Online Appendix

A Data and features of earnings and wages

A.1 The British Household Panel Survey (BHPS)

Starting in 1991 (and continuing until 2010, when it was discontinued to be included within the wider Understanding Society survey), the BHPS sampled 5,500 households and 10,300 individuals, which were then followed over time, hence generating a long panel. If an individual in the initial sample separated from his/her original household, all members of his/her new household were also interviewed. Children were interviewed once they reached the age of 16. These features imply that this survey should remain representative of the UK population.

As most household surveys, it has the limitation that all answers are self-reported and thus potentially subject to measurement error. However, the design of the survey suggests that measurement error in earnings is likely to be lower than in other surveys, such as the PSID in the US, because instead of just being asked about their total labor earnings in the last twelve months, respondents were asked to check their last pay slip and report about it. Furthermore, in a relevant proportion of the observations (around 30%), the interviewer saw the pay slip.

A.2 Sample and variable construction

We include individuals between 25 and 60 years of age. Given that our focus is on labour market income dynamics, we drop individuals from the sample if they have income from self-employment. We deflate earnings and wages with the CPI (2015=100).

We reconstruct age whenever the change of date in the interview implies that the individual is reported to be the same age in two consecutive years, but only when reported age does not differ by more than one year from one's expected age.

We use this broader sample to compute all of our measures of labour market participation. For the estimation of the earnings processes, the correlations of wages between members of a couple, etc. we impose further sample selection criteria related to the availability of earnings and wage data as follows.

Men's earnings. For men, our main variable of interest is total annual earnings. We construct this measure by adding up earnings for all jobs held over the past year (1st September to 31st August) for each worker. We do so by using the information available on the start and end months of all employment spells, together with the "usual" payment per unit of time. We exclude jobs that were held for a period shorter than a calendar month because many individuals do not report the exact day when the employment or unemployment spell began. We also drop observations in which the respondent does not report their usual payment per unit of time and self-employed men, retirees, full-time students, and the long-term disabled.

After excluding these and using the complete BHPS sample period (1991-2009), we have 42,659 person-year observations for men's earnings. Typically, the literature on earnings dynamics (e.g. Kaplan, 2012; Guvenen et al., 2021) further excludes observations below some minimum threshold, that is those below 5% of yearly median earnings, or £1,300 a year in our data. There are 2,259 (5.2%) male-earnings observation below such threshold in our dataset, of which 2071 (4.8%) display earnings which are exactly zero. The vast majority of individuals in the latter group report being unemployed. Rather than excluding these observations, we bottom-code them to £1,300 and check for the robustness of our results to changing this threshold.

We make this choice because for our question it is important to include the most unfavorable earnings outcomes, such as staying out of work for a long time, for which government insurance is likely particularly valuable. However, the Arellano, Blundell, and Bonhomme (2017) procedure that separates persistent and transitory earnings and estimates their rich dynamics, requires taking logs of earnings. Bottom coding allows for the inclusion of all observations. Although the choice of a lower bound is somewhat arbitrary, our bottom coding is low enough (around £100 per month) to capture the really high marginal utility of consumption in this situation, and yet reflects other sources of insurance which are likely under-reported, such as help from family and friends, private charities, informal work, and so on.

Women's wages. For women, our main variable of interest is their hourly wage. For simplicity, we focus on the current job being held by the individual rather than an annual average. As we describe in Section 2.1, we focus on potential rather than observed wages. Therefore, we keep in our sample both the women who are currently working and those who are not, as long as we can impute a wage to the latter, for which we require that we observe them working at least once in the sample. We eliminate outliers that most likely reflect recording errors and missing values and drop individuals whose total working hours exceed 80 hours per week and those who have negative earnings or wages. Because wages are computed by dividing earnings by hours worked, we eliminate extreme changes ($|\log w_t - \log w_{t-1}| > 2$) that likely are errors in recording hours of work. After excluding all of these, we have 58,116 observations for women, of which 43,198 correspond to women for which we observe positive hours worked, and thus wages, and the remainder correspond to women for whom we can impute a wage.

While we estimate our wage process on imputed wages, the statistics that we report for female wages in the BHPS (for example, those in Section 2 or Appendix A.4) refer to female wages for labor market participants.

For both men and women, when we apply the Arellano, Blundell, and Bonhomme (2017) we increase our sample size by performing a rolling-sample transformation similar to that used by De Nardi, Fella, and Paz-Pardo (2020) for the PSID data.

We decompose potential wages into a deterministic age-efficiency profiles η_t^{gp} which varies by gender g and marital status p and a stochastic residual component. To estimate the profiles η_t^{gp} more precisely, we expand our sample to include the Understanding Society survey (2010-2016). We report the resulting profiles in Appendix C.4.

A.3 Comparing the BHPS and NESPD data

A.3.1 The New Earnings Survey Panel Dataset (NESPD)

A National Insurance Number (NIN) number is randomly issued to all UK residents at age 16 and kept constant throughout one's lifetime. Individuals whose NIN ends in a certain set of digits are automatically selected for the NESPD sample. Its data is available for the years between 1975 and 2015. Every April, all employers whose employees qualify for the sample receive a form (currently online, although it was on paper in the early years of the sample) where they must provide payroll data about those employees.

This implies that, for individuals included in the survey, the NESPD contains complete information on their working life from the first year they started working (or 1975) until retirement age (or 2015), for all years during which the individual was working with the last recorded employer in April *and* the employer returned the questionnaire.

The most important limitation of the NESPD is that it is has a 25-30% employer non-response rate, implying that it only gathers 0.7% of all UK workers rather than 1%. Moreover, valid responses fell from 75% in the 1980s to 60% in 2012 (Adam, Phillips, Roantree 2016). This generates two main problems. First, endogenous non-responses might affect the randomness of the sample. Second, we cannot distinguish individuals who are not working from individuals whose employers do not respond to the survey.

As a result, the population covered by the BHPS' earnings measure is more comprehensive than that covered by NESPD. This is due to the fact that the latter is filled by the employer, so individuals who happen to be unemployed or out of the labor force in the week of reference will not appear in the sample. In contrast, the BHPS, being a household survey, can capture people who are non-employed but have worked at some point during the previous year.

A.3.2 Sample selection in the NESPD

We drop cases for which there are two records with the same identifier (year pair), as well as individuals whose hours worked or weekly pay are missing, or for whom age evolves unexpectedly, which can reflect, in the case of the NESPD, errors in recording NINs (as stated in the documentation for the data).

We apply the same transformation and sample selection criteria to the NESPD data as those for our main BHPS sample that we have described in Appendix A.2, with three main differences, which are motivated by the characteristics of the NESPD. For the purposes of this comparison, we apply the same screens to the BHPS data.

First, the NESPD only considers the highest-paid job for each individual, for which a direct measure of "annual earnings" is reported (earnings go from April 7th to April 6th, consistent with the tax year in the UK). Thus, for our comparisons with the NESPD, we also only keep the highest-paid job in the BHPS.

Second, we only consider men who have received at least 5% of median earnings (around £1,300 (2015)) in the year up to the moment when they are observed. This choice is motivated by the fact that the NESPD does not capture workers who spend all year out of the labor force.

Third, we use data from 1996 to 2006 because of three considerations. First, annual earnings only start being available in the NESPD after 1996. Second, up to the mid-90s there were many changes in the UK labor market (e.g. de-unionization) that could confound the analysis. Third, in the years 2007 and 2008 the New Earnings Survey suffered a budget cut that implied non-random attrition of part of the sample (those in smaller businesses which were still filling paper-based forms), and this was immediately followed by the financial crisis, whose specific effects are not the object of our study.

A.3.3 Comparison

Figures 13 and 14 show that the key implications of the BHPS and NESPD data are very similar. The most salient difference is that average persistence is higher in the NESPD, which could reflect the presence of larger measurement error in the BHPS. Luckily, the econometric procedure we use and describe in Section 2.2 separately identifies the persistent and transitory components of earnings and wage changes. Hence, whenever present, measurement error is mostly captured by the transitory component, which we do not include in our structural model.

