SIMULATING HOUSEHOLD SAVINGS AND LABOUR SUPPLY: AN APPLICATION OF DYNAMIC PROGRAMMING

James Sefton* and Justin van de Ven**

This paper describes a fully behavioural microsimulation model that has recently been developed at the National Institute for considering responses to changes in pension policy of household savings and labour supply. The model generates household decisions regarding labour/leisure, and consumption/savings by solving a dynamic programming problem over the simulated lifetime. This analytical framework incorporates a degree of complexity that is usually omitted from econometric analyses that are common in the literature.

1. Introduction

Economic analyses are often complicated by the impracticality of controlled experimentation. This has meant that simulation models are frequently useful for advancing our understanding of complex social systems, and for inferring the likely effects of policy counterfactuals. Within the field of models that have been developed by economists, microsimulation models of the household – models that generate data for individual households – are of particular value when considering the distributional implications of alternative government policies. Nevertheless, the economic literature has paid little attention to microsimulation models, which is partly attributable to the fact that such models have traditionally omitted behavioural responses – a cardinal sin for a *behavioural* science. Importantly, by failing to take household behaviour into account, microsimulation models fall foul of the Lucas Critique, although the latter has had an important bearing on the design of simulation models in macroeconomics for more than 25 years.[1]

Microsimulation models that incorporate behavioural effects are rare primarily because simulating behaviour is computationally demanding. It is, however, this aspect that presents today's analyst with a tremendous opportunity. Recent advances in personal computing power and software design mean that fully behavioural microsimulation models are now practicable, and anticipated advances mean that such models are likely to become increasingly sophisticated in the near future. The current paper describes a fully behavioural microsimulation model that has recently been developed at the National Institute of Economic and Social Research (NIESR). The model is designed to consider household labour/leisure and consumption/savings decisions – two issues of fundamental economic concern – at annual intervals during the simulated lifetime.

Microsimulation models were first used for economic analysis by Orcutt (1957), and are now commonly employed to undertake policy analyses in many countries around the world.[2] Microsimulation models are traditionally classified as either dynamic or static, depending upon how (and whether) the population is aged. Static microsimulation models, as their name suggests, determine the impact of counterfactual conditions upon a population of agents at a point in time. They usually consist of two parts; a reference database that details the characteristics of each agent in a population, and a procedure for calculating the impact on each agent of counterfactual conditions. Consequently, the range of policies that can be analysed by static microsimulation models is determined by the degree of detail that is provided by the reference database used. Given the demographic and income characteristics of families, for example, static microsimulation models are often used to determine the impact effects of alternative benefits policies on the income distribution, and upon the budgetary cost of the transfer system.

Static microsimulation models 'age' a population by reweighting the reference database using statistical projections to reflect an alternative time period. In contrast, dynamic microsimulation models age each individual described by the reference database in response to stochastic variation and an accumulated history. A dynamic microsimulation model that is designed to consider the effects of fiscal policy may, for example, generate characteristics that include marital

* The Business School, Imperial College London, and National Institute of Economic and Social Research. e-mail: jsefton@niesr.ac.uk. ** National Institute of Economic and Social Research. e-mail: jvandeven@niesr.ac.uk.

status, parenthood, income, and mortality at annual intervals for each person described by a reference database. The income of each individual at any given year is often simulated based on characteristics such as the individual's past income, their demographic characteristics, and upon a stochastic term that accounts for unexplained variation. This type of procedure builds up a life history for each individual in a population, which significantly increases the range of questions that can be explored, relative to static models. Most dynamic microsimulation models are designed specifically to consider the intertemporal and long-term effects of counterfactual conditions, rather than the impact effects with which static models are usually concerned.

Most microsimulation models that are currently in use are static. Prominent examples of these include, STINMOD (Australia; refer to the STINMOD Technical Series, NATSEM, Australia), POLIMOD (UK; see Redmond *et al*., 1998), EUROMOD (15 member states of the European Union; see Sutherland, 2001), TRIM2 (US; see Giannarelli, 1992), SPSP (Canada; refer to Statistics Canada), SWITCH (Ireland; see Callan *et al*., 1996), LOTTE (Norway, see Fjaerli *et al*., 1995), and FASIT (Sweden; refer to the Swedish Ministry of Finance).[3] Advances in computing power, analytical techniques, and the availability of increasingly detailed survey data have led to an increase in both the number and sophistication of dynamic microsimulation models. Some recent examples of these include ASPEN (US; see Basu *et al*., 1998), CORSIM (US; see Caldwell, 1997), DYNACAN (Canada; refer to Statistics Canada, based on DYNASIM, see Orcutt *et al*., 1976), HARDING (Australia; see Harding, 1993), DYNAMOD-2 (Australia; King *et al*., 1999), MICROHUS (Sweden; see Andersson *et al*., 1992), and SESIM (Sweden; refer to the Swedish Ministry of Finance).

In addition to the static-dynamic dichotomy, microsimulation models can also be distinguished by the extent to which they incorporate agent specific behavioural responses. Given the ageing populations and reduced rates of economic growth that have been observed in many industrialised countries, attention has focused in recent years on the responsiveness of labour supply, savings, and fertility to alternative tax and benefit systems.[4] Behavioural response may be modelled using statistical projections estimated from survey data (see, for example, CORSIM), or an explicit consideration of how decisions are made. The former of these methods is relatively easy to apply, but suffers from inherent inconsistencies (which are discussed at length in Section 2). The latter method usually involves assuming that reference units make their decisions to maximise an assumed objective (utility) function, subject to various practical constraints (such as the available funds that a household can spend). It is the most complex computationally, and therefore used only rarely.

The model described by this paper falls into the last of the categories described above. Specifically, household decisions regarding labour and consumption are simulated by assuming that the household maximises an intertemporal utility function, subject to a budget constraint. This approach is particularly useful when considering counterfactuals that are likely to affect agent behaviour. If, for example, an analysis of alternative pension policies holds household savings and labour supply fixed, then the conclusions derived are likely to be systematically in error – the analysis will fall foul of the Lucas Critique. Behavioural microsimulation models are motivated by the view that important insights may be obtained by allowing households to adapt their behaviour in response to the incentives of policy counterfactuals.

It would, however, be disingenuous to suggest that there are no disadvantages to using the type of model that is described here. Behavioural models need to assume that households behave in some well defined manner. Although it is possible to test the assumptions inherent in the behavioural framework assumed – and a great deal of work has been devoted to this (see Deaton, 1993) – the validity of such models can never be verified positively (we might reject, but can never accept). Consequently, the predictions made by behavioural models always remain subject to the uncertainty that underlies the analytical framework adopted. Furthermore, these models are complicated to solve, and so remain highly stylised (subject to existing computing technology). Importantly, this computational complexity means that it is difficult to describe statistically the uncertainty that is associated with observations derived from such models, unlike common econometric analyses. These models also do not yet attempt to capture realistically the learning process as people adjust to a new policy environment. They are therefore better at modelling the long-term impact of a policy change.

Alternative methods of simulating household behaviour – with particular reference to the retirement decision – are described and compared in Section 2. NIESR's fully behavioural microsimulation model is described in

Sections 3–5. Practical applications of the model described by the current paper are discussed in Section 6. Directions for further research are discussed in a concluding section.

2. Alternative models of household behaviour

The model described by this paper uses dynamic programming methods to simulate household behaviour. It is consequently useful to contrast the dynamic programming approach with econometric methods that are commonly employed in the economic literature.

Econometric models are useful tools for analysis because they describe complex interactions in a highly accessible form. The specifications used for estimation can help to stylise a prohibitively complex relation in such a way as to focus attention upon those aspects that are of immediate concern. Furthermore, the explicit nature of the error associated with econometric regressions emphasises the limitations of a stylised specification in a way that is inherently appealing – the error structure endows econometric estimates with an air of honesty. The principal limitation of any regression model, however, is the data that are available for estimation.

The issues are particularly evident when analysing the impact on household retirement behaviour of pension policy counterfactuals. The timing of retirement is an aspect of household behaviour that depends heavily upon expectations. Most econometric analyses reflect this fact by making explicit assumptions regarding expectations of future income and pension benefits, and of mortality rates. Data regarding household income have improved substantially during recent decades, which facilitates the specification of associated expectations. Data on other important variables – wealth being one prominent example – remain relatively scarce. Most importantly, however, the variation of fiscal policy that is described by current data is highly limited. This often hamstrings the econometric approach for analysing policy counterfactuals.

Consider, for example, a sample of households that must choose when to retire from one of two periods. A late retiring household works during period 1 ($L_1 = 0$) to earn a wage Y, and receives a pension P_L in period 2. An early retiring household retires in period 1 ($L_1 = 1$), and receives a pension P_E in each of periods 1 and 2. Furthermore, households are assumed to make their retirement decision to maximise their intertemporal utility:

$$U = \frac{c_1^{1-\gamma}}{1-\gamma} + \alpha L_1 + \frac{c_2^{1-\gamma}}{1-\gamma} \tag{1}$$

subject to a budget constraint:

$$W_0 + (Y + P_L)(1 - L_1) + 2P_E L_1 \le c_1 + c_2 \tag{2}$$

where c_t denotes consumption in period t, and W_0 defines a household's initial wealth.

