

Levell, Peter; Roantree, Barra; Shaw, Jonathan

Working Paper

Mobility and the lifetime distributional impact of tax and transfer reforms

IFS Working Papers, No. W17/17

Provided in Cooperation with:

The Institute for Fiscal Studies (IFS), London

Suggested Citation: Levell, Peter; Roantree, Barra; Shaw, Jonathan (2017) : Mobility and the lifetime distributional impact of tax and transfer reforms, IFS Working Papers, No. W17/17, Institute for Fiscal Studies (IFS), London

This Version is available at:

<https://hdl.handle.net/10419/200338>

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.

Mobility and the lifetime distributional impact of tax and transfer reforms

IFS Working Paper W17/17

Peter Levell

Barra Roantree

Jonathan Shaw

Mobility and the Lifetime Distributional Impact of Tax and Transfer Reforms

Peter Levell⁺, Barra Roantree⁺, and Jonathan Shaw^{†,*}

This paper examines the distributional impact of increases to out-of-work transfers, increases to work-contingent transfers, and increases in higher rates of income tax over the whole of life. We find that, in contrast to what is implied by standard snapshot analyses, increases to work-contingent benefits are just as effective at redistributing resources to the lifetime poor as increases to out-of-work benefits. This has important implications for the equity-efficiency trade-off typically thought to apply to work-contingent transfers. However, we find that higher rates of tax on annually assessed income are an effective way of targeting the lifetime rich, as incomes are more persistent towards the top of the distribution. Our results illustrate the importance of moving beyond an exclusively snapshot perspective when analysing tax and transfer reforms.

(JEL D31, H20, H24)

Keywords: inequality; redistribution; income mobility; lifetime; tax and transfer reform

⁺ Institute for Fiscal Studies, 7 Ridgmount Street, London WC1E7AE, United Kingdom and University College London, 30 Gordon Street, London WC1H0AX, United Kingdom (Levell: peter_1@ifs.org.uk (corresponding author), Roantree: barra.roantree@ifs.org.uk).

[†] Alan Turing Institute, 96 Euston Rd, London NW1 2DB, United Kingdom and Institute for Fiscal Studies, 7 Ridgmount Street, London WC1E7AE, United Kingdom (jonshaw@turing.ac.uk).

*The authors gratefully acknowledge a grant from the Nuffield Foundation (OPD/40976) and co-funding from the European Research Council (reference ERC-2010-AdG-269440—WSCWTBDS) and the ESRC-funded Centre for the Microeconomic Analysis of Public Policy at the Institute for Fiscal Studies (IFS) (RES-544-28-5001). The Nuffield Foundation is an endowed charitable trust that aims to improve social well-being in the widest sense. It funds research and innovation in education and social policy and also works to build capacity in education, science and social science research. More information is available at <http://www.nuffieldfoundation.org>. We would also like to thank Stuart Adam, Richard Blundell, James Browne, Carl Emmerson, Rachel Griffith, Paul Johnson, Stephen Jenkins, Robert Joyce, Henrik Kleven, Guy Laroque, Cormac O’Dea, Ian Preston and advisory group participants for comments on earlier drafts of this paper. The English Longitudinal Study of Ageing (ELSA) data were made available through the UK Data Archive (UKDA). ELSA was developed by a team of researchers based at the National Centre for Social Research, University College London and the IFS. The data were collected by the National Centre for Social Research. The funding is provided by the National Institute of Aging in the United States, and a consortium of UK government departments co-ordinated by the Office for National Statistics. Data from the Family Expenditure Survey, Expenditure and Food Survey and Living Costs and Food Survey (LCFS) are Crown Copyright and reproduced with the permission of the Controller of HMSO and the Queen’s Printer for Scotland. The British Household Panel Survey (BHPS) is produced by the ESRC UK Longitudinal Studies Centre, together with the Institute for Social and Economic Research at the University of Essex. Data for the BHPS were supplied by the UK Data Service. The BHPS is also crown copyright material and reproduced with the permission of the Controller of HMSO and the Queen’s Printer for Scotland. Any errors and omissions are the responsibility of the authors.

The distributional impact of tax and transfer reforms plays a central role in policy debates. This is typically assessed by comparing the gains and losses across different income groups. Reforms are classified as progressive if they result in a greater proportional gain, or a smaller proportional loss, for households with low incomes, and regressive if the other way around.¹ Such analyses suffer from an important limitation: the incomes used to rank households and the characteristics used to assess transfer entitlements or tax liabilities are assessed using cross-sectional data covering a single point in time. However, individuals' circumstances – and therefore tax and transfer payments – can vary a great deal across life: for example, those with low current incomes may have high lifetime incomes and vice versa, meaning the distributional impact of reforms assessed using a snapshot in time may give a partial or misleading picture.

This paper examines the lifetime distributional impact of three tax and transfer reforms which have been the subject of intensive policy debate in many countries: increases to out-of-work transfers, increases to work-contingent transfers, and increases in higher rates of income tax. In each case we compare our results to those implied by a standard cross-sectional or “snapshot” analysis.

Taking a lifetime perspective, we find that increases in work-contingent benefits are just as effective as increases to out-of-work benefits at redistributing resources to the lifetime poor. This finding has important implications for the perceived equity-efficiency trade-off between the two types of transfer. Work-contingent transfers (like the Earned Income Tax Credit available to low-income families in the US) are typically thought to be less effective at reaching low-income households than out-of-work benefits, which have a more distortionary effect on labour supply (e.g. Eissa and Hoynes, 2011). Our results show that this apparent trade-off is much less stark when a lifetime perspective is taken. This finding is driven by the fact that the lifetime poor spend most of their working lives in (low-paid) work. In contrast to our findings on the effectiveness of these transfers to low income households, we also find that higher rates of tax on annually assessed income are an effective way of targeting the lifetime rich. This is because higher earners tend to exhibit less income mobility, with the result that high current incomes are indicative of high lifetime incomes. Our results illustrate the importance of moving beyond an exclusively snapshot perspective when analysing tax

¹ For recent examples see Congressional Budget Office (2015) which looks at the impact of possible changes to the Supplemental Nutritional Program, and HM Treasury (2015) which undertakes an assessment of the impact of the UK 2015 Budget.

and transfer reforms, although there are of course reasons to think that insights from a cross-sectional perspective remain important. In particular, credit market imperfections can mean that redistribution towards periods of temporarily low current incomes can be welfare-improving.

In order to assess the distributional impact of reforms from a lifetime perspective, we require longitudinal data, which not only tracks individuals over the entirety of their adult lives, but which also includes all the necessary information for computing individuals' tax liabilities and for determining their eligibility for different transfers. Unfortunately, even in the few countries where sufficiently long-running administrative or longitudinal survey data are available, to the best of our knowledge they do not include all the information needed to calculate individuals' disposable incomes under counterfactual tax and transfer systems.

Instead, we build a statistical model of earnings, employment, health, housing tenure and family composition over the life-cycle, and then use this to simulate lifetime profiles. Our approach is to first estimate transition equations for annual outcomes, conditioning on a rich set of demographics and outcomes for previous years using an 18-wave panel survey. We then use the resulting estimates to construct a simulated dataset for individuals born in the 'baby-boom' cohort born between 1945–54 that contains all the relevant characteristics for calculating tax and transfer payments. In order to ensure that our resulting simulated life-courses are representative of the experiences of our cohort of interest, we combine estimates from our relatively short panel with information from a much longer-running cross-sectional survey. Specifically, we adjust the transitions estimated using our panel data to match cross-sectional averages for the baby-boom cohort at each age. This approach allows us to minimise the risks of conflating age, period and cohort effects that would otherwise occur when generalising processes estimated across different cohorts using a relatively short panel.

In simulating life-cycle profiles for individuals, our paper relates to previous work examining the lifetime impact of tax and transfer changes, that has largely focused on indirect tax changes (e.g. Davies, St-Hilaire and Whalley, 1984; Poterba, 1989; Caspersen and Metcalf, 1994; Lyon and Schwab, 1995; Metcalf, 1999) or taxes on personal and corporate income (e.g. Fullerton and Rogers, 1991). Our work builds on this earlier literature in two main ways. First, we greatly extend the characteristics that are modelled, allowing us to examine the lifetime distributional impact of changes to out-of-work and work-contingent transfers, which have become an important policy tool in advanced economies (Blundell, 2006). Eligibility for these payments depends on time-varying factors such family composition (including partnering), health, housing tenure and hours worked in addition to

earnings, making a more comprehensive model vital. Second, our simulation approach incorporates insights from recent research on earnings dynamics: including the stylised facts that negative (positive) earnings shocks are more persistent for high (low) earners (Guvenen et al., 2015) and that the variance of shocks to income varies with age (Blundell, 2014). Incorporating these more realistic income dynamics has been shown to have important implications for understanding the variance of lifetime earnings (Altonji, Smith and Vidangos, 2013), observed patterns of wealth inequality (Castañeda, Díaz-Giménez and Ríos-Rull, 2003; De Nardi, Fella and Gonzalo Paz-Pardo, 2016), and consumption decisions (Arellano, Blundell and Bonhomme, 2015). In this paper, we account for their role in determining lifetime tax and transfer payments.

More generally, our paper contributes to a growing literature that highlights the important long-run implications of tax and transfer policies for individual welfare and behaviour: for example, Hoynes and Luttmer (2011) consider the long-run value of the tax and transfer system to individuals both through redistribution and insurance. Hoynes, Schanzenbach and Almond (2016) investigate the long-run effects of Food Stamps on health and economic outcomes in the US. Blundell et al. (2016) look at how welfare programmes subsidising the wages of low-earning individuals affect the careers of women. Finally, Dahl, Kostøl and Mogstad (2014) consider the effects of participation in welfare programmes on the participation of subsequent generations.

The remainder of this paper is structured as follows. In the next section, we describe how we construct our simulated lifetime profiles. Section II presents some descriptive statistics on lifetime incomes and how they compare to cross-sectional measures. This forms the backdrop to our results on the effects of tax and transfer reforms on lifetime and cross-sectional inequality in Section III. Section IV concludes.

I. Methodology

In this section we describe the data we use and how our lifetime processes are estimated and simulated for use in our tax policy analysis. Our aim throughout is to simulate life profiles that are representative of individuals born in the United Kingdom in the period 1945-54 (roughly corresponding to the ‘baby-boom’ cohort). Here we provide an overview of our method and how we validate it; for further details, we refer the interested reader to the online appendix and to Levell and Shaw (2016).

