individual can survive is 100, and to ease the computational burden, the last age at which employment is an option is 85.

If the individual is old enough to be eligible for a Social Security benefit (Old Age and Survivors Insurance; abbreviated here as SS) and has not yet claimed the benefit, he makes a claiming decision in period t. Employment and 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 as a lump sum and transfer it into household savings, where it is taxable.[7] There is no separate claiming decision for DB pensions: age and years enrolled at the time of exit from the pension-providing job

---

[5]Most men are married. Behavior of the wife is not modeled. For simplicity the husband and wife are assumed to have the same age. Solving for the behavior of both spouses is conceptually straightforward but very computation-intensive. See van der Klaauw and Wolpin (2008) for an example, in a context with no private pensions.

[6]Modeling pensions on new jobs is difficult because it greatly increases the size of the state space. Job offers would have to be characterized by whether a pension is offered, and the type and characteristics of the plan.

[7]There is a tax penalty of 10% of the account balance if the pension is claimed before age 59 1/2 (60 in the model). The balance must be claimed no later than age 70, consistent with legal requirements, even if the

fully determine whether and when the individual will receive a benefit and the benefit amount.

The logarithm of the hourly wage offered by the current employer is quadratic in age and subject to an independent and identically distributed (i.i.d.) normal shock.[8] The log wage process for new job offers is also quadratic in age and is subject to a different shock. The wife's earnings enter the model as an exogenous stochastic process, described by two equations: one governing whether she works and the other her earnings. Both are quadratic in age and subject to i.i.d. normal shocks. The logarithm of total family out-of-pocket medical expenditure, $\ln(m_t)$, is a linear function of age and an i.i.d. normal shock. The probability of death in period t, conditional on survival to t, is denoted π_t and is assumed to correspond to standard life table mortality rates.[9] The probability that an employed individual is laid off at the end of period t, λ_t, is a function of age.

An individual is eligible to claim his SS benefit if he has reached the early retirement age (62).[10] The benefit, s_t, is a real annuity determined by the function $s_t = s(\text{AIME}_{\text{fe}}, \text{fe}, E_t, t)$, where AIME_{fe} is average indexed monthly earnings, fe is the age at which the individual claims the benefit (first entitles), and E_t is annual earnings in period t. Age and current earnings matter because there is an earnings test at some ages.[11]

The nominal annuity provided by a DB pension plan, b_t, depends on age, years of enrollment in the plan, and earnings history at the date of exit from the pension-providing firm. The formula can be written in general as $b_t = b(E_p, a_e, y_e, t)$, where E_p is a summary statistic for the worker's earnings history at the pension-providing firm at the time of exit (e.g., average earnings in the last 5 years of employment at the firm), and a_e and y_e are age and years of enrollment in the plan at the time of exit. The DB benefit is nominal, so its real value at age t depends on the inflation rate and years since the benefit

---

individual remains employed at the pension-providing firm. There is no installment payment 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-providing 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.

[8]For computational reasons described below, the earnings process is stationary. Similarly, allowing work experience and job tenure to affect wages is infeasible, as the state space becomes unmanageably large.

[9]Health is not included in the model, and the possibility of becoming disabled and enrolling in the Social Security Disability Insurance program is also ignored. A previous version of the model incorporated these features, but they made little difference to the outcomes, so I dropped them so as to focus on the elements of the model that are crucial for pensions. The model excludes health insurance for similar reasons. Finally, I assume that the wife dies at the same time as the husband to avoid adding additional state variables.

[10]I assume that the individual accumulates the minimum required years of covered employment (10) by the time he is 62, eliminating the need to make work experience a state variable. This is never binding in practice.

[11]AIME is the average of the highest 35 years of wage-indexed SS-covered earnings. Covered earnings are capped at a real value of $55,500 in 1992 dollars. AIME is recalculated each period until the individual claims the benefit. The 1992 SS rules are used in the analysis, with one exception: the elimination of the earnings test for beneficiaries who have reached full retirement age, which took place in 2000, is assumed to apply throughout the analysis. See Social Security Administration (2010) for description of the benefit determination rules. In principle, the wife's earnings should enter the formula as well, but for simplicity I assume that the wife's spousal benefit exceeds her retired worker benefit, in which case her earnings do not matter, and the total household benefit is 1.5 times the husband's benefit.

began. There is no risk of default by the pension plan. It is computationally infeasible to allow both E_p and $AIME_{fe}$ to be state variables, so as in French (2005) I assume that $E_p = AIME_p$, the value of AIME at the time of exit from the pension-providing firm. The DB pension benefit cannot be received until the individual leaves the pension job. The benefit formula depends on the specific pension plan in which the worker is enrolled. I describe the pension data in the following section.

DC pension plans are characterized by 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 pre-tax earnings to the pension account. These fractions are taken as given and fixed in the model. If the individual remains with the pension-providing firm at the beginning of period $t + 1$ and has not reached age 70, the account balance is given by $D_{t+1} = (D_t + E_t(\text{wc} + \text{ec}))(1 + r_{t+1})$, where wc and ec are the worker and employer contribution rates, respectively, and r_{t+1} is the rate of return on assets held at the beginning of period $t + 1$. If the individual has left the firm but has not yet claimed the account balance, then $D_{t+1} = D_t(1 + r_{t+1})$. Borrowing from a DC pension account is not allowed.

The rate of return earned on assets held at the end of period t, A_t^*, is realized at the beginning of period $t + 1$. The rate of return is determined by a stationary stochastic process specified as $1 + r_{t+1} = (1 + r_m) \exp\{\theta_{t+1}\}$, where r_m is the mean rate of return and θ_t is an idiosyncratic individual-specific shock, distributed i.i.d. normal. Returns are defined to include capital gains, so $r_{t+1} < 0$ corresponds to a capital loss. The rate of return is assumed to be the same for the DC pension account and household wealth.

The law of motion for assets held outside the DC account is $A_{t+1} = A_t^*(1 + r_{t+1})$. There is assumed to be a borrowing constraint ($A_t \geq 0$ $\forall t$) and a consumption floor, $c_F > 0$. The consumption floor is a simple approximation to income-and-asset-tested government programs such as Supplemental Security Income, Food Stamps, and Medicaid, that allow individuals with no other sources of income to survive (Hubbard, Skinner, and Zeldes (1995), van der Klaauw and Wolpin (2008)). If cash on hand is less than c_F, the government provides a grant sufficient to allow consumption of c_F.[12]

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

$$I_t = E_t + P_t E_{wt} + b_t + s_t - m_t - \tau(E_t, P_t E_{wt}, b_t, s_t, m_t, A_t, t, \text{wc}),$$

where $P_t = 1$ if the wife works and 0 otherwise, E_{wt} is the wife's earnings offer, and τ includes federal income and payroll taxes, calculated using the rules in effect for 1992, and assuming the household takes the standard deduction. The tax computation accounts for the tax-sheltered nature of the worker's contribution to the DC account, and for the rules governing taxation of SS 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).

---

[12]A large medical expenditure shock could cause end-of-period assets to be negative. In this case, the debt is forgiven before the beginning of the next period. This is intended to roughly mimic the Medicaid program, which provides means-tested health insurance for the poor.

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,[13]

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

where $W = 1$ if employed and 0 otherwise, $NJ = 1$ if a new job is chosen and 0 otherwise, and ε_{ct} and $\varepsilon_{\ell t}$ are i.i.d. normal shocks to the utility from consumption and employment, respectively.[14] Parameter α is the coefficient of relative risk aversion, γ_1 is the disutility of employment, γ_2 is the additional disutility of employment if the individual was not employed in the previous period, and γ_3 is the additional disutility of employment in the period in which the individual changes jobs. The dynamic features of the utility function are important so as to avoid excessive churning in employment choices in response to transitory earnings and preference shocks.

The bequest function takes the form

$$B_t = \beta_0\big(1 - \exp\{-\beta_1 A_t\}\big),$$

where B_t is the utility received in period t from leaving a bequest of A_t in the event of death at the beginning of $t + 1$, and β_0 and β_1 are nonnegative parameters.

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

This model is very flexible with respect to pension crowd out. There are combinations of preferences and constraints that could result in very small crowd out or even "crowd in." For example, a strong preference for leisure could induce both early retirement and a high rate of private saving to finance consumption during retirement until eligibility for a pension benefit, as noted by Feldstein (1974). The high rate of saving would occur during the same part of the life cycle when the implicit pension value is growing rapidly, resulting in a positive relationship between pension wealth and saving. High risk aversion could drastically limit the extent to which individuals are willing to substitute an illiquid pension for household saving, yielding very small crowd out. Patient individuals will tend to accumulate a lot of saving for retirement, and are unlikely to face either a binding liquidity constraint or a significant risk of hitting the consumption floor. In this case, one form of saving for retirement might be a very good substitute for another, resulting in one-for-one crowd out.

The model has limitations as well. Pensions are complicated, and a number of potentially important features have been omitted. These include the choice over asset allocation and the form of benefit (lump sum, installments, or an annuity),[15] borrowing

---

[13]I experimented with a nonseparable utility function, but it proved unnecessary for fitting the employment and asset profiles in the data, so I opted for the more parsimonious specification.

[14]I assume that all of the disturbances in the model are independent. The timing convention is that shocks are realized at the beginning of the period, before choices are made.

[15]The model does not allow the option of purchasing an annuity in the insurance market. This is in principle a significant limitation, but in practice the private annuity market is not widely used. For example,

against the DC pension balance, the risk of bankruptcy in DB pensions, the enrollment decision, and the possibility of pension coverage on a new job. Some features of the environment that are relevant to saving and retirement behavior have been simplified considerably or omitted, so as to focus on pensions. These include health, disability and health insurance, time inconsistency, and other nonstandard aspects of preferences and decision-making. Persistence in unobservables is likely to be an important source of variation in saving behavior, so the absence of persistence in the wage and medical expenditure innovation processes makes the model poorly suited to explain heterogeneity in life cycle wealth profiles. As a result, the analysis focuses on the central tendency of the asset distribution. The results should be interpreted with these limitations in mind.