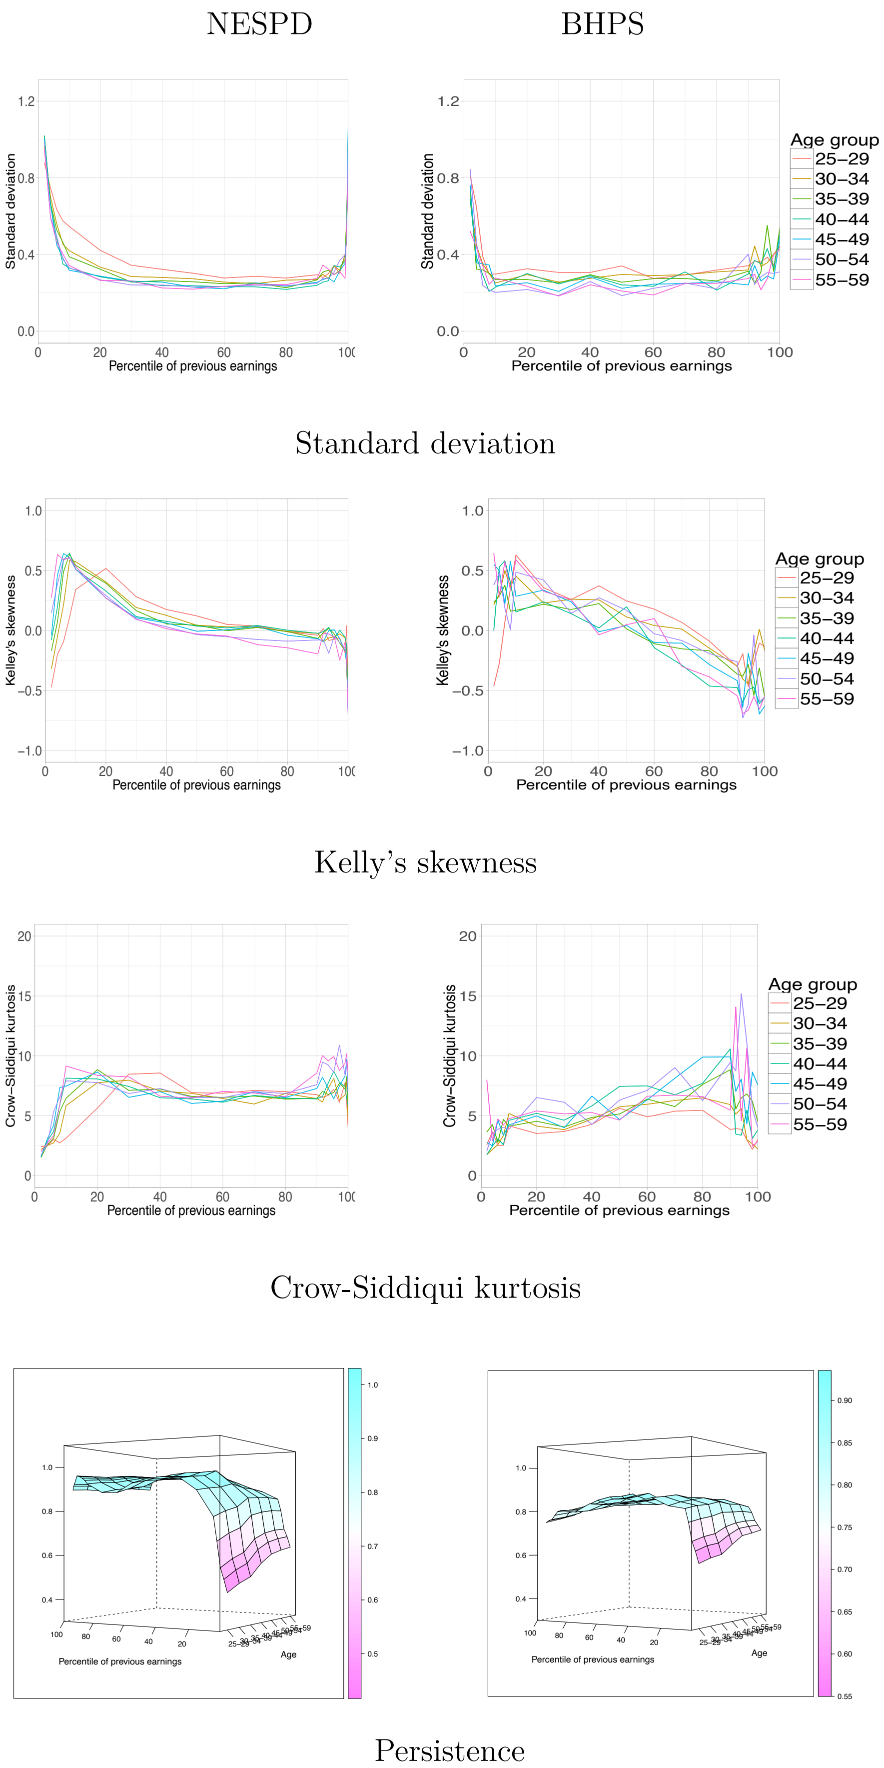

Figure 13: Moments of male earnings changes in the BHPS and NESPD. Top three panels: by previous earnings. Bottom panel, by previous earnings and age.

NESPD BHPS

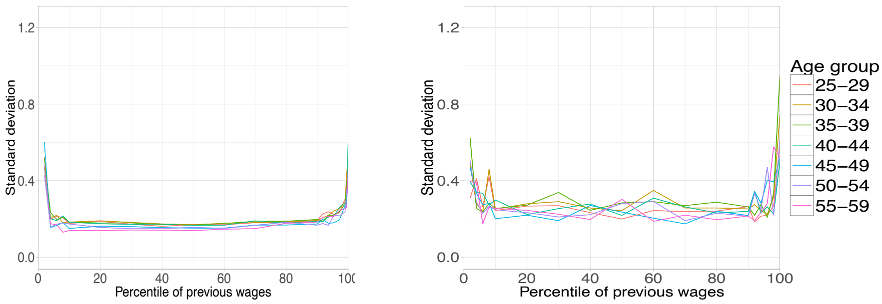

Standard deviation

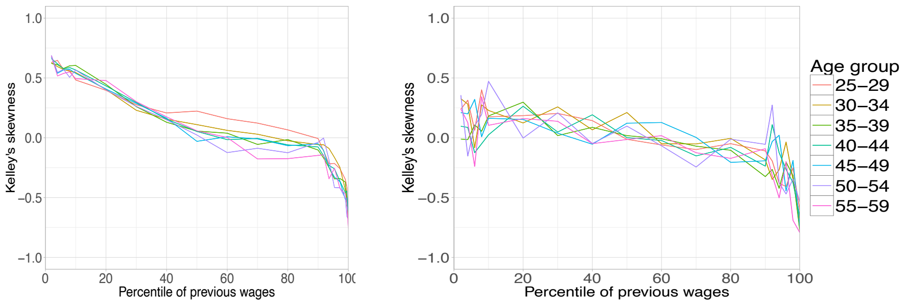

Kelly's skewness

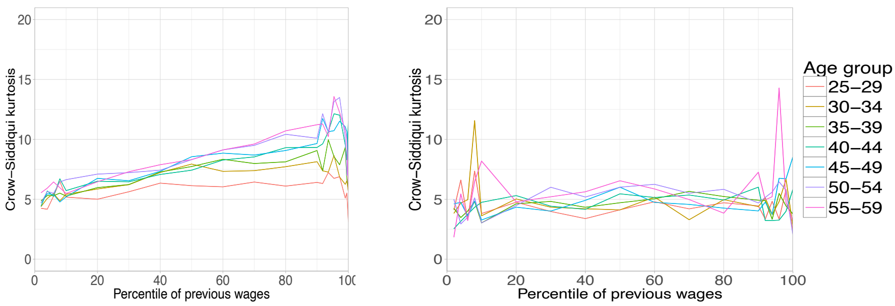

Crow-Siddiqui kurtosis

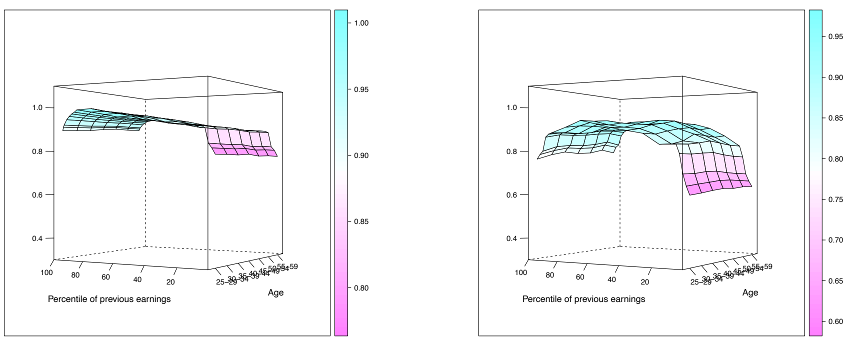

Persistence

Figure 14: Moments of female wage changes in the BHPS and NESPD. Top three panels: by previous earnings. Bottom panel, by previous earnings and age.

A.4 Comparing the earnings dynamics of singles and married

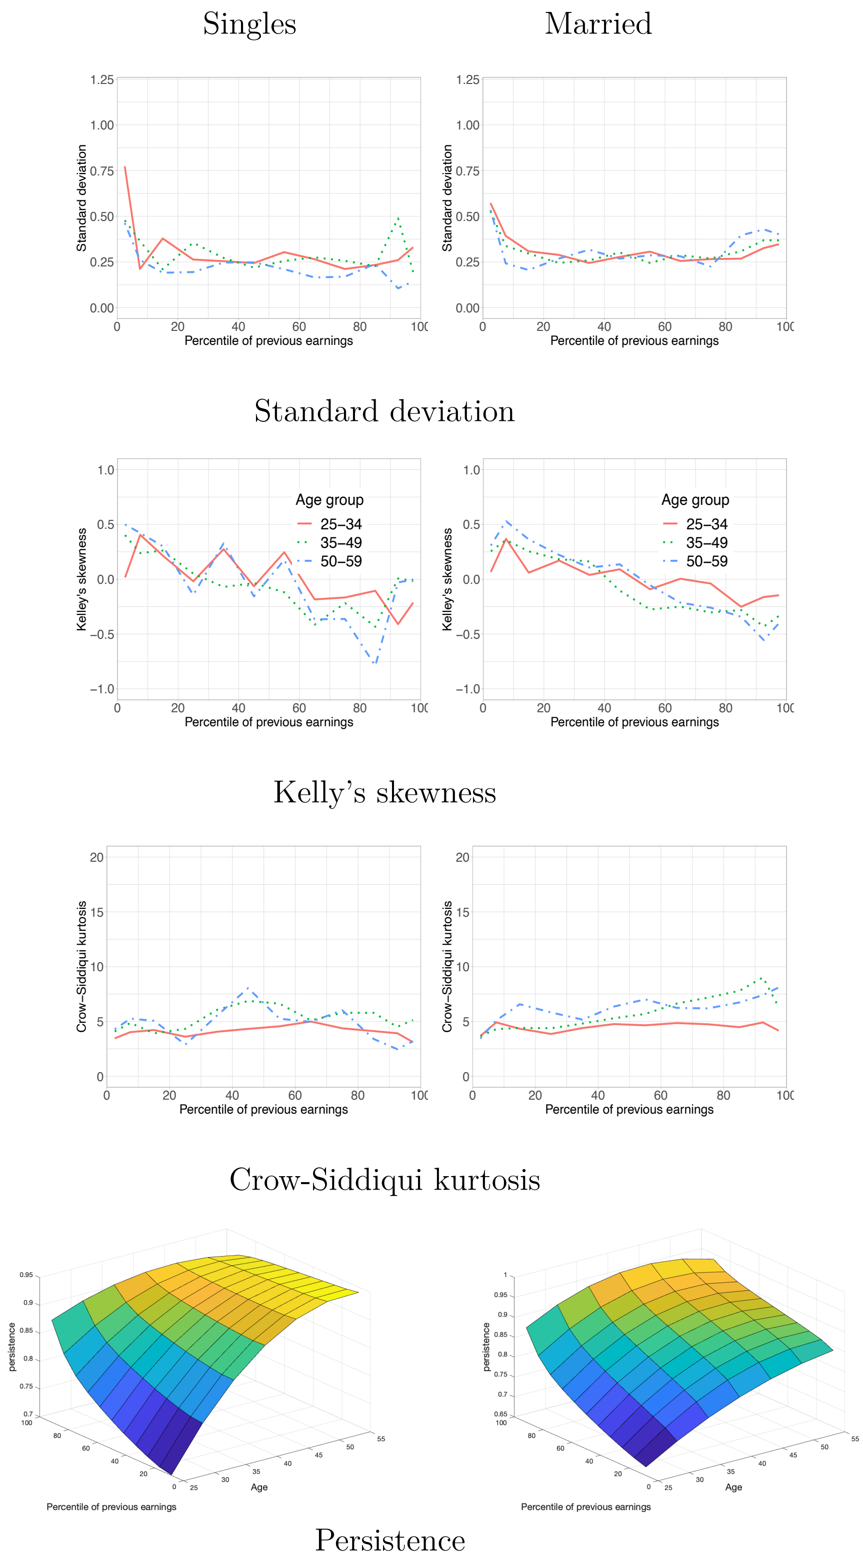

Figure 15: Moments of male earnings changes in the BHPS by marital status. Singles/married are defined as those observed single/cohabiting in both t and t+1. Top three panels: by previous earnings. Bottom panel, by previous earnings and age.

