Bank of England

Modelling income risk dynamics in the UK: a parametric approach

Staff Working Paper No. 1,129
May 2025

Marco D'Amico and Martina Fazio

Bank of England

Staff Working Paper No. 1,129

Modelling income risk dynamics in the UK: a parametric approach

Marco D'Amico[(1)] and Martina Fazio[(2)]

Abstract

This paper uses rich, administrative-quality data on earnings in the UK from the Annual Survey of Hours and Earnings (ASHE) to provide a detailed analysis of income risk and its patterns across individuals and over time. We develop a model of income dynamics that accounts for the broader state of the economy and successfully captures key features of the UK earnings growth distribution, including: a cyclical variance, procyclical skewness (more frequent negative earnings shocks during recessions), and a distribution that combines sharp peaks with long, heavy tails. The model is simple enough to be integrated into broader macroeconomic frameworks, such as heterogeneous agent models, and could be used to support policy scenario analysis.

Key words: Idiosyncratic risk, income dynamics, inequality, income process.

JEL classification: C15, C63, D3, J01.

---

(1) Uppsala University. Email: marcodamico11@gmail.com
(2) Bank of England. Email: martina.fazio@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England or its committees. We are grateful to Stephen Burgess, Ken Deeley, Spyridon Lazarakis, Ricardo Reis, Marek Rojicek and Arthur Turrell for useful comments and suggestions. This project was started while Marco D'Amico was an intern in the Financial Stability Strategy and Risk directorate at the Bank of England, which he gratefully acknowledges for hospitality.

The Bank's working paper series can be found at www.bankofengland.co.uk/working-paper/staff-working-papers

Bank of England, Threadneedle Street, London, EC2R 8AH
Email: enquiries@bankofengland.co.uk

ISSN 1749-9135 (on-line)

1 Introduction

Understanding how income heterogeneity both shapes and is shaped by macroeconomic fluctuations is fundamental for designing effective policy interventions. While extensive research has shown that features of the income distribution alters the transmission of aggregate shocks through consumption, investment, and financial channels (e.g., Heathcote et al., 2009), there is growing evidence that business-cycle dynamics themselves materially reshape the distribution of labour income and its growth. An accurate characterization of this reverse channel is crucial: shifts in the tails of the income shock distribution can exacerbate crises Ravn and Sterk (2017), influence aggregate demand Amromin et al. (2018), and alter the efficacy of monetary policy Auclert (2019). Moreover, persistent negative income shocks are strong predictors of debt repayment difficulties Ganong and Noel (2023), and have motivated macroprudential Loan-To-Income and Debt-Service-Ratio limits, highlighting the need to accurately model the left tail of the income-risk distribution.

This paper addresses the question: How do business-cycle fluctuations affect the distribution of labour income growth in the UK, and can a simple, tractable parametric model replicate these dynamics? Building on previous work documenting income-risk cyclicality in the US, Europe and the UK (e.g., Guvenen et al., 2014; Busch et al., 2022; Bell et al., 2022), we document some key empirical regularities of the UK income growth distribution; we then estimate a parametric process that captures its most important features.

The process is sufficiently simple to be integrated in larger macroeconomic models, and has properties that make it suitable to be used for policy projections and scenario analysis. As an example, given a path of weak GDP and high unemployment, the model will produce a distribution where sharp income losses become more prevalent. Policymakers can use this information to assess the implications of changes in the distribution for other variables of interest, such as debt affordability, as well as for a deeper understanding of how monetary and macroprudential policies may transmit to the economy.

Another key benefit of our approach is the use of nearly fifty years of administrative-quality data, which we use to systematically document properties of the income growth distribution and its evolution over the business cycle. From this exercise, we select some key moments to serve as targets for a tractable income process. Unlike other survey sources, this dataset provides a very large cross-sectional panel, covering the period from 1975 to 2023. This is particularly important for detecting relevant patterns of cyclical variations, as well as to estimate distributional features relating to the tails, where small sample sizes would typically impede inference. The data comes from the Annual Survey of Hours and Earnings (ASHE), one of the most comprehensive sources of information on the structure and distribution of earnings in the UK.[1] It is based on a 1% representative sample of employee jobs taken from HM Revenue and Customs' (HMRC's) Pay As You Earn (PAYE) records, with around 170,000 annual observations. The information is not self-reported by employees, instead, it is a legal requirement for employers to complete the survey.

In line with previous findings, we confirm that the UK household income process has fat tails, with some degree of asymmetry, such that a normal distribution represents a poor approximation. The standard deviation of income shocks shows no relationship with the business cycle, while the skewness is pro-cyclical, becoming more negative during economic downturns. This implies that during a recession, the spread of income shocks remains approximately unchanged, while large negative income shocks become more likely. Based on these findings, we compute a set of statistics aimed at summarizing the income risk distribution for several earnings' growth horizons, in relation to its shape, tails, and time series dynamics. By pooling information from the earnings of both male and female employees in the population, our unit of observation aligns with a two-earner household.[2] We then use these statistics as targets in a Simulated Method of Moments (SMM) approach à la McFadden (1989), to estimate a simple parametric model that replicates key features of the data.

To model income dynamics in a way consistent with the data, while preserving tractability, we follow the approach of McKay (2017) and Guvenen et al. (2023). We start from a canonical "persistent plus transitory" model but implement two key changes. First, we replace the transitory Gaussian process with a non-employment shock, allowing for job losses and gains with persistent effects on individuals' income profiles. This helps both to capture the "scarring" effects of recessions (Davis et al., 2011), as well as to reproduce the high kurtosis and long and fat tails of earnings growth. Second, we employ a more versatile support for the persistent shock distribution, which is driven by aggregate factors. This allows the earnings distribution to change with the cycle, and in particular to replicate the procyclicality of income risk. The model identifies output growth and unemployment as important aggregate drivers of income cyclicality.

[1] See also the Global Repository of Income Dynamics (GRID) for more details.

[2] There may be patterns of risk-sharing within the household which we are unable to capture fully (see e.g. Ortigueira and Siassi, 2013). Appendix B discusses patterns and provide estimates for males and females separately.

Relation to the literature

The work of Storesletten et al. (2004) first proposed an income process with countercyclical variance, paving the way for research on the cyclicality of income risk. More recently, Guvenen et al. (2014) showed the importance of higher-order risk in US data, also confirmed in Swedish, French and German data (Busch et al., 2022), as well as in UK data (Angelopoulos et al., 2022; Bell et al., 2022). Our results on the properties of the income distributions align with these findings. But differently from the existing literature, this paper also provides estimates of a parametric income growth process. As such, the closest related studies are Guvenen et al. (2023) and McKay (2017), both focusing on US data. This paper uses UK data to estimate a process for household income risk, capturing its specific features and dynamics. For example, we find no evidence of unequal exposure to the business cycle across the income distribution in the UK: poorer households do not appear disproportionately more likely to experience negative shocks during a recession. This contrasts with US data, where such patterns are well-documented (Guvenen et al., 2014).

The model is tractable enough to be used in heterogenous agents (HA) settings, while allowing realistic income risk properties and dynamics. Different studies have highlighted the importance of specific features of the income risk distribution in the context of HA models. For example, countercyclical income risk induces countercyclical precautionary savings among households, amplifying consumption volatility (McKay, 2017; Challe and Ragot, 2016). It is also an important amplifier of business cycles in estimated medium-scale DSGE models (D'Amico, 2024; Bilbiie et al., 2022) and can magnify demand shocks (Auclert et al., 2024). These works however abstract from the high kurtosis of idiosyncratic income risk. Kaplan et al. (2018) showed that income risk and households heterogeneity affect the transmission mechanism of monetary policy, though they focus on an income process that yields a leptokurtic earnings distribution without cyclical variation. Bhandari et al. (2021) study optimal monetary and fiscal policies in a HA setting with a factor structure on idiosyncratic risk, abstracting from high kurtosis and procyclical skewness.

The rest of the paper is structured as follows. In Section 2, we describe our data, the measure of income we use and how we select our sample, as well as briefly explaining key characteristics of the income distribution. In Section 3 we describe our modelling choices and estimation procedure. Section 4 presents the results and Section 5 concludes.

2 Data

In order to analyse the properties of income risk in the UK, we employ the Annual Survey of Hours and Earnings (ASHE), providing administrative-quality data on the earnings of UK employees.[3] The data is based on a 1% sample of employee jobs from the HMRC's PAYE records, comprising approximately 300,000 annual observations. Importantly, the information is not self-reported, and it is a legal requirement for employers to complete the survey. ASHE data are collected every year in April, and they are not top-coded.

The ASHE survey in its current form started in 2004, but the New Earnings Survey (NES) was its predecessor.[4] Compared to the older version, the ASHE introduced some methodological improvements over the NES. These included the use of weighting for job representativeness and a new methodology to allow for job changes between sample selection and survey date.[5] Nevertheless, we choose to adopt the longest panel available, from 1975 to 2023. The availability of employees' National Insurance Numbers across waves of the survey allows researchers to track the same individual over time and as they change jobs.

A key limitation of the ASHE data relates to the inability of studying why people drop in and out of sample, as unemployment or self-employment spells are not identified in the survey. While these are clearly important, we are not aware of other sufficiently long and wide panel datasets that could help us overcome this obstacle. We refer to Bell et al. (2022) for a broader discussion of the data and its limitations.

2.1 Income measure and sample selection

We mostly focus on gross weekly earnings, as this is the earnings measure with the longest time series available (1975-2023). The definition used includes overtime pay, basic pay and other pay. We annualize this information based on 52 weeks of earnings. When this information is not available, we combine this with data from annual income, which affects just over 3% of the entire sample.[6]

[3] Office for National Statistics, released 11 May 2023, ONS SRS Metadata Catalogue, Annual Survey of Hours and Earnings Longitudinal dataset for the UK, https://doi.org/10.57906/nz2h-kc10.

[4] New Earnings Survey https://doi.org/10.57906/69bn-3853.