Given the utility maximisation problem defined by equations (1) and (2), a household will select early retirement if:

$$\frac{(W_0 + 2P_E)^{1-\gamma}}{1-\gamma} + \frac{\alpha}{2^{\gamma}} \ge \frac{(W_0 + Y + P_L)^{1-\gamma}}{1-\gamma} \tag{3}$$

Linearising inequality (3) around $W_0 = 0$ and the income $\bar{Y}$ that makes a household indifferent between early and late retirement, it can be shown that a household will prefer early retirement if:

$$Y - \left(1 - \left(1 + \frac{\alpha(1-\gamma)}{2(P_E)^{1-\gamma}}\right)^{\frac{-\gamma}{1-\gamma}}\right) W_0 - \bar{Y} \le 0 \tag{4a}$$

$$\bar{Y} = \left((2P_E)^{1-\gamma} + \frac{\alpha(1-\gamma)}{2^{\gamma}}\right)^{\frac{1}{1-\gamma}} - P_L \tag{4b}$$

It is possible to estimate equation (4a) as a probit model, and this method has been adopted by, for example, Blau and Gelleski (2001), Blau (1994), and Gruber and Madrian (1985). Furthermore, probit models are commonly used to simulate labour supply in statistical dynamic microsimulation models; see, for example, CORSIM, SESIM, and DYNAMOD-2. Although the coefficients estimated from such a probit model provide useful information regarding retirement behaviour *given* existing pension policy, they are not suitable for considering the behavioural response to policy counterfactuals. This is because the coefficients of the model depend upon pension policy parameters (defined by P_E and P_L in equations (4a) and (4b).

To permit econometric estimation of a retirement model that can be used to analyse the behavioural effects of pension counterfactuals, data are required that describe the retirement decisions of a sample population for whom some variation in the relevant pension characteristics is observed. A prominent example of this type of study is by Stock and Wise (1990). The data used by Stock and Wise provide a detailed description of the income and employment history for a sample of older salesmen from a large Fortune 500 company in the United States. Stock and Wise also had detailed information regarding the occupational pension administered by the firm, which enabled aspects of the pension rights accrued by individuals in the sample to be imputed. Importantly, these pension rights exhibited variation amongst the individuals of the sample.

Stock and Wise (1990) modelled the retirement decision as an irreversible exit from the labour market by way of its Option Value (OV). The model abstracts from the effects of intertemporal consumption smoothing by focusing upon the utility of an individual's income stream – wealth is omitted from the model by assuming that consumption at any time is equal to income. Individuals are assumed to delay retirement at any time, *t*, if the discounted expected value of their future utility from income is improved by doing so. In terms of the example given above, an individual is considered to prefer early retirement using the OV model if:

$$\left(Y^{1-\gamma} + (kP_L)^{1-\gamma}\right) - 2(kP_E)^{1-\gamma} \leq 0 \qquad (5)$$

where the parameter *k* accounts for the different contribution to utility made by unearned relative to earned income.

The specification defined by equation (5) can be combined with household specific characteristics to form a probit model for econometric analysis. The results that Stock and Wise (1990) present are compelling, and consequently provide a strong argument in favour of the OV framework.

The principal simplification of the OV model is its omission of wealth from the retirement decision – it focuses upon the foregone opportunities associated with retirement, and not the historical provisions made for retirement.[5] The fact that Stock and Wise (1990) report plausible econometric estimates is attributable, at least in part, to the homogeneity of the sample that they used. In contrast, when Blundell and Emmerson (2003) estimated the OV model using data for a nationally representative sample of the UK population, they found that wealth had a positive and highly significant effect on the probability of retirement. This suggests that early accumulation of wealth tends to encourage early retirement. Notably, the coefficient on the OV variable estimated by Blundell and Emmerson ceased to be significant (at any reasonable confidence interval) following the addition of wealth to the probit regression.

When individuals are free to choose the timing of their retirement, it is intuitive that they take into consideration the consumption that they will be able to finance during retirement. The OV model reflects one aspect of this consideration – an evaluation of the foregone opportunities associated with selecting a particular date for retirement. Wealth is clearly another important aspect that has a bearing on the future consumption that an individual can afford. Omission of either of these considerations from an analysis of the retirement decision is likely to result in systematic error, as indicated by the findings reported by Blundell and Emmerson (2003). The scarcity of microdata that describe holdings of wealth is consequently an important limitation for econometric analyses of retirement.

Even if extensive wealth data were readily available, however, the practical implications for retirement behaviour of pension policy counterfactuals would remain difficult to infer. This is because the observed distribution of wealth held by a population will reflect expectations regarding the policy environment. To estimate econometrically the implications for retirement behaviour of pension policy counterfactuals, data are consequently required that describe variation regarding expectations of pension policy.[6] Such regression models – and the data required to estimate them – are obviously demanding.

Nevertheless, a number of econometric studies have attempted to estimate behavioural response to changes in pension policy. Country-specific econometric studies of the behavioural effects of pension policy usually consider data that describe known policy experiments. The estimates reported by such studies often indicate that fiscal policy tends to have a statistically insignificant effect on savings and retirement behaviour.[7] However, this finding can often be attributed to the subtlety of the policy change, and to delays in the behavioural response. In contrast, studies that report econometric estimates calculated using cross-country data usually find statistically significant behavioural effects, consistent with the wider variation of pension policy that is observed between countries.[8] Even so, cross-country data describe a limited range of policy alternatives, and suffer from undesirable population heterogeneity that may be difficult to control for. This undesirable population heterogeneity arises, for example, due to the institutional differences that exist between countries.

In summary, it is useful to refer to the following data continuum. On one extreme, all of the survey data that are required to consider a particular issue are readily available; and on the other, no survey data can be obtained. In the former extreme, regression methods will provide accurate estimates for any well-specified model, in which case econometric estimation presents a useful tool for describing complex relations. In the latter extreme, regression methods will produce biased estimates for a well-specified model, which may confuse an already complicated debate. When all of the data required to consider a problem do not exist (and there are some examples when this will *always* be the case), then it is useful to impose a framework of analysis that focuses upon outcomes that might reasonably be expected. This is the objective of dynamic programming.[9]

The dynamic programming (DP) model differentiates between two types of variables; *state variables*, which define the existing characteristics of an agent, and *control variables*, which define the set of decisions that the agent can make. Behaviour is described within this framework by an *optimal decision rule*,[10] which indicates the control variables that an agent would select, given any combination of state variables. The optimal decision rule characterised by the DP model maximises the agent's *value function*, subject to defined practical constraints. In the case of the simple example described above, W_0, and Y are state variables, c_1, c_2, and L_1 are control variables (which are subject to the constraint defined by equation (2)), and the value function is defined by equation (1).

The preceding discussion highlights the importance of the practical constraints imposed on a DP problem. If, for example, the simple two-period retirement problem described above is amended to eliminate household wealth, such that $W_0 = 0$ and consumption in each period is exactly equal to income, then the OV and DP models would provide very similar descriptions of retirement behaviour. In contrast, imposing a budget constraint that permits the accumulation of wealth allows the retirement decisions of households to be affected by both the value of foregone opportunities (which is the focus of the Option Value model), and the value of consumption financed from accumulated assets (consistent with the findings of Blundell and Emerson, 2003).

The cost of the additional flexibility afforded by the DP framework is its computational complexity. In the two-period example discussed above, an analytical solution is easily obtained. However, as additional choice variables and time periods are added, the complexity of the DP problem rapidly increases to the point where analytical solutions become impractical and numerical solution methods are necessary. The limitations attributable to this complexity are of immediate practical concern. It would, for example, be desirable to use the DP framework to impose 'sensible' restrictions upon an econometric problem when insufficient data are available for estimation. Alternatively, econometric estimation of a DP framework could help to ensure sensible policy simulations. Unfortunately, the complexity of most DP problems complicates attempts to bridge the gap between the DP and regression frameworks.[11] Consequently, DP microsimulation models are usually calibrated to stylised observations, rather than econometrically estimated.

3. The current microsimulation model

A partial equilibrium dynamic microsimulation model has been constructed to explore household savings and labour decisions. The decision unit in the model is the household. Each household is aged by annual increments, from 20 to 90 based upon the age of the household's reference person.[12] In every year, the household decides whether to work full time or not at all (households are treated as having an aggregate labour supply),[13] and how much to consume given its economic situation, under the constraint that its net worth must remain positive. A broad definition is assumed for the economic situation of a household, which includes the household's age, its size, the wealth that it has managed to accumulate, the interest rate, the level of means tested income support available, and the wage that it can command for its labour. This wage rate evolves stochastically.

