Pension Wealth and Household Saving: Evidence from Pension Reforms in the United Kingdom

By ORAZIO P. ATTANASIO AND SUSANN ROHWEDDER*

Using three major U.K. pension reforms as natural experiments we investigate the relationship between pension saving and discretionary private savings. Unlike most differences-in-differences approaches which rely on average differences between control and treatment group, we use economic theory to model the response of each individual household. The empirical analysis, based on the Family Expenditure Survey, uses both time-series and cross-sectional variation to identify the behavioral response. The earnings-related tier of the pension scheme is found to have a negative impact on private savings with relatively high substitution elasticities; the impact of the flat-rate tier is not significantly different from zero. (*JEL* H55, H91)

The relationship between pension wealth and household savings is crucial for important policy issues, such as establishing the effects of changes in pension legislation on saving behavior. From a theoretical point of view, the life-cycle framework suggests that the provision of public benefits after retirement constitutes a negative incentive to accumulate wealth during one's working life. Quantifying the resulting relationship between pension wealth and saving, even within a simple theoretical framework, is difficult, however, as the degree of substitutability between pension and nonpension wealth depends on a variety of factors, ranging from the presence of liquidity constraints that might be binding, to the importance of bequest motives, to the size of discount factors and rates of return, to the possibility of distortionary effects of pension contributions on labor supply behavior. The complexity of the theoretical relationship makes its empirical quantification all the more important.

Given the importance of the issues involved, the empirical evidence on the relationship between pension wealth and household saving is relatively scarce. Moreover, the results are far from being conclusive. Martin S. Feldstein (1974) was among the first to investigate the relationship between public pension wealth and household saving empirically. His study, based on the time-series behavior of aggregate saving rates and pension wealth indicates a large negative effect of pension wealth on saving rates. Subsequently, Mervin A. King and Louis A. Dicks-Mireaux (1982) provided evidence from microdata analyzing the relationship between the stock of pension wealth and saving rates. They interpreted the coefficient of pension wealth in a regression for financial wealth as a measure of the degree of substitutability between the latter and the former. They find that an increase of a dollar in social security wealth decreases financial wealth by 25 cents, while an increase in private pension wealth decreases

* Attanasio: Department of Economics, University College London, Gower Street, London WC1E 6BT, United Kingdom, and Institute for Fiscal Studies, 7 Ridgmount Street, London WC1E 7AE, United Kingdom (e-mail: o.attanasio@ucl.ac.uk); Rohwedder: RAND, 1700 Main Street, Santa Monica, CA 90407, and Department of Economics, University College London (e-mail: susannr@rand.org). Both authors thank the TMR Network (Contract No. ERBFMRXCT960016) on "Structural Analysis of Household Savings and Wealth Positions over the Life-Cycle" for financial support. Attanasio also acknowledges the funding received from the Leverhulme Trust under the project "The Changing Distribution of Consumption, Economic Resources, and the Welfare of Households." Material from the Family Expenditure Survey made available by the Office for National Statistics (ONS) through the ESRC data archive has been used by permission of the controller of Her Majesty's Stationers Office. Moreover, the authors wish to thank James Banks, Richard Blundell, Richard Disney, Carl Emmerson, Michael Hurd, Paul Johnson, Costas Meghir, and seminar participants at UCL, RAND, and at the TMR conference in Deidesheim for helpful comments. Any remaining errors are, of course, our own.

financial wealth by 10 or 18 cents, depending on the methodology.

Peter A. Diamond and Jerry A. Hausman (1984) and R. Glenn Hubbard (1986), using U.S. data, find even smaller offsets than King and Dicks-Mireaux (1982). Agar Brugiavini (1987) and Tullio Jappelli (1995) have carried out exercises similar to those of King and Dicks-Mireaux (1982) on Italian data, obtaining larger estimates of the degree of substitutability between private wealth and public pension provisions.

More recently, the relationship between pension wealth and savings has received renewed attention. William G. Gale (1998) uses U.S. microdata from the *Survey of Consumer Finances* to assess the extent to which changes in pension wealth are offset by changes in other forms of savings. Gale (1998), as we do below, stresses the importance of adjusting pension wealth for the age of the individual. He finds that, for his favorite specification, the offset is substantial, between 0.39 and 0.82, depending on the estimator used. Alan L. Gustman and Thomas L. Steinmeier (1999), on the other hand, reach different conclusions using data from the Health and Retirement Study (HRS). They find very small and insignificant offsets of pension wealth to total wealth.[1]

The scarcity of empirical evidence on the relationship between pension wealth and saving behavior is in part due to the severe data limitations that researchers face in this area. Aggregate time-series information, while it covers potentially long time periods and therefore changes in pension wealth, is of limited use because of aggregation issues. Data sets containing individual-level information on income and consumption (and therefore saving) are few and far between. Moreover, data sets that contain this type of information very rarely contain information on financial wealth and even less on pension entitlements. In addition to the limitations in available data sets, pension wealth does not exhibit much exogenous variation, especially because microdata sets typically cover relatively short time periods. Therefore, differences in pension wealth among individuals could reflect differences in tastes for saving.

The approach we use in this paper differs from those used in the literature in two important dimensions. First, the data of our study allow us to identify the parameter of interest in a much more robust way than in previous studies. Rather than using data on wealth, we consider the implications of different levels of pension wealth for *saving rates.* We compute these from high-quality household data that contain detailed information on income and consumption and go back to the early 1970's. Because we have a long series of cross sections we can exploit not just cross-sectional variation, but also variation over time induced by several important reforms to the U.K. public pension system which affected different groups of individuals differently. Our period of study contains as many as three major reforms that induced substantial exogenous variation in pension wealth. On the basis of these data we exploit the variability in pension wealth induced by the reforms to identify the relationship between pension wealth and saving rates. Second, we combine the differences-in-differences estimations with important implications from a simple theoretical framework on how to relate stock and flow variables and on how to control for timing effects of the reforms with respect to the age of the household and with respect to when the household is observed in the data. As a result we are able to derive an empirical specification in which the coefficients of interest can be directly interpreted as the degree of substitutability between financial and pension wealth.

The conceptual framework we use is that of a life-cycle model in which individuals save, among other reasons, to finance their retirement. In a simple version of the model, pension wealth is a perfect substitute of financial wealth. Our approach, like the one in Gale (1998), recognizes that the effect of a change in pension wealth on saving behavior depends on the position of an individual over her life cycle. Moreover, consistently with the model, the effect on observed saving rates depends on when in the past the individual experienced an (unexpected) pension reform. As Gale, we recognize that several factors, such as uncertainty and liquidity

[1] The literature on the effectiveness of tax incentives on saving, such as IRAs and 401(k)s, is also relevant to whether pension wealth constitutes a good substitute for other forms of wealth. See, for instance, Eric M. Engen et al. (1996), James M. Poterba et al. (1996), and B. Douglas Bernheim (1997) and the references therein.

constraints, might result in less than a full offset. However, our empirical specification is more flexible than Gale's in that we let the degree of substitutability between pension and other forms of wealth vary with age, recognizing that the importance of these factors is likely to be different for individuals at different stages of their life cycle. The paper most similar to ours is Attanasio and Agar Brugiavini (2003), which has investigated the impact of the 1992 reform of the Italian pension system. Using two household surveys collected before and after 1992 they assess the effect that pension wealth has on household saving. In the Italian case, as in our study, the reform had different effects for different groups of the population.

The data we use are from the Family Expenditure Survey (FES) covering a period in which three major pension reforms were implemented. We use the variation induced by these reforms, and in particular the fact that they affect different groups of individuals differently, to identify the impact of changes in pension wealth on household saving rates. Our study goes beyond a simple differences-in-differences approach because we impose on the empirical analysis the structure that we derive from economic theory. We use the life-cycle framework to model the behavioral response of every household. If pension reforms are not fully anticipated, individuals' reactions will vary depending on when in their life cycle they experience the reform: for example, younger individuals will have more time to absorb the changes entailed by the reform. In this context the life-cycle model gives us a framework in which we can specify a regression equation and in which we can interpret the estimated coefficients.

While the relationship between pension wealth and saving is relatively simple in a stylized model, reality is likely to be much more complex. In particular, the presence of uncertainty (about income and rates of return), liquidity constraints, the lack of liquidity of pension wealth, the interaction of saving for old age with other savings motives, and the interaction with labor supply choices make the optimization problem of a typical agent very complex. Trying to model all these features is extremely hard, and typically requires using numerical methods in order to obtain a solution for the optimization problem. In this paper, we use a simplified theoretical setting to inform the specification of our regression equation and interpret the results we obtain. At the same time we keep our empirical specification flexible to take into account the fact that our simple model abstracts from very important features of reality.

The paper is organized as follows. In Section I we give a brief overview of the U.K. pension system and of the main reforms that have changed its operation in the last 25 years. In Section II, we describe the data set we use and how we construct for each individual household in our sample an estimate of pension wealth. We do this by using information on the individual household members and the pension legislation, including its changes over the years. In Section III we discuss the theoretical framework that informs our econometric specification and present some preliminary evidence on the relationship between pension wealth and saving. Section IV deals with econometric issues and reports our main results. Section V concludes the paper.

