

Bourquin, Pascale; Joyce, Robert; Sturrock, David

Research Report

Inheritances and inequality over the life cycle: what will they mean for younger generations?

IFS Report, No. R188

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: Bourquin, Pascale; Joyce, Robert; Sturrock, David (2021) : Inheritances and inequality over the life cycle: what will they mean for younger generations?, IFS Report, No. R188, ISBN 978-1-80103-035-9, Institute for Fiscal Studies (IFS), London, <https://doi.org/10.1920/re.ifs.2021.0188>

This Version is available at:

<https://hdl.handle.net/10419/235074>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Model Appendix to ‘Inheritances and inequality over the life cycle: what will they mean for younger generations?’

Pascale Bourquin

Robert Joyce

David Sturrock

April 22, 2021

Abstract

This appendix formally sets out the lifecycle model of savings and consumption in the presence of parental wealth and inheritances that is employed in the IFS report ‘Inheritances and inequality over the life cycle: what will they mean for younger generations?’.

1 Model

Here, we formally set out the heterogeneous agents life cycle model that we employ in Bourquin et al. (2021). The agent is a household that is endowed with a level of education (either attains compulsory schooling attendee only i.e. GCSEs; completed high school i.e. achieved A-levels, or completed higher education) and begins life at age 26, faces survival risk from age 65 and dies at latest at age 110. In the case where a household is a couple, this education level can be interpreted as the highest educational level of the two couple members and household death as the death of the second member of the couple to die. The household receives exogenously given earnings from work and a state pension in retirement and makes consumption (equivalently, savings) decisions in each period. Each household has two parental households. The parental households have a level of wealth whose value evolves according to an exogenous process. The timing of the death of the parental households is uncertain and when they die they may leave an inheritance which is split equally between their heirs.

We now set out more detail on each aspect of the model. Many of the model parameters vary by decade of birth. For ease of notation, we suppress this dependence on decade of birth when describing the model. Where components vary by decade of birth, this is made clear in the description of the model estimation.

1.1 Preferences and demographics

The household is assumed to have constant relative risk aversion utility, defined as a the following function of consumption, $c_{i,t}$, within each period:

$$u_t(c_{i,t}) = \theta_t \frac{(c_{i,t}/\theta_t)^{1-\gamma} - 1}{1-\gamma} \quad (1)$$

where θ_t is the equivalisation factor. The household faces mortality risk in each year from age 65 onwards. The unconditional probability of survival to age t is given by S_t and the

probability of survival to t conditional on survival to $t - 1$ is given by s_t . The household receives warm glow utility from bequests, b_i according to the function (following De Nardi (2004))

$$\phi(b_i) = \phi_1 \frac{(b_i + \phi_2)^{(1-\gamma)}}{1 - \gamma} \quad (2)$$

where ϕ_1 determines the strength of the bequest motive and ϕ_2 the degree to which bequests are a luxury. Lifetime utility is therefore given by

$$U_i = E_0 \left[\sum_{t=0}^{t=110} S_{t-1} \beta^t [s_t u_t(c_{i,t}) + (1 - s_t) \phi(b(a_t))] \right] \quad (3)$$

where β is the discount factor.

1.2 Exogenous processes

Earnings

The household receives earnings in each period from age 26 to the latest retirement age, K . Earnings household i in period t are denoted $e_{i,t}$. Earnings can be zero, representing voluntary or involuntary unemployment or early retirement. We denote the household's binary employment status as $E_{i,t}$. Log earnings conditional on working are the sum of an education-specific deterministic age component, $f_{ed}^e(age_{i,t})$, a permanent household fixed effect and a persistent stochastic component, $\eta_{i,t}$:

$$\begin{aligned} \ln(e_{i,t}) &= f_{ed}^e(age_{i,t}) + \zeta_i + \eta_{i,t} & E_{i,t} &= 1 \\ e_{i,t} &= 0 & E_{i,t} &= 0 \end{aligned}$$

Following the framework of Arellano et al. (2017), we assume that the stochastic component follows a first order Markov process. The persistent earnings component is drawn from a distribution that varies with age and with its lagged value and is given by the series of conditional quantile functions $Q_t(\cdot | \eta_{i,t-1})$. Employment status is assumed to be drawn from a

distribution that varies by age and lagged employment status and lagged persistent earnings component. The draws of the persistent component and employment status are independent across time and independent of each other. These conditions can be written as