The household is forced to retire by state pensionable age (65 for the UK), if it has not already chosen to do so. In retirement the household pays for its consumption out of its savings, or out of income derived via pensions and investments.

Simulated households are described by seven characteristics:

i) the number and age of household members
ii) time of death

iii) the human capital of the household
iv) the labour supply of the household
v) household consumption
vi) household wealth
vii) household (mandatory) defined benefit (DB) pension entitlement.

The following sections describe the methods used to generate each of the seven household characteristics defined above. To provide a practical example, statistics are reported for a specification that has been calibrated to reflect UK survey data.

4. Non-behavioural characteristics

Of the seven characteristics defined in the preceding section, three are simulated exogenously. That is, household behaviour with regard to household size, time of death, and human capital does not respond to policy counterfactuals. The procedures used to generate these characteristics, and the associated modelling considerations are described below.

4.1 Demographic size and composition

The size of each household varies with time to reflect the coupling of individuals, and the birth and aging of children who eventually leave home. Household size is, however, modelled in a predetermined (exogenous) fashion, such that each reference person is assumed to know at the very beginning of his/her simulated lifetime (age 20), when he/she will marry, when and how many children he/she will have, and when his/her children will leave home. This is a strong assumption, particularly when compared with statistical (non-behavioural) dynamic microsimulation models that commonly simulate much more demographic heterogeneity than is considered here, using random allocation methods that reflect real world uncertainty (see, for example, CORSIM, DYNACAN, and DYNAMOD-2). The simplified framework used to simulate demographics is made necessary by the current state of the art in personal computing technology. With regard to analysis of savings and retirement for which the model has been constructed, the methods used to simulate demographics will fail to capture shocks that households experience in practice due, for example, to divorce, unplanned childbirth, or unanticipated changes in health.[14] Chart 1 describes the numbers of adults and children by age of reference person – the smoothed data are used for simulations.[15]

Chart 1. Household size by age of reference person

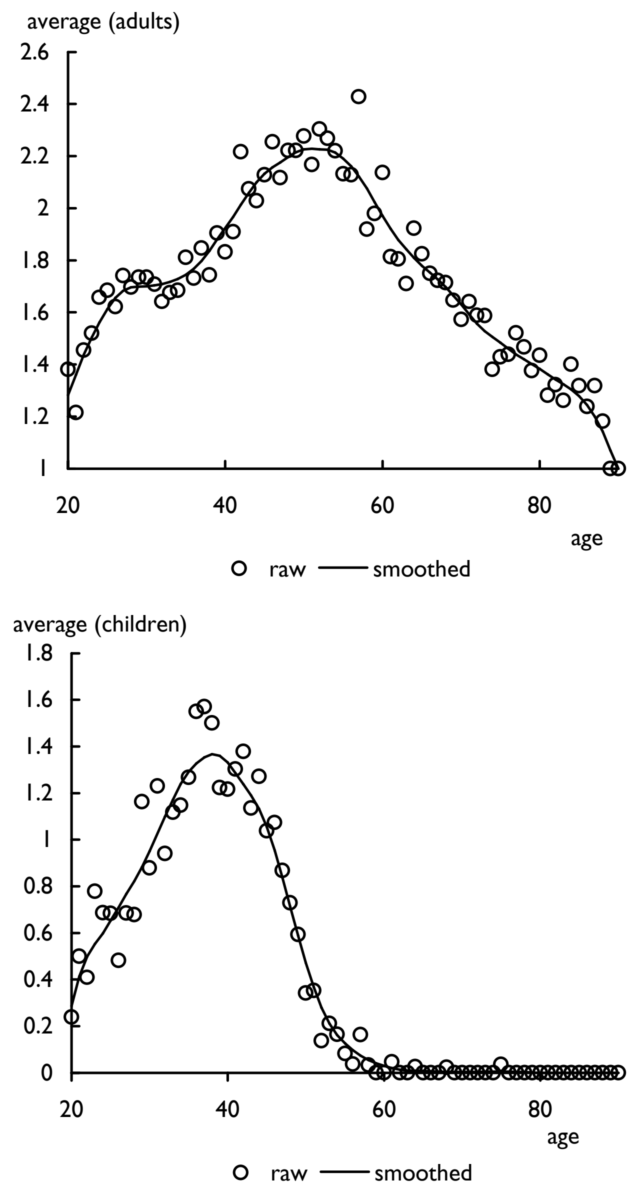

Source: ECHP data and authors' calculations, see footnote 15.

4.2 Household mortality

Each household is selected to die, based upon an exogenously defined survival function that varies with age. Importantly, households do not know *a priori* when they will die, they only know the simulated survival function. This means that the model is able to capture the precautionary savings that households are likely to accrue in practice to offset the effects of uncertain life expectancy. The model does not, however, include any

Chart 2. Mortality rates by age

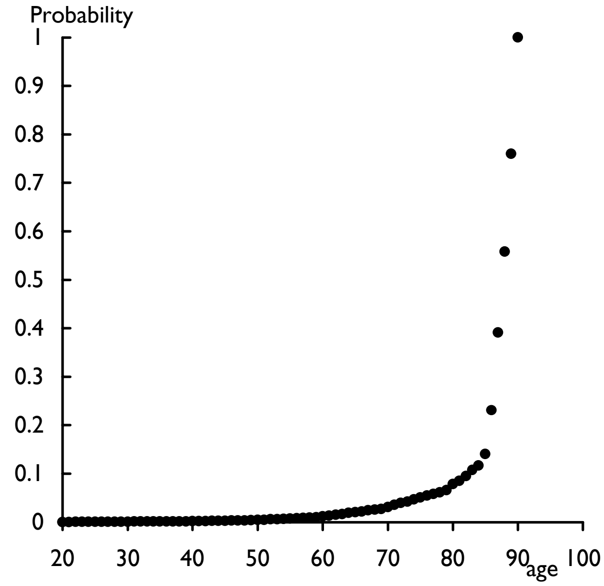

Source: See footnote 16.

endogeneity of the survival probabilities that might be expected to exist with, for example, household wealth. Chart 2 displays mortality rates by age of household reference person.[16]

4.3 Human capital

A household's labour income is equal to its human capital multiplied by its labour supply. The human capital of a household is simulated as a stochastic process, described by:

$$h_{it} = \beta h_{it-1} + \theta w_{it-1} + f(t) + \varepsilon_{it} \tag{6}$$

where h_{it} defines (log) human capital of household i at age t, w_{it} is a dummy variable that takes a value of one if individual i was working at age t and zero otherwise, and ε_{it} is an individual specific error term (which is assumed to be identically and independently distributed across all i and t). In each period, human capital consequently depends upon human capital in the preceding period (where β accounts for some depreciation), labour in the preceding period (to include a learning-by-doing effect), an underlying age trend (that is the same for all simulated households), and a random disturbance term. Observations derived during model calibrations suggest that the learning-by-doing effect plays an important role in motivating households to supply labour at the beginning of their working lives when the instantaneous returns to labour are low.

The specification of the model used to generate human capital has been selected with some care. Its principal advantage is that the non-random inputs to equation (6), $(t,\ h_{it-1},\ w_{it-1})$, are all variables of the microsimulation model that extend no more than one period into the past. This helps to simplify the analytical problem considerably, which is discussed at greater length in the following section.

In addition to making the microsimulation model analytically tractable, the model of human capital described by equation (6) also bears close similarities to alternative wage equations that have been considered in the literature. Consider, for example, the simple 'regression-toward-the-mean' (RTM) model of human capital evolution that is studied in detail by Atkinson *et al*. (1992), and used by Huggett (1996) in his equilibrium model of the US economy.[17] The RTM model of human capital is described by:

$$z_{it} = \beta z_{it-1} + \varepsilon_{it} \tag{7}$$

where $z_{it} = (h_{it} - \bar{h}_t)$ is the deviation of household i's human capital from the population's geometric mean $(\bar{h}_t = \frac{1}{n}\Sigma_i^n h_{it})$. Including a learning-by-doing effect into the model defined by equation (7) and rearranging:

$$\begin{aligned} h_{it} &= \beta h_{it-1} + \theta w_{it-1} + (\bar{h}_t - \beta\bar{h}_{t-1} - \theta\bar{w}_{t-1}) + \varepsilon_{it} \\ &= \beta h_{it-1} + \theta w_{it-1} + g(t) + \varepsilon_{it} \end{aligned} \tag{8}$$

where $\bar{w}_t$ is the proportion of the population employed at age t, the learning-by-doing effect is described by $\theta(w_{it} - \bar{w}_t)$, and $g(t)$ defines the bracketed term in the first line of equation (8). Comparison of equation (6) with (8) reveals the similarities between the model used to generate the evolution of human capital and the RTM model described by equation (7). Furthermore, van de Ven (1998) suggests that there exists a close relationship between the RTM model of human capital, and the classical model of income dynamics advocated by Mincer (1974).