I. The U.K. Pension System

Since the end of the last world war the public pension system in the United Kingdom has been subject to repeated changes. Each decade has seen at least one substantial pension reform. As a result, entitlements to future benefits exhibit a fair amount of variation over time. Moreover, many of these reforms have affected different individuals in different ways, so that the induced changes in pension wealth (i.e., the expected present value of future benefits less future contributions) also exhibit a considerable amount of cross-sectional variability. In most other industrialized countries pension reforms have only recently been included in the policy agenda so that the United Kingdom stands out in this respect as an interesting case to study. The reform process has led to pension arrangements where both unfunded public and funded private pensions are integral parts of the system.

While the rules that govern the public component of the system are relatively straightforward and homogeneous, the private sector is very heterogeneous. People are enrolled in very different schemes whose rules vary greatly. Without knowing which pension plan someone belongs to it is not possible even to approximate the person's entitlements. To our knowledge there is no representative household survey in the United Kingdom that would provide this information, forcing us to confine our analysis to the impact of changes in public sector schemes. Nevertheless, it is important to note that this constraint does not interfere with the findings of our study for two reasons. Firstly, until 1987 participation in private pensions could conceivably be assumed to be exogenous with respect to saving choices, as occupational pensions were available only for those employees whose employer provided such a scheme. Moreover, the introduction of SERPS (State Earnings-Related Pension Scheme) in the 1970's did not change this and policy makers paid particular attention in the design of SERPS not to crowd out private pension provisions.[2] Coverage rates of occupational pensions were stable over that time: overall it was between 49 percent and 52 percent throughout our sample period (Richard F. Disney et al., 1999). Coverage increased gradually for women by about 6 percent, which goes along with women increasingly participating in the labor market. In 1988 the way public and private pensions were integrated changed, making the assumption of exogeneity less plausible.[3] Therefore, we limit our analysis to the period before 1988. Furthermore, the heterogeneity among private pension schemes, which hindered their inclusion in the analysis in the first place, now becomes a virtue. Taken together with the fact that there have not been major regulatory changes in private pension arrangements between 1974 and 1987, we argue that households observed in our sample do not experience any common shocks of considerable magnitudes in their private pension wealth. Hence as long as we control for membership in private and public sector pensions our results should not suffer from the exclusion of a measure of private pension wealth.

In the remainder of this section we describe the evolution of the different components of the public pension scheme, giving greater details for the period we use in our study, that is the 1970's to late 1980's. We start by describing the two main components of the U.K. public pension system, the flat-rate Basic State Pension (BSP) and the so-called "State Earnings-Related Pension Scheme" (SERPS). We then explain the possibility of opting out of the public system. This is followed by a discussion of the variations in pension wealth induced by the reforms that we use in our empirical analysis to identify the effect of pension wealth on saving.

A. *Basic State Pension (BSP)*

The structure of the BSP was put in place in 1948. It is a flat-rate benefit scheme that is compulsory for all workers and employees with earnings above a threshold, the lower earnings limit. When reaching state pension age—currently 60 for women and 65 for men—an individual who has contributed at least nine-tenths of his or her working life is entitled to a full-rate pension that will be paid for the remainder of his or her life. Virtually all men acquire entitlements to a full pension whereas until recently many women did not. In that case their husbands can claim a dependant's addition, which amounts to 60 percent of the full rate.[4]

Until the early 1970's benefits were up-rated in a rather ad hoc way, on average making up for slightly more than average earnings growth. From 1975 onwards, they were increased

[2] Andrew Dilnot et al. (1994) state: "The Social Security Act of 1975 introduced SERPS. To avoid substituting for private sector provision, occupational schemes were allowed to contract out of SERPS." Contracting out was allowed only if the private pension was at least as generous as SERPS. In that case employers and employees were granted a reduction in National Insurance Contributions, which was designed to be comparable to the contributions to private pensions.

[3] We discuss this change and its implications in further detail below when explaining the possibilities of opting out of public sector second-tier pensions. A good discussion of the current U.K. pension system can be found in Disney et al. (2001).

[4] The reasons for fewer women acquiring their own pension are twofold. First, until 1978 women had the possibility to pay reduced National Insurance Contributions (the U.K. term for Social Security Contributions) in exchange for forgoing benefits to the Basic State Pensions in their own right. Furthermore, many women have incomplete contribution records, in a large part due to childbearing. The reform of 1978 dealt with this issue by explicitly recognizing periods of "home responsibilities" as contribution years in the pension scheme. It also introduced rules that would lead more divorced women to be able to claim their own pension. By 2002 all women will have had the chance to take full advantage of these changes.

roughly in line with gross earnings. In 1980–1981 the government decided to link them to prices only—a step that would subsequently reduce considerably the growth of public pensions, given that real earnings have exhibited positive growth on average. As a result, the value of the flat-rate benefit, which equaled 20 percent of average earnings at its peak in the second half of the 1970's, has been eroded. Now the value stands at just under 15 percent of average earnings and is projected to fall further to be worth about 7 percent of average earnings by 2050.[5]

These indexation changes are the most important reforms of the BSP contained in the period we study. Note that this kind of reform is relevant for everyone in the population, and the effects differ across households depending mainly on the age of the head of household and, in the case of married couples, on the age of the spouse.

B. *State Earnings-Related Pension Scheme (SERPS)*

Legislation to introduce SERPS was passed in 1975, and the scheme was implemented in 1978. It pays an additional pension that is linked to earnings. More precisely, it pays 25 percent of those earnings that lie between the lower and the upper earnings limit—two thresholds defined by the scheme. Benefits are calculated on the basis of the best 20 years of earnings and are payable together with the BSP.[6] The same retirement ages apply. Until 1988 the only possibility to opt out of SERPS and to pay reduced National Insurance Contributions (NIC) was if the employer provided an approved private pension. This was not at the discretion of the employee, but was decided by the employer, a fact that is important in our empirical analysis and is discussed in further detail in the following subsection.

C. *Opting Out*

When SERPS was introduced, the government aimed to design an opting-out scheme to ensure that existing secondary pensions provided by the private sector would not be crowded out by the new scheme. As a result, membership in SERPS was compulsory only for those employees whose employer did not offer an occupational pension. It was not at the worker's discretion whether to participate in the new government scheme. In this setup, which remained in place until 1987, we argue that membership in SERPS was not a choice variable for the employee, but an exogenous event. This implicitly assumes that people with different attitudes towards saving do not systematically select into jobs that provide private pension schemes.[7]

From 1988 onwards, this framework no longer applied. The new legislation ruled that an employee could choose to opt out of SERPS even if the employer did not offer a private occupational pension scheme. In this case the employee could choose to join a so-called "Approved Personal Pension," a form of individual retirement account that met the minimum criteria set out by the government. The new policy sought to increase further private pension coverage in the United Kingdom. During the first years, the government offered an extra 2-percent incentive rebate on top of the reduced rate of National Insurance Contributions that a worker who opted out would be entitled to.[8] An important implication of this change, especially in the context of this study, is that from 1988 onwards workers were able to choose themselves whether to opt out of SERPS. Enrollment in SERPS could no longer be considered an exogenous event for the worker after the date the new legislation was implemented. For this reason we limit our analysis to data before 1988.

[5] This projection assumes real earnings growth of 1.5 percent per year and does not take into account the recent off-the-rule increase of the Basic State Pension benefit in April 2001.

[6] These are the rules applying to the period we study. The 1986 reform announced changes to the computations for people retiring in April 1999 or thereafter.

[7] To our knowledge there is no firm evidence suggesting that people systematically choose jobs with private pension provision.

[8] This incentive rebate was reduced to 1 percent in April 1993 and only granted to those workers aged 30 and above.

D. *Variation in Pension Wealth and the Differential Impact of the Pension Reforms*

The U.K. pension system underwent many reforms. Our estimates of pension wealth reflect all the changes that were implemented during the period we consider. However, it is the variation in the pension wealth variables caused by the three most important and clear-cut reforms that will mainly drive the identification of our estimates in the empirical analysis. These three major reforms include two indexation changes of the BSP in 1975 and in 1981, and the introduction of SERPS in 1978. While the indexation changes decreased everybody's future entitlements, SERPS increased the public pension wealth of those individuals who were not covered by a private pension at the time and earned above the lower earnings limit.

The indexation changes in the BSP affected everybody in the population since this part of the public pension scheme was universal. The magnitudes depended mainly on the age of the head of household and on the age of the spouse in the case of married couples. While the impact of this reform varied mostly across the different date-of-birth cohorts, it also varied within these groups due to differences in household composition and age-pairs in the households.

The introduction of SERPS affected cohorts born after 1913 in the case of men (1918 for women), and targeted those workers who did not yet have any private pension coverage. Some people were not affected by this new scheme at all: older cohorts who were retired by the time the new scheme was implemented; those who already had a private pension; and those who have very low earnings. Among those who could take advantage of SERPS, the expected entitlements were generally higher for later cohorts. This is because people who were already close to retirement when SERPS came into being did not have many years remaining to accumulate entitlements. Furthermore, later cohorts usually have higher earnings due to economic growth and therefore acquire higher entitlements on average. Because SERPS is earnings-related there exists variation within cohorts due to differences in individual-level earnings. To exploit these differences, we classify our sample into four occupational groups that differ in the shape of their lifetime earning profiles.

In addition to the three important reforms in our period of study, there are a number of smaller changes in pension wealth. These were mostly due to minor updates in the information available to the individuals when forming an idea about their expectations. This concerns in particular observing the actual up-rated benefit values for the BSP and revaluation factors for SERPS which were published every year, and which individuals use to replace previously anticipated values with the actual realizations. Even though these are minor changes we incorporate them in our calculation of pension wealth.