A. Data

We rely primarily on two datasets: the British Household Panel Survey (BHPS) and the Living Costs and Food Survey (LCFS).

The BHPS is a panel survey that ran for 18 waves from 1991 to 2008, collecting a wide range of demographic and socio-economic information. The survey followed individuals and their descendants over successive waves. The original sample comprised around 10,000 individuals in 5,500 households and was nationally representative. Booster samples were introduced for Scotland and Wales in 1999. In each wave, the survey aimed to interview all individuals aged 16+ in each household, including children who reached adulthood after the survey began and adults who moved into households that were previously surveyed.

The Living Costs and Food Survey (LCFS) is the latest name for a long-running, annual (for most of its history), cross-sectional survey of household spending patterns in the United Kingdom. It was known as the Family Expenditure Survey (FES) between 1957 and 2001 and the Expenditure and Food Survey (EFS) between 2001 and 2008. For simplicity in what follows we shall refer to all these surveys simply as the LCFS. The LCFS collects data on household incomes from various sources over the past 12 months, employment, family characteristics (including years of education from 1978 onwards) and expenditures. We make use of the LCFS between 1968 and 2012.

B. Overview of Simulation Approach

The first step in our approach is to estimate the conditional probabilities associated with different transitions at each age using the BHPS panel data. The processes we model are those that are central to determining taxes and benefits: mortality, partnering, separation, child arrival and departure, movements into and out of disability, movements in and out of employment, movements between full-time and part-time work, movements across ranks in the earnings and rent distributions, movements into and out of rented accommodation, and movements between local property tax bands (known in the United Kingdom as council tax). We do not model capital or asset incomes, or receipts of transfers and inheritances from other households. We also ignore benefits in kind, such as health and education spending. In most cases, we estimate transition probabilities using binary and multivariate logit models with a detailed set of covariates. A summary of the exact specifications we use in the estimation stage is set out in the online appendix.

Once we have obtained estimates, the next step is to use these to simulate a set of lifetime profiles. These simulations allow for correlations between outcomes in a sequential manner. First we determine whether or not the agent lives or dies in the period. We then assign births to individuals according to probabilities of child arrival that we have estimated, and determine whether children between ages 16 and 18 leave the household. Individuals then partner or separate. Childbirth is determined prior to partnering so that it will depend on lagged rather than current partner status (thus allowing for a nine month gestation period). We then determine whether or not individuals receive disability benefits, before assigning an employment status, and a location in the earnings distribution. We impose that all those who are disabled are unemployed. Finally we determine whether or not the individual is a renter, and the household's council tax band, before incrementing individuals' ages and repeating the process.

We start simulating from 1960 when the baby-boomers are in childhood. Initial conditions (education levels, likelihood of being a renter and so on) are set using data on the baby-boom cohort from the LCFS. We simulate 5,000 lifecycles. Consumption and private pension profiles are imputed to individuals once the simulations are complete. Consumption is imputed on a year-by-year basis using regressions run on the cross-sectional LCFS. For private pensions we use pension incomes projected for future years that were calculated using the English Longitudinal Study of Ageing for real world members of the baby-boom cohort (details of the methodology to construct these profiles can be found in Crawford, 2012, and an example of their use in Banks, Emmerson and Tetlow, 2014). These are calculated given observed pension wealth and for different possible retirement dates for each individual. The profiles are matched to our simulated individuals within cells defined by cohort, year, age and sex according to estimated ranks in the private pension wealth distribution.

For some of our variables, the process of simulation is quite straightforward. Probabilities are estimated and transitions then drawn from the relevant distribution. For other variables further explanation is required. In what follows we give more detail on how we model employment and earnings (for details of how we have modelled other variables, and how we impute pensions and consumption, we refer the reader to the online appendix). Before doing this however, we discuss how we ensure that the transition probabilities estimated using our panel data for the period 1991-2008 are consistent with the experiences of our cohort of interest as described by the cross-sectional data.

C. Scaling Transition Probabilities

Typically two different approaches are used to construct lifetime profiles from panels that cover a relatively short period of time. One is to estimate processes using a sample that pools together data on different cohorts (see for example Falkingham and Hills, 1995). Another is to adopt a splicing approach that joins individuals observed at different ages with the same short period together to form a complete lifecycle (Bovenberg, Hansen and Sørensen, 2008).² The downside of both of these approaches is that processes for earlier ages will be estimated using data taken from cohorts that are younger than the baby-boom cohort and those for later in life from older cohorts. This creates difficulties in interpreting the resulting estimates, since the simulated life-cycles may not correspond to the actual experiences of any real-world cohort of individuals. In addition, the approach does not make use of the information we do have on the evolution of average employment, fertility and so on across years from longer-running cross-sectional surveys.

To make best use of the data we have available, we adopt a different approach. We first estimate transition probabilities for each of the relevant variables using panel data as we described above, and then we adjust these probabilities so that the resulting cross-sectional averages for our simulated cohort match those in the cross-sectional data for the same cohort from 1960 onwards.³

To be more precise we aim to find transition probabilities for, say, employment such that we move from the (observed) cross-sectional distribution given at time $t-1$ to the cross-sectional distribution at time t . For a general binary outcome for some individual i in year t and group g , y_{it}^g , let

$$(1) \quad p_{it}^g = \text{Prob}(y_{it}^g = 1)$$

$$(2) \quad \pi_{it}^{g,01} = \text{Prob}(y_{it}^g = 1 | y_{i,t-1}^g = 0)$$

$$(3) \quad \pi_{it}^{g,11} = \text{Prob}(y_{it}^g = 1 | y_{i,t-1}^g = 1)$$

Then by the law of total probability, we get the flow equation

$$(4) \quad p_{it}^g = \pi_{it}^{g,01}(1 - p_{i,t-1}^g) + \pi_{it}^{g,11}p_{i,t-1}^g$$

It is clear that $\pi_{it}^{g,01}$ and $\pi_{it}^{g,11}$ cannot be uniquely identified with knowledge of p_{it}^g and $(1 - p_{i,t-1}^g)$ alone. This is because a variety of transition probabilities could in principle

² For a comparison of our current method with that of the splicing approach using in Bovenberg, Hansen and Sørensen (2008) see Levell and Shaw (2015).

³ The latest year we observe individuals in 2012. For years beyond this we project variables using a combination of data taken from older cohorts, estimated rates of change by age, and external forecasts.

satisfy (4). Cross-sectional data alone is therefore not sufficient to identify these and panel data is required. However, the transition probabilities that we have estimated need not satisfy this expression – not least because they are partly based on data for different cohorts and years. Let these probabilities be denoted by $\hat{\pi}_{it}^{g,01}$ and $\hat{\pi}_{it}^{g,11}$. We proceed by choosing values for $\pi_{it}^{g,01}$ and $\pi_{it}^{g,11}$ so as to minimise the average distance to our estimated transition probabilities

$$(5) \quad \left(\frac{1}{N_g} \sum_i \pi_{it}^{g,01} - \frac{1}{N_g} \sum_i \hat{\pi}_{it}^{g,01} \right)^2 + \left(\frac{1}{N_g} \sum_i \pi_{it}^{g,11} - \frac{1}{N_g} \sum_i \hat{\pi}_{it}^{g,11} \right)^2$$

subject to satisfying equation (4). Here N_g is the number of individuals in group g . Since our probabilities are estimated using logit models, we can equivalently achieve this by choosing scalars $z^{g,1}$ and $z^{g,2}$ to minimise

$$(6) \quad \left(\frac{1}{N_g} \sum_i \Lambda(X_{it}\beta^{g,01} + z^{g,1}) - \frac{1}{N_g} \sum_i \Lambda(X_{it}\beta^{g,01}) \right)^2 + \left(\frac{1}{N_g} \sum_i \Lambda(X_{it}\beta^{g,11} + z^{g,2}) - \frac{1}{N_g} \sum_i \Lambda(X_{it}\beta^{g,11}) \right)^2$$

again subject to satisfying (4). $\Lambda(\cdot)$ is the cdf of the logistic function and $\beta^{g,01}$ for example denotes the coefficients estimated for transitions from $y_{i,t-1}^g = 0$ to $y_{it}^g = 1$. Choosing terms to add to the index of individual's logistic functions rather than scaling overall probabilities ensures that all adjusted probabilities remain between zero and one.

The above scaling procedure is used to adjust transition probabilities for couple status, moving from being a renter to an owner (and vice versa) and employment. It was not found necessary to scale the probability of child arrival even though earlier cohorts tended to have more children. This is because the number of children among our simulated individuals tended to reach the right level once the probabilities for partnering and separation had been adjusted. We allow the scaling factors $z^{g,1}$ and $z^{g,2}$ we choose to differ across subgroups of the population g , allowing us for example to ensure that we match male and female employment rates implied by our repeated cross sections. The precise subgroups used differ depending on the variable we are considering and are given in Table 1.

TABLE 1 HERE

Scaling within subgroups means we will not necessarily exactly match the probability of being employed or in a couple in each year (since it is possible that the number of individuals with children differ from the numbers for the population for example). As we show in our online appendix however, we match the cross-sectional probabilities extremely well.

We scale mortality rates for individuals using a simpler method. For this we take data from the UK Office for National Statistics Life Tables which provide projected mortality rates for men and women at different ages for different birth years. We then use the difference between these and average within-sample mortality rates for individuals in the BHPS to scale mortality rates predicted using a logit regression on income, disability benefit receipt, education and couple status.

D. Modelling Employment and Earnings

We model earnings and rent *ranks* rather than attempting to model the levels of these variables. The ranks are defined as relative positions in the distribution by age and year. Once our simulations are complete we are then able to ‘fill-in’ earnings and rent levels from the actual cross-sectional distributions of earnings as observed for our particular birth cohorts in the LCFS. By using this approach, we exactly match the distributions across individuals at a point in time in terms of means, variances and higher order moments by construction.