Note, however, that the RTM model described by equation (8) and the model used to simulate the evolution of human capital (described by equation (6)) are not equivalent. Importantly, equation (8) reveals that augmenting the model described by equation (7) to include a learning-by-doing effect implies that the specification of $g(t)$ depends upon the policy regime. This is in contrast to the specification assumed for $f(t)$ in equation (6), which is policy invariant.

Estimates and calibration

Difficulties were encountered when estimating equation (6) econometrically due to two principal factors; equation (6) is a highly stylised specification for the evolution of human capital, and the data used for estimation provide insufficient information to describe adequately lifetime income dynamics. These practical complications provide a pertinent example of the limitations to econometric analysis that are discussed in Section 2, and it is consequently useful to describe them at some length here.

Equation (6) was estimated using a sample selection model of individual full-time employment wages. This model takes into consideration the fact that wages are only observed for individuals who are working, and that there is likely to be a relationship between the probability of working and the wage rate. The regression was undertaken using the 'Sampsel' procedure in TSP, full details of which can be obtained from the 'TSP 4.4 User's Guide' (see http://elsa.berkeley.edu/wp/tsp_user/tspugpdf.htm). The data used to estimate equation (6) were derived from the ECHP, which provides panel data for a period of up to seven years (at the time of writing).

The sample selection model involves estimating two equations, a probit to identify individuals who are employed, and a (log) wage equation. The probit equation used predicts the probability of an individual's employment status with regard to various demographic, health, and economic variables. These are not of immediate interest here, and are consequently reported in Appendix B.[18] Estimation of the specification defined by equation (6) requires data regarding an individual's human capital in a given period, and in the immediately preceding period. Since an individual's human capital is observed only if he/she is working, and since the current analysis considers only full-time employment, estimation of the specification defined by equation (6) results in multicollinearity between the employment identifier, w_{it-1}, and the regression constant. To overcome this problem, it is necessary to resort to a reduced form of equation (6):[19]

$$h_{it} = \beta^R h_{it-R} + \sum_{s=1}^{R} \beta^{s-1} (\theta w_{it-s} + f(t+1-s) + \varepsilon_{it+1-s}) \quad (9)$$

Two sets of regression estimates for equation (9) are reported in table 1, one in which $R = 3$, and another in which $R = 6$. Both of these regressions use observations drawn from the ECHP for 2000/01 to describe h_{it}. Table 1 also includes 'restricted estimates', which are described at length below.

When the data used to estimate equation (9) describe limited temporal variation, the ability to capture important aspects of persistence in the evolution of human capital, and consistency over the working lifetime are compromised. The practical relevance of these limitations is reflected by the two sets of unrestricted estimates that are reported in table 1. With regard to the issue of persistence, note that extending the temporal dimension of the data used for estimation from $R = 3$ to $R = 6$ results in a significantly higher estimate for β. With regard to consistency over the lifetime, the top panel of chart 3 plots the profiles of human capital over the working lifetime that are implied by the unrestricted regression estimates for an individual with geometric mean income at age 20, assuming full-time employment over the entire working lifetime. The top panel of chart 3 also plots average income by age derived from Family Expenditure Survey[20] and ECHP data.

The top panel of chart 3 highlights the difficulties that may arise when the data used for econometric

Table 1. Econometric estimates of human capital equation

	R = 3		R = 6		Restricted ests.	
	Estimate	Std error	Estimate	Std error	Estimate	Std error
β	0.829272	8.36E–03	0.881672	6.49E–03	0.975000	NA
full-time (t–1)	0.080523	3.40E–02	0.108149	2.21E–02	0.056670	2.08E–02
c	1.548770	1.80E–01	1.062290	1.19E–01	0.631478	9.72E–02
age	1.96E–02	1.33E–02	5.16E–03	9.04E–03	–2.44E–02	7.49E–03
age^2	–5.41E–04	3.26E–04	–2.26E–04	2.24E–04	4.64E–04	1.88E–04
age^3	4.42E–06	2.59E–06	2.30E–06	1.80E–06	–2.91E–06	1.53E–06
inverse mills	–0.143250	2.99E–02	–0.157948	3.75E–02	–0.028813	4.21E–02
R^2	0.509533		0.383955		0.362528	
std error	0.333069		0.37255		0.424956	
adj std error	0.154155		0.106433		0.148882	

estimation provide insufficient information to infer important aspects of the relationship of interest. Including three lagged periods in the regression of equation (9) results in an implied human capital profile that is quite different to the profile obtained when six lagged periods are included – neither of which bear a particularly close resemblance to the lifetime profile described by the survey data.

In response to the above observations, the six-lagged period specification of equation (9) was used to obtain

Chart 3. Estimates of human capital by age

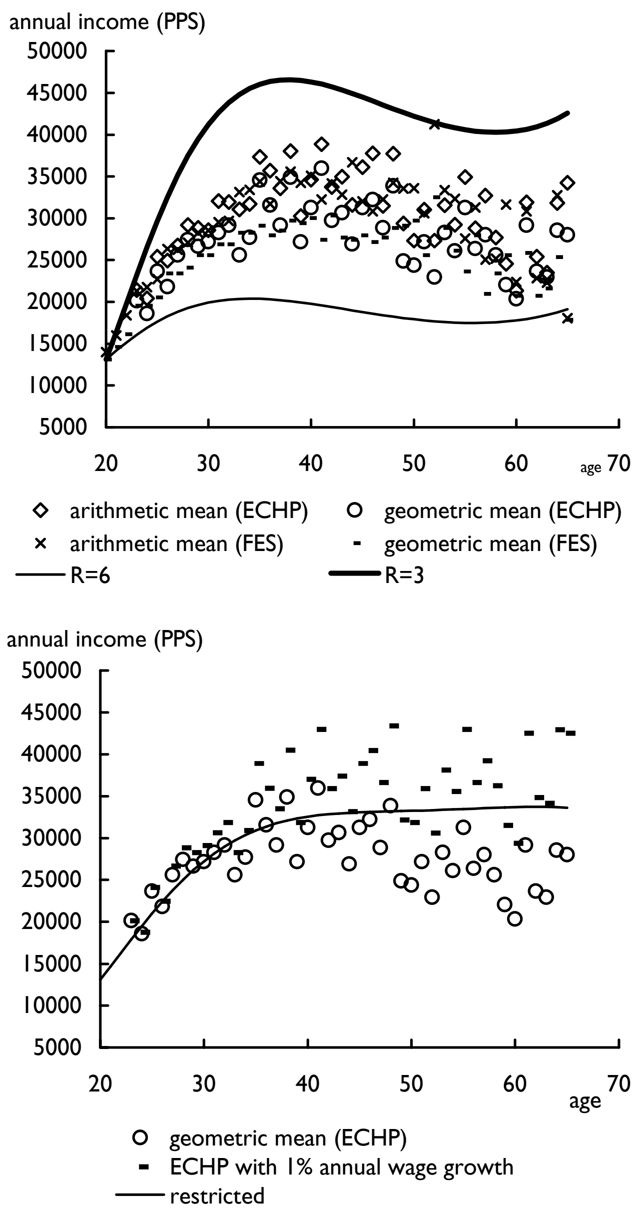

Source: Authors' calculations.

estimates for the model of human capital evolution, where β was restricted to take a value of 0.975. Regression estimates are displayed in table 1, and the implied lifetime profile of human capital can be compared against associated survey data in the bottom panel of chart 3. The population averages of annual income reported in chart 3 are derived from a cross-section, which fails to reflect the impact on human capital evolution of wage growth. Consequently, a series that reflects a conservative wage growth of 1 per cent per year is also included for comparisonin the bottom panel of chart 3.

Comparing the regression statistics reported in table 1 for the restricted and unrestricted specifications of the six lagged period model indicate that fixing β to 0.975 has a small, though significant, debilitating effect on the model's ability to reflect variation observed in the survey data. Nevertheless, the progression of human capital over the lifetime that is implied by the restricted model, as displayed in chart 3, appears to exhibit a closer relation to the survey data. Consequently, the restricted estimates reported in table 1 are used for the microsimulation model.

All that remains to characterise fully the model used to simulate human capital evolution is a description of population heterogeneity. Individual heterogeneity of human capital enters the model in two ways, through the initial dispersion imposed at age 20, and through transitory terms, ε_{it}, for each subsequent period of the working lifetime. The dispersion of human capital assumed at age 20 is based upon the relationship between the standard deviation of (log) full-time wages and age described by ECHP data for 2000/01, as displayed in chart 4. Following the trend displayed in chart 4, and the associated relation with FES data, a value of 0.40 was selected for the initial dispersion of human capital. Furthermore, the geometric mean assumed for human capital at age 20 is 13,091.88 (PPS), which is equivalent to £9,426 (GBP) described by the FES.[21] The standard deviation of the temporal variation term ε_{it}, is assumed to equal 0.148882, which was calculated from the econometric estimate for the standard deviation of the restricted regression reported in table 1.