The expected pension wealth variables vary mainly with age and occupation of the adult members of the households, and their opting-out status with respect to the earnings-related scheme. This is the variation we will use in our empirical exercise to identify the effect of pension wealth on savings.

II. Data Sources and Estimation of Pension Wealth

Our main source of information is the U.K. Family Expenditure Survey (FES), a time series of cross-sectional household surveys comprising roughly 7,000 households per year. As discussed above, we focus our analysis on the three most important and clear-cut reforms of the 1970's and 1980's, i.e., the introduction of SERPS and the two indexation changes of benefits from the BSP. We use observations spanning the years from 1974 until 1987, focusing on people born between 1909 and 1968.

We exclude from our sample households headed by self-employed individuals, whose income and consumption are often misreported. We also exclude composite households, that is households containing other adults than the head of household and the spouse. To reduce the influence of a few outliers, we trim the households that report income in the top and bottom 2 percent of the sample in each year. This leaves us with a sample of about 4,000 households per year.

In our data, saving is measured as the residual between disposable household income and total household expenditure, where the latter includes spending on durable goods. The information on both income and expenditure is of good quality, as documented in James Banks and Paul Johnson (1998). However, our definition of saving is not ideal, as, for instance, it excludes capital gains on real estate and financial assets.

In addition to detailed information on expenditure and family income which allows us to compute saving, the FES is also a rich source for variables like earnings, occupation, and several demographic variables. Combining this with our knowledge of the pension legislation in place in each year, we construct estimates of the present value of future pension benefits, net of future contributions.

A. *Pension Wealth*

The FES data do not provide information on subjective pension wealth expectations so that this variable has to be estimated. We assume that people have a reasonable understanding of the working of the scheme. Given that everyone is enrolled in the public pension system in some form, general knowledge should be quite good. We will compute the expected present value of net pension wealth, where pension wealth is defined as the sum of future benefits, assuming continued participation until retirement, minus future contributions, and use this measure as an estimate of perceived pension wealth.

For each household observed during the chosen period, we first approximate the present value of future benefits both from the BSP and from SERPS. For this purpose, we use information on age, sex, and marital status of the adult household members, and the relevant legislation in the year of observation. In a second step we compute the present value of current and future contributions to public pensions schemes. We deduct these contributions from the values obtained in the first step, that is from the total anticipated benefit receipts. The result is what we refer to as pension wealth at the time of observation.

There are many conceivable ways to compute pension entitlements. We refer the interested reader to Appendix A, where we discuss some conceptual issues involved in more detail. Here we shall only state the assumptions we make and the main steps we use in our computations.

- We use the entitlements that people will have acquired by the time they retire *according to the current legislation,* assuming they keep contributing to the schemes they are enrolled in just as they do at the time they are observed in the data. We take into account any reforms and future up-rating rules that have been legislated up to the time of observation. We assume that people expect the current legislation to persist.
- We net out current and future National Insurance Contributions (NICs).[9]
- We express all values in constant prices and we assume *perfect foresight about inflation rates* when computing future expected benefits.
- We assume that, when forming their expectations, people take their current *characteristics,* such as marital status and their (non) participation in SERPS as given and fixed.
- We account for *uncertainty about longevity* by applying survival probabilities to each period considered in the computations. The maximum attainable age, denoted T, is fixed at 100.
- We calculate lifetime earnings profiles that are needed to compute entitlements for SERPS separately for groups defined by cohort, occupation, and sex using earnings information from 32 years of cross sections from the FES and synthetic cohort techniques. We describe the details of this procedure, including the extrapolation over parts of the life cycle of each cohort not covered by the survey, in Appendix C.
- We assume that the age at which individuals expect to retire is the official state pension age of currently 65 and 60 for men and women, respectively.[10]

[9] Note that NICs do not only pay for pension benefits, but also provide the funds for other social security programs including unemployment, statutory sick and maternity pay, incapacity benefit, maternity allowance, widow's payment, widowed mother's allowance, and widow's pension. No fixed percentage is, however, assigned to these different "other" programs, and compared to the payment of public pension entitlements they are rather small.

[10] In face of the observation that on average individuals retire increasingly early, it is worth noting that receipt of public pension benefits is conditional on reaching statutory retirement age. Hence, early retirement trends affect public pension wealth only in an indirect way: early retirement will

We should stress that our computations reflect only the knowledge of the period when the expectation is formed. That is, for a household whose saving rate is observed at time t, we estimate its pension wealth using information on the current benefit rates, earnings growth factors, inflation, etc. Changes in the rules will be reflected in our calculation of pension wealth.

Given the assumptions above, the formula for the expected value of pension wealth at time t for a single individual i, $E_t(PW_{it})$, is defined as the present value of expected social security benefits reduced by the present value of current and future social security contributions:

$$(1)\quad E_t(PW_{it}) = \sum_{k=R_i}^{T} \frac{b_{itk} \cdot s_{tk}}{(1+r)^{k-t}} - \sum_{j=t}^{R_i-1} \frac{w_j \cdot c_j \cdot s_{tj}}{(1+r)^j}.$$

b_{itk} denotes the annual level of benefits expressed in constant prices that will be received at time k according to the legislation at time t. s_{tk} is the gender-specific probability in period t of surviving until year k conditional on having survived until period t, R_i is the expected date of first receipt of pension benefits, hence individual i's retirement date, r stands for the real interest rate, and T for the maximum attainable age. Contributions in any one period are calculated by multiplying the expected level of earnings at age j, w_j, with the applicable contribution rate in place at the time, c_j, taking into account survival probabilities.

The formula above applies to a single person. For married couples we take into account inheritance rules that are relevant for pension entitlements. This implies placing probabilities on the events of joint survival as well as single survival of the head of household and the spouse. Hence the formula for married couples becomes:

$$(2)\quad E_t(PW_{tc_i}) = \sum_{k=\min(R_1,R_2)}^{T} \frac{b_{c_i tk} \cdot s_{mtk} \cdot s_{ftk} + [(s_{ftk} - s_{mtk} \cdot s_{ftk}) + (s_{mtk} - s_{mtk} \cdot s_{ftk})] \cdot b_{s_i tk}}{(1+r)^{k-t}} - \sum_{j=t}^{R_m-1} \frac{w_{1j} \cdot c_t \cdot s_{mtj}}{(1+r)^j} - \sum_{l=t}^{R_f-1} \frac{w_{2l} \cdot c_t \cdot s_{ftl}}{(1+r)^l}$$

where $b_{c_i tk}$ is the benefit rate for couple i in period t to be paid in period k, $b_{s_i tk}$ the benefit rate that applies to a surviving single person i in period t, and T the maximum attainable age. The benefit calculations run from the year when the first member of the couple retires, $\min(R_1, R_2)$, while the contribution summations span from the current period of observation until the end of the working life of each partner. s_{mtk} is the probability in period t for a man to survive until year k, while s_{ftk} is the corresponding probability for a woman.[11] As in the case for singles, the information on survival probabilities for men and women is taken from English Life Tables. For the probability of joint survival we use the product of the individual survival probabilities in that we do not have available any joint life tables.

These formulas are straightforward to apply once the appropriate benefit rate is known. In the case of the BSP the current benefit rate is published by the Department of Social Security every year. For SERPS, the benefit rates have to be computed applying the contemporary set of rules.

B. *Cohort-Age Profiles of Pension Wealth*

Having computed pension wealth for the households observed in the FES, we now study its variation over time and in the cross section. The pension reforms during our sample period have a differential impact on individuals depending on birth cohort and (for SERPS) occupation groups. In this subsection, we focus mainly on differences across cohorts and over

not reduce entitlements to the BSP as long as the individual has contributed for the minimum number of years. In the case of SERPS, benefit computations were originally based on the best 20 years of earnings. As long as a worker has paid NICs for at least 20 years the impact of not working until age 65 will be minor.

[11] All the relevant quantities, such as the probabilities of survival, are adjusted for the actual age of the individuals in the households considered. No attempt is made, however, to take into account differential mortality by economic status.

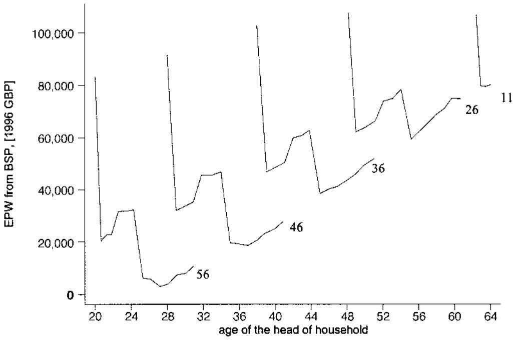

FIGURE 1. COHORT-AGE PROFILES FOR PENSION WEALTH FROM THE BSP, 1974–1987

time. The exact definition of year-of-birth cohorts is given in Appendix B. We should stress that the variation in pension wealth we illustrate here is not the only one we use in our regression analysis. To identify the parameters of interest, we also exploit the variation across occupation groups within cohorts. Groups were chosen to maximize variation in both the level of pension entitlements and in the effects of pension reforms on these entitlements.

We start by calculating cohort means for each year in the survey, where a cohort is defined by the year of birth of the head of household. We plot the cohort-age profiles of these means in Figure 1 for BSP and in Figure 2, panels (a) and (b), for SERPS. The expected value of pension wealth (EPW) is set in real terms expressed in 1996 prices and is discounted to the year in which the expectation is formed. The numbers can be interpreted as the lump sum that has the same present value as the person's future benefits net of future contributions. Note that the way in which we compute pension wealth, in the absence of legislative changes, EPW increases over the life cycle for three reasons. First, older individuals need to discount over fewer periods, so that the same entitlements are worth more in present value terms. Second, with every extra year that an individual survives, some uncertainty about mortality is resolved, again increasing the expected values. Third, our pension wealth computations net out future contributions to the public pension scheme so that each year fewer remaining contributions are deducted.