2.2 *Restrictions imposed in empirical analysis*

The stylized version of the life cycle model that is the basis for empirical analysis of pension crowd out is a special case of this model, with no uncertainty, no liquidity constraint, no bequest motive, no consumption floor, and no employment and claiming choices. This restricted version of the model can be solved analytically for household wealth or the saving rate, and implies a regression specification for nonpension wealth of the form

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

where

$X_{1t} = (1 - \kappa_t)(A^*_{t-1} + D^*_{t-1}) =$ adjusted previous-period wealth plus the DC balance,

$\kappa_t \in [0, 1)$ is an adjustment factor for remaining length of life ($d\kappa_t/dt > 0$, $\kappa_T = 1$),

$X_{2t} = (1 - \kappa_t)E_t(1 - wc - \tau)$ is adjusted net annual earnings, assuming a flat tax rate of τ,

$X_{3t} = \kappa_t \, \text{PDV}_t(E)$ is the adjusted present discounted value (PDV) of future earnings,

$X_{4t} = \kappa_t \, \text{PDV}_t(b)$ is the adjusted PDV of future DB pension benefits,

$X_{5t} = \kappa_t \, \text{PDV}_t(s)$ is the adjusted PDV of future SS benefits,

$X_{6t} = D_t$ is the DC balance in period t.

The restrictions used to derive this specification deliver strong predictions: $\beta_1 = \beta_2 = 1$ and $\beta_3 = \beta_4 = \beta_5 = \beta_6 = -1$. The parameters β_4, β_5, and β_6 measure pension crowd out: the amount by which household wealth is reduced as a result of a 1 dollar increase in pension wealth. The virtue of this framework is its empirical tractability: under the

---

Inkmann, Lopes, and Michaelides (2011) report that fewer than 6% of retired individuals in the English Longitudinal Study of Ageing voluntarily purchase annuities.

assumptions of the model, the right hand side variables can be computed and the regression can be estimated with suitable data. I estimate regressions of this form using data simulated from the unrestricted model.

2.3 *Model solution*

The model is formulated as a dynamic program and solved numerically by backward recursion on the value function, using Monte Carlo integration to approximate integrals. There are as many as three continuous state variables in a given period, depending on age and pension coverage: household assets, the DC account balance, and AIME. There are up to five discrete state variables depending on age: employment status in the previous period, the age at which the individual first entitles to SS (fe, set equal to zero if not yet entitled), the number of years worked after claiming SS and before the full retirement age (FRA) (to account for benefit readjustment upon reaching the FRA for benefits lost to the earnings test), years enrolled in the DB pension plan, and a categorical indicator of current pension status: whether covered, which type of plan, and whether claimed. The model is solved for grids of values of the continuous state variables along with all feasible values of the discrete state variables at a given age. Multidimensional local linear interpolation is used to approximate the expected value at t of the period $t + 1$ value function for each selected point in the period t state space and each alternative in the period t choice set. The pension wealth measures are also approximated by interpolation. Appendix A in the Supplement provides further details on the solution method. For a given set of discrete choices, optimal consumption is found by grid search.

3. Data, pensions, calibration, and simulation

3.1 *Data*

Data from the HRS and SIPP are used to estimate some parameters of the model, and to compute age profiles of employment, assets, and SS claiming to which the model simulations can be compared. The HRS also provides benefit formulas for a sample of DB pension plans, 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 from birth cohorts 1931–1941, and their spouses. Additional cohorts have been added periodically, and I use individuals from all of the recent cohorts, including birth years 1921–1953. Data from the SIPP are used to construct profiles for ages 25–50. The 2004 SIPP panel (covering the years 2003–2008) is used to measure employment and earnings, and the 1996, 2001, and 2004 panels (covering 1996–2008) are used to measure assets and medical expenditure. The data sources are combined, and various subsamples are used to estimate the parameters of the earnings, medical expenditure, and layoff functions, and to measure age profiles of the key outcomes.[16]

---

[16]The SIPP data are used for ages 25–50, and the HRS for ages 51 and above. The age profiles from the HRS and SIPP for all variables are very close to each other at the seam age, with one exception: the job switch rate. As shown in Figure A1, there is a large drop in the job switch rate at age 50, which is the seam age between the SIPP and HRS data. It is difficult to determine which is more accurate. For purposes of

I use a narrowly defined population so as to make the sample as homogeneous as possible, while also maintaining a reasonable sample size (as noted above, the model solution is very computation-intensive, precluding the possibility of solving for a heterogeneous sample). The population is white, non-Hispanic married men who graduated from high school and may have attended college but did not receive a 4 year college degree. Furthermore, I only use cases that do *not* have pension coverage (but are covered by SS). The model is calibrated using only the no-pension cases, and the calibrated model is then used to simulate the impact of pensions. The reason for this approach is that the distribution of worker characteristics differs substantially by pension coverage. Simultaneously attempting to fit profiles for no-pension, DB, and DC subsamples is much more difficult, and runs the risk of calibrating to differences across these groups that result from unobserved heterogeneity.[17] As noted above, my approach is not well suited to explain heterogeneity in wealth profiles or features of the wealth distribution other than the central tendency.

An individual is defined as employed in a given year if he worked at least 6 months of the year, regardless of hours worked per week. This annual measure of employment was constructed from monthly employment histories available in both the HRS and SIPP. Assets are measured by total net worth, including home equity. I use imputed and cleaned measures from the RAND version of the HRS and from the constructed variable reported in SIPP. Other variables include the age at which the SS retirement benefit was claimed (from the HRS) and earnings, layoffs, and out-of-pocket household medical expenditure (from both sources).

3.2 *Pensions*

The HRS asks respondents who report being enrolled in a DB pension plan a substantial battery of questions about the plan, including the ages of early and normal retirement, expected benefits if the respondent were to leave 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 (Chan and Stevens (2008), Engelhardt and Kumar (2011)), but the information is not sufficient to calculate benefits for all possible retirement ages and earnings realizations. The HRS also obtained Summary Plan Descriptions and other relevant information about the pension for a subsample of HRS respondents. The benefit formulas and other plan features derived from these documents were coded by HRS staff and

---

not generating excessive job turnover in the DB simulations, I decided to target the much lower HRS job switching rate. The Panel Study of Income Dynamics (PSID) provides data for the entire age range, but the sample sizes are much smaller at each age. I compared asset data from the PSID to the data from SIPP and HRS and found a very close correspondence by age, but the PSID profiles are noisier.

[17]Pension holders are better educated, have higher earnings, are more likely to belong to a union and work in the public sector, are in better health, have greater attachment to the labor force, and are much more likely to have employer-provided health insurance compared to workers without pensions. See Blau (2008) and Gustman, Steinmeier, and Tabatabai (2010) for extensive description of pensions in the HRS. An alternative approach is to calibrate the model to the behavior of DB or DC pension holders, and simulate the impact of removing the pension. I tried this approach but found it difficult to fit the data on pension holders, perhaps because pensions are so heterogeneous and I can solve the model using only one plan at a time.

made available in a data base, along with pension calculator software. The data and software allow one to compute the benefit to which an individual would be entitled for any combination of age, years of enrollment, and salary profile. Unfortunately, it was not feasible to integrate the pension calculator software with the model solution software. Instead, I developed a very flexible plan-specific regression approximation of the benefit formulas for use in solution and simulation of the model. The approach is described in Appendix B in the Supplement. The pension data base also contains DC plans, but the only relevant DC plan characteristics in the model are contribution rates and the initial balance. These are described below.

3.3 *Calibration*

The model was calibrated in three steps. First, the HRS and SIPP data were used to estimate the parameters of the wage, medical expenditure, wife's earnings, and layoff functions.[18] The estimates are shown in the Supplement in Tables A1 and A2, and the estimated variances used in solution of the model are reported in the top panel of Table 1. Supplement Tables A3 and A4 display summary statistics for the wage, wife earnings, and medical expenditure from model simulations based on these estimates.

TABLE 1. Key parameters used in calibration.

Parameter	Value
Derived from estimates using HRS and SIPP data	
Variance of log out-of-pocket medical expenditure shock	0.826
Variance of log husband wage offer shock	0.043
Variance of log wife earnings offer shock	0.340
Annual layoff rate	0.006
Fixed arbitrarily	
Consumption floor (c_F) (thousands of 1992 dollars)	5
Mean real rate of return (r_m)	0.03
Variance of log rate of return shock	0.0002
Annual inflation rate	0.02
Calibrated to match employment and asset patterns	
Disutility of employment (γ_1)	−0.0001
Additional disutility of employment if not employed in $t-1$ (γ_2)	−0.18
Additional disutility of employment if in a different job in $t-1$ (γ_3)	−0.010
Variance of disutility of employment shock ($\sigma^2_{\varepsilon_\ell}$)	0.0016
Variance of log utility of consumption shock ($\sigma^2_{\varepsilon_c}$)	0.03
Bequest parameters (β_0, β_1)	1.60, 0.01
Rate of time preference (δ)	0.08
Coefficient of relative risk aversion (α)	2.5

Note: See Tables A1 and A2 in the Supplement for the full set of regression parameter estimates.

---

[18]The log wage, log wife earnings, and log out-of-pocket husband+wife medical expenditure functions are specified as linear fixed effect models. I use the fixed effects estimates to compute variances of the transitory shocks, which are used in the model solution. The layoff and wife's labor force participation models are estimated by probit. The notes to Table A1 provide details on the estimates. Selectivity correction terms estimated from a probit are included in the wage and wife's earnings equations.

Second, the values of several other parameters were set arbitrarily. These include the consumption floor ($c_F = 5$K) and the mean rate of return ($r_m = 0.03$). The second panel of Table 1 shows these and the values of other arbitrarily chosen parameters. All monetary amounts are expressed in thousands of 1992 dollars (multiply by 1.69 to convert to 2014 dollars).

