



Parameterising a detailed dynamic programming model of savings and labour supply using cross-sectional data

Justin W. van de Ven

NIESR, 2 Dean Trench Street, London SW1P 3HE, UK
MIAESR, The University of Melbourne, Australia
jvandenven@niesr.ac.uk

ABSTRACT: Dynamic programming methods are now commonly used to describe behaviour in contexts where uncertainty is likely to have an important bearing on decision making. Using a publicly available structural dynamic microsimulation model, LINDA, this paper provides new insights into how unobservable preference parameters – particularly those associated with risk aversion – can be coherently identified on broad-based moments of decision making observed for a population cross-section. Preference parameters identified on UK data are found to be in-line with those reported in the wider econometric literature.

KEYWORDS: Dynamic Programming, Savings, Labor Supply, Empirical Identification

JEL classification: C51, C61, C63, H31

1. INTRODUCTION

Empirical analysis of intertemporal decision making is complicated by the effects of uncertainty on incentives. Where uncertainty is considered sufficiently important to warrant a central place in a structural model, then dynamic programming methods are now commonly employed. Studies of savings behaviour in this vein often limit the computational burden by focussing upon the evolving circumstances of individual birth cohorts. The computational advantage that is gained by limiting a dynamic programming model to focus on a single birth cohort is, however, off-set by at least two attendant complications. The first is that the savings behaviour of a single birth cohort is only revealed over a substantial period of time, which complicates the task of capturing time-varying incentives described by the evolving economic context. Furthermore, it is questionable whether empirical results obtained for a single birth cohort will generalise to the wider population. This paper reports empirical results for a dynamic programming model that avoids both of these problems by projecting the circumstances of a population cross-section through time, which permits identification of saving preferences on cross-sectional survey data. The results obtained demonstrate the feasibility and advantages for empirical analysis of the cross-sectional approach for modelling savings decisions in context of uncertainty.

A complex two-dimensional relationship exists between time, cohort, and age effects that characterise differences between heterogeneous population subgroups. Focussing upon the evolving circumstances of a single birth cohort is a useful way for empirical studies to cut through this complexity, as age, time and cohort effects are then described by a single dimension. Such a simplification is particularly appealing where the central subject of interest is complex, as is often the case for decision problems that have no closed form solution. This is a principal reason why the dynamic programming literature that explores savings behaviour in the context of uncertainty has focussed predominantly upon empirical analysis of cohort-specific structural models, following the seminal study by Gourinchas and Parker (2002).

Alternative data options exist for empirical analyses that focus on cohort-specific structural models of savings behaviour. An obvious choice is to parameterise a cohort-specific model with reference to data observed for a single birth cohort (O. Attanasio, Low, & Sanchez-Marcos, 2005).¹ This approach imposes a somewhat heavy burden on the time-frame of survey data required for analysis, and is usually complicated by the associated difficulty of obtaining an adequate description of the evolving policy context. An additional problem is that it is uncertain how far results obtained for a single birth cohort can be generalised to the wider population. These drawbacks stem from fundamental features of the empirical problem in relation to savings behaviour. An empirical analysis of savings decisions in context of uncertainty requires for identification data observed for an appreciable period of life. The longer is the period from which data for analysis are drawn, the greater is the scope for substantive variation of the policy environment underlying observed behaviour. The greater is the variation of the policy environment over multiple dimensions, the stronger is the proposition that such variation is likely to be an important determinant of observed behaviour.

Aspects of the policy environment that typically exhibit substantial variation with time, and which are likely to influence savings decisions of the household sector, include taxes and benefits, (pre-transfer) rates of return, variation of employment opportunities, and the changing nature of family demographics. Obtaining comprehensive (pseudo) panel data regarding all of these factors usually represents a significant challenge, and integrating

these data into a structural model in a coherent fashion is more challenging still. Furthermore, allowing for such variation can work to off-set any computational advantage that is derived from focussing on the circumstances of a single birth cohort. I am not aware, for example, of any dynamic programming model of household sector savings that includes an explicit account of reforms to tax and benefits policy implemented during the period of estimation.² Such complications hamper efforts to reflect adequately the savings and employment incentives that individuals face.

One popular way to identify results that generalise to the wider population is to conduct sensitivity analysis by exploring data reported for alternative birth cohorts (as in O. Attanasio, Low, and Sanchez-Marcos (2008)). Such an approach complicates the challenges involved in adequately describing the evolving policy context. Alternatively, empirical techniques can be used that are designed to estimate age-specific moments which control for time and cohort effects (Sefton, van de Ven, & Weale, 2008). Collinearity between age, cohort and time effects requires an additional restriction to permit identification. One common restriction, suggested by Deaton (1997), is to assume that time effects average out over the long run. This assumption produces estimated age profiles that represent an average taken over all cohorts included in the panel data used for estimation. The averaging that such methods apply obscures the nature of the underlying policy environment, so that it is difficult – if not impossible – to ensure that the assumed structural specification provides an adequate representation of the incentives underlying observed behaviour.

A third approach that has been applied in the literature is designed to simplify identification of the incentives that underly observed behaviour, which is the principal drawback associated with the two alternatives referred to above. In this case, empirical analysis is based upon cross-sectional data that are adjusted to reflect assumptions about the relationship between the characteristics of the population cross-section and those of a single birth cohort (van de Ven & Weale, 2010). Focussing on cross-sectional data limits the incentives underlying observed behaviour to those that applied at a single point in time, which are relatively simple to document. The drawback of this approach, however, is that strong assumptions are required to derive a stylised relationship between the characteristics of the population cross-section and those of a single birth cohort; assumptions that are unlikely to hold in practice.³

This paper explores the proposition that an overlapping generations (OLG) model structure can simplify coherent identification of the preference parameters of a dynamic programming model of savings and labour supply. This proposition is based on the observation that an OLG structural model can describe behaviour observed throughout the life-course at a single point in time, albeit for individuals drawn from different birth cohorts. If it is assumed that preference parameters are stable across generations, then this implies that the parameters of an OLG modelling approach can be identified entirely on cross-sectional survey data. This simplifies both the task of gathering the data necessary for empirical analysis, and the model description of the policy environment underlying the considered survey data.

Although OLG models of savings in context of uncertainty are not new (Livshits, MacGee, & Tertilt, 2007; G. D. Hansen & Imrohoroglu, 2008; Feigenbaum, 2008; Hairault & Langot, 2008), most of the associated literature focus on behavioural implications of alternative theoretical frameworks, rather than on the task of empirical identification. This study focusses squarely on the issue of empirical identification. The study is based on LINDA – the Lifetime INcome Distributional Analysis – model, designed to reflect the UK policy context as

observed in 2011. LINDA is a full structural dynamic microsimulation model that uses dynamic programming methods to project savings and employment behaviour in a way that takes risk aversion fully into account. This makes the model particularly useful for exploring preferences concerning risk. The model accommodates a rich description of agent specific heterogeneity, and is publicly available so that results should be replicable by the reader. The analysis that is reported demonstrates the feasibility of the proposed estimation strategy in context of contemporary computing technology, and the model parameters obtained are found to relate closely to those reported in the broader empirical literature.

Section 2 provides an overview of the LINDA model upon which the analysis is based. The analytical mechanics that underly the empirical approach considered here are described in Section 3. Data are described in Section 4, and results are reported in Sections 5 and 6. Discussion of results focusses on drawing out the ways that preference parameters influence alternative observable margins, which are crucially important for parameter identification. A concluding section provides a summary and directions for further research.

2. THE LINDA MODEL

LINDA is a dynamic programming model of household sector savings and labour supply decisions. LINDA is a country-specific adaptation of a flexible model framework, referred to as the Simulator of Individual Dynamic Decisions (SIDDD), which has been developed to make current best practice economic methods of analysis of savings and labour supply available to non-specialists. SIDDD is currently free for all practitioners to download via the website: www.simdynamics.org. The remainder of this text refers solely to LINDA, to avoid any potential confusion. This section provides a brief overview of the subset of model features that are employed for empirical identification; see van de Ven (2016) for a full technical description of the model.

The decision unit of the model is the family, defined as a single adult or partner couple and their dependent children. LINDA considers the evolving circumstances of a sample of reference adults and their benefit units, organised into annual snap-shots during the life-course. Allocations within families are ignored. Decisions regarding consumption, labour supply, pension scheme participation, and timing of pension access are endogenous, and are assumed to be made to maximise expected lifetime utility, given a family's prevailing circumstances, its preference relation, and beliefs regarding the future. Preferences are described by a nested Constant Elasticity of Substitution utility function. Expectations are 'substantively-rational' in the sense that they are either perfectly consistent with, or specified to approximate, the intertemporal processes that govern individual characteristics. The model assumes a small open economy (appropriate for the UK), for which rates of return to labour and capital are exogenously given. Heterogeneous circumstances of reference adults are limited to the following twelve characteristics:

- | | | |
|--------------------------------|-----------------------------|----------------------------|
| - year of birth | - age | - relationship status |
| - number of dependent children | - age of dependent children | - student status |
| - education status | - private pension wealth | - timing of pension access |
| - non-pension wealth | - wage potential | - survival |

Seven of the characteristics listed here are considered to be uncertain and uninsurable from one year to the next

when evaluating expected lifetime utility (relationship status, number and age of dependent children, student status, education status, wage potential, and time of death). This specification for the model was carefully selected to ensure adequate margins for empirical identification of unobserved preference parameters. Including year of birth in the list of heterogeneous family characteristics introduces the overlapping generations framework that is necessary to reflect the circumstances of a population cross-section. Age, wage potential, measures of wealth, and survival are all centrally important for any empirical analysis of savings and labour supply. Relationship status and children are important for reflecting the influence of tax and benefits policy, and for capturing labour supply and consumption decisions. Finally, as discussed in Section 3, education status and pension scheme participation decisions feature in the empirical identification strategy employed in this paper.