5. The dynamic programming problem

Labour supply, consumption, wealth, and DB pension entitlement are all modelled endogenously as a DP problem. Households are considered to choose their

Chart 4. Standard deviation of full-time wages by age

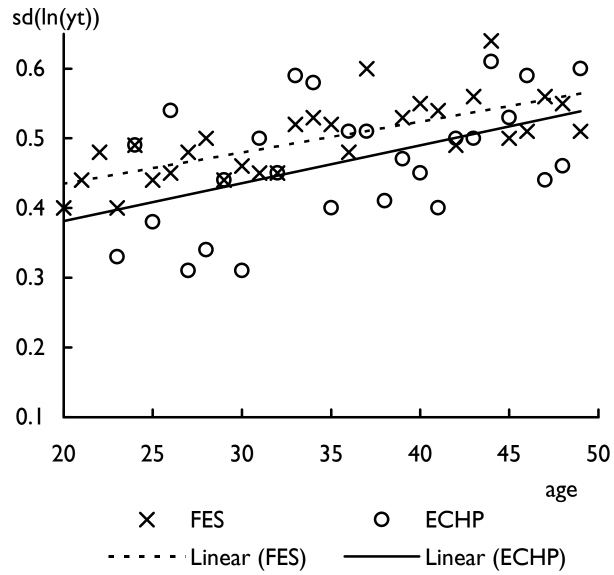

Source: Authors' calculations.

labour supply and consumption in every period to maximise their expected utility, subject to a budget constraint.[22] Expected lifetime utility is described by the additively separable function:

$$U_t = E_t\left(\sum_{i=t}^{70} u\left(\frac{c_i}{m_i}, l_i\right)(1+\pi)^{t-i}\right) \quad (10)$$

where E_t is the expectation operator, $c_t \in R^+$ is household consumption, $m_t \in R^+$ is the household's adult equivalent size, and $l_t \in [0,1]$ is the proportion of the household's time devoted to leisure at time t = (0,...,70).[23] The parameter π is a discount factor, which is assumed to be time independent.[24] During the working lifetime – defined between t = 1 and the relevant state pensionable age t_P – a household's labour choice is restricted to full-time employment, and not employed. After t_P, the household is forced to retire.

A Constant Elasticity of Substitution (CES) utility function is assumed, which is defined by:

$$u(C_t, l_t) = \frac{1}{(1-\frac{1}{\gamma})}\left(C_t^{(1-\frac{1}{\rho})} + \alpha^{\frac{1}{\rho}}(\delta_t l_t)^{(1-\frac{1}{\rho})}\right)^{\frac{1-1/\gamma}{1-1/\rho}} \quad (11)$$

where γ is the inter-temporal elasticity of substitution, and ρ is the elasticity of substitution between $C_t = c_t/m_t$ and l_t. δ_t is a scaling parameter that adjusts for lifestyle changes that arise during retirement (and can also be varied to reflect changes in health). The higher the value of ρ, the higher the proportional change between consumption and leisure for a given proportional change in prices. Similarly, the larger the value of γ, the higher the proportional substitution between consumption today and consumption tomorrow for a given proportional change in interest rates. Wealth in any period, W_t, is constrained to be non-negative, and is given by:

$$W_{t+1} = W_t - c_t + y_t^{DI}(y_t, W_t, m_t) \quad (12)$$

where $y_t^{DI}(y_t, W_t, m_t)$ is the post-tax and benefit income obtained by a household of age t, given pre-tax income y_t, wealth W_t, and equivalence size m_t. Pre-tax income is derived via the real return on household wealth, RW_t, through a household's labour supply during the working lifetime ($t < t_P$), and via DB pension entitlements during the retired lifetime ($t \geq t_P$), such that: $y_{it} = RW_{it} + h_{it}(1-l_{it}) + (S_{it}|t > t_P)$, where S_{it} defines household i's DB pension entitlement at age t. The real interest rate is assumed to be 5 per cent. DB pension entitlements are calibrated to reflect the rules of the respective tax and benefits system. Full details of the tax and benefits system simulated for the UK are provided in Appendix C.

Solving the model

Given a household's wealth, human capital, and DB pension entitlement at any time t, the value function associated with equation (10), can be defined as:

$$V(W_t, h_t, S_t) = \max_{c_t, l_t}\left\{u(C_t, l_t) + (1+\pi)^{-1} E_t V(W_{t+1}, h_{t+1}, S_{t+1})\right\} \quad (13)$$

We require the optimal values of the control variables, (c_t, l_t).

First, grids are constructed that define every practicable combination of the state variables – wealth, W_t, human capital, h_t, and pension entitlement, S_t – for each period in the simulated lifetime, t = 0,...,70. Solution of the utility maximisation problem then proceeds by backward induction. In period t = 70, the specification of the model implies that l_{70} = 1 (all households are assumed to be retired), and $c_{70} = W_{70} + y_{70}^{DI}(S_{70} + RW_{70}, W_{70}, m_{70})$ (there is no bequest motive, and death following period 70 is certain). This gives the value of $V(W_{70}, h_{70}, S_{70})$ associated with every grid point defined in period 70.

Maximisation of equation (10) with respect to c_t, subject to the budget constraint defined by equation (12), gives rise to the following Bellman equation for internal solutions:

$$\frac{\partial u}{\partial C_t} m_t^{-1} - (1+\pi)^{-1} R'_{t+1} m_{t+1}^{-1} E_t \frac{\partial u}{\partial C_{t+1}} = 0 \qquad (14)$$

where R'_{t+1} is the post-tax real rate of return at time t+1. The solution method involves solving equation (14), subject to the non-negativity constraints defined above for each grid reference point, where the required expectations are evaluated with reference to higher age grid solutions given the defined relations for W_t, h_t, and S_t, and the relevant probability of survival.[25] Where leisure is a choice variable ($t < t_P$), the alternative leisure choices are considered separately, and a selection is made on the basis of the implied value function V.

The above discussion reveals the importance of ensuring that the specification adopted for the evolution of human capital depends only upon variables that extend one period back in the microsimulation model. If, for example, the human capital specification described by equation (6) was augmented to include a learning-by-doing effect with a two-period lag, then the labour choice in any given period would affect the household's human capital in both the immediately succeeding period and two periods hence. This would substantially complicate calculation of the expected marginal utilities that are required to determine the solution to the model, as described by equation (14).

6. Practical applications

The current model provides a useful tool for considering a range of policy relevant issues. Sefton *et al*. (1998), for example, consider the distributional implications of alternative pension policies using a model that is similar to the one described here, with the exception that the model they used assumes an inelastic labour supply. The results obtained by Sefton *et al*. (1998) suggest that means testing of old age pensions can increase wealth inequality among the retired because of the effects that it has on the incentives to save. Means testing reduces the benefits of saving for affected households, and the analysis reported in Sefton *et al*. (1998) indicates that this may reduce the provisions made for retirement by the poorest households. This is an interesting finding, because advocates of means testing are often also concerned about inequalities in wealth holdings and would like these to be reduced.

The study by Sefton *et al*. (1998) is complemented by another study by Sefton and Weale (2003), which considers the likely effects of a reduction in the pension taper rate (rate of withdrawal in response to private income) from 100 per cent to 50 per cent. This scenario reflects the policy change that occurred in the UK with the introduction of the Pension Credit in October of 2003. The analysis presented by Sefton and Weale (2003) focuses upon the income and substitution effects of the simulated policy change. Notably, substitution effects can only be captured by a model that includes behavioural responses. The study by Sefton and Weale (2003) suggests that there is likely to be popular support for means-tested rather than flat-rate benefits in the short term, but that voters would prefer flat-rate benefits in the long run.

In the UK, various tax incentives are offered to encourage individuals to save for retirement. The most important of these are the 25 per cent tax free lump sum payment received at retirement, and the opportunity for high tax rate payers to receive relief at the high rate on pension contributions but pay tax at the low rate on the pension benefits. To take advantage of these tax incentives, however, individuals must effectively 'lock away' their savings in approved pension savings accounts during their working life, thereby removing any liquidity on these assets. They must also annuitise the majority (75 per cent) of these assets on retirement. Cantor and Sefton (2002) examined how this loss of liquidity on pension savings might affect the incentives of individuals to save in approved pension accounts. They found that the impact depended significantly on an individual's income.

The findings of Cantor and Sefton (2002) can be explained as follows. Individuals on low incomes generally have a lower ratio of savings to income, and face a higher level of income uncertainty than individuals on higher incomes. Individuals with low incomes consequently have a stronger motive to keep their savings liquid, should their future income fall unexpectedly, and the need to finance consumption from savings arise. This makes individuals on low incomes less willing to 'lock away' their savings, which means that they fail to take full advantage of the tax incentives offered on pension accounts.