$$\eta_{i,t} = Q_t(u_{i,t}|\eta_{i,t-1}) \quad (4)$$

$$E_{i,t} = \mathbb{1}\{v_{i,t} > \bar{v}(E_{i,t-1}, \eta_{i,t-1})\} \quad (5)$$

$$u_{i,t} \stackrel{iid}{\sim} U(0, 1) \quad (6)$$

$$v_{i,t} \stackrel{iid}{\sim} U(0, 1) \quad (7)$$

$$u_{i,t} \perp\!\!\!\perp v_{i,s}, \quad \forall \quad t, s \quad (8)$$

Number of living parental households

Each household may start life with up to two living parental households. These parental households face a probability of death in each year that is a function of their age and the (child) household's education level. We assume that the age gap between living parents and their children is constant, conditional on education, such that parental survival is equivalently a function of the household's age. The realisation of the parental households' mortality is drawn independently between the two households (if there are two). We denote the unconditional survival curve for the parental household j of household i with the set of probabilities $\{S_t^{pj}\}_{t=0}^{t=111}$. These assumptions then define a law of motion for the number of parents, $P_t \in \{0, 1, 2\}$, denoted

$$P_{t+1} = \Phi_t^{age_{i,t}, ed_i}(P_t).$$

Parental wealth

Living parental households have a level of wealth that evolves according to an exogenous stochastic process. We denote the level of wealth of parental household j of household i at time t as $a_{i,t}^{pj}$. As in the case of earnings, log parental wealth is assumed to be the sum of a

deterministic component which is a function of the parental household's age and education level, $\tilde{f}_{ed}^a(age_{i,t}^{p_j})$ and a stochastic component, $\mu_{i,t}^j$:

$$\ln(a_{i,t}^{p_j}) = \tilde{f}_{ed}^a(age_{i,t}^{p_j}) + \mu_{i,t}^j$$

The stochastic component is assumed to follow a first order Markov process as given by a series of conditional quantile functions:

$$\mu_{i,t}^j = Q_t(v_{i,t}^j | \mu_{i,t-1}^j) \quad (9)$$

$$v_{i,t}^j \stackrel{iid}{\sim} U(0, 1) \quad (10)$$

Inheritances

The inheritances received by a household depend on the processes for parental wealth and parental survival in the following way. When a parental household dies, their wealth is left as a bequest, which is taxed according to the function $b(\cdot)$ and split between their heirs, n_k^{ed} , to yield an inheritance. The number of heirs of parental household is a function of the level of education of the (child) household. Denoting the inheritance received by household i from parental household j as $H_{i,t}^j$, we can write the level of inheritance received as a function of the above processes as follows:

$$H_{i,t}^j = \begin{cases} 0 & \text{if } P_t = P_{t-1} \\ b(a_{i,t}^j)/n_k^{ed} \text{ w.p. } 0.5 & \text{if } P_t = P_{t-1} - 1 \\ b(a_{i,t}^j)/n_k^{ed} & \text{if } P_t = P_{t-1} - 2 \end{cases} \quad (11)$$

Initial conditions

The household earnings fixed effect and persistent earnings shock are drawn independently from education-specific distributions. The initial number of parental households is drawn

from an education-specific distribution. The initial level of each persistent parental wealth shock are drawn independently from an education-specific distribution which varies by education and the level of the household earnings fixed effect.

1.3 Choices and constraints

In each period the household chooses its level of consumption or, equivalently, its level of net saving, $z_{i,t}$. Saving is into a single, riskless asset whose returns vary with the household's age. Agents may borrow but must repay all borrowing with certainty by age 75 (they face the implied natural borrowing constraint for ages younger than 75). Inheritances received add to the household's assets.

A government levies a labour tax on income less net savings, according to the tax function $T_t(\cdot)$, and provides a public pension, $sp_t(e_{i,K-1})$, that is a function of education and earnings in period $K - 1$.

This can be formalised through the following budget constraints:

$$y_{i,t} = e_{i,t} + sp_{i,t} + H_{i,t}^1 + H_{i,t}^2 \quad (12)$$

$$c_{i,t} = T_t(y_{i,t} - z_{i,t}) \quad (13)$$

$$a_{i,t+1} = (a_{i,t} + z_{i,t})(1 + r_{t+1}) \quad (14)$$

$$a_{i,t+1} \geq \underline{a} \quad (15)$$

where $y_{i,t}$ is gross income in period t , $a_{i,t}$ is the level of assets held in time t , r_{t+1} is the net rate of return on assets held from time t to time $t + 1$ and \underline{a} is the natural borrowing limit for period $t + 1$.

1.4 Timing and household problem

We now formally define the household problem. We first clearly state the within-period timing assumptions. The timing of events each period is as follows:

1. The household may die, potentially leaving a bequest
2. If the household does not die, their earnings of evolve
3. Any remaining parental households have their wealth evolve
4. One or more of the remaining parental households may die, resulting in an inheritance
5. Household makes their consumption/savings choice