2.1. Preference relation

Expected lifetime utility of reference adult i , with birth year b , at age a is described by the time separable function:

$$U_{i,a} = \frac{1}{1-\gamma} \left\{ u \left(\frac{c_{i,a}}{\theta_{i,a}}, l_{i,a} \right)^{1-\gamma} + E_{a,b} \left[\sum_{j=a+1}^A \beta^{j-a} \left(\phi_{j-a,a}^b u \left(\frac{c_{i,j}}{\theta_{i,j}}, l_{i,j} \right)^{1-\gamma} + (1 - \phi_{j-a,a}^b) \zeta B_{i,j}^{1-\gamma} \right) \right] \right\} \quad (1a)$$

$$u \left(\frac{c_{i,a}}{\theta_{i,a}}, l_{i,a} \right) = \left(\left(\frac{c_{i,a}}{\theta_{i,a}} \right)^{(1-1/\varepsilon)} + \alpha^{1/\varepsilon} l_{i,a}^{(1-1/\varepsilon)} \right)^{\frac{1}{1-1/\varepsilon}} \quad (1b)$$

Observable characteristics of the preference relation are $\phi_{j-a,a}^b$ the probability that a reference adult with birth year b will survive to age j given survival to age a ; $c_{i,a} \in R^+$ discretionary composite (non-durable) consumption; $l_{i,a} \in [0, 1]$ the proportion of family time spent in leisure; $\theta_{i,a} \in R^+$ adult equivalent size based on the “revised” or “modified” OECD scale⁴; and $B_{i,a} \in R^+$ the legacy that reference adult from family i would leave if they died at age a . Unobserved preference parameters are $\gamma > 0$ the (constant) coefficient of relative risk aversion; β an exponential discount factor; ζ the “warm-glow” model of bequests⁵; $\varepsilon > 0$ the (intra-temporal) elasticity of substitution between equivalised consumption ($c_{i,a}/\theta_{i,a}$) and leisure ($l_{i,a}$); and $\alpha > 0$ the utility price of leisure. $E_{a,b}$ is the expectations operator and A is the assumed maximum age that any individual can survive to.

Although the preference relation defined by equation (1) is popular in the associated literature, it has also been the subject of considerable criticism. Four points can be singled out here. First, the assumption of time separability suppresses behavioural persistence, and has been the subject of an extensive debate (e.g. Deaton and Muellbauer (1980), pp. 124–5; Hicks (1939), p. 261). Any empirical study concerned with inter-relations between decisions through time would need to consider data observed thorough time, in contrast to the cross-sectional data that are the focus of the empirical application considered here. Second, it is now increasingly common to allow preference parameters, including discount rates, to vary with individual specific characteristics (e.g. Gustman and Steinmeier (2005), who consider variation in relation to discount rates, and the relative attractiveness of alternative employment options). The associated literature suggests that suppressing this form of

variation in an empirical analysis of preferences can be interpreted as a form of omitted variable bias. Third, the assumption that preferences are time consistent – as is implied by the preference relation defined by equation (1a) – has been criticised for failing to reflect a growing body of empirical evidence (Thaler, 1981; G. Ainslie & Haslam, 1992; Green, Fry, & Myerson, 1994; Kirby, 1997; G. W. Ainslie, 1992). Adapting preferences to accommodate time-inconsistency has also been shown to affect behavioural margins that are used for empirical identification in this study (Laibson, Repetto, & Tobacman, 2007; van de Ven & Weale, 2010). Finally, the assumption of a CES specification for intertemporal preferences has been criticised for the restrictions that it imposes upon the relationship between relative risk aversion and the intertemporal elasticity of consumption.

The preference relation defined by equation (1) remains predominant in the associated literature despite the limitations set out in the preceding paragraph. This is because relaxing the model along any one of the four points referred to above would expand the state-space and/or the number of preference parameters. Expanding the state space of the decision problem implies a (geometric) increase in the computational burden of the utility maximisation problem, which exaggerates the limitations of existing computing technology. Increasing the dimensionality of the (unobservable) preference parameters of the model places an increased burden on the data and numerical methods used for parametrisation. The preference relation defined by equation (1) is consequently a trade-off between parsimony and computational burden, which remains sensible given contemporary computing resources and the most common analyses that the model has been devised to explore.

2.2. Labour income dynamics

Earnings are modelled at the family level, and are described by:

$$\begin{aligned} g_{i,a} &= \max(h_{i,a}, h_{a,t}^{\min}) \lambda_{i,a} & (2) \\ \lambda_{i,a} &= \lambda_{i,a}^o \lambda_{i,a}^{emp} \lambda_{i,a}^{ret} \end{aligned}$$

where $h_{i,a}$ defines family i 's latent wage at age a , $h_{a,t}^{\min}$ is the (statutory) minimum wage, λ^o is an adjustment factor to allow for uncertain wage offers, λ^{emp} adjusts for (endogenous) labour supply decisions, and λ^{ret} is the impact on earnings of taking up private pension income.

In most periods, latent wages h are assumed to follow a random-walk with drift:

$$\log\left(\frac{h_{i,a}}{m_{i,a}}\right) = \log\left(\frac{h_{i,a-1}}{m_{i,a-1}}\right) + \omega_{i,a-1} \quad (3a)$$

$$m_{i,a} = m(n_{i,a}, ed_{i,a}, a, b) \quad (3b)$$

$$\omega_{i,a} \sim N(0, \sigma_{\omega}^2(n_{i,a}, ed_{i,a})) \quad (3c)$$

where the parameters $m(\cdot)$ account for wage growth, which in turn depend on relationship status $n_{i,a}$, education $ed_{i,a}$, age a , and birth year b , and $\omega_{i,a}$ is an identically and independently distributed family specific disturbance term. The variance σ_{ω}^2 is defined as a function of relationship status and education. The only exceptions to equation (3) are when a reference adult changes their education status (see Section 2.5), in which case a new random draw is taken from a log-normal distribution, the mean and variance of which are specific to the family's age, birth year, relationship, and education status.

The minimum wage h^{\min} allows a floor to be imposed, with reference to the hourly wage rate. This floor is specified so that it can differ relative to four age thresholds. Each age-specific minimum wage rate can be defined to growth through time at different rates.

Wage offers, λ^o , are included in the model to allow for the possibility of (involuntary) unemployment among employees, which we have found to be important in matching the model to rates of employment during peak working years. Receipt of a wage offer is modelled as uncertain between one period and the next, subject to age, education, and relationship specific probabilities $p^o(n_{i,a}, ed_{i,a}, a)$. If a wage offer is received, $\lambda_{i,a}^o = 1$, then family income responds to the labour supply decision of all adults in the unit. If a wage offer is not received, $\lambda_{i,a}^o = 0$, then any labour that one adult supplies returns no labour income to the family, implying non-employment where working incurs a leisure penalty.

The solution to the lifetime decision problem assumes that families expect that the probability of a low wage offer is age, relationship, and education specific, but is time invariant (as p^o is defined above). When a population is simulated through time, however, allowance is made for historical variation in unemployment rates to reflect observed fluctuations through the economic cycle.

Each discrete labour alternative $l_{i,a}$ is associated with its own factor, $\lambda^{emp}(l_{i,a})$. λ^{emp} is defined to be an increasing function of labour supply, and is scaled so that full-time employment of all adult members implies $\lambda^{emp} = 1$. The form assumed for λ^{emp} treats spouses as symmetric, and permits each adult's share of family labour income, $g_{i,a}^j$, to be evaluated from total family labour income: $g_{i,a}^j = \lambda^{emp}(l_{i,a}^j) / \lambda^{emp}(l_{i,a}) \cdot g_{i,a}$.

Wage penalties are imposed on families that have started to draw upon their private pension wealth. This is allowed for through a fixed factor adjustment applied to the family's latent wage, $\lambda_{i,a}^{ret} < 1$ if the family has accessed their pension wealth.

2.3. The wealth constraint

Equation (1) is maximised, subject to an age specific credit constraint imposed on liquid (non-pension) net wealth, $w_{i,a} \geq D_a$ for family i at age a . D_a is set equal to minus the discounted present value of the minimum potential future income stream, subject to the condition that all debt be repaid by age 70. Intertemporal variation of $w_{i,a}$ is, in most periods, described by the simple accounting identity:

$$w_{i,a} = w_{i,a-1} + \tau_{i,a-1} + ur_{i,a-1}^h - c_{i,a-1} - ndc_{i,a}^x \quad (4)$$

where τ denotes disposable income, ur^h is un-realised returns to owner-occupied housing, c is discretionary non-durable composite consumption, and ndc^x is non-discretionary expenditure. Non-discretionary costs (sometimes referred to as "committed expenditure") are disaggregated into child care, housing (rent and mortgage interest), and 'other' categories to facilitate simulation of welfare benefits that make explicit reference to any one of these expenditure categories. Simulated child care costs, ndc^c , are described as a function of the number and age of dependent children, and of the employment status of the least employed adult family member. Non-discretionary housing expenditure is comprised of rent and mortgage payments, $ndc^{hg} = rent + mort$, and is described below. 'Other' non-discretionary expenditure, ndc^o , is loosely designed to reflect the minimum

expenditure required to participate in society, consistent with standard definitions of poverty. Consumption on other basic necessities is defined in terms of equivalised (non-housing / non-child care / non-health) consumption, and varies by age and year.

The only potential departures from equation (4) occur when a family is identified as accessing pension wealth, or when a reference adult is identified as getting married or incurring a marital dissolution. Wealth effects at the time a family accesses its pension wealth are discussed in Section 2.4. In relation to marital transitions, spouses are identified from within the simulated sample. A marriage between two simulated singles consequently results in the liquid net wealth of each being combined in the common family unit. A divorce is assumed to see liquid net wealth split evenly between each divorcee, whereas widowhood sees all liquid net wealth bequeathed to the surviving spouse. Solutions to the utility maximisation problem are evaluated on the assumption that marriages are between identical clones.

Disposable income

As the model has been designed to undertake public policy analysis, particular care was taken concerning formulation of the module that simulates the effects of taxes and benefits. The model allows the measures of income accruing to each adult family member to be accounted for separately, so that it can reflect the taxation of individual incomes that is applied in the UK. The tax function assumed for the model is represented by:

$$\tau_{i,a} = \tau \left(\begin{array}{l} b, a, n_{i,a}, n_{i,a}^c, l_{i,a}^j, g_{i,a}^j, hh_{i,a}, mh_{i,a}, w_{i,a}^h, rent_{i,a}, \\ mort_{i,a}, rr_{i,a}^h, w_{i,a}^{nh,j}, r_{i,a}^{nh}, w_{i,a}^{nh,j}, pc_{i,a}^{c/nc,j}, py_{i,a}^j, ndc_{i,a}^c \end{array} \right) \quad (5)$$

which depends on the birth year of the reference adult b ; age of the reference adult, a ; number of adults (relationship status), $n_{i,a}$; number and age of all dependent children, represented by the vector $n_{i,a}^c$; labour supply of each adult j in the family, $l_{i,a}^j$; the labour income of each adult, $g_{i,a}^j$; indicator variables for home-owners, $hh_{i,a}$, and mortgage holders, $mh_{i,a}$; the net owner-occupied housing wealth held by the family, $w_{i,a}^h$; the rent paid by non-home-owners, $rent_{i,a}$; the mortgage interest paid by mortgage holders, $mort_{i,a}$; the realised returns to (gross) housing wealth, $rr_{i,a}^h$; the non-housing net liquid wealth held by each adult, $w_{i,a}^{nh,j}$; the investment return on liquid net wealth of each adult in the family, $r_{i,a}^{nh} w_{i,a}^{nh,j}$ (which may be negative); the pension contributions made by each adult, $pc_{i,a}^{c/nc,j}$; the (retirement) pension income received by each adult in the family, $py_{i,a}^j$; and non-discretionary child care costs, $ndc_{i,a}^c$. All of the inputs to the tax function are described in other subsections of this paper.