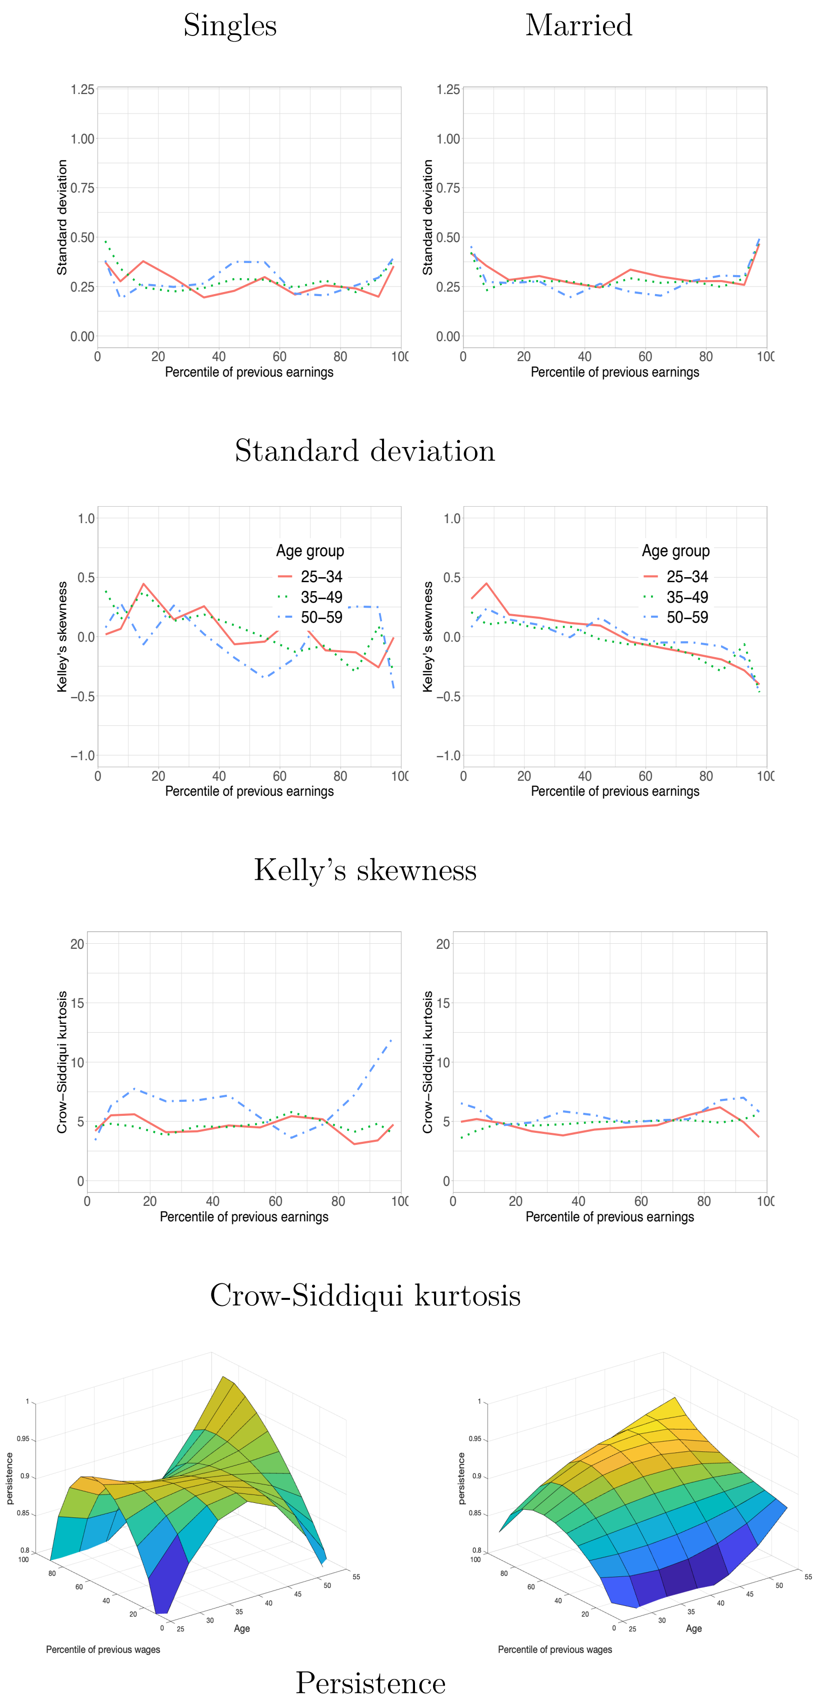

Figure 16: Moments of female wage changes in the BHPS by marital status. Singles/married are defined as those observed single/cohabiting in both t and $t+1$. Top three panels: by previous earnings. Bottom panel, by previous earnings and age.

A.5 Estimating the distribution of potential female wages

As described in Section 2.1, our preferred imputation for the potential wage of women not working in a given period uses out-of-work benefit income as an instrument for selection into employment.

In this appendix, we first provide more details about the excluded instrument in the selection equation for our preferred imputation procedure. We then contrast its implications with those of the alternative imputation procedures that we discuss in Section 2.1 in the main body of the paper. Finally, we report the results for all our imputation regressions.

Computing out-of-work benefit income To compute potential out-of-work welfare income, we use the UK tax-benefit simulator FORTAX, developed by Shephard (2009) and Shaw (2011). More precisely, we utilize the code by Blundell et al. (2016). We rely on many observables from the BHPS and Understanding Society, including marital status, earnings and hours worked by the partner, age of both partners, number of children in the household and their ages, housing tenure, region, rents paid, childcare expenses, etc. We assume that all homeowners are in council tax band D. FORTAX captures the variations in the tax and benefit system over our sample period.

Implications for potential wages Figures 17 and 18 compare the implied profiles of average earnings over the life cycle and implied distributions of potential wages of our alternative imputation procedures. They show that they are remarkably similar.

Imputation regressions Tables 9-15 report below the participation and imputation regressions for all cases. For **H** and **H-Child** we effectively estimate one participation equation for every year. Here, in the interest of space, we report one equation for all years together and omit some of the lengthier interaction coefficients.

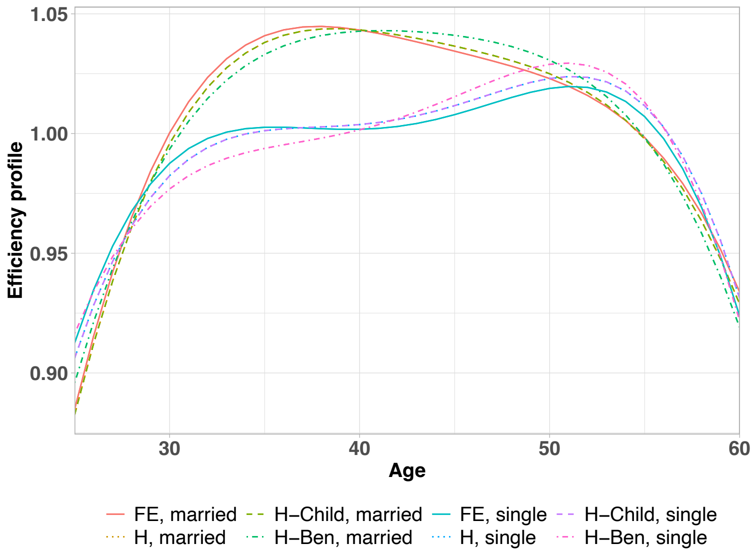

Figure 17: Age-efficiency profile for single and married women (estimated on the Understanding Society survey). **H-Child** and **H** mostly overlap.

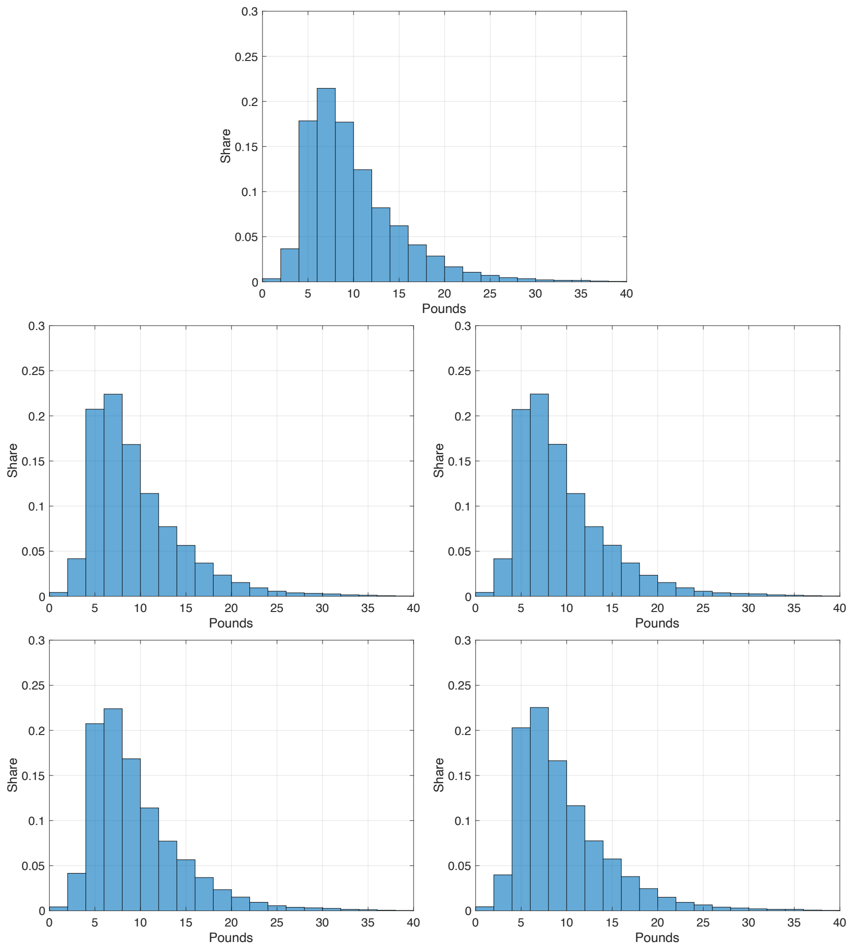

Figure 18: Distribution of women's wages. Top left: observed wages in the data. Middle left: **FE**. Middle right: **H**. Bottom left: **H-Child**. Bottom right: **H-Ben**