The remaining parameters were calibrated so as to generate simulated life cycle employment, asset, and SS claiming patterns that match the patterns observed in the data for the subsample of men without pension coverage. The coefficient of relative risk aversion (CRRA) (α), rate of time preference (δ), disutility of employment parameters (γ), variances of the preference shocks, and the bequest parameters (β) were chosen in this manner. The resulting parameter values are shown in the last panel of Table 1.[19] It is notable that simultaneously fitting employment, asset, and SS claiming patterns requires a relatively high discount factor (1.08), a relative large bequest motive, and a small disutility of employment. The high discount factor is needed to explain the prevalence of early Social Security claiming, while the large bequest motive is needed to explain why assets fail to decline much after retirement. The CRRA is calibrated to 2.5, a fairly typical value. Smaller or larger values provided a significantly worse fit.

3.4 *Simulation*

Simulations of the model compare four pension scenarios: (1) no pension, (2) a DB pension, (3) a DC pension, with the household eligible for SS in all three cases, and (4) no SS (and no payroll tax) and no pension. The calibration and simulations use a randomly chosen DB pension plan from the HRS pension provider data base.[20] The DC plans are characterized by the employer and employee contribution rates. These are set to 0.06 for the worker and 0.09 for the firm.[21] I simulate 1500 cases from age 25 to 100 for each scenario, and average the results across simulations. Individuals face mortality risk in the simulations, but for ease of interpretation I assume that no deaths actually occur. The model is solved and simulated for a married white man with a high school diploma but no Bachelors degree, born in 1937, age 25 in the first period, employed at the beginning

---

[19]The calibration targets are the age profiles shown below in Figures 1–3 and A1–A3. The calibration method was informal and involved trying large numbers of parameter combinations until reasonable fits were obtained. I used this approach because a method-of-simulated-moments approach failed to converge regardless of starting values. This may be a result of an underidentified model. The calibration exercise was not intended to explain heterogeneity in saving, so this does not limit the usefulness of the analysis for understanding the determinants of crowd out.

[20]Each DB pension plan has a unique benefit formula, so the model must be solved for each plan used in the analysis. I attempted to develop a value function approximation approach similar to the one used by van der Klaauw and Wolpin (2008) that would be flexible enough to encompass the formulas of a wide variety of plans. The formulas are so heterogeneous that this proved impossible. Results for other randomly selected DB plans are discussed below.

[21]A 6% worker contribution rate is the modal value reported in the National Compensation Survey (NCS; US Bureau of Labor Statistics (2005, Tables 90, 94)). A 100% employer match rate is the modal value reported in the NCS. I use a 150% match rate because a lower rate resulted in an excessively high job-changing rate in some specifications. However, in the final specification the simulated DC crowd out was very similar with a 100% employer match.

of the first period, and in the DB and DC pension simulations, enrolled in the plan at age 25. Other initial conditions are assets of $5000 (the sample median at age 25 in the SIPP in the no-pension case), DC balance $= 0$, and AIME $= \$1000$.

4. Results

4.1 *Model fit*

Figures 1–3 illustrate the fit of the model to employment, asset, and SS claiming profiles in the data for men with no pension. The model provides an exceptionally good fit to the employment profile up to age 70, with the sharp decline in employment beginning around age 60 well predicted by the model, without including age in the utility function.[22] The simulated asset profile in Figure 2 closely matches the median asset profile in the data until about age 80. The simulated profile declines less steeply than the actual profile after age 80, but the data are noisy. Figure 3 shows that SS claiming in the data is highly concentrated at age 62, with almost 50% claiming at the earliest possible age and 60% by age 63. The model predicts that about 20% claim at age 62 and another 30% at age 63. The model matches the claiming profile in the data reasonably well at all ages except 62. The failure of the model to capture the sharp spike in claiming at age 62 suggests that the model may inappropriately treat all assets as liquid, while in reality many household assets are tied up in houses, cars, and other relatively illiquid forms. Illiquidity is often suggested as a reason for the spike in claiming at age 62 (e.g., Coile, Diamond, Gruber, and Jousten (2002)). The relatively high calibrated discount rate (8%)
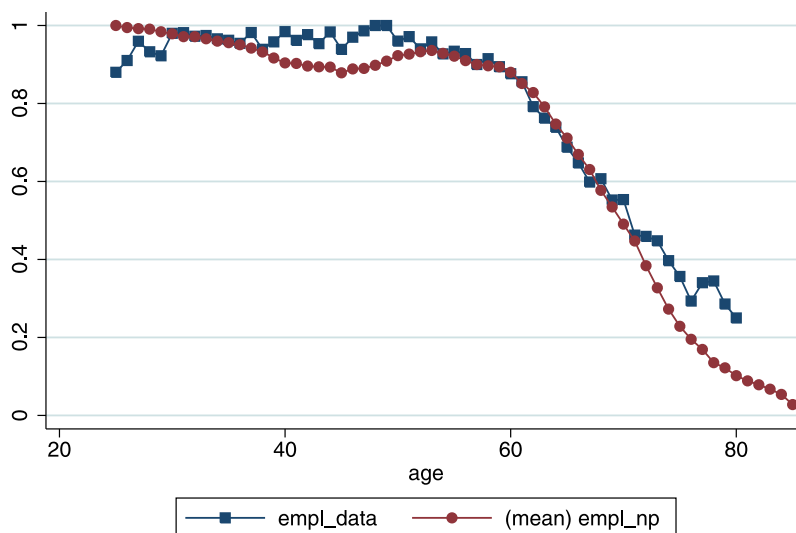

Figure 1. Actual and simulated employment rate.

---

[22]I do not discuss employment dynamics here because they are not central to the issue of crowd out. Figures A1–A3 in the Supplement show the fit of the model to employment dynamics.
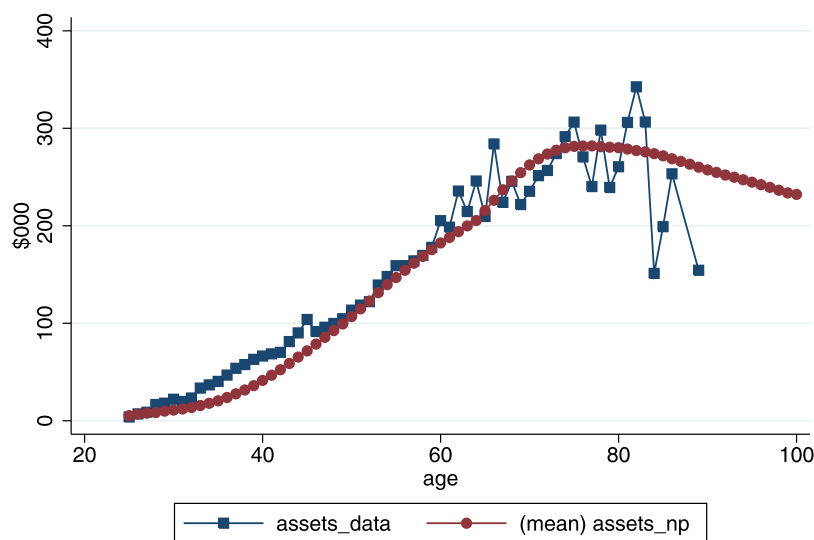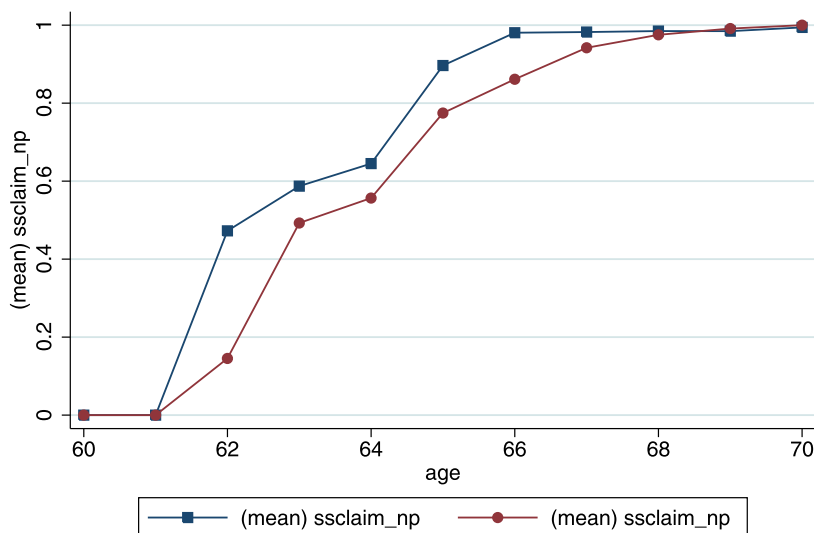
FIGURE 2. Actual and simulated assets.

FIGURE 3. Actual and simulated Social Security claiming age.

is needed to get even this close to the age 62 SS claiming spike. A still higher discount rate enables the model to fit the claiming profile quite well, but the fit to employment and asset profiles deteriorates.[23]

[23]The fit of the model to employment, asset, and claiming data for DB and DC pension holders, using the same population as for the nonpension case (white, married, high school graduate, not a college graduate) is shown in the Supplement. In general, the model predicts employment patterns quite well. Assets are underpredicted for DB holders by 10–20K and by a much larger amount for DC holders. As noted above, this likely reflects heterogeneity, despite a relatively narrowly defined population. The fit for Social Security claiming in the DB case is about the same as for NP case, but is underpredicted for DC.