Modelling ranks in the way we do means that we assume that movements across the earnings distribution by age are the same across cohorts and periods. This assumption is not completely innocuous. One could for instance imagine that the degree of relative mobility has changed over time between cohorts or that recessions may differentially affect the positions of some individuals relative to others. However, if we had instead modelled earnings and rents levels, we would face the problem of having to disentangle age, period and cohort effects when predicting values from our panel data. Failure to specify the correct model in this regard would potentially severely bias the shape of our estimated age profiles. In addition, we would have faced the challenge of ensuring that our simulated individuals were subject to a realistic sequence of business cycles by for example modelling a separate process for period effects. Taking values from actual cross-sectional distributions as we do ensures that our simulated individuals automatically face a real-world process for aggregate shocks.

A standard approach to modelling earnings dynamics would be to estimate an earnings process including both an individual-level fixed effect, an AR(1) error and a normally distributed innovation. However, such models are now known to have problems capturing key features of real-world earnings mobility, including aspects which are likely to be particularly important in understanding the lifetime impact of different tax and benefit reforms. For instance, these models assume that the persistence of earnings is independent of both age and individual earnings histories. Moreover, in such models, positive and negative

shocks to earnings are equally likely, and equally persistent, regardless of individuals' initial locations in the earnings distribution. However, recent work using detailed administrative data (Guvenen et al., (2015)) has shown that the distribution and persistence of shocks differs in important ways both over the life-cycle and according to current earnings.

To model movements across the earnings distribution accurately, we must be careful to allow for key features of real-world transitions. To do so, we adopt the relatively flexible approach of directly estimating transition matrices across earnings quantiles and part-time or full-time employment (for other examples see Buchinsky and Hunt, 1999; Bowlus and Robin, 2012). In particular, we proceed using the following three-step parametric approach.

1. Determine employment status: We estimate transition matrices for employment status separately for males and females and according to individuals' employment status over the previous two waves. The probabilities making up this matrix are estimated through a set of logit models which include several lags of employment status (and interactions thereof) to help us match the high persistence of employment status observed in the data. We also estimate separate logits according to the individual's employment in the previous two periods and by sex. These transition probabilities are scaled so as to match the observed unemployment rates at different ages for the baby-boom cohort in the cross-sectional LCFS data (as discussed above).

2. Place the individual in an earnings bin: Once an individual's employment status is determined, we then place the individual in one of five possible bins: in part-time work, or in full-time work and in one of four different earnings quartiles. Distinguishing between part- and full-time work is important in our case as it determines eligibility for receipt of tax credits in the United Kingdom. We assume that part-time work corresponds to 20 hours per week and full-time work to 40 hours. To determine which bin an individual should be placed in we estimate multinomial logits from each of the six possible prior states i (which include unemployment)

$$(7) \quad \Pr(i, j | X_{it}) = \frac{\exp(X_{it}\beta^{ij})}{\sum_{m=0}^N \exp(X_{it}\beta^{im})}$$

X_{it} is a set of covariates which includes a cubic in age, education, a dummy for whether individuals have children or not, and a dummy for whether they have children under the age of five (and various interactions of these) their current earnings rank (entering linearly) as well as five lags of full-time and employment status and lagged earnings quartiles. Including

these lags relaxes what would otherwise be a Markov assumption that next period's transition depends only on current circumstances. Differences in the coefficients attached to these allow for differences in the persistence of earnings ranks across the distribution. We run logistic regressions separately for each initial bin, and separately for men and women.⁴

3. Determine the individual's precise earnings rank: The results from these models can be used to estimate the probability of moving between unemployment, part-time and full-time work and the different income quartiles. However, it does not place individuals precisely within these quartiles. One approach is to deal with this is to match simulated individuals to real-world individuals who made the same employment and income quintile transitions as they did and use these individuals' new ranks to determine the simulated individual's new locations (the approach adopted by Bowlus and Robin, 2012). We adopt an alternative parametric method. This involves predicting ranks using regressions of the following form (8)

$$\Phi^{-1}(r_{ikt}) = \sum_{\tau=1}^4 \sum_j \delta_0^{\tau j} D_{Q_{t-\tau}=j}^i + \sum_{\tau=1}^4 \sum_j \delta_1^{\tau j} D_{Q_{t-\tau}=j}^i \times r_{ij,t-1} + \sum_{\tau=1}^4 \sum_j \delta_2^{\tau j} D_{Q_{t-\tau}=j}^i \times r_{ij,t-1}^2 \dots$$

where r_{ikt} is the within-bin rank of individual i in period t with bin k . $D_{Q_{t-\tau}=j}^i$ is a dummy which equals one if the individual was located in bin j in period $t - \tau$. $\Phi^{-1}(\cdot)$ is the inverse of the CDF of the normal distribution. These regressions are run separately for each destination bin, allowing us to capture the asymmetric nature of persistence over the income distribution. They are also run separately for males and females. Linearly predicting $\Phi^{-1}(r_{ikt})$ (and then feeding this prediction through $\Phi(\cdot)$) ensures that the predicted within-bin rank always lies between 0 and 1. The polynomial of past ranks (up to a cubic) included in this regression is also interacted with a cubic in age in order to help us match the differing persistence of earnings over the life-cycle. In our simulations, we add a normally distributed noise term with the variance of residuals seen in the data to the linear prediction made using (8). The raw ranks that we predict for our simulated individuals need not have a uniform distribution. If this were not corrected before imputing earnings, the quantiles of earnings for our simulated individuals would differ from those in our cross-sectional data. To make the distribution uniform, we assign new ranks to individuals based on their relative positions

⁴ For those under 22, we predict bins and employment using a simpler model. See the online appendix for details.

within-sample as a final step. We see this additional step as analogous to our other cohort scaling adjustments.

Younger individuals (who do not have all the lags needed to be included in the regression models mentioned so far) have their status determined by a multinomial logit (across our five possible bins and unemployment) and a within-bin regression involving only one lag.

Once earnings ranks have been predicted, we then impute actual earnings levels using the distributions observed in the LCFS. For years when our cohort is not observed in the LCFS (prior to 1968 and after 2012) we uprate or downrate earnings distributions for the relevant ages from different cohorts using historical data on earnings growth or forecasts taken from the UK Office for Budget Responsibility.

Our approach allows for the persistence of earnings to differ by age, and for transition probabilities to vary depending on individual's earnings histories. To assess how well we capture those patterns of mobility that we observe in the data, panels (a) and (b) in Figure 1 plot the joint distribution of earnings ranks for workers in the current period and one period ahead in the BHPS data and for our simulated individuals. In Figure 2, we plot the joint distributions for individual locations in the current period and five years ahead. In both of these graphs, the x and y axes show earnings vingtiles while the z axis shows relative frequencies. The left-hand panels show that, in the data, individuals are highly likely to remain in their current vingtile, with persistence increasing towards the tails of the distribution. Persistence is however greater for those at the top of the earnings distribution than for those at bottom. The right-hand panel shows that we replicate these patterns in our simulated data well, though we may understate the overall persistence in individual incomes.

FIGURE 1 HERE

FIGURE 2 HERE

The fit of our model is a considerable improvement over a more standard ARIMA earnings process. In panel (c) of Figures 1 and 2, we plot the joint distributions for simulated results based on an alternative a fixed-effects earnings model with a normally distributed AR(1) error (we discuss the details of the estimation of this model in Appendix B). Overall, the fit of the ARIMA alternative is relatively poor. The model underestimates the probability that individuals will remain in the current income vingtile. Moreover, as this model does not allow for different dynamics in the tails of the earnings distribution, it also greatly overstates mobility for the highest and lowest earners.

A second test of our model concerns its ability evolution of earnings mobility over the life-cycle. Figures 3 plots autocorrelations for earnings ranks separated by one year, five years and ten years for males and females aged 16-65. It is clear that in the data these autocorrelations differ over different stages of life. Specifically, autocorrelations tend to be lower at younger and older ages. Our simulated data captures these patterns reasonably well. While the fit for one and five year horizons is good however, ranks in our simulated earnings distribution are less persistent at middle age for the longer 10-year horizon than earnings ranks in the BHPS, and a little more persistent at older ages.

Figure 4 plots the equivalent statistics using the alternative fixed-effects earnings model. In contrast, to our preferred approach, we find that simulations based on the fixed-effects earnings model tend to understate persistence over short horizons, while overstating it for the longer periods. This reflects the role that the fixed-effects parameter plays in accounting for the persistence in earnings, swamping other influences for longer time horizons. Indeed there is very little difference in estimated autocorrelations for the five and ten year horizons. Moreover, unlike the data, this model predicts that earnings persistence is constant over the life-cycle.

FIGURE 3 HERE

FIGURE 4 HERE

In the online appendix, we compare the persistence of employment in our preferred simulations with our cohort of interest in the BHPS over a 10 year horizon. The proportions always employed and never employed within this period are very similar, lending support to our approach for modelling employment transitions. Overall, we conclude our lifetime earnings process captures several key features of mobility and shows superior performance to more traditional approaches.

Autocorrelations across different ages and different time horizons for other variables we simulate fit the patterns in the BHPS well. These are presented in the online appendix.

E. Taxes and Benefits

The rich set of outcomes we model gives us all the key variables we need to calculate tax and benefits for each of our simulated individuals. We assign tax and benefit payments to our simulated individuals using TAXBEN, a detailed tax and benefit microsimulation model for the United Kingdom. We model all major benefit payments – both universal and means-tested – including unemployment insurance, payments to families with children, disability

benefits and tax credits.⁵ We assume that all those eligible for benefits take them up. The taxes included are income tax, payroll taxes (“employee National Insurance contributions”), VAT, excise and fuel duties, and local property taxes (“council tax”). We do not model the effects of capital taxes, inheritance duties or business taxes which are difficult to attribute to individual households (such as corporation tax). In addition, our approach does not account for any differences in behaviour that might result from our individuals being exposed to the tax systems we impose on them rather than the actual tax and benefit systems they faced.

F. Behavioural Responses

The approach we use does not incorporate possible behavioural responses to tax and transfer payments. This is for two reasons. First, the envelope theorem implies that, for optimising individuals, the welfare effect of small reforms like the ones we consider is captured by the increase or decrease in lifetime income absent behavioural response. Any additional changes in income that are due to behavioural responses will (in the limit) be associated with zero net welfare gains or losses and so can be disregarded for the relatively small changes we consider. Second, it allows us to model a much more complete set of characteristics that are relevant for calculating tax and transfer payments, such as detailed family composition, joint employment and earnings in couples, health status, housing tenure and house value. Typically, a number of these variables are omitted from dynamic models that incorporate behavioural responses (usually in labour supply or consumption) in order to maintain tractability.