(Intercept)	6.7750×10^8** (2.1679×10^8)
year	-1.3536×10^6** (4.3365×10^5)
year^2	1.0142×10^3** (3.2530×10^2)
year^3	−0.3377** (0.1085)
year^4	0.0000** (0.0000)
Born 40s	0.1993*** (0.0366)
Born 50s	0.5359*** (0.0371)
Born 60s	0.6492*** (0.0380)
Born 70s	0.8170*** (0.0417)
Born 80s	1.2436*** (0.1043)
College	0.2928*** (0.0181)
Married	0.0710+ (0.0408)
No. children under 4	−0.4451*** (0.0147)
No. children	−0.1756*** (0.0076)
Potential benefit	−0.0014*** (0.0002)
married:pot.benefit	−0.0002 (0.0002)
Num.Obs.	58 124
RMSE	0.37

Table 9: Participation equation - H-AB

year	-2.4594×10^3*** (2.2703×10^2)
year^2	1.2296*** (0.1135)
year^3	−0.0002*** (0.0000)
College	0.0780*** (0.0222)
Married	0.0118 (0.0098)
No. children under 4	0.0687*** (0.0131)
No. children	−0.0078 (0.0062)
Inv. Mills Ratio	−0.1662** (0.0593)
Num.Obs.	43 198
R2	0.754
R2 Within	0.159
RMSE	0.25
FE: individual	X

Table 10: Wage imputation - H-AB

(Intercept)	−0.660*** (0.027)
Homeowner	0.677*** (0.012)
Age youngest child	0.008*** (0.001)
No. children	−0.255*** (0.007)
No. children under 4	−0.386*** (0.013)
College	0.378*** (0.016)
Married	0.121*** (0.013)
Born 40s	0.395*** (0.026)
Born 50s	0.859*** (0.026)
Born 60s	1.136*** (0.027)
Born 70s	1.246*** (0.029)
Born 80s	1.341*** (0.068)
Num.Obs.	70 165
Log.Lik.	−38 301.102
F	814.399
RMSE	0.43

Table 11: Participation equation - H

Age	0.2037* (0.0913)
Age^2	−0.0052 (0.0034)
Age^3	0.0001 (0.0001)
Age^4	−0.0000 (0.0000)
Experience	0.0250*** (0.0042)
$Experience^2$	0.0005** (0.0002)
$Experience^3$	−0.0000*** (0.0000)
No. children	−0.0352*** (0.0062)
Age youngest child	0.0001 (0.0004)
Married	0.0128 (0.0096)
Inv. Mills Ratio	0.0105 (0.0186)
Num.Obs.	43 250
R2	0.755
R2 Within	0.165
RMSE	0.25
FE: individual	X

Table 12: Wage imputation - H

(Intercept)	−0.513*** (0.030)
Years since child:1	−0.083 (0.133)
Years since child:2	−0.013 (0.109)
Years since child:3	−0.191+ (0.101)
Years since child:4	−0.230* (0.092)
Years since child:5	−0.133 (0.087)
Years since child:6	−0.190* (0.084)
Years since child:7	−0.085 (0.082)
Years since child:8	0.001 (0.081)
Years since child:9	0.052 (0.081)
Years since child:10	0.016 (0.081)
Years since child:11	0.011 (0.078)
Years since child:12	0.028 (0.078)
Years since child:13	0.068 (0.076)
Years since child:14	0.181* (0.075)
Years since child:15	0.064 (0.075)
Years since child:16	0.045 (0.074)
Years since child:17	0.033 (0.076)
Years since child:18	0.080 (0.050)
Children in hh	−0.023 (0.040)
Grandparents in hh	−0.130*** (0.027)
Husband has job	0.464*** (0.014)
Homeowner	0.628*** (0.013)
Age youngest child	0.004*** (0.001)
No. children	−0.299*** (0.010)
No. children under 4	−0.320*** (0.018)
College	0.358*** (0.016)
Married	−0.258*** (0.021)
Born 40s	0.331*** (0.026)
Born 50s	0.747*** (0.026)
Born 60s	1.027*** (0.027)
Born 70s	1.156*** (0.030)
Born 80s	1.282*** (0.069)
Y since child:1:married	0.073 (0.136)
Y since child:2:married	0.094 (0.112)
Y since child:3:married	0.117 (0.103)
Y since child:4:married	0.190* (0.093)
Y since child:5:married	0.095 (0.088)
Y since child:6:married	0.195* (0.085)
Y since child:7:married	0.105 (0.082)
Y since child:8:married	0.127 (0.081)
Y since child:9:married	0.139+ (0.080)
Y since child:10:married	0.200* (0.080)
Y since child:11:married	0.237** (0.077)
Y since child:12:married	0.291*** (0.077)
Y since child:13:married	0.309*** (0.074)
Y since child:14:married	0.222** (0.073)
Y since child:15:married	0.199** (0.073)
Y since child:16:married	0.174* (0.073)
Y since child:17:married	0.145+ (0.076)
Y since child:18:married	−0.034 (0.036)
Num.Obs.	70 165
Log.Lik.	−37 552.444
F	201.911
RMSE	0.42

Table 13: Participation equation - H-Child

Age	0.2082* (0.0908)
Age^2	−0.0053 (0.0034)
Age^3	0.0001 (0.0001)
Age^4	−0.0000 (0.0000)
Experience	0.0250*** (0.0042)
Experience^2	0.0005** (0.0002)
Experience^3	−0.0000*** (0.0000)
No. children	−0.0336*** (0.0058)
Age youngest child	0.0001 (0.0004)
Married	0.0121 (0.0095)
Inv. Mills Ratio	0.0006 (0.0145)
Num.Obs.	43 250
R2	0.755
R2 Within	0.165
RMSE	0.25
FE: individual	X

Table 14: Wage imputation - H-Child

Age	0.2110* (0.0908)
Age^2	−0.0054 (0.0034)
Age^3	0.0001 (0.0001)
Age^4	−0.0000 (0.0000)
Experience	0.0249*** (0.0042)
Experience^2	0.0005** (0.0002)
Experience^3	−0.0000*** (0.0000)
No. children	−0.0336*** (0.0052)
Age youngest child	0.0001 (0.0004)
Married	0.0188+ (0.0104)
Husband has job	−0.0101 (0.0069)
Num.Obs.	43 250
R2	0.755
R2 Within	0.165
RMSE	0.25
FE: individual	X

Table 15: Wage imputation - FE

B Estimation and features of the earnings processes

B.1 Comparing the non-linear and canonical processes

As described in De Nardi, Fella, and Paz-Pardo (2020), the canonical process, described in Equation 4, can be specified as a restricted version of the NL process in Equation 5, where:

$$Q_{z_{i,t}}(v_{it}|z_{i,t-1},t) = \rho z_{i,t-1} + \sigma_{\nu}\phi^{-1}(v_{it}) \tag{19}$$

$$Q_{\epsilon}(e_{it}) = \sigma_{\epsilon}\phi^{-1}(e_{it}), \tag{20}$$

where $\phi^{-1}(.)$ is the inverse of the cumulative density function of a standard, normal distribution. This specification allows to clearly see the restrictions the canonical process imposes on the earnings process:

1. *Age-independence* (stationarity) of the autoregressive coefficient ρ and of the shock distributions (both normal with constant standard deviations σ_{ν} and σ_{ϵ}), which imply age-independence of the second and higher moments of the conditional distributions of both the transitory and the persistent component.

2. *Normality* of the shock distributions ($\phi^{-1}(\cdot)$).

3. *Linearity* of the process for the persistent component, which can be seen in the additive separability of equation 19 into the conditional expectation—the first addendum—and an innovation independent of $z_{i,t-1}$, and (b) the linearity of the conditional expectation in $z_{i,t-1}$. Under separability, deviations of z_{it} from its conditional expectation are just a function of the innovation ν_{it}. As a consequence, all conditional centered second and higher moments are independent of previous realizations of z.

One further way to understand the role of nonlinearity is in terms of a generalized notion of persistence

$$\rho(q|z_{i,t-1},t) = \frac{\partial Q_z(q|z_{i,t-1},t)}{\partial z_{i,t-1}} \tag{21}$$

which measures the persistence of $z_{i,t-1}$ when it is hit by a shock that has rank q. In the canonical model, $\rho(q|z_{i,t-1},t) = \rho$, independently of both the past realization of $z_{i,t-1}$ and of the shock rank q. Instead, the general model allows persistence to depend both on the past realization $z_{i,t-1}$, but also on the sign and magnitude of the shock realization. Basically, in the nonlinear model shocks are allowed to wipe out the memory of past shocks or, equivalently, the future persistence of a current shock may depend on future shocks.

Of course, a similar unrestricted representation can be used for the transitory component ϵ_{it} and the initial condition η_1, with the only difference that they are not persistent.

We proceed in two steps. First, we use the quantile-based panel data method proposed by Arellano, Blundell, and Bonhomme, 2017 to estimate a non-parametric model that allows for age-dependence, non-normality and nonlinearity, and that can be applied in datasets of moderate sample size like the PSID. This step gives us quantile functions for both the two (persistent and transitory) component of earnings (see the next section, Appendix B.2 for details on the estimation). Second, we use the two quantile functions to simulate histories for the two earnings components and proceed to estimate, for the persistent component, a discrete Markov-chain approximation, which can then be easily introduced in a structural model.

B.2 Estimation

Following Arellano, Blundell, and Bonhomme, 2017, we parameterize the quantile functions for the three variables as low order Hermite polynomials

$$Q_\epsilon(q|age_{it}) = \sum_{k=0}^{K} a_k^\epsilon(q)\psi_k(age_{it}) \tag{22}$$

$$Q_{z_1}(q|age_{i1}) = \sum_{k=0}^{K} a_k^{z_1}(q)\psi_k(age_{i1}) \tag{23}$$

$$Q_z(q|z_{i,t-1}, age_{it}) = \sum_{k=0}^{K} a_k^z(q)\psi_k(z_{i,t-1}, age_{it}) \tag{24}$$

where the coefficients $a_k^i(q)$, $i = \epsilon, z_1, z$, are modeled as piecewise-linear splines in q on a grid $\{q_1 < \ldots < q_L\} \in (0,1)$.[11] The intercept coefficients $a_0^i(q)$ for q in $(0, q_1]$ and $[q_L, 1)$ are specified as the quantiles of an exponential distribution with parameters λ_1^i and λ_L^i.

If the two earnings components ϵ_{it} and z_{it} were observable one could compute the

[11] Following Arellano, Blundell, and Bonhomme, 2017 we use tensor products of Hermite polynomials of degrees (3,2) in $z_{i,t-1}$, and age for $Q_z(q|z_{i,t-1}, age_{it})$ and second-order polynomials in age for $Q_\epsilon(q|age_{it})$ and $Q_{z_1}(q|age_{i1})$.

polynomial coefficients simply by quantile regression for each point of the quantile grid q_j. To deal with the latent earnings components, the estimation algorithm starts from an initial guess for the coefficients and iterates sequentially between draws from the posterior distribution of the latent persistent components of earnings and quantile regression estimation until convergence of the sequence of coefficient estimates.