4.2 *Simulated effects of pensions and Social Security*

Figures 4–7 show simulated life cycle patterns of employment, consumption, pension claiming, and SS claiming for the four scenarios of interest: no pension (NP), DB pension, and DC pension, all with SS, and no-pension–no-SS (NPNSS). Employment declines more rapidly in the DB and DC scenarios in Figure 4 than in the NP case, consistent with the early retirement incentives of DB plans (e.g., Stock and Wise (1990)) and the wealth effects of DC plans. The employment decline starts at about the same age in
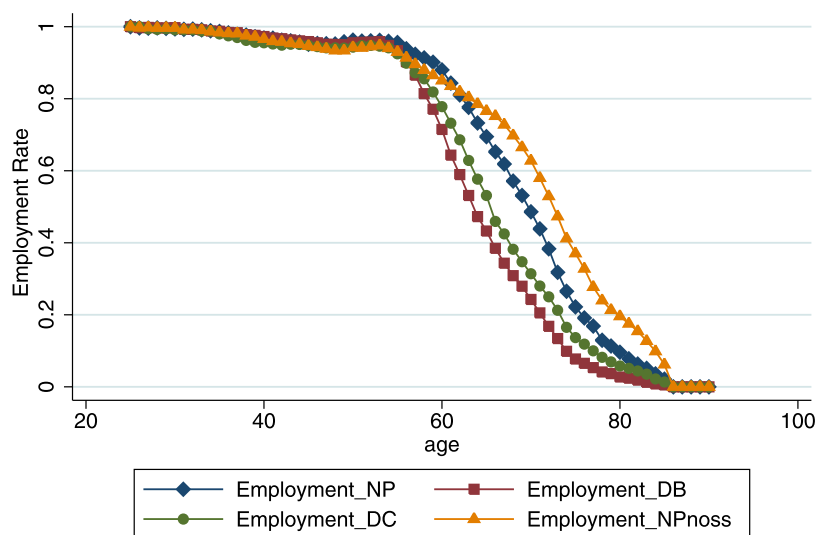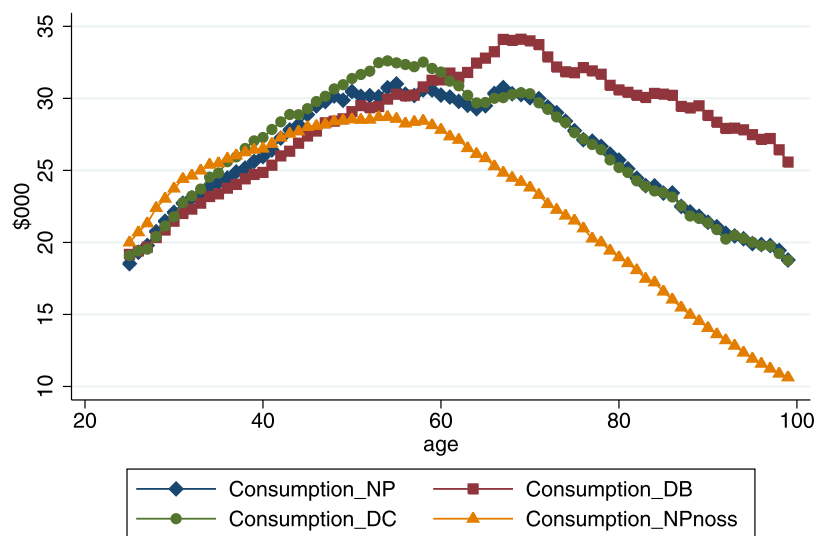

FIGURE 4. Simulated employment profiles.

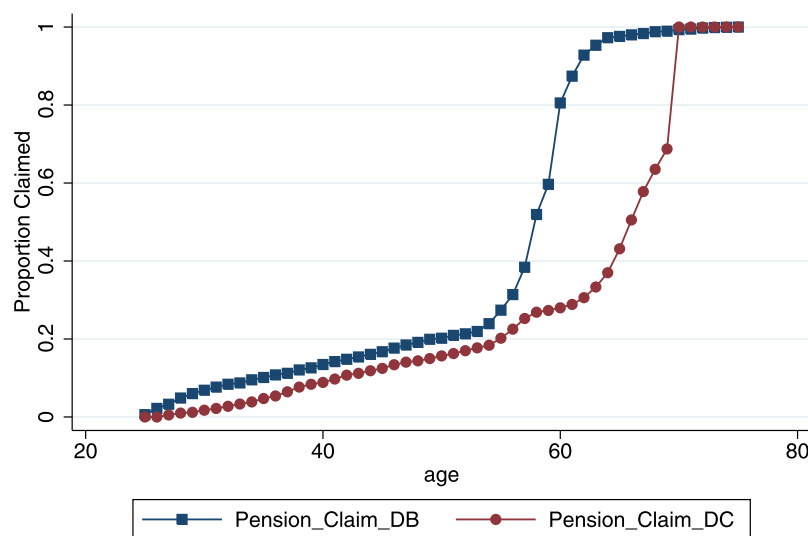FIGURE 5. Mean simulated consumption profiles.

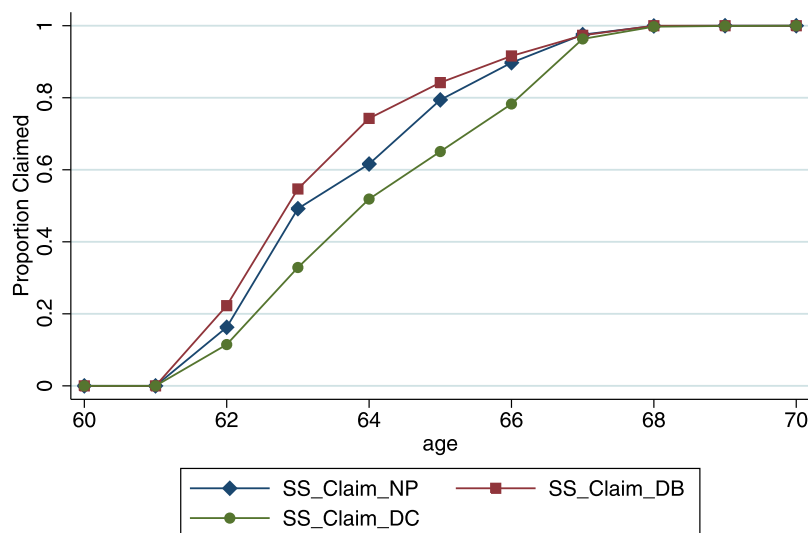FIGURE 6. Simulated pension claiming profiles.

FIGURE 7. Simulated social security claiming profiles.

the NPNSS and NP scenarios, but employment remains higher in the NPNSS case. This is likely due to both wealth and substitution effects.[24] The simulated consumption profiles illustrated in Figure 5 show that consumption is much lower in the NPNSS scenario than in the scenarios with SS. Consumption in the DB pension scenario is significantly higher

[24]A marginal cut in benefits would induce a pure wealth effect, which is different from the scenario simulated here. See Section 5 for analysis of this case. The empirical literature shows that SS has complicated employment effects, increasing employment at some ages, and reducing it at other ages. See Behaghel and Blau (2012), French (2005), Mastrobuoni (2009), and van der Klaauw and Wolpin (2008) for recent evidence.

at ages 60+ than in the NP case. This is not too surprising since DB pension recipients receive two annuities during retirement.

Figure 6 shows that DB benefit claiming is concentrated at ages 54–59, consistent with the abrupt changes in retirement incentives in DB plans at benchmark ages.[25] DC claiming is more gradual. The tax-sheltered nature of the DC account provides an incentive to do as much saving as possible in this account, and the results show that about one fifth of simulated cases delay claiming until the mandatory age of 70. Figure 7 shows that SS claiming occurs earlier in the DB scenario than in the NP scenario, consistent with earlier labor force exit by DB pension holders. SS claiming occurs later for DC pension holders. This may be a result of the relatively liquid nature of the DC account, especially after age 59, when the tax penalty ends. Access to the DC account balance may enable a household with a DC plan to take advantage of the relatively high rate of SS benefit increase resulting from delaying claiming.

4.3 *Crowd out*

Figures 8–12 illustrate the crowd out patterns implied by the simulations. The wealth trajectories shown in Figure 8 are similar in shape for the DB, DC, and NP cases, but the trajectory for the NPNSS case is much higher and has a more typical hump shape. Figure 9 displays the life cycle crowd out pattern in the DB scenario, measured by the difference between the DB and NP asset profiles in Figure 8. Note that all households begin with the same initial wealth, so crowd out is zero by construction at age 25. Figure 9 shows that assets in the DB case decline very gradually relative to the NP case, with the difference reaching about −20K in the late 50s. Figure 9 also shows the relevant pension wealth measure: the EPDV of future DB benefits, computed as part of the model solution (accounting for employment and claiming choices, the liquidity constraint, etc.). DB pension wealth rises from about 10K at age 25 to more than 150K at age 61. Figure 10 shows the proportional DB crowd out profile, calculated as the dollar magnitude of crowd out divided by DB pension wealth. Proportional DB crowd out never exceeds −0.15.[26] Figure 11 presents the crowd out profile for the DC scenario, along with the DC account balance. Crowd out by DC pensions reaches −75K in the early 60s, equal to about one third of the DC balance (see Figure 10). Figure 12 displays SS crowd out, measured by the difference between the NPNSS and NP asset profiles in

---

[25]Quitting the DB job is rare from ages 33 to 52: more than 99% of job exits at these ages are due to layoffs. There are some quits at younger ages, when the EPDV of DB benefits is low, but past a certain point individuals are effectively locked in to their DB job. Employment on the DB job declines rapidly after age 53. Excluding cases laid off from the pension job, employment on the DB job is 90% at age 54, 87% at 55, 82% at 56, and 74% at 57. Employment on the DB job drops to 56% at 58, 49% at 59, and 24% at 60. Comparing these figures to the DB employment rate in Figure 4 indicates that many workers covered by a DB plan switch jobs so as to collect the pension while continuing to work. For example, at age 60, 80% are employed while only 24% remain on the DB job. In the data this is less common, but not rare: 21% of individuals who were covered by a DB pension in the first wave of the HRS and receiving a benefit from the pension at a subsequent wave were employed (at a different job) while receiving the benefit.

