



Mobility and the lifetime distributional impact of tax and transfer reforms

Peter Levell¹ · Barra Roantree^{2,3}  · Jonathan Shaw⁴

Accepted: 28 August 2020
© The Author(s) 2020

Abstract

This paper examines the lifetime distributional impact of changes to the tax and transfer system. We find that—in contrast to 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. We also show that increases to higher rates of income tax are an effective way of targeting the lifetime rich because higher earners tend to exhibit greater persistence in their incomes. Our results illustrate the importance that moving beyond an exclusively snapshot perspective can have when analysing tax and transfer reforms.

Keywords Inequality · Redistribution · Income mobility

JEL codes D31 · H20 · H24

✉ Barra Roantree
barra.roantree@esri.ie

Peter Levell
peter.levell@ifs.org.uk

¹ Institute for Fiscal Studies, London, UK

² Economic and Social Research Institute, Dublin, Ireland

³ Trinity College Dublin, Dublin, Ireland

⁴ Financial Conduct Authority, London, UK

1 Introduction

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 converse is true.¹ However, such analyses suffer from an important limitation: the incomes used to rank households—and the changes in transfer entitlements or tax liabilities—are usually snapshot measures derived from cross-sectional data. These could provide a poor guide to the distributional implications of reforms because snapshot incomes are not always a good proxy for well-being. Individuals' circumstances often vary a great deal across life, and resources can be transferred across periods of life through borrowing and saving. As a result, well-being may be better measured using longer-run income measures (see, for example, Aaberge and Mogstad 2015; Bönke et al. 2015).

In this paper, we contrast the snapshot and lifetime distributional impacts of two sets of reforms, one targeting the bottom of the distribution and the other targeting the top. At the bottom, we investigate reforms to the working-age benefit system that have played an important role in the policy response to widening inequality (Blundell 2006): an expansion in work-contingent transfers (like the EITC in the United States) and an increase to out-of-work benefits (like jobseekers benefit in most European economies). It has been well-established that work-contingent transfers are an extremely effective means of reducing poverty (especially among children: Sherman 2009; Joyce and Sibieta 2013) and raising the employment of lone parents (Hotz and Scholz, Hotz and Scholz; Brewer et al. 2006). Yet they are also seen as involving an equity-efficiency trade off, with Eissa and Hoynes (2011), for example, arguing that they raise a “glaring” distributional issue by transferring resources only to working families and “providing minimal support to the poorest families”, who for the most part are not in paid work.

A key contribution of our work is to show that while increasing out-of-work benefits looks far more progressive than equally costly increases to work-contingent benefits at a snapshot in time, increases to work-contingent benefits are—on average—just as effective as increases to out-of-work benefits at redistributing resources to the lifetime poor. The reason for this is that the lifetime poor spend most of their working lives in low-paid work, whereas many of those who benefit from increases to out-of-work benefits are experiencing temporary periods of joblessness.²

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.

² Blundell et al. (2016) consider the relative welfare impacts of out of work and in work benefit payments. However, they are not able to provide reliable distributional estimates for the full population because they focus only on females and do not consider families with more than one child. It is also unclear which cohort their results relate to since they estimate transitions from data across all cohorts and do not correct to match aggregate patterns for a specific cohort or set of cohorts as we do (see Sect. 2)

Such patterns of mobility suggest that the perceived equity-efficiency trade-off that policymakers face in deciding between increasing work-contingent and out-of-work transfers is less stark than typically imagined. Alongside the well-documented fact that income is far more equally distributed on a lifetime than annual basis (e.g. Blomquist 1981; Björklund 1993; Creedy 1999; van de Ven 2006; Kopczuk et al. 2010; Brewer et al. 2012), these patterns also raise the question as to whether personal income taxes—which are typically assessed on monthly or annual income—are effective at targeting those who are lifetime rich. Like Bengtsson et al. (2012) we find that they are, with increases to the main rates of income tax primarily affecting those with high levels of income over their lifecycle. This is because higher—though not necessarily the highest—earners tend to exhibit greater persistence in their incomes over time. As a result, annually assessed personal income taxes appear well-targeted at the top half of the lifetime income distribution. While our analysis uses UK data and the UK tax and benefit system, our results are likely to generalise to other countries with similar labour market structures.³

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, they do not include all the information needed to calculate individuals' disposable incomes under counterfactual tax and transfer systems. Administrative datasets that do contain such information typically are not available for sufficiently long periods of time to cover individuals for the whole of their lives.⁴

Instead, we use statistical models of earnings, employment, health, housing tenure and family composition over the life-cycle to simulate lifetime profiles for individuals born in the 'baby-boom' cohort (1945–54). The approach we take deploys a method proposed in Levell and Shaw (2016), who compare the performance of alternative lifetime simulation methods. It involves first estimating transition equations for annual outcomes, conditioning on a rich set of demographics and outcomes for previous years using an 18-wave panel survey. The resulting estimates are then used to construct a simulated dataset for individuals that contains all the relevant characteristics for calculating tax and transfer payments. In order to ensure that our

³ Differences between cross-sectional and lifetime impacts for work-contingent benefits are likely to be even greater in the US, where shorter unemployment spells are more common, implying workers are more likely to cycle in and out of work. For example, in 2019 OECD statistics suggest that 19% of unemployment spells in the UK lasted under one month compared to 35% in the US (see OECD 'Unemployment by Duration' Labour Force Statistics series at https://stats.oecd.org/Index.aspx?DataSetCode=DUR_D, accessed 23rd July 2020). Fullerton and Rao (2019) also find that while 68% of individuals in the US did not pay federal taxes for at least one year in their lives, 21% of these individuals paid federal income tax in the following year and 45% paid some tax within the following five years.

⁴ For example, the comprehensive Swedish register data used by Bengtsson et al. (2012) covers those who are aged 20 in 1968 until they are 62 in 2009, excluding retirement. Conversely, while the US Social Security Administration data used by Kopczuk et al. (2007) follows cohorts from age 18 to 70, it contains only information on income subject to US Social Security taxes, excluding benefit income.

resulting simulated life-courses are representative of the experiences of our cohort of interest, estimates from the relatively short panel we use are combined with information from a much longer-running cross-sectional survey. Specifically, the transitions estimated using our panel data are adjusted 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.

