

De Nardi, Mariacristina; Fella, Giulio; Paz Pardo, Gonzalo

Working Paper

Earnings risk, government policy, and household welfare

IFS Working Paper, No. W22/17

Provided in Cooperation with:

Institute for Fiscal Studies (IFS), London

Suggested Citation: De Nardi, Mariacristina; Fella, Giulio; Paz Pardo, Gonzalo (2022) : Earnings risk, government policy, and household welfare, IFS Working Paper, No. W22/17, Institute for Fiscal Studies (IFS), London,
<https://doi.org/10.1920/wp.ifs.2022.1722>

This Version is available at:

<https://hdl.handle.net/10419/267948>

Standard-Nutzungsbedingungen:

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

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

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

Terms of use:

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

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

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

Mariacristina De Nardi
Giulio Fella
Gonzalo Paz-Pardo

22/17

Working paper

Earnings risk, government policy, and household welfare

Earnings Risk, Government Policy, and Household Welfare

Final Report

January 2022

Mariacristina De Nardi

Giulio Fella

Gonzalo Paz-Pardo



Earnings Risk, Government Policy, and Household Welfare

About the authors of this report:

Mariacristina De Nardi (denar001@umn.edu) is the Thomas Sargent Professor at the University of Minnesota and the Principal Investigator of this research project. She is also a consultant at the Federal Reserve Bank of Minneapolis, a Research Fellow at the CEPR (Centre for Economic Policy Research) and NBER (National Bureau of Economic Research), and an International Research Fellow at IFS (Institute for Fiscal Studies).

Giulio Fella (g.fella@qmul.ac.uk) is a Professor in the School of Economic and Finance at Queen Mary University of London and Co-Investigator on this project. He is also a Research Fellow of the CEPR (Centre for Economic Policy Research) and a Research Associate at the CfM (Centre for Macroeconomics) and IFS (Institute for Fiscal Studies).

Gonzalo Paz-Pardo (Gonzalo.Paz_Pardo@ecb.europa.eu) is a Senior Economist at the European Central Bank and was employed as a research assistant on this project.

The views expressed in this report are those of the authors and not necessarily the Foundation, or the NBER, CEPR, any agency of the federal government, the Federal Reserve Bank of Minneapolis, the European Central Bank, or the IFS.

Published: December 2020

Contact information:

Giulio Fella

School of Economics and Finance - Queen Mary University of London

Mile End Road, E1 4 NS London, UK

Phone: +44 20 7882 8823

Email: g.fella@qmul.ac.uk

Website: <http://giuliofella.net>

Acknowledgements

This report summarizes the output of a two-year research project on “Earnings Risk, Government Policy, and Household Welfare” that has been supported by a generous grant from the Nuffield Foundation (<https://www.nuffieldfoundation.org/earnings-risk-government-policy-and-household-welfare>). We are extremely indebted to the Nuffield Foundation for their continuous economic and non-economic support that has made this research project possible. The Nuffield Foundation (www.nuffieldfoundation.org) is an independent charitable trust with a mission to advance social well-being. It funds research that informs social policy, primarily in Education, Welfare, and Justice. It also funds student programmes that provide opportunities for young people to develop skills in quantitative and scientific methods. The Nuffield Foundation is the founder and co-funder of the Nuffield Council on Bioethics and the Ada Lovelace Institute. The Foundation has funded this project, but the views expressed in this report are our own views and not necessarily those of the Foundation.

We are also extremely grateful to the members of **our Advisory Board** – Richard Blundell (UCL), Mike Daly (DWP), and Robert Joyce (IFS) – for providing extremely useful feedback and invaluable policy insight on our research and for taking part in the conference on “Earnings, Risk, and Insurance” that was held at the IFS, London, in April 2019. (see conference programme in Appendix A and further information at: <https://www.ucl.ac.uk/economics/events/2019/apr/conference-earnings-risk-and-insurance>).

We also thank **the administrative staff in the Department of Economics at UCL** for helping us with the management of the grant and the organization of the final conference.

Finally, we would like to thank **all participants to seminars and conferences** where we presented our research paper for their comments, suggestions and encouragement.

Table of Contents

1	Introduction	10
2	Our research	11
2.1	Empirical findings on labour income dynamics in the UK	11
2.1.1	Empirical question	11
2.1.2	Data	12
2.1.3	Comparing labour income processes	14
2.2	Implications of the wage process for the evaluation of welfare policies . .	18
2.2.1	Optimal benefit reform	20
2.2.2	Universal Credit	27
3	Conclusions	30
A	Conference Programme	35

Executive summary

Introduction

The desirability, efficacy, and cost-effectiveness of government welfare policies depend crucially on the income risk that households face and the actions that they can take to reduce consumption fluctuations, for instance by adjusting their saving and labor supply. Shocks to labour earnings are a key driver of income risk. A recent literature has documented that the dynamics of labour earnings are much richer than typically assumed in the class of models used to evaluate the desirability of tax and welfare policies. This raises the question of how taking into account this richer dynamics affects the evaluation of such policies.