Before formally defining the solution to the household problem, we note that the household problem can be re-written in a way that will be convenient when solving the model by defining the variable “cash in hand” as the sum of assets and income: $M_{i,t} = e_{i,t} + sp_{i,t} + H_{i,t}^1 + H_{i,t}^2 + a_{i,t}$. The solution to the household problem solves the following Bellman equation for the household problem (subscript i is committed for clarity)

$$\begin{aligned}
V_t(\eta_t, M_t, P_t, \mu_t^1, \mu_t^2; \zeta, ed, cohort) = & \max_{c_t} \{u(c_t; \theta_t) + \\
& \beta s_{t+1} \int V_{t+1}(\eta_{t+1}, M_{t+1}, P_{t+1}, \mu_{t+1}^1, \mu_{t+1}^2; \zeta, ed, cohort) \\
& dF(\eta_{t+1}, M_{t+1}, P_{t+1}, \mu_{t+1}^1, \mu_{t+1}^2 | \eta_t, M_t, P_t, \mu_t^1, \mu_t^2; \zeta, ed, cohort) \\
& + (1 - s_{t+1}) \beta \phi(b(a_{t+1}))\}
\end{aligned}$$

subject to the budget constraints and laws of motion for earnings, parental wealth and number of parental households. The time-invariant state variables for this problem are cohort, education and the household earnings fixed effect. The time-varying state variables of this problem are: time (equivalently age), education, the persistent component of earnings, cash in hand, the number of living parental households and the levels of the persistent component of wealth for the parental households. This problem has no analytical solution and must be solved numerically.

2 Identification and estimation of exogenous processes

In this section, we set out the identification of the exogenous processes for household earnings and parental wealth that feed into the model of household behaviour. We also set out the data and procedure used to estimate the model parameters.

2.1 Earnings process

2.1.1 Empirical specification and identification

The earnings model set out above places minimal restrictions on the form of the deterministic earnings component and the process for the persistent earnings component. In order to take this process to the data, we will must specify the form of the deterministic component, conditional quantile functions and initial conditions in a way can be estimated.

Deterministic component

The deterministic component of earnings in our model is a function of education, age and time. In the description of our model above, we abstracted from differences in earnings across cohorts. As the estimation of our earnings model will require us to pool data from multiple cohorts in order to have observations of households of all working ages, we allow for differences by 10-year cohort in the deterministic component of earnings. We specify the deterministic function as

$$f_t^{c,ed}(age_{i,t}) = \beta_{ce} + \sum_{a=16}^{64} \gamma_{ae} \mathbb{1}\{age_{i,t} = a\} + \delta_t + \sum_{k=1}^{45} DP_{it}^k \quad (16)$$

where β_{ce} is a 10-year birth cohort and education-specific intercept, $\{\gamma_{ae}\}$ is an education-specific set of individual year-of-age effects, δ_t is a linear time trend and DP^k are a series of time dummies constrained to sum to zero.¹ We are allowing for levels of earnings that differ

¹This restriction, proposed in Deaton and Paxson (1994), means we do not have collinearity of age, time and cohort controls.

by the interaction of 10-year birth cohort and education level. The age profile of earnings is flexibly given by single year of age dummies that vary by education level. Time effects are assumed to take the form of common deviations around a linear trend. This specification differs from Arellano et al. (2017), which includes only a set of age dummies.

The identification of the parameters of type of specification with cross-sectional data is standard. We can separately identify the cohort and age effects for each education group so long as we have cohort overlap, given the assumption of a constant age profile across cohorts, within education groups. The time trends are separately identified from the age and cohort effects given the Deaton-Paxson restriction, which imposes a particular form on the effects of time. The assumption around the commonality of the age profile across cohorts may seem strong, but produces a close fit for each cohort in the first stage of the estimation process.

Stochastic component

We follow the empirical specification of Arellano et al. (2017) and specify the quantile functions as the sum of a set of products of low-order hermite polynomial functions of age, lagged value of the persistent earnings component and vary over a grid of the shock distribution. We also allow for a transitory component to earnings as in the setting of Arellano et al. (2017). We denote this $\epsilon_{i,t}$. This component can be interpreted as in part measurement error. It is for this reason that we discard this part of the earnings process when feeding it into our model. This transitory component is specified as an age-varying quantile function, with an empirical specification set out in an analogous way to that for the persistent component.² We make the addition, compared to Arellano et al. (2017), of allowing the quantile functions to vary fully flexibly by education group.

Arellano et al. (2017) show that the parameters of an earnings process of the form set out are identified given panel data with at least 5 periods. We give a brief summary here. The marginal distributions of ϵ_t and η_1 and the conditional distributions of η_t given η_{t-1} (for each t

²We also follow Arellano et al. (2017) in other aspects of the empirical specification not discussed here for brevity. See the appendix for further details.

from 1 to $T - 1$) are identified given $T \geq 3$ (i.e. 4 periods of data) and some conditions on the process, including a completeness condition. Intuitively, the completeness condition requires that there is some dependence of the η 's such that they can be distinguished from the ϵ 's. The intuition for why we require $T \geq 3$ is as follows. Consider any 3 period subpanel of data $(e_{i,0}, e_{i,1}, e_{i,2})$. The model implies that conditional on $\eta_{i,1}$, $e_{i,0}$, $e_{i,1}$ and $e_{i,2}$ are independent. This allows us to identify the distribution of $\eta_{i,1}$ and hence $\epsilon_{i,1}$. These can be identified for the middle periods in any panel with $T \geq 2$. With $T \geq 3$ we can therefore identify the joint distribution of $(\epsilon_{i,1}, \dots, \epsilon_{i,T-1})$ due to the serial independence of the transitory shocks. This means we can identify the joint distribution of $(\eta_{i,1}, \dots, \eta_{i,T-1})$. Therefore if $T - 1 > 1$ i.e. $T \geq 3$ (i.e. 4 periods of data available), we can identify one transition of η . For any panel with $T \geq 3$ we can identify all other than the first and last transitions. Intuitively, separately identifying the household fixed effect then requires one additional period in the panel.