B.3 Persistent and transitory earnings

In this section, we compare the non-linear and non-normal features of the BHPS data and the persistent and transitory components that result from the Arellano, Blundell, and Bonhomme (2017) decomposition.

Starting with male earnings, persistence is lowest for the young and for the lowest earners both for the BHPS data and the persistent component (Figure 19). As expected, the persistent component displays a larger overall persistence than the data, but shows the same patterns by age and over the earnings distribution.

Figure 20 shows the standard deviation, skewness, and kurtosis of earnings changes for the BHPS data and their persistent component. Their persistent component preserves most of the features of non-normality that are present in the data and the dependence on previous earnings realizations. The main difference lies in the Crow-Siddiqui kurtosis, which is significantly larger for the persistent component than in the raw data.

Transitory shocks, that we consider to be measurement error, are very leptokurtic, in particular for male earnings, and display negative skewness (see Figure 21).

Women's wages display similar patterns (see Figures 22, 23, and 24). The most noticeable difference is that the persistence of the persistent component is relatively high and close to 1, but still replicates the inverted U-shape by previous wages that we observe in the data.

Finally, in Figures 25 and 26 we show that most of the differences in dynamics between men's earnings and women's wages are also present if we compare male and female earnings. For example, the profile of persistence over the earnings distribution is much flatter for women than for men.

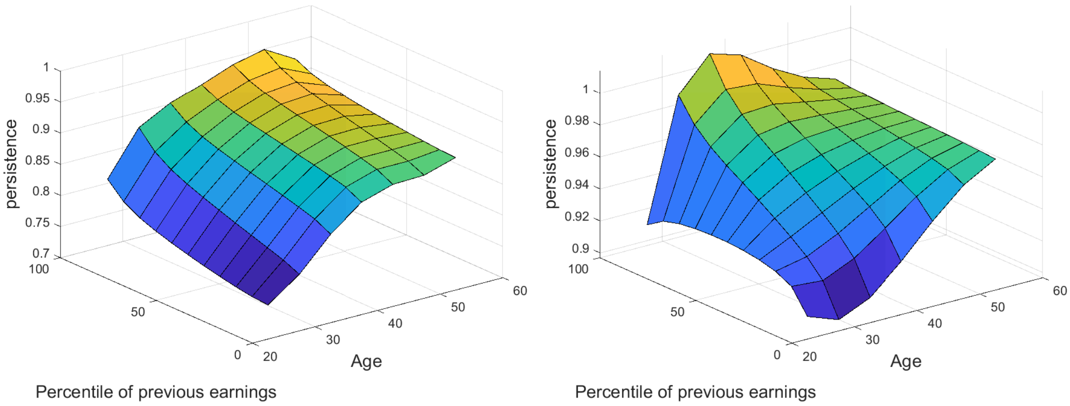

Figure 19: Non-linear persistence of male earnings by age and previous earnings in the BHPS. Left, data; right, persistent component

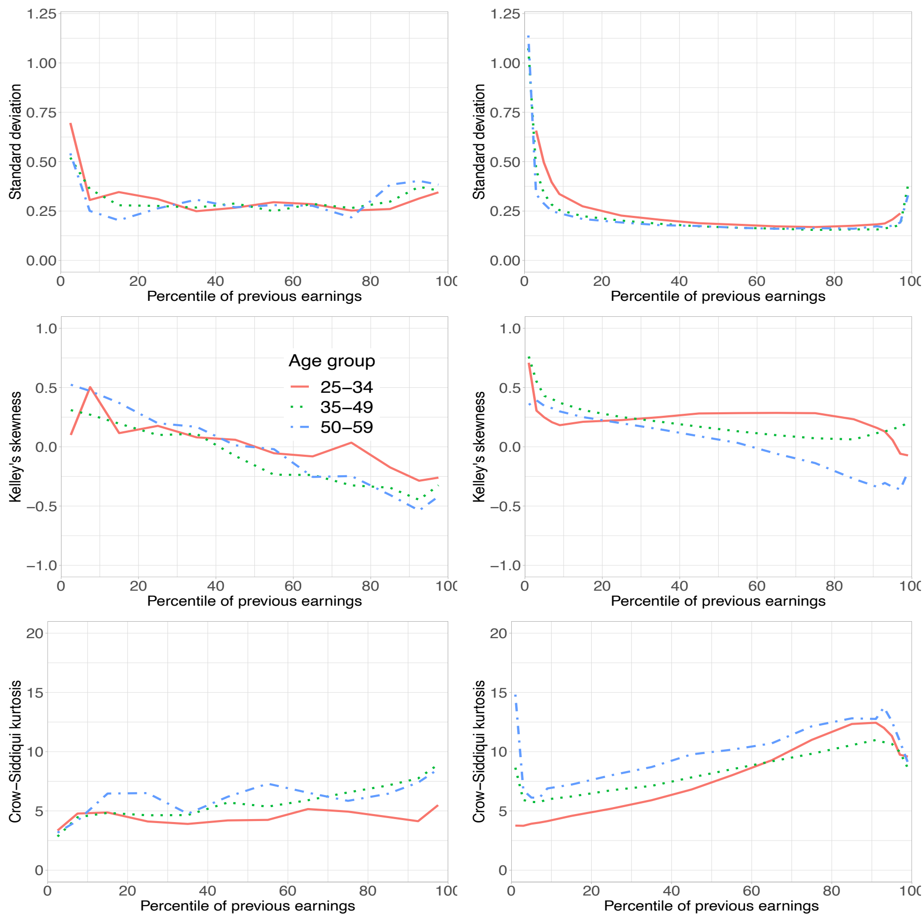

Figure 20: Standard deviation (top), skewness (middle) and kurtosis (bottom) of male earnings changes in the BHPS. Left, data; right, persistent component

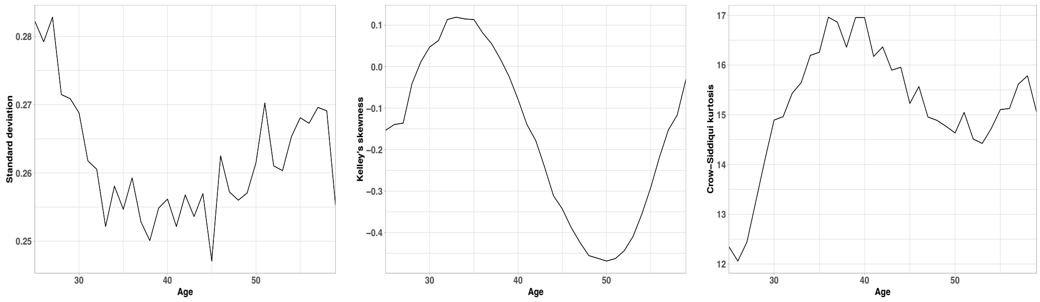

Figure 21: Transitory shock to male earnings: standard deviation, skewness and kurtosis by age

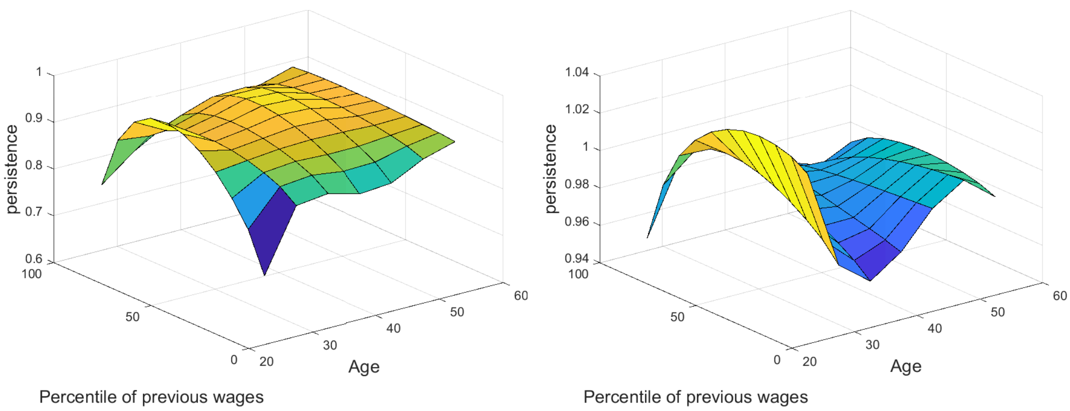

Figure 22: Non-linear persistence of female wages by age and previous wages in the BHPS. Left, data; right, persistent component