Two important features emerge from Figures 1 and 2. First, both pension wealth components vary a great deal as a consequence of the pension reforms we are considering. Second, this variation is very different across cohorts.

In Figure 1 the two marked downshifts in the BSP are due to the indexation changes: first in 1974–1975 from an ad hoc regime to earnings indexation of benefits, and in 1981 a further change towards price indexation. Table 1 gives an idea of the magnitudes involved, focusing on the two large indexing changes and showing the level as well as the percentage changes in pension wealth. The first indexation change of the mid-1970's led to an average decrease in EPW of 50 percent for all the cohorts considered here. However, the effects differ greatly across cohorts. The oldest cohort which was born between 1909 and 1913 suffered a reduction in EPW of 25 percent while the youngest cohorts lost around 80 percent of their expected wealth. The retirement period lies further in the future, and the slower growth rate of benefits created a much larger gap between previously expected benefits and adjusted expectations.

Table 1 displays the impact of the indexation change of the BSP in 1981, which linked benefit growth to prices. It caused an overall reduction of wealth in the flat-rate pension of 43 percent. The reductions range from 20 percent for those born between 1909 and 1913 to as much as 81 percent for the youngest cohort. Both indexation changes show, as expected, large variations in pension wealth loss depending on when in the life cycle the change happened.

The same kind of graphs for SERPS is displayed in Figure 2, where panel (a) refers to those households where only one adult is part of SERPS, and panel (b) refers to couples where both adults participate in SERPS.

In both panels we notice that the life-cycle experiences of the cohorts differ tremendously. The oldest cohort had just a few years to build up entitlements to the earnings-related scheme, so their EPW from SERPS is small [see cohort (11) in particular]. In addition, later cohorts have higher earnings as a result of economic growth and also additional female earnings, both of which lead to higher entitlements per household unit. As shown in Table 2 the importance of SERPS varied by cohort: in 1975 SERPS were only 0.4 percent of public pension wealth for the cohort of 1909–1913 whereas for

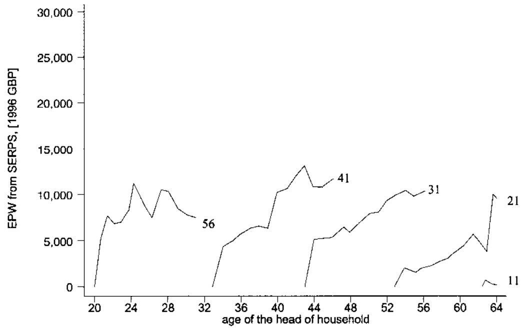

(a) Couples with one adult opted out, one in

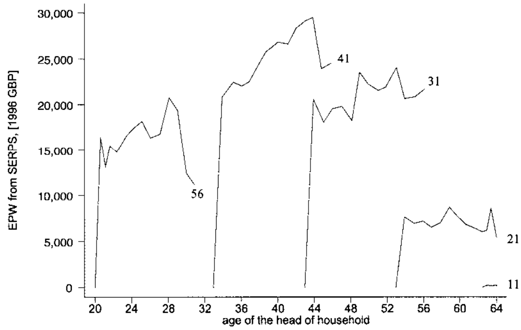

(b) Couples both in SERPS, singles in SERPS

Figure 2. Cohort-Age Profiles for Pension Wealth from SERPS, 1974–1987

Table 1—Expected Pension Wealth from the BSP (Amounts in 1996 British Pounds)

	Cohort (year of birth)										
Year	9–13	14–18	19–23	24–28	29–33	34–38	39–43	44–48	49–53	54–58	All
1974	106,723	111,578	108,864	107,442	105,580	102,823	99,514	91,497	87,146	83,200	101,927
1975	79,819	78,536	71,522	62,249	55,826	46,834	39,068	31,914	24,719	20,420[a]	50,492
Percent change	−25	−30	−34	−42	−47	−54	−61	−65	−72	−75[a]	−50
1980	n/a	85,576	87,738	78,483	70,302	62,940	54,746	46,806	39,247	32,257	58,461
1981	n/a	68,624	71,110	59,348	47,837	38,494	28,644	19,643	11,508	6,214	33,160
Percent change	n/a	−20	−19	−24	−32	−39	−48	−58	−71	−81	−43

Notes: All summary statistics in this table are derived from the same sample that we also use in our final estimation. Details were provided at the beginning of the data description in Section II.

[a] The number of observations in this cell is only 19, so that the numbers might give a distorted picture.

the cohort of 1954–1958 SERPS were 40 percent of total public pension wealth. These figures show again the importance of the timing of the policy changes.

While Figures 1 and 2 and Tables 1 and 2 show differences across cohorts, the data for EPW from both the BSP and from SERPS contain substantial variation within cohorts, reflecting occupation, opting-out status, marital status, and different age matches among couples. To give an idea of the extent of this variation, in Table 3 we display the fraction of total public pension wealth that the earnings-related scheme accounted for in 1975, but this time stratifying not only by cohort, but also by occupation of the head of household and opting-out status of the household members. We distinguish four occupation groups: professional employees, white collar workers, skilled workers and other occupations, and unoccupied.[12]

Table 3 shows large differences in the importance of SERPS according to SERPS enrollment. For example, in the cohort of 1949–1953 SERPS was about 59 percent of public pension

[12] "Unoccupied" means "not working." This status is self-reported and could be permanent in the case of, say, a housewife, or temporary in the case of an unemployed person. For this group the question arises how to estimate their earnings profiles if we do not observe their wages. Certainly some of these people will have worked at some point in their life spans. Nevertheless, they are most likely to belong to the low income groups in the population who mostly do not earn any entitlements to SERPS. We therefore set the EPW from SERPS to zero for individuals that belong to this group. Note that the spouse of such an individual can still be earning SERPS entitlements of any amount.

TABLE 2—EPW FROM SERPS BY COHORT: PDV IN 1996 BRITISH POUNDS, PERCENT OF TOTAL EPW IN 1975

	Cohort (year of birth)										
Year	9–13	14–18	19–23	24–28	29–33	34–38	39–43	44–48	49–53	54–58	All
PDV in 1975	349	1,787	4,869	8,828	11,881	13,315	12,790	13,594	14,114	13,669	10,231
Percent of total PW	0.4	2.2	6.4	12.4	17.5	22.1	24.7	29.9	36.3	40.1	16.8

wealth among professionals when both spouses were enrolled but just 25 percent of pension wealth when only one spouse was enrolled. We also see large variations in the importance of SERPS across occupation categories.

III. A Theoretical Framework and Its Empirical Implications

A. *A Simple Model*

The conceptual framework we use to investigate the relationship between pension wealth and household savings is the life-cycle model. It guides our choice of the econometric specification as well as the interpretation of the results. In this section, we present the simplest version of the model we can use to make the basic points about our econometric specification. We keep the model simple in order to allow for a closed-form solution, and so we do not explicitly consider uncertainty, changes in rates of return, labor supply, and many other important elements.

We analyze a four-period model so that we can study how a pension reform affects people of different ages and how a pension reform affects saving a number of years after the reform. We assume that in the first three periods of their life, individuals work and receive an exogenous income w_i, $i = 1, 2, 3$. In the last period they retire and receive pension benefits denoted by b. During retirement, in addition to their pension benefits, they can use their savings that they might have accumulated during the first three periods of their life. These savings are assumed to appreciate over time at an exogenous and fixed interest rate equal to r. The individuals in our model have a log-utility function, face no income or interest rate uncertainty, and have no bequest motive. Finally, they do not face liquidity constraints. The optimization problem takes the following form:

(3)

$$\max_{\{c_t, t=1,\dots,4\}} \sum_{t=1}^{4} \beta^{t-1} \log(c_t)$$

$$\text{s.t.} \sum_{t=1}^{4} \frac{c_t}{(1+r)^{t-1}} \le \sum_{t=1}^{3} \left(\frac{w_t}{(1+r)^{t-1}} \right) + \frac{b}{(1+r)^3}$$

where β is the factor by which future utility is discounted. Note that the lifetime budget constraint results from collapsing the budget constraints of the four single periods:

$$(4) \quad A_t = A_{t-1}(1+r) + w_t - c_t, \quad t = 1, 2, 3;$$

$$A_0 = A_4 = 0$$

where A_t is the amount of assets held at the end of period t. It is in the absence of liquidity constraints that these four single constraints can be collapsed into one.