If we impose further restrictions then we can identify $\eta_{i,0}$, $\eta_{i,T}$, $\epsilon_{i,0}$, $\epsilon_{i,T}$ and the first and final transitions for η . In particular, as we have data on multiple cohorts for whom $t = 0$ and $t = T$ occur at different ages, then by the assumption of the invariance of these distributions over time/cohorts, we can identify these objects at all ages other than the oldest and youngest ages in our samples. Further if we make parametric restrictions (as we do in the empirical specification set out above) then we can extrapolate to the oldest and youngest ages.

Probabilities of employment

We specify the probabilities of being in employment as a function of education, age and the persistent stochastic component of earnings. As with the level of earnings, these probabilities are in practice allowed to vary also by 10-year birth cohort. The empirical specification that we implement allows these probabilities to vary non-parametrically by age, education, 10-year cohort and by a fixed number of quantiles of the lagged persistent component.

2.1.2 Data and sample

We now set out the data used to estimate this model. We draw upon data from the Family Expenditure Survey (FES) and its successor surveys (the Expenditure and Food Survey and the Living Costs and Food Survey), which are cross-sectional surveys covering the years 1968 to 2018. We will refer to this collection of surveys as “the FES”. We also use the UKHLS, a household panel survey running 1991 to 2018. The “British Household Panel Survey” waves of the data cover 1991 to 2008 and these survey waves are annual, with fieldwork in September to December of each year. The “Understanding Society” waves cover 2009 to 2018 and are rolling 2-year periods. In all surveys, individuals are interviewed at approximately 12 month intervals and data can be treated as annual.

We use two different datasets for the following reason. The FES covers a longer time period, enabling us to more precisely disentangle age, time and cohort effects. However, to estimate the stochastic processes, we require panel data and must use the UKHLS. Since the estimation of the deterministic profile only requires cross-sectional data, we use the FES to estimate the parameters of this function. We then strip out the deterministic component of earnings from the UKHLS data using the same method and estimate the stochastic process for earnings using this data.

We are interested in modelling processes for total household (or, more accurately, “benefit unit”, corresponding to the fiscal unit i.e. a couple or single individual) pre-tax earnings. We drop dependent children from our sample. When modelling benefit unit earnings, we define a couple’s age as the mean of their ages and their level of education as the highest achieved by the couple. The 10-year birth cohort variable is defined based on the mean birth year of the couple. We keep observations from the 1930s through to the 1980s birth cohorts. We define individual education using a three-way education categorisation based on the highest qualifications achieved by the individual: low: up to and including GCSEs (or no qualifications for those born before 1958), mid: A-level or equivalent (or GCSEs for those born before 1958), high: higher education degree.

When selecting a balanced panel of benefit units as required for the estimation of the stochastic components, we define a benefit unit as the same benefit unit in another period if it has the same members (excluding dependent children).³

The measure of earnings that we use is annual gross real earnings, including self-employment income.⁴

2.1.3 Estimation

To estimate the parameters of the deterministic component of earnings, we pool all observations of benefit unit earnings from the FES for those born between the 1930s and 1980s and run OLS estimation of the empirical counterpart to Eq. (16):

$$\ln(e_{it}) = \sum_{c=30s}^{80s} \sum_{e=1}^3 \beta_{ce} cohort_{it} \times ed_{it} + \sum_{a=16}^{64} \sum_{e=1}^3 \gamma_{ae} age_{it} \times ed_{it} + \delta_t + \sum_{k=1}^{40} DP_{it}^k + v_{i,t} \quad (17)$$

The second step is to use obtain the residuals from the equivalent OLS regression on the UKHLS. We then take all non-overlapping sets of 6 consecutive observations of a benefit unit. We estimate the parameters of the empirical specification of the conditional quantile functions by using the EM algorithm. A description of the algorithm is given in the appendix and in Arellano et al. (2017). Note that in this step we pool all observations of a given education group together (combining cohorts) and estimate sets of parameters separately by education group. We assume here that the process for the stochastic component of earnings varied by education group but not by cohort.