Our modelling 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 vary 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 et al. 2013), observed patterns of wealth inequality (Castaneda et al. 2003; De Nardi et al. 2016), and consumption decisions (Arellano et al. 2017).

Our approach also allows us to realistically model complex tax liabilities and benefit entitlements over entire life-courses, enabling us to characterise the lifetime distributional effects of different tax and transfer reforms. However, it has two key limitations. The first is that our analysis does not incorporate behavioural responses to tax reforms, as is common in cross-sectional analyses of tax reforms.⁵ In our lifetime analysis we impose two counterfactual tax systems on individuals (a baseline tax system as if they were under the current tax and benefit system for the whole of their lives, and a reformed tax system) but do not model how these counterfactual systems might change individuals' earnings and employment, partnering and fertility. This allows us to model an almost 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.⁶

A second limitation concerns coverage of those with the very highest incomes. Those in the top 1% of the income distribution are often not captured well in household surveys (Atkinson et al. 2011; Burkhauser et al. 2018) but may have different earnings dynamics even when compared to those in neighbouring percentiles. For this reason, we do not examine reforms that are only likely to affect those at the very top of the income distribution (such as reforms to the top rate of income tax), and examine the distributional effects by cross-sectional and lifetime income deciles, rather than percentiles.

⁵ A justification often given for this is an appeal to the envelope theorem which implies that, for optimising individuals, the welfare effect of small reforms like the ones we consider is captured by the increase or decrease in incomes absent behavioural response (e.g. Piketty and Saez 2013).

⁶ See Blundell et al. (2016) and Haan et al. (2017) for examples of structural life-cycle models that examine tax reforms.

In contrasting the snapshot and lifetime distributional impact of these reforms, we add to a recent literature that has highlighted the importance of considering the long-run impacts of taxes and transfers. These include studies of their impact on individuals' education (Blundell et al. 2016), employment (Hoyne et al. 2016), earnings (Dahl et al. 2014; Neumark and Shirley 2017) and health (Evans and Garthwaite 2014), as well as that of their children (Bastian and Micheltore 2018).

We also contribute to a literature examining the lifetime distributional impact of policy that has tended to focus on indirect taxes (Davies et al. 1984; Poterba 1989; Caspersen and Metcalf 1994; Lyon and Schwab 1995; Metcalf 1999). This showed that while such taxes on consumption look highly regressive when assessed against annual income, the same is not true when lifetime income or—as a proxy—annual consumption is used. Similarly, we show the importance of moving beyond an exclusively snapshot perspective when analysing direct tax and transfer reforms.

Our results suggest that policymakers seeking to redistribute resources to the long-run poor might favour doing so through increases to work-contingent benefits, given that these typically have less damaging effects on work incentives than out-of-work benefits. However, redistributing to those with few lifetime resources is by no means the only goal of policy. A growing literature has shown the consumption floor provided by benefits provides significant insurance value to individuals, both in retirement (O'Dea 2018) and working life (Landais and Spinnewijn 2019). In addition, credit market imperfections mean that redistribution towards periods of temporarily low current incomes can be particularly welfare-improving even if policymakers' primary concern is the distribution of lifetime resources (Roantree and Shaw 2018). Rather, our results illustrate the importance of taking both a lifetime and snapshot perspective and of being clear about what one is trying to achieve: the alleviation of short-term hardship or the redistribution of lifetime resources.

The remainder of this paper is structured as follows. In the next section, we describe how we construct our simulated lifetime profiles, with further details provided in Appendices. Section 3 presents some descriptive statistics on lifetime incomes and how they compare to cross-sectional measures. This forms the backdrop to our results on the distributional effects of tax and transfer reforms on lifetime and cross-sectional income in Sect. 4. Section 5 concludes.

2 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, see Appendix A and Levell and Shaw (2016), who compare our microsimulation approach to others used in the literature.

2.1 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 5500 households and was nationally representative south of the Caledonian canal. 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.

In addition to these datasets, we also make use of the English Longitudinal Study of Ageing (ELSA) to impute private pension incomes to our simulated individuals. ELSA is a panel survey of the over 50s in England, with rich information on private pension wealth.

2.2 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 Appendix A.

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 (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 5000 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.⁷ These are calculated for different possible retirement dates for each individual given observed pension wealth. 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. The only additional complication is that after estimating transition probabilities using the BHPS our results are scaled in order to match cross-sectional patterns of demographics and employment observed for our cohort of interest. We describe this scaling procedure in Sect. 2.4. Further details of the exact approaches we take are presented in Appendix A. The way we model earnings is a little more complicated however, and so we provide more detail here.

To model movements of individuals across the earnings distribution accurately, we must be careful to allow for key features of real-world transitions. 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. Recent work using detailed administrative data (Guvonen et al. 2015) has shown that

⁷ Details of the methodology to construct these profiles can be found in Crawford (2012) and an example of their use in Banks et al. (2014)

the distribution and persistence of shocks differs in important ways both over the life-cycle and according to current earnings.

The approach we take is to directly estimate 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 model individuals' earnings ranks (relative positions in the distribution by age and year) in a three-step approach. First, we model movements into and out of employment. Second, for those employed, we then model movements across five earnings 'bins' (corresponding to quartiles of earnings if in full-time employment and a separate bin for part-time employment). Third, we predict individuals' exact earnings ranks within each bin and 'fill in' earnings levels from the actual cross-sectional distributions of earnings observed for our particular birth cohorts in the LCFS. More details on the procedure is given in the following subsection. But the key points here are that we capture earnings and employment transitions in a highly flexible way and match cross-sectional distributions (means, variances and higher order moments) exactly 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.

2.3 Details on modelling employment and earnings

In this subsection, we provide more detail on our approach to modelling employment and earnings. We use the following three-step parametric approach.

1. *Determine employment status* We estimate transition matrices for employment status separately for males and females using multinomial logit models 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 below).

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)

$$\Pr(i, j | X_{it}) = \frac{\exp(X_{it}\beta^{ij})}{\sum_{m=0}^N \exp(X_{it}\beta^{im})} \tag{1}$$

where 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.