[5] See also Ma et al. (2006) for more details.

[6] Annual data is only available since 1998. However, this measure reflects only the time the an employee has spent with the current employer, which may, in some circumstances, be less than a full year. Similarly, the annualisation of weekly earnings may introduce measurement errors in case of job changes or unemployment spells.

We focus on male and female employees aged 22-70. We use an average of both males and females' earnings for our estimates, although males' income tends to be more cyclical. We provide estimates also for females and males separately in Appendix B. We trim the bottom 0.1% and the top 99.9% in order to remove outliers. In order to study income risk for those who have a strong labour market attachment, we remove those who earn less than seven times the minimum wage.[7] Minimum wage data for the UK are available from the Office for National Statistics (ONS) from 1998.[8] Before 1998, we follow the approach of Bell et al. (2022) and compute a pseudo minimum wage for 1975-1997 using median wage growth in those years and scaling the 1998 minimum wage by the respective cumulative median wage growth rate. We then deflate the measure of nominal earnings by the annual CPI (also available from the ONS).[9] We then compute 1-,3- and 5-years log-income changes. For 1-year earnings changes, we use earnings information one or two years prior to the year of reference, where the previous year is preferred where available. This leaves us with just over 122,000 observations on average per year. For 3-years earnings changes, we use earnings information three years prior to the year of reference, with just under 94,000 observations on average per year. For the 5-years growth rate, we use earnings information between four and seven years prior to the year of reference, giving us slightly less than 110,000 observations on average per year.[10]

2.2 Summary measures used to describe the earnings growth distribution.

We assess income risk using a mix of parametric and non-parametric moments. Below, we provide details on some key metrics.

To capture the symmetry of the earnings growth distribution, we use Kelley's skewness, which compares the distances from the median to the upper and lower deciles, normalized by the inter-decile range:

$$KS = \frac{(P90 - P50) - (P50 - P10)}{P90 - P10},$$

[7]This would correspond to a weekly salary of less than one day of work for a minimum-wage worker.

[8]Data can be accessed here.

[9]The CPI time series for 1988-2023 can be retrieved at this link, while the one for 1975-1988 at this link. The two series are spliced together, with 2015 as the year of reference.

[10]All these measures are expressed as growth rate over the entire reference period (adjusted for the period covered) rather than annualised.

where Pn is the n-th percentile of the income growth distribution. A positive value, as we observe in the data, indicates that the upper tail is wider than the lower tail, or equivalently that positive income shocks are more likely than negative ones.

To gain information about the tails of the distribution, we use the Excess Crow–Siddiqui kurtosis of income changes, computed as:

$$CSK = \left(\frac{P97.5 - P2.5}{P75 - P25} \right) - 2.91.$$

The first term is the Crow–Siddiqui kurtosis, the second its value for a normal distribution. An excess kurtosis of 0 therefore indicates similarity to a normal distribution in terms of tail heaviness. A positive value, as found in the data, indicates that the empirical earnings growth distribution has longer and thicker tails than a normal distribution.

We also show properties for the cross-sectional standard deviation of earnings changes, to measure dispersion around the mean.

2.3 Properties of earning growth distribution in the UK

The left panel of Figure 1 shows that the standard deviation of income shocks is mostly acyclical and constant over time, implying that the spread of shocks during recessions and expansions in the UK is more or less constant. The right panel of Figure 1 shows that the Kelley's skewness of the 1-year income growth rate distribution is positive on average and strongly procyclical, turning negative during crisis periods. On average, positive income shocks are more likely than negative ones, however this is reversed during periods of recessions.

Figure 2 illustrates the shape of the income growth distribution, examining the log-density of 1-year income growth rates between the years 1999-2000. A normal distribution with the same mean and standard deviation as the data is superimposed for comparison. The distribution is characterised by a high peak in the centre and by long and thick tails. Both tails are approximately linear, spanning between annual log growth rates of 1 and 3 on the right side, and between -1 and -3 on the left. This would correspond, respectively, to 3-fold and 20-fold increases in income and 63% and 95% decreases in income. The high peak in the centre and the long and thick tails suggest that the normal distribution is not a good approximation of the UK earnings risk distribution. This is also summarised for example by an excess Crow-Siddiqui kurtosis of 5.44.

The figure also suggests that the tails have asymmetric slopes: the left tail is slightly

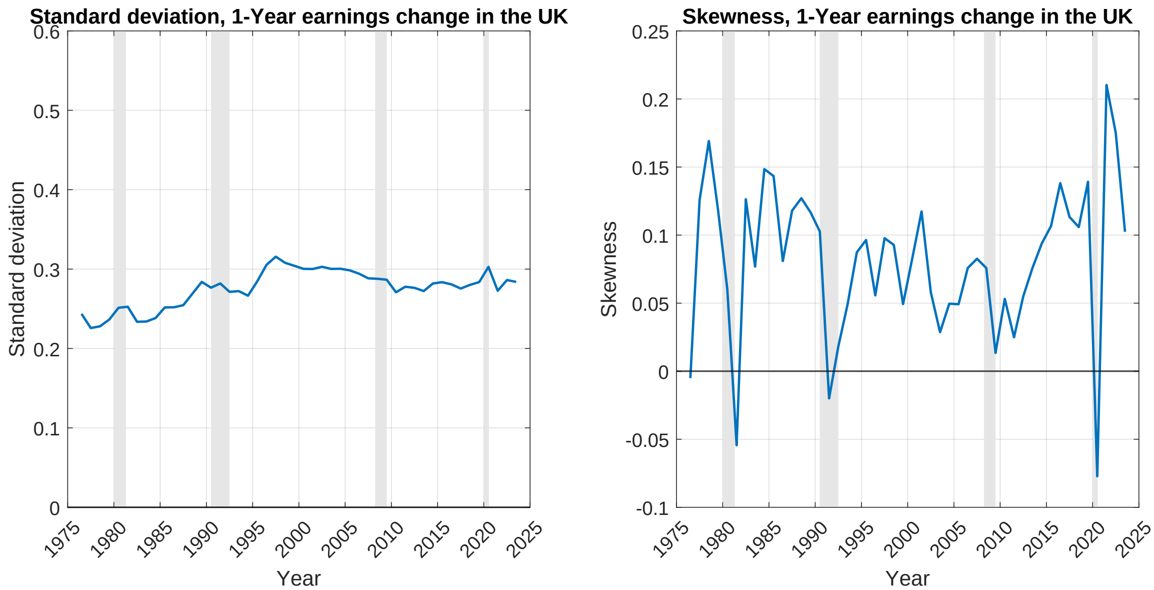

Figure 1: The figure shows the standard deviation (left) and Kelley's skewness (right) of 1-year log-income growth rates in the UK. Grey bars denote ONS recessions. The time period is 1976-2023. Source: ASHE and ONS.

flatter than the right, indicating that negative tail risk is more likely than positive tail risk. The distribution decreases more gradually on the left side compared to the right, so large income drops are more spread out (or occur across a broader range of values) than large income gains, which are more concentrated near the centre. However, this asymmetry is not pronounced, as the sample includes roughly half cases where the asymmetry is present and half where it is not.

To simplify the estimated process for integration into heterogeneous agent models or policy analysis, we abstract from certain cross-sectional features of the distribution,such as the age-variance profile. While these can be meaningfully captured in life-cycle models, they are not the focus of our analysis, and we do not attempt to reproduce them. For completeness, figure A1 in the appendix reports the cross-sectional variance of log-income for our main sample, as well as for samples of males and females separately.

To summarise, the acyclical standard deviation of income shocks indicates that their unpredictability remains largely unaffected by economic booms or recessions. On the other hand, the procyclical skewness of income shocks implies that the asymmetry of their distribution changes over the cycle, with negative shocks becoming more likely during recessions. A positive excess Crow-Siddiqui kurtosis shows that the income risk distribution has a higher peak and longer and fatter tails than a normal distribution. Finally, the earnings risk distribution features approximately linear asymmetric tails. In

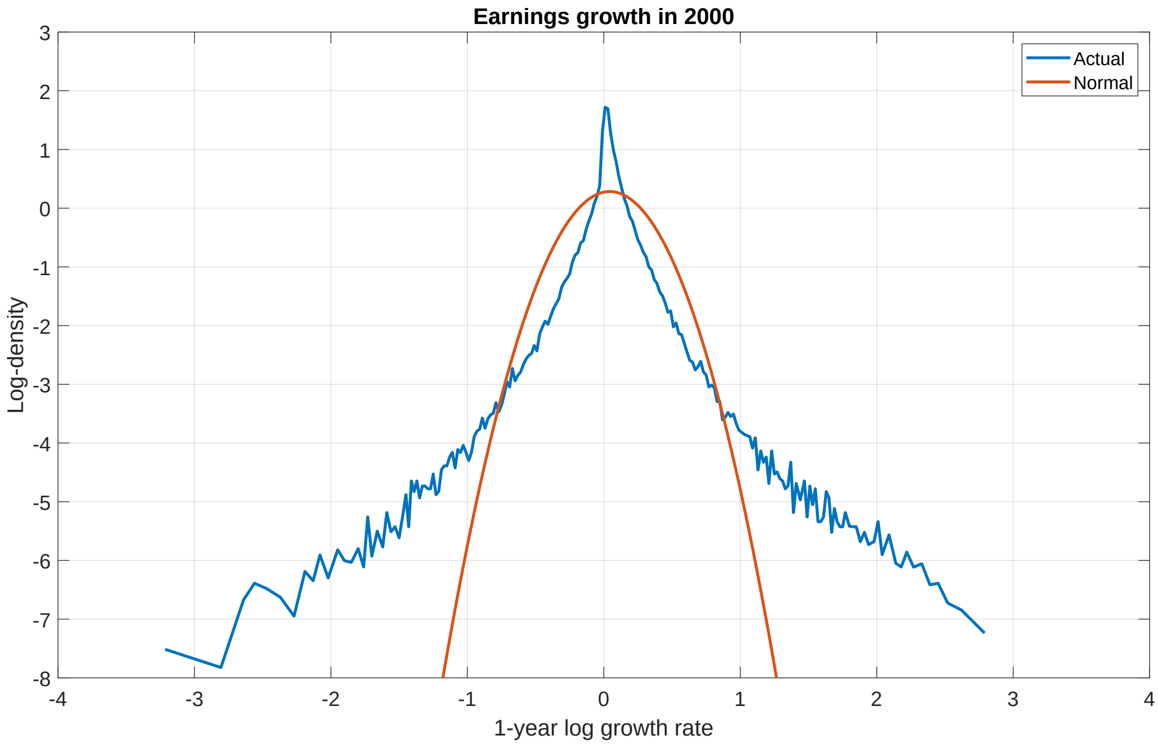

Figure 2: The figure shows the log-density of 1-year log income changes in 2000 (in blue), and a normal log-density with the same mean and standard deviation (in yellow).

the next section, we will set up a procedure to estimate a simple process to replicate all these features.