The final step is to estimate the employment probabilities. Our approach here is somewhat ad-hoc since the existing literature on earnings process estimation does not yet model non-employment jointly with the evolution of earnings in the form set out above. We pool all

³When selecting balanced panels of both individuals and benefit units, we enforce that a panel cannot run across the ‘seam’ between the BHPS and Understanding Society surveys. When selecting balanced panels we also drop observations where consecutive interviews are less than 9 months or more than 15 months apart (this is 3% of interviews, causing us to drop 14% of otherwise usable subpanels).

⁴We use the IFS’ Household’s Below Average Incomes “before housing costs” variant of the consumer price index (estimated based on the RPI for years before 1997-98) to convert nominal values to real terms.

observations in the UKHLS of benefit units observed for two consecutive periods conditional on being in employment in the first period. We then calculate, conditional on age, education, 10-year birth cohort and quantile of lagged earnings the proportion of observations that are in employment in the second observation. Given the observed employment rate for each group at each age, we can then back out the implied probability of being in employment in the following period conditional on being out of work in the previous period, for each quantile of the persistent component of earnings.

We show assessments of model fit in the appendix.

2.2 Parental wealth process

We now set out our empirical specification, identification and estimation of the process for parental households' wealth. This takes a similar form to the process for earnings.

2.2.1 Empirical specification and identification

In the model set out in section 1, wealth is the sum of a deterministic component that is a function of education and age, and a stochastic component whose process is given by age-varying conditional quantile functions. We set out the empirical specification given to these two components here.

Deterministic component

We specify the function that defines the deterministic component of wealth as the sum of an education and cohort-specific intercept, an education-specific fourth order polynomial in age and time effects that are the sum of a linear trend and a series of time effects constrained to sum to zero (i.e. we again make the Deaton and Paxson (1994) restriction).

$$f_{ed}^a(age_{i,t}^{p_j}) = \alpha_{ce} + \phi_{ae}^1 age_{i,t}^{p_j} + \phi_{ae}^2 (age_{i,t}^{p_j})^2 + \phi_{ae}^3 (age_{i,t}^{p_j})^3 + \phi_{ae}^4 (age_{i,t}^{p_j})^4 + \delta_t + \sum_{k=1}^7 DP_{it}^k \quad (18)$$

We note that this function depends on the (child) level of education. One might expect a more natural specification to be that the function vary by parental household education. We make this specification as parental households' level of education is not a state variable of our lifecycle model. Intuitively, the intercept captures the average level of parental wealth for households of a particular level of education and the age coefficients capture the average age trend of parental wealth (in parental age) within that group. We note also that as parental age is not a state variable of our lifecycle model, when we feed the parental process in the model we will assume a constant age gap between parents and children conditional on birth cohort and education level and hence assume away heterogeneity in the parental wealth process that comes from heterogeneity in this age gap.

The identifying assumptions include those analogous to in the case of earnings. We require that, conditional on education, the age-profile of parental wealth is the same across cohorts. One threat to this might be if different cohorts of parental households draw down on their wealth at different rates on average. This might be expected if, for example, the parents of later-born households expect to live longer and so have a more gradual path of decline of wealth at older ages. To some extent, changes in the average rate of wealth drawdown across cohorts are allowed for to the extent that these are captured by a different educational composition of the households in the cohort. Given we are examining parental households at older ages, differential mortality according to wealth means that we will observe a selected sample (only those who survive) in the cross-section. We can augment the above specification in two ways to deal with this. First, we can use data only from balanced panels of observations and secondly, we can interact the intercepts with a set of dummy variables indicating the first age at which an individual was observed. Under the assumptions that (1) level but not the age profile of wealth varies systematically by age of death and (2) we observe each cohort at some age before there has been any differential mortality by wealth level then we will recover the common age profile and intercept terms for each cohort and education group. These assumptions of course rule out that parental households might draw

down their wealth faster in response to news about the resolution of uncertainty over the timing of their death.

Stochastic component

Here we follow the same empirical specification as used for the household earnings process with the same identifying assumptions applying. One exception is that we restrict the quantile functions by imposing a no-crossing condition in all years other than every 10 years. We do this because the model without these restrictions fails to fit the long-run persistence of parental wealth. We assume that the process for parental wealth varies by child education level but does not vary across cohort and is not affected by survivor bias.

2.2.2 Data and sample

We draw on data from the English Longitudinal Study of Ageing (ELSA) a biennial household panel that began in 2002-03. There are currently 9 waves available. We use the measured level of household non-pension wealth (the overwhelming majority of pension wealth for the cohorts we examine is non-bequeathable defined benefit pension wealth) which consists of the sum of net property wealth (including second homes), business, physical and net financial assets.