This report summarizes the findings of a two-year research project on “**Earnings Risk, Government Policy, and Household Welfare**,” funded by a grant from the Nuffield Foundation. The project documents that the dynamics of male and female labour earnings in the UK is substantially richer than typically assumed and shows that accounting for this richer dynamics is important for the evaluation of welfare benefit reforms.

Background and methodology

- We use both administrative and survey data to document that the dynamics of both male and female labour earnings in the UK is substantially richer than typically assumed in positive and normative studies of inequality and welfare policies.
- These richer dynamics also apply to (hourly) wages, which means that they are a genuine feature of labor income risk rather than being driven by individuals’ choice of hours worked.
- We model the dynamics of men’s earnings and women’s wages, as accounting for a separate dynamics of wages and hours is most meaningful for women who have a higher labour supply elasticity. For easiness of exposition we refer to men’s *earnings* as wages in what follows.
- We estimate two alternative stochastic processes for wages: (1) a traditional one, of the form typically assumed in the literature on the evaluation of welfare policies;

and (2) a richer, more flexible, one which can accommodate the features of the data we document.

- We build and estimate a structural, life-cycle model featuring households of different composition and an active female labour supply margin.
- We use the model to compare the implications of the richer wage process that we estimate to those of the traditional, more restrictive, one.

New findings on wage dynamics in the UK

- The typical (canonical) approach to the statistical modelling of wage dynamics assumes that the probability and size of shocks to wages are: (1) the same at all ages; (2) independent from current wage. Using novel, flexible, statistical techniques that do not impose these restrictions reveals that these assumptions are violated.
- Rather than being constant, the dispersion of shocks to wages is substantially higher for individuals at the bottom and top of the distribution of current wages.
- Persistence – the extent to which future wages depend on current ones – is also not constant; wages are less persistent for younger than for older individuals. Furthermore, persistence is lower for individuals with very low wages than for those with substantially higher wage levels.
- Comparing our estimated flexible process to the canonical one typically used, we find that the latter overestimates average wage persistence for men and underestimates it for women.

A framework to study the implications of our new findings

- We construct a life-cycle model to study how our new findings on the dynamics of labour income in the UK affect the evaluation of tax and welfare policies compared to the canonical labour income process typically used.
- The model features households of different composition, in order to capture the following important channels. First, both the need for resources and the level of welfare benefits in the UK depend on family structure; namely, the presence of a

spouse and the number of dependent children. Secondly, allowing for both single and married households is crucial since pooling of labour income within families and adjustment of labour supply by secondary workers are important margins to reduce the impact of individual income fluctuations on household income and consumption level.

- The model features progressive taxes and captures, in a stylized way, the range of welfare benefits in the UK before the introduction of Universal Credit in 2016. We approximate the full range of benefits by two main ones: income support (an income floor) and in-work benefits. Both feature a maximum award and a (tapering) rate at which benefit are reduced as market income increases. While both benefits are means-tested, in-work benefits are conditional on working a minimum-number of hours.
- We estimate the model – alternatively under the canonical and flexible wage process – using household survey data. We target a number of statistics of the UK economy.

Evaluating alternative benefit reforms

- We use our model to evaluate alternative benefit reforms and understand how the evaluation is affected by the way the wage process is modelled.
- The first, hypothetical, reform chooses the structure of the two main benefits – income support and in-work benefits – to maximize (utilitarian) welfare in the economy. Namely, for either type of benefit, the maximum benefit level for single individuals and the rate at which benefits are withdrawn as income increases are chosen to maximize welfare, while leaving unchanged the spouse- and children-related component of benefits.
- The reform entails small welfare gains compared to the pre-reform, benchmark UK benefit configuration. The welfare gains are negligible for most types of households, with the exception of single women, who are the main beneficiaries. Although these welfare results are robust to the way the wage process is modelled, the optimal benefit configuration is very different under the two processes. Under the flexible process, which captures the rich wage dynamics in the data, the optimal benefit

configuration is very similar to the pre-reform one. In contrast, if one were to impose the counterfactual canonical wage process, one would conclude that optimal benefits during the same period should have been very different. In particular, the optimal policy would incorrectly prescribe a trebling of in-work benefits and their withdrawal at a much faster (from 40 to 100 per cent) rate. As a result, the canonical wage process implies a much larger increase in female labour force participation when moving to the optimal policy. The intuition is that the canonical process underestimates the average persistence of shocks to female wages, relative to the richer process. Since the more persistent shocks are, the more difficult it is to buffer their impact on consumption through borrowing, the optimal policy under the canonical process is skewed towards providing incentives to work rather than insurance against low labour income realisations.

- The second reform we consider mimics the switch to Universal Credit which was introduced in 2016 and completed in 2018. Universal Credit includes an earnings disregard for households with children and thus does not belong to the class of simple, linear benefit functions that we consider for optimality. We find that, irrespective of the wage process, the move to Universal Credit implies overall welfare gains which are larger than those of the optimal benefit system. However, this average improvement masks heterogeneous effects. The main beneficiaries are households with children, while singles without children lose out, and dramatically so in the case of women.