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 earning quartiles. However, it does not place individuals precisely within these quartiles. One approach to deal with this is to match simulated individuals to real-world individuals who made the same employment and earnings 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

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

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 $t - 1$. $\Phi^{-1}(\dots)$ 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 $\Phi(\dots)$] 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 Eq. 2. 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 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 individuals' earnings histories. In Appendix A.1, we assess how well we capture patterns of mobility that we observe in the data compared to a more traditional earnings process. We find that our approach is much better at matching the tendency of individuals to remain at the same location in the earnings distribution as well as the greater persistence of earnings in the tails of the distribution (and particularly for high earners).⁸

2.4 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 Harding 1996). Another is to adopt a splicing approach that joins individuals

⁸ Levell and Shaw (2016) compare autocorrelations of earnings ranks using this approach with those in the BHPS for time periods of one year, five years and ten years. The autocorrelations in the data differ over different stages of life. Specifically, they tend to be lower at younger and older ages. Our simulated data capture these patterns reasonably well. While the fit for one and five year horizons is good, we note that 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. In addition, Levell and Shaw compare the persistence of employment in our simulations with those from our cohort of interest as observed in the BHPS over a 10 year horizon. They find that the proportions always employed and never employed within this period are very similar.

observed at different ages within the same short period together to form a complete lifecycle (Bovenberg et al. 2008).⁹ 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 modify the simulation approach in the following way. 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.¹⁰

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:

$$p_{it}^g = \text{Prob}(y_{it}^g = 1) \tag{3}$$

$$\pi_{it}^{g,01} = \text{Prob}(y_{it}^g = 1 | y_{i,t-1}^g = 0) \tag{4}$$

$$\pi_{it}^{g,11} = \text{Prob}(y_{it}^g = 1 | y_{i,t-1}^g = 1) \tag{5}$$

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

$$p_{it}^g = \pi_{it}^{g,01} (1 - p_{i,t-1}^g) + \pi_{it}^{g,11} * p_{i,t-1}^g \tag{6}$$

It is clear that $\pi_{it}^{g,01}$ and $\pi_{it}^{g,11}$ cannot be uniquely identified with knowledge of p_{it}^g and $p_{i,t-1}^g$ alone. This is because a variety of transition probabilities could in principle satisfy Eq. 6. 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

⁹ For a comparison of our current method with that of the splicing approach used in Bovenberg et al. (2008), see Levell and Shaw (2016).

¹⁰ The latest year we observe individuals is 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.

$$\left(\frac{1}{N_g} \sum_i \pi_{it}^{g,01} - \frac{1}{N_g} \sum_i \hat{\pi}_{it}^{g,01}\right)^2 + \left(\frac{1}{N_g} \sum_i \pi_{it}^{g,11} - \frac{1}{N_g} \sum_i \hat{\pi}_{it}^{g,11}\right)^2 \tag{7}$$

subject to satisfying Eq. 6. 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 $z^{g,1}$ to minimise

$$\left(\frac{1}{N_g} \sum_i \Lambda(X_{it}\beta^{g,01} + z^{g,1}) - \frac{1}{N_g} \sum_i \Lambda(X_{it}\beta^{g,01})\right)^2 + \left(\frac{1}{N_g} \sum_i \Lambda(X_{it}\beta^{g,11} + z^{g,2}) - \frac{1}{N_g} \sum_i \Lambda(X_{it}\beta^{g,11})\right)^2 \tag{8}$$

again subject to satisfying Eq. 6. Λ 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 individuals' 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 $z^{g,1}$ 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.

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). However, as we show in Appendix A, 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

Table 1 Cells within which transition probabilities are scaled

State	Cells
Couple	Age, year, sex, number of children
Renter	Age, year
Employed	Age, year, sex, has children

individuals in the BHPS to scale mortality rates predicted using a logit regression on income, disability benefit receipt, education and couple status.

2.5 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 (Waters 2017). We model all major benefit payments—both universal and means-tested—including unemployment insurance, payments to families with children, disability benefits and tax credits.¹¹ 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, we are limited in our ability to accurately model taxes on the very highest (i.e. top 1% of) earners because they are not captured well in household surveys (Atkinson et al. 2011; Burkhauser et al. 2018). For this reason, although we incorporate the additional rate of income tax in our simulations, we do not examine reforms to this rate of tax. Finally, 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.

2.6 Behavioural responses

Our microsimulation approach does not incorporate possible behavioural responses to tax and transfer reforms. While not without limitations, it allows us to model the impact of a much more complete set of characteristics relevant for calculating tax and transfer payments, including 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.¹² The approach we take allows us to model the effects of more realistic and intricate tax and transfer reforms than would otherwise be the case.

¹¹ 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.

¹² For example, Blundell et al. (2016) build a rich dynamic lifecycle model of women's education, employment and savings decisions, but impose exogenous partnering and fertility dynamics - as here - and abstract altogether from housing tenure and health.

3 Simulated 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.¹³

The course of incomes over the lifecycle is shown in Fig. 1. It 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.¹⁴ The figure shows that, on average, earnings rise for men until the late 40s and then



Fig. 1 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

¹³ 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 an earlier version of this paper (Levell et al. 2015), we compare this synthetic cross-section to cross-sections drawn from the 1978 and 2012 LCFS, and show that one obtains broadly similar results using these.

¹⁴ 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.




Fig. 2 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 equalised net income ignoring indirect taxes (annualised net income for lifetime deciles)

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 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.

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 living short lives but with higher average incomes. Such outcomes are counterintuitive. Instead, we use a measure of the discounted average lifetime income

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

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 equalised 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 2 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 mobility is non-symmetric. 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%).¹⁵

Figure 3 shows how employment varies across the net income distribution. The cross-section series shows the proportion of working-age individuals who are




Fig. 3 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)

¹⁵ A similar result is obtained when we consider the BHPS data itself (see Figure 2.8 in Levell et al. 2015). Analysis using US 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 et al. 2010).




Fig. 4 Distribution of gross income. 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 Eq. 9). We exclude the top 1% of incomes and those with zero incomes

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.

Finally, Fig. 4 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.

4 Results

In this section, we contrast the snapshot and lifetime distributional impact of two sets of changes to the tax and transfer system. We first consider three reforms which have been seen as alternative ways of redistributing resources towards lower income families, with different trade-offs between equity and efficiency. We

then consider changes to the basic and higher rates of income tax, key policy instruments for financing such increases in expenditure.

The analysis assumes throughout 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, tax liabilities 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 average earnings growth. This brings us close to ensuring that the tax and benefit system raises the same revenue each year.¹⁶

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), while 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 described for the EITC in Eissa and Hoynes (2011). How does a lifetime analysis affect this assessment?