Solving the maximization problem yields the optimal consumption levels for each period:

(5)

$$c_1 = \frac{1}{1+\beta+\beta^2+\beta^3} \times \left[w_1 + \frac{w_2}{1+r} + \frac{w_3}{(1+r)^2} + \frac{b}{(1+r)^3} \right]$$

$$c_2 = \frac{\beta}{1+\beta+\beta^2+\beta^3} \times \left[w_1(1+r) + w_2 + \frac{w_3}{(1+r)} + \frac{b}{(1+r)^2} \right]$$

TABLE 3—EPW FROM SERPS EXPRESSED AS A PERCENTAGE OF TOTAL PUBLIC PENSION WEALTH BY COHORT AND OCCUPATION IN THE YEAR OF LEGISLATION, 1975

	Cohort (year of birth)									
	9–13		14–18		19–23		24–28		29–33	
	SERPS enrollment									
Occupation	Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in
Professional	n/a	0.7	1.2	6.5	3.2	19.0	8.9	28.0	10.0	40.4
White collar	0.8	0.5	1.0	4.9	2.7	14.1	6.1	22.9	8.9	29.5
Skilled & other	0.8	0.2	0.9	2.9	2.1	8.6	3.9	14.4	6.5	21.5
Unoccupied etc.	0.0	0.1	0.0	0.2	0.0	0.5	0.0	0.3	n/a	4.4
All	0.7	0.3	1.0	3.2	2.6	10.3	5.8	19.0	8.3	27.3

$$= \frac{1}{1 + \beta + \beta_2} \times \left[A_1(1 + r) + w_2 + \frac{w_3}{(1 + r)} + \frac{b}{(1 + r)^2} \right]$$

$$c_3 = \frac{\beta^2}{1 + \beta + \beta^2 + \beta^3} \times \left[w_1(1 + r)^2 + w_2(1 + r) + w_3 + \frac{b}{(1 + r)} \right]$$

$$= \frac{\beta}{1 + \beta + \beta^2} \times \left[A_1(1 + r)^2 + w_2(1 + r) + w_3 + \frac{b}{(1 + r)} \right].$$

Consumption in period 4 will equal the total amount of remaining resources. Notice that for c_2 we present more than one expression. The first expression displays the solution as seen from period one. The second expression shows the solution as a function of current and future resources as seen from the perspective of period two. Similarly, c_3 is written to depend on resources from the perspective of periods one, two, and three. As long as no changes occur during the four periods that were not anticipated in period one, the different expressions for c_2 and c_3 are identical. This can be verified by substituting the intertemporal budget constraint in (5). However, should there be an unexpected change in period t in some of the determinants of the solution the equivalence breaks down: saving decisions taken before that date cannot be changed, and the individual has to reoptimize over resources remaining from previous periods (A_{t-1}) and future income. To evaluate the effect of an unexpected change in pension benefits occurring some time in period one on consumption (and saving) for an individual aged 2, one can use the second expression for c_2, which takes the asset level at the beginning of the period as given and determined by decisions in the first period. For an individual aged 3 at the time of reform, we will use the third expression for c_3.

The simple point we want to make is that while an increase in the exogenous level of benefits will decrease savings and saving rates during working life, the effect will be different for individuals at different stages of their life cycle. The impact depends on when during the individual's lifetime the household learned about the change because it matters over how many remaining periods the individual can distribute the readjustment. This corresponds to the adjustment discussed in Gale (1998). Furthermore, the magnitude of the effect we observe depends on how long the reform occurred before the period of observation. The combination of these two factors determines the impact of a change in pension wealth on an individuals' consumption and savings choices in any one period.

TABLE 3—*Continued.*

Cohort (year of birth)											
34–38		39–43		44–48		49–53		54–58		All	
SERPS enrollment											
Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in	Couples one in, one out	Couples both in, single in
12.0	47.1	14.2	48.1	16.4	52.7	24.9	58.6	n/a	68.4	11.5	38.4
10.6	36.4	8.6	39.6	13.9	42.7	25.0	49.0	23.2	53.4	7.6	30.5
10.6	25.5	8.8	27.9	15.6	33.8	15.1	32.4	15.1	37.4	6.1	17.9
0.0	1.5	0.0	3.4	0.0	4.67	0.0	1.7	0.0	0.0	0.0	1.0
10.8	32.2	9.7	35.5	14.7	39.8	23.0	41.8	20.0	44.8	7.7	24.2

Note: n/a—no observations in that cell.

B. *Empirical Specification*

Rather than basing our empirical analysis on consumption as in equation (5), we will base it on saving rates, as this formulation is likely to lead to a better error structure. From equation (5) one can derive an expression for saving rates as a function of pension wealth. For example, the saving rate in period 2 is

$$(6)\quad \frac{y_2 - c_2}{y_2} = 1 - \Psi \cdot \frac{b}{(1+r)^2 y_2} - \Psi \cdot \frac{1}{y_2}\left[w_1(1+r) + w_2 + \frac{w_3}{1+r}\right]$$

where Ψ is a function of β given by one of the two expressions on the right-hand side of equation (5), depending on when we observe the household in the survey. For example, if an age 2 household is observed in a reform year, we would use the second expression for c_2 in equation (5), while if it is observed the year after a reform we would use the first expression. Given β, r, and future pension entitlements b, one can calculate Ψ and $\Psi \cdot (b/((1 + r)^2 y_2))$. It is the latter expression, generalized to an N-period problem, that we enter as a regressor in our empirical specification for saving rates.

The empirical specification generalized to a multiperiod framework is

$$(7)\quad SR_{it} = \mathbf{X}_{\mathbf{it}}\theta + \gamma\left(\Psi(t_i, tr_i) \cdot \frac{EPW_{it}}{y_{it}}\right) + \varepsilon_{it}$$

where SR_{it} is the saving rate of household i observed at time t, EPW_{it} is the present value of (earned) expected pension wealth (whose computation was described in Section III), and y_{it} is current income. $\Psi(t_i, tr_i)$ is the normalization factor, equivalent to Ψ in equation (6), by which we multiply each household's expected pension wealth. In a multiperiod model with log-utility, $\Psi(t_i, tr_i)$ is given by the expression:[13]

$$(8)\quad \Psi(t_i, tr_i) = \frac{1-\beta}{1-\beta^{E(T)-tr_i}} \cdot \beta^{t_i - tr_i - 1}$$

where t_i denotes the t-th period in household i's life cycle, and tr_i the number of the life-cycle period in which the last reform experienced by household i occurred. t runs from 1, when the household begins working life, until $E(T)$ (set to 56 in our case), the number of the last period that households expect to reach on average. The vector $\mathbf{X}$ in equation (7) represents a number of control variables, including group and time effects, that are meant to capture various

[13] The derivation of $\Psi(t_i, tr_i)$ for a multiperiod model, which is clearly related to the adjustment factor in Gale (1998), is included in Appendix D.

determinants of saving other than pension wealth. The groups are defined by year of birth, occupation, and SERPS enrollment.

Given the structure of the model, it would be desirable to include a measure of household wealth, but our data do not contain information on this variable. We can proxy household wealth by a flexible (group-specific) function of age. As our groups are defined (among other things) by year-of-birth cohort, including indicator variables for group and time is equivalent to considering a categorical variable in age. A similar argument applies to including the present discounted value of future earnings in the regression, which also appears in the equation for saving derived from the structural model. In some of the specifications we tried, we included estimates of future earnings in **X** and applied the same adjustment as to pension wealth variables.

A literal interpretation of our model, in which pension wealth is a perfect substitute of financial wealth, would imply that the coefficient γ in equation (7) is -1.0. In what follows we estimate γ and interpret it as a measure of the substitutability between financial and pension wealth. Our coefficient γ is therefore equivalent to the coefficient that Feldstein (1974), King and Dicks-Mireaux (1982) and many other authors have tried to estimate on micro- and macrodata.[14] To allow for more flexibility in our specification, we let the coefficient γ vary with age, reflecting the fact that the degree of substitutability between financial and pension wealth might be different for individuals at different points of their life cycle. After experimenting with various polynomial specifications, we decided to let the coefficient be a step function of age. The results using polynomials were similar and are available upon request.

As we described in Section I, the public pension system in the United Kingdom consists of two different tiers, the BSP and SERPS. These two components vary substantially in their structure. Therefore, we let γ take different values for each of these components and estimate the following relationship:

$$(9)\quad SR_{it} = \mathbf{X_{it}}\theta + \gamma_1\left(\Psi(t_i, tr_i^{SERPS}) \cdot \frac{EPW_{it}^{SERPS}}{y_{it}}\right) + \gamma_2\left(\Psi(t_i, tr_i^{BSP}) \cdot \frac{EPW_{it}^{BSP}}{y_{it}}\right) + \varepsilon_{it}.$$

C. *Preliminary Evidence on the Relationship Between Pension Wealth and Saving*

Before taking equation (9) to the data or considering slightly more complicated models that allow for different degrees of substitutability between financial and pension wealth for different ages, we investigate whether we can find any evidence in the raw data of a relationship between EPW and saving rates and whether such evidence is consistent with our theoretical approach. We focus on the two episodes that induced the largest changes in EPW: the introduction of SERPS which coincided with the first change in the indexation of the Basic State Pension in 1975 and the next indexation change of the Basic State Pension in 1981.

In order to control in the simplest possible way for aggregate shocks and group differences without imposing any further structure, we regress saving rates on pension wealth adding as controls group and time dummies. We define groups in the same way as in Section II, that is on the basis of cohort, occupation, and enrollment in SERPS. Notice that, as we condition on year-of-birth cohort and time, we also implicitly control for age.

The results we obtain are promising: the regression of individual saving rates on the PV of EPW as well as group and year dummies yields a negative relationship that is significant at the 1-percent level. Carrying out the same exercise for the level of consumption, the choice variable of our theoretical model, leads to equivalent conclusions; it provides evidence of a positive relationship between consumption and EPW which is also significant at the 1-percent level. We report the estimated coefficients on EPW in Table 4.

[14] Feldstein studied the time-series relationship between aggregate saving rates and pension wealth. Dicks-Mireaux studied the relationship between the stock of financial wealth and pension wealth in microdata.