Disaggregating liquid net wealth

Liquid net wealth includes all assets other than private pensions. Importantly for the UK, this includes owner occupied housing. Although formal modelling of housing investment decisions is analytically feasible (O. Attanasio, Bottazzi, Low, Nesheim, & Wakefield, 2012), it is also computationally burdensome. Computational feasibility of the model is maintained by adopting an exogenous procedure for identifying selected housing-related features: home owners (hh), mortgage holders (mh), net housing wealth (w^h), mortgage debt (md^h),

gross housing wealth ($w^h + md^h$), realised returns on gross housing equity (rr^h), unrealised return on gross housing equity (wr^h), mortgage interest costs ($mort$), and rent ($rent$).

The exogenous procedure starts with a logit equation that describes the incidence of home-ownership as a function of age, marital status, and liquid net wealth. A similar logit equation is used to identify the incidence of mortgage holders among all home-owners. An age specific proportion of liquid net wealth is assumed to be held in housing for all home owners. Each mortgage holder's mortgage value is defined as a linear function of non-negative (log) liquid net wealth. The slope and intercept of this function are allowed to vary between singles and couples, and the multiple is restricted to values between 0 and 20. The annual interest change on mortgage debt is evaluated by multiplying the (gross) mortgage value by an (exogenously assumed) fixed rate of mortgage interest. Similarly, gross housing wealth is assumed to attract an exogenous rate of return. Fixed rates of return are applied for solving the lifetime decision problem, and year-specific rates are accommodated when projecting the population through time. The total return to gross housing wealth is then disaggregated into realised and unrealised components using an age specific ratio, based on the age of the family reference person. Rent is imputed for non-home-owners, based on the number of bedrooms required by the constituent family members: one bedroom is assumed for the reference adult and their spouse (if married); one bedroom is assumed for each dependent child aged 13 or over; and one bedroom for every two children aged under 13 years.

If non-housing liquid net wealth is non-negative, $w^{nh} \geq 0$, then the assumed rate of return is r^I . Otherwise, the return to w^{nh} is designed to vary from r_l^D at low measures of debt to r_u^D when debt exceeds the value of working full time for one period (g^{ft}):

$$r^{nh} = \begin{cases} r^I & \text{if } w^{nh} \geq 0 \\ r_l^D + (r_u^D - r_l^D) \min \left\{ \frac{-w^s}{g^{ft}}, 1 \right\}, r_l^D < r_u^D & \text{if } w^{nh} < 0 \end{cases} \quad (6)$$

Specifying $r_l^D < r_u^D$ reflects a so-called 'soft' credit constraint in which interest charges increase with loan size. The model parameters r^I , r_l^D , and r_u^D take fixed values when solving for utility maximising decisions, and are allowed to vary when simulating the intertemporal evolution of a population.

2.4. Private Pensions

All private pensions are modelled at the family level, and are Defined Contribution in the sense that every family is assigned an account into which their respective pension contributions are (notionally) deposited. Contributions to private pensions are defined as fixed rates of employment income (implying that they are limited to families that work), and are distinguished by whether they are made by the employer, π_{er} , or the employee, π_{ee} : $pc_{i,a} = (\pi_{ee} + \pi_{er}) g_{i,a}$. All employer pension contributions are assumed to be exempt from taxation, and labour income is reported net of these. Employee contributions up to a year-specific cap are also exempt from income tax, consistent with the EET nature of the UK pension system.⁶ Any employee contributes in excess of the cap are subject to income tax. Labour income is reported gross of all employee contributions. A cap is also imposed on the maximum size of the aggregate pension pot, which remains fixed through time.

Until the year in which a family accesses its pension wealth, intertemporal accrual of private pension wealth,

w^p , is described by equation (7):

$$w_{i,a}^p = \max \left\{ 0, \min \left[w^{p,\max}, r_{t-1}^p w_{i,a-1}^p + pc_{i,a}^p \right] \right\} \quad (7)$$

where $w^{p,\max}$ defines the maximum size of a pension pot. Equation (7) holds in all periods prior to pension receipt except following relationship transitions, in which case associated fluctuations in pension rights are modelled in a similar fashion as described for liquid net wealth.

The age at which pension dispersals are accessed, a^P , is determined endogenously from a defined range of ‘pensionable ages’. At the time that pension wealth is accessed, a fixed fraction of accrued pension wealth is received as a tax-free lump-sum cash payment, and the remainder converted into a level life annuity that is subject to income tax. Annuity rates are calculated to reflect birth cohort-specific survival probabilities in the model, subject to assumed rates of investment returns, real growth, and transaction costs levied at time of purchase.

When the timing of pension dispersals is exogenously imposed, then all families are assumed to access their pension wealth at their respective state pension ages (a exogenously defined policy parameter). When the timing of pension dispersals is endogenously determined, then this decision can be made subject to minimum thresholds on age and annuity income.

2.5. Education

Each adult is allocated an education state at entry into the model, $ed_{i,a}$, referring to the highest qualification held, distinguishing between those with and without graduate level qualifications. Individuals with tertiary education are distinguished from non-tertiary educated in relation to employment offers, the age specific evolution of latent wages (h in Section 2.2), and transition probabilities governing marriage and divorce.

Individuals who do not enter the simulated population with tertiary education may be identified as tertiary students, $stud_{i,a}$. Any individual who first appears as a tertiary student is assumed to leave tertiary education at an exogenously defined age (assuming that they survive), at which time they may transition to tertiary educated, depending on a stochastic process that represents whether they pass their final exams. At the time an individual leaves tertiary education, they receive a new random draw for their wage potential from a log-normal distribution, where the terms of the distribution differ for graduates and non-graduates. All processes that govern transitions between alternative education states when projecting a population through time are assumed to be fully consistent with the associated expectations adopted to solve the lifetime decision problem.

2.6. Allowing for demographics

Three demographic characteristics are considered for calibrating model parameters: mortality, relationship status, and dependent children.

Modelling mortality

The model focusses upon survival with respect to reference adults only; the mortality of the spouses of reference adults is aggregated with divorce to obtain the probabilities of a relationship dissolution (discussed below). Survival in the model is governed by age and year specific mortality rates, which are commonly reported components of official life-tables.

Modelling relationship status

A 'relationship' is defined as a cohabiting partnership (including formal marriages and civil partnerships). The relationship status of each reference adult in each prospective year is considered to be uncertain. The transition probabilities that govern relationship transitions depend upon a reference adult's existing relationship status, their education, age, and birth year. These probabilities are stored in a series of 'transition matrices', each cell of which refers to a discrete relationship/education/age/birth year combination.

Modelling children

The model takes explicit account of the number and age of dependent children of reference adults. The birth of dependent children is assumed to be uncertain in the model, and described by transition probabilities that vary by the age, birth year, relationship status, and previously born children of a reference adult. These transition probabilities are stored in a series of transition matrices, in common with the approach used to model relationship status (described above). Having been born into a family, children are assumed to remain dependants until an exogenously defined age of maturity. A child may, however, depart the modelled family unit prior to attaining maturity, if the reference adult experiences a relationship dissolution (to account for the influence of divorce).

The model is made computationally feasible by limiting child birth to three 'child birth' ages. Realistic family sizes are accommodated by allowing up to two children to be born at each child birth age. Restricting the number of ages at which a child can be born in the model raises a thorny problem regarding identification of the transition probabilities that are used to describe fertility risks. The model calculates the required probabilities internally, based upon the assumed birth ages and fertility rates reported at a highly disaggregated level. This approach has been adopted both because statistical agencies tend to publish data at the disaggregated (annual age band) level, and because it facilitates associated sensitivity analyses to be conducted around the number and precise birth ages assumed.

3. BASIC MECHANICS OF THE EMPIRICAL APPROACH

In common with the existing dynamic programming literature, a two stage procedure was used to identify parameters that match our structural model to survey data.⁷ The first stage identified a subset of parameters exogenously from the model structure, using methods that have changed little since the advent in the 1960s of

‘classical’ dynamic microsimulation models. Most of the parameters identified in the first stage are directly observable – e.g. marital transition rates, contribution rates to private pensions, the functional forms assumed for taxes and benefits – and were evaluated from publicly available data sources. Given the model parameters evaluated in the first stage, remaining model parameters were adjusted in a second stage so that selected ‘simulated moments’ implied by the structural model matched to ‘sample moments’ estimated from survey data. Conceptually, the second stage of the procedure involves adjusting unobserved model parameters to ensure that observable endogenous characteristics implied by the assumed theoretical framework best reflect a selected set of moments estimated from survey data. The remainder of this section describes the second stage of the parametrisation.

3.1. Evaluating simulated moments

The approach taken to evaluate population moments implied by the assumed theoretical framework is now well established in the related literature. This section consequently provides a brief overview of the techniques employed; for further detail see, for example, Adda and Cooper (2003) or Christensen and Kiefer (2009).

The model considered here does not generate projections for a representative agent, but for a population of heterogeneous agents designed to reflect an observed cross-section. The population moments implied by the model under a given set of model parameters were consequently evaluated by: (i) solving the lifetime decision problem for any feasible combination of family specific characteristics; (ii) using the solutions obtained in (i) to project endogenous decisions for each heterogeneous agent; and (iii) aggregating up the agent-specific decisions generated by the model to obtain the moments of interest. The first two of these steps are described below.

Solving the decision problem

No analytical solution exists to the utility maximisation problem considered here, and numerical solution routines were consequently employed. These solution routines are structured around a ‘grid’ that over-lays all feasible combinations of individual specific characteristics (the state space).⁸ As noted in Section 2.1, the model assumes that there is a maximum potential age to which any individual may survive, A . At this age, the decision problem is deterministic, and trivial to solve. The solution routine that we employed starts by solving for utility maximising decisions at all intersections of the grid that correspond to this final period of life, and stores both the maximising decisions and optimised measures of utility (the value function). These solutions at grid intersections for age A are used to approximate solutions at age A more generally, via the linear interpolation routine that is described in Keys (1981).