II. Lifetime Incomes

To provide some background for interpreting our subsequent results, this section presents descriptive statistics for our simulated lifetime profiles and how these compare to statistics from a synthetic 2015/16 cross-section that is also based on our simulated lifecycles. This cross-section describes what the 2015/16 population would look like if all cohorts were the same as the baby-boom cohort. As a result, any differences relative to the lifetime will be due to the lifetime perspective. An alternative would have been to use a real-world cross-sectional distribution, but this would leave the reader in doubt as to whether our results were driven by the cross-sectional perspective or sample differences in the cohorts being considered. In

⁵ More precisely, the benefits included are income support, jobseeker’s allowance, housing benefit, council tax benefit, child benefit, family credit, tax credits (working families’ tax credit, working tax credit and child tax credit), and the state pension.

Levell, Roantree and Shaw (2015), we examine how close our synthetic cross-section is to the 1978 and 2012 LCFS cross-sections. Results using the different approaches tell a very similar story.

The course of incomes over the lifecycle is shown in Figure 5. This plots mean annual earnings (not conditional on working) and mean private pensions (not conditional on receiving a pension) across life for men and women for our simulated profiles. We do not plot these against survey data, as average earnings match those in the LCFS by construction and private pension information is taken directly from our cohort of interest. The figure shows that, on average, earnings rise for men until the late 40s and then decline steadily thereafter. For women, earnings flatten off during the late 20s, associated with taking time out of the labour market for child-rearing. They then rise again, reaching a peak at around age 50 before falling again towards retirement. As earnings decline around retirement, individuals start receiving private pensions, though at a much lower level, on average, than earnings during working life.

FIGURE 5 HERE

In order to compare individuals with different lifetime incomes, we must first specify a way to convert income profiles across life into a single figure. One way to do this would be to set lifetime income equal to the present discounted sum of incomes across all ages. However, this definition could mean that those living long lives with low living standards in each period will appear to have higher incomes than those who living short lives but with higher average incomes. Such outcomes are counterintuitive. Instead, we use a measure of the discounted average lifetime income

$$(9) \quad Y_i = \frac{1}{\sum_{a=16}^{A_i} ((1+r_a)^{a-16})} \sum_{a=16}^{A_i} \frac{y_a}{(1+r_a)^{a-16}}$$

where a is age, r_a is the nominal interest rate (taken from the yields on UK government consols) and A_i is the maximum age reached by individual i . This converts lifetime income into an annualised figure. Incomes are equivalised using the modified OECD scale in each year prior to doing this. The discounted value of lifetime income, Y_i , is then inflated using the Retail Price Index (RPI) to 2015 terms to ensure comparability across cohorts.

Figure 6 gives an indication of how persistent net income is over time in our simulations at different points in the income distribution. This shows that those lifetime poor (rich) spend quite a substantial fraction of life outside the poorest (richest) cross-sectional decile; in other words, there is a substantial degree of mobility around the distribution across life. It also shows that our simulations capture the non-symmetric nature of mobility. Persistence is

greater at the top of the distribution than at the bottom: individuals in the richest lifetime decile spend more of their lives in the richest cross-sectional decile than individuals in the poorest lifetime decile spend in the poorest cross-sectional decile (34% compared to 21%). A similar result is obtained when we consider the BHPS data itself (see Figure 2.8 in Roantree and Shaw, 2014). Analysis using administrative data also indicates a similar difference in persistence across income groups, with the those in the top earnings quintile most likely to remain in their current relative position (Kopczuk, Saez and Song, 2007).

FIGURE 6 HERE

Figure 7 shows how employment varies across the net income distribution. The cross-section series shows the proportion of working-age individuals who are employed, split by cross-sectional net income decile. The lifetime series shows the average fraction of working life that individuals are employed, split by annualised lifetime net income decile. From this graph, it is clear that relatively few individuals in the bottom cross-sectional decile are employed (22%), but from a lifetime perspective, individuals in the bottom lifetime decile are employed for the majority of working life (an average of 63%). This has implications for the relative impact of in-work and out-of work benefits on lifetime inequality, which we discuss below.

FIGURE 7 HERE

Finally, Figure 8 plots the distributions for gross lifetime and cross-sectional incomes. The cross-sectional distribution exhibits positive skew, with a long tail of individuals with high incomes, while the lifetime distribution is more symmetric. This reflects the impact of income mobility on the lifetime distribution – a point we also return to in what follows.

FIGURE 8 HERE

III. Results

In this section we present results for the effect of the UK tax system on inequality and consider the impact of various reforms that have either been argued for on efficiency grounds or else play in an important role in current policy debates around the world.

The analysis below assumes that – barring the reforms we consider – individuals face the 2015/16 UK tax and benefit system for the entirety of their lives. This is because we are primarily interested in the characteristics of given tax and benefit systems from a lifetime perspective rather than, say, the experiences of a particular cohort under the systems they were actually exposed to. We hold behaviour fixed following reforms, meaning that labour

supply and other variables are assumed to be the same as they were for the baby-boom cohort under all the different tax systems we consider. To determine individual incomes and tax and benefit payments we assume equal sharing of resources between members of couples.

In order to apply a given tax and benefit system to data from earlier or later years, we uprate in line with earnings. This brings us close to ensuring that the tax and benefit system raises the same revenue each year. An alternative would be to uprate in line with prices. Results from this alternative scheme tell a similar story (see Levell, Roantree and Shaw, 2015).

B. The Effects of the Tax and Benefit System on Inequality

Before discussing the impact of particular reforms, we start by showing how the 2015/16 system itself affects lifetime and cross-sectional inequality.

Table 2 shows gross and net income Gini coefficients calculated for the synthetic cross-section and on a lifetime basis for the baby-boom cohort. The first thing to notice is how much lower inequality is over the lifetime: the cross-section Gini coefficient for gross income is 0.493 compared with 0.258 across the whole of adult life. This indicates that a lot of the income inequality before taxes and benefits between individuals is temporary, either reflecting the stage of life they are at (such as differences in work experience and family structure) or reflecting transitory shocks individuals experience (such as unemployment). Our figure for lifetime income inequality is roughly in line with related figures from other studies. Using German administrative data, Bönke, Corneo and Lüthen (2015) calculate a lifetime earnings inequality measure. They find that for the cohort of individuals born in 1949 it is 0.212 (calculated up to age 60).⁶

TABLE 2 HERE

The second thing to notice is that the tax and benefit system is effective at reducing inequality, but more so in the cross-section: when including the effect of indirect taxes, the cross-section Gini falls from 0.493 to 0.337, a reduction of 0.155 (or 31.4%) while the lifetime Gini falls from 0.258 to 0.195, a reduction of 0.046 (or 17.8%, which is just over half the corresponding cross-sectional fall).

⁶ The fact that our figure is slightly greater may partly be due to there being greater inequality in incomes than earnings. Bönke et al. find that inequality of long-run earnings are around two thirds of those for cross-sectional earnings. Our lifetime measure of inequality is less than half. Bönke et al. are however comparing within-cohort lifetime and cross-sectional inequality measures at different ages. Their measure of cross-sectional inequality is correspondingly lower.

Why should this be the case? The explanation is that much of the redistribution undertaken by the tax and benefit system is *intrapersonal*: redistribution between periods of life that nets off over the whole lifecycle. This sort of redistribution will be effective at reducing cross-sectional inequality but not lifetime inequality.

The proportion of redistribution that are intra or interpersonal can be independently quantified using a decomposition similar to Bovenberg, Hansen and Sørensen (2008). Let redistribution towards individual i at age a be given by

$$(10) \quad R_{i,a} = B_{i,a} - T_{i,a} - K_{i,a}$$

where $T_{i,a}$ is the taxes paid by individual i at age a , $B_{i,a}$ are the benefits received and $K_{i,a}$ is a ‘no redistribution’ baseline, all defined in PV 2015 terms. The baseline ensures that total redistribution across all individuals sums to zero even if the tax system raises net revenue overall. We consider two different definitions of a non-redistributive tax system to calculate $K_{i,a}$. Under the first, every individual’s contribution is a constant amount in each period (a ‘lump-sum baseline’). The second is where each individual pays a constant proportion of gross income in each period (a ‘proportional baseline’). Redistribution towards or away from individuals is said to occur when net taxes paid are greater than or smaller than these baselines. Using these definitions it is possible to show (see Levell, Roantree and Shaw (2015)) that total redistribution can be decomposed into intra- and interpersonal components using the following relationship

$$(11) \quad \underbrace{\sum_i \sum_a |R_{ia}|}_{\text{Total redistribution}} = \underbrace{\sum_i |R_i|}_{\text{Total interpersonal}} + \underbrace{2 \sum_i \min \{ \sum_a [1(R_{ia} > 0) R_{ia}], - \sum_a [1(R_{ia} \leq 0) R_{ia}] \}}_{\text{Total intrapersonal}}.$$

where R_i is the sum of redistribution for individual i across life and $|\cdot|$ is the absolute value operator. When we calculate average proportions of intra- and interpersonal redistribution, we find that a majority of redistribution is intrapersonal i.e. across periods of life rather than across individuals. For the 1950, cohort these figures are 58.8% under the lump-sum baseline and 61.7% under a proportional baseline.⁷ The reason for this is that most taxes and transfers are assessed over periods of a year or less, making it much easier to target outcomes over short horizons.

We now turn to using our simulations to answer questions on the lifetime impact of changing particular elements of the tax and benefit system, and how this might matter for more general design issues.

⁷ Results from other birth cohorts are similar.

C. How Effective are Increases in In-work and Out-of-work Benefits at Targeting the Lifetime Poor?

Several countries provide significant income top-ups to low-income in-work families through tax credits. For example in the United States there is the Earned Income Tax Credit (EITC). In the United Kingdom these take the form of the Working Tax Credit (supplemented by Child Tax Credit for in and out-of-work families with children). Governments also attempt to help low-income households by removing low earning individuals from the direct tax net. Such reforms have the potential to provide help to low-income households while at the same time improving work incentives. However, in a cross-sectional analysis, they often appear less progressive than increases in out of work benefits, creating an equity-efficiency trade-off similar to that described for the EITC in Eissa and Hoynes (2011). How does a lifetime analysis affect this assessment?