To answer this question we consider three reforms relative to the 2015–16 tax and benefit system with similar cross-sectional revenue consequences (around £3 billion/\$4.2 billion per year), namely:

- 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;
- an increase in in-work (i.e. work-contingent) benefits: an 18% increase in the maximum working tax credit award;
- 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).

Figure 5 shows the cross-sectional effect of the three reforms listed above.¹⁷ Unsurprisingly, the most progressive reform is the increase to out-of-work

¹⁶ An earlier version of this paper (Levell et al. 2015) shows that adopting an alternative approach of uprating tax and benefit thresholds by prices leads to similar results.

¹⁷ Appendix C shows versions of Figures 5, 6 and 7 and 8 using an alternative measure of the distributional impacts: the proportions of the total gains and losses accruing to households in different income deciles. The results are qualitatively similar to those presented here though suggest a larger share of total gains from the income tax cut accrue to the lifetime richest decile.

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.

The pattern of gains over lifetime income is different, as shown in Fig. 6. 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. This stems from the pattern of worklessness over the lifecycle shown in Fig. 3. Although the poorest individuals at a point in time are frequently out of work, this is often a temporary state with many moving in and out of work over the course of their lifetimes.

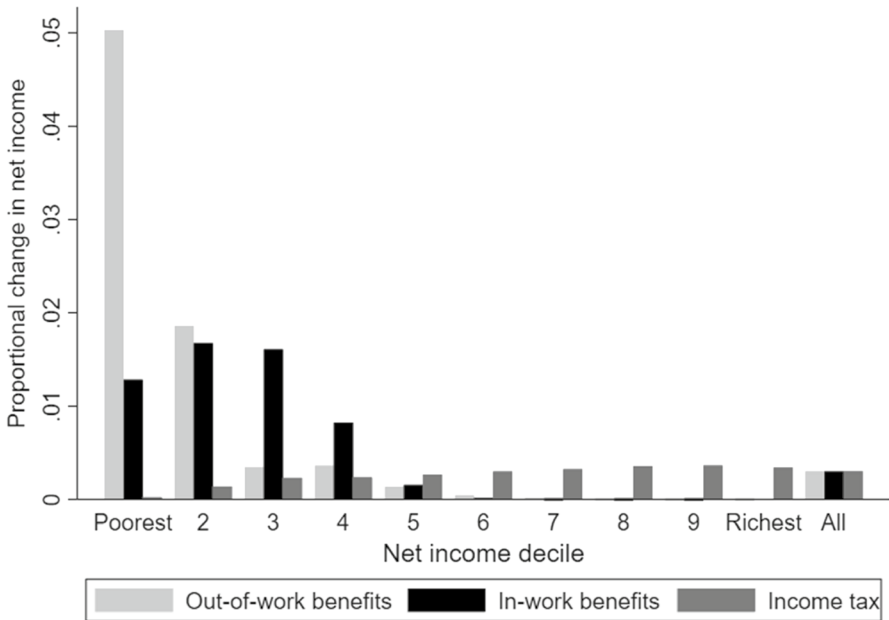


Fig. 5 Cross-sectional distributional impact of reforms. Note: Deciles are defined on the basis of cross-sectional equivalised net household income. The height of the bars is the gain or loss as a percentage of the relevant decile’s total net (unequalised) household income. The ‘Out-of-work benefits’ 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




Fig. 6 Lifetime distributional impact of reforms. Note: Deciles are defined on the basis of average discounted equivalised disposable income, as described in the text. The height of the bars is the gain or loss as a percentage of the relevant decile's total net (unequalised) household income. To aid comparison, we have scaled gains proportionally such that the 'All' bars are the same across reforms. 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

In addition, when in work, they are relatively likely to be in low-paid work and qualify for in-work support while also being liable for income tax.

That increases to out-of-work and in-work benefits have such a similar distributional pattern is a particularly striking result. In-work benefits have the advantage 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.




Fig. 7 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 updated in line with average earnings (AEI)

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. However, recent experience suggests that, at least in the UK, the lifetime poor among younger cohorts are so because of low levels of earnings, rather than long periods out of the labour market (Belfield et al. 2014).

Changes to income tax are among the primary policy instruments that policymakers have to finance increased expenditure. We therefore consider the snapshot and lifetime distributional impact of 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).¹⁸

¹⁸ The 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.




Fig. 8 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)

Figure 7 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 few individuals in these 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 Fig. 1). 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 (Fig. 8). 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.

5 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 two transfer reforms which have been the subject of intensive policy debate in many countries.

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 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 (O'Dea 2018; Landais and Spinnewijn 2019; Roantree and Shaw 2018). Rather, our results illustrate the importance of taking both a lifetime and snapshot perspective and of being clear about what one is trying to achieve: the alleviation of short-term hardship or the redistribution of lifetime resources.

Appendix A provides additional details on how we construct and validate our simulated lifetime profiles. This appendix draws partly on the description of the model in Levell and Shaw (2016) but also includes additional validation statistics. 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.

Acknowledgements JS the authors gratefully acknowledge a grant from the Nuffield Foundation (OPD/40976), the European Research Council (reference ERC-2010-AdG-269440-WSCWTBDS) and the ESRC-funded Centre for the Microeconomic Analysis of Public Policy (RES-544-28-5001). 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, advisory group participants and referees for comments on earlier drafts. The English Longitudinal Study of Ageing (ELSA) data were made available through the UK Data Archive (UKDA) and were collected by the National Centre for Social Research. 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, and supplied by the UK Data Service. The BHPS is 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.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A: Methods and validation

Tables 2 and 3 show the specifications we run to estimate transition probabilities for each of our processes. In Sect. A.1, we assess how well-modelled earnings and employment over the life cycle matches the patterns observed in real data. Section A.2 then discusses how we model partnering behaviour. Section A.3 discusses how we model rents, which bears some similarities to the way we model earnings. In Sect. A.4 we describe how we impute private pension profiles. Section A.5 discusses the imputation of consumption. For details on the performance of our model in matching average rates of coupling, number of children and housing tenure relative to our cohort of interest as well as in matching the persistence of earnings over time, see Levell and Shaw (2016).