ELSA contains information about the number and year of birth of all children of sample members. This allows us to use ELSA where the level of observation is the ‘child’. ELSA does not contain information on the educational attainment of the children of sample members. We therefore impute child education in a 2-step procedure. This procedure draws on two datasets that have data linking parents and their children and contain parent characteristics and child educational attainment. First, we use the UKHLS. The UKHLS follows household ‘split offs’ of original sample members, including those who are originally children in a household and leave home to form their own household at older ages. There is also a wealth module in selected waves of the UKHLS. This allows us to determine, for each percentile of

the distribution of parental wealth, the distribution across child education levels. Specifically, we make a non-parametric estimation, for each parental wealth percentile, of the percentage of children with each education level. Here the unit of observation is again the ‘child’. We make rankings of parental wealth levels based on the parental wealth observation from when parental household was aged closest to 50 (this is approximately the starting age of parents in our model).⁵ Secondly, we use estimates from Bourquin et al. (2020) of the relationship between parental characteristics (including their education, social class, housing tenure and region, but not including wealth) and child education. These estimates are taken from probit regressions estimated by where the outcome variable is the child education level and the explanatory variables are parent characteristics. These models were estimated separately for each parent birth-decade and child birth-decade combination.

We use these two sets of estimates to impute child education within ELSA in the following way. Firstly, within parental wealth ranks in ELSA (where rankings were calculated based on a household’s first wealth observation - which will be that closest to age 50 - and within wave and cohort), we rank observations according to the predicted probability that they are high-educated, using the estimated coefficients from the models estimated using the LS in Bourquin et al. (2020). Second, we assign the first X% to be highly educated, the next Y% to be mid-educated and the remaining 1-X-Y% to be low educated, according to the proportions estimated for that percentile using the UKHLS.

In our sample, we include observations of benefit units born between 1950 and 1989 that have wealth information and the required covariates. From these, we select all observations that are observed for 6 consecutive waves or more.

⁵We make these rankings separately by parent decade of birth and wave in order to account for age and time effects.

2.2.3 Estimation

Estimation of the deterministic component of the wealth process is by means of OLS estimation of the below equation.

$$\ln(a)_{it} = \sum_{c=50s}^{80s} \sum_{e=1}^3 \sum_{p=1}^4 \alpha_{ce} \text{cohort}_{it} \times ed_{it} \times p_{i,t} + \sum_{e=1}^3 \phi_{ae} \tilde{f}(\text{age}_{it}) \times ed_{it} + \delta_t + \sum_{k=1}^7 DP_{it}^k + v_{i,t} \quad (19)$$

where $\tilde{f}(\text{age}_{it})$ is a quartic in age. Note that the cohort-and education-specific intercepts are interacted with a series of dummy variables denoted by $p_{i,t}$. These dummies record which wave a household is first observed in the survey (given that we restrict to observations that are present for at least 6 waves and there are 9 waves of ELSA, this means that observations are first observed in either wave 1, 2, 3 or 4). This estimation using only observations of households with positive levels of wealth (around 90% of observations). Note that the interpretation of the estimated age profile of log wealth is the expected level of log wealth, conditional on wealth being positive.

The second stage estimation uses the residuals from the first stage estimation. For households with negative levels of wealth, assign them a level of log wealth of zero and assign them a residual equal to the negative of the predicted level of log wealth from the estimated first stage relationship. In effect, we are bottom-coding the wealth distribution at zero for use in the second stage. This creates a mass point of low negative levels of the residual component. This means that in our second stage estimation, we cannot model any dynamics within those who have negative wealth levels and any heterogeneity in the evolution of their subsequent levels of wealth, if positive. Given that debts are not heritable, this is not a concern for us.

3 Parameterising and solving the model

This sections sets out the parameterisation of some elements of the model and how it is solved numerically.

3.1 Parameterisation

There are several other components that must be parameterised in the model. Here, we briefly outline the data sources used for each.

Tax and benefit system

Taxation of labour incomes in the UK is at the individual level. Given that we are using a household model, we therefore estimate a household level tax function. We do this by using the FES data from 1968 to 2018 , which includes measures of pre-tax income and post-tax income. Using these, we estimate for each year, the following tax function using nonlinear least squares:

$$T_t(y_{i,t}) = \psi_t^a + \psi_t^b(y_{i,t})^{\psi_t^c} \quad (20)$$

For future, years we assume that tax system remains unchanged in its 2018 form.

Estate tax

The estate tax is of a form designed to capture the features of the UK inheritance tax over the relevant period. The UK inheritance tax is set at a 40% rate on the value of estates over a threshold. There are a number of exemptions and additional allowances. The most relevant of these in the vast majority of cases of these is for owner-occupied housing. As we model only total wealth, we assume that a certain share of total wealth is held as housing and apply the inheritance tax system, assuming that the remainder of wealth attracts no exemptions.

Public pensions

Public pensions in the UK are based on an individual’s full history of earnings and employment.⁶ We are restricted to include in our model a public pension that is a function of the model’s state variables. In this case, the relevant variables are education (relevant here as proxy for lifetime average earnings) and earnings in period $K - 1$, the period before the pension is received.