To answer this question we compare reforms of these kinds in a cross-section and over the lifetime. We consider are three measures with similar cross-sectional revenue consequences (around £3 billion/\$4.2 billion per year), namely:

- (i) an increase in out-of-work benefits: a 16.5% increase in the maximum income support, (income-based) jobseeker's allowance, pension credit and (non-contributory) employment support allowance awards;
- (ii) an increase in in-work (i.e. work-contingent) benefits: an 18% increase in the maximum working tax credit award;
- (iii) an income tax cut: a 4% increase in the income tax personal allowance. This is the threshold below which income tax is not paid. In 2015/16 this stood at £10,600 (around \$16,000).

Reforms are applied using the 2015/16 tax and benefit system as a base.

Figure 9 shows the cross-sectional effect of the three reforms listed above. Unsurprisingly, the most progressive reform is the increase to out-of-work benefits: gains are concentrated in the bottom two income deciles, with the largest average gain experienced by the lowest income decile (5.0%), in which there is a high share of non-working individuals. Next most progressive is the increase to working tax credit. Here, the bottom four deciles are gainers, with gains peaking at an average of 1.7% in decile two. The bottom decile gains by less because fewer individuals here are entitled to working tax credit (because they are not in work). The least progressive (indeed a regressive) reform is the income tax cut in the form of

an increase in the personal allowance. Gains are concentrated among the upper half of the income distribution, reflecting the fact that the poorest adults have income that is too low to benefit from the giveaway, while dual-income couples – who tend to have higher family incomes – can benefit twice over.

FIGURE 9 HERE

The pattern of gains over lifetime income is different, as shown in Figure 10. In particular, increases in out-of-work and in-work benefits are strongly progressive, with very similar distributional patterns, while the income tax cut (personal allowance increase) is close to distributionally neutral.

What is most interesting here is the fact that increases in out-of-work and in-work benefits have such a similar distributional pattern. This result stems from the pattern of worklessness over the lifecycle shown in Figure 7. While the poorest individuals in the cross-section are often out of work, this is often a temporary state and many of the poorest individuals in a lifetime sense move in and out of work. In addition, when in work, they are relatively likely to be in low-paid work and therefore qualify for in-work support. This makes it possible to reach many of the lifetime poor through either out-of-work or in-work benefits.

FIGURE 10 HERE

The lifetime analysis is therefore relatively more favourable to a policy of increasing in-work benefits. The advantage of these sorts of payments is that – in general – they have much less of a negative impact on work incentives at the extensive margin (where labour supply responses may be particularly large; see Eissa and Liebman (1996) for a discussion in the context of the EITC). Out-of-work benefits reduce the net-financial gain to being in work, while in-work benefits in the United Kingdom are contingent on working a certain number of hours. Furthermore, as Blundell et al. (2016) highlight, the two types of benefits have different effects on incentives to accumulate human capital: while increasing out-of-work benefits provides a high level of insurance, they are associated with strong moral hazard effects and are less effective in improving overall welfare than in-work benefits.

Policymakers looking to target the lifetime poor might therefore favour doing so through in-work benefits. The disadvantage of such an approach is that it would do less to help the lifetime poor in the particular periods that they were not working, which could matter if they did not have access to savings or borrowing facilities. It would also do less to help the minority of the lifetime poor who do experience sustained periods without work. That said, recent experience suggests that the lifetime poor among younger cohorts may increasingly

have substantial amounts of work over their lives but low levels of earnings rather than long periods out of the labour market (for cross-sectional evidence, see Belfield et al., 2014).

D. Income Taxes and the Lifetime Rich

So far we have tended to find that measures that would otherwise be thought of as progressive: appearing less redistributive once a longer-term view is taken. In light of this, we ask how well progressive income taxes target the lifetime rich. We consider increases in both the ‘higher rate’ of income tax which in 2015/16 applied at a rate of 40% to individual incomes above £31,786 (roughly \$48,000) and the ‘basic rate’ of 20% applied to incomes between this and the personal allowance of £10,600 (around \$16,000).⁸

Figure 11 shows the distributional impact of a one percentage point increase in the higher rate of income tax, from both a cross-sectional and lifetime perspective. The reform is extremely progressive in the cross-section: the bottom four deciles are completely unaffected (because these individuals do not earn enough to pay the higher rate), and it is only the top two deciles that experience a hit to incomes of more than 0.1%, with losses peaking at 0.38% for the top decile. Over the lifetime, the reform remains strongly progressive, but there is slightly more of an impact further down the distribution. Those in the top four deciles experience a loss exceeding 0.1% but, as before, it is the very top decile that stands out, with a 0.23% fall. This reflects the greater persistence in earnings at the top of the distribution in our simulations (as we showed in Figure 5). Thus, changes to the higher rate of income tax are reasonably effective at targeting the lifetime rich.

For comparison, we also present the distributional impact of a one percentage point increase in the basic rate of income tax (Figure 12). This shows that the cross-sectional impact is progressive: the bottom decile is largely unaffected (because most individuals in this decile do not earn enough to pay income tax) and the average loss peaks at 0.60% for the ninth decile. The top decile loses by slightly less because a smaller share of income for these individuals is subject to the basic rate. Over the lifetime, the impact remains progressive, but much less so because many more individuals will pay basic rate income tax at some point in life than in any one year. The bottom decile suffers a loss of 0.35%, rising to 0.53% for the top decile.

⁸The United Kingdom also has an additional rate of 45% which applies to incomes greater than £150,000 (\$220,000). We would ideally also like to consider this but the survey data used to construct our simulations does not capture incomes well at the very top of the distribution and so we are unlikely to be able to model this accurately.

FIGURE 11 HERE

FIGURE 12 HERE

IV. Conclusion

Mobility in employment, earnings, housing tenure, health status and family composition means that in the longer-run, the distributional impact of tax and transfer reforms can look very different to those implied by snapshot analyses. This paper examines the lifetime distributional impact of three tax and transfer reforms which have been the subject of intensive policy debate in many countries.

In summary, we find that increases in work-contingent benefits are just as effective as increases to out-of-work benefits at redistributing resources to the lifetime poor, with important implications for the perceived equity-efficiency trade-off between the two types of transfer. We also find, by contrast, that higher rates of tax on annually assessed income are an effective way of targeting the lifetime rich. These results are both driven by differing patterns of mobility over the earnings distribution. Those at the bottom of the lifetime income distribution tend to see greater year to year variation in their earnings, as they move in and out of (low-paid work), whereas the earnings of those at the top of the distribution see greater persistence.

Our results illustrate the importance of moving beyond an exclusively snapshot perspective when analysing tax and transfer reforms. However, there remain good reasons to also consider cross-sectional impacts: in particular, the presence of credit market imperfections or significant uncertainty which inhibit consumption smoothing means that redistribution towards periods of temporarily low current incomes can be particularly valuable. Quantifying the long run insurance- and consumption-smoothing value that different forms of redistribution provide would be an important avenue for future research.

References

- Altonji, Joseph., Anthony Smith, and Ivan Vidangos, 2013. “Modeling Earnings Dynamics.” *Econometrica* 81, 1395–1454.
- Arellano, Manuel., Richard Blundell and Stéphane Bonhomme, 2015. “Earnings and consumption dynamics: a nonlinear panel data framework.” IZA Discussion Paper No. 9344.
- Atkinson, Anthony Joseph Stiglitz, 1976. “The design of tax structure: direct versus indirect taxation.” *Journal of Public Economics* 6, 55-75.

Banks, James, Carl Emmerson, and Gemma Tetlow, 2014. “Effect of pensions and disability benefits on retirement in the UK.” NBER Working Paper No. 19907.

Barrett, Garry, Kevin Milligan, and Peter Levell, 2015. “A comparison of micro and macro expenditure measures across countries.” In Chris Carroll, Thomas Crossley, and John Sabelhaus *Improving the measurement of consumer expenditures*, Studies in Income and Wealth 74, National Bureau of Economic Research, University of Chicago Press, Chicago.

Belfield, Christopher, Jonathan Cribb, and Andrew Hood, Robert Joyce, 2014. “Living Standards, Poverty and Inequality in the UK: 2014.” Report R96, Institute for Fiscal Studies, London.

Blundell, Richard, 2014. “Income dynamics and life-cycle inequality mechanisms and controversies” *Economic Journal* 124, 289-318.

Blundell, Richard, 2006. “Earned Income Tax Credit Policies: Impact and Optimality, The Adam Smith Lecture 2005”, *Labour Economics* 13, 423-443.

Blundell, Richard, Monica Costa-Dias, Costas Meghir and Jonathan Shaw, 2016. “Female Labour Supply, Human Capital and Welfare Reform.” *Econometrica* 84, 1705-1753.

Bönke, Timm, Giacomo Corneo, and Holger Lüthen, 2015. “Lifetime earnings inequality in Germany”, *Journal of Labor Economics* 33, 171-208.

Bovenberg, A. Lans, Martin I. Hansen, and Peter B. Sørensen, 2008. “Individual savings accounts for social insurance: rationale and alternative designs.” *International Tax and Public Finance* 15, 67–86.

Bowlus, Audra J. and Jean-Marc Robin, 2012. “An international comparison of lifetime inequality: how continental Europe resembles North America.” *Journal of the European Economics Association* 10, 1236-1262.

Brewer, Mike and Cormac O'Dea, 2012. “Measuring living standards with income and consumption: evidence from the UK.” IFS Working Papers W12/12, Institute for Fiscal Studies.

Buchinsky, Moshe and Jennifer Hunt, 1999. “Wage mobility in the United States.” *Review of Economics and Statistics* 81, 351-368.

Caspersen, Erik and Gilbert Metcalf, 1994. “Is a Value-Added tax regressive? Annual versus lifetime incidence measures.” *National Tax Journal* 47, 731-746.

Castañeda, Ana, Javier Díaz-Giménez, and José-Víctor Ríos-Rull, 2003. “Accounting for the U.S. Earnings and Wealth Inequality.” *Journal of Political Economy* 111, 818– 857.

Congressional Budget Office, 2015. “The Effects of Potential Cuts in SNAP Spending on Households With Different Amounts of Income.” Congressional Budget Office Report No. 49978 (<https://www.cbo.gov/publication/49978>)

Crawford, Rowena, 2012. “ELSA pension wealth derived variables (waves 2 to 5): Methodology.” UK Data Archive Study Number 5050 - English Longitudinal Study of Ageing.