[26]I use the convention of reporting crowd out as a negative number so as to be consistent with the empirical literature, which reports coefficient estimates from a regression of assets on pension wealth.
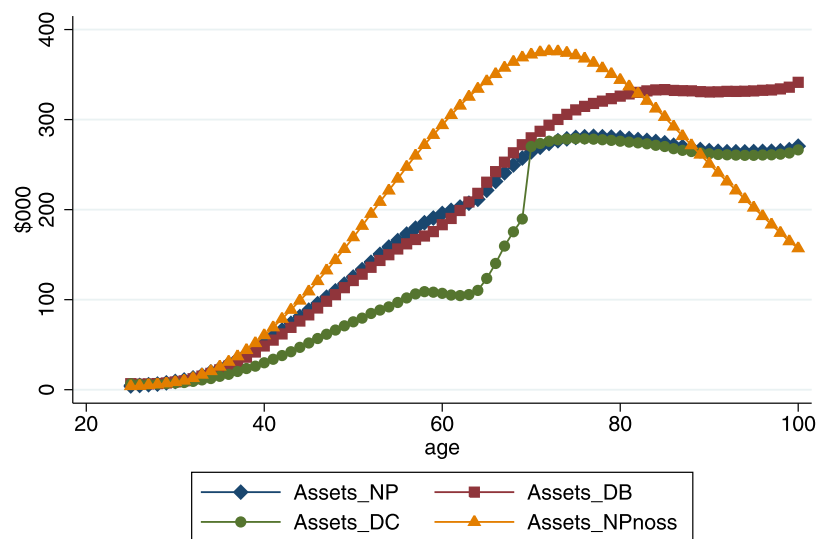
FIGURE 8. Mean simulated asset profiles.
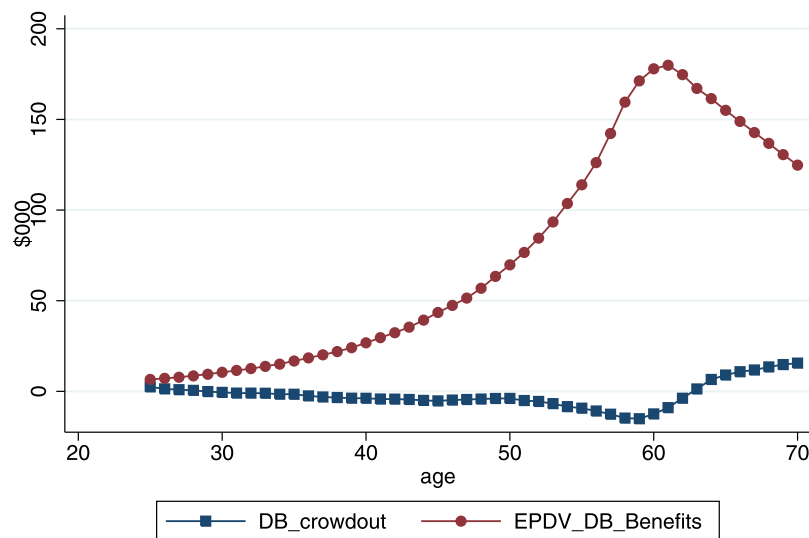
FIGURE 9. Defined benefit crowdout and EPDV of DB benefits.

Figure 8. Crowd out by SS is very large, reaching −150K at age 64, or −0.64 as a proportion of SS wealth (gross of taxes). The life cycle pattern of crowd out is increasing until around age 60 and then is mildly declining. The small amount of empirical evidence on the life cycle pattern of crowd out shows a somewhat similar pattern (see footnote 28 below).

The graphs are useful for illustrating crowd out patterns, and the age variation shown in the figures makes it clear that there is no single correct age at which to measure crowd out. Nevertheless, it is useful to summarize the magnitude of crowd out with a
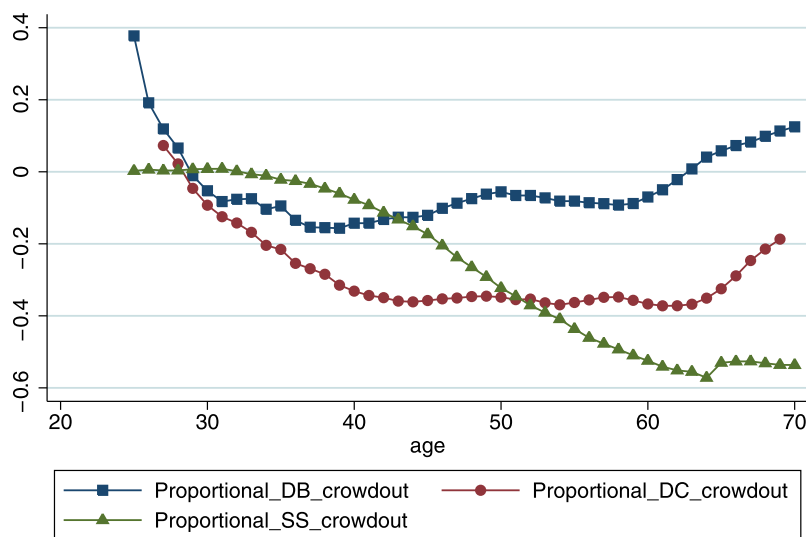
FIGURE 10. Proportional DB, DC, and SS crowdout.
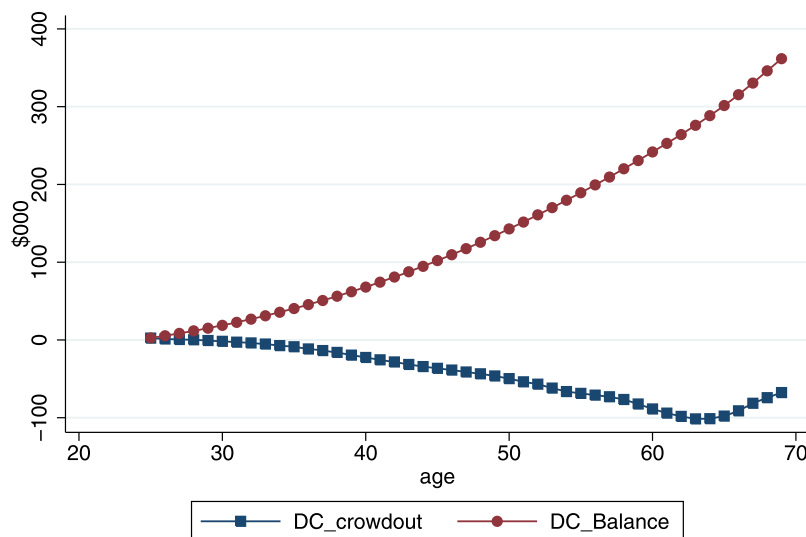
FIGURE 11. Defined contribution crowdout and DC balance.

single number, as in much of the empirical literature. I use an age at which pensions and SS have not yet been claimed by most individuals, both because this is the typical approach in the literature and because crowd out behavior can be quite different in the asset decumulation phase. I arbitrarily measure crowd out at two specific ages: the last age at which no more than 25% and 50% of simulated cases have claimed the pension or SS benefit. Table 2 shows the results. The upper panel indicates that the last age at which at least 75% of simulated individuals remain on the DB job is 54, and at age 54, DB crowd out is −8K, or −0.08 as a proportion of DB pension wealth at age 54.

FIGURE 12. Social security crowdout and EPDV of SS benefits.

TABLE 2. Simulated pension crowd out computed directly from simulations.

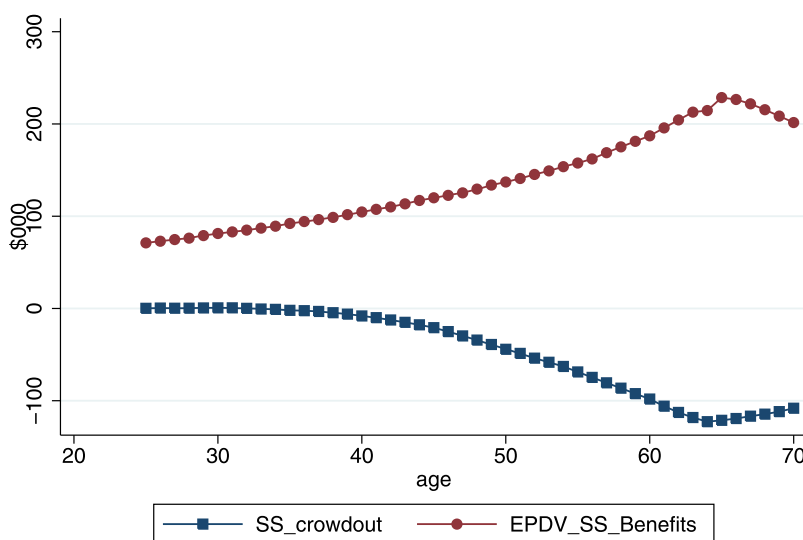|	Defined Benefit	Defined Contribution	Social Security
A. 75% have not yet claimed			
Age at which crowd out is measured	54	56	62
Dollar magnitude of crowd out	−8	−71	−113
Pension/SS wealth	104	197	204
Proportional crowd out	−0.08	−0.36	−0.55
B. 50% have not yet claimed			
Age at which crowd out is measured	57	62	63
Dollar magnitude of crowd out	−13	−98	−118
Pension/SS wealth	142	263	213
Proportional crowd out	−0.09	−0.37	−0.56

Note: Proportional crowd out equals dollar crowd out divided by pension/SS wealth. Monetary amounts are in thousands of 1992 dollars. In panel A, the age at which crowd out is measured is the last age at which at least 75% of simulated individuals remain on the pension job or have not yet claimed Social Security. In panel B, the criterion is 50%. Initial assets = 5K. Initial DC balance = 0.

The last age at which at least 75% of DC pension holders remain on the pension job is 56, and crowd out at that age is −0.36 as a proportion of the DC balance. New jobs in the model do not provide a pension, so it may seem strange that so many simulated cases would leave the pension job and thereby give up the option to accumulate savings tax free in a DC account. However, many of these cases result from layoffs, which occur at a 0.6% annual rate in the model. The last age at which no more than 25% of cases have claimed SS is 62, and crowd out at that age is −113K, or −0.55 as a proportion of SS wealth measured at age 62. The lower panel of the table shows that using the alternative criterion of the last age at which at least half the simulated cases remain on

the pension job (or have not claimed SS) has very little impact on proportional crowd out.[27]

The magnitude of simulated crowd out by SS is similar to many empirical estimates in the literature, although as noted in the Introduction such estimates are not directly comparable to those presented here.[28] The relatively small crowd out by DB pensions is consistent with some empirical findings, and is plausible because DB pensions are highly illiquid. But SS is illiquid as well, so it is somewhat surprising that SS crowd out is so much larger than DB crowd out. A unique feature of SS is its inflation-protected insurance against longevity and medical expenditure risk. SS also increases the reward to working past age 62, so crowd out predictions based on life cycle saving considerations alone can be quite misleading, as demonstrated by Samwick (2003). DC account balances are fairly liquid, so it is not surprising to find that DC crowd out is larger than DB crowd out.