Policy conclusions

We find that the modelling of the wage process is important for the evaluation of benefit reform. Estimating a canonical process that does not allow the distribution of wage shocks to depend on age and past wage realisations implies an economically and statistically significant downward bias in the estimated average persistence of female wages. Underestimating the persistence of wage shocks overestimates the possibility of self-insure against them through borrowing. As a result, the optimal policy under the misspecified, restrictive process is much more skewed towards providing incentives to work rather than insurance, compared to the optimal policy under the richer wage process.

We also evaluate the 2016 wholesale benefit reform that introduced Universal Credit. Unlike both the pre-reform system and an optimized version of that system which uses the same policy tools, Universal Credit features an earnings disregard for households with children. Thanks to this extra margin, it implies welfare gains with respect to both the pre-reform and the optimized system. The average welfare improvement from Universal Credit are driven by households with children, who benefit from the earnings disregard. Single households without children lose out.

1 Introduction

The necessity, efficacy, and cost-effectiveness of government welfare policies depends both on the risks that households face and the actions that they can take to self-insure, for instance by adjusting their saving and labor supply. Wage risk is a key driver of household risk, while being single rather than in a couple is an important factor affecting a household's sources of risk and tools for self-insurance. This is because single earners are solely exposed to their own wage risk and can only use their own savings and labor supply to smooth consumption fluctuations. In contrast, couples face the wage risk of both household members, but can use their joint savings and the labor supply of both partners to at least partly counteract wage fluctuations. In addition, couples benefit from sharing resources, which gives rise to economies of scale in their consumption.

Better understanding the dynamics of wage and earnings risk is key to study the ability of households to self-insure and to properly design an efficient benefit system. In addition, explicitly modeling couples and singles, as well the dynamics of fertility and saving over the life-cycle, is crucial to understand how wage and earnings risks interact with self-insurance depending on family structure.

Our work builds on the literature that studies the effects of taxation and welfare policies taking into account household composition. A robust finding of this literature is the importance of the response of female labor supply to understand these effects. Keane and Wolpin (2010) study the effect of the US welfare system on women's welfare participation, labor supply, marriage, fertility, and schooling. Blundell, Costa Dias, Meghir and Shaw (2016) study how the UK tax and welfare system affects the career choices of women. Guner, Kaygusuz and Ventura (2012) and Bick and Fuchs-Schündeln (2017) investigate the effect of taxation on household labor supply while Borella, De Nardi and Yang (2016) study the effects of marriage-related taxes and Social Security rules for different cohorts of women whose labor supply behavior has been changing.

All previous work addressing these questions has maintained the assumption that labour income risk follows a "canonical" statistical process which assumes that the distribution of innovations to labour income is independent of age and the history of past shocks. A recent literature, though, has documented that these assumptions are violated and that earnings dynamics are much richer than typically assumed in this class of models. Guvenen, Karahan, Ozkan and Song (2016) document rich dynamics for pre-tax

individual earnings in the US, Arellano, Blundell and Bonhomme (2017) for household pre-tax earnings in the US and Norway, De Nardi, Fella and Paz-Pardo (2020) for household disposable earnings in the US and Ozkan, Storesletten, Holter and Halvorsen (2017) for household earnings in Norway both before and after taxes. De Nardi, Fella, Knoef, Paz-Pardo and Van Ooijen (2021) study the relative contributions of wages and hours to male earnings dynamics. De Nardi et al. (2020) is the only paper which studies household insurance against earnings shocks but in a model that abstracts from labor supply decisions and family structure.

Our work brings together the two strands of the literature discussed above by **investigating the importance of allowing for a rich dynamics of earnings and wages – together with heterogeneous family composition and labor supply – for the evaluation of welfare policies.**

2 Our research

This project makes two contributions. First, it documents, using two different datasets, that the rich features of earnings dynamics documented for other countries apply also to earnings and wage dynamics in the UK (Section 2.1).

Second, we build and estimate a structural life-cycle model with heterogeneous households and a labour income process that allows for the rich features of the data. We use the model to assess how the evaluation of alternative welfare policies is affected by the more sophisticated modelling of labour income dynamics (Section 2.2).

2.1 Empirical findings on labour income dynamics in the UK

2.1.1 Empirical question

The vast majority of the literature on the evaluation of welfare policies assumes that labour earnings, which are one of the main sources of income risk, follow a canonical process with a stationary – i.e. independent of age and past realizations – distribution. A recent and expanding empirical literature discussed in the previous section has documented that this is not the case. The dynamics of labour income is much richer than implied by the canonical model and displays substantial heterogeneity in the variance and persistence of shocks by age and past labour income. This is true across countries.

No such study has yet been conducted for the UK. Furthermore, most of the existing studies are concerned with the dynamics of earnings. Yet earnings reflect both exogenous shocks to wages and job loss, both actual sources of risk, and endogenous choices of hours worked which are not shocks and indeed may be adjusted in response to shocks. The distinction is particularly important in the case of women, who have who have a significantly larger labour supply elasticity than men.

In the empirical part of our project we investigate whether moments of the distribution of male earnings and female wages in the UK display richer dynamics than that implied by a canonical model. Further, we use novel statistical techniques to estimate rich stochastic processes for earnings and wages that allow for deviations from the canonical process. Finally, we compare estimates of relevant moments of earnings and wages under the richer process and under the canonical one.