Davies, David G., 1960. “Progressiveness of a Sales Tax in Relation to Various Income Bases.” *American Economic Review* 50, 987-995.

Davies, James, France St-Hilaire, and John Walley, 1984. “Some calculations of lifetime tax incidence.” *American Economic Review* 74, 633-649.

Dahl, Gordon, Andreas Kostøl and Magne Mogstad, 2014. “Family welfare cultures.” *Quarterly Journal of Economics* 129, 1711-1752.

Disney, Richard, Carl Emmerson, and Gemma Tetlow, 2007. “What is a public sector pension worth.” Institute for Fiscal Studies (IFS) Working Paper no. W07/17. (<http://www.ifs.org.uk/publications/4051>).

Di Nardi, Mariacristina, Giulio Fella and Gonzalo Paz-Pardo, 2016. “The Implications of Richer Earnings Dynamics for Consumption, Wealth, and Welfare.” NBER working paper no. 21917

Eissa, Nada and Hilary Hoynes, 2011. “Redistribution and tax expenditures: the earned income tax credit.” *National Tax Journal* 64, 689-730.

Eissa, Nada, and Jeffrey Liebman, 1996. “Labor Supply Response to the Earned Income Tax Credit,” *Quarterly Journal of Economics* 111, 605– 637.

Falkingham, Jane and John Hills, J. (eds), 1995. *The Dynamic of Welfare: the welfare state and the life cycle*, Hemel Hempstead: Prentice Hall.

Fullerton, Don and Diane L. Rogers, 1991. “Lifetime versus annual perspectives on tax incidence.” *National Tax Journal* 44, 277-287.

Guvenen, Fatih, Fatih Karhan, Ozkan Serder and Jae Song, 2015. “What do data on millions of U.S. workers reveal about the life-cycle earnings risk?” NBER Working Paper No. 20913, National Bureau of Economic Research, (<http://www.nber.org/papers/w20913>)

HM Treasury, *Budget 2015*, Controller of Her Majesty’s Stationery Office, London, 2015.

Hoynes, Hilary, and Erzo Luttmer, 2011. “The Insurance Value of State Tax-and-transfer Programs”, *Journal of Public Economics* 95, 1466-1484.

Hoynes, Hilary, Diane Whitmore Schanzenbach and Douglas Almond, 2016. “Long Run Economic and Health Impacts of Participation in the Food Stamp Program.” *American Economic Review* 106, 903-934.

Kaplow, Louis, 2006. “On the undesirability of commodity taxation even when income taxation is not optimal.” *Journal of Public Economics* 90, 1235-1250.

Kay, John, 2010. “The Base for Direct Taxation.” In Mirrlees, James, Stuart Adam, Timothy Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, James Poterba (eds). *Dimensions of Tax Design: The Mirrlees Review*. Oxford University Press, Oxford.

Kopczuk, Wojciech, Emmanuel Saez and Jae Song, 2007. “Uncovering the American Dream: Inequality and Mobility in Social Security Earnings Data since 1937.” National Bureau of Economic Research Working Paper 13345.

Levell, Peter and Jonathan Shaw, 2016. “Constructing full adult life-cycles from short panels.” *International Journal of Microsimulation* 9(2), 5-40.

Levell, Peter, Barra Roantree and Jonathan Shaw, 2015. “Redistribution from a lifetime perspective.” Institute for Fiscal Studies (IFS), Working Paper W15/27 (<http://www.ifs.org.uk/publications/7986>).

Lyon, Andrew and Robert Schwab, 1995. “Consumption taxes in a life-cycle framework: are sin taxes regressive?” *Review of Economics and Statistics* 3, 389-405.

Metcalf, Gilbert, 1999. “A distributional analysis of green tax reforms.” *National Tax Journal* 52, 655-681.

Mirrlees, James, Stuart Adam, Timothy Besley, Richard Blundell, Stephen Bond, Robert Chote, Malcolm Gammie, Paul Johnson, Gareth Myles, James Poterba (eds), 2011. *Tax By Design: The Mirrlees Review*. Oxford University Press, Oxford.

Poterba, James, 1989. “Lifetime Incidence and the Distributional Burden of Excise Taxes.” *American Economic Review* 79, 325-39.

Office for National Statistics (ONS), 2015. “The Effects of Taxes and Benefits on Household Income, Financial Year Ending 2014.” Office for National Statistics, available at <http://www.ons.gov.uk/ons/rel/household-income/the-effects-of-taxes-and-benefits-on-household-income/2013-2014/etb-stats-bulletin-2013-14.html>.

Roantree, Barra and Jonathan Shaw, 2014. “The Case for Taking a Life-Cycle Perspective: Inequality, Redistribution, and Tax and Benefit Reforms.” Report R92, Institute for Fiscal Studies, London.

TABLES

Table 1: Cells within which Transition Probabilities are Scaled

<i>State</i>	<i>Cells</i>
Couple	Age, year, sex, number of children
Renter	Age, year
Employed	Age, year, sex, has children

Table 2: Gross and Net Income Gini Coefficients

<i>Horizon</i>	<i>Gross income</i>	<i>Net income</i>	<i>Net income less indirect taxes</i>
Cross-section	0.493	0.298	0.337
Lifetime	0.258	0.195	0.212

Note: Taxes and benefits are calculated on an annual basis and are equivalised using the Modified OECD equivalence scale. The 'Net income' column excludes the effect of indirect taxes, while the 'Net income less indirect taxes' column subtracts them. Individuals face the 2015/16 tax and benefit system throughout life uprated in line with average earnings (AEI).

FIGURES

Figure 1: Joint Distribution of Earnings: Year t and Year $t+1$

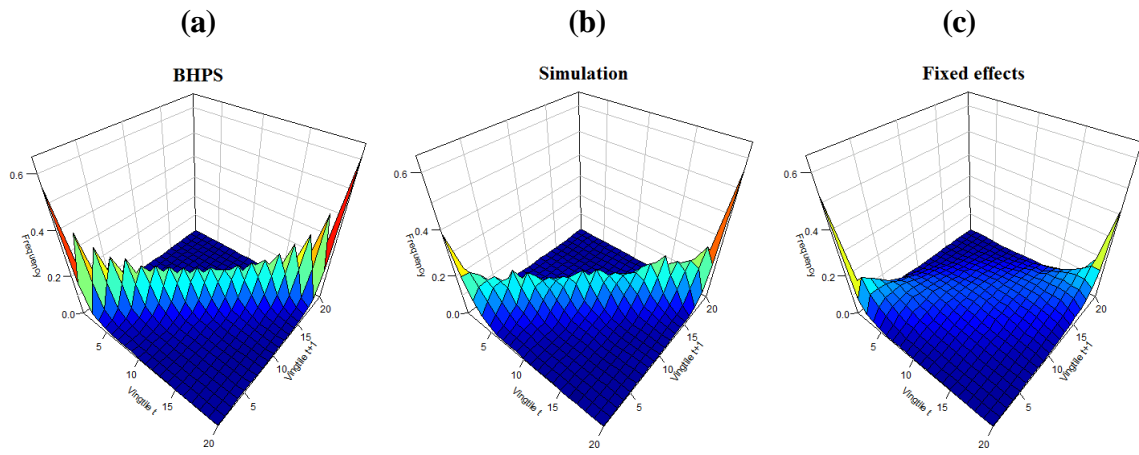


Figure 2: Joint Distribution of Earnings: Year t and Year $t+5$

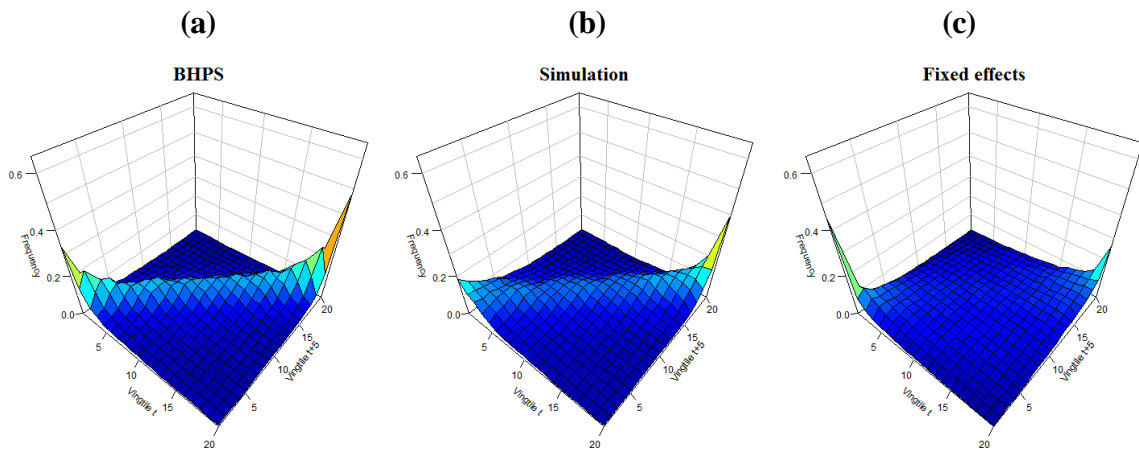


Figure 3: Autocorrelations Earnings Ranks: Simulations versus Data




Figure 4: Autocorrelations Earnings Ranks: Simulations versus Data (Fixed Effects Earnings Process)






Figure 5: Mean Earnings and Pensions by Age and Sex




Note: Series show mean earnings and private pensions across life. Values are expressed in real 2015 terms (deflated by the Retail Prices Index, or RPI). Earnings are zero for the unemployed and for those not in receipt of a private pension.

Figure 6: Proportion of Life Spent in each Cross-Sectional Decile, by Lifetime Net Income Decile




Note: The series show the proportion of life spent in each cross-section net income decile by individuals in the poorest/richest lifetime net income decile. Deciles are defined on equivalised net income ignoring indirect taxes (annualised net income for lifetime deciles).

Figure 7: Employment among Working Age Individuals by Net Income Decile




Note: The cross-section series shows the fraction of working-age individuals in each net income decile who are employed (deciles are defined using the whole population, not just working-age individuals). The lifetime series shows the average fraction of working life that individuals are employed for. Working age is defined as under 63 for women and under 65 for men. Deciles are defined on equivalised net income ignoring indirect taxes (annualised net income for lifetime deciles).