Given results for age A , the solution routine that we used then considers decisions at intersections corresponding to the penultimate age, $A - 1$. Here, expected lifetime utility is comprised of the utility enjoyed at age $A - 1$, and the impact that decisions taken at age $A - 1$ have on circumstances – and therefore utility – at age A . Given any decision set at age $A - 1$, d_{A-1} , the solution routine projects forward the set of individual specific characteristics at age A , z_A , that is implied by the processes assumed to govern intertemporal transitions (e.g. equation 4 for wealth, equation 3 for wage potential). If characteristics at age A are uncertain (e.g. as in equation 3), then each potential characteristic vector z_A^p is projected forward with an assigned probability

pr_A^p . Uncertainty in the model is either between a discrete set of alternatives (relationship status, wage offers, and death), or over a continuous normal distribution (wage potential). Expectations over normal distributions are approximated at 5 discrete points, using weights and abscissae implied by the Gauss-Hermite quadrature (implemented following Press, Flannery, Teukolsky, and Vetterling (1986)). These terms, combined with a von Neuman Morgenstern preference relation, allow the expected lifetime utility associated with any decision set d_{A-1} to be evaluated. A numerical routine (described below) is used to search over the set of feasible decisions to maximise expected lifetime utility at each intersection of the grid corresponding to age $A - 1$. These solutions, and the associated measures of optimised utility are stored, and the solution routine then considers the next preceding age. Repeated application of this procedure obtained a numerical approximation of the solution to the lifetime decision problem at all intersections of the grid spanning the feasible state space.

The numerical search routine that was employed for this study is adapted to the decisions that are considered for analysis. As described in Section 2, families are assumed to decide over one continuous domain relating to the consumption/savings margin, and a series of discrete alternatives relating to labour supply, pension participation, and the take-up of pension benefits. The search routine used considers each potential discrete alternative in turn, and searches for a local optimum in relation to consumption. Of all feasible alternative solutions, the one associated with the maximum numerical approximation of expected lifetime utility is taken as the solution to the lifetime decision problem.

As the value function of the utility maximisation problem considered here is not smooth, three alternative search routines over the eligible consumption domain were employed to test the robustness of model projections. The first uses Brent's method as described in Press et al. (1986); the second uses the simplex method of Lagarias, Reeds, Wright, and Wright (1998); and the third employs the multi-level coordinate search method described in Huyer and Neumaier (1999) (as implemented by the NAG library). All three approaches generated very similar results, and the fastest search routine (Brent's method) was consequently used for the calibrations by default.

Calculating simulated moments

The simulated moments used to guide adjustment of the model's parameters were calculated using data generated by the model for a population of reference adults drawn from a nationally representative cross-sectional survey. The circumstances of each reference adult described by the survey were used to locate them within the grid structure that is referred to above. Given their respective grid co-ordinates, the linear interpolation methods that are also mentioned above were used to approximate each reference adult's utility maximising decision set, as implied by the numerical solutions identified at grid intersections. Given a family's characteristics (state variables) and behaviour, its characteristics were projected through time following the processes that are considered to govern their intertemporal variation. Where these processes depend upon stochastic terms, random draws were taken from their defined distributions in a process that is common in the microsimulation literature (sometimes referred to as Monte Carlo simulation).

3.2. Adjusting model parameters

The second stage of the model parametrisation involved identifying the parameters of the assumed preference relation, simulated rental charges, and a selected set of parameters governing intertemporal evolution of latent wages. Preference parameters are unobservable, and are consequently prime candidates for the second stage of the model parametrisation. Although rental charges are observable, distributional considerations complicate identification of appropriate charges to assume for the model. Similarly, although wages are observable, a subset of wage parameters were included in the second stage of the parametrisation to account of associated selection effects.

Adjustment of the parameters to match the simulated moments implied by a dynamic programming model to associated sample moments is commonly conducted either by manual calibration or optimisation of a loss function using an econometric criterion.⁹ The results reported here were obtained via a series of manual adjustments of model parameters, guided by graphical representations and sums of squared errors for a set of age specific population moments, following the approach described by Sefton et al. (2008).

Calibrated parameters

The assumed preference relation (see Section 2.1) includes five parameters: relative risk aversion, γ ; an exponential discount factor, δ ; a parameter for the warm-glow model of bequests, ζ ; the intra-temporal elasticity, ε ; and the utility price of leisure, α . There are 12 rental rates included with the model, which distinguish families by income, demographic size, and age. Finally, the specification adopted for wages (see Section 2.2) includes a very large number of parameters. The persistence of latent wages, ψ , and the factor effects of alternative labour supply decisions, λ^{emp} , were identified in the first stage of the model parametrisation. This left the parameters governing wage growth $m(\cdot)$, earnings volatility $\sigma_w^2(\cdot)$, and the factor effects of pension take-up λ^{ret} to be identified in the second stage of the parametrisation.

Calibration procedure

Following extensive experimentation, we settled upon a step-wise procedure based on concentric cycling over three sets of model parameters; A , B , and C . Parameters in a higher set were re-adjusted each time the parameters in a lower set were altered, so that those in set C were subject to the most frequent adjustment, and those in set A the least frequent. Set A comprises simulated rental rates, set B the parameters governing wage growth and earnings volatility (of employees), and set C all other endogenously calibrated parameters. We began by setting rental charges to average market rates, all wage growth parameters $m(\cdot) = 1$, and made initial guesses for earnings volatility, $\sigma_w^2(\cdot)$. Given these assumptions for parameter sets A and B , and the model parameters identified exogenously from the model structure in the first stage of the analysis, the calibration procedure began by adjusting parameters in set C . We found that it was necessary to pass through a small number of cycles to obtain convergence, a property attributable to the invariance of the cross-sectional population characteristics upon which the calibration was based, as explained below.

Identification of parameter set C

All five preference parameters of the model and the factor effects of pension take-up λ^{ret} , were identified by matching the model to moments evaluated on survey data reported for a single (reference) population cross-section. This is notable, given that preference parameters are often a central focus of interest in the related literature. It is also extremely useful because it simplifies specification of the policy context underlying the behaviour considered for identification, and omits the feed-back effects that can otherwise complicate parameter adjustments.

The feed-back effects that are mentioned here complicate any empirical analysis that refers to dynamic behaviour described for an appreciable period of time. Suppose, for example, that we were interested in matching a structural model of savings and retirement to data observed during the life-course of a single birth cohort. If a given set of model parameters implied savings early in the life course that over-stated observed data, then this might suggest that preferences should reflect greater impatience. Adjusting preferences in this way might then imply lower wealth later in life, and thereby influence model implications for the timing of retirement. Such feed-back effects can be ignored in an empirical analysis of household sector savings that focuses exclusively on behaviour described for a single point in time, as population characteristics such as wealth holdings are then exogenously defined. This considerably simplifies the identification problem.

The approach taken to calibrate parameters identified on cross-section survey data started with the assumption of a high value for relative risk aversion ($\gamma = 5$), a high value for the exponential discount factor ($\delta = 1$), a low value for the warm-glow model of bequests ($\zeta = 0$), and a moderate value for intra-temporal elasticity ($\varepsilon = 0.5$). Parametrisation then proceeded in four concentric ‘loops’.

(1) The inner-most loop, which was repeated most frequently, focussed on adjusting α and λ^{ret} . Increasing the utility price of leisure α tends to decrease labour supply throughout the working lifetime. Exaggerating the wage discount for families that have previously accessed their private pensions tends to decrease labour supply late in the working lifetime. These two model parameters provide a high degree of control over the employment profile throughout the life-course, and were jointly adjusted to match the model to age and relationship-specific means for employment participation.

(2) The second loop jointly adjusted δ and ζ to reflect age and relationship specific geometric means for consumption. Increasing the discount factor δ makes families more patient, and consequently tends to decrease consumption throughout the working lifetime. Increasing ζ exaggerates the bequest motive, which tends to lower consumption late in the life course when the probability of imminent mortality becomes appreciable. Taken together δ and ζ provide a high degree of control over the age profile of consumption implied by the structural model.

(3) The third loop of the calibration strategy adjusted γ , by focussing on the associated influence on savings incentives. Raising γ *ceteris paribus* exaggerates precautionary savings motives, implying lower consumption and *lower* pension scheme participation (due to the illiquidity of pension wealth). In contrast, raising δ tends to imply lower consumption and *higher* pension scheme participation as families are made more patient. Hence, if the rates of pension scheme participation implied by the model following the second loop of the calibration were too low (high), we reduced (increased) γ and returned to the inner-most loop. Otherwise we proceeded

to the outer-most loop.

(4) ε was adjusted to match the model to distributional variation described by data for the ratio between equivalised consumption and leisure. If the utility maximisation problem was separable, and labour supply was a decision on a continuous domain, then the preference relation defined by equation (1) would imply the following relationship between the decision variables c and l in the region of the optimum:

$$\alpha \frac{\hat{c}_{i,a}}{l_{i,a}} = \hat{h}_{i,a}^\varepsilon \quad (8)$$

where \hat{c} denotes equivalised consumption and \hat{h} is the equivalised post-tax and benefit wage rate. This relationship will approximately hold late in the simulated working lifetime, when benefit units exhibit substantial variation over labour supply decisions and continue to possess multiple periods over which they can choose between (discrete) labour supply alternatives. The relationship defined by equation (8) can be used to compare the decisions taken by any two benefit units, 0 and 1, as described by the ratio:

$$\left(\frac{\hat{c}}{l}\right)_1 / \left(\frac{\hat{c}}{l}\right)_0 = \left(\frac{\hat{h}_1}{\hat{h}_0}\right)^\varepsilon \quad (9)$$

Equation (9) indicates that increasing ε will tend to shift period specific expenditure in favour of (equivalised) consumption, relative to leisure, for families with relatively high (equivalised) wage rates.

Model implications were consequently evaluated for the ratio between equivalised consumption and leisure for every family with a reference adult aged 55 to 60 in the reference population cross-section. Two separate averages were calculated over these ratios, distinguishing families with and without reference adults educated to graduate level. If the value of the ratio of the graduate average divided by the non-graduate average was too low (high), then ε was increased (decreased). The calibration then proceeded back to the inner-most loop, and the entire process repeated until a convergence was obtained.

Section 6.1 quantifies the effects of parameter variation on simulated moments that are discussed above.