Table 2 Estimation equations for demographics and rents

Outcome	Method	Subsamples	Independent variables
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 lagged employed, dummies for number of kids in household (0, 1, and 2 or more), dummies for couple status in previous three periods, dummy for lagged single status 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: age quartic, 4 lags of employment status (interacted), 4 lags of DLA status (interacted) lagged earnings quartile and 2 lags 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

Table 2 (continued)

Outcome	Method	Subsamples	Independent variables
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

Table 3 Estimation equations for employment and earnings

Outcome	Method	Subsamples	Independent variables
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

Appendix A.1: Assessing modelled earnings over the life cycle

Our approach to modelling earnings, described in Sect. 2.3, 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, Table 4 shows the proportion of workers in different earnings deciles in the following year given their current earnings deciles in the BHPS data and for our simulated individuals. In Table 4, we show the same for earnings deciles five years ahead. Both Tables 4 and 5 show individuals are highly likely to remain in their current decile, with persistence increasing for the top and bottom deciles. Persistence is however greater for those at the top of the earnings distribution than for those at bottom. Comparing results in the first and second rows for each decile shows that we replicate these patterns in our simulated data well, though we may understate the overall persistence in individual incomes.

The fit of our model is a considerable improvement over a more standard estimated earnings process. In the third row for each earnings decile in Tables 4 and 5, we report earnings deciles one and five years ahead 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 this alternative is relatively poor. The model underestimates the probability that individuals will remain in the current earnings decile. 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 earners.

A second test of our model concerns its ability evolution of earnings mobility over the life cycle. Plots of the autocorrelations of earnings ranks separated by one year, five years and ten years for males and females aged 16-65 are presented in Levell and Shaw (2016). The autocorrelations in the data differ over different stages of life. Specifically, they 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.

For comparison, Fig. 9 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.

Table 4 Earnings deciles in years t and $t + 1$

	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	Decile 6	Decile 7	Decile 8	Decile 9	Decile 10
Decile 1										
Actual	0.61	0.20	0.06	0.03	0.02	0.02	0.02	0.01	0.01	0.01
Preferred simulation method	0.51	0.24	0.08	0.05	0.03	0.02	0.02	0.01	0.02	0.02
FE AR(1)	0.48	0.21	0.11	0.06	0.04	0.03	0.02	0.02	0.02	0.02
Decile 2										
Actual	0.16	0.48	0.20	0.06	0.03	0.02	0.02	0.01	0.01	0.01
Preferred simulation method	0.23	0.32	0.21	0.10	0.04	0.03	0.02	0.01	0.02	0.02
FE AR(1)	0.14	0.22	0.20	0.15	0.10	0.07	0.04	0.03	0.03	0.02
Decile 3										
Actual	0.05	0.17	0.42	0.20	0.07	0.04	0.02	0.02	0.01	0.01
Preferred simulation method	0.07	0.20	0.34	0.20	0.08	0.04	0.03	0.01	0.01	0.02
FE AR(1)	0.06	0.15	0.18	0.16	0.14	0.11	0.08	0.05	0.03	0.02
Decile 4										
Actual	0.03	0.05	0.17	0.38	0.21	0.07	0.04	0.02	0.01	0.01
Preferred simulation method	0.04	0.09	0.20	0.28	0.20	0.09	0.05	0.02	0.02	0.02
FE AR(1)	0.03	0.10	0.14	0.17	0.15	0.13	0.12	0.08	0.05	0.02
Decile 5										
Actual	0.02	0.03	0.06	0.18	0.37	0.20	0.08	0.03	0.02	0.01
Preferred simulation method	0.03	0.05	0.08	0.19	0.32	0.19	0.08	0.03	0.02	0.01
FE AR(1)	0.02	0.07	0.11	0.14	0.16	0.15	0.14	0.11	0.07	0.04
Decile 6										
Actual	0.01	0.02	0.03	0.06	0.19	0.36	0.21	0.07	0.03	0.02
Preferred simulation method	0.01	0.02	0.03	0.08	0.18	0.32	0.22	0.08	0.04	0.02
FE AR(1)	0.02	0.05	0.08	0.11	0.15	0.15	0.15	0.14	0.10	0.06
Decile 7										
Actual	0.01	0.01	0.02	0.04	0.07	0.18	0.37	0.20	0.07	0.03

Table 4 (continued)

	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	Decile 6	Decile 7	Decile 8	Decile 9	Decile 10
Preferred simulation method	0.01	0.02	0.02	0.05	0.08	0.20	0.29	0.21	0.09	0.03
FE AR(1)	0.02	0.03	0.06	0.08	0.12	0.14	0.16	0.16	0.14	0.09
Decile 8										
Actual	0.01	0.01	0.01	0.02	0.03	0.06	0.19	0.41	0.21	0.05
Preferred simulation method	0.01	0.01	0.02	0.02	0.03	0.07	0.19	0.38	0.22	0.05
FE AR(1)	0.01	0.02	0.04	0.06	0.09	0.12	0.16	0.18	0.19	0.13
Decile 9										
Actual	0.01	0.01	0.01	0.01	0.02	0.03	0.06	0.18	0.50	0.17
Preferred simulation method	0.01	0.01	0.01	0.02	0.02	0.04	0.09	0.21	0.38	0.22
FE AR(1)	0.01	0.02	0.02	0.04	0.06	0.09	0.13	0.18	0.23	0.23
Decile 10										
Actual	0.01	0.01	0.01	0.01	0.01	0.02	0.02	0.04	0.15	0.71
Preferred simulation method	0.01	0.01	0.01	0.02	0.02	0.02	0.03	0.05	0.20	0.63
FE AR(1)	0.01	0.01	0.02	0.02	0.03	0.04	0.07	0.12	0.21	0.47

Notes: Table shows the probability of being in different earnings deciles next year conditional on current earnings decile in the data and in our simulations. FE AR(1) are simulated results estimated from a model with fixed effects and a normally distributed AR(1) error. Preferred simulation method refers to results from the approach used to obtain results in the main text