3 The Income Process

In this section, we describe the assumptions necessary to generate a stochastic income process that can reproduce the features of income growth observed in the data: (i) acyclical variance, (ii) procyclical skewness and (iii) leptokurtic log-density. In doing so, we follow the approach of McKay (2017) and Guvenen et al. (2023).

We adopt a flexible parametric approach that employs a mixture of different distributions to approximate the distribution of income changes. A parametric model is particularly useful for projecting incomes into the future, as it ensures regularity conditions even over long time horizons. Moreover, the model is relatively simple and features a limited number of state variables, such that it could be incorporated in larger macro heterogeneous agent models.

We present the most general model and then specify three different variants, switching on separate components. The process for the log of income for individual i at time t is

$$y_{i,t} = \gamma_i + \theta_{i,t} + (1-\psi)\,\xi_{i,t} + \mathbb{1}\alpha w_t \tag{1}$$

where γ_i is an individual fixed effect, normally distributed with zero mean and standard deviation σ^γ; $\xi_{i,t}$ is a transitory shock; $\theta_{i,t}$ is a persistent idiosyncratic state; w_t is the average wage growth; $\alpha > 0$ a scaling factor and $\mathbb{1}$ is an indicator function which takes value 1 in the baseline specification and 0 otherwise. The persistent state is defined as a random walk:[11]

$$\theta_{i,t} = \theta_{i,t-1} + \eta_{i,t} + \psi\xi_{i,t} \tag{2}$$

where $\theta_{i,0} = 0$ is the initial condition and $\eta_{i,t}$ is the innovation on the persistent component.

There are two features that make this process different from a classic "persistent plus transitory" specification. First, the transitory shock $\xi_{i,t}$ and the persistent component of the income process $\theta_{i,t}$ are correlated, with the degree of correlation governed by the parameter $\psi \in [0,1]$. Allowing for correlation between the transitory and persistent elements captures the idea of "scarring" effects from transitory shocks (Davis et al., 2011). Second, the support from which innovations to both $\xi_{i,t}$ and $\eta_{i,t}$ are drawn is non-standard. In particular, we assume $\xi_{i,t}$ follows a "non-employment" process as in Guvenen et al. (2023) where

$$\xi_{i,t} = \begin{cases} 0 & \text{with probability } p^\xi \\ (1-l_{i,t}) & \text{with probability } \left(1-p^\xi\right) \end{cases} \tag{3}$$

where $l_{i,t} \overset{i.i.d.}{\sim} \exp\left(\frac{1}{\lambda}, \frac{1}{\lambda^2}\right)$ and is normalized to be in the interval $[0,1]$ in every period. This specification allows us to mimic the effects of job-losses and gains without explicitly modelling job-market flows. Intuitively, when an individual is hit by the transitory shock (*e.g.* job-loss), her income is cut by a factor $1-l_{i,t}$ and generates a long-left tail of income risk. Equivalently, when the shock vanishes (*i.e.* job-gain), it generates a long right-tail of income risk. When there is correlation between the transitory shock and persistent process (*i.e.* $\psi > 0$), the job-loss leaves scars on the income of the affected individual, which will not return back to the same level when the shock vanishes, generating the steeper right

[11] We assume $\theta_{i,t}$ to encompass permanent income changes, and as such, could be defined as a "permanent" component. In the rest of the paper, we however maintain the more common language of "persistent" shocks. Assuming that persistent shocks do not vanish has the benefit of reducing the number of state variables when incorporating the process into broader macroeconomic settings, at the cost of potentially missing some life-cycle features of the income process. See also Guvenen et al. (2023) for a discussion.

tail of income growth. This allows the model to match the long and asymmetric tails of the income shock distribution.

We assume the persistent shock is drawn from a mixture of normal distributions:

$$\eta_{i,t} \sim \begin{cases} \mathcal{N}\left(\mu^{\eta}_{1,t}, \left(\sigma^{\eta}_{1}\right)^2\right) & \text{with probability } p^{\eta}_1 \\ \mathcal{N}\left(\mu^{\eta}_{2,t}, \left(\sigma^{\eta}_{2}\right)^2\right) & \text{with probability } p^{\eta}_2 \\ \mathcal{N}\left(\mu^{\eta}_{3,t}, \left(\sigma^{\eta}_{3}\right)^2\right) & \text{with probability } p^{\eta}_3 \end{cases}$$

where $p^{\eta}_1 + p^{\eta}_2 + p^{\eta}_3 = 1$. The means are driven by the latent variable x_t and thus change over time:

$$\begin{aligned} \mu^{\eta}_{1,t} &= \overline{\mu}^{\eta}_t \\ \mu^{\eta}_{2,t} &= \overline{\mu}^{\eta}_t + \mu^{\eta}_2 - x_t \\ \mu^{\eta}_{3,t} &= \overline{\mu}^{\eta}_t + \mu^{\eta}_3 - x_t \end{aligned}$$

where $\overline{\mu}^{\eta}_t$ ensures $\mathbb{E}\left[\exp\left(\eta_{i,t}\right)\right] = 1 \; \forall t$ implying that x_t does not affect mean income. We follow Catherine (2021) and Guvenen et al. (2023) and assume $x_t = \beta' \Delta \mathbf{C}_t$, where β is a factor loading on the growth rate of the aggregate cyclical variable $\Delta \mathbf{C}_t$ and captures the exposure of income risk to the aggregate cycle. Importantly, the aggregate cyclical variable $\mathbf{C}_t$ can either be a scalar or a vector, depending on the degrees of freedom required by the complexity of the underlying data. In our baseline specification, we focus on the case in which the aggregate cyclical component is a vector, namely $\mathbf{C}_t = [GDP_t; u_t]$, where GDP_t and u_t are, respectively, demeaned and detrended log output and changes in the stock of unemployed. This specification captures the countercyclicality of income risk (procyclical skewness) keeping the structure computationally simple.

The latent variable x_t is crucial in yielding procyclical skewness as it affects the behavior of the left and right tail of income shocks. In particular, focus for a moment on the case $C_t = GDP_t$. Suppose that $\beta < 0$ and, say, a recession hits, so that the aggregate cyclical variable displays negative growth $\Delta C_t < 0$; then both $\mu^{\eta}_{2,t}$ and $\mu^{\eta}_{3,t}$ decline, meaning that those who will draw from these distributions will face larger income risk (i.e. both left and right tails are shifted to the left).

Notice that the process in (1) collapses to the standard persistent plus transitory specification when $\psi = 0$, $\mathbb{1} = 0$, $\xi_{i,t} \overset{i.i.d.}{\sim} \mathcal{N}\left(0, \left(\sigma^{\xi}\right)^2\right)$ and $\eta_{i,t} \overset{i.i.d.}{\sim} \mathcal{N}\left(0, \left(\sigma^{\eta}\right)^2\right)$. We will refer to this model interchangeably as the "canonical" or "Gaussian" one. Also, from the structure of the latter, replacing the transitory shock $\xi_{i,t}$ drawn from the i.i.d. normal with the non-employment shock defined in (3), it is possible to get an intermediate model between the Gaussian and the full blown one, which we will refer to as the "non-employment" model.

3.1 Estimation

We use the Simulated Method of Moments (SMM) to estimate the model on UK data from 1977 to 2023. This involves repeatedly simulating the model under the above assumptions to minimize the distance between the observed data moments and those generated by the model. The model is annual and we simulate a panel of 300,000 individuals per year with mortality shocks at a rate of one every forty years. This allows for a finite cross-sectional variance of income despite the fact that innovations are permanent. When an individual dies, they are replaced by a newborn with a zero persistent shock ($i.e.$ $\theta = 0$).

The selected moments used as targets are:

1. The 10th, 50th and 90th percentiles of the distributions for 1-year, 3-years and 5-years earnings changes (9 moments), averaged over the entire time series.

2. The Excess Crow–Siddiqui kurtosis of 1-year and 5-years income growth (2 moments).

3. The mass of 1-year changes in log income in 2006 above 1.2 and below -1.2 (2 moments).

4. The slope of two lines fitted on the log density on the tails intervals [-4.0, -1.2] and [1.2, 4.0] (2 moments)

5. The Kelley's skewness for 1-year, 3-years and 5-years income, averaged over the entire time series, together with its standard deviation over the time series, and a measure of correlation between the skewness time series in the simulated model and in the data (9 moments).

The first two sets of moments help us target the shape of the income growth distribution, sets of moments (3) and (4) help us target the tails of the distribution, and finally the last set of moments help us match the cyclicality of income risk.

We compute the sum of squared differences between the simulated and data moments. We subtract from the objective function the correlation between the simulated Kelley's skewness time series and the data. We also add the Kolmogorov-Smirnov statistic as a measure of distance between the histogram generated by the model and from the data.[12] All moments are weighted equally, except for the tails' mass and the correlation between the simulated Kelley's skewness time series and data, so as to put them on a more equal scale with other moments.