Figure 8: Gross Income Distributions




Note: The series show the densities of gross equivalised household incomes over the lifetime and in a cross-section. Lifetime incomes are expressed in annualised terms and are discounted to the year when individuals turned 16 (see equation (9)). We exclude the top 1% of incomes and those with zero incomes.

Figure 9: Cross-sectional Distributional Impact of Increases to Out-of-work and In-work Benefits and of the Income Tax Personal Allowance




Note: Deciles are defined on the basis of cross-sectional equivalised net household income. The height of the bars is the gain or loss as a percentage of the relevant decile's total net (unequalised) household income. The 'Out-of-work benefits' series shows the effect of a 16.5% increase in maximum income support, (income-based) jobseeker's allowance and (non-contributory) employment support allowance. The 'In-work benefits' series shows the effect of an 18% increase in maximum working tax credit. The 'Income tax' series shows the effect of a 4% increase in the income tax personal allowance. In all cases, the baseline tax and benefit system is the 2015/16 system.

Figure 10: Lifetime Distributional Impact of Increases to Out-of-work and In-work Benefits and of the Income Tax Personal Allowance




Note: The 'Out-of-work benefits' series shows the effect of a 16.5% increase in maximum income support, (income-based) jobseeker's allowance and (non-contributory) employment support allowance. The 'In-work benefits' series shows the effect of an 18% increase in maximum working tax credit. The 'Income tax' series shows the effect of a 4% increase in the income tax personal allowance. In all cases, the baseline tax and benefit system is the 2015/16 system. All individuals face the same system throughout life uprated in line with average earnings (AEI). To aid comparison, we have scaled gains proportionally such that the 'All' bars are the same across reforms.

Figure 11: Distributional Impact of a One Percentage Point Increase in the Higher Rate of Income Tax



Note: Deciles are defined on the basis of equivalised net household income (cross-section income for the 'Cross-section' series and lifetime income for the 'Lifetime' series). The height of the bars is the gain or loss as a percentage of the relevant decile's total net (unequalised) household income (cross-section income for the 'Cross-section' series and lifetime income for the 'Lifetime' series). The baseline tax and benefit system is the 2015/16 system. For the 'Lifetime' series, all individuals face the same system throughout life uprated in line with average earnings (AEI).

Figure 12: Distributional Impact of a One Percentage Point Increase in the Basic Rate of Income Tax



Note: Deciles are defined on the basis of equivalised net household income (cross section income for the 'Cross-section' series and lifetime income for the 'Lifetime' series). The height of the bars is the gain or loss as a percentage of the relevant decile's total net (unequalised) household income (cross-section income for the 'Cross-section' series and lifetime income for the 'Lifetime' series). The baseline tax and benefit system is the 2015/16 system. For the 'Lifetime' series, all individuals face the same system throughout life uprated in line with average earnings (AEI).

ONLINE APPENDIX

Appendix A provides additional details on how we construct and validate our simulated lifetime profiles. Details on how we model employment and earnings as well as how we calculate taxes and benefits are given in the main body of the text. In Appendix B, we describe how we implemented an alternative fixed effects earnings process to compare with our main results.

Appendix A: Methods and Validation

Tables A1 and A2 show the specifications we run to estimate transition probabilities for each of our processes. Section A.II then discusses how we model partnering behaviour. Section A.III discusses how we model rents, which bears some similarities to the way we model earnings. In Section A.IV we describe how we impute private pension profiles. Section A.V discusses the imputation of consumption. In Section A.VI we provide some statistics on how well our simulations capture important moments of the data.

Table A1: Estimation Equations for Demographics and Rents

<i>Outcome</i>	<i>Method</i>	<i>Subsamples</i>	<i>Independent variables</i>
Mortality	Logit		Cubic in age, dummy for receipt of disability benefits, couple status, education dummies and earnings quintile
Child arrival	LPM	Run separately for women in couples and single women	For childless women: quadratic in age, dummy for ever had kids, number of kids ever had For women in couples: as for childless but also banded number of kids (0,1,2, and 3 or more) in household, age of youngest child, age of youngest child interacted with age
Child departure	LPM	Run separately by age of child (16-19)	Dummies for mother's and father's education
Partnering	Logit	Run separately for 3 education groups and sex	Quartic in age, dummy for employed last period, dummies for number of kids in household (0,1,and 2 or more), dummies for couple status in previous three periods, dummy for single status last period interacted with age
Separating	Logit	Run separately for own education and sex	Quartic in age, employed last period, partner employed last period, dummies for banded number of kids in household (0,1,and 2 or more), cubic in current relationship length, age of youngest child, dummy for education same as partner
Health (IB and DLA receipt)	Logit		For IB: quartic in age, 4 lags of employment status (interacted), 4 lags of IB status (interacted) earnings quartile last period For DLA: quartic in age, 4 lags of employment status (interacted), 4 lags of DLA status (interacted) earnings quartile last period and 2 lags of IB status
Renter (21 and over)	Logit	Run separately for current owners and current renters and for over and under 21s	Age of head of household, education of head of household, earnings quintile last period of head of household, banded number of kids (0,1,2 or 3 or more), couple status, relationship length, dummy for rented last period, 4 lags of ownership status
Rank in rent distribution (21 and over)	OL	Run separately for owners, and renters in each of 5 rent quintiles	Age of head of household, education of head of household, earnings quintile last period of head of household, banded number of kids (0,1,2 or 3 or more), couple status, relationship length dummy for rented last period, 4 lags of ownership status
Renter status and rank (under 21)	MNL		Age of head of household, age of head of household squared
Council tax band	OL	Run separately for each of 8 possible prior bands	Cubic in age, banded number of children (0,1,2,3, 4 or more), renter status, earnings quartile of household head, employment status

Notes: LPM = Linear probability model, OL = ordered logit, MNL = multinomial logit.

Table A2: Estimation Equations for Employment and Earnings

<i>Outcome</i>	<i>Method</i>	<i>Subsamples</i>	<i>Independent variables</i>
Employment (22 and over)	Logit	Run separately for males and females, by employment in prior wave and by employment 2 waves ago	Education dummies, quartic in age, age-education interactions, dummy for over state pension age, dummy for having kids, dummy for couple status, dummy for having kids under 5, kids under 5 interacted with cubic in age, 3 lags of full-time status, banded number of kids (0,1,2 and 3 or more), couple status, couple-age interaction, lagged full-time status, lagged earnings rank, dummies for earnings quartiles (and 5 lags), employment status 3, 4,5 and 6 waves ago (and interactions), lagged disability status
Earnings quartile and part-time/full-time status (22 and over)	MNL	Run separately for each of 5 possible prior states: in part-time work, in full-time work and in 4 earnings quartiles and separately for males and females	Education dummies, quartic in age, age-education interactions, dummy for over state pension age, dummy for has kids, couple status, dummy for kids under 5, 3 lags of full-time status, current earnings rank (and 3 lags), 3 status (interacted) lags of earnings quartile dummies, 3 lags of employment
Employment and earnings (under 22)	MNL	Run separately for each of 6 prior possible states: unemployment, in part-time work, in full-time work and in 4 earnings quartiles	Sex, education dummies, dummy for has kids and age
Earnings rank within `bin' (20 and over)	OLS	Run separately by prior state and sex	Cubics in 4 lagged (within bin) ranks interacted with cubic in age, education dummies, dummies for `bin' in previous 4 periods
Earnings rank within `bin' (under 20)	OLS	Run separately by prior state and sex	Cubics in lagged (within bin) ranks interacted with cubic in age, education dummies

Notes: LPM = Linear probability model, OL = ordered logit, MNL = multinomial logit.

A.II Partnering

Individuals select partners within our simulated sample. Thus all matches are assumed to take place within the same (nine year) birth cohort. We allow for assortative matching in the choice of partners on the basis of education level, such that university-educated individuals are more likely to match with other university-educated individuals than those with high school qualifications only. In order to implement this, we match potential partners based on an index that depends on education level and a random shock:

$$(A1) \quad I = ed_2 + \beta ed_3 + u$$

with $u \sim N(0, \sigma_u^2)$. The values of the unknown parameters β and σ_u^2 are chosen such that the distance between the simulated three-by-three matrix of education group against partner education group is as close to the empirical one as possible.

Which potential couples are realised, and which existing couples are dissolved, depends on partner arrival and departure probabilities estimated from our panel data and scaled so as to match the LCFS cross-sectional proportion of couples. Probabilities are scaled separately for those with and without children to allow for cohort differences in the partnering and separating behaviour of parents between our sample and the baby-boomers (in particular, there is a large secular increase in the proportion of single parents). The lower of the two male and female probabilities are used to calculate the probability of separation for couples. This is to allow us to better match the persistence of couples observed in the data. New couples and newly single individuals do not return to the partnering market until the following period.

Each couple requires a male and a female, and so a mismatch in the numbers of each can lead to too few matches being formed relative to what our estimated probabilities would imply. To avoid this happening, probabilities of partnering are again scaled to achieve the expected number of matches. Matches can only occur between individuals who are both aged 16 or older.

We wish to allow for the fact that males in couples in the 1945-54 cohort seen in the LCFS are on average just over two years older than females. This is important because it has a knock-on effect on the ages at which children are born. To achieve this, our simulated males are born in the years 1945-52 while females are born between 1947-1954. This means in each period that the marriage market will be composed of females that are on average two years younger than their male counterparts.

A.III Rents

For rental payments and ownership status, we adopt a very similar procedure to that for earnings. We first use the estimates of a logit to determine whether an individual is an owner or a renter. For those who are renters, we then use the estimates of an ordered logit to predict their rent quintile; controlling for education of the household head (assumed to be the male in any couple), a cubic in age for the household head, couple status, relationship length, banded number of children and several lags of past renter status and past quintiles of the rent distribution. Placement within rental quintiles is random. The variance of the rental distribution is not as great as that of earnings meaning the exact placement within quintiles matters less. If the lagged variables differ between two members of a couple, they are taken from the household head. For younger individuals for whom we do not have a complete set of lags (those under 21), we run a simpler multinomial logit to determine transitions across all the possible states.

The probability of owners becoming renters and renters become owners are scaled to match the LCFS proportions. For rental status this scaling is particularly important, as historically in the UK the proportion of renters was much higher than what we observe within the time frame covered by the BHPS panel.