Table 5 Earnings deciles in years t and $t + 5$

	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	Decile 6	Decile 7	Decile 8	Decile 9	Decile 10
Decile 1										
Actual	0.38	0.25	0.12	0.07	0.05	0.03	0.03	0.02	0.02	0.02
Preferred simulation method	0.29	0.22	0.16	0.10	0.07	0.05	0.04	0.03	0.03	0.03
FE AR(1)	0.40	0.16	0.12	0.09	0.07	0.05	0.04	0.03	0.03	0.02
Decile 2										
Actual	0.15	0.30	0.22	0.12	0.06	0.05	0.03	0.03	0.02	0.02
Preferred simulation method	0.19	0.21	0.19	0.13	0.09	0.06	0.05	0.04	0.03	0.03
FE AR(1)	0.15	0.16	0.15	0.12	0.11	0.09	0.08	0.06	0.04	0.03
Decile 3										
Actual	0.08	0.15	0.24	0.21	0.13	0.07	0.05	0.03	0.03	0.02
Preferred simulation method	0.11	0.15	0.21	0.17	0.13	0.09	0.06	0.04	0.03	0.02
FE AR(1)	0.10	0.14	0.13	0.13	0.12	0.11	0.09	0.08	0.07	0.04
Decile 4										
Actual	0.05	0.07	0.14	0.24	0.22	0.12	0.07	0.04	0.03	0.02
Preferred simulation method	0.07	0.10	0.15	0.18	0.18	0.14	0.08	0.05	0.04	0.03
FE AR(1)	0.07	0.11	0.12	0.13	0.12	0.11	0.11	0.09	0.08	0.06
Decile 5										
Actual	0.04	0.05	0.07	0.14	0.22	0.21	0.12	0.08	0.05	0.02
Preferred simulation method	0.06	0.07	0.09	0.15	0.19	0.17	0.12	0.07	0.05	0.02
FE AR(1)	0.05	0.09	0.10	0.12	0.12	0.12	0.11	0.11	0.10	0.07
Decile 6										
Actual	0.03	0.04	0.05	0.08	0.14	0.22	0.22	0.12	0.06	0.03
Preferred simulation method	0.04	0.05	0.06	0.09	0.14	0.19	0.18	0.13	0.08	0.04
FE AR(1)	0.04	0.08	0.10	0.10	0.11	0.12	0.12	0.12	0.11	0.09
Decile 7										
Actual	0.03	0.03	0.04	0.05	0.08	0.14	0.23	0.22	0.13	0.05

Table 5 (continued)

	Decile 1	Decile 2	Decile 3	Decile 4	Decile 5	Decile 6	Decile 7	Decile 8	Decile 9	Decile 10
Preferred simulation method	0.03	0.04	0.05	0.07	0.10	0.15	0.19	0.19	0.12	0.06
FE AR(1)	0.04	0.06	0.09	0.10	0.11	0.12	0.13	0.12	0.13	0.11
Decile 8										
Actual	0.03	0.02	0.03	0.04	0.05	0.08	0.16	0.26	0.25	0.09
Preferred simulation method	0.03	0.03	0.04	0.05	0.06	0.10	0.15	0.23	0.20	0.11
FE AR(1)	0.03	0.05	0.07	0.08	0.10	0.11	0.13	0.13	0.15	0.14
Decile 9										
Actual	0.02	0.02	0.02	0.03	0.04	0.05	0.09	0.15	0.33	0.25
Preferred simulation method	0.03	0.03	0.03	0.04	0.04	0.07	0.11	0.17	0.25	0.24
FE AR(1)	0.02	0.04	0.06	0.07	0.09	0.10	0.12	0.15	0.16	0.19
Decile 10										
Actual	0.03	0.02	0.02	0.02	0.03	0.03	0.04	0.06	0.16	0.59
Preferred simulation method	0.02	0.03	0.03	0.03	0.03	0.04	0.06	0.09	0.20	0.47
FE AR(1)	0.02	0.03	0.04	0.05	0.06	0.08	0.11	0.13	0.18	0.31

Notes: Table shows the probability of being in different earnings deciles five years ahead conditional on current earnings decile in the data and in our simulations. FE AR(1) are simulated results estimated from a model with fixed effects and a normally distributed AR(1) error. Preferred simulation method refers to results from the approach used to obtain results in the main text




Fig. 9 Autocorrelations of earnings ranks: fixed effects earnings process simulations versus data

Appendix A.2: 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:

$$I = ed_2 + \beta ed_3 + u \tag{10}$$

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.

Appendix A.3: 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.

Appendix A.4: 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. (2009). 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. 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.¹⁹

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:

- We 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.
- 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.
- 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.

¹⁹ The authors are indebted to Rowena Crawford, Soumaya Keynes and Gemma Tetlow for producing these projections and sharing them with us. Details of their methodology can be found in Crawford (2012) with an example of their use in Banks et al. (2014).

Appendix A.5: 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. 2013). 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 would however have to be borne in mind when considering our results on the impacts of indirect taxation.

Appendix B: Fixed effects earnings model

As an alternative to the earnings estimation procedure outlined in Sect. 2, we also consider a more traditional fixed effects earnings regression. In particular, for those aged 22 and above we estimate a model of the form²⁰

$$\begin{aligned} \ln e_{it} &= X_{it}\beta + FT_{it} + \lambda(Z_{it}\delta) + \alpha_i + u_{it} \\ u_{it} &= \rho u_{i,t-1} + v_{it} \end{aligned} \quad (11)$$

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.²¹

²⁰ For those under 22, we predict earnings with a simple OLS specification, omitting fixed effects and the AR(1) process in the error term.

²¹ 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.

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 Eq. 11.

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

- 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)
- 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.²²

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).²³ 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.

Appendix C: Alternative decile charts

In this Appendix we plot versions Figs. 5, 6 and 7 and 8 using an alternative measure of the distributional impacts: the proportion of gains and losses accruing to different income groups.

Figures 10 and 11 show similar patterns to those shown in Figs. 5 and 6. 47% of the cross-sectional gains from the increases in out-of-work benefits we consider go to the poorest 10% of households. Those in the bottom decile of lifetime incomes receive only 9% of the lifetime gains from this reform however. The benefits from increases in in-work benefits are much more evenly distributed across lifetime income deciles than cross-sectional income deciles. Those in the top income decile also receive a 5% of the lifetime gains from an increase in out of work benefits, even though, as Fig. 6 shows, this accounts for a very small share of their lifetime incomes.

²² Taxes are imputed conditional on a given system: that in April 1995.

²³ Due 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.