I solved and simulated the model under some alternative assumptions, as a robustness analysis. Using different DB plans had a relatively minor impact on the results in most cases. For example, simulated crowd out was -0.07 and -0.10 for two other randomly selected plans compared to -0.09 for the plan used in the main analysis. Varying the mean interest rate made little difference to the results. Using an alternative bequest function proposed by Lockwood (2012), which encompasses many functional forms used in previous studies, did not affect the results. Increasing the number of points in the asset, AIME, and DC balance grids made little difference. A smaller employer match rate in the DC plan resulted in very similar crowd out (-0.40 versus -0.37 in Table 2).

4.4 *Sensitivity to model specification*

There are four key features that distinguish the model from the simpler version implicit in the empirical literature: the liquidity constraint, uncertainty, the choice of retirement and claiming ages, and the bequest motive. Here, I examine the sensitivity of the results to these features.

---

[27]An alternative approach to measuring crowd out as a single number, suggested by a referee, is to compute the "steady state" value: a weighted average across ages, with weights given by the survival probabilities. Crowd out measured this way is -0.04 for DB, -0.21 for DC, and -0.21 for SS, using joint survival probabilities of the husband and wife (assuming they are the same age) and using ages 25–69 only as in Figures 9–12. Crowd out often becomes positive at older ages, so I do not use them. These figures are noticeably smaller than the figures in Table 2.

[28]Estimates of SS crowd out from several well known studies include Attanasio and Brugiavini (2003): -0.49 (ages 45–54), -0.21 (ages 56–60), -0.11 (ages 61–65); Attanasio and Rohwedder (2003): 0.01 (ages 20–31), -0.55 (ages 32–42), -0.65 (ages 43–53), and -0.75 (ages 55–64); Hubbard (1986): -0.33; Hurd, Michaud, and Rohwedder (2012): -0.44; Kapteyn, Alessie, and Lusardi (2005): -0.11. Aguila (2011) and Gelber (2011) use a treatment effect approach, so comparable estimates of crowd out are not available in their studies. Chetty et al. (2014) estimate very small crowd out in response to policy changes that do not require an active saving response, and much larger crowd out in response to policy changes that require an active response by all individuals. Pension crowd out estimates include Engelhardt and Kumar (2011), -0.50; Gale (1998), -0.77; Gustman and Steinmeier (1999), 0.012; Hubbard (1986), -0.16. Some of these studies estimate a single crowd out parameter for pensions and SS combined.

One might expect larger crowd out in the absence of a liquidity constraint: the ability to borrow against future benefits reduces the need to save so as to finance consumption until the benefit can be claimed. However, in the presence of uncertainty a liquidity constraint may not be binding, since risk aversion will induce individuals to save for precautionary reasons. To study this issue, I solved a version of the model with no liquidity constraint. Two issues complicate the interpretation, however. Allowing individuals to take on an arbitrarily large amount of debt at young ages often results in a consumption splurge in the first period. To avoid this sharp and unrealistic discontinuity, I impose a lower bound on wealth of −50K at age 25, rising linearly to 0 at age 100.[29] The second issue is the consumption floor, which dampens the impact of relaxing the liquidity constraint since going into debt will never cause starvation. Thus, relaxing the liquidity constraint has very little impact with the consumption floor in place. Therefore, when the liquidity constraint is relaxed, the consumption floor is set to a much lower value: $c_F = 0.5$K, rather than 5.0K in the baseline specification.

Table 3 shows results from crowd out simulations for a variety of model specifications. Column 1 repeats the results from Table 2 for the baseline specification, using the 50% claiming criterion. The second column shows results from the no-liquidity-constraint case. Relaxing the liquidity constraint has very little impact on crowd out. This is the case in many other specifications I examined, and reflects the fact that the liquidity constraint is rarely binding.

Column 3 of Table 3 shows results for a model specification with no uncertainty. Eliminating uncertainty increases DB crowd out from −0.09 in column 1 to −0.39. This suggest a significant precautionary role for saving. Eliminating uncertainty decreases DC crowd out slightly and raises SS crowd out from −0.56 to −0.73, consistent with the pattern for the DB case.

To examine the importance of employment and pension claiming choices, I respecified the model with a fixed age of exit from the labor force (65), no possibility to change jobs or reenter employment, and no SS claiming decision.[30] Column 4 reports results for this specification. Eliminating the employment choice does not affect DB crowd out, causes a modest increase in DC crowd out, and causes a large increase in SS crowd out (compare columns 4 and 1). When the employment option is eliminated, the option of using additional pension wealth to finance earlier retirement is eliminated, resulting in an increase in crowd out. This reasoning is consistent with the results for DC and SS, but not for DB. However, pensions also have substitution effects that can be difficult to characterize given the complexity of pension formulas.

Column 5 reports results from a specification without a bequest motive. Eliminating the bequest motive has little impact on DB crowd out, and results in drastic declines in DC and SS crowd out. Saving is much lower without the bequest motive, so there are fewer household assets to be crowded out by pensions. The importance of the bequest motive for crowd out is rather unexpected. A large bequest motive may or may not be

---

[29]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 as I do here.

[30]The SS benefit must be claimed at 65. The DC claiming decision is still included in this specification, but the results are very similar if DC claiming is also mandatory following exit from employment.

TABLE 3. Simulated pension crowd out under alternative model specifications.

	1	2	3	4	5	6	7	8
Model Feature								
Employment choice	Yes	Yes	Yes	No	Yes	No	No	Yes
Uncertainty	Yes	Yes	No	Yes	Yes	No	No	Yes
Liquidity constraint	Yes	No	Yes	Yes	Yes	No	No	Yes
Bequest motive	Yes	Yes	Yes	Yes	No	No	No	Yes
Income tax	Yes	Yes	Yes	Yes	Yes	Yes	No	Yes
Defined Benefit, with SS								
Age	57	57	59	64	59	64	64	57
Proportional crowd out	−0.09	−0.09	−0.39	−0.07	−0.13	−0.26	−0.21	−0.10
Defined Contribution, with SS								
Age	62	61	69	64	55	64	64	62
Proportional crowd out	−0.37	−0.35	−0.32	−0.48	+0.01	−0.32	−0.39	−0.38
Social Security, no pension								
Age	63	63	64	65	62	63	64	63
Proportional crowd out	−0.56	−0.55	−0.73	−0.89	−0.13	−0.91	−1.06	−0.56
Social Security, with DB								
Age	62	62	63	64	62	64	64	62
Proportional crowd out	−0.24	−0.22	−0.29	−0.46	−0.24	−0.46	−0.42	−0.24
Social Security, with DC								
Age	63	63	62	66	63	63	63	63
Proportional crowd out	−0.25	−0.23	−0.41	−0.49	−0.11	−0.50	−0.39	−0.26

Note: "Age" refers to the last age at which at least 50% of the simulated cases have not yet left the pension job (DB), claimed the pension (DC), or claimed Social Security. This is the age at which crowd out is measured. In the scenarios with no employment choice, the mandatory age of retirement is 65, so the last age at which at least 50% of individuals remain on the pension job is 64. DC pension holders cannot claim until they leave the pension job, but they are not required to claim until age 70. In the no-employment scenario in column 4, there is also no job switching or reentry. Column 8 shows the substitution crowd out effect: see the text for explanation.

realistic, but something causes households to avoid spending down assets as rapidly as expected based on the stylized life cycle model. The other plausible explanation is medical expenditure risk, but this is incorporated here. See Lockwood (2012) and De Nardi, French, and Jones (2010) for evidence on the bequest and medical expenditure explanations for the slow rate of asset spend down at older ages.

Column 6 reports results from a specification that combines all four of the restrictions discussed above. Compared to the baseline results in column 1, DB crowd out increases from −0.09 to −0.26, DC crowd out is very similar at −0.32 compared to −0.37 in column 1, and SS crowd out increases substantially from −0.56 to −0.91. The model specification reported in column 6 is close to the stylized model described at the end of Section 2. Yet crowd out is far less than 1 in absolute value in the DB and DC cases. One additional feature that distinguishes the model from the stylized version is the income tax. This could be important because of the tax advantages of pensions and SS. Eliminating the income tax, and therefore the tax benefit of saving in a DC plan, might be expected to increase substitutability between DC saving and household wealth. SS benefits are taxed differently from other income, so SS crowd out may be affected as

well.[31] To determine whether this is the case, I solved and simulated a specification like the one in column 6, but eliminated taxes. As shown in column 7, DB and DC crowd are barely affected by this change (compare columns 6 and 7), while SS crowd out increases still more to −1.06. It is not obvious why this specification of the model fails to produce crowd out closer to −1 for the DB and DC cases, but this is a very robust finding.

SS crowd out is calculated here under the assumption that there is no employer-provided pension. The bottom two panels in Table 3 repeat the SS crowd out calculations for the cases with a DB or DC pension. This makes a big difference: SS crowd out is much smaller when SS is not the only source of retirement income. For the main specification in column 1, SS crowd out falls from −0.56 with no pension to −0.24 with a DB pension and −0.25 with a DC pension.

4.5 *Compensating variation*

The value of a pension is measured here by the compensating variation (CV): the amount by which the initial assets of a household without a pension must be increased so as to equate its EPDV of optimized lifetime utility (value function) at age 25 to that of the same household with a pension. Dividing the CV by initial pension wealth at age 25 provides a measure with the same scale as the proportional crowd out measure. Table 4 reports CV figures for the baseline specification. The lifetime-utility-equivalent value of the DB pension is −1.7K, which is −26% of the 6.6K EPDV of the DB benefit. This indicates that, conditional on the availability of a real annuity from SS, households in the

TABLE 4. Compensating variation (CV) associated with pensions and Social Security.