Identification of parameter set B

The drift parameters, $m(\cdot)$, and the dispersion parameters, $\sigma_\omega^2(\cdot)$, were calibrated against historical data by projecting the reference population cross-section backward through time. The drift parameters were adjusted to reflect geometric means of employment income, distinguished by age, year, relationship status, and education status. The model includes a separate drift parameter for each age, year, education, and relationship combination, so that a close match could be obtained to the associated sample moments. Given the large number of model parameters involved, this stage of the parametrisation was undertaken using an automated procedure. First, age, year, education, and relationship specific means of log employment income implied by the model under any given parameter combination were calculated from simulated panel data projected back in time for the reference population cross-section. These simulated moments were subtracted from associated sample moments estimated from survey data. The differences so obtained were then multiplied by a ‘dampening factor’, set equal to 0.9 in the first instance, and subsequently reduced to 0.4.¹⁰ The exponent of the result was taken, and

multiplied by the prevailing drift parameter to obtain an updated value for the parameter. This procedure was repeated until the average absolute variation of parameters over ages for any year, education, and relationship combination fell below 5 percentage points.

Similarly, the variance parameters were adjusted to reflect age, year, and relationship specific variances of log employment income calculated from survey data. Unlike the drift parameters, however, only four parameters – distinguish singles from couples, and graduates from non-graduates – were adjusted to reflect the dispersion of employment income. These model parameters were adjusted manually.

Identification of parameter set A

Two sets of rents (*rent*) are supplied to the model: Local Housing Allowance (LHA) rates are assumed for families with equivalised incomes below 60% of median gross full-time earnings in 2011, and “market rents” for those with equivalised incomes above 120% of median earnings. Between these two income thresholds, rental charges are assumed to vary linearly between the LHA and market rents. Both sets of rents are described in terms of numbers of bedrooms, where one bedroom is allowed for each single adult / cohabiting couple, another for each child aged 13 or over, and another for every 2 children aged under 13 years, subject to a maximum of four bedrooms. The only exception is in relation to single adults aged 30 or under without children, who are assumed to share their accommodation, and who consequently incur lower rental charges.

The simulated benefits system includes a scheme to subsidise rental charges. Furthermore, rents in the model are disproportionately incurred by individuals toward the bottom of the income/wealth distribution (reflecting observed data). Increasing rents consequently tends to increase disposable incomes toward the lower end of the distribution on a before housing costs basis, and to reduce disposable incomes on an after housing costs basis, with associated implications for simulated poverty rates. Rental charges were consequently adjusted to match poverty rates generated by the model to survey data.

4. SURVEY DATA

This section defines the cross-sectional data selected for analysis, before describing the sample moments used to conduct the second stage of the model calibration.

4.1. The reference population cross-section

Data for the reference population cross-section were drawn from the Wealth and Assets Survey (WAS), which is currently the micro-data set that provides the most complete description of household demographics, income, and wealth that is available for the UK. The sample frame for the WAS is the small users Postcode Address File covering residential addresses in Great Britain, excluding regions north of the Caledonian Canal and the Isles of Scilly. At the time of writing, the WAS is comprised of three waves. Wave 1 data were collected between July 2006 and June 2008. The achieved sample from Wave 1 was re-interviewed two years later in Wave 2 (conducted between July 2008 and June 2010). In Wave 3, responding and non-contact households in Waves 1 and 2 were

re-surveyed two years later, between July 2010 and June 2012. The Wave 3 sample was also augmented by a new random sample. All surveys were administered by the Office of National Statistics via computer assisted personal interviewing

The sample of households for Wave 1 was spread evenly across the sample period, and was selected to be geographically representative of the population of Great Britain, subject to over-sampling of high wealth households. Of the 55,835 households invited to participate in wave 1 of the survey, responses are reported for 30,511 households, implying an achieved response rate of 55 per cent. Wave 2 reports data for 20,009 households (response rate 68 per cent), and Wave 3 reports data for 21,251 households (response rate of 61 per cent).

The model parameters have been updated using a reference population comprised of 10,771 households reported by wave 3 of the WAS between January and December 2011. This sample period aligns with the period of data reported by the 2011 Living Costs and Food Survey, which is the second most important data source used to parameterise the model. The representativeness of the WAS sample for the British population was considered, by comparing the distribution of total gross household income reported by the WAS against associated data reported by the Family Resources Survey. Results of this analysis, which are available from the author upon request, suggest that a close degree of agreement exists between these two data sources.

Formatting the data for use in the model was performed by a single STATA “do” file, which is provided with the model available on the website (www.simdynamics.org). This do file merges data from the household level WAS file with data reported in the personal level file for wave 3. Each individual is then allocated to a family, with each family comprised of a single adult or cohabiting couple and their dependent children. All individuals under age 16, or under age 19 and full-time students, are identified as dependent children. Age, relationship status (single/couple), and the number and age of all dependent children in the family are evaluated.

Indicator variables are evaluated for each adult to distinguish those with graduate education, and those currently enrolled in a tertiary level education.

All earnings from employment are evaluated for each adult, and identifiers calculated to distinguish self-employed, full-time employees, part-time employees, and unemployed. Furthermore, identifiers are calculated to distinguish those with non-contributory pension schemes (predominantly public sector employees), those eligible for a (contributory) occupational pension scheme, members of occupational pensions, and whether defined benefit and defined contribution pensions are held. For those who are identified as members of an occupational pension, their private contribution rate to the pension is also evaluated. Furthermore, the value of all state pensions currently in payment is recorded.

Regression models for log earnings are evaluated that adjust for sample selection via a Heckman correction, separately for men and women, and for individuals aged under 50 and those aged 50 and over. These regression models are calculated on the full set of data reported by wave 3 of the WAS and include, in addition to the range of data saved for loading into the LINDA model, identifiers for health status, stated preferences for saving, and housing tenure. The regression results are used to adjust the earnings of part-time employed adults to their full-time equivalents, and to impute predicted wages for those identified as not employed.

To impute a full-time wage for each adult not identified as employed in the WAS data, predicted values at the coefficient estimates for the wage equation, xb , and the selection equation, zg , are evaluated by STATA. These

data are transferred to LINDA, along with regression estimates for sigma (the standard deviation of the residual of the wage regression) and lambda (the standard deviation of the wage regression times the correlation between the residuals of the target and selection equations). These data permit the model to generate a full-time wage for each without a wage reported in the WAS, after generating a random draw from a standard normal distribution.

The value of savings held in Individual Savings Accounts, own businesses, property other than the main home, financial and non-financial assets and pensions are identified at the family level by aggregating up the value of each asset class held by all family members. The principal exception to this approach is the value of equity held in the main home, which is allocated entirely to the family of the household reference person.

The model is designed to track the evolving household circumstances of a sample of “reference people”. Each adult aged 18 or over in the WAS is represented as a reference person of a benefit unit in the reference cross-section, so that the families of couples are represented twice in the base data – once for each spouse. An indicator variable is included in the model, which identifies which reference adults are married to one-another. This is consistent with the approach taken to simulate the evolution of relationship status in the model, where marriages are considered to be between individuals represented in the simulated population.

Two adjustments were applied to the data reported by the WAS to obtain the base data-set from which model projections are made. First, the cross-sectional weights reported by the WAS are designed to aggregate up to the 2010 principal projections for the population of Great Britain. These estimates have subsequently been revised upward in light of data reported by the 2011 Census. The cross-sectional weights reported for wave 3 data were consequently adjusted to align the aggregate weighted population reported by the WAS to ONS mid-year estimates for Great Britain in 2011.

Secondly, a pseudo population for Northern Ireland was imputed, by randomly selecting observations reported by the WAS for Great Britain, structured to reflect the age, relationship status, and income distributions of Northern Ireland. This was achieved, by first analysing data reported by the Family Resources Survey (FRS) for the UK in 2011/12.¹¹

A measure of total gross income, comprised of earnings, self-employment, social benefits, pensions, investments, and other income, was extracted from the FRS for each reported family, along with the age group of the reference adult distinguished by 10 year intervals. Age and relationship specific quintile thresholds for total gross income were evaluated for the sample of families reported to be living in Great Britain, weighting each family by their respective sample weight (GROSS3) and the number of members of the family (ADULTB+DEPCHLDB). This produced 80 mutually exclusive and collectively exhaustive population sub-groupings, distinguishing 5 income-ranges for each of 8 age-bands, separately for singles and couples. The proportion of the Northern Irish population reported by the FRS as corresponding to each of the 80 population subgroups was then evaluated (available from the author upon request).

Comparable age bands and income measures to those considered for the FRS were evaluated for each family reported by the WAS. Each family reported by the WAS was then sorted into a unique age-band and relationship-specific quintile group, based on the considered measure of income. Random selections (based on the WAS household identifier) from each of the age-band/income quintile groups were then taken to match the distribution of the Northern Irish population, as calculated using FRS data.

The base data set derived as described here is comprised of 20,247 adults and 5,177 children in 13,592 families. This includes the pseudo population for Northern Ireland comprised of 657 adults and 212 children in 428 families. Associated weighted populations are 61,371,315 for Great Britain (the ONS mid-year population estimate), and 1,812,671 for Northern Ireland (compared with the ONS mid-year population estimate of 1,810,863).

4.2. Sample moments

The calibration strategy, described in Section 3.2, was implemented with reference to the following sample moments:

1. The proportion of adult family members employed, by age and relationship status; estimated on data for the population cross-section observed in 2011.
2. The geometric mean of family employment income, by age, education, and relationship status; estimated on data for population cross-sections observed from 1978 to 2012.
3. The variance of family log employment income, by age, education, and relationship status; estimated on data for the population cross-sections observed from 1978 to 2012.
4. The geometric mean of family consumption, by age and relationship status; estimated on data for the population cross-section observed in 2011.
5. The proportion of families reporting to contribute to private pensions, by age and relationship status; estimated on data for the population cross-section observed in 2011.
6. The proportion of all individuals, and individuals above state pension age, with less than 60% of median equivalised disposable income, measured on both a before and after housing costs basis in 2011.

These sample moments were estimated on survey data from the Living Costs and Food Survey (LCFS) and the Family Resources Survey (FRS). In common with the WAS, the LCFS and FRS are conducted by the Office for National Statistics, use similar sampling frames and methods, and typically achieve similar response rates to the WAS. The most significant departures between the sampling approaches implemented by the three surveys are the over-sampling of high wealth households by the WAS, and the time periods covered by the respective surveys: while we focus on the WAS data reported for the calendar year ending December 2011 in common with the sample from for the LCFS, the FRS reports data for the UK financial year (starting in April). We ignore this mismatch between the time frames covered by the alternative data sources.

The LCFS is the principal source of micro-data for domestic expenditure in the UK. In addition to expenditure, it provides detailed information regarding family demographics, employment, and earnings, and covers a relatively long time-series, reporting at annual intervals from 1978. Most of the sample moments used for calibrating the model parameters were consequently estimated on LCFS data. The exception concerns participation rates in private pensions, which are more adequately described by the FRS than the LCFS.