A.IV Private Pensions

For private pensions we combine information from two datasets. The first consists of estimates of the discounted value of future private pension incomes for individuals in the BHPS survey from Disney et al. (2007). These estimates give the present value of future incomes for individuals if had they retired in 2001 or earlier, as well as projections for the future value of private pension wealth if individuals had continued in their present employment status until state retirement age. They are calculated using information from the special module of questions on private pensions included in the 2001 wave of the survey.⁹ The second is a set of predicted future private pension incomes for individuals seen in 2008 of the English Longitudinal Study of Ageing (ELSA). These include projected income streams conditional on individuals beginning to draw their private pensions in different years from 2008 onwards. The authors are indebted to Rowena Crawford, Soumaya Keynes and Gemma Tetlow for producing these projections and sharing them with us. Details of their

⁹ The data itself has been deposited in the UK Data Archive.

methodology can be found in Crawford (2012) with an example of their use in Banks et al. (2014).

The approach we follow allows us to match real-world private pension income profiles to our simulated individuals on the basis of their labour market histories and other characteristics. We implement it in the following steps once our simulations have completed.

1. We first estimate a probability that a simulated individual will ever receive a private pension using an individual's characteristics in 2001. We do this by estimating a logit model in the BHPS for that year. This regresses a dummy for positive projected private pension wealth in 2001 on sex and education dummies (and interactions of these), dummies for the number of the previous five years the individual was employed and dummies for the individuals' decile of a five-year moving average of previous earnings ranks.

2. We then predict the 2001 private pension 'wealth' (defined in here as the discounted value of future private pension incomes) for those simulated individuals who are to receive private pensions. This is done using the results of a regression of pension wealth in 2001 on a cubic in age, education dummies (and interactions of these) sex, years employed and a moving average of past earnings in the BHPS to which we add a normally distributed noise term.

3. Finally we calculate the simulated individuals' ranks in this distribution within cells defined by age and year. We can then use these ranks to match individuals to a one of a set of future streams of private pension income from the ELSA data within cells defined by cohort, sex and couple status in 2008 (or earlier if they retire before this). Ranks for our simulated individuals (estimated for 2001) are used to match individuals to private pension profiles ranked according to their present value in 2008.

An individual's retirement age is defined as the maximum of the final age at which they stopped working and 55. The ELSA data only predicts pension income for those who retire from 2008 onwards. For those who retire earlier than this, we deflate pension profiles associated with their retirement age using average earnings growth between 2008 and the year of their retirement. Earnings growth is what would determine private pension income for prior years from a defined benefit final salary scheme. The matching procedure works well,

with on average 100 potential matches for each individual and an average distance between the ranks of donors and recipients of less than 1 percentage point.

A.V Consumption

Including consumption spending in our simulations is important because it will help us calculate the value of indirect taxes individuals pay at different life stages. Detailed consumption expenditure is necessary because different spending items are subject to different tax treatments in the UK. VAT is not charged on food for example, and a reduced rate is charged on energy spending. Consumption is imputed to our individuals separately by spending categories defined by tax treatment using regressions estimated in the LCFS over the period 1978-2012. We are not able to sort consumption by tax category before this. As with other national consumption surveys, spending as recorded by the LCFS has tended to fall over time relative to national account measures (Brewer and O’Dea, 2012; Barrett et al., 2015). To offset this we scale spending categories by a common factor for each individual in such a way that total spending matches the national accounts figures in each year. A similar (but smaller) scaling is also applied to average earnings.

The manner in which we impute consumption captures variation by income, age and demographics, but does not allow for autocorrelations in shocks to spending. Other things equal, this will mean that those with volatile incomes in our data will have the same propensity to save as those with steady incomes. Such issues are unlikely to prevent us from drawing conclusions about the broad distributional impacts of consumption taxes over the lifetime. They will however have to be borne in mind when considering our results on the impacts of indirect taxation.

A.VI Validation

Figures A1-A6 show age profiles for males and females from our simulated individuals compared to those observed for the baby-boom cohort in the LCFS for couple status, employment, parenthood, single parenthood, number of children, and housing tenure. Despite our scaling procedure, cross-sectional averages for our simulations need not automatically match those in the LCFS, because the scaling only occurs within population subgroups. For instance, even if we accurately reproduced probabilities of being in a couple for those who have children and those who don’t, the proportion of couples would not match those in the LCFS if we did not also have the correct probabilities of being a parent at each age.

Nonetheless, the match between the simulated individuals and cross-sectional averages in the data is excellent for all variables and both sexes.

Age profiles show familiar “hump” shapes for parenthood and fertility. Employment tends to be higher for males than females and to decline with age. For females the age profile in employment shows a dip in the main child-rearing years just before age 30. The proportion of households who are renters declines steeply with age reflecting a secular increase in ownership rates over this period (partly driven by a series of “right to buy” reforms which allowed tenants in social housing to purchase their properties at reduced cost).

Some differences arise because of particular restrictions we impose. A difference in employment rates between the simulations and the data for younger ages is due to the fact that we impose that all those who have not completed full-time education are unemployed. A similar difference in the proportion of parents who are single in Figure A4 is due to the fact that, for years when cohorts are unobserved, we set the marriage rate for under 18s to be zero.

Figure A1: Proportion in Couples: Simulations versus Data




Figure A2: Employment: Simulations versus Data




Figure A3: Proportion of Parents: Simulations versus Data




Figure A.4: Proportion of Parents who are Single: Simulations versus Data




Figure A5: Number of Children: Simulations versus Data





Figure A6: Proportion of Renters: Simulations versus Data



Since we use our simulated individuals for distributional analysis of lifetime outcomes it is important that the persistence of variables such as income match those of the data, as well as cross-sectional average. Unfortunately we are not able to compare the persistence of variables for our simulated individuals directly with individuals from the baby-boom cohort throughout the whole life-cycle, because we do not have access to a panel dataset covering the whole of the adult life-cycle for the baby boomers. Instead, we plot autocorrelations for our spliced individuals against those individuals seen in the BHPS for the period 1991-2008. These may show whether the transitions we obtain are plausible even if they cannot be used for direct validation. Figures A7-A9 plot autocorrelations for 1 year ahead, 5 years ahead and 10 years ahead for males and females from ages 16-65 for couple status, parent status, and employment. The persistence of renter status is very different in the BHPS from the baby-boom cohort as a result of the much steeper declines experienced by the baby-boomers relative to those in later years. As a result autocorrelations for this variable are unlikely to be very informative and so we do not show them.

The figures show that employment, couple, and parent status have similar persistence in our simulations to the data for much of life- even over 10 a year horizon – but tend to have

greater persistence at the end of life. This latter feature may partly be due to differences between older individuals in the BHPS and in our cohort.

Figure A7: Autocorrelations Couple: Simulations versus Data




Figure A8: Autocorrelations Parent: Simulations versus Data





Figure A9: Autocorrelations Employment: Simulations versus Data



As a final check on the performance of our simulations, we can compare our simulated individuals to those from the same cohort in the waves they are observed in the BHPS. Table A3 shows the proportions always employed, always unemployed, always in a couple and always single over 10 years from 1995 to 2004 (inclusive) in the BHPS and in our simulations. Our simulations come very close to matching these proportions.

Table A3: Persistence of Employment and Couple Status for 1945-54 Cohort in BHPS and Simulations

	<i>BHPS</i>	<i>Simulations</i>
Always employed	55.9%	56.3%
Always unemployed	12.4%	11.4%
Always couple	73.7%	73.6%
Always single	15.6%	12.9%
N	676	4666

Appendix B: Fixed Effects Earnings Model

As an alternative to the earnings estimation procedure outlined in Section D, we also consider a more traditional fixed effects earnings regression. In particular, for those aged 22 and above we estimate a model of the form¹⁰

$$(B2) \quad \ln e_{it} = X_{it}\beta + FT_{it} + \lambda(Z_{it}\delta) + \alpha_i + u_{it}$$

$$u_{it} = \rho u_{it} + v_{it}$$

where e_{it} is individual earnings, FT_{it} is a dummy for full time status, α_i is an individual-level fixed effect, and X_{it} is a matrix of covariates (including a quartic in age, a dummy for being over state pension age, a cubic in age interacted with education, a dummy for having children, a dummy for couple status, three lags of employment status, and three lags of a dummy indicating whether or not the worker was working full-time or not).

$\lambda(Z_{it}\delta)$ are the generalised residuals from an ordered probit estimated for the outcomes unemployment, part-time work and full-time work. We include this to correct for selection into employment and full-time work.¹¹

The matrix Z_{it} contains the variables included in X_{it} as well as imputed net earnings variables. In addition we include the imputed net incomes individuals would be expected to receive if they were either unemployed, in part-time work or in full-time work. These variables act as instruments that should affect the probability of being in different employment states, while not entering the earnings equation in (B2).

To impute these we use the following two-step procedure:

- 1) We firstly impute gross earnings to individuals in different by employment states using the results of regressions of earnings on the covariates in X_{it} . These are estimated separately by employment status (part-time and full-time)
- 2) We then use these to impute tax and transfer payments (and thus net incomes) for individuals if they changed from their current employment status to unemployment, part-time or full-time work using the TAXBEN microsimulation model.¹²

¹⁰ For those under 22, we predict earnings with a simple OLS specification, omitting fixed effects and the AR(1) process in the error term.

¹¹ We do not make use of such corrections in our preferred approach to modelling earnings ranks. This is because we model transitions across all states rather than earnings levels, and under this approach there is no analogous selection issue.

We run this procedure separately for men and women. Employment transitions are estimated using the results from the ordered probit model.

Once we have obtained our estimates, we then initialise the simulations by randomly drawing from the estimated fixed effects (conditioning on education group), and selecting v_{i0} for each individual i by taking random draws from normal distribution with the variance of log earnings for 23 years olds taken from a the National Child Development Study (conditional on sex and education).¹³ In the simulations, we then estimate each individual earnings given the result of our model, and then rank individuals to obtain their locations in the earnings distribution. Finally, we can impute actual earnings conditional on simulated ranks using the LCFS data as we do in our main approach.

¹²Taxes are imputed conditional on a given system (1995).

¹³Due to data limitations, these variances are taken for those born in the 1958 cohort, which is slightly younger than the baby-boom cohort we would ideally want to use.