TABLE 4—REGRESSION OF SAVING RATES AND CONSUMPTION ON EPW, CONDITIONING ON GROUP AND TIME EFFECTS, CONSUMPTION, AND EPW IN THOUSAND POUNDS

Variable	Coefficient on EPW	p-value
Saving rate	−0.00087 (0.00017)	0.000
Consumption	0.08347 (0.00234)	0.000

Note: Standard errors are in parentheses.

The coefficients in Table 4 give the population average of the observed response to the changes in EPW, conditioning on time and group effects. On average an increase in EPW of £1,000 leads to an increase in annual consumption spending of 80 pounds, or an average increase of 0.08 percentage points in households' saving rates.

Having checked the existence of a simple relationship between saving (or consumption) and pension wealth, we move on to estimating more structural models based on equation (9). This analysis will allow us to estimate the degree of substitutability between financial and pension wealth, which cannot be inferred from the numbers in Table 4.

IV. Results

A. *Identification: Using the Differential Effects of the Pension Reforms*

There are several reasons why ordinary least-squares (OLS) estimation of equation (9) on household data would yield biased and inconsistent estimates. First, the subjective expected value of pension wealth is unlikely to be equal to actual pension wealth; furthermore, we only have an approximation to pension wealth. Therefore, we need to allow for the presence of measurement error. Moreover, it is possible that unobserved heterogeneity in the taste for savings is related to the individual stock of (pension) wealth. For these reasons we take an instrumental variable approach and use the pension reforms to identify a number of instruments. As the effects of these reforms were different for different groups, we use the interaction of group dummies and year dummies as instruments for our measure of pension wealth. The group indicators distinguish between cohorts, occupations, and enrollment in SERPS. The use of such a differences-in-differences estimation strategy is legitimate if two conditions are satisfied: once we control for group and year dummies, their interaction does not enter the equation for saving in its own right;and pension wealth has variation over and above that captured by group and year dummies.

The first assumption is an identification assumption and therefore is not testable. We can test the second by finding whether pension reforms affect different groups in different ways. The analysis presented in Section III suggests that this is the case, and an F-test for the significance of the interactions between year and group dummies shows this formally. That is we reject at any sensible level of significance the hypothesis that the variation in pension wealth is fully explained by time and group effects. Detailed results are available upon request.

B. *Empirical Specifications and Results*

The empirical specification we use is slightly different from the one in equation (9), in that we allow γ_1 and γ_2 to vary with age. The rationale for this is that our model omits aspects of reality that are bound to have an impact on the substitutability between pension wealth and other assets. While these simplifications kept our theoretical framework analytically tractable, there are reasons to believe that the direct implementation of equation (9) would be too restrictive. For example we do not model liquidity constraints, which are likely to be more important among the younger population and are bound to affect the elasticity we are trying to estimate. Therefore, the specification we use allows the coefficients on (adjusted) pension wealth to vary with age.

Equation (9) includes, in addition to future pension wealth, future earnings. There are two possible approaches to deal with this variable. The first is to estimate it using a methodology similar to the one used for EPW. Alternatively, as the present discounted value of future earnings is a function of age, cohort, and occupation, we can control for such a variable using

TABLE 5—ESTIMATES OF THE SUBSTITUTABILITY OF PRIVATE SAVING FOR EPW FROM SERPS AND BSP

Age group	SERPS	BSP
20–31	0.0135	−0.3061
	(0.334)	(0.133)
32–42	−0.5472	0.0060
	(0.277)	(0.139)
43–53	−0.6511	0.0432
	(0.269)	(0.087)
54–64	−0.7487	0.0351
	(0.243)	(0.040)

Note: Standard errors are in parentheses.

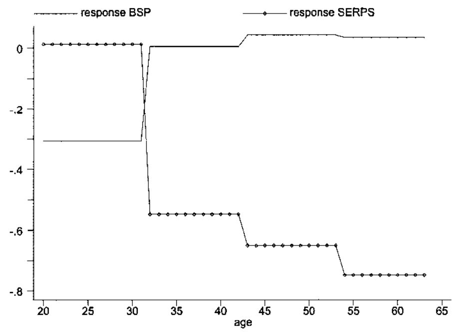

FIGURE 3. ESTIMATED AGE PROFILES FOR THE EFFECTS OF SERPS AND BSP

group and age effects.[15] We have implemented both strategies and found that the results for the effect of pension wealth are not affected by this choice. However, we prefer the second strategy because, in contrast to the EPW variable, we have no exogenous variation that we could use to instrument future earnings in order to deal with measurement error and correlation with unobserved heterogeneity.

In the pension wealth calculations and in the computation of Ψ, we need to make an assumption about the discount factor β and the interest rate r. We tried several combinations of values. Those we report are for $r = 0.03$, and $\beta = 0.98$. The results did not change substantially with different values of the discount factor and the interest rate. We use annual data from 1974–1987.[16]

In Table 5 we report the results from estimating equation (9). While our regressions include time and group dummies and some demographic variables, we only report the coefficients on pension wealth from SERPS and from the BSP for different age groups. These coefficients are estimates of the degree of substitutability between the two types of pension wealth and other wealth holdings. As such they answer the question of the following underlying thought experiment: If pension wealth from SERPS (BSP) were increased by 1 pound, what fraction of this amount would result in a reduction in other wealth and, eventually, show up as (appropriately discounted) consumption? We plot these coefficients in Figure 3.

For SERPS, we see that for individuals above the age of 31, pension wealth is a good substitute for financial savings. Indeed, for the individuals in the top age group, the coefficient is as high as −0.75, implying that for an increase in SERPS pension wealth of 100 pounds, they would decrease nonpension wealth by 75 pounds. For the youngest individuals the estimated coefficient is not significantly different from zero, indicating no substitutability between future pension benefits and financial savings. This result might be explained by binding liquidity constraints for the youngest individuals.

For BSP, we only find a significant effect for the youngest group, and even for this group the effect is relatively small (−0.3). For all the other age groups, the effect is not significantly different from zero.

These results are quite robust: we have tried a number of different specifications, changing the controls in the regression and the sample over which we estimate it. In particular, we have added estimates of future earnings, age polynomials to capture differences in the age effect within cohorts, and demographic variables. All these experiments, whose detailed results are available on request, confirmed the basic pattern of age-specific coefficients reported in Table 5.

[15] As noted earlier, by controlling for group and time effects we are implicitly controlling for age, because one dimension of the group definitions is date-of-birth cohort.

[16] When we introduce future earnings we use the same discount rate as for future pension benefits, that is, a real rate of 3 percent.

C. *Interpretation of the Regression Results*

Our empirical results indicate that SERPS and financial wealth are good substitutes for all but the youngest individuals. This finding is consistent with the life-cycle model. The absence of an effect for individuals younger than 32 may be an indication of binding liquidity constraints for this group.

Our estimates are substantially higher than those found in various empirical studies, such as King and Dicks-Mireaux (1982), Brugiavini (1987), Jappelli (1995), or Gustman and Steinmeier (1999), and they are more in line with the evidence presented in Gale (1998). This result might be an indication of the fact that the last set of authors failed to correct for the differential effects of pension wealth by age and/or of the importance of measurement error and simultaneity biases. Our methodology addresses both of these issues. Furthermore, we exploit exogenous variation induced by policy reforms. Disney et al. (2001) discuss the likely effects of the pension reforms of the late 1980's and 1990's on saving behavior in the United Kingdom. While they do not use a formal econometric analysis, they argue for the degree of substitutability between (private) pension wealth and financial wealth to be close to the one we estimate.

The evidence on BSP is quite different, as we do not find any significant effect of changes in this component of pension wealth on saving rates, except for the youngest individuals. One possible explanation is that the variation we use to identify the coefficient of interest is exclusively induced by changes in the indexation rules. Even though the effects proved quite substantial, as shown in Section III, maybe people did not fully understand the implications at the time the reform was implemented. While the SERPS legislation was the outcome of a public debate that lasted many years, the indexation changes to the BSP were quite different in nature: in 1975 explicit up-rating rules took the place of implicit ones, and in 1981 price indexation was discussed to be a temporary measure at first, but was then turned into a permanent one.

Another possible explanation for the difference between the effect of changes in SERPS and BSP might be due to the fact that the poorer part of the population does not earn pension wealth from SERPS, while still being entitled to the BSP; SERPS have a bottom threshold for earning benefits, but BSP does not.[17] Substitution of pension wealth for private wealth in the poorer part of the population is likely to be small due to liquidity constraints. Differential enrollment combined with a lack of substitution would lead qualitatively to our results.

V. Conclusions

In this paper we use an estimate of pension wealth from the public pension scheme to investigate its impact on household saving behavior. We focused on the time period that encompassed three major reforms, the introduction of SERPS in 1978, which was legislated in 1975, and the two indexation changes of the BSP in 1975 and 1981. The pension wealth profiles show substantial differences across and within cohorts and occupation groups. The variation induced by these reforms combined with implications from economic theory allowed us to analyze the link between pension wealth and household savings rates in a way that goes beyond a simple differences-in-differences strategy.

The results we obtain indicate a considerable degree of substitutability of SERPS for financial wealth: for example, among 43–53-year-olds we estimate an elasticity of substitution of about -0.65, and -0.75 for the 54–64-year-olds. However we estimate little, if any, substitutability between the BSP and private wealth.

Our results have important implications. They confirm qualitatively Feldstein's (1974) findings, which were based on time-series data. We find that, once we allow the effect of pension wealth to be age dependent, for large fractions of the population, the substitutability between pension and financial wealth is relatively high. This result is in accordance with the basic prediction of the life-cycle model. However, it does not hold for the youngest consumers, who are likely to be affected by liquidity constraints, and for the Basic State Pension.