Solution algorithm. We solve our model in Matlab via the SMM. For each simulation, we build the objective function as described above and minimize it with Matlab function patternsearch, subject to some regularity constraints on the parameters. We perform sensitivity checks to the initial starting point by building a grid of initial points, and choose the parameter vector that performs best out of 100.

4 Estimation results

We estimate three model specifications, starting from a simple "persistent plus transitory" income process and progressively adding features to assess how each improves the data fit. Estimated parameter values are reported in Table 2, while data and simulated moments are in Table 1. Figure 3 illustrates how the estimated model matches the cross-sectional income growth distribution in the data across different horizons. Figure 4 compares the time-series properties of skewness in the estimated model with those observed in the data.

[12] An alternative to targeting historical mean, volatility and correlation of skewness might be to target the time series of skewness year by year. We don't find any noticeable differences between these two approaches, so chose the former for simplicity. The correlation is subtracted from the sum of squared residuals so as as to maximise over the value of this indicator.

Table 1: Model simulation results

		Model specification		
Moments	Data	(1)	(2)	(3)
P10, 1-year change	-0.187	-0.403	-0.257	-0.230
P10, 3-years change	-0.242	-0.617	-0.208	-0.280
P10, 5-years change	-0.354	-0.766	-0.156	-0.322
P50, 1-year change	0.018	-0.002	0.029	0.025
P50, 3-years change	0.070	0.006	0.088	0.080
P50, 5-years change	0.121	0.022	0.150	0.144
P90, 1-year change	0.259	0.398	0.315	0.311
P90, 3-years change	0.430	0.650	0.386	0.491
P90, 5-years change	0.640	0.861	0.458	0.656
Mean wage, 1-year change	0.030	-0.002	0.029	0.030
Kurtosis, 1-year change	5.439	0.002	3.555	5.730
Kurtosis, 5-years change	2.898	0.303	2.386	3.001
Skewness, 1-year change	0.082	0.000	-0.000	0.058
Skewness, 3-years change	0.071	0.018	0.002	0.065
Skewness, 5-years change	0.043	0.031	0.002	0.045
Skewness, 1-year, correlation	-	-0.078	0.349	0.735
Skewness, 3-years, correlation	-	0.122	0.312	0.717
Skewness, 5-years, correlation	-	0.352	0.515	0.596
Left-tail mass	0.007	0.000	0.005	0.004
Right-tail mass	0.005	0.000	0.006	0.004
Left-tail slope	1.557	-	1.605	1.745
Right-tail slope	-1.741	-	-1.544	-1.810
Skewness st.dev. 1-year change	0.056	0.002	0.003	0.048
Skewness st.dev. 3-years change	0.055	0.008	0.003	0.057
Skewness st.dev. 5-years change	0.048	0.008	0.003	0.054
St. dev., 1-year change*	0.276	0.313	0.338	0.309
St. dev., 5-year change*	0.516	0.690	0.348	0.476
Objective value		22.940	9.977	-7.481

Notes: Variables not targeted in estimation are denoted with *. Model 1 refers to the standard Gaussian persistent plus transitory model. Model 2 refers to the non-employment specification while Model 3 refers to the baseline.

4.1 The canonical model

The canonical 'persistent plus transitory' model is well known for its parsimonious structure and ability to match standard moments, such as historical mean and standard deviation, as shown in Table 1. However, the simple model performs worse with regards to other features of the distribution such as skewness, kurtosis and tails of earnings growth. The canonical model generates zero skewness for 1-year earnings changes and slightly positive one at 3- and 5-years changes (respectively 0.018 and 0.031), against positive and slightly larger values in the data (Table 1). The kurtosis that emerges from the model, at 0.002 and 0.303, is also too low given data values of 4.578 and 2.566 for the 1- and 5-years change respectively. Looking at Figure 4, it is evident that the Gaussian model does not generate any cyclical variation in the skewness, while Figure 3 shows that the tails of the 1-year log change are short and thin, in stark contrast with those in the data. The time series of the standard deviation of earnings growth that emerges from the model is close to the data and acyclical.

4.2 Non-employment model

Replacing the transitory shock in the canonical model with a non-employment shock as in (3), allows the model to perform better in terms of Kurtosis and moments related to the tails of the distribution, as can be seen from Table 1 and Figure 3. However, the model is still not able to match the level and cyclicality of earnings growth skewness. Like the canonical model, the non-employment model is good at matching the standard deviation of earnings changes, replicating the acyclical pattern seen in the data.

4.3 Baseline model

Unlike the canonical and non-employment models, where the persistent shock is drawn from a normal distribution with constant mean, the baseline model features a persistent shock drawn from a mixture of normal distributions with time-varying means. These means are driven by an aggregate cyclical component – a combination of changes in the stock of unemployed and GDP growth in our case. This allows the model to match the average level of skewness, at 0.058, 0.065 and 0.045 in the model against 0.082, 0.071 and 0.043 in the data, for 1-, 3- and 5-years earnings changes respectively. Moreover, the skewness cyclicality and its standard deviation are also accurately reproduced, with the correlation between model's skewness time series and data at 0.735, 0.717 and 0.596 respectively for 1-, 3- and 5-years earnings changes. The baseline model is slightly less successful in matching the tails properties of the income growth distribution, but the strong peakedness of the data is captured appropriately. Indeed, it outperforms the non-employment model in terms of kurtosis, estimated at 5.730 and 3.001 against 5.440 and

2.898 in the data respectively for 1- and 5-years earnings changes. The model produces tails with slope 1.745 to the left, against a 1.557 in the data, and -1.810 to the right, against -1.741 in the data. The model's mass in both the left and right tails is 0.004, against 0.007 and 0.005 in the data. Overall, there appears to be a tension in the model's ability to match both the cyclical properties as well as all cross-sectional features of the data at the same time.

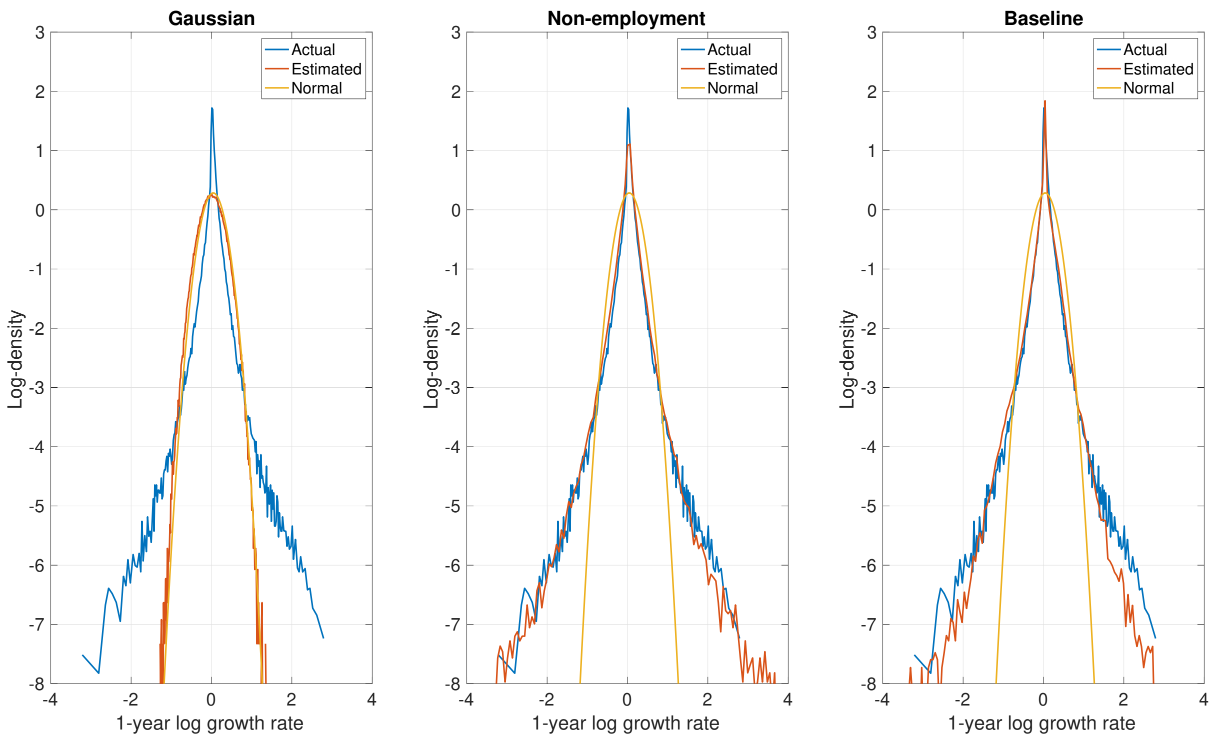

Figure 3: The figure shows data (solid, blue) and simulated 1-year log change income (log) distributions (solid, orange) for the three different model specifications in 2000, together with a normal distribution with the same mean and standard deviation.

4.4 Parameters of the estimated process

For the persistent process, most individuals $\left(p_1^\eta = 88\%\right)$ draw from a "normal-times" distribution with small shocks $\left(\sigma_1^\eta = 1\%\right)$. Skill-loss episodes are infrequent $p_2^\eta = 5.7\%$ but associated to large shocks $\left(\sigma_2^\eta = 55\%\right)$. On the other hand, skill-gain episodes are slightly more common $\left(p_3^\eta = 6.3\%\right)$ but are accompanied by relatively milder shocks $\left(\sigma_3^\eta = 12\%\right)$. The model predicts individuals in the UK to be in employment for a full-year with probability $p^\xi = 55\%$. The remaining 45% of individuals receive a transitory shocks ξ which drags income down to $15\% \simeq \frac{1}{\lambda}$ of its value, and $43\% = \psi$ of the shock persists

Figure 4: The figure show data (solid, blue) and simulated Kelley's skewness time series (solid, orange) for the three different model specifications. The first row refers to skewness of 1-year changes (1977-2020), the second row refers to skewness of 3-year changes (1979-2020) and the third row refers to skewness of 5-year changes (1981-2020). Grey bars denote ONS recessions. Source: ASHE and models' simulations.

after every year. The factor loading on the aggregate cyclical component GDP_t is negative, while the loading on changes in the stock of unemployed u_t is smaller in size and positive, suggesting that income risk rises during recessions.