Furthermore, the paper by Cantor and Sefton (2002) investigates a specific proposal to reduce the liquidity trap for those on low incomes. All annual contribution limits to pension accounts were removed and replaced by a single lifetime contribution limit. This effectively allowed those on low incomes to save in liquid assets and just before retirement switch these assets into pension accounts to take advantage of the tax benefits. The results of the model suggested that this could boost savings of those on low incomes by as much as 30–50 per cent, whilst leaving those on high incomes unaffected. The changes to pension arrangements being implemented after the 2004 Budget reflect this. They replace annual contribution limits by something much closer to lifetime contribution limits, thereby reducing the illiquidity of retirement saving.

In a related study, Dutta *et al*. (2000) attempt to explain the observed fall in the level of UK capital taxation since the Second World War as a response to the changing demographic structure. Generally, the old prefer lower levels of capital taxation and proportionally higher levels of labour income taxation compared to the young. This is because an increase in capital taxation shifts resources from the old to the young, as the old own a far larger proportion of the assets. One might consequently expect that capital tax rates will fall as a population ages. The model used by Dutta *et al*. (2000) is calibrated to the UK economy as observed in in 1951 and 1991. In each case, the model population is asked to vote for a one-off change in the capital tax rate. The labour income tax rate then adjusts so that the government budget is balanced. The majority of the 1951 model population chose a 40 per cent capital tax rate relative to a 20 per cent rate, whereas the majority of the 1991 population chose the 20 per cent tax rate.

The model that is described by the current paper is also being used to consider the likely effects on retirement behaviour of alternative pension policy counterfactuals. Preliminary analysis, for example, suggests that the Pension Credit referred to above may strike an acceptable balance between redistributive objectives and the disincentive effects to savings. Consistent with recent research on voting behaviour (see Conde-Ruiz and Profeta, 2003), preliminary analysis suggests that low income and high income households prefer a limited means-tested pension system, while middle income groups prefer a universal pension system. In another recently considered application, the model has been used to undertake preliminary analysis into the extent to which international differences in retirement behaviour can be explained by differences in the pension policies of the respective countries, and to consider the associated distributional implications.

All of these issues analyse behavioural responses. They therefore cannot be coherently investigated except by using the framework we describe here.

7. Directions for further research

The data that are required to evaluate econometrically household responses to alternative pension policy environments are limited in many respects. This has motivated the development of microsimulation tools for analysis. The model that is described by the current paper was created to analyse household savings and labour responses to pension policy counterfactuals. Simulation models are rarely developed to consider behavioural responses due to the computational complexities that are involved. In contrast, the model described here uses a behaviourally consistent dynamic programming framework to simulate the consumption/savings, and the labour/leisure decisions of households.

The programming architecture of the microsimulation model has been designed with computational efficiency in mind, and this has achieved a manageable run-time of 3.5 hours on a personal computer.[26] Nevertheless, the computational intensity of the problem has meant that it was necessary to impose non-trivial stylisations on the simulated characteristics. One of the most important simplifications assumed by the model concerns the simulation of household demographic characteristics, which evolve in an exogenous fashion, and so omit the possibility of behavioural responses to a substantial aspect of real-world uncertainty. It is with regard to these demographic characteristics that future research effort regarding the model described here is likely to focus. There is good reason to be confident that, with modest improvements on existing computer technology, the full range of demographic heterogeneity that is currently described by non-behavioural dynamic microsimulation models can be incorporated into the current fully behavioural framework.

Appendix A. European Community Household Panel

The European Community Household Panel Survey (ECHP) is 'the most closely co-ordinated component of the European system of social surveys' that are collected by Eurostat (the statistical office of the European communities). The ECHP provides detailed panel data for households that are drawn from fifteen European Community counties, spanning the period between 1994 and 2001 (the most recent year for which data have been made available). The data are collected at annual intervals, and so build up an historical record of 60,500 nationally representative households.[27]

The ECHP data that are considered in this paper have been sourced by Eurostat from the British Household Panel Survey (BHPS). The BHPS is a panel survey of households that were originally selected to provide a nationally representative sample of the UK population.[28] The first wave of the survey was undertaken in 1990, and includes information for 13,840 individuals drawn from 5,511 households. Subsequent waves have been undertaken annually, to provide a survey history for individuals who were approached in the original wave (and their subsequent households). The most recent wave released by the Office for National Statistics (ONS) supplies data for the year 2000/01 (the tenth consecutive wave). The variables used to undertake the analysis presented here were extracted from the ECHP using SPSS programs. The authors may be contacted for further details.

Appendix B. Probit analysis of employment status

Table 2. Probit regression of full-time employment

Parameter	Estimate	Std error
C	–12.17570	(3.47E+00)
AGE	0.99082	(3.48E–01)
AGE2	–0.03158	(1.26E–02)
AGE3	4.60E–04	(1.95E–04)
AGE4	–2.69E–06	(1.10E–06)
NC	–0.33892	(3.24E–02)
MALE	1.01702	(5.81E–02)
COUPLE	0.03831	(6.83E–02)
CAR	0.48719	(8.45E–02)
ROOMST	0.03914	(1.87E–02)
H45TTM3	–1.43907	(2.28E–01)
H12TTM3	0.23211	(5.72E–02)
correct predictions	0.78782	

Appendix C. Simulating UK tax and benefits policy

UK pension policy

The current UK pension system is comprised of three tiers. The first tier consists of the Basic State Pension, BSP; the second tier of all government run contributory pension benefits (the State Earnings Related Pension Scheme, SERPS, and the Second State Pension, S2P); and the third tier of all private pension schemes. Furthermore, Incapacity Benefit is a commonly used vehicle to fund early retirement.

The following describes each system, as it stood in 2003. Simulated households are assumed to draw upon the BSP and the S2P, as they are described here, from age 65 (the State Pensionable Age, SPA). The potential role of incapacity benefit to fund early retirement is modelled using a stylised specification that is described in the following subsection.

- **Basic State Pension**: The full BSP, equal to £77.45 per week for a single person and £123.80 for a couple (in 2003), is paid to individuals who have been accredited with qualifying years for approximately 90 per cent of their working lives. A qualifying year is defined as one in which an individual has earned an annual income that exceeds the Lower Earnings Limit, equal to £4,004, and also includes years of unemployment, or incapacitation. This implies that most households qualify for the full BSP. For simplicity, the BSP is consequently modelled as a universal benefit. BSP is funded by PAYGO contributions of current employees. Specifically, annual income earned between £4,628 (the Employees' Earnings Threshold, EET, as at 2003), and £30,940 (the Upper Earnings Limit, UEL, as at 2003) is subject to National Insurance Contributions (NICs) of 8.95 per cent to fund the BSP.[29] The taxation of income during the working lifetime is discussed in the following subsection.

- **Second State Pension:** Until recent reforms, the benefit payable under the second tier of the UK pension system was entirely related to an individual's average earnings over his/her working lifetime. Membership to the second tier state pension is compulsory for all employees (but not the self-employed), unless the employee has contracted out into a private pension scheme. The second tier system was administered under the SERPS until April 2002, when it became the S2P. Upon reaching SPA, the wages earned by an individual during each year of his/her working life are rescaled by average wage growth, and the average determined. The average wages earned between £4,004 and £11,200 (in 2003) are multiplied by 0.46, wages between £11,201 and £25,600 are multiplied by 0.115, and wages between £25,601 and £30,940 are multiplied by 0.23.[30] The aggregate of these values determine the individual's annual S2P benefit. Unlike SERPS, individuals with incomes below the lower earnings threshold (£11,200 per year in 2003) earn S2P entitlements as if their income was at the lower earnings threshold. The S2P is PAYGO, funded through contributions of current workers at a rate of 1.6 per cent on income earned between the UEL and the EET.[31]

Underlying the BSP and the S2P is the Pension Credit (PC), which guarantees anyone aged 60 or over an income of at least £102.10 per week or £155.80 per week for a couple (including the BSP). The PC applies a taper rate of 40 per cent on gross private income in excess of the full BSP. The PC is also subject to an assets test. The first £6,000 of assets are ignored, but thereafter an income is imputed to any savings above this threshold at a rate of 10 per cent a year.

- **Private and Occupational Pensions**: The third tier of the UK pension system is comprised of private pension schemes, of which there are two types: occupational pensions and personal pensions. Contributions into these schemes are made out of pre-tax income, so that contributions are effectively subsidised (at the basic tax rate) by the Government. An occupational pension can usually be classified as either a 'defined benefit' scheme (where the benefits are earnings related), or as a 'defined contribution' scheme (where the benefits are related to the value of the accumulated contributions). Personal pensions are always run on a defined contribution basis. Occupational pensions play an important role in the UK pension system – forming one half of the so-called public-private partnership they account for approximately 50 per cent of total pension entitlements.[32] Private and Occupational pensions are simulated as a form of discretionary saving, with 55 per cent subject to forced annuitisation from age 65.[33]

The working lifetime

The model is specified to focus attention upon the behavioural effects of state provided pensions. Consequently, stylised specifications are used to simulate the impact of tax and benefits policy during the working lifetime. It should be noted, however, that the stylised methods used to simulate tax and benefits policy were not adopted in response to limitations of the Dynamic Programming framework – indeed more complex specifications based upon the official rates and thresholds of the UK transfer system have been considered elsewhere.[34] The stylised specifications described here were assumed to facilitate cross-country comparisons for which the model was constructed.