Our approach to construct a household level state pension function is as follows. First, we simulate our household earnings process 10,000 times for each of our cohort and education groups. We then assign household earnings within members of a couple by using the mean shares of earnings of received by the first and second-highest earning members of couples by using shares estimated from the FES.⁷ We then use a pensions calculator which calculates entitlements for each simulated individual. Couples have their entitlements recombined to give a household level of pension income. We then estimate, separately for each education and cohort group, entitlements as a linear function of final period earnings.

Our estimated functions capture the main features of the UK state pension system: higher educated individuals receive higher entitlements but the system is progressive. For the later cohorts, the system becomes less related to earnings but more generous for lower earners, reflecting reforms to the UK state pension system over time.

Survival curves

We estimate individual survival curves for each cohort and education group. We combine UK Office for National Statistics data for estimated survival curves that are specific to year of birth and sex with mortality data from ELSA to create survival curves that vary also by individual education level.⁸

⁶Entitlement can also be gained for some other activities including the receipt of out of work benefits but we abstract from this.

⁷We estimate a the share of earnings amongst couples as cubic function of age interacted with education plus the interaction of cohort and education.

⁸We follow the method set out in Sturrock and O’Dea (2020).

With these individual survival curves, we are able to construct household survival curves for each 10-year birth cohort and education group by taking the observed distribution of birth years and education levels for both individuals and couples and assuming independence of timing of death within couples. This gives us a set of survival curves for the final member of the household that we use for both households and their parental households.

The joint distribution of the initial level of earnings and parental wealth

The joint distribution of the initial level of earnings and parental wealth is an important input into our model. We draw on the UKHLS which contains information on linked parent-child pairs. We take all observations where parental wealth is observed and child earnings are observed (while in their 20s). We rank parental wealth as described in section 2.2.2. We rank child earnings by removing age effects from each earnings observation and the taking the mean level of observed earnings from all of the child’s earnings observations and ranking children within cohort according to this earnings rank. We then non-parametrically estimate the joint distribution of quantiles of initial earnings and parental wealth within each education group.

Behavioural parameters

We parameterise the coefficient of relative risk aversion at 3 following Crawford and O’Dea (2020). To parameterise the bequest function, we use the estimates of Lockwood (2018). Specifically, we convert the parameter governing the extent to which bequests are a luxury into 2018 pounds. We then use the implied marginal propensity to consume out of final period assets implied by the parameter estimates in that paper, and back out the weight on the bequest function which gives the same value for our model. The discount factor is set to 0.955. This value was selected using a grid-search method, comparing the median wealth levels predicted by the model to the median levels recorded in the Wealth and Assets Survey.

3.2 Solving the model

We solve the model numerically. In order to do this, we need to make our earnings and parental wealth processes into discretised processes with Markov transition matrices. We do this by following the method of De Nardi et al. (2020). We simulate the processes for the stochastic components of earnings/wealth a large number of times, divide the distributions of earnings/wealth at each age into a certain number of quantiles and calculate the transition probabilities between quantiles in the simulations. We set the level of earnings/wealth for that quantile as equal to the median simulated level of earnings/wealth in each quantile. This grid of levels of earnings/wealth and transition probabilities defines a discrete Markov process.⁹ We note that the periodicity of the wealth process is 10 years. Therefore in 9 out of 10 years it is defined by a transition matrix equal to the identity matrix.

We construct a grid of values for cash-in-hand based on the borrowing constraint and the maximum attainable level of cash in hand at each age. We then begin in the terminal period and, for each combination of the state variables, we calculate the optimal choice of end-of-period assets, taking into account the possible evolution of the household's own mortality, earnings, parental mortality and parental wealth. We solve the model recursively, working back until the initial period.

References

- Arellano, M., R. Blundell, and S. Bonhomme (2017). Earnings and Consumption Dynamics: A Nonlinear Panel Data Framework. *Econometrica* 85(3), 693–734.
- Bourquin, P., R. Joyce, and D. Sturrock (2020). Inheritances and Inequality Within Generations. Technical report, IFS.

⁹We use 7 quantiles of the following sizes for wealth: 10%, 15%, 15%, 20%, 15%, 15%, 10%. We use 3 quantiles of the following sizes for the household fixed effect: 33%, 33%, 33%. We use 3 quantiles of the following sizes for the stochastic earnings component: 25%, 50%, 25%.

- Bourquin, P., R. Joyce, and D. Sturrock (2021). Inheritances and inequality over the life cycle: what will they mean for younger generations? Technical report, IFS.
- De Nardi, M. (2004). Wealth inequality and intergenerational links. Review of Economic Studies 71(3), 743–768.
- De Nardi, M., G. Fella, and G. Paz-Pardo (2020). Nonlinear Household Earnings Dynamics, Self-Insurance, and Welfare. Journal of the European Economic Association 18(2), 890–926.
- Deaton, A. S. and C. H. Paxson (1994). Saving, Growth, and Aging in Taiwan. In Studies in the Economics of Aging, pp. 331–362. University of Chicago Press.