The starting point for calibrating LHA rents were the rental averages reported over all Local Authorities for June 2011 by the Valuation Office Agency. Market rates were set to twice the assumed LHA rates. The rental charges were adjusted to match the model to poverty rates reported by the Households Below Average Income publication issued by the Department for Work and Pensions.

5. EXOGENOUSLY IDENTIFIED MODEL PARAMETERS

The methods used to evaluate the exogenously identified parameters of the model are commonly employed in the associated literature. Furthermore, the specific parameter estimates generated for the analysis described here are unlikely to be of particular value to readers in their own right. This section consequently provides an abbreviated description of the methods and data sources used to evaluate exogenous model parameters; full details are available from the author upon request.

Parameters were exogenously evaluated for seven key features of the model: transfer payments, returns to non-pension wealth, a subset of wage parameters, description of private pensions, probabilities governing relationship transitions, fertility rates, and mortality rates. Each of these factors is discussed in turn.

Transfer policy implemented in the model is designed to reflect UK tax and benefits policy as it applied in April 2011. As behaviour in 2011 is generated on the assumption of forward looking expectations, a description of policy beyond 2011 is also required. Expectations assume that all tax thresholds are indexed to real wage growth, assumed to be 1.5% p.a. In contrast, benefits values and thresholds are frozen (in nominal terms) from 2016 to 2020, and indexed to prices thereafter. These assumptions broadly reflect the policy environment at the time of writing.

The evolution of **non-pension wealth** involves distinguishing housing from non-housing wealth, and evaluating returns to housing and non-housing, non-pension assets. The (reduced form) regression equations used to distinguish housing wealth were estimated from data reported by the WAS for 2011. Returns to gross housing wealth ($rr^h + ur^h$) were calculated from the ONS mix adjusted house price index reported between 1970 and 2010, and discounted to real terms by the National Accounts final consumption deflator. The calibration assumes that the return to housing wealth in forward projections is equal to the mean return of the observed time-series, equal to 3.65% p.a.. Mortgage interest was set to the maximum rate eligible for subsidy through the UK Income Support system, equal to 3.85% p.a. nominal. The return to positive balances of non-pension/non-housing wealth (r^I) was set equal to the average real return on long-term treasury bills reported between 1970 and 2010, equal to 1.52% p.a. (real). The lower bound interest charge on unsecured debt (r_l^D) was set equal to annual averages of the monthly interest rates on sterling personal loans up to £10,000 to households reported by the Bank of England (code IUMHPTL) between 1995 and 2010; 8.36% p.a.. Similarly the upper bound interest charge on unsecured debt (r_u^D) was set equal to annual averages of the monthly interest rates for sterling credit card lending to households reported by the Bank of England (code IUMCCTL) between 1995 and 2010; 15.37% p.a..

Exogenously defined **wage parameters** assume that full-time employment of all adult family members reduce family leisure time by 40%, and part-time employment is equivalent to 50% of a full-time job.

Private pensions in the model depend upon five parameters: the rate of employer pension contributions π_{er} , the rate of employee contributions to pensions out of employment income π_{ee} , the rate of return to pension wealth r^P , the return assumed for calculating the price of pension annuities, and the fixed capital charge associated with purchasing a pension annuity. Although there is a wide diversity of private pension schemes in the UK, data from the Annual Survey of Hours and Earnings reported between 2005 and 2009 indicate that the distribution of employer contribution rates to such schemes is dominated by a single mode between 12.5 and 15 per cent of employee wages; an employer contribution rate of 14% is consequently assumed for the model. The rate of employee contributions to private pensions is set equal to 8%, which is the ‘normal’ contribution rate stated in the guidance to interviewers for the FRS. The real rate of return assumed for pension wealth during the accrual phase, r^P , was set equal to 3.5% p.a., based on a typical 60:40 split of pension wealth between equities and bonds, and estimates for rates of return for equities and Gilts observed between 1899 and 2009 reported by Barclays Capital¹², of 5% p.a. and 1.2% p.a. respectively. The capital return assumed for calculating the price of pension annuities was set equal to 1.5%, reflecting the average rate of return to long-term government debt observed between 1970 and 2010, and the associated capital charge was set to 4.7% based on “typical” pricing margins reported in the pension buy-outs market (see Lane, Clark, and Peacock (2008), p. 22).

The model requires age and year specific probabilities governing **relationship transitions**. These probabilities were evaluated using a simple statistical projection to ensure that simulated rates of cohabitation reflect data reported between 1978 and 2012 by repeated issues of the LCFS. The statistical projection was initialised using age specific rates of marriage and divorce reported by the ONS between 1971 and 2033 (including official forward projections for the population), and a standard implementation of Newton’s method for adjusting rates of relationship formation and dissolution.

The model requires **fertility rates** by age, year, relationship status, and number of previous births to simulate dependent children. These rates are not readily available for the UK, and so were evaluated using a similar statistical approach to that applied to evaluate rates of relationship formation and dissolution.

Age and year specific **mortality rates** assumed for the model were based on components of life tables reported by the ONS. The ONS mortality rates are based on observed data between 1951 and 2012, and the 2012 principal population projections for the UK from 2013 to 2062. ONS mortality rates are reported to age 100, and were extended to 1.0 at the assumed maximum age of life of 130 using a smooth sigmoidal function.

6. ENDOGENOUSLY CALIBRATED MODEL PARAMETERS

This section reports calibrated model parameters that were adjusted endogenously to the structural model, and which were identified using data observed for a reference population cross-section. As discussed in Section 3.2, this includes all of the parameters of the assumed preference relation, the factor effects of pension take-up, and simulated rental charges. The wage parameters that were also identified endogenously to the model structure are too numerous to report here. These parameters can be obtained from the author upon request and are also included with the model in the form of Excel spreadsheets, which can be downloaded from www.simdynamics.org.

The preferred parameter set is reported in Table 1.

Table 1: Calibrated model parameters adjusted to match behaviour reported for the British population cross-section in 2011

	singles	couples
relative risk aversion, γ	1.55	1.55
intra-temporal elasticity, ε	0.6	0.6
utility price of leisure, α	2.2	1.034
discount factor, δ	0.97	0.93
bequest motive, ζ	30000	10000
factor effects of pension take-up		
at age 55	1	1
from age 65	0.6	0.6
rents	lower	upper
shared	63.89	63.89
1 bedroom (from spa)	55.00	100.00
1 bedroom (pre spa)	101.48	106.82
2 bedrooms	69.25	132.54
3 bedrooms	81.82	156.59
4 bedrooms	107.12	205.01

Note: 'spa' refers to rents from 'state pension age'.

The calibrated value for the parameter of relative risk aversion $\gamma = 1.55$ is within the broad range identified by the associated literature. Simulations undertaken by Auerbach and Kotlikoff (1987), for example, are based upon a coefficient of risk aversion of 4, while Cooley and Prescott (1995) consider a value of 1. Grossman and Shiller (1981) and Blundell, Browning, and Meghir (1994) report estimates just over 1.0, while P. Hansen and Singleton (1983), Mankiw, Rotemberg, and Summers (1985), and Ziliak and Kniesner (2005) report estimates of approximately 1. Values of the coefficient of risk aversion required to explain the equity premium puzzle (Mehra and Prescott (1985)) are high by comparison, supported by econometric estimates reported by Mankiw (1985) and Hall (1988). Nevertheless, evidence from attitudinal surveys suggest that the value is unlikely to be greater than 5 (Barsky, Kimball, Juster, & Shapiro, 1997).

The inverse of the calibrated value for the intra-temporal elasticity of substitution is slightly greater than the parameter of relative risk aversion, implying that consumption and leisure are (weak) direct complements.¹³ The calibrated value for the utility price of leisure for single adults, at 2.2, is appreciably higher than that for couples, equal to 1.03, reflecting the relatively low age-specific rates of employment described by survey data for single adults. Furthermore, calibrated parameters for both the discount factor and bequest motive indicate stronger preferences for savings among singles than couples, all else being equal. The calibrated wage discount factor for pension take-up, which is the same for both singles and couples, increases linearly from 0% at age 55, to 40% at age 65.

Numerical simulations indicate that the calibrated model parameters imply an inter-temporal elasticity of consumption of 0.610 measured at the population means.¹⁴ This statistic is within the wide band of estimates that are reported in the associated literature. The meta-analysis by Havranek, Horvath, Irsova, and Rusnak (2013), for example, includes 34 studies that report 242 estimates for the intertemporal elasticity of substitution calculated on UK data, with a mean of 0.487 and a standard deviation of 1.09. The influential study by Hall (1988)

reports that the inter-temporal elasticity may not be very different from zero; see also Dynan (1993), Grossman and Shiller (1981), and Mankiw (1985). In contrast, O. P. Attanasio and Weber (1993) finds that focussing upon cohort data for individuals who are less likely to be liquidity constrained than the wider population obtains an estimate for the inter-temporal elasticity of consumption of 0.8 on UK data, and O. P. Attanasio and Weber (1995) report estimates between 0.6 and 0.7 for the US. Other empirical studies that support higher rates for the inter-temporal elasticity include Blundell et al. (1994) (0.75), Engelhardt and Kumar (2007) (0.75), P. Hansen and Singleton (1983) and Mankiw et al. (1985) (just over 1).

6.1. Identification

Table 2 reports a set of summary statistics that indicate the influence of selected model parameters on the moments considered for calibration. All statistics reported in the Table were obtained by perturbing the preferred parameter calibration reported in Table 1, calculating the simulated moments associated with the perturbation, and then subtracting the calibrated simulated moments.

The statistics for the ‘high alpha’ simulation report the effects on simulated moments of inflating the assumed values for alpha by 20%, relative to their calibrated values. These statistics indicate that, as preferences for leisure are strengthened, the proportion of the population choosing non-employment in the model during the reference cross-section rises, with relatively weak effects reported at the beginning of the simulated working lifetime (ages 20-29). The reduction in work that the rise in leisure implies, also reduces disposable income during the working lifetime, resulting in lower consumption. Notably, the reduction in consumption extends into early life, as individuals anticipate the influence of stronger leisure motives into the future. The reduced rates of employment result in lower rates of pension participation from age 30, as pension participation is limited to those in work in the model. In contrast, rates of pension participation early in the simulated working life are higher when alpha is exaggerated, which is again attributable to the forward looking nature of behaviour in the model.