2.1.2 Data

The analysis is based on two datasets with different strengths.

The first dataset is the New Earnings Survey Panel Dataset for the United Kingdom. It is a mandatory yearly, administrative survey, filled by employers, whose aim is to gather data on earnings, wages and hours for a 1% random sample of the UK population.¹ The fact that it is the employer rather than the employee who replies to the survey implies that the data is of high quality and less subject to measurement error or misreporting than standard household surveys. The most important limitation of the NESPD is that it has a 25-30% employer non-response rate, which implies that it may not be fully representative, if non-responses are non-random. Second, it does not allow to distinguish individuals who are not working from individuals whose employers do not respond to the survey.

Our second dataset is the British Household Panel Survey (BHPS), one of the most widely known household surveys in the UK. As all household survey, it has the limitation that all answers are self-reported and thus potentially subject to measurement error.²

¹Individuals whose National Insurance Number (NIN) ends in a certain set of two digits are automatically selected for the sample.

²However, the design of the survey suggests that measurement error in earnings is likely to be lower than in other surveys, like the PSID in the US, because instead of 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 himself saw the pay slip.

On the plus side, the BHPS includes a wide variety of information (such as off-sample labor market histories) in addition to income data. Furthermore, unlike the NESPD, it collects information on all household members, rather than a single individual, and is thus more suited for the study of family and government insurance. Importantly, despite taxation being at the individual level, most subsidies and benefits in the UK are linked to the family unit.

For the purpose of comparing the two datasets we use data from 1996 to 2006.³

The top panel of Figure 1 plots the standard deviation of male earnings changes against the percentile of last period's earnings. In both data sets, the standard deviation follows a U-shaped pattern which is inconsistent with the assumption of constant variance which underpins the canonical model.

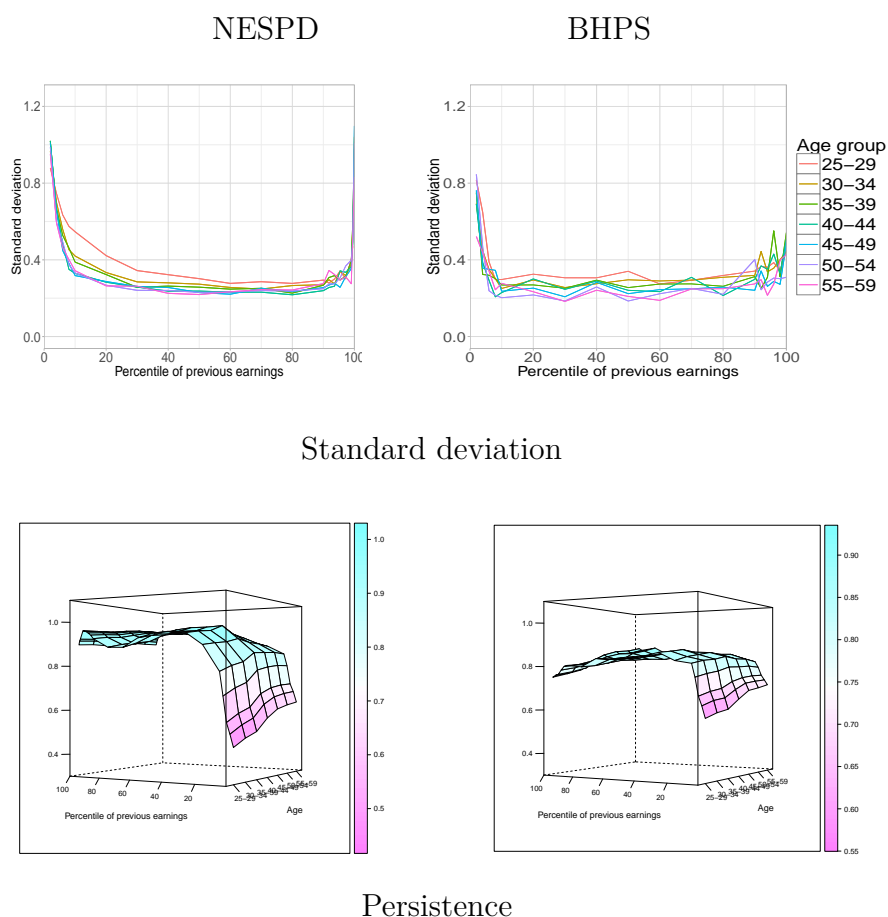


Figure 1: Moments of male earnings changes in the BHPS and NESPD. Top panel: by previous earnings. Bottom panel, by previous earnings and age