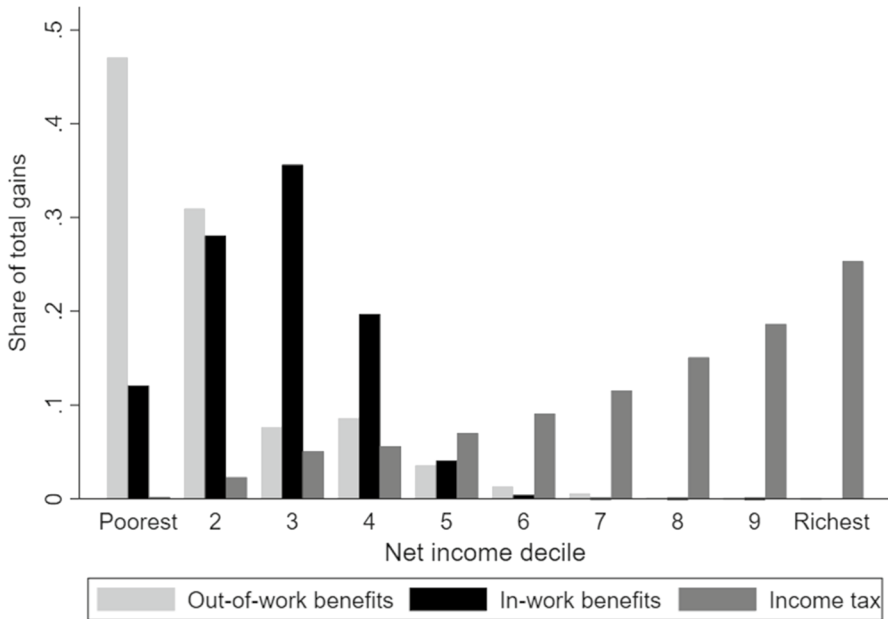


Fig. 10 Cross-sectional distributional impact of reforms (proportion of gains). Note: Deciles are defined on the basis of cross-sectional equivalised net household income. The height of the bars is the share of total gains from reforms accruing to households in different income groups. The ‘Out-of-work 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

Figures 12 and 13 show the proportion of revenue from increases in the higher and basic rate of income tax accruing to different cross-sectional and lifetime deciles. These show similar patterns to Figs. 7 and 8. The share of revenue raised from the top lifetime income deciles is in both cases lower than the proportion of revenue raised from those with the highest incomes in a cross-sectional snapshot.




Fig. 11 Lifetime distributional impact of reforms (proportion of gains). Note: Deciles are defined on the basis of average discounted equivalised disposable income, as described in the text. The height of the bars is the share of total gains from reforms accruing to households in different lifetime income groups. 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




Fig. 12 Distributional impact of a one percentage point increase in the higher rate of income tax (proportion of losses). 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 proportion of total revenue raised from different income groups (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 updated in line with average earnings (AEI)




Fig. 13 Distributional impact of a one percentage point increase in the basic rate of income tax (proportion of losses). 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 proportion of total revenue raised from different income groups (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 updated in line with average earnings (AEI)

References

- Aaberge, R., & Mogstad, M. (2015). Inequality in current and lifetime income. *Social Choice and Welfare*, *44*(2), 217–230.
- Altonji, J. G., Smith, A. A., & Vidangos, I. (2013). Modeling earnings dynamics. *Econometrica*, *81*(4), 1395–1454.
- Arellano, M., Blundell, R., & Bonhomme, S. (2017). Earnings and consumption dynamics: A nonlinear panel data framework. *Econometrica*, *85*, 693–734.
- Atkinson, A. B., Piketty, T., & Saez, E. (2011). Top incomes in the long run of history. *Journal of Economic Literature*, *49*, 3–71.
- Banks, J., Emmerson, C., & Tetlow, G. C. (2014). Effect of pensions and disability benefits on retirement in the UK. NBER working paper 19907, National Bureau of Economic Research, Boston.
- Barrett, G., Levell, P., & Milligan, K. (2013). A comparison of micro and macro expenditure measures across countries using differing survey methods. NBER working paper 19544, National Bureau of Economic Research, Boston.
- Bastian, J., & Michelmore, K. (2018). The long-term impact of the earned income tax credit on children’s education and employment outcomes. *Journal of Labor Economics*, *36*(4), 1127–1163.
- Belfield, C., Cribb, J., Hood, A., & R. Joyce (2014). Living standards, poverty and inequality in the UK: 2014. IFS report 96, Institute for Fiscal Studies, London.
- Bengtsson, N., Holmlund, B., & Waldenström, D. (2012). Lifetime versus annual tax progressivity: Sweden, 1968–2009. IZA discussion paper series No 6641, Institute of Labor Economics (IZA), Bonn.