When ‘no retirement effects’ are included that suppress wages following pension take-up, then Table 2 indicates that simulated rates of employment increase substantively late in the working lifetime. As discussed in Section 3.2, the calibration strategy adopted for this study exploits the disproportionate influence of retirement (wage) effects on employment late in life, by jointly adjusting α and retirement effects to match simulated moments of employment to survey data reported for the reference cross-section. The higher employment generated when retirement (wage) effects are suppressed produces higher labour income, increasing disposable income, and permitting higher consumption late in life. At the same time, rates of pension scheme participation fall, as pensions can only be contributed to prior to pension draw-down, and the retirement effect on wages acts as a dis-incentive to pension draw-down.

The ‘low delta’ simulation reduces the exponential discount rates of both singles and couples by 3 percentage points. This increase in impatience results in higher simulated leisure (more adults not employed), higher consumption, and lower pension scheme participation throughout the simulated working lifetime. The effects on simulated leisure are most pronounced during peak working years (30-54), while the effects on consumption are fairly level through the life-course. An important feature of the calibration strategy adopted here is that, whereas a low value of δ can have a similar effect on leisure and pension scheme participation as a high value

for α , the two parameters have opposing implications concerning consumption as indicated in Table 3.2. In practical terms, the inter-dependence of the effects on fitted moments that this discussion reveals between alternative assumptions concerning α and δ indicates the need to iterate between the respective calibration ‘loops’ to identify a preferred parameter combination.

The direction of the effects of omitting a warm-glow bequest motive from analysis are the same as those of increasing impatience, raising leisure and consumption in the reference cross-section, and reducing participation in private pensions. This observation is consistent with the fact that both parameter alternatives weaken savings incentives, relative to the calibrated parameter combination. However, Table 2 emphasises the extent to which incentives associated with the bequest motive are skewed toward later life. The influence on rates of non-employment under the ‘low delta’ simulation are highest during peak working years (30-54), and decline substantively into later life. In contrast, omitting a bequest motive has the strongest effects in the 55-74 age band. Furthermore, whereas the ‘low delta’ simulation resulted in similar increases in consumption among 30-54 year-olds as among 55-74 year-olds, omitting a bequest motive had an impact on consumption in the age band 55-74 that is over 3.5 times as great as among 30-54 year-olds. As discussed in Section 3.2, the disproportionate influence that bequest motives have on simulated consumption late in life is exploited by the calibration strategy considered here, by jointly adjusting δ and the bequest motive to match the model to age specific means of consumption observed for the reference cross-section.

The ‘low epsilon’ simulation reduces the intra-temporal elasticity from its calibrated value of 0.6 to 0.3. Table 2 indicates that this adjustment tends to reverse and dampen the effects reported for the ‘low delta’ simulation, increasing employment, reducing consumption, and (weakly) increasing pension scheme participation. The stand-out feature of the statistics reported for the ‘low epsilon’ simulation is the bearing on the equalised consumption to leisure ratio of graduates to non-graduates. As discussed in Section 3.2, reducing ε theoretically reduces the consumption to leisure ratio of high income families, relative to low income families. Table 2 indicates the extent of this effect, and reveals that the moment selected for calibration is well targeted, in the sense that it is broadly insensitive to variation of other model parameters.

Sensitivity of the model fit to the assumed value for γ is reported in the ‘high gamma’ series. The high gamma series assumes a value of γ equal to 2.0, up from 1.55 in our preferred specification. As discussed in Section 3.2, the three parameters γ , δ , and ζ all have an important bearing on simulated moments for both consumption and pension participation. Furthermore, the above discussion reveals that these parameters also influence preferences concerning labour supply. As the parameters adjusted in each of the calibration loops discussed above were identified taking the value of γ as given, they are re-specified in the high gamma series to clarify the effects underlying the assumed identification strategy. This involved reducing δ (implying less patience) to off-set the heightened precautionary savings motive associated with greater risk aversion. The ζ parameters were increased to force down consumption late in life (where modelled uncertainty is less pronounced), and α was reduced to off-set associated employment effects. This combination of adjustments ensures that simulated age profiles for geometric mean consumption and rates of non-employment are approximately the same as under the preferred parameter combination, as indicated by the small shifts associated with these statistics in Table 2.

The measures of fit for pension participation that are reported for the high gamma series indicate that high relative risk aversion discourages pension participation. The revised parameter specification based on increasing

γ from 1.55 to 2.0 reduces simulated participation in private pensions by a margin of approximately 15 percentage points for singles aged 18 to 54, and by approximately 20 percentage points for couples. Simulated participation rates in private pensions fall by less than 10 percentage points late in the working lifetime under the high gamma scenario, which is attributable to the coincident reduction in the time at which pension wealth can be accessed and the muted uncertainty that families with older reference adults face.

The 'high rents' simulation increases simulated rental charges by 50%, relative to their calibrated values. The statistics reported for this simulation in Table 2 indicate that raising rental charges has a negligible impact on rates of employment and private pension scheme participation. Consumption increases slightly, as (exogenous) rental expenditure rises. The rows second from the bottom of the table indicates that increasing rental charges lowers poverty rates on a before housing costs (BHC) basis among those in excess of state pension age (pensioners), and the population more generally. In contrast, the increased rental charges has a mixed impact on poverty measured on an after housing costs basis (AHC), falling slightly among the full population, but rising among pensioners. As discussed in Section 3.2, these effects on poverty are due to public subsidies for housing costs, which are targeted at low income households, and which rise with the value of rental charges.

Table 2: Effects of model parameters on simulated moments considered for calibration

	high alpha		no retirement effects		low delta		no bequest		low epsilon		high gamma		high rents	
	singles	couples	singles	couples	singles	couples	singles	couples	singles	couples	singles	couples	singles	couples
proportions of population not employed														
age 18-29	0.005	0.003	0.000	-0.001	0.008	0.008	0.006	0.005	-0.003	0.001	0.001	-0.002	0.000	-0.001
age 30-54	0.018	0.037	0.002	0.005	0.037	0.076	0.090	0.096	-0.015	-0.006	0.035	-0.004	0.000	0.000
age 55-74	0.024	0.028	-0.098	-0.117	0.004	0.026	0.159	0.141	-0.038	-0.023	0.052	0.002	0.000	0.000
geometric mean consumption (£2011 per week)														
age 18-29	-2.871	-7.022	0.228	0.741	10.279	36.642	4.593	15.421	-4.252	-2.864	9.681	8.327	8.155	11.605
age 30-54	-2.754	-12.967	1.248	8.072	13.229	68.993	25.536	92.553	-9.079	-34.765	19.582	21.433	30.011	14.241
age 55-74	-1.560	-4.169	3.168	22.309	12.742	68.909	139.796	340.303	-10.600	-21.884	29.616	2.075	16.485	7.192
proportions of population participating in private pensions														
age 18-29	0.004	0.014	-0.006	-0.004	-0.159	-0.112	-0.053	-0.002	0.015	0.016	-0.158	-0.117	-0.001	-0.002
age 30-54	-0.013	-0.012	-0.002	-0.010	-0.104	-0.197	-0.072	-0.039	0.022	-0.002	-0.132	-0.278	-0.001	-0.003
age 55-64	-0.009	-0.016	-0.003	-0.026	-0.043	-0.031	-0.048	-0.034	0.000	-0.003	-0.081	-0.062	-0.012	0.000
equivalised consumption to leisure ratio of graduates to non-graduates - aged 55-60														
	-0.066	-0.015	-0.025	-0.025	0.127	0.127	-0.181	0.123	-0.062					
proportion of population below 60% of median income (%)														
	all	retirees	all	retirees	all	retirees	all	retirees	all	retirees	all	retirees	all	retirees
BHC	-1	-2	1	-1	-2	-4	-2	-4	0	1	-2	-3	-2	-1
AHC	0	-1	0	-1	-1	-2	-1	-2	0	1	0	-2	-1	1

Source: Author's calculations on simulated data

Notes: All simulations based on the calibrated model parameters, except where explicitly indicated

'high alpha' simulation increases calibrated values of alpha by 20 percent

'no retirement effect' simulation suppresses wage penalties of pension take-up (from 55)

'low delta' simulation reduce exponential discount factors by 3 percentage points

'no bequest' simulation omits warm glow bequest motive

'low epsilon' simulation reduces intratemporal elasticity by half

'high gamma' simulation assumes: gamma = 2.0, delta(single) = 0.95, delta(couple) = 0.89, alpha reduced by 12%, zeta (single) = 3300000, zeta (couple) = 2650000

'high rents' simulation increases all simulated rental charges by 50%

7. CONCLUSIONS

This paper reports results of a calibration of LINDA, a structural dynamic microsimulation model of families that can be freely downloaded from the internet. The model is based upon a preference relation that is standard in the literature, and behavioural solutions are obtained using dynamic programming techniques. Margins of uncertainty that are explicitly included in the solution to the lifetime decision problem include wages, unemployment, relationship formation and dissolution, student status, education status, the number and timing of birth of dependent children, and the time of death. In contrast to previous empirical studies that focus on individual birth cohorts, LINDA takes an overlapping generations form that is well adapted for identifying model parameters on data reported for a representative population cross-section.

The paper reports calibration of LINDA on data reported for a representative cross-section of the UK population observed in 2011. Discussion focusses upon the information used to identify unobservable model parameters, including preference parameters that are the focus of extensive empirical debate. The calibration strategy that is described is designed to identify sets of model parameters on specific – and important – behavioural moments via a hierarchical procedure. Importantly, results obtained support the proposition that preference parameters for a structural model of savings and employment – including the parameter of relative risk aversion – can be identified on behavioural margins observed for a population cross-section at a single point in time. It is argued here that the additional complications involved in extending a dynamic programming model of savings to allow for heterogenous birth cohorts are more than off-set by the conceptual advantages derived when bringing such a model to survey data.

Parameterising a structural dynamic programming model of savings on data observed for a population cross-section at a point in time opens up a range of exciting empirical possibilities. One such possibility is to consider whether the intertemporal elasticity of substitution exhibits systematic variation with the economic cycle. This might help to explain the wide diversity of estimates that have previously been reported for this important preference parameter, with important behavioural and policy implications. Combined with on-going improvements in that accessibility of high performance computing technology, and associated advancements in empirical methods, it is hoped that such analyses will substantively improve our understanding of the decisions that people make during the next few decades.

REFERENCES

- Adda, J., & Cooper, R. W. (2003). *Dynamic Economics: Quantitative Methods and Applications*. Cambridge, Massachusetts: MIT Press.
- Ainslie, G., & Haslam, N. (1992). Hyperbolic discounting. In G. Loewenstein & J. Elster (Eds.), *Choice over Time*. New York: Russell Sage Foundation.
- Ainslie, G. W. (1992). *Picoeconomics*. Cambridge: Cambridge University Press.