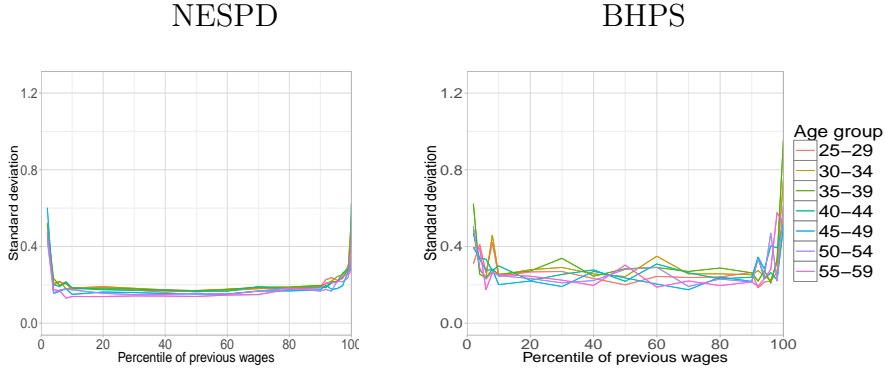
³The choice of sample period is dictated by three considerations. First, annual earnings only start being available in the NESPD after 1996. Second, up to the mid-90s there were many changes in the UK labor market (e.g. deunionization) that could confound the analysis.

The bottom panel plots the persistence of male earnings as a function of age and the percentile of the previous earnings realization. As for the variance, persistence is not independent of previous earnings levels (or age) which again is inconsistent with the constant persistence of the canonical model. More specifically, the picture shows that the persistence of male earnings is lowest at young ages and low earnings levels, with a persistence of about 0.55, compared with a persistence above 0.85 for people of the same age but above the median of previous earnings. The data indicate lower average persistence in the BHPS, which could be the result of larger measurement error in the BHPS earnings data than in the NESPD. Our estimation procedure will extract this measurement error from our data.

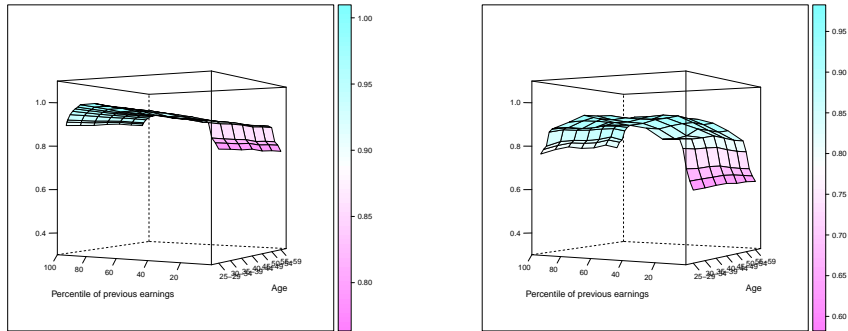
Turning to female wage changes⁴, the top panel of Figure 2 plots the standard deviation as a function of the rank of the previous period's realization. Qualitatively, their properties are remarkably similar to those of male earnings changes: the variance has a U-shaped pattern. The bottom panel of the same picture, instead, plots the persistence of female wages as a function of age and the percentile of the previous wage realization. Similarly to male earnings, the pattern of persistence is inconsistent with the standard, canonical model. Persistence is hump shaped as a function of a previous wage realization, though it displays much less variability with respect to age than in the case of male wages. Again, the data indicate lower average persistence in the BHPS, which could be the result of larger measurement error in the BHPS earnings data than in the NESPD.

Two key messages emerge from these pictures. First, both male earnings and female wages display strong deviations from the assumptions of the canonical model. The fact that these patterns apply to both earnings and wages suggest that they are genuine features of earnings risk rather than just reflecting labour supply response. Second, our findings are very similar across both datasets. Because of this and because the BHPS has the advantage of containing demographics information which is crucial for the calibration of our model which features the household as the relevant economic unit, we use the BHPS (in its full sample, 1991-2008) in the analysis that follows.

⁴We document that similar features hold for female earnings changes.



Standard deviation



Persistence

Figure 2: Moments of female wage changes in the BHPS and NESPD. Top panel: by previous earnings. Bottom panel, by previous earnings and age

2.1.3 Comparing labour income processes

As we have discussed in the previous section, the canonical assumption that the distribution of shocks to labour income is independent of age and previous history is at odds with the data. For this reason, we use a methodology proposed in Arellano et al. (2017) to estimate an alternative *flexible* process that can capture the rich features of men’s and women’s wages we have described. For comparison we also estimate a typical *canonical* process on the same data.

Both processes share the common assumption, that the stochastic wage component y_{it} for individual i of age t^5 can be decomposed as

$$y_{it} = \eta_{it} + \epsilon_{it}. \quad (1)$$

⁵We omit the gender superscript in what follows to streamline notation.




Figure 3: Variance of log earnings for men (left) and log wages for women (right).

The first component, η_{it} , is assumed to be *persistent* and to follow a first-order Markov process. The second component, ϵ_{it} , is assumed to be *transitory* and independent over time and of η_{is} for all s .

The *canonical* (linear) model, first proposed by Abowd and Card (1989) and commonly used in macroeconomics, assumes that the distributions of the two components have the following functional forms

$$\begin{aligned} \eta_{i,t} &= \rho\eta_{i,t-1} + \zeta_{it}, \\ \eta_{i1} &\overset{iid}{\sim} N(0, \sigma_{\eta_1}), \quad \zeta_{it} \overset{iid}{\sim} N(0, \sigma_{\zeta}), \quad \epsilon_{it} \overset{iid}{\sim} N(0, \sigma_{\epsilon}). \end{aligned} \tag{2}$$

The above formulas imply that both components have an *age-independent, normal distribution* and that the persistent component η_{it} is a *linear* autoregressive process, which implies that its conditional second and higher order moment are independent of its previous realization $\eta_{i,t-1}$.