Table 2: Estimated parameter values

		Model specification		
Parameter	Description	(1)	(2)	(3)
σ^{γ}	St. dev. fixed effects	4.360	7.417	2.497
σ^{ξ}	St. dev. transitory shock	0.131	-	-
λ	Transitory exponential parameter	-	6.132	6.625
ψ	Scarring effect of transitory shock	-	0.685	0.429
p^{ξ}	Probability of full year employment	-	0.374	0.543
p_2^{η}	Mix. prob. persistent innov. 2	-	-	0.057
p_3^{η}	Mix. prob. persistent innov. 3	-	-	0.063
σ_1^{η}	St. dev. persistent innov. 1	0.252	0.037	0.010
σ_2^{η}	St. dev. persistent innov. 2	-	-	0.550
σ_3^{η}	St. dev. persistent innov. 3	-	-	0.120
μ_2^{η}	Center for persistent component 2	-	-	-0.189
μ_3^{η}	Center for persistent component 3	-	-	0.212
α	Center of the distribution	-	-	1.375
β^{u}	Loading on unemployment	-	-	0.230
β^{gdp}	Loading on GDP	-	-	-1.537

Notes: Model 1 refers to the standard Gaussian persistent plus transitory model. Model 2 refers to the non-employment specification while Model 3 refers to the baseline.

4.5 Validation

The model is also successful in matching several empirical features of the income-risk distribution not targeted in estimation. Figure A2 in Appendix A shows that the model matches the time series of 1-year mean wage growth very well. Figures A4 and A5 present the log-density plot of 1-year log income changes for two additional non-targeted years, 2009 and 2012, demonstrating the model's ability to replicate the empirical distribution successfully in both years.

The model matches the time series of the 10th, 50th and 90th percentile of the distribution of 1-year, 3-years and 5-years income changes particularly well, although none of these are targeted in the estimation. In particular, Figure A3 first and last column, show the estimated and actual time series for, respectively, the 10th and 90th percentile the distribution of the 1-year, 3-years and 5-years income changes. The fit is remarkably good for both the 10th and 90th percentile, although the model slightly overestimates the volatility of the 10th percentile of growth at 3 and 5 years. The middle column of Figure A3 shows the same series for the median. The difference between actual and model's produced series is almost unnoticeable.

The model also closely matches the historical standard deviation of 1- and 5-year income growth, neither of which are targeted in estimation (Table 1). Similarly to the canonical and non-employment model, the baseline model also produces acyclical standard deviation for earnings growth, in line with the data (Figure A7 in Appendix A).

Even though the model correctly matches the historical kurtosis of income changes distribution at the considered changes-lengths, and to some extent its cyclicality, it fails to produce the observed upward trend observed in the data (Figure A6 in the Appendix A).[13] This was to be expected, given that there is no mechanism in the model to reproduce such pattern. At present, it is unclear whether this pattern will persist, and in fact it looks like a reversal of the trend started after 2011. As such, we have not included this among our targets. If further evidence confirms this as a long-lasting feature of the UK labour market, our model could be extended to capture this.

5 Conclusion

In this project, we estimated a tractable income process using almost 50 years of UK administrative-quality micro data. Our model accurately reproduces key time-series and cross-sectional properties of the observed distribution. Patterns in the data confirmed that income risk in the UK exhibits significant cyclicality, chiefly driven by negative shocks becoming more prevalent during recessions. This has important implications for how shocks transmit to the economy and for designing policies to moderate downturns effectively.

We hope this will enable future researchers to incorporate country-specific parameters that capture the cyclicality of UK income risk. This is especially relevant in the context of imperfect-insurance, heterogeneous-agent models, where the transmission of a variety of shocks is shaped by the link between macroeconomic and microeconomic variables.

This work should also prove valuable to policymakers seeking to project micro-level income distributions into the future starting from aggregate variables. Such projections

[13]See Bell et al. (2022) for a discussion of this trend, which does not feature in other countries' data.

could be used to support scenario analyses, stress testing exercises – where the evolution of household income is a key element, for example in the mortgage market – and the evaluation of cyclical labour market interventions.

By providing an estimated income process for the UK, this work facilitates quantitative studies of these topics and highlights key differences between the UK and other countries. Policymakers should consider these distinctions when developing strategies to smooth the business cycle.

References

Amromin, G., M. De Nardi, and K. Schulze (2018). Household inequality and the consumption response to aggregate real shocks. *Economic Perspectives 42*(1).

Angelopoulos, K., S. Lazarakis, and J. Malley (2022). Cyclical labour income risk in great britain. *Journal of Applied Econometrics 37*(1), 116–130.

Auclert, A. (2019, June). Monetary policy and the redistribution channel. *American Economic Review 109*(6), 2333–67.

Auclert, A., M. Rognlie, and L. Straub (2024, September). Fiscal and monetary policy with heterogeneous agents. *National Bureau of Economic Research* (32991).

Bell, B., N. Bloom, and J. Blundell (2022). Income dynamics in the united kingdom and the impact of the covid-19 recession. *Quantitative Economics 13*(4), 1849–1878.

Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent (2021). Inequality, business cycles, and monetary-fiscal policy. *Econometrica 89*(6), 2559–2599.

Bilbiie, F., G. E. Primiceri, and A. Tambalotti (2022). Inequality and business cycles. Technical report, Technical report, mimeo.

Busch, C., D. Domeij, F. Guvenen, and R. Madera (2022, April). Skewed idiosyncratic income risk over the business cycle: Sources and insurance. *American Economic Journal: Macroeconomics 14*(2), 207–42.

Catherine, S. (2021, 12). Countercyclical Labor Income Risk and Portfolio Choices over the Life Cycle. *The Review of Financial Studies 35*(9), 4016–4054.

Challe, E. and X. Ragot (2016). Precautionary saving over the business cycle. *The Economic Journal 126*(590), 135–164.

D'Amico, M. (2024). Inequality and the medium-term business cycle. *Working Paper*.

Davis, S. J., T. V. Wachter, R. E. Hall, and R. Rogerson (2011). Recessions and the costs of job loss [with comments and discussion]. *Brookings Papers on Economic Activity*, 1–72.

Ganong, P. and P. Noel (2023). Why do borrowers default on mortgages? *The Quarterly Journal of Economics 138*(2), 1001–1065.

Guvenen, F., G. Kaplan, and J. Song (2014, May). How risky are recessions for top earners? *American Economic Review 104*(5), 148–53.

Guvenen, F., A. McKay, and C. B. Ryan (2023). A tractable income process for business cycle analysis. Working Paper 30959, National Bureau of Economic Research.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

Heathcote, J., K. Storesletten, and G. L. Violante (2009). Quantitative macroeconomics with heterogeneous households. *Annual Review of Economics 1*(1), 319–354.

Kaplan, G., B. Moll, and G. L. Violante (2018, March). Monetary policy according to hank. *American Economic Review 108*(3), 697–743.

Ma, A. H., E. Roberts, R. F. Elliott, D. Bell, and A. Scott (2006). Comparing the new earnings survey (nes) and the labour force survey (lfs): an analysis of the differences between the data sets and their implications for the pattern of geographical pay in the uk. *Regional Studies 40*(6), 645–665.

McFadden, D. (1989). A method of simulated moments for estimation of discrete response models without numerical integration. *Econometrica: Journal of the Econometric Society*, 995–1026.

McKay, A. (2017). Time-varying idiosyncratic risk and aggregate consumption dynamics. *Journal of Monetary Economics 88*, 1–14.

Ortigueira, S. and N. Siassi (2013). How important is intra-household risk sharing for savings and labor supply? *Journal of Monetary Economics 60*(6), 650–666.

Ravn, M. O. and V. Sterk (2017). Job uncertainty and deep recessions. *Journal of Monetary Economics 90*, 125–141.

Storesletten, K., C. Telmer, and A. Yaron (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy 112*(3), 695–717.

Appendix

A Additional figures and tables

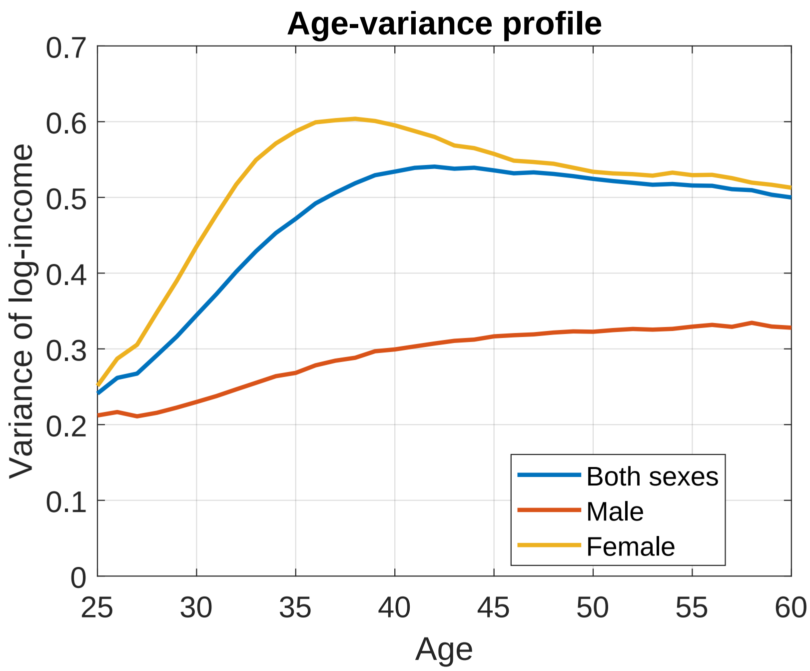

Figure A1: The figure shows the cross sectional variance of the log of income for both sexes combined (blue), males only (orange) and females only (yellow).