Three functions are used to simulate the impact of tax and benefits policy during the working lifetime; one for the employed, one for those not-employed under age 51, and another for those not-employed under age 65. Older not-employed are distinguished from younger not-employed, to take into consideration the effects of early retirement vehicles (such as incapacity benefit as discussed in the preceding section).[35] All three functions are specified with respect to the number of adults, to the number of children, and to pre-tax and benefit (hereafter referred to as pre-tax) income. Equation (15a) defines the specification adopted for the employed, equation (15b) for the not-employed under age 51, and equation (15c) for the not-employed between ages 51 and 64.

$$y_i = (\beta_{00} + \beta_{01}na_i + \beta_{02}nc_i + \beta_{03}(na_i + nc_i)^2) + (\beta_{04} + \beta_{05}na_i)x_i \quad (15a)$$

$$y_i = (\beta_{10} + \beta_{11}na_i + \beta_{12}nc_i) + \beta_{14}x_i \quad (15b)$$

$$y_i = (\beta_{20} + \beta_{21}na_i + \beta_{22}d_{63i} + \beta_{23}d_{64i}) + \beta_{24}x_i \quad (15c)$$

where y_i denotes the post-tax income of household i, na_i the number of adults, nc_i the number of children, and x_i the pre-tax income. d_τ are dummy variables that take the value one if age equals τ and zero otherwise. These specifications were selected after trialling various alternatives.[36] Estimates for the coefficients of equations (15a) to (15c) were obtained using UK household level microdata for 2000/01 derived from the ECHP, and are reported in table 3.

The functions adopted to simulate the tax and benefits system during the working lifetime are obviously highly stylised. Nevertheless, they manage to capture much of the variation described by the ECHP survey data, as indicated by the high R-square statistics.

Table 3. Country specific tax and benefit estimates

Parameter	Estimate	Std error
Employed		
C	510.62	(801.5)
NA	7776.90	(481.4)
NC	2638.40	(653.3)
(NA+NC)^2	–299.43	(96.1)
X	0.47559	(2.60E–02)
NA*X	–0.03221	(8.85E–03)
R-square	0.88914	
std error	7001.51	
Not employed – age 20–50		
C	5824.84	(1175.6)
NA	2346.85	(812.1)
NC	1726.32	(301.8)
X	0.34482	(1.46E–02)
R-square	0.69807	
std error	5957.37	
Not employed – age 51–64		
C	1702.07	(1759.2)
NA	6181.22	(1186.0)
D63	6555.53	(2609.2)
D64	5838.84	(1708.9)
X	0.53963	(1.39E–01)
R-square	0.63396	
std error	9502.58	

NOTES

1 Lucas (1976).
2 For macro-based models that study the impact of policy changes, see Dervis *et al.* (1982), Taylor (1990), and De Janvry *et al.* (1991). These are examples of Computable General Equilibrium models. Most micro-based models are constructed using a partial equilibrium framework. For examples of micro-based models that use a general equilibrium framework, see Meagher (1993), and Cogneau and Robilliard (2000).
3 For useful surveys, refer to Zaidi and Rake (2001), Sutherland (1995), and Merz (1991).
4 See Macunovich (1998), and Hotz *et al*. (1997) for surveys of the fertility literature, Auerbach (1997) on savings, and Debelle and Swann (1998) on trends in the Australian labour market.
5 This assumption is made by Stock and Wise (1990) in view of the fact that the data set they used for estimation does not describe wealth holdings. As noted in Section D of Stock and Wise (1990), the Option Value model also imposes a more technical simplification on the retirement decision. Specifically, it assumes that individuals consider the value of future utility in terms of the maximum of the expected value of future alternatives, rather than the expectation of the maximum of future alternatives. This essentially means that an individual fails to take into consideration the fact that he/she can adapt his/her choices to new information obtained in the future when evaluating their expectations regarding the utility value of future options.
6 This point is not new. See, for example, Moffitt (1987, p. 185), cited by Kruger and Meyer (2002).
7 See, for example, review by Kruger and Meyer (2002).
8 See, for example, Gruber and Wise (1999).
9 It is clear that for most analytical problems, neither extreme of the data continuum referred to above is likely to be observed in practice. In such circumstances it seems pertinent to compare information that is drawn from alternative analytical approaches.
10 Also referred to as the *policy rule*.
11 See, for example, Rust (1987) for an example of an attempt to econometrically estimate a DP model that includes decisions regarding labour supply and consumption. See also, Rust and Phelan (1997), and Gustman and Steinmeier (2001). In preliminary work, van der Klaauw and Wolpin (2003) suggest an interesting hybrid econometric/calibration approach.
12 See the *Family Expenditure Survey 2000–2001 User Guide*, Vol. I for the definition of a household reference person.
13 An alternative version of the model allows households to work part-time. This option is omitted here to focus attention upon the issue of retirement.
14 For models of endogenous fertility, see Nerlove *et al*. (1984), and Barro and Becker (1989).
15 The estimates are based on household arithmetic means derived from European Community Household Panel (ECHP) data, recorded for the UK during 2000/01. See Appendix A for details regarding the data used. The raw averages were smoothed using non-parametric methods (the KSM procedure, bandwidth = 0.2, in STATA).
16 The mortality rates used were calculated using the proportion of female reference people by age recorded in the 2000/01 Family Expenditure Survey (FES), and mortality rates by age and sex recorded in the *Annual Abstract of Statistics*, Table 5.21, The Stationary Office. Mortality rates after the age of 84 are subject to manual adjustment to ensure simulated death by age 90. The use of backward induction to solve the DP problem that is considered by the microsimulation model necessitates truncation of the simulated lifetime at a terminal age.
17 See also, Kalecki (1945) and Creedy (1985).
18 The specification of the probit model was arrived at after considering a range of alternatives.
19 Equation (9) is obtained from equation (6) by recursive substitution of h_{it-s}, s = 1, ..., R–1.
20 For details regarding the Family Expenditure Survey (FES) see *Family Expenditure Survey 2000–2001 User Guide*. London: Office for National Statistics. Statistics calculated from FES data are reported for comparison.
21 FES data were used to specify the simulated geometric mean for 20 year olds due to small samples observed in ECHP data.
22 As such, involuntary unemployment is not considered by the model.
23 See, for example, Balcer and Sadka (1986), and Muellbauer and van de Ven (2003) on the use of this form of adjustment for household size in the utility function. The model uses the McClements scale to equivalise consumption. See, for example, ONS *Social Trends* 28 (1998) on the McClements equivalence scale.
24 It is a simple matter to incorporate temporal variation for π.
25 The Gauss-Hermite discrete approximation is used to calculate expectations. See Sefton (2000) for technical details regarding the numerical solution algorithm used.
26 Pentium 4 1.5 GHz processor with 512 MB of RAM.
27 See Eurostat (2003) for further details regarding the ECHP.
28 Due to the repeated survey methods employed, the most recent wave of the BHPS no longer provides a representative

sample of the UK population. See Taylor (2002) for further details regarding the BHPS.

29 The total NIC charged is 11 per cent, 2.05 per cent of which is used to fund the National Health Service (NHS). Employers are also required to pay NICs above the EET, at a rate of 12.8 per cent (of which, 10.9 per cent is used to fund the BSP).

30 For example, if an individual earned the equivalent of £40,000 in one year, then he/she would be credited with £6,194.02 = 0.46*(11,200–4,004) + 0.115*(25,600–11,201) + 0.23*(30,940–25,601).

31 Contracting-out' rebate on NICs.

32 See, for example, Blake and Orszag (1999, table 12).

33 Based upon wealth data calculated from the BHPS.

34 See Sefton and van de Ven (2003).

35 The age threshold was selected with reference to observations drawn from survey data, which suggest that early retirement becomes increasingly prevalent from age 51 (see, for example, Blundell *et al*., 2002), and Figure 15a in Section 5.

36 The authors may be contacted for details.

REFERENCES

Andersson, I., Brose, P., Flood, L., Klevmarken, N.A., Olovsson, P. and Tasizan, A. (1992), 'MICROHUS – a microsimulation model of the Swedish household sector. A progress report, paper presented at the International Symposium on Economic Modelling, Gothenburg.

Atkinson, A.B., Bourguignon, F. and Morrisson, C. (1992), *Empirical Studies of Earnings Mobility*, Reading, Harwood Academic Publishers.