Instead, the methodology proposed in Arellano et al. (2017) which we use does not impose any of these assumptions, but lets the data determine the shape of the distribution.

Figure 3 compares the age profiles for the cross-sectional variance of log earnings for men and log wages for women to those implied by the canonical and flexible processes. It is clear that the flexible process (red line) tracks the life-cycle evolution of the cross-sectional variance in the data (blue line) significantly better than the canonical one (yellow one). This is particularly true in the case of female wages.

Finally, Figure 4 compares the estimated persistence and variance of the persistent component of male earnings and female wages for the canonical and the flexible process.

The main takeaway is that in the case of male earnings the average estimated persistence is substantially lower under the flexible process while the innovation variance is higher. These findings are reversed in the case of female wages which feature much larger average persistence and lower innovation variance. Importantly, the differences between the estimates under the two processes are statistically significant.

The differences in the estimated persistence of shocks implied by the two methods are potentially important, not only from a statistical, but also from an economic perspective. More persistent shocks are more difficult to self-insure through household borrowing and therefore imply a bigger role for complementary forms of insurance, such as public insurance. Our findings suggest that the canonical process overestimates labour income risk for men and underestimates it for women. This raises the question of the extent to which these differences are important for the evaluation of welfare policies aimed at insurances against income risk. It is this question that we address in the second part of the paper.

2.2 Implications of the wage process for the evaluation of welfare policies

As we have noted above, existing studies of the welfare cost of labour income risk and the role of government in providing partial insurance against it through the benefit system have assumed that labour income follows a canonical process. We want to understand how a more realistic modelling of the labour income process quantitatively affects the answers to those questions. In order to do so, we build and estimate of structural life-cycle model in the vein of those typically used in this literature.

The model features households going through a realistic life-cycle. Agents start their working life at age 25 and are heterogeneous in gender, marital status (single or married), number of children, and the initial wage realization. The dynamics of earnings and wage shocks is described, alternatively, by the two processes we have discussed in the previous section. Children are born stochastically into single-female and married household. The probability that children arrive and leave a household depends on their mother's age, marital status, and the number of children already in the household. Children increase household consumption needs, entail child care costs if their mother works, and matter for benefit eligibility. In each period, households choose how much to consume and save.




Figure 4: Persistence by age (top), by earnings and quantile of the shock (middle), and standard deviation of shocks (bottom), for male earnings (left) and women's wages (right), NL process vs canonical process, persistent component. Dotted lines and transparent surfaces represent 95% confidence intervals computed by bootstrapping

Women of working age also choose whether to work part-time, full-time or not at all, while men work a fixed number of hours.

This setup is meant to capture, in a reasonably tractable way, a number of important features of reality. First, both the need for resources and the level of welfare benefits in the UK depend on family structure; namely, the presence of a spouse and the number of dependent children. Secondly, allowing for both single and married households is crucial since pooling of labour income within families is an important margin of insurance at the household level. The decision to restrict an active labour supply choice to women reflects the fact that the vast majority of men work full time, while women have a lower attachment to the labour market and substantially higher labour supply elasticity. Third,

we allow the initial wage realization for people in couples to be correlated to capture the extent of marital sorting on the basis of income observed in the data.

We capture the progressive structure of income taxation in the UK by estimating a tax function with enough flexibility to capture the degree of progressivity in the data. We approximate the UK benefit system, before the 2016 reform that introduced Universal Credit, through two main types of benefits: income support, which is unrelated to labor market participation and in-work benefits, which are conditional on working at least part-time. Both types of benefits have a baseline (maximum) level which depends on the presence of a spouse and the number of children in the household and are withdrawn at some tapering rate as pre-tax labour income increases.

For in-work benefits in our benchmark benefit system, we follow the statutory rules of the Working Tax Credit. The base benefit level is £1,960 plus an extra £2,100 for households with children. Benefits are withdrawn at a rate of 41p per each pound of pre-tax labour income.

Our income-support programme is meant to replicate many benefits available to low-income households. These programs have differential take-up rates and eligibility criteria which would be very complicated to explicitly model individually. Therefore, we use benefit data available in the BHPS and in the BHPS Derived Net Household Income Variables to estimate maximum benefit levels. More specifically, we look at average benefit receipts for households whose labor income in a given year is close to zero (below £2,000, although results are robust to changing the threshold to £1,000 or £3,000). This approach allows us to average across various types of benefits and weighting by the cross-sectional distribution of benefit receipts within this subset of the population. However, we cannot use the same approach to directly compute the tapering rate from benefit data. Most benefits have weekly or monthly eligibility criteria, while our data are annual, so the actual relationship between income and benefits received gets attenuated in the data due to time aggregation. For this reason, we estimate the withdrawal rate used in the model as a weighted average of the statutory tapering rates of the relevant benefits taking into account cross-eligibility criteria and legal thresholds. The associate estimates of the maximum benefit level for income support are £4,574 for singles plus an extra £1,366 for couples and additional £907 per child. The benefit withdrawal rate is 70p per pound of pre-tax income.