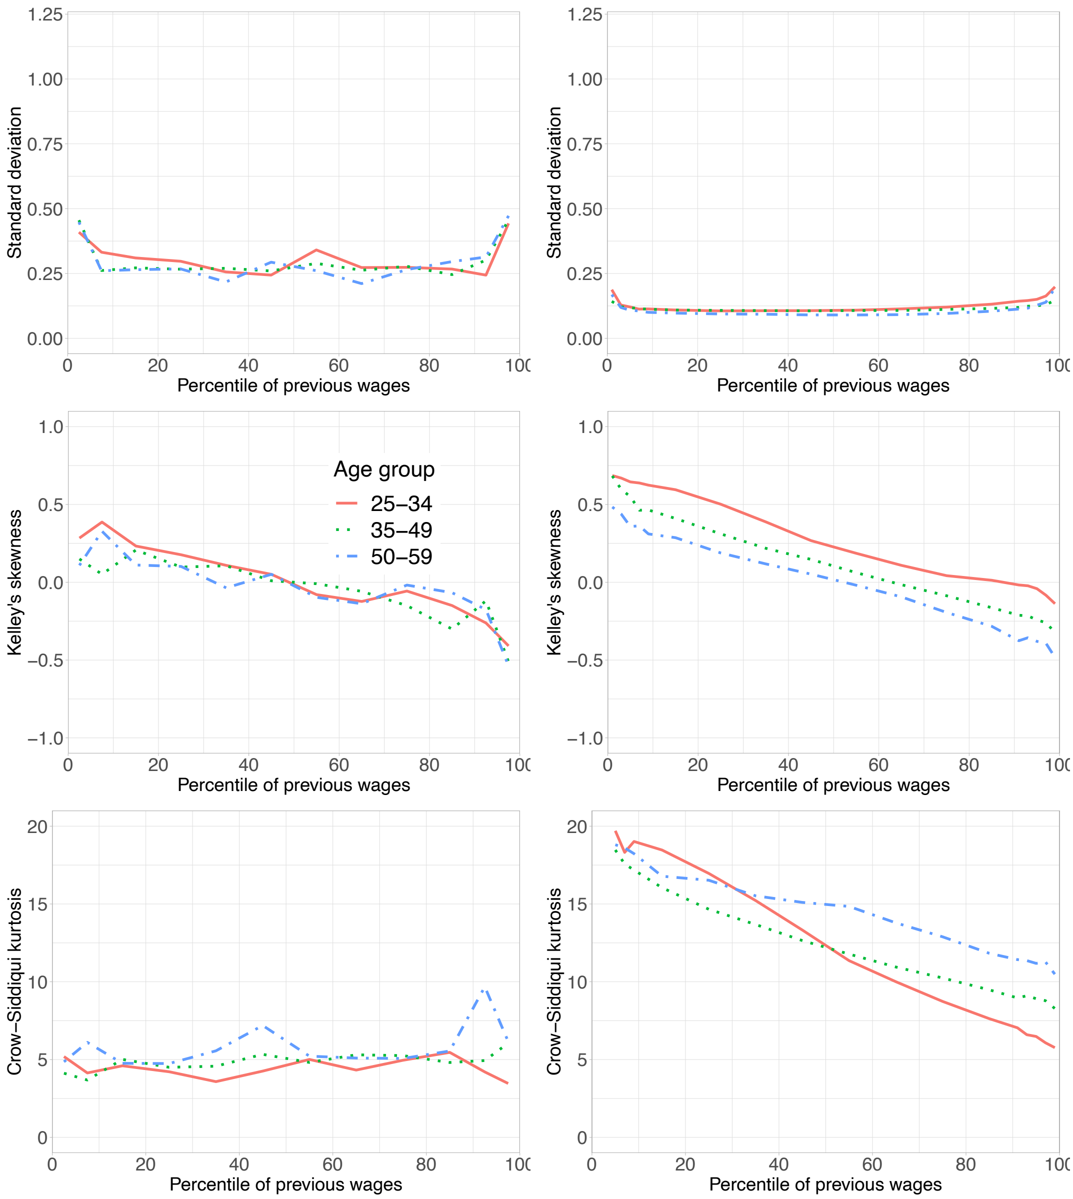

Figure 23: Standard deviation (top), skewness (middle) and kurtosis (bottom) of female wage changes in the BHPS. Left, data; right, persistent component

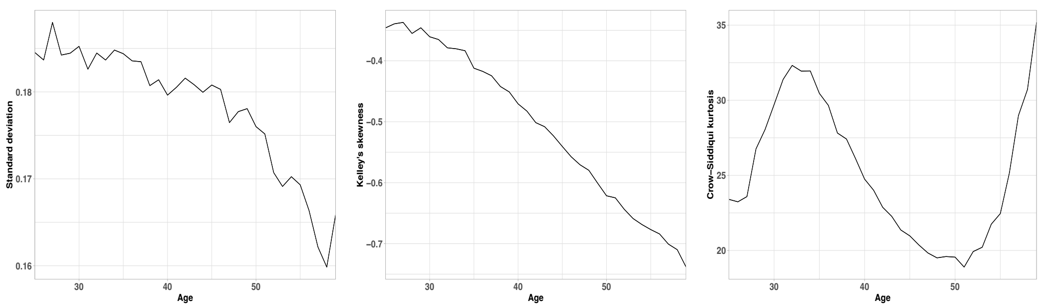

Figure 24: Transitory shock to female wages: standard deviation, skewness and kurtosis by age

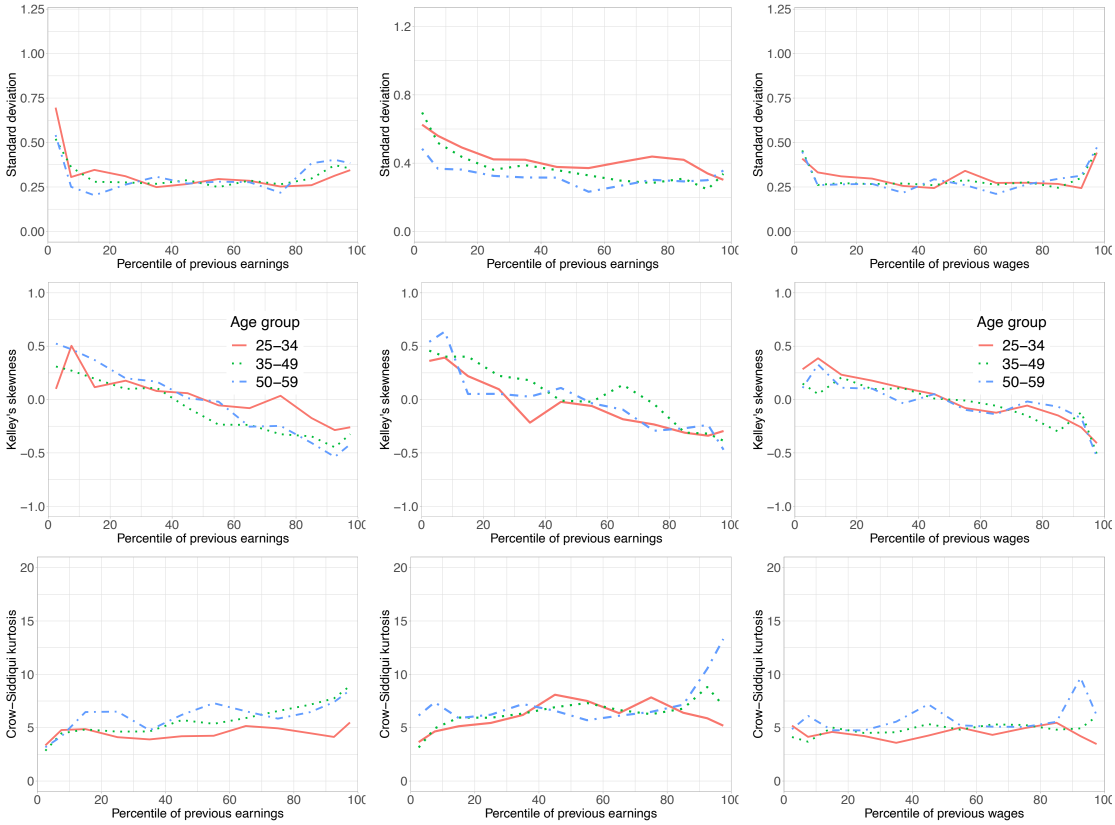

Figure 25: Standard deviation (top), skewness (middle) and kurtosis (bottom). Left: male earnings; middle: female earnings; right: female wages

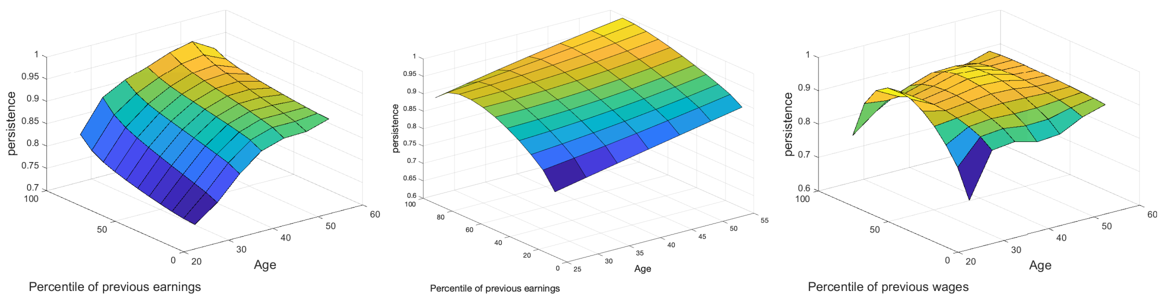

Figure 26: Non-linear persistence of male earnings (left), female earnings (middle) and female wages (right), by age and percentile of previous wages, BHPS data

C Other model inputs

C.1 Marriage and divorce

	Dependent variable:	
	marriage	divorce
	(1)	(2)
Age	-0.032^{***}	-0.018^{***}
	(0.002)	(0.003)
(log) Wife's imputed wage	0.041	-0.017
	(0.035)	(0.046)
(log) Husband's income		-0.007^{***}
		(0.002)
Constant	-0.161	-1.156^{***}
	(0.103)	(0.136)
Observations	11,349	22,013
Log Likelihood	-3,358.402	-1,898.700
Akaike Inf. Crit.	6,722.805	3,805.399
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 16: Probability of marriage and divorce (probit regressions) between t-1 and t, conditional on income at t-1

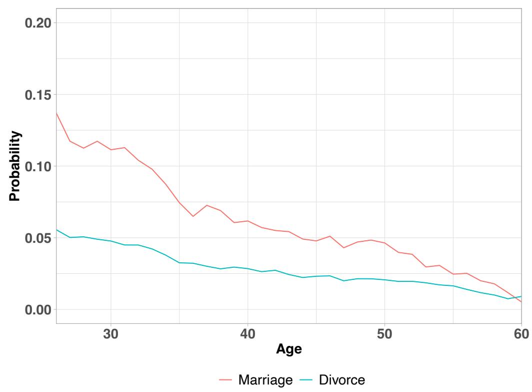

Figure 27: Marriage probabilities for single women, and divorce probabilities for married women, by age (BHPS data)

	Dependent variable:
	(log) Earnings of husband in t
(log) Woman's wage in t	0.325***
	(0.019)
Constant	9.359***
	(0.041)
Observations	3,728
R^2	0.076
Adjusted R^2	0.075
Residual Std. Error	0.506 (df = 3726)
F Statistic	304.354*** (df = 1; 3726)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 17: Correlation of husband's earnings and wife's wages before 30

	Dependent variable:
	(log) Earnings of husband in t
(log) Woman's wage in t-1	0.272***
	(0.055)
Constant	9.480***
	(0.123)
Observations	386
R^2	0.059
Adjusted R^2	0.056
Residual Std. Error	0.524 (df = 384)
F Statistic	23.987*** (df = 1; 384)
Note:	*p<0.1; **p<0.05; ***p<0.01

Table 18: Correlation of husband's earnings and wife's wages at marriage

	Dependent variable:	
	log wealth of partner	
	(1)	(2)
Age	−0.003	−0.033
	(0.058)	(0.062)
(log) Woman's wage	2.362**	
	(1.074)	
(log) Men's income		1.688**
		(0.743)
Constant	3.701	−9.011
	(2.929)	(7.458)
Observations	86	117
R^2	0.055	0.044
Adjusted R^2	0.032	0.027
Residual Std. Error	4.257 (df = 83)	4.705 (df = 114)
F Statistic	2.424* (df = 2; 83)	2.625* (df = 2; 114)
Note:		*p<0.1; **p<0.05; ***p<0.01

Table 19: Correlation between partner's wealth before marriage and income of reference person at marriage year