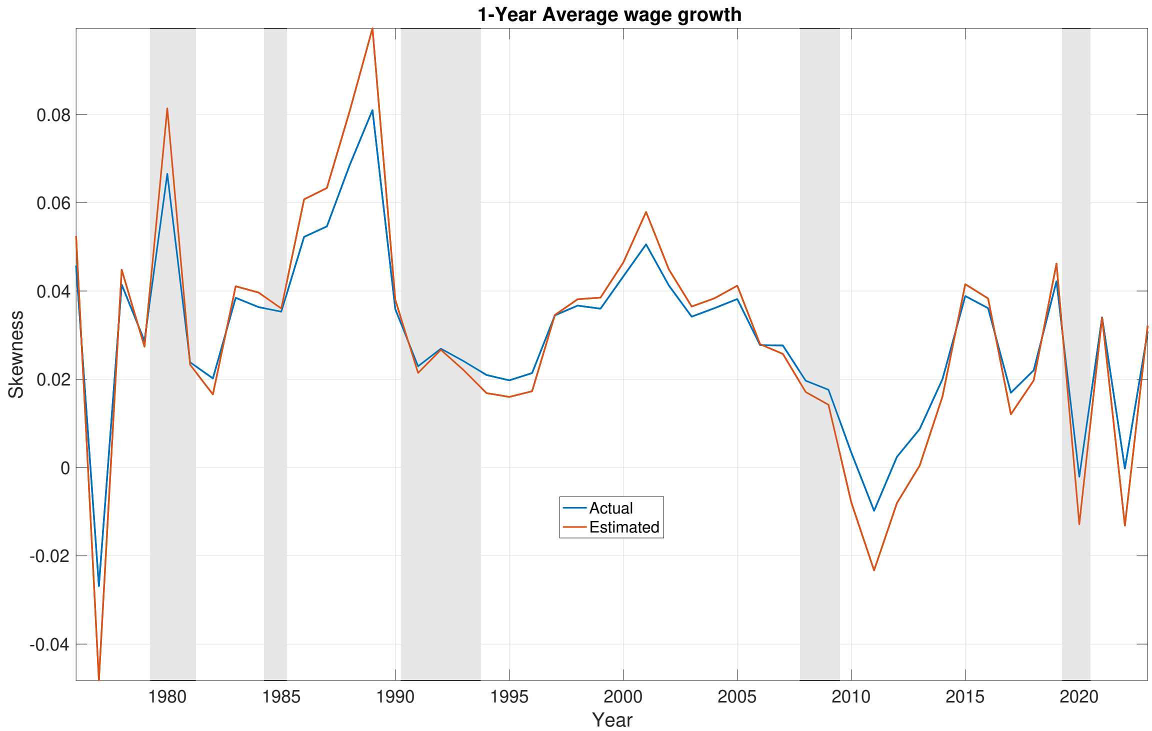

Figure A2: The figure shows data (solid, blue) and simulated 1-year average wage growth (solid, orange) for the baseline model.

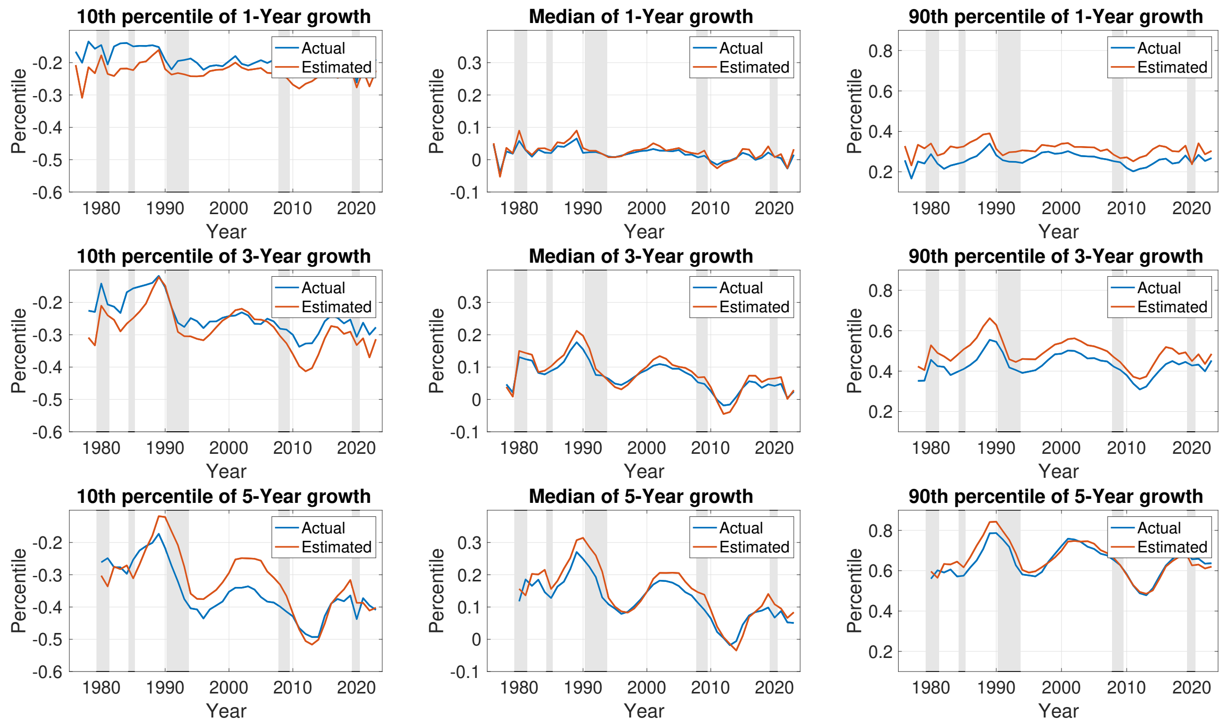

Figure A3: The figure shows the time series for the 10th, 50th and 90th percentile of 1-year, 3-years and 5-years income growth rates distribution in the UK. Grey bars denote ONS recessions. Source: ASHE and ONS.

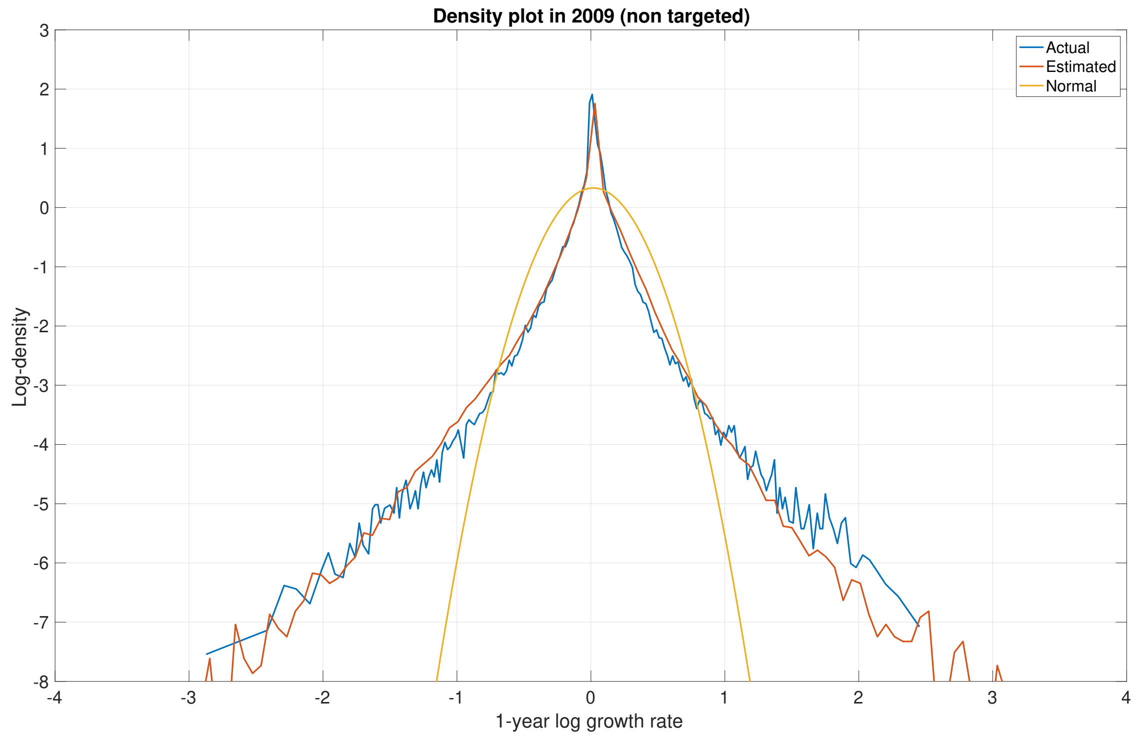

Figure A4: The figure shows data (solid, blue) and simulated 1-year log change income (log) distributions (solid, orange) for the baseline model in 2009, together with a normal distribution with the same mean and standard deviation.