	Compensating Variation (CV)	Pension/SS Wealth at Age 25	Proportional CV
Defined Benefit			
With Social Security	−1.7	6.6	−0.26
Without Social Security	4.4	6.8	0.65
Defined Contribution			
With Social Security	18.6	78.9	0.24
Without Social Security	21	81.1	0.26
Social Security			
No pension	−5.0	71.1	−0.07
With defined benefit	−9.7	57.6	−0.17
With defined contribution	−4.8	74.5	−0.06

Note: Initial assets are 5K in the baseline case. The first data column shows the amount by which initial assets must be increased so as to equate the EPDV of lifetime utility with and without the pension or Social Security. The second data column shows pension/SS wealth at age 25 in the baseline. Since the DC balance at age 25 is 0 by assumption, the EPDV of the DC plan given optimal behavior is reported in the second data column, computed as described in the text. The third data column shows compensating variation as a fraction of baseline pension/SS wealth (first data column divided by second data column). Monetary amounts are in thousands of 1992 dollars.

---

[31]A portion of the SS benefit is excluded from taxable income, with a higher proportion excluded for low income households. The portions excluded are 100% for low income households, 50% for medium income households, and 15% for higher income households. In the no-income-tax scenario the SS payroll tax remains part of the model.

model place no value on the DB pension. The main reason for this surprising finding is the small EPDV of the DB pension, which is itself a consequence of the high calibrated discount rate (8%). In a scenario without SS, the DB pension is slightly more valuable: 4.4K.

In the DC case the CV is 18.6K, but the DC balance is 0 at age 25 by construction, so a proportional CV cannot be computed using this approach. Instead, I computed the EPDV of the DC plan at age 25 given optimal future behavior. This turned out to be 78.9K, yielding a CV for the DC plan of 0.24. The DC plan is valuable because of the employer match, but it is a form of forced saving and is costly to tap into until age 59. In this case, the compensating variation is only modestly higher without SS.

For SS, the CV is −5.0K, −7% of the value of SS wealth at age 25. SS is evidently of no *ex ante* value to a household that would otherwise have to fully finance its consumption expenditure in retirement as well as self-insure against a long life. This is surprising, but is a consequence of the low rate of return on tax payments and the high calibrated value of the discount factor. This is consistent with the fact that for the 1937 birth cohort the net present value of SS accounting for the payroll tax is small: evaluated at the mean simulated earnings profile, the net EPDV of SS is 17K at a 3% real interest rate for claiming at ages 62–65, and is negative evaluated at the 8% discount factor.[32] In a scenario with a DB plan, SS is even less valuable, with a proportional CV of −0.17. With a DC plan, SS has about the same value as in the no-pension scenario, suggesting that DC plans are not a good substitute for SS.

The CV calculations can be used to decompose the saving response to a pension into two components. The first is the wealth effect: the response of household saving to the increase in lifetime wealth associated with the pension. The second is the substitution effect: holding lifetime utility constant, how does the pension alter saving incentives? The substitution effect is a result of the longevity insurance provided by pensions, reducing the need for precautionary saving. To compute the substitution effect, I added the dollar CV to the initial wealth of the no-pension (or no-SS) household and computed crowd out by comparing saving in this case to saving in the pension (or SS) case. The results are shown column 8 of Table 3, which can be compared to the first column. The substitution crowd out effects are virtually identical to the total crowd out effects. This is not surprising given the low and even negative value of pensions shown in Table 4, indicating very small and even negative wealth effects.

4.6 *Regression estimates of crowd out*

Here I explore whether regression estimates of crowd out are a reasonably accurate guide to behavior when the data are generated from a process that does not obey the strong assumptions of the stylized model. The simulation approach used to generate the data for the regression analysis is somewhat different from the approach used above. The results presented so far are based on simulations of four cases (NP, DB, DC, and NPNSS), so the sample size is rather small for the regression analysis. Instead I simulate

---

[32]Note that the benefit includes the spousal benefit, equal to 50% of the husband's benefit. Excluding the spousal benefit, the EPDV of lifetime benefits at age 25 is approximately equal to zero at a 3% interest rate.

TABLE 5. Regression models of wealth using simulated data.

Pension/SS Wealth Measure	Lagged Dependent Variable Included		Lagged Dependent Variable Excluded	
	Full Model	Standard	Full Model	Standard
EPDV DB benefit	−0.38	−0.13	−0.39	−0.14
DC balance	−0.64	−0.45	−0.69	−0.53
EPDV SS benefit	−0.47	−0.33	−0.50	−0.35
R^2	0.98	0.94	0.97	0.93

Note: "Full model" indicates that the pension/SS wealth measures were calculated as part of the model solution, as described in the text. "Standard" indicates that pension/SS wealth measures were calculated using the actual pension and SS claiming ages, assuming no uncertainty (except over date of death), no employment choice, and no liquidity constraint. The entries are coefficient estimates on the indicated variables in regression models of household wealth. EPDV is expected present discounted value; DB is defined benefit; DC is defined contribution; SS is Social Security. The other explanatory variables, with coefficient estimates from the specification in the first column, are current net annual earnings (−21), EPDV of remaining lifetime earnings (−0.22), and (in the first two columns) lagged wealth (0.85) and the lagged DC balance (0.13). Initial assets are 5K. Sample size is 3581. Each observation used in the regressions is the mean over 1000 simulations, using the four pension/SS cases (NP, DB, DC, and NPNSS) in combination with alternative values of initial assets, the initial DC balance, and AIME. The sample includes observations at ages at which no more than 50% of simulated cases have claimed the pension or Social Security.

several different scenarios for each of the four pension/SS cases, varying initial assets, the initial DC balance, and initial AIME. This resulted in 176 simulated scenarios. In each scenario I include in the regression sample observations at all ages at which no more than 50% of simulated cases have claimed the pension or SS benefit. This resulted in a sample of 3581 observations, each of which is the average over 1000 simulated cases at a given age (instead of 1500, to reduce computation time).

Table 5 reports coefficient estimates on pension and SS wealth in household wealth regressions. I computed the life cycle adjustment factor (κ_t; see Section 2) using a continuous time approximation from Gale (1998). The other regressors are those implied by the theory: the EPDV of remaining lifetime earnings, current period earnings net of payroll tax and worker DC contribution, and household wealth and the DC balance in the previous period, all adjusted by κ_t. The estimates in the first data column in Table 5 using measures of pension wealth derived from the model solution indicate crowd out of −0.38 in the DB case, −0.64 for DC pensions, and −0.47 for SS. The DB and DC crowd out estimates are substantially larger in absolute value than the corresponding simulated values of −0.09 and −0.37 reported in Table 3. The regression estimate of SS crowd out is relatively close to the simulated value of −0.56. The results reported in the second data column of Table 5 use a more standard measure of the EPDV of future DB and SS benefits, calculated using the observed claiming age and benefit, and ignoring uncertainty except over mortality. The regression estimates of DB and DC crowd out in this case are −0.13 and −0.45, much closer to the simulated values, but the regression estimate of SS crowd out moves further away from the simulated value. For comparison with the more common cross section and cohort regression models used in the literature, the third and fourth data columns present estimates from a specification that does not control for lagged wealth or the lagged DC balance. This specification yields very similar results.

The results shown in Figures 9–12 are essentially nonparametric estimates of crowd out. To explore the difference between these estimates and the parametric regression

estimates, I used the small simulated sample from Figures 9–12 to estimate a nonlinear wealth regression: quadratic in pension wealth, controlling for a quadratic in age. Using the same criteria to select the observations to use in the regression, the sample size is 151. The advantage of using this sample is to ensure that any differences are not the result of the alternative simulation approach used for the analysis reported in Table 5. The fitted values from the regression evaluated at the same ages as in Table 3 yield crowd out results quite close to those in column 1 of Table 3: DB: −0.06, DC: −0.36, SS: −0.59. Thus, the quadratic specification provides a reasonably good approximation to the simulation results. Using these data to estimate the specification in the first data column of Table 5 yields crowd out coefficients of −1.04 for DB, −0.60 for DC, and −0.51 for SS. The DB and SS estimates are far from the nonparametric estimates. The column 3 specification yields −1.03 for DB, −0.54 for DC, and −0.50 for SS. In this case DB and DC estimates are far from the corresponding simulation results. These results suggest that the linear regression model of household wealth may be seriously misspecified. Future empirical research using this approach should explore alternative functional forms.

5. An application of the model to Social Security reform

An interesting application of the model is to analyze the impact of an important SS policy reform, the increase in the full retirement age (FRA). The 1983 amendments to the Social Security Act increased the FRA from 65 for birth cohorts up to 1937, to 66 for cohorts 1943–1954, and 67 for cohorts born in 1960 and after. A 1 year increase in the FRA is equivalent to a 6.67% cut in the benefit, holding claiming age constant.[33] Here I simulate the effects of these two policy changes, along with the effects of further hypothetical increases in the FRA to 68 and 69.

The first data column of Table 6 shows simulation results for the baseline case with a FRA of 65, separately for the NP, DB, and DC scenarios. The table shows results for four outcomes: assets at age 63,[34] the EPDV of lifetime consumption and lifetime utility, and mean age of exit from employment ("retirement," for brevity). The remaining columns show the percentage effects on these outcomes of changes in the FRA, except for the change in the age of retirement, which is measured in years.

Increasing the FRA from 65 to 66 is predicted to cause a small decline in lifetime consumption (−0.3 to −0.5%) and a small increase in the age of retirement (0.2 to 0.3 years). The resulting change in lifetime welfare ranges from −0.4 to −1.2% depending on the pension scenario. Assets at age 63 rise by 2.3% and 2.4% in the no-pension and DC cases, and by 0.4% in the DB case. Empirical evidence on the impact of the increase in the FRA is scarce. Banerjee (2010) used the increase in the FRA as a quasi-experiment to estimate the magnitude of SS crowd out. His estimates imply that the increase in the FRA from 65 to 66 caused saving to increase by 1.5%. Mastrobuoni (2009) estimates that

---

[33]This is true for any claiming age between 63 and the FRA. The benefit cut is 5% if the benefit is claimed at 62. The increase in the FRA from 65 to 66 was phased in at a rate of two months per birth year from 1938–1943, and similarly for the increase from 66 to 67 for cohorts 1955–1960.