C.2 Children

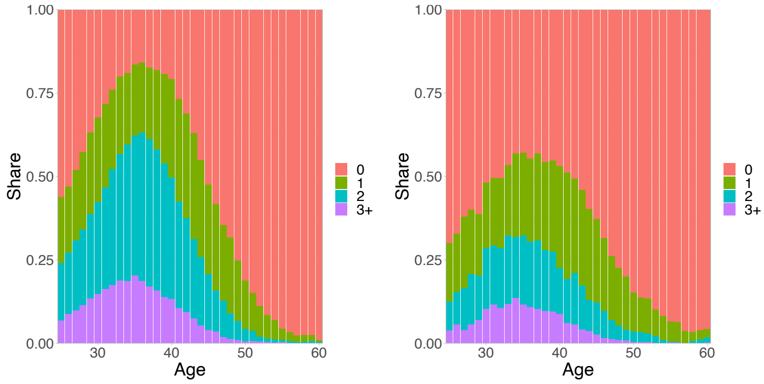

Figure 28: Distribution of number of children in the household, by age of the mother. Left: married mothers; right: single mothers

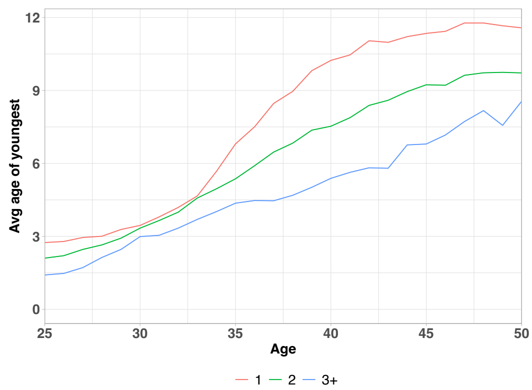

Figure 29: Average age of youngest cohabiting child by age of mother and number of children

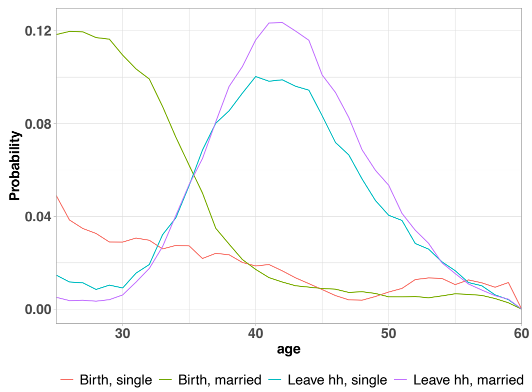

Figure 30: Probability that a child arrives (birth) and leaves (leave hh) a household, by age and marital status (of the woman), unconditionally on today's child status (BHPS data)

C.3 Mortality risk

Figure 31 shows the mortality risk by age and marital status in the model. Although there is information about mortality in the BHPS data, its sample size is too small to obtain reliable estimates for death probabilities that are age, gender, and marital-status specific. Thus, we turn to the life tables data from the Human Mortality Database (1980-2010), which are reported separately by gender and age. Then, to incorporate the increased mortality risk for singles, we estimate the average gap in mortality probabilities between singles and married people in the BHPS during the retirement period. We assume that this gap is constant during adult life, and compute the death probabilities for single and married people that are consistent with this gap and the observed mortality probabilities in the life tables. We report them in Figure 31.

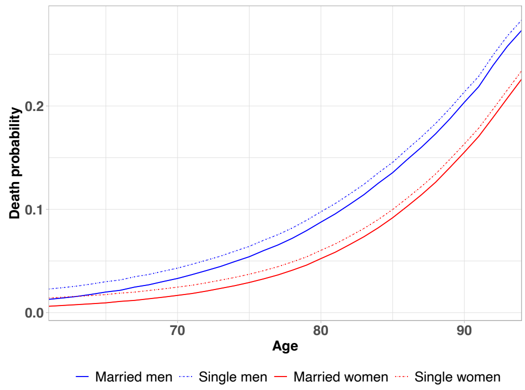

Figure 31: Annual death probabilities by age, gender, and marital status. Source: Human Mortality Database and BHPS data.

C.4 Average male earnings and female wages

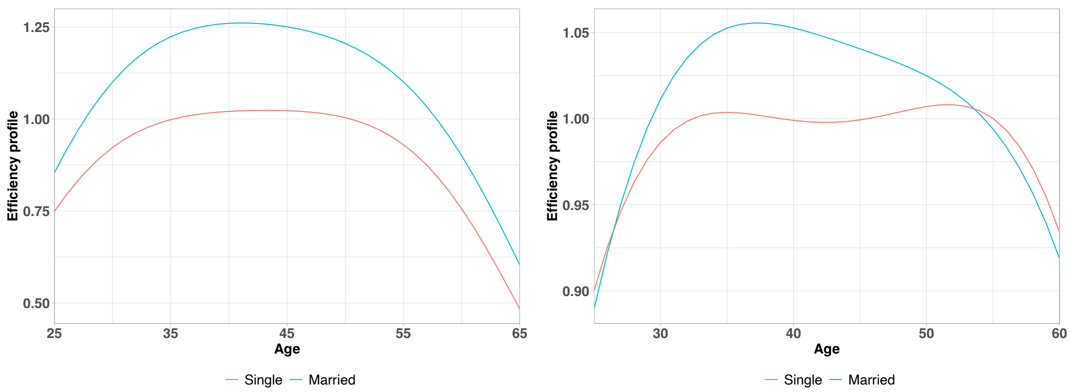

Figure 32: Age-efficiency profiles, left: men's earnings; right: women's wages. For this representation, both are individually normalized so that their average is 1

C.5 Population shares, data vs. model

In our sample, after excluding the retirees, long-term disabled, and full-time students, 9% of single men and 4% of married men display zero earnings in a given calendar year. The corresponding shares are 20% and 19.9% for single and married women, respectively. Within the male working population, 4.3% of singles and 3.2% of married people work part-time. The corresponding shares for women are 25% and 41%, respectively. As a result, we have chosen to model the labor supply decision of women explicitly and to assume that men always work full time.

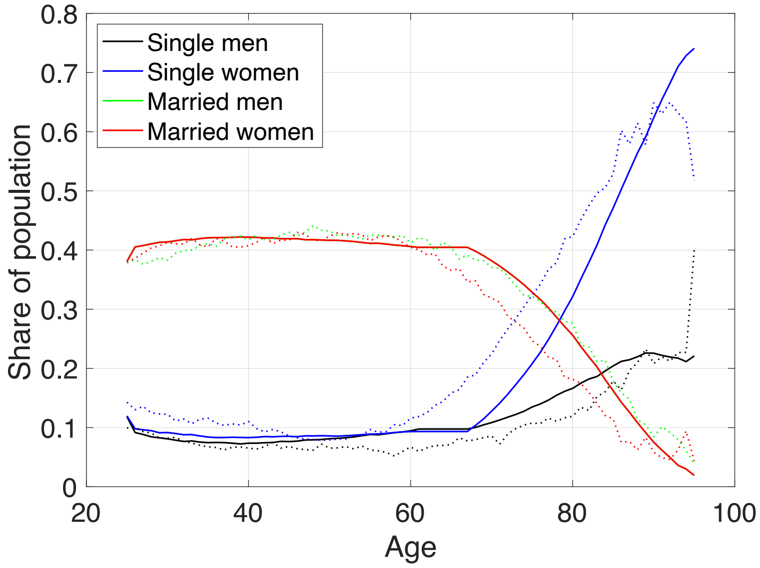

Figure 33: Share of people by age, gender, and marital status. Solid lines: model outcomes, dotted lines: data. In the model, the number of married men and married women at each age is identical by construction. Data: BHPS, whole sample

D UK Benefit system, details

Table 20 provides a brief overview of the main benefits for the working age population in the United Kingdom before the introduction of Universal Credit in 2016.[12]

In our model, in-work benefits are meant to capture the Working Tax Credit, while income support replicates a variety of benefits that low-income people receive, including Income-based Jobseeker's Allowance, Income Support, Housing Benefits, Child Benefits, and Child Tax Credits.

The tapering rate for in-work benefits corresponds to the statutory tapering rate for the Working Tax Credit (0.41). For income support, we compute an average tapering rate ω of the different benefits it summarizes, considering their respective sizes, tapering rates, and eligibility criteria, including how access to one of the benefits impacts the entitlement to the others. We do so in the following way. First, we calculate the benefit entitlement B_i^k by demographic group k (gender, marital status, and number of children) and household labor income y_i. We do so under the assumption that the household is eligible for all of the benefits that compose our income support, also taking into account that a household can only claim either Income-based Jobseeker's Allowance or Income Support, but not both at

[12]Given the gradual and too recent phase-in of Universal Credit, it would not have been appropriate to calibrate our steady-state benchmark economy to the post-2016 period.

the same time. We additionally assume that the household would be getting Working Tax Credit whenever eligible, which affects their eligibility criteria for other benefits (namely, Child Tax Credits and the Working Tax Credit are considered as income for purposes of computing eligibility for Income Support and Housing Benefits).

We then find the β_0^k and β_1^k that minimize:

$$\sum_i (B_i^k - max(\beta_0^k - \beta_1^k y_i^{hk}, 0))^2 \tag{25}$$

where the sum i is taking over all possible income levels between 0 and £100,000. We then obtain our estimate of ω by weighing the different β_1^k by the relative sizes in the population of each k group. The average tapering rate is then $-\beta_1$ is 0.70, which also corresponds to the tapering rate for couples with zero children.

E Additional model implications

E.1 Observed wages in the data and in the model

Figure 34 reports the distribution of potential wages that we use in our model, computed using our Heckman selection correction (left), the implied distribution of observed wages that the model delivers, under the assumption that we can only observe wages for women who choose to work (center), and the distribution of female wages in the data, which we can only observe for those who are actively participating (right).

Our model-implied distribution of observed wages is closer to the data than the distribution of potential wages, thus suggesting that the model replicates the patterns of selection in the data well. For instance, looking at the second bar of these histograms (which are computed in such a way that the binning is identical for all three), one can observe that it is taller in the potential wage distribution (a lot of women have low potential wages), but lower and closer to the data in the model-implied observed wage distribution (thus suggesting that many women select out of the labor force when they receive a low wage realization).