Parameters	Benchmark	Optimum (NL)	Optimum (Ca)
Income floor			
Base level (singles)	4574	4504	3106
Couple supplement	1366	1366	1366
Per-child supplement	907	907	907
Withdrawal rate	70%	62%	56%
In-work benefit			
Base level	1960	2550	7500
Child supplement	2100	2100	2100
Withdrawal rate	41%	44%	100%

Table 1: Income floors and in-work benefits: benchmark vs optimum under NL and canonical processes

The model is estimated under the two alternative stochastic processes for labour income discussed in the previous section: our rich, flexible process and the canonical process. We use the resulting two versions of the model to assess the importance of the labour income process for the evaluation of welfare policies. We evaluate two policy reforms. The first one is a hypothetical reform that sets the maximum benefit level for singles and the tapering rate for both types of benefits so as to maximize aggregate welfare in the economy (Section 2.2.1). The second one is the actual reform, introduced in 2016 and completed in 2018, which saw the unification of various benefits under a single one, Universal Credit. (Section 2.2.2).

To evaluate whether benefit reforms have different implications under the canonical and nonlinear processes for earnings and wages, we look at both outcomes and welfare. Our welfare criterion is given by the utilitarian, un-weighted, average of the lifetime utilities of newborns. We report results both behind the full veil of ignorance and after the realization of gender, marital status and number of children.

2.2.1 Optimal benefit reform

In this section, we evaluate a welfare-maximizing reform of the UK welfare system before the 2016 reform that saw introduced Universal Credit. Namely, we optimize over the parameters of two stylised benefits that capture the UK welfare system before 2016: an income floor and an in-work benefit. For each of the two benefits, we optimize over the withdrawal rate and maximum benefit level for singles, leaving unchanged the supplements for married couples and children. We choose the value of these four policy

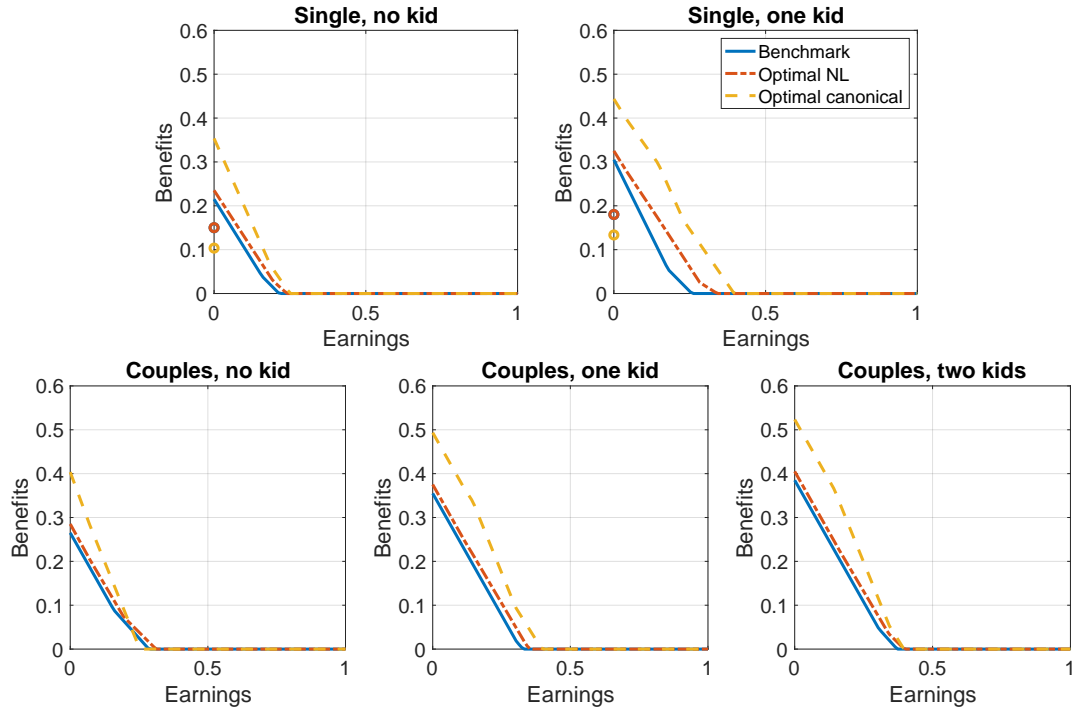


Figure 5: Implied total level of benefits, by pre-tax income levels, marital status, and number of children. For singles, circles represent benefits for households where everyone is out of work, while lines represent benefits for households in which at least one member works. Benefits are expressed as a share of average male earnings

instruments that maximize ex-ante welfare (under the veil of ignorance) while maintaining the tax function unchanged and keeping total tax revenues minus total benefit outlays constant (i.e., the change is budget neutral for the government).