A Details of earnings process estimation and model fit

A.1 Further details of empirical specification

We use the same specification of the order of the hermite polynomials for each component of the process as used by Arellano et al. (2017) (tensor products of polynomials of degree 3 (in η_{t-1}) and 2 (in age) for the transition of η , and polynomials of degree 2 in age for η_0 and ϵ). Again following Arellano et al. (2017), we use quantile regressions of earnings on their lagged values and age to set initial parameter values. We use the same variances for the random walk proposals in the Metropolis-Hastings sampler as Arellano et al. (2017), which yields an acceptance rate of around 0.20-0.25.

A.2 Further details of estimation

The estimation procedure is based on quantile regressions corresponding to the restrictions embodied in the empirical specification. We begin with a balanced panel (in our case $T = 6$) dataset of observations of log of gross earnings and age. Log earnings have been purged of cohort, age and time effects in the initial regression step. Denoting the parameter vector θ , the posterior density of $(\eta_{i,0}, \dots, \eta_{i,T})$ given $(y_{i,0}, \dots, y_{i,T})$, the ages at which the individuals is observed and θ , as $f(\eta_i|y_i, age_i; \theta)$, we carry out the following steps:

1. **Select initial values for the parameter vector, $\hat{\theta}^{(0)}$.** These are selected by running a series of quantile regressions. For the parameters for the distribution of η_0 , we estimate quantile regressions of observed earnings on hermite polynomials in age. For the distribution of ϵ_i we estimate quantile regressions of observed earnings on a different set of hermite polynomials in age. For the distribution of η_t given η_{t-1} we estimate quantile regressions of observed earnings on hermite polynomials in lagged earnings and age (and their products).
2. **Draw $\{\eta_i\}_{i=1}^N$ from $f(\eta_i|y_i, age_i; \hat{\theta}^{(s)})$, using the current parameter vector $\hat{\theta}^{(s)}$.**

A set of permanent components $\{\eta_i\}_{i=1}^N$ clearly implies a set of transitory components $\{\epsilon_i\}_{i=1}^N$, given observed earnings.

3. **Update the parameter vector to $\hat{\theta}^{(s+1)}$** by computing, for each $l = 1, \dots, L$:

$$(a_{1,l}^{Q(s+1)}, \dots, a_{K,l}^{Q(s+1)}) = \underset{(a_{1,l}^Q, \dots, a_{K,l}^Q)}{\operatorname{argmin}} \sum_{i=1}^N \sum_{t=1}^T \rho_{\tau_l} \left(\eta_{i,t} - \sum_{k=1}^K a_{k,l}^Q(\tau) \varphi_k(\eta_{i,t-1}, age_{i,t}) \right) \quad (21)$$

$$(a_{1,l}^{\epsilon(s+1)}, \dots, a_{K,l}^{\epsilon(s+1)}) = \underset{(a_{1,l}^\epsilon, \dots, a_{K,l}^\epsilon)}{\operatorname{argmin}} \sum_{i=1}^N \sum_{t=0}^T \rho_{\tau_l} \left(y_{i,t} - \eta_{i,t} - \sum_{k=1}^K a_{k,l}^\epsilon(\tau) \varphi_k(age_{i,t}) \right) \quad (22)$$

$$(a_{1,l}^{\eta(s+1)}, \dots, a_{K,l}^{\eta(s+1)}) = \underset{(a_{1,l}^\eta, \dots, a_{K,l}^\eta)}{\operatorname{argmin}} \sum_{i=1}^N \rho_{\tau_l} \left(\eta_{i,0} - \sum_{k=1}^K a_{k,l}^\eta(\tau) \varphi_k(age_{i,0}) \right) \quad (23)$$

$$\hat{\lambda}_-^{Q(s+1)} = - \frac{\sum_{i=1}^N \sum_{t=1}^T \mathbb{1}\{\eta_{i,t} < \sum_{k=1}^K a_{k,l}^Q(\tau) \varphi_k(\eta_{i,t-1}, age_{i,t})\}}{\sum_{i=1}^N \sum_{t=1}^T \left(\eta_{i,t} - \sum_{k=1}^K a_{k,l}^Q(\tau) \varphi_k(\eta_{i,t-1}, age_{i,t}) \right) \mathbb{1}\{\eta_{i,t} < \sum_{k=1}^K a_{k,l}^Q(\tau) \varphi_k(\eta_{i,t-1}, age_{i,t})\}} \quad (24)$$

and analogously for the other tail parameters, where $\rho_\tau(u) = u(\tau - \mathbb{1}\{u \leq 0\})$ is the check function.

4. **Iterate steps 2 and 3 for $s = 1, \dots, S$.**

5. **Set the final parameters as the mean values over the last \tilde{S} iterations:** $\theta =$

$$\frac{1}{\tilde{S}} \sum_{s=S-\tilde{S}+1}^S \hat{\theta}^{(s)}$$

The draws from the posterior distribution of η_i are made using the Metropolis-Hastings algorithm.