[34]Assets are measured at age 63 because that is the last age at which no more than 50% of simulated individuals have claimed SS in the FRA = 65 case.

TABLE 6. Simulated impact of alternative changes to the Social Security full retirement age.

		Changes Due to Increase in the FRA to			
	Baseline FRA (65)	66	67	68	69
No pension					
Assets at age 63	177.3	2.3%	5.7%	9.2%	10.9%
Age of retirement	58.5 (years)	0.3	0.5	0.5	1.1
Lifetime welfare	−0.02128	−0.7%	−1.4%	−2.0%	−2.7%
EPDV life cons.	685	−0.5%	−0.9%	−1.7%	−1.8%
DB pension					
Assets at age 63	205.0	0.5%	0.8%	1.1%	1.7%
Age of retirement	65.0 (years)	0.2	0.4	0.6	0.9
Lifetime welfare	−0.02252	−0.4%	−0.8%	−1.1%	−1.5%
EPDV life cons.	679	−0.3%	−0.5%	−0.7%	−1.0%
DC pension					
Assets at age 63	180.7	2.4%	4.2%	7.8%	9.2%
Age of retirement	58.6 (years)	0.2	0.5	0.3	1.0
Lifetime welfare	−0.01184	−1.2%	−2.5%	−3.7%	−5.0%
EPDV life cons.	711	−0.4%	−0.8%	−1.5%	−1.6%

Note: Assets are measured at age 63 because that is the last age at which no more than 50% of simulated individuals have claimed Social Security in the FRA = 65 case. Lifetime welfare is the value function in the first period. The simulation results shown here are based on smaller grid sizes than in the other results in the paper. Computer memory constraints made it impossible to simulate the impact of increasing the FRA beyond 66 without reducing the size of the state space. The grid sizes used here are 60 for assets and 13 for AIME, compared to 70 and 15 for the results in Tables 2–5. Monetary amounts are in thousands of 1992 dollars.

the increase in the FRA from 65 to 66 caused the mean age of retirement to increase by about 0.5 years. Behaghel and Blau (2012) find a similar effect on retirement.

The remaining columns of Table 6 show the simulated effects of further increases in the FRA. The hypothetical increase to age 69 results in increases in assets of 9–11% for NP and DC, and only 1.7% for DB. Retirement age increases by about 1 full year, consumption declines by 1–2%, and welfare declines by about 3% in the NP case, 1.5% in the DB case, and 5% in the DC case, all relative to the baseline FRA of 65. Raising the FRA from 65 to 69 is equivalent to a 27% benefit cut (for claiming at ages 63+). The welfare impact is an order of magnitude smaller. This is perhaps not surprising given concavity of the utility function and the option to substitute between leisure and consumption.

6. Conclusions

Empirical studies of crowd out of household saving by pensions have focused on identification issues, which are obviously important for obtaining useful estimates. I take a different and complementary approach, studying crowd out behavior in a model that incorporates choice alternatives and constraints more realistically than in the stylized life cycle model that is the implicit or explicit basis for empirical analysis. Broadly speaking, the results show that modeling choices matter: imposing the strong assumptions of the stylized life cycle model has large effects on the magnitude of simulated crowd out in several cases. Specifically, crowd out by SS almost doubles, from −0.56 to −1.06, between

the least and most restrictive specification. DB crowd out also is sensitive to specification, but never exceeds −0.40. DC crowd out is relatively insensitive to specification. Regression estimates of crowd out using the simulated data are also quite sensitive to specification.

The results in this paper are of more than academic interest. There have been many public pension reforms around the world, ranging from pure benefit cuts to complete restructuring. Empirical estimates of crowd out using such reforms for identification are of great value but provide an incomplete picture of the effects of alternative approaches to reform. This paper provides a framework that can be used to develop a richer analysis of the behavioral effects and welfare implications of pension reforms. The framework needs further development before it would be useful for serious policy analysis, but it is an initial step.

An obvious next step in this line of research is to estimate the model structurally. I started this project intending to do so, but I did not fully appreciate the difficulties caused by pension heterogeneity. Each DB pension is different from others, and a global approach to approximating the value function was unsuccessful in capturing this variation. As a result, I used a local approximation method, which greatly limits modeling flexibility due to computer memory constraints and computation time. As computation power becomes cheaper, structural estimation of a model with multiple pension plans should become feasible in the future. And it should become feasible to estimate the model on a more heterogeneous sample. Both of these advances would greatly improve the usefulness of the framework.

References

Aguila, E. (2011), "Personal retirement accounts and savings." *American Economic Journal—Economic Policy*, 3, 1–24. [214]

Alessie, R. J. M., A. Kapteyn, and F. Klijn (1997), "Mandatory pensions and personal saving in the Netherlands." *De Economist*, 145 (3), 291–324. [194]

Attanasio, O. P. and A. Brugiavini (2003), "Social security and household saving." *Quarterly Journal of Economics*, 188 (3), 1075–1120. [196, 214]

Attanasio, O. P. and S. Rohwedder (2003), "Pension wealth and household saving: Evidence from pension reforms in the United Kingdom." *American Economic Review*, 93 (5), 1499–1521. [194, 196, 214]

Attanasio, O. P. and G. Weber (2010), "Consumption and saving: Models of intertemporal allocation and their implications for public policy." *Journal of Economic Literature*, 48 (3), 693–751. [194]

Banerjee, S. (2010), "Does social security affect household saving?" Working paper, Ohio State University. [220]

Behaghel, L. and D. M. Blau (2012), "Framing social security reform: Behavioral responses to changes in the full retirement age." *American Economic Journal—Economic Policy*, 4, 1–30. [209, 221]

Blau, D. M. (2008), "How do pensions affect household wealth accumulation?" Working paper, Ohio State University. [203]

Chan, S. and A. H. Stevens (2008), "What you don't know can't help you: Pension knowledge and retirement decision-making." *Review of Economics and Statistics*, 90 (2), 253–266. [203]

Chetty, R., J. N. Friedman, S. Leth-Petersen, T. H. Nielsen, and T. Olsen (2014), "Active vs. passive decisions and crowd-out in retirement savings accounts: Evidence from Denmark." *Quarterly Journal of Economics*, 129 (3), 1141–1219. [196, 214]

Coile, C. C., P. Diamond, J. Gruber, and A. Jousten (2002), "Delays in claiming social security benefits." *Journal of Public Economics*, 84 (3), 357–385. [206]

De Nardi, M., E. French, and J. B. Jones (2010), "Why do the elderly save? The role of medical expenses." *Journal of Political Economy*, 118 (1), 39–75. [216]

Engelhardt, G. V. and A. Kumar (2011), "Pensions and household wealth accumulation." *Journal of Human Resources*, 46 (1), 203–236. [203, 214]

Engen, E. M., W. G. Gale, and C. E. Uccello (1999), "The adequacy of household saving." *Brookings Papers on Economics Activity*, 2, 65–165. [196]

Feldstein, M. (1974), "Social security, induced retirement, and aggregate capital accumulation." *Journal of Political Economy*, 82 (5), 905–926. [194, 200]

French, E. (2005), "The effects of health, wealth, and wages on labour supply and retirement behavior." *Review of Economic Studies*, 72, 395–427. [196, 199, 209]

Gale, W. G. (1998), "The effects of pensions on household wealth: A reevaluation of theory and evidence." *Journal of Political Economy*, 106 (4), 706–723. [194, 214, 219]

Gelber, A. M. (2011), "How do 401(k)s affect savings? Evidence from changes in 401(k) eligibility." *American Economic Journal—Economic Policy*, 3, 103–122. [214]

Gustman, A. L. and T. 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. [214]

Gustman, A. L., T. L. Steinmeier, and N. Tabatabai (2010), *Pensions in the Health and Retirement Study*. Harvard University Press, Cambridge. [203]

Hubbard, R. G. (1986), "Pension wealth and individual saving." *Journal of Money, Credit, and Banking*, 18 (2), 167–178. [214]

Hubbard, R. G., J. Skinner, and S. P. Zeldes (1995), "Precautionary saving and social insurance." *Journal of Political Economy*, 103 (2), 360–399. [199]

Hurd, M., P.-C. Michaud, and S. B. Rohwedder (2012), "The displacement effect of public pensions on the accumulation of financial assets." *Fiscal Studies*, 33 (1), 107–128. [196, 214]

Inkmann, J., P. Lopes, and A. Michaelides (2011), "How deep is the annuity market participation puzzle?" *Review of Financial Studies*, 1, 279–319. [201]

Kapteyn, A., R. Alessie, and A. Lusardi (2005), "Explaining the wealth holdings of different cohorts: Productivity growth and social security." *European Economic Review*, 49, 1361–1391. [196, 214]

Laibson, D. I., A. Repetto, and J. Tobacman (1998), "Self-control and saving for retirement." *Brookings Papers on Economics Activity*, 1, 91–172. [196]

Lockwood, L. M. (2012), "Incidental bequests: Bequest motives and the choice to self-insure late life risks," Working paper, Northwestern University. [214, 216]

Mastrobuoni, G. (2009), "Labor supply effects of the recent social security benefit cuts: Empirical estimates using cohort discontinuities." *Journal of Public Economics*, 93, 1224–1233. [209, 220]

Samwick, A. (2003), "Patience, pensions, and savings." Working paper, Dartmouth College. [196, 214]

Scholz, J. K., A. Seshadri and S. Khitatrakun (2006), "Are Americans saving optimally for retirement?" *Journal of Political Economy*, 114 (4), 607–643. [196]

Social Security Administration (2010), "Annual statistical supplement to the Social Security Bulletin, 2009." SSA Publication 13-11700. [198]

Stock, J. H. and D. A. Wise (1990), "Pensions, the option value of work, and retirement." *Econometrica*, 58 (5), 1151–1180. [208]

US Bureau of Labor Statistics (2005), "National Compensation Survey: Employee benefits in private industry in the United States, 2002–2003." Bulletin 2573. [205]

van der Klaauw, W. and K. I. Wolpin (2008), "Social security and the retirement and saving behavior of low income households." *Journal of Econometrics*, 45 (1–2), 21–42. [196, 197, 199, 205, 209, 215]