[17] As discussed in Dilnot et al. (1994) the poorer part of the population tends not to have either SERPS or a private pension, and relies exclusively on the BSP for retirement income.

The implications of our results are also important for the debate on pension reform and on the adequacy of individual savings: should public pensions be reduced, households will increase their savings to make up for a large part of the loss. There is, however, one important modification to these conclusions: the results might not extend to the poorer part of the population since they were derived from the earnings-related tier of the public pension scheme. Poor households generally do not earn any entitlements from SERPS, so for them the relationship between pension wealth and savings might look different. In this sense, our results are in accordance with the studies that have expressed concern that the reforms to the U.K. pension system during the 1980's and 1990's have left some (relatively small) sectors of the population without sufficient provisions for old age.

Appendix A: Construction of the Variable "Expected Pension Wealth" (EPW)

In the absence of survey data for expectations of PW we will use an estimate of actual anticipated entitlements.

A. Discussion of Some Conceptual Issues

There exist many conceivable ways of measuring expected pension wealth. Obviously, in order to determine how to proceed, a number of choices have to be made each of which has important implications for the final results. Naturally, the choices we make will be driven either by considerations of feasibility or by coherence with the ultimate purpose of this study. What this means in detail should become clear in the course of the discussion.

In face of the ultimate purpose of our computations, that is to investigate the impact of pension wealth on people's saving decisions, one would ideally want to enter subjective expectations of pension wealth into the analysis. These perceived pension rights are likely to differ from the true values. The 1989 Retirement Survey has collected this information. However, it only covers very few cohorts and with only one observation we are not in the position to follow individuals over time to identify our parameter of interest. Hence, we need to find a concept that is likely to come close to these subjective expectations, but can be systematically computed from individual characteristics.

Anticipated Versus Acquired Rights.—In the case of *anticipated rights,* individuals form expectations about the benefits that they should receive during retirement if they continue to contribute to the pension schemes as they currently do. The discounted values of these expected benefit flows constitute the anticipated rights that might influence—among other things—individual saving behavior. Another possibility would be that individuals rather consider those pension rights as important that they already have acquired instead of basing their saving decisions on some expectation about the usually distant future. It is worth laying out what this roughly means in the case of the U.K. public pension scheme. For the case of the flat-rate part of the public pension scheme: as long as contribution records are not completed, reduced rates apply. Given that almost all male workers end up receiving the full rate it is unlikely that they would work out their currently accrued rights. The portion of women entitled to a full pension in their own rights has increased over the past two decades, especially since the recognition of "Home Responsibility Periods" was introduced.[18] Nevertheless, also for them working out reduced rates seems an unlikely approach. People will hence rather have these full amounts on mind instead of working out the fraction that they are so far entitled to, in some period before retirement.

For the earnings-related tier of the public pension scheme (SERPS), the arguments look similar. Especially if we believe that people have at least a rough idea about their needs in retirement then in order to determine how much extra saving they should involve in, apart from the compulsory amounts, they have to have some expectations of what they will get from the public pension schemes in total.

Therefore, we base our computations on the notion of *anticipated rights.*

[18] Home Responsibility Periods (HRP) were introduced in 1978. This change is also taken into account in our computations of EPW.

To predict future benefit flows people also need to make assumptions about the future path of a number of factors involved in the concept of expected pension wealth.

Pension Reforms.—One central issue is the evolution of the pension scheme itself and to what extent people anticipate changes to the system. Static expectations where individuals take the pension scheme "as is," ignoring the possibility of reforms, seems extremely rigid at first sight; particularly when considering the "tradition" of pension reform in the United Kingdom over the past decades. Nevertheless, the look at the system's history reveals at the same time that the outcome of reforms strongly depends on which part is in government in a way that even anticipating the qualitative direction of the system in the long run might be difficult if not impossible. That is why any particular assumption about the future evolution of the system would be arbitrary; therefore, the static assumption may be simplistic, but a prudent one.

Concerning reforms that have been passed into law but have not been implemented yet, we assume that people integrate these changes in their expectations. This seems a reasonable assumption since this type of legislation is preceded by a fair amount of public debate.

Demographics and other Individual Information.—Some individual characteristics that matter for the calculation of pension benefits might change over time. Consider the example of a single male worker in his mid-30's. The magnitude of his pension benefits depends on the outcomes of a number of variables:

- potential spells of unemployment and the entailed interruption NIC payments,
- possible change in expected earnings due to job change,
- change in marital status,
- life expectancy,
- ... .

A way of dealing with these uncertainties would be to place probabilities on the various possible outcomes, ideally cohort- and occupation-specific probabilities, but this is beyond what we are doing in this study. For marital status we assume that people form their expectations on the basis of their current marital status ignoring the possibility that it might change, i.e., we use the information recorded in the FES at the time of observation. We adopt the same strategy also for employment status and occupation.

B. *The Actual Pension Scheme: A Feasible Set of Simplified Rules*[19]

Computation of Entitlements to the BSP.—Given that the BSP is a flat-rate scheme there are not many different possible outcomes that need to be considered: either an individual meets the requirements to claim a full-rate pension, or—if her contribution records are incomplete—a reduced rate applies. Should a woman not qualify for a pension in her own right then her husband can claim a dependant's addition for her on his contribution records.

It is hard to know who is likely to be able to apply for a full pension at retirement age on a single observation basis as one would need contribution histories. It will therefore be necessary to base the computations on simplifying assumptions.

The simplifications listed below yield a reasonable approximation of the actual portions of full rates and married couples' rates that are paid out in reality:

- Men always expect to complete their contribution records and to receive the full-rate BSP.
- Single women tend to have a similar working life as single men, and therefore we assume that they also acquire the entitlements for a full-rate BSP.
- Married women who retired before 1978 never acquired entitlements to a full-rate pension—mainly due to breaks in their working lives for childbearing or because they chose to pay reduced National Insurance Contributions and to forgo pension entitlements in their own rights. They therefore usually had to rely on the dependants' rate which they

[19] *Tolley's Social Security & State Benefits Handbook* (various years) describes the rules governing the U.K. public pension scheme in great detail.

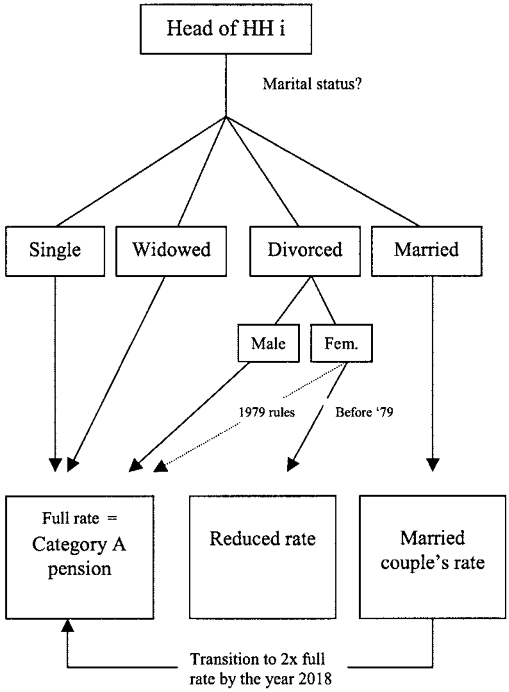

The simplified structure presented in the figure above, indicates that there are four different scenarios to be distinguished that lead to the following three outcomes:

1. Full rate (singles, widow(er)s, male divorced)
2. Reduced rate (female divorced)
3. Married couple's rate

FIGURE A1. ENTITLEMENTS TO THE BSP: A SIMPLIFIED SCHEME

could claim on their husband's contributions; that is the rate I apply in this case.

- Women who retire in 2018 or later will all be entitled to a full-rate pension on their own rights due to the recognition of home responsibility periods (HRP) as contribution years. Note that by 2018 the 1978 HRP rules will apply to the entire working life of all women.

Women retiring between the above bounds (i.e., between 1978 and 2018) increasingly profit from the HRP rules. To what extent this is the case for each single woman is, however, impossible to know. We therefore suggest the following approximation: the dependants' rate amounts to 60 percent of the full-rate BSP. The transition period lasts at most 40 years; hence, add 1 percent of the full rate to the dependants' rate each year to take account of the increasing recognition of HRPs.

Based on these simplifications a rather simple situation for the BSPs emerges that is illustrated in Figure A1.

We addressed the difficulty of dealing with reduced rates earlier, i.e., that there is no way of knowing for how many years an individual has actually contributed to the scheme and without that information we are unable to calculate the appropriate reduced rate. Furthermore, the FES does not distinguish singles from divorced people until 1978. As a consequence, we are left with only two outcomes, that is the full rate and the married couple's rate, not counting extra the transitional outcomes for married couples between 1979 and 2018.

Computation of SERPS Entitlements.—To work out the benefit formula for SERPS below, lifetime earnings profiles are required. Their estimation is explained in Appendix C.

The SERPS formula:

$$\text{(A1)} \quad b_{SERPS} = \sum_{t=1978}^{R} \left(W_t \frac{Y_{R-1}}{Y_t} - LEL_{R-1} \right) x_{Rt}$$

$$\text{if } LEL_t \leq W_t$$

$$\text{and } W_t = UEL_t \text{ if } W_t > UEL_t$$

Wt—weekly gross earnings at time *t*

LEL—Lower Earnings Limit

UEL—Upper Earnings Limit

R—retirement age

x_{Rt}—accrual factor.