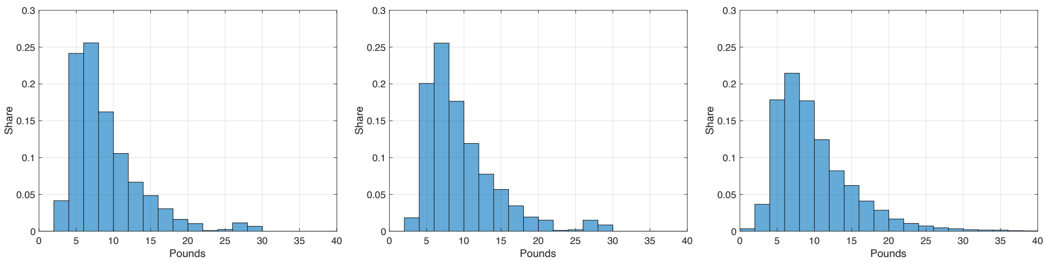

Figure 34: Distribution of women's wages. Left, potential wages in the model; middle: observed wages in the model; right: observed wages in the data.

Benefit	Time period	Eligibility (income)	Tapering	Wealth test	£M (2016)
Benefits for the unemployed					
Jobseeker's Allowance (Contributory)	1996-today	Work < 16h/week	100%	No	306
Jobseeker's Allowance (Income-based)	1996-today	Work < 16h/week	100%	Yes	2000
Benefits for low-income people					
Income Support		Work < 16h/week	100%	Yes	2700
Housing benefit		Tapering starts after JSA amount	65%	16k	24300
Council Tax Benefit	-2013	Being on IS, JSA, etc.	No	Yes	
Benefits for families					
Child benefit		Income < £50k	No	No	11300
Statutory Maternity Pay		None	No	No	2300
Maternity Allowance (Contributory)		Min £30 pw	No	No	443
Tax credits					
Child Tax Credit	2003-	Taper from £16,105 (2014)	41%	No	21700
Working Tax Credit	2003-	Working FT, taper from £6,420	41%	No	5900
Benefits for the sick and disabled					
ESA	2011-today	Work <16h/week	100%	No	14300
Personal Independence Payment	2013-	Work capability assessment	-	No	3000
Disability Living Allowance	-2013	Unable to work	-	No	13200
Carer's Allowance		No	No	No	2600
Industrial Injuries Benefits		Depends on disablement rate	No	No	869

Table 20: Main benefits for working age population in the UK (source: Hood and Norris Keiller (2016))

E.2 Labor market participation by number of children

Figure 35 shows the share of women that are working, conditional on the number of children they have, in the data and in the two versions of our model. We find that both are good at capturing the general patterns of participation by number of children, particularly for the largest groups (those with 0, 1 or 2 children). Our model is also successful at capturing the larger decrease in labor market participation for singles than for couples as the number of children increases. However, the model overestimates the share of women with 3 children or more who are working for both earnings process. This mismatch is related to the assumption that the maximum number of children is 3 in our model, but in the data the number of children might be larger than 3, which further discourages female labour supply. Furthermore, older cohorts in the data are both more likely to have more children and stay at home, and in our model we abstract from cohort heterogeneity.

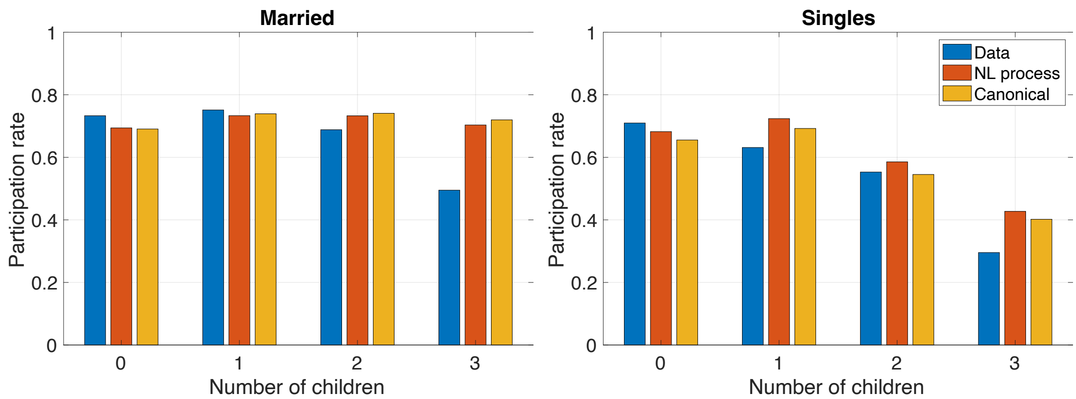

Figure 35: Women's labour market participation, by number of children.

E.3 Persistence of female labor force participation

In our model, we introduce heterogeneity in the disutility of work in line with previous literature on female labor supply, including Keane and Wolpin (2010), Blundell et al. (2016), or Adda, Dustmann, and Stevens (2017), and as a parsimonious way to capture the large amount of heterogeneity in the data. For simplicity and transparency, we assume that there are two types of women, one with higher disutility from work than the other.

A way of evaluating whether the size of our fixed costs of work and their heterogeneity are quantitatively reasonable is to look at the dynamics of female labor force participation.

Aspects of the data that are pertinent for these purposes are the persistence of (a) being on benefits and (b) being unemployed or out of the labor force. In the data, the persistence of benefit receipt is 0.78; in the model it is also 0.78. In the data, the persistence of the unemployed/out-of-labor-force status for women is 0.80; in the model it is 0.88. Both are non-targeted moments by our estimation strategy.

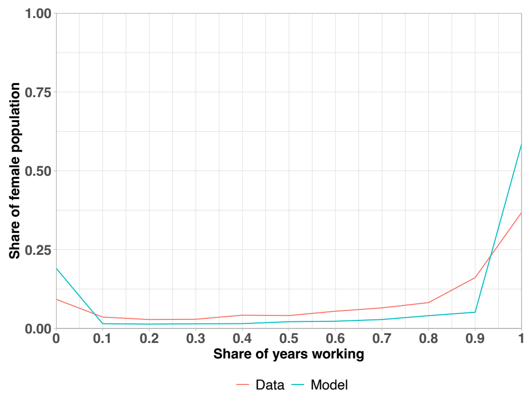

Figure 36: Distribution of the female population over the number of years worked in a 10-year period. Data: BHPS balanced sample of women who are observed for 10 years in a row.

In addition, Figure 36 shows, in a balanced 10-year sample, how many women work every year (1.0), don't work at all (0.0), or intermediate cases during that 10 year span. Both in the data and in the model there is large heterogeneity: some women work all the time, while others never work at all. The model does a reasonably good job of matching this untargeted distribution.

E.4 Universal Credit, canonical process

In this section, we report the welfare effects of the introduction of Universal Credit under the canonical wage process. As described in Section 4, in our main results with the NL process, we keep the change to Universal Credit budget neutral by multiplying all allowances with a proportional scaling factor of 0.9. For the purposes of this section, we keep budget neutrality under the canonical process, which implies that we scale these allowances by 0.82

Under the canonical wage process, the switch to Universal Credit generates a drop in full-time labor force participation and a large rise in part-time labor force participation, particularly at older ages (Figure 37).

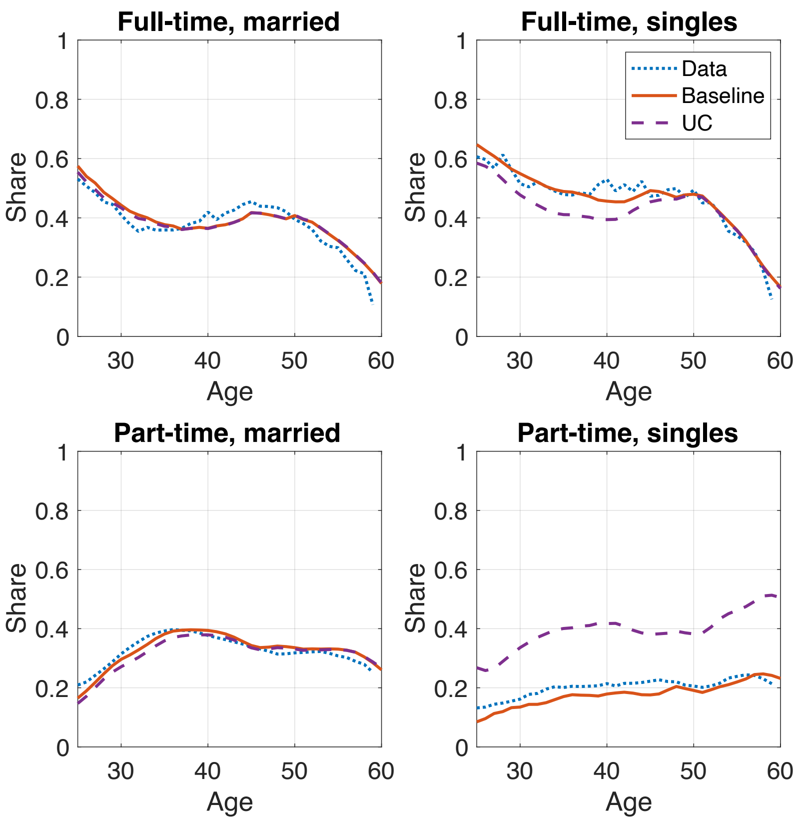

Figure 37: Labor force participation under canonical process: Universal Credit vs baseline, universal credit.

References

Adda, Jérôme, Christian Dustmann, and Katrien Stevens (2017). "The career costs of children". In: *Journal of Political Economy* 125.2, pp. 293–337.

Arellano, Manuel, Richard Blundell, and Stéphane Bonhomme (2017). "Earnings and Consumption Dynamics: A Non-Linear Panel Data Framework". In: *Econometrica* 85.3, pp. 693–734.

Blundell, Richard et al. (2016). "Female labor supply, human capital, and welfare reform". In: *Econometrica* 84.5, pp. 1705–1753.

De Nardi, Mariacristina, Giulio Fella, and Gonzalo Paz-Pardo (2020). "Nonlinear household earnings dynamics, self-insurance, and welfare". In: *Journal of the European Economic Association* 18.2, pp. 890–926.

Guvenen, Fatih et al. (2021). "What Do Data on Millions of U.S. Workers Reveal About Lifecycle Earnings Dynamics?" In: *Econometrica* 89.5, pp. 2303–2339.

Hood, Andrew and Agnes Norris Keiller (2016). "A survey of the UK benefit system". In: *London: Institute for Fiscal Studies.*

Kaplan, Greg (2012). "Inequality and the life cycle". In: *Quantitative Economics* 3.3, pp. 471–525.

Keane, Michael P. and Kenneth I. Wolpin (2010). "The Role of Labor and Marriage Markets, Preference Heterogeneity, and the Welfare System in the Life Cycle Decisions of Black, Hispanic, and White Women". In: *International Economic Review* 51.3, pp. 851–892.

Shaw, Jonathan (2011). *FORTAX: UK tax and benefit system documentation.* IFS Working Paper 08/11.

Shephard, Andrew (2009). *FORTAX: Reference Manual.* Unpublished Manuscript.