- Andreoni, J. (1989). Giving with impure altruism: applications to charity and Ricardian equivalence. *Journal of Political Economy*, 97, 1447-1458.
- Attanasio, O., Bottazzi, R., Low, H., Nesheim, L., & Wakefield, M. (2012). Modelling the demand for housing over the life cycle. *Review of Economic Dynamics*, 15, 1-18.
- Attanasio, O., Low, H., & Sanchez-Marcos, V. (2005). Female labor supply as insurance against idiosyncratic risk. *Journal of the European Economic Association*, 3, 755-764.
- Attanasio, O., Low, H., & Sanchez-Marcos, V. (2008). Explaining changes in female labor supply in a life-cycle model. *American Economic Review*, 98, 1517-1552.
- Attanasio, O. P., & Weber, G. (1993). Consumption growth, the interest rate and aggregation. *Review of Economic Studies*, 60, 631-649.
- Attanasio, O. P., & Weber, G. (1995). Is consumption growth consistent with intertemporal optimization? *Journal of Political Economy*, 103, 1121-1157.
- Auerbach, A. J., & Kotlikoff, L. J. (1987). *Dynamic Fiscal Policy*. Cambridge: Cambridge University Press.
- Barsky, R. B., Kimball, M. S., Juster, F. T., & Shapiro, M. D. (1997). Preference parameters and behavioral heterogeneity: an experimental approach in the health and retirement survey. *Quarterly Journal of Economics*, 112, 537-579.
- Blundell, R., Browning, M., & Meghir, C. (1994). Consumer demand and the life cycle allocation of household expenditure. *Review of Economic Studies*, 61, 57-80.
- Christensen, B. J., & Kiefer, N. M. (2009). *Economic Modeling and Inference*. Princeton, New Jersey: Princeton University Press.
- Cooley, T. F., & Prescott, E. C. (1995). Economic growth and business cycles. In T. F. Cooley (Ed.), *Frontiers of Business Cycles Research*. Princeton: Princeton University Press.
- Deaton, A. (1997). *The Analysis of Household Surveys*. Baltimore: John Hopkins University Press.
- Deaton, A., & Muellbauer, J. (1980). *Economics and Consumer Behavior*. Cambridge: Cambridge University Press.
- Dynan, K. (1993). How prudent are consumers? *Journal of Political Economy*, 101, 1104-1113.
- Engelhardt, G. V., & Kumar, A. (2007). *The elasticity of intertemporal substitution: new evidence from 401(k) participation*. (Federal Reserve Bank of Dallas Working Papers 0812)
- Feigenbaum, J. (2008). Information shocks and precautionary saving. *Journal of Economic Dynamics and Control*, 32, 3917-3938.

- French, E. (2005). The effects of health, wealth, and wages on labour supply and retirement behaviour. *Review of Economic Studies*, 72, 395-427.
- Gallant, A. R., & Tauchen, G. E. (1996). Which moments to match? *Econometric Theory*, 12, 657-681.
- Gourieroux, C., Monfort, A., & Renault, E. (1993). Indirect inference. *Journal of Applied Econometrics*, 8, S85-S118.
- Gourinchas, P., & Parker, J. (2002). Consumption over the Life Cycle. *Econometrica*, 70, 47-89.
- Green, L., Fry, A. F., & Myerson, J. (1994). Discounting of delayed rewards: A life span comparison. *Psychological Science*, 5, 33-36.
- Grossman, S. J., & Shiller, R. J. (1981). The determinants of the variability of stock market prices. *American Economic Review*, 71, 222-227.
- Gustman, A. L., & Steinmeier, T. L. (2005). The social security early entitlement age in a structural model of retirement and wealth. *Journal of Public Economics*, 89, 441-464.
- Hairault, J. O., & Langot, F. (2008). Inequality and social security reforms. *Journal of Economic Dynamics and Control*, 32, 386-410.
- Hall, R. E. (1988). *Intertemporal substitution in consumption*. (NBER Working Paper W0720)
- Hansen, G. D., & Imrohoroglu, S. (2008). Consumption over the life cycle: The role of annuities. *Review of Economic Dynamics*, 11, 566-583.
- Hansen, P., & Singleton, K. (1983). Stochastic consumption, risk aversion and the temporal behaviour of asset returns. *Journal of Political Economy*, 91, 249-266.
- Havranek, T., Horvath, R., Irsova, Z., & Rusnak, M. (2013). *Cross-country heterogeneity in intertemporal substitution*. (William Davidson Institute Working Paper Number 1056)
- Hicks, J. R. (1939). *Value and Capital*. Oxford: Clarendon Press.
- Huyer, W., & Neumaier, A. (1999). Global optimization by multi-level coordinate search. *Journal of Global Optimization*, 14, 331-355.
- Keys, R. G. (1981). Cubic convolution interpolation for digital image processing. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 29, 1153-1160.
- Kirby, K. N. (1997). Bidding on the future: Evidence against normative discounting of delayed rewards. *Journal of Experimental Psychology-General*, 126, 54-70.
- Lagarias, J., Reeds, J. A., Wright, M. H., & Wright, P. E. (1998). Convergence properties of the Nelder-Mead simplex method in low dimensions. *SIAM Journal of Optimization*, 9, 112-147.

- Laibson, D., Repetto, A., & Tobacman, J. (2007). *Estimating discount functions with consumption choices over the lifecycle*. (Oxford University Department of Economics Discussion Paper 341)
- Lane, Clark, & Peacock. (2008). *Pension Buyouts 2008*. (www.lcp.uk.com)
- Lee, B. S., & Ingram, B. (1991). Simulation estimation of time-series models. *Journal of Econometrics*, 47, 197-205.
- Livshits, I., MacGee, J., & Tertilt, M. (2007). Consumer bankruptcy: A fresh start. *American Economic Review*, 97, 402-418.
- Mankiw, N. G. (1985). Consumer durables and the real interest rate. *Review of Economics and Statistics*, 62, 353-362.
- Mankiw, N. G., Rotemberg, J. J., & Summers, L. H. (1985). Intertemporal substitution in macroeconomics. *Quarterly Journal of Economics*, 100, 225-251.
- Mehra, R., & Prescott, E. C. (1985). The equity premium: a puzzle. *Journal of Monetary Economics*, 15, 145-161.
- Nelissen, J. H. M. (1998). Annual versus lifetime income redistribution by social security. *Journal of Public Economics*, 68, 223-249.
- Press, W. H., Flannery, B. P., Teukolsky, S. A., & Vetterling, W. T. (1986). *Numerical Recipes: the art of scientific computing*. Cambridge: Cambridge University Press.
- Sefton, J., van de Ven, J., & Weale, M. (2008). Means testing retirement benefits: fostering equity or discouraging savings? *Economic Journal*, 118, 556-590.
- Stern, S. (1997). Approximate solutions to stochastic dynamic programs. *Econometric Theory*, 13, 392-405.
- Thaler, R. (1981). Some empirical evidence on dynamic inconsistency. *Economic Letters*, 8, 201-207.
- van de Ven, J. (2016). *LINDA: A dynamic microsimulation model for analysing policy effects on the evolving population cross-section*. (NIESR Discussion Paper 459)
- van de Ven, J., & Weale, M. (2010). *An empirical investigation of quasi-hyperbolic discounting*. (National Institute of Economic and Social Research Discussion Paper 355)
- Ziliak, J. P., & Kniesner, T. J. (2005). The effect of income taxation on consumption and labor supply. *Journal of Labor Economics*, 23, 769-796.

NOTES

¹French (2005) applies a similar procedure, but uses regression techniques to improve estimated age profiles for his reference cohort by drawing upon data observed for near-by cohorts.

²An explicit allowance for evolving tax and benefits policy has, however, been implemented in dynamic microsimulation models based on analytical functional forms for behaviour; see, e.g., Nelissen (1998).

³Adjusting age profiles of income and consumption by trend growth, for example, rests upon the assumption that the economy is in a steady-state equilibrium, characterised by a stable growth path. This assumption is highly unlikely to hold for any modern economy.

⁴The modified OECD scale assigns a value of 1.0 to the family reference person, 0.5 to their spouse (if one is present), and 0.3 to each dependent child. The OECD scale is currently the standard scale for adjusting before housing costs incomes in European Union countries, and is included here to reflect the impact that family size has been found to have on the timing of consumption (e.g. O. P. Attanasio and Weber (1995) and Blundell et al. (1994)).

⁵See, for example, Andreoni (1989) for details regarding the warm-glow model.

⁶EET is short for Exempt-Exempt-Tax, referring, respectively to pension contributions, pension investment returns, and pension dispersals.

⁷This two-step procedure is well adapted to the extended computation times required to determine the implications of a given parameter combination and the large number of parameters upon which the model depends.

⁸The grid assumed for analysis has the following dimensions: 26 points for non-pension wealth between ages 18 and 74, and 151 points between ages 75 and 130; 26 points for earnings potential between ages 18 and 74; 21 points for private pension rights from age 18 to 74, and 151 points between ages 75 and 130; 2 points for wage offers between ages 18 and 74; 2 points for pension receipt from age 55 to 75; 2 points for education status from age 18 to 74, 2 points for relationship status from age 18 to age 89. Hence, the grid considered for analysis comprised 10,409,209 individual cells. This problem was solved in 19.6 minutes on a desktop workstation purchased in 2011.

⁹Econometric methods include Simulated Minimum Distance (Lee & Ingram, 1991), Method of Simulated Moments (Stern, 1997), Indirect Estimation (Gourieroux, Monfort, & Renault, 1993), and Efficient Method of Moments (Gallant & Tauchen, 1996).

¹⁰A dampening parameter often improves convergence properties of iterative search routines like the one considered here.

¹¹Publicly available microdata reported for the UK commonly refer to the 'families' considered by LINDA as 'benefit units'.

¹²Barclays Bank reports an annual 'Equity Gilt Study', which is a principal source of data concerning long-run returns in the UK.

¹³The preference relation described by equation (1) implies that $U_{cl} = (1/\varepsilon - \gamma) U_c U_l / u^{1-\gamma}$, which is positive when $1/\varepsilon > \gamma$.

¹⁴This statistic was estimated by numerically calculating the derivative $d(\Delta \ln c_{i,t}) / d \ln r_{i,t}$, where $\Delta \ln c_{i,t} = \ln c_{i,t} - \ln c_{i,t-1}$, for the reference population cross-section. The derivative was taken by perturbing interest rates up by 0.5 percentage points (giving an elasticity estimate of 0.263), and down by 0.5 percentage points (giving an estimate of 0.285). The average between these two estimates is reported here.