Auerbach, A.J. (1997), *Fiscal Policy: Lessons from Economic Research*, London, MIT Press.

Balcer, Y. and Sadka, E. (1986), 'Equivalence scales, horizontal equity and optimal taxation under utilitarianism', *Journal of Public Economics*, 29, pp. 79–97.

Barro, R.J. and Becker, G.S. (1989), 'Fertility choice in a model of economic growth', *Econometrica*, 57, pp. 481–501.

Basu, N., Pryor, R. and Quint, T. (1998), 'ASPEN: a microsimulation model of the economy', *Computational Economics*, 12, pp. 223–41.

Blake, D. and Orszag, M. (1999), 'Annual estimates of personal wealth holdings in the UK since 1948', *Applied Financial Economics*, 99, pp. 397–421.

Blau, D.M. (1994), 'Labour-force dynamics of older men', *Econometrica*, 62, pp. 117–56.

Blau, D.M. and Gilleskie, D.B. (2001), 'Retiree health insurance and the labor force behavior of older men in the 1990s', *Review of Economics and Statistics*, 83, pp. 64–80.

Blundell, R. and Emmerson, C. (2003), 'Fiscal effects of reforming the UK state pension system', Institute for Fiscal Studies Working Paper, WP03/13.

Blundell, R., Meghir, C. and Smith, S. (2002), 'Pension incentives and the pattern of early retirement', *Economic Journal*, 112, pp. C153–70.

Caldwell, S. (1997), *Corsim 3.0 User and Technical Documentation*, New York, Ithaca.

Callan, T., O'Donoghue, C. and O'Neill, C. (1996), *Simulating Welfare and Income Tax Changes: The ESRI Tax-Benefit Model*, Dublin, The Economic and Social Research Institute.

Cantor, A. and Sefton, J. (2002), 'Economic applications to actuarial work: personal pensions and future rates of return', *British Actuarial Journal*, 8, 1, pp. 91–131(41).

Cogneau, D. and Robilliard, A.S. (2000), 'Growth, distribution and poverty in Madagascar: learning from a microsimulation model in a general equilibrium framework', Trade and Macroeconomics Division, International Food Policy Research Institute, Discussion Paper no. 61.

Conde-Ruiz, J.I. and Profeta, P. (2003), 'What social security: Beveridgean or Bismarckian?', FEDA Working Paper no. 2003–16.

Creedy, J. (1985), *Dynamics of Income Distribution*, Oxford,Basil Blackwell.

Deaton, A. (1993), *Data and Econometric Tools for Analysis*, Princeton, Research Program in Development Studies.

Debelle, G. and Swann, T. (1998), 'Stylised facts of the Australian labour market', Reserve Bank of Australia Research Discussion Paper no. 9804.

Dervis, K., De Melo, J. and Robinson, S. (1982), *General Equilibrium Models for Development Policy*, Cambridge, Cambridge University Press.

de Janvry, A., Sadoulet, E. and Fargeix, A. (1991), 'Politically feasible and equitable adjustment: some alternatives for Ecuador', *World Development*, 19, pp. 1577–94.

Dutta, J., Sefton, J. and Weale, M. (2000), 'Capital taxation and public choice', National Institute of Economic and Social Research Discussion Paper no.

Eurostat (2003), *European Community Household Panel Longitudinal Users' Database*, Eurostat, DOC.PAN 168/2003–06.

Fjaerli, J.A.E., Gravningsmyhr, H., Holmøy, A.M.K. and Lian, B. (1995), 'The Norwegian Microsimulation Model LOTTE: applications to personal and corporate taxes and social security benefits', Microsimulation Unit Discussion Paper MU9504.

Giannarelli, L. (1992), *An Analyst's Guide to TRIM2, the Transfer Income Model, Version 2*, Washington DC, Urban Institute Press.

Gruber, J. and Madrian, B. (1995), 'Health insurance availability and early retirement decision', *American Economic Review*, 85, pp. 938–48.

Gruber, J. and Wise, D.A. (1999), *Social Security and Retirement Around the World*, Chicago, Chicago University Press.

Gustman, A.L. and Steinmeier, T.L. (2001), 'Retirement and wealth', NBER Working Paper no. 8229.

Harding, A. (1993), *Lifetime Income Distribution and Redistribution: Applications of a Microsimulation Model*, London, North-Holland.

Hotz, V.J., Klerman, J.A. and Willis, R.J. (1997), 'The economics of fertility in developed countries', in *Handbook of Population and Family Economics*.

Huggett, M. (1996), 'Wealth distribution in life-cycle models', *Journal of Monetary Economics*, 38, pp. 469–94.

Kalecki, M. (1945), 'On the Gibrat distribution', *Econometrica*, 13, pp. 161–70.

King, A., Baekgaard, H. and Robinson, M. (1999), 'DYNAMOD-2: an overview', NATSEM Technical Paper, 19.

Kruger, A. and Meyer, B.D. (2002), 'Labor supply effects of social insurance', NBER Working Paper no. 9014.

Lucas, R.E. (1976), 'Econometric policy evaluation: a critique', *Carnegie-Rochester Conference Series on Public Policy*, 1, pp. 19–46.

Macunovich, D.J. (1998), 'Fertility and the Easterlin hypothesis: an assessment of the literature', *Journal of Population Economics*, 11, pp. 53–111.

Meagher, G.A. (1993), 'Forecasting changes in the income distribution: an applied general equilibrium approach', in Harding, A. (ed.), *Microsimulation and Public Policy*, Amsterdam, Elsevier.

Merz, J. (1991), 'Microsimulation – a survey of principles, developments, and applications', *International Journal of Forecasting*,

7, pp. 77–104.
Mincer, J. (1974), *Schooling, Experience, and Earnings*, New York, Columbia University Press.
Moffitt, R.A. (1987), 'Life-cycle labour supply and social security: a time-series analysis', in Burtless, G. (ed.), *Work, Health, and Income among the Elderly*, Washington DC, Brookings Institute.
Muellbauer, J. and van de Ven, J. (2003), 'Equivalence scales and taxation', NIESR Discussion Paper no. 229.
Nerlove, M., Razin, A. and Sadka, E. (1984), 'Income distribution policies with endogenous fertility', *Journal of Public Economics*, 24, pp. 221–30.
Orcutt, G. (1957), 'A new type of socio-economic system', *Review of Economics and Statistics*, 58, pp. 773–97.
Orcutt, G., Caldwell, S.B. and Wertheimer, R. (1976), *Policy Exploration Through Microanalytic Simulation*, Washington DC, Urban Institute.
Redmond, G., Sutherland, H. and Wilson, M. (1998), *The Arithmetic of Tax and Social Security Reform: A Users' Guide to Microsimulation Methods and Analysis*, Cambridge, Cambridge University Press.
Rust, J. (1987), 'A dynamic programming model of retirement behaviour', NBER Working Paper no. 2470.
Rust, J. and Phelan, C. (1997), 'How social security and medicare affect retirement behaviour in a world of incomplete markets', *Econometrica*, 65, pp. 781–831.
Sefton, J.A. (2000), 'A solution method for consumption decisions in a dynamic stochastic general equilibrium model', *Journal of Economic Dynamics and Control*, 24, pp. 1097–119.
Sefton, J., Dutta, J. and Weale, M. (1998), 'Pension finance in a calibrated model of saving and income distribution for the UK', *National Institute Economic Review*, 166, pp. 97–107.
Sefton, J. and van de Ven, J. (2003), 'Does means testing exacerbate early retirement?', mimeo, NIESR.
Sefton, J. and Weale, M. (2003), 'State pensions, means testing and public choice', in Weale, M. (ed.), *Pension Reform: Redistribution and Risk*, London, National Institute of Economic and Social Research Occasional Paper, 56.
Stock, J.H. and Wise, D.A. (1990), 'Pensions, the potion value of work, and reitrement', *Econometrica*, 58, pp. 1151–80.
Sutherland, H. (2001), 'EUROMOD: an integrated European benefit-tax model', EUROMOD Working Paper, EM9/01 Final Report.
—(1995), 'Static microsimulation models in Europe: a survey', University of Cambridge Department of Applied Economics Working Paper, Amalgamated Series no. 9523.
Taylor, L. (1990), *Socially Relevant Policy Analysis. Structural Computable General Equilibrium Models for the Developing World*, Cambridge, Mass., MIT Press.
Taylor, M.F. (2002) *British Household Panel Survey Users Manual*, Colchester, University of Essex.
van der Klaauw, W. and Wolpin, K.I. (2003), 'Social security, pensions and the savings and retirement behavior of households', mimeo, University of North Carolina.
van de Ven, J. (1998), 'A dynamic cohort microsimulation model', University of Melbourne Department of Economics Working Paper.
Zaidi, A and Rake, K. (2001), 'Dynamic microsimulation models: a review and some lessons for SAGE', SAGE Discussion Paper, SAGEDP/02.