To calculate the weekly SERPS entitlements for each relevant tax year the individual weekly earnings are revalued to the year prior to retirement age. The Department of Social Security publishes these revaluation factors every year which are captured by Y_{R-1}/Y_t above. These are supposed to preserve the value of "past" amounts along with average earnings growth. From the resulting amount the lower earnings limit of that very preretirement year is deducted to obtain the so-called "excess earnings." These are multiplied with the accrual factor as x_{Rt}, that was originally fixed at 1/80 for each year of service. Summing up the calculated amounts for the 20 best years of earnings during which the individual paid National Insurance Contributions will yield the weekly benefit that an individual earned in this scheme. The accrual factor as well as the calculation of excess earnings was changed in subsequent reforms reducing the generosity of the scheme.

Appendix B: Data Details and Cohort Definitions

We use a sequence of cross sections from the Family Expenditure Survey spanning the period 1974 to 1987 (see Table B1). Given the purpose of our study, we exclude households whose main earner is self-employed or reported to be retired at the date of observation. We further exclude composite households.

This leaves us with roughly 4,000 household observations per sample year.

Table B1—Cohort Definition

Cohort number[a]	Date of birth	Years observed in the sample
11	1909–1913	1974–1977
16	1914–1918	1974–1982
21	1919–1923	1974–1987
26	1924–1928	1974–1987
31	1929–1933	1974–1987
36	1934–1938	1974–1987
41	1939–1943	1974–1987
46	1944–1948	1974–1987
51	1949–1953	1974–1987
56	1954–1958	1974–1987
61	1959–1963	1979–1987
66	1964–1968	1984–1987

[a] Cohort numbers reflect the midpoint of the five-year interval spanned by each cohort group.

Appendix C: Earnings Profiles

The lifetime earnings profiles are required to work out SERPS entitlements. We construct these for each sex and occupation group of different cohorts.

Basic Strategy.

- Use the 32 years of FES data to compute average cohort earnings for the years where each sex/occupation/cohort group is observed. We distinguish four different occupation groups: "professional," "white collar," "skilled and other occupations," "rest" (mainly the unoccupied).[20]
- Despite the length of the time series of FES surveys, each of these groups still exhibit

[20] Note that the last group mainly consists of zero- or close to zero-earners. Being unoccupied is usually a transitory state for heads of households and hence also not a representative state in the life cycle of a head of household. Most people included in this last group are individuals on very low incomes who mostly do not earn any or only very small entitlements to SERPS. Hence we do not estimate any earnings profiles for this group and set their SERPS entitlements to zero.

some missing values in its lifetime profile. These are extrapolated using the approach suggested in Attanasio and Banks (1998) where they construct lifetime profiles of various variables by using the information of adjacent cohorts that are observed at ages where another cohort in question is not. The missing parts in a profile are hence extrapolated by applying a weighted average of the differences observed for other cohorts. The resulting profiles are smoothed using a Kernel estimator.

Sample for Estimating Earnings Profiles.—The earnings profiles will be used to compute the SERPS entitlements of individuals. Hence the question arises whether the earnings profiles should only be estimated on the sample of those to whom these earnings will actually be assigned in the pension wealth computations (i.e., people contracted-in to the second tier of the public pension scheme), or should we rather use the entire population.

We decide to use the entire population for two main reasons. First, we are interested in lifetime profiles, i.e., we also have to work out earnings observations for the years before 1978 when SERPS did not exist and opting out did not exist. We could of course use private pension contributions as a proxy. This strategy would, however, lead to considerably reduced cell sizes.

Another detail is related to employment status and how to deal with unemployed individuals. Some of these may be unoccupied throughout their lives while others might be unemployed only in the period we observe them, but are participating in the labor force otherwise. We include all zero-income individuals in the sample on which we estimate the average earnings profiles so that the portion of zero earners is well taken into account. Our approach is guided by the desire to obtain the correct cohort average of lifetime earnings.

Appendix D: Derivation of $\Psi(t_i, tr_i)$

In the derivation we will use the following:

$$\sum_{j=0}^{n-1} r^j = \frac{1 - r^n}{1 - r}. \tag{D1}$$

The starting point is equation (5). In the case of no reforms the generalized term for Ψ at time t follows from considering the factor in front of the parenthesis in the first expressions for each time period in equation (5):

$$\Psi(t, tr = 0) = \frac{1}{1 + \beta + \beta^2 + \ldots + \beta^{T-1}} \cdot \beta^{t-1}. \tag{D2}$$

Using (D1) this can be rewritten as:

$$\Psi(t, tr = 0) = \frac{1 - \beta}{1 - \beta^T} \cdot \beta^{t-1}. \tag{D3}$$

To generalize further to accommodate the occurrence of a reform at time tr, let us assume that in the four-period model a reform took place at the end of period one. The consumer needs to reoptimize her choices over the remaining periods two to four. In equation (5), the second expressions for periods two and three apply. The following pattern for the factor in front of the parenthesis emerges for any time period t after the reform in tr:

$$\Psi(t, tr) = \frac{1}{1 + \beta + \beta^2 + \ldots + \beta^{T-tr-1}} \cdot \beta^{t-tr-1} \tag{D4}$$

for $t > tr$.

Using (D1) this can be rewritten as:

$$\Psi(t, tr) = \frac{1 - \beta}{1 - \beta^{T-tr}} \cdot \beta^{t-tr-1}. \tag{D5}$$

Replacing T, the maximum attainable period in a person's life cycle with the expectation, and indexing the time periods to reflect that individuals experience a reform in a given year at different stages of their life cycle, yields the generalized form of the adjustment factor $\Psi(t_i, tr_i)$ stated in equation (8) and that we use in our analysis. Note that in the case of multiple reforms over time, the value of tr in the expressions for the adjustment factor is determined by the timing of the most recent reform.

REFERENCES

Attanasio, Orazio P. and Banks, James. "Trends in Household Saving: A Tale of Two Countries." Working Paper No. 98/15, Institute for Fiscal Studies, 1998.

Attanasio, Orazio P. and Brugiavini, Agar. "Social Security and Households' Saving." *Quarterly Journal of Economics*, 2003 (forthcoming).

Banks, James and Johnson, Paul, eds. *How reliable is the family expenditure survey? Trends in incomes and expenditure over time*. London: Institute for Fiscal Studies, 1998.

Bernheim, B. Douglas. "Rethinking Saving Incentives," in A. A. Auerbach, ed., *Fiscal policy: Lessons from economic research.* Cambridge, MA: MIT Press, 1997, pp. 259–313.

Brugiavini, Agar. "Empirical Evidence on Wealth Accumulation and the Effects of Pension Wealth: An Application to Italian Cross Section Data." Financial Markets Group, D.P. 20, LSE, London, 1987.

Diamond, Peter A. and Hausman, Jerry A. "Individual Retirement and Savings Behavior." *Journal of Public Economics*, February/March 1984, *23*(1–2), pp. 81–114.

Dilnot, Andrew; Disney, Richard; Johnson, Paul and Whitehouse, Edward. *Pension policy in the U.K.: An economic analysis.* London: Institute for Fiscal Studies, May 1994.

Disney, Richard F.; Emmerson, Carl G. and Tanner, Sarah. "Partnership in Pensions: An Assessment." London: Institute for Fiscal Studies, Commentary No. 78, 1999.

Disney, Richard F.; Emmerson, Carl G. and Wakefield, Matthew. "Pension Reform and Saving in Britain." *Oxford Review of Economic Policy*, Spring 2001, *17*(1), pp. 70–94.

Engen, Eric M.; Gale, William G. and Scholz, John Karl. "The Illusory Effect of Saving Incentives on Saving." *Journal of Economic Perspectives*, Fall 1996, *10*(4), pp. 113–38.

English Life Tables. Government Actuary's Department, various issues.

Family Expenditure Survey. Waves spanning the years 1974–1987; made available by the Office for National Statistics (ONS) through the ESRC data archive.

Feldstein, Martin S. "Social Security, Induced Retirement, and Aggregate Capital Accumulation." *Journal of Political Economy*, September/October 1974, *82*(5), pp. 905–26.

Gale, William G. "The Effects of Pensions on Household Wealth: A Reevaluation of Theory and Evidence." *Journal of Political Economy*, August 1998, *106*(4), pp. 706–23.

Gustman, Alan L. and Steinmeier, Thomas L. "Effects of Pensions on Savings: Analysis with Data from the Health and Retirement Study." *Carnegie-Rochester Conference Series on Public Policy*, June 1999, *50*(99), pp. 271–324.

Hubbard, R. Glenn. "Pension Wealth and Individual Saving: Some New Evidence." *Journal of Money, Credit, and Banking*, May 1986, *18*(2), pp. 167–78.

Jappelli, Tullio. "Does Social Security Reduce the Accumulation of Private Wealth? Evidence from Italian Survey Data." *Ricerche Economiche*, March 1995, *49*(1), pp. 1–31.

King, Mervin A. and Dicks-Mireaux, Louis A. "Asset Holdings and the Life-Cycle." *Economic Journal*, June 1982, *92*(366), pp. 247–67.

Poterba, James M.; Venti, Steven F. and Wise, David A. "How Retirement Saving Programs Increase Saving." *Journal of Economic Perspectives*, Fall 1996, *10*(4), pp. 91–112.

Tolley's Social Security & State Benefits Handbook. London: Tolley Publishing Company, Ltd., various years.