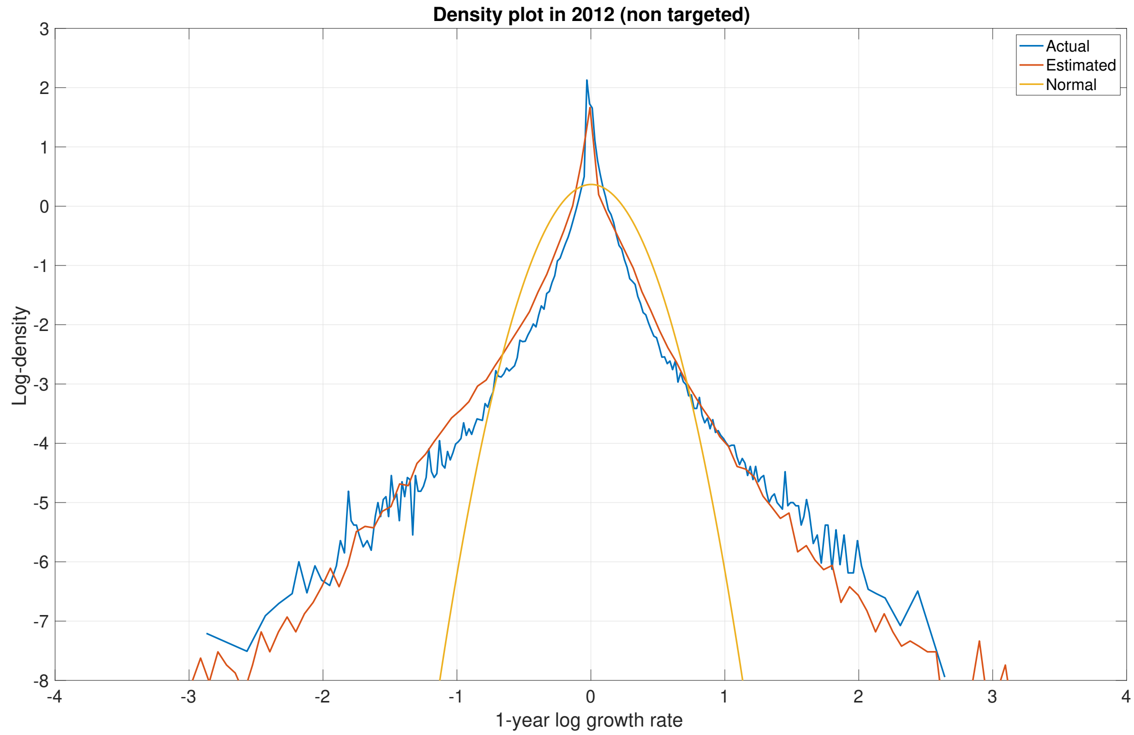

Figure A5: The figure shows data (solid, blue) and simulated 1-year log change income (log) distributions (solid, orange) for the baseline model in 2012, together with a normal distribution with the same mean and standard deviation.

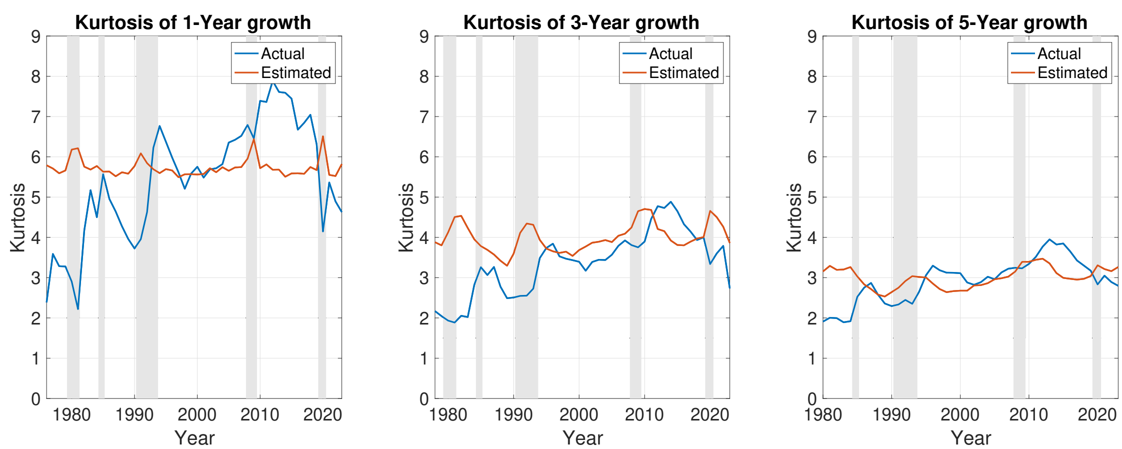

Figure A6: The figure shows data (solid, blue) and simulated Crow-Siddiqui Kurtosis of the 1-, 3- and 5-year log change income distribution (solid, orange) for the baseline model.

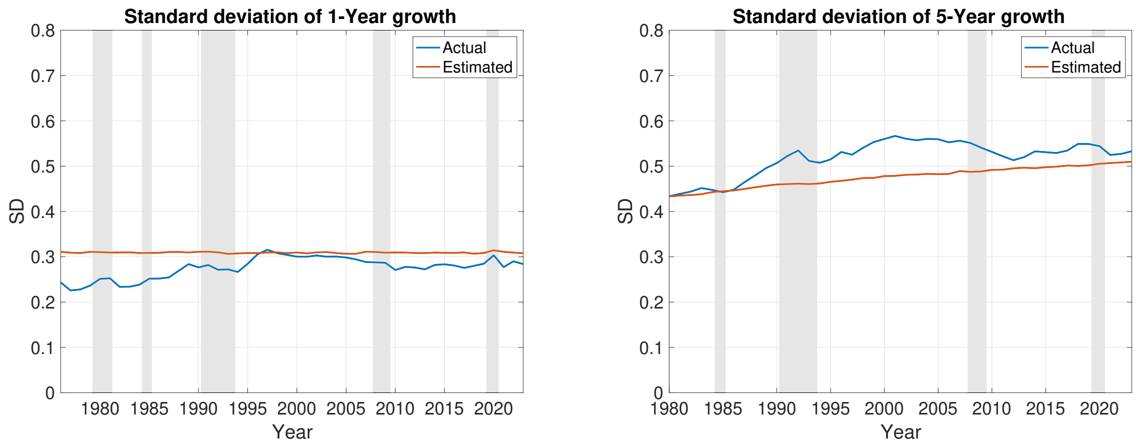

Figure A7: The figure shows data (solid, blue) and simulated Standard deviation of the 1- and 5-year log change income distribution (solid, orange) for the baseline model.

B Further results

We estimated all model (1) variations also for subsamples consisting of only females or males. Results are discussed below.

B.1 Females

In this section we discuss estimation results for females. Estimated parameter values as well as estimated moments are reported, respectively, in Tables 3 and 4.

The main difference in the data moments between the females and households subsamples lies in the shape of the income risk distribution and in its dynamics. In particular, females display much higher kurtosis (6.64) than households do (5.44) and the skewness of 1-year income changes appears to be less driven by the UK business cycle (Figure A8. This poses a difficulty for the baseline model, which is reflected in an overestimation of the kurtosis (7.316 in the model against a value of 6.638 in the data) and a poor correlation of the skewness time series (0.48, 0.55 and 0.55 respectively for 1-, 3- and 5-years income changes). The remaining estimated moments represent a relatively good fit to the data. However, the poor performance regarding the shape and dynamics yield a higher loss function (-4.833) than in the households sample (-7.481).

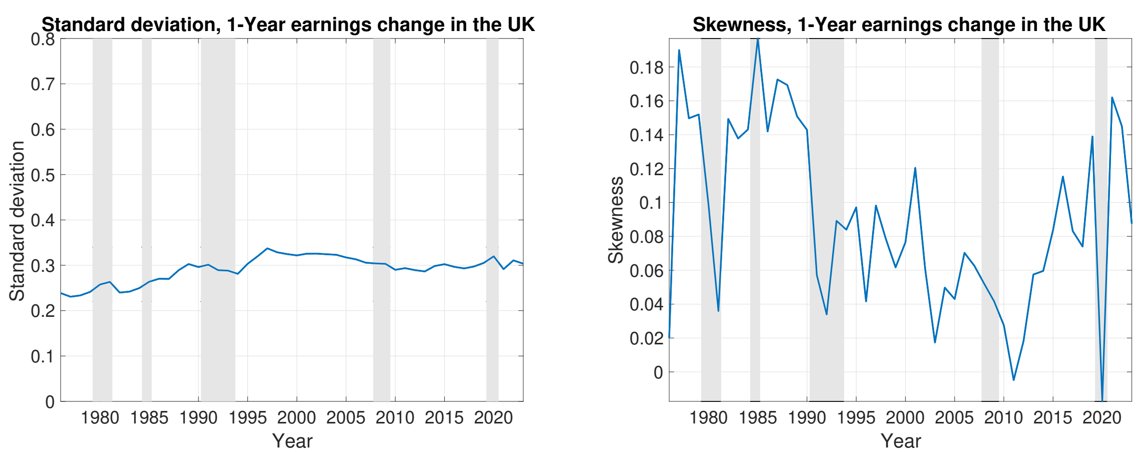

Figure A8: The figure shows the standard deviation (left) and Kelley's skewness (right) of 1-year income growth rates distribution in the UK for females. Grey bars denote ONS recessions. Time period is 1966-2023. Source: ASHE and ONS.

Regarding parameter values, the model predicts females in the UK to be in employment for a full-year with $p^{\xi} = 57\%$, close to the households value of 55%. The remaining 43% of individuals receive a transitory shock ξ which drags income down to $17\% \simeq \frac{1}{\lambda}$ of its value and 50% $= \psi$ of the shock persists every year. These values are close to estimates for households. In relation to the persistent process, most females $\left(p_1^\eta = 91\%\right)$ draw from a "normal-times" distribution with small shocks $\left(\sigma_1^\eta = 1\%\right)$, skill-loss episodes are rare $p_2^\eta = 5.8\%$ and associated to large shocks $\left(\sigma_2^\eta = 53\%\right)$, and skill-gain episodes are less common $\left(p_3^\eta = 3.6\%\right)$ and accompanied by tiny shocks $\left(\sigma_3^\eta = 3\%\right)$. The factor loading on the aggregate cyclical component GDP_t is negative, while the loading on changes in the stock of unemployed u_t is smaller in size and positive, suggesting that income risk rises during recessions. Compared to households value, females income appears to be much more driven by GDP_t $\left(\beta^{gdp} = -2.474\right)$ than by unemployment, which is almost irrelevant with a factor loading of $\beta^u = 0.118$.

This analysis suggests that women labor dynamics are not totally accounted for by our baseline model (1) and further research is needed in this area to model them.