- Björklund, A. (1993). A comparison between actual distributions of annual and lifetime income: Sweden 1951–89. *Review of Income and Wealth*, 39(4), 377–386.
- Blomquist, N. S. (1981). A comparison of distributions of annual and lifetime income: Sweden around 1970. *Review of Income and Wealth*, 27(3), 243–264.
- Blundell, R. (2006). Earned income tax credit policies: Impact and optimality: The Adam Smith Lecture, 2005. *Labour Economics*, 13(4), 423–443.
- Blundell, R. (2014). Income dynamics and life-cycle inequality: Mechanisms and controversies. *The Economic Journal*, 124(576), 289–318.
- Blundell, R., Dias, M. C., Meghir, C., & Shaw, J. M. (2016). Female labour supply, human capital and welfare reform. *Econometrica*, 84(5), 1705–1763.
- Bovenberg, A. L., Hansen, M. I., & Sørensen, P. B. (2008). Individual savings accounts for social insurance: Rationale and alternative designs. *International Tax and Public Finance*, 15(1), 67–86.
- Bowlus, A. J., & Robin, J.-M. (2012). An international comparison of lifetime inequality: How continental Europe resembles North America. *Journal of the European Economic Association*, 10(6), 1236–1262.
- Brewer, M., Costa Dias, M., & Shaw, J. (2012). Lifetime inequality and redistribution. IFS working paper W12/23, Institute for Fiscal Studies, London.
- Brewer, M., Goodman, A., & Leicester, A. (2006). *Household spending in Britain: What can it teach us about poverty?*. Bristol: The Policy Press.
- Brewer, M., & O’Dea, C. (2012). Measuring living standards with income and consumption: Evidence from the UK. IFS working paper W12/12, Institute for Fiscal Studies, London.
- Buchinsky, M., & Hunt, J. (1999). Wage mobility in the United States. *The Review of Economics and Statistics*, 81(3), 351–368.
- Burkhauser, R. V., Hérault, N., Jenkins, S. P., & Wilkins, R. (2018). Survey under-coverage of top incomes and estimation of inequality: What is the role of the UK’s SPI adjustment? *Fiscal Studies*, 39(2), 213–240.
- Bönke, T., Corneo, G., & Lüthen, H. (2015). Lifetime earnings inequality in Germany. *Journal of Labor Economics*, 33(1), 171–208.
- Caspersen, E., & Metcalf, G. (1994). Is a value added tax regressive? Annual versus lifetime incidence measures. *National Tax Journal*, 47(4), 731–46.
- Castaneda, A., Diaz-Gimenez, J., & Rios-Rull, J.-V. (2003). Accounting for the US earnings and wealth inequality. *Journal of Political Economy*, 111(4), 818–857.
- Congressional Budget Office. (2015). The effects of potential cuts in SNAP spending on households with different amounts of income. CBO report 49978.
- Crawford, R. (2012). ELSA pension wealth derived variables: Methodology. UK data service SN5050, UK data service, Colchester.
- Creedy, J. (1999). *Lifetime versus annual income distribution. Handbook of income inequality measurement* (pp. 513–534). Berlin: Springer.
- Dahl, G. B., Kostøl, A. R., & Mogstad, M. (2014). Family welfare cultures. *The Quarterly Journal of Economics*, 129(4), 1711–1752.
- Davies, J., St-Hilaire, F., & Whalley, J. (1984). Some calculations of lifetime tax incidence. *The American Economic Review*, 74(4), 633–649.
- De Nardi, M., Fella, G., & Pardo, G. P. (2016). The implications of richer earnings dynamics for consumption and wealth. NBER working paper 21917, National Bureau of Economic Research, Boston.
- Disney, R., Emmerson, C., & Tetlow, G. (2009). What is a public sector pension worth? *The Economic Journal*, 119(541), F517–F535.
- Eissa, N., & Hoynes, H. (2011). Redistribution and tax expenditures: The earned income tax credit. *National Tax Journal*, 64(2), 689–729.
- Eissa, N., & Liebman, J. B. (1996). Labor supply response to the earned income tax credit. *The Quarterly Journal of Economics*, 111(2), 605–637.
- Evans, W. N., & Garthwaite, C. L. (2014). Giving mom a break: The impact of higher EITC payments on maternal health. *American Economic Journal: Economic Policy*, 6(2), 258–290.
- Falkingham, J., & Harding, A. (1996). *Poverty alleviation versus social insurance systems: A comparison of lifetime redistribution*. Amsterdam: North-Holland.
- Fullerton, D., & Rao, N. L. (2019). The lifecycle of the 47 percent. *National Tax Journal*, 72(2), 359–396.

- Guvenen, F., Karahan, F., Ozkan, S., & Song, J. (2015). What do data on millions of US workers reveal about life-cycle earnings risk? NBER working paper 20913, National Bureau of Economic Research, Boston.
- Haan, P., Kemptner, D., & Prowse, V. (2017). Insurance, redistribution, and the inequality of lifetime income. Technical report.
- HM Treasury. (2015). *Budget 2015*. London: Controller of HMSO.
- Hotz, V. J., & Scholz, J. K. The earned income tax credit. In: Moffitt, R. A. (ed). Means-tested transfer programs in the United States. Chicago, IL: National Bureau of Economic Research, University of Chicago, 2003.
- Hoynes, H., Schanzenbach, D. W., & Almond, D. (2016). Long-run impacts of childhood access to the safety net. *American Economic Review*, 106(4), 903–934.
- Joyce, R., & Sibieta, L. (2013). An assessment of labour's record on income inequality and poverty. *Oxford Review of Economic Policy*, 29(1), 178–202.
- Kopczuk, W., Saez, E., & Song, J. (2007, August). Uncovering the American dream: Inequality and mobility in social security earnings data since 1937. Technical report 13345, National Bureau of Economic Research, Inc.
- Kopczuk, W., Saez, E., & Song, J. (2010). Earnings inequality and mobility in the United States: Evidence from social security data since 1937. *The Quarterly Journal of Economics*, 125(1), 91–128.
- Landais, C., & Spinnewijn, J. (2019). The value of unemployment insurance. CEPR discussion paper No. DP13624, CEPR, London.
- Levell, P., Roantree, B., & Shaw, J. (2015). Redistribution from a lifetime perspective. IFS working paper W15/27, Institute for Fiscal Studies, London.
- Levell, P., & Shaw, J. (2016). Constructing full adult life-cycles from short panels. *International Journal of Microsimulation*, 9(2), 5–40.
- Lyon, A. B., & Schwab, R. M. (1995). Consumption taxes in a life-cycle framework: Are sin taxes regressive? *The Review of Economics and Statistics*, 77(3), 389–406.
- Metcalfe, G. E. (1999). A distributional analysis of green tax reforms. *National Tax Journal*, 52(4), 655–681.
- Neumark, D., & Shirley, P. (2017). The Long-Run Effects of the Earned Income Tax Credit on Women's Earnings. NBER Working Paper 24114, National Bureau of Economic Research, Boston.
- O'Dea, C. (2018). Insurance, efficiency and the design of public pensions. Working paper, Yale University Economics Department, New Haven, CT.
- Piketty, T., & Saez, E. (2013). *Optimal labor income taxation*. *Handbook of public economics* (pp. 391–474). Amsterdam: Elsevier.
- Poterba, J. M. (1989). Lifetime incidence and the distributional burden of excise taxes. *The American Economic Review*, 79(2), 325–330.
- Roantree, B., & Shaw, J. (2018). What a difference a day makes: Inequality and the tax and benefit system from a long-run perspective. *The Journal of Economic Inequality*, 16(1), 23–40.
- Sherman, A. (2009). *Safety net effective at fighting poverty but has weakened for the very poorest*. Washington, DC: Center on Budget and Policy Priorities.
- van de Ven, J. (2006). *The distributional effects of taxation in Australia and the United Kingdom: Evidence from microsimulations*. *Dynamics of inequality and poverty, volume 13 of research on economic inequality* (pp. 81–122). Bingley: Emerald Group Publishing Limited.
- Waters, T. (2017). TAXBEN: The IFS tax and benefit microsimulation model. User guide 15 Nov 2017, Institute for Fiscal Studies, London.