Table 3: Estimated parameter values: females

		Model specification		
Parameter	Description	(1)	(2)	(3)
σ^γ	St. dev. fixed effects	0.555	8.904	1.169
σ^ξ	St. dev. transitory shock	0.167	-	-
λ	Transitory exponential parameter	-	5.895	5.844
ψ	Scarring effect of transitory shock	-	0.384	0.505
p^ξ	Probability of full year employment	-	0.374	0.574
p_2^η	Mix. prob. persistent innov. 2	-	-	0.058
p_3^η	Mix. prob. persistent innov. 3	-	-	0.036
σ_1^η	St. dev. persistent innov. 1	0.200	0.022	0.010
σ_2^η	St. dev. persistent innov. 2	-	-	0.534
σ_3^η	St. dev. persistent innov. 3	-	-	0.003
μ_2^η	Center for persistent component 2	-	-	-0.292
μ_3^η	Center for persistent component 3	-	-	0.434
α	Center of the distribution	-	-	1.438
β^u	Loading on unemployment	-	-	0.118
β^{gdp}	Loading on GDP	-	-	-2.474

Notes: Model 1 refers to the standard Gaussian persistent plus transitory model. Model 2 refers to the non-employment specification while Model 3 refers to the baseline.

Table 4: Model simulation results: females

		Model specification		
Moments	Data	(1)	(2)	(3)
P10, 1-year change	-0.189	-0.385	-0.265	-0.259
P10, 3-years change	-0.270	-0.505	-0.204	-0.329
P10, 5-years change	-0.420	-0.587	-0.139	-0.379
P50, 1-year change	0.020	0.012	0.032	0.030
P50, 3-years change	0.076	0.045	0.097	0.093
P50, 5-years change	0.132	0.084	0.165	0.162
P90, 1-year change	0.269	0.410	0.329	0.361
P90, 3-years change	0.473	0.610	0.398	0.590
P90, 5-years change	0.719	0.787	0.469	0.740
Mean wage, 1-year change	0.033	0.012	0.032	0.033
Kurtosis, 1-year change	6.638	0.003	3.821	7.316
Kurtosis, 5-years change	3.165	0.208	3.282	2.858
Skewness, 1-year change	0.090	0.001	-0.000	0.067
Skewness, 3-years change	0.071	0.013	0.001	0.081
Skewness, 5-years change	0.031	0.023	0.001	0.032
Skewness, 1-year, correlation	-	-0.086	0.294	0.480
Skewness, 3-years, correlation	-	-0.291	0.278	0.552
Skewness, 5-years, correlation	-	0.173	0.453	0.496
Left-tail mass	0.008	0.000	0.006	0.006
Right-tail mass	0.006	0.000	0.007	0.006
Left-tail slope	1.409	-	1.486	1.730
Right-tail slope	-1.746	-	-1.552	-1.664
Skewness st.dev. 1-year change	0.053	0.002	0.003	0.044
Skewness st.dev. 3-years change	0.056	0.007	0.003	0.056
Skewness st.dev. 5-years change	0.042	0.012	0.003	0.044
St. dev., 1-year change*	0.276	0.310	0.355	0.352
St. dev., 5-year change*	0.518	0.570	0.359	0.535
Objective value		22.249	11.111	-4.833

Notes: Variables not targeted in estimation are denoted with *. Model 1 refers to the standard Gaussian persistent plus transitory model. Model 2 refers to the non-employment specification while Model 3 refers to the baseline.

B.2 Males

In this section we discuss estimation results for males. Estimated parameter values as well as estimated moments are reported, respectively, in Tables 5 and 6.

Males earnings changes are characterised by lower kurtosis, thinner tails and more volatile skewness, with respect to households data. In particular, males kurtosis is 4.5 against 5.4 for households, the left tail slope is 1.003 against 1.5 for households and skewness volatility is 0.069 against 0.056 for households.

The baseline model is very successful in capturing these features, as reflected by the overall good fit (the loss function is -6.79). Kurtosis in the model is 4.6 against 4.5 in the data and skewness volatility is 0.071 in the model against 0.069 in the data. The matching of tails slopes is somewhat less accurate, with a left tail slope of 1.15 and a right one of -1.27 in the model against 1.003 and -1.12 in the data respectively for the left and right tails slopes.

This poses a difficulty for the baseline model, which is reflected in an overestimation of the kurtosis (7.316 in the model against a value of 6.638 in the data) and a poor correlation of the skewness time series (0.48, 0.55 and 0.55 respectively for 1-, 3- and 5-years income changes). The remaining estimated moments represent a relatively good fit to the data. However, the poor performance regarding the shape and dynamics yield a higher loss function (-4.833) than in the households sample (-7.481).

Regarding parameter values, the model predicts males in the UK to be in employment for a full-year with $p^{\xi} = 96\%$, far from the households value of 55%. The remaining 4% of individuals receive a transitory shock ξ which drags income down to $50\% \simeq \frac{1}{\lambda}$ of its value and $55\% = \psi$ of the shock persists every year. The size of the show is much larger than for estimated values for households, with similar persistence though. In relation to the persistent process, most males $\left(p_1^{\eta} = 82\%\right)$ draw from a "normal-times" distribution with small shocks $\left(\sigma_1^{\eta} = 7\%\right)$, skill-loss episodes are rare $p_2^{\eta} = 6\%$ and associated to large shocks $\left(\sigma_2^{\eta} = 26\%\right)$, and skill-gain episodes are more common $\left(p_3^{\eta} = 12\%\right)$ and accompanied by big shocks $\left(\sigma_3^{\eta} = 21\%\right)$. The factor loading on the aggregate cyclical component GDP_t is negative, while the loading on changes in the stock of unemployed u_t is smaller in size and positive, suggesting that income risk rises during recessions.

This analysis suggests that men labor dynamics are well accounted for by our baseline model (1).

Table 5: Estimated parameter values: males

		Model specification		
Parameter	Description	(1)	(2)	(3)
σ^{γ}	St. dev. fixed effects	2.606	8.028	4.423
σ^{ξ}	St. dev. transitory shock	0.033	-	-
λ	Transitory exponential parameter	-	1.703	2.027
ψ	Scarring effect of transitory shock	-	0.067	0.554
p^{ξ}	Probability of full year employment	-	0.966	0.965
p_2^{η}	Mix. prob. persistent innov. 2	-	-	0.066
p_3^{η}	Mix. prob. persistent innov. 3	-	-	0.118
σ_1^{η}	St. dev. persistent innov. 1	0.151	0.123	0.073
σ_2^{η}	St. dev. persistent innov. 2	-	-	0.261
σ_3^{η}	St. dev. persistent innov. 3	-	-	0.210
μ_2^{η}	Center for persistent component 2	-	-	-0.251
μ_3^{η}	Center for persistent component 3	-	-	0.196
α	Center of the distribution	-	-	1.336
β^{u}	Loading on unemployment	-	-	0.329
β^{gdp}	Loading on GDP	-	-	-0.420

Notes: Model 1 refers to the standard Gaussian persistent plus transitory model. Model 2 refers to the non-employment specification while Model 3 refers to the baseline.

Comparison with the US. The above patterns have been extensively studied in the US (Guvenen et al., 2014); it is thus useful to summarise the main differences between UK and US data. As far as the standard deviation of shocks is concerned, in both countries it is acyclical, without noticeable trends, with the UK featuring a value that is half as big as it is in the US. The skewness instead is on average positive in the UK while it is negative in the US, with both countries displaying the similar cyclicalities in terms of uncertainty variation. Regarding the shape of the distribution, the UK displays far less peakedness than the US do (kurtosis is 5.4 in the UK against 20 in the US), and the tails are much more symmetric in the UK than in the US (the left tail slope in the UK is 90% of the right one while in the US this number is 60%). Also, in the UK the domain of income changes is much different: a 2% income drop is half as likely in the UK than it is in the US.

Table 6: Model simulation results: males

		Model specification		
Moments	Data	(1)	(2)	(3)
P10, 1-year change	-0.182	-0.187	-0.152	-0.115
P10, 3-years change	-0.219	-0.297	-0.238	-0.248
P10, 5-years change	-0.305	-0.360	-0.279	-0.310
P50, 1-year change	0.017	0.016	0.020	0.021
P50, 3-years change	0.066	0.054	0.064	0.076
P50, 5-years change	0.113	0.096	0.112	0.138
P90, 1-year change	0.250	0.219	0.192	0.186
P90, 3-years change	0.397	0.413	0.373	0.453
P90, 5-years change	0.579	0.570	0.517	0.632
Mean wage, 1-year change	0.028	0.016	0.020	0.025
Kurtosis, 1-year change	4.541	0.002	0.573	4.640
Kurtosis, 5-years change	2.558	0.223	0.543	1.309
Skewness, 1-year change	0.082	-0.000	-0.000	0.094
Skewness, 3-years change	0.076	0.011	0.011	0.076
Skewness, 5-years change	0.053	0.019	0.018	0.047
Skewness, 1-year, correlation	-	0.131	0.259	0.716
Skewness, 3-years, correlation	-	0.293	0.362	0.731
Skewness, 5-years, correlation	-	0.506	0.589	0.645
Left-tail mass	0.005	0.000	0.005	0.004
Right-tail mass	0.004	0.000	0.005	0.005
Left-tail slope	1.003	-	1.142	1.146
Right-tail slope	-1.126	-	-1.089	-1.270
Skewness st.dev. 1-year change	0.069	0.002	0.002	0.071
Skewness st.dev. 3-years change	0.067	0.006	0.005	0.070
Skewness st.dev. 5-years change	0.060	0.010	0.009	0.049
St. dev., 1-year change*	0.276	0.158	0.283	0.284
St. dev., 5-year change*	0.518	0.390	0.404	0.457
Objective value		32.241	11.953	-6.789

Notes: Variables not targeted in estimation are denoted with *. Model 1 refers to the standard Gaussian persistent plus transitory model. Model 2 refers to the non-employment specification while Model 3 refers to the baseline.