Table 1 shows the results of this optimization. Column 2 reports the parameter values for the two benefit functions in the benchmark economy, while columns 3 and 4 report the optimal values under, respectively, the richer, non-linear (NL) and canonical (Ca) wage process. Under the NL wage process, the optimal benefit system is quite close to the one in the benchmark economy. The main difference is that it features a 30 per cent higher level of in-work benefits, as well as a slightly higher withdrawal rate for in-work benefits and a lower one for the income floor. The difference between the optimal and the benchmark benefit policies is possibly best appreciated with the help of Figure 5 which plots the relationship between benefit levels and pre-tax labor income for single men, women and couples in the benchmark (solid blue lines) and under the optimal system under the NL (dash-dotted red lines) and canonical (dashed yellow lines) wage processes. The lines plot benefit levels for working individuals, while the circles in the

top two panels denote benefits for non-working individuals (single women in the model). Under the NL wage process, benefits for working households are marginally higher than in the benchmark and are exhausted at a slightly higher level of disposable income due to the fall in the tapering rate for income-support. Single, working women with one child between the 10th and 30th percentile of labor income are the only group for which the switch to the optimal policy implies a significant increase in total benefits.

Under the canonical wage process, instead, the optimal benefit system is substantially different from the one in the benchmark, as can be seen by comparing columns 2 and 4 in Table 1. In particular, the optimal system implies a 30 per cent reduction in benefit levels for non-working individuals, from £4,500 to £3,100, accompanied by a more than three-fold increase in the level of in-work benefits. As a result, the net return to the first pound of labor income—the difference between the vertical intercept of the straight line and of the corresponding circle in Figure 5—is three times as large as in both the benchmark and the optimal system under the NL process. Similarly, the reduction in the withdrawal rate for the income floor and its increase for in-work benefits are larger compared the NL case. In particular, the withdrawal rate for in-work benefits increases from 41 to 100 per cent. The net effect is a substantial increase in benefits for individuals working at least part-time.

In order to understand the substantial difference in the policy prescription between the two labour income processes it is important to keep in mind that the optimal benefit system reflects a trade-off between insurance provision, which implies redistribution towards those who have low income and/or high needs, and work incentives.

To understand how the two types of benefits feature in this trade-off, it is useful to consider their introduction, one at time, starting from income support. Consider first the case in which income support takes the form of a social dividend to all individuals, what is often referred to as Universal (i.e. not means-tested) Basic Income, while keeping the income tax function unchanged.⁶ Since the social dividend would reduce incentives to work, through the positive income effect, it would not be self-financing. Therefore, in order to provide a meaningful income floor to low income individuals *and* be self-financing, the transfer has to be mean-tested rather than universal, that is phased out at higher

⁶The main objection to a Universal Basic Income is that the tax rate necessary to finance a reasonable Basic Income would be so high as to significantly reduce incentives to work for many individuals. For this reason, we take the tax function as given here.

income levels according to some, sufficiently high, withdrawal rate. The withdrawal rate, though, is effectively a marginal tax rate on low-income individuals and reduces incentives to work through a negative substitution effect. This is the trade-off associated with the income support benefit studied here. If it is not withdrawn fast enough it is not self-financing. If, instead, it tapers quickly it creates significant disincentives to work for the poor.

Consider now the introduction of a means-tested in-work benefit, conditional on working at least part-time, which is withdrawn at its own rate. The benefit increases incentives to work compared to not working but reduces incentives to work full-time as opposed to part-time. Individuals working full-time face both a positive income effect and a negative substitution effect, the latter stemming from the withdrawal rate of in-work benefits further increasing the effective marginal tax rate for benefit claimants. Hence, in-work benefits trade off higher labour market participation for a reduction in full-time work.

Note that, in the absence of income support, in-work benefits provide very limited insurance because, by being conditional on work, they do not compensate individuals those market wages are too low relative to the disutility of work.

Finally, the desirability of government-provided insurance depends on individuals' ability to self-insure against negative income shocks through, for example, dissaving. If being on low labour income is mostly a transitory state then self-insurance through dissaving is an effective means to smooth consumption in the face of labour income fluctuations. Vice versa, if shocks are fully permanent they cannot be smoothed through self-insurance and, in the absence of government insurance, consumption fluctuates one-to-one with labour income. Therefore, the ability to smooth consumption in the face of labour income shocks is decreasing in the persistence of such shocks.

The intuition for why, relative to the benchmark, the optimal system is substantially tilted towards in-work benefits in the case of the canonical process but not the rich one is that the canonical process underestimates wage persistence for women. Thus, the welfare cost of reducing insurance to low-wage women by cutting income support is lower under the canonical process, because it implies low wages are a more transitory state, against which it is easier to self-insure. On the other hand, the richer wage process replicates the fact that low-wage status is a relatively persistent state. Thus, reducing income support for women on low wages in order to induce them to work would really reduce their welfare,

Group	Average	0 kids	1 kid	2 kids	3+ kids
NL process					
Overall	0.15				
Single men	0.17				
Single women	0.42	0.34	0.59	0.55	0.47
Couples	0.06	0.05	0.06	0.07	0.07
Canonical process					
Overall	0.19				
Single men	0.11				
Single women	0.52	-0.16	2.32	1.48	0.85
Couples	0.10	0.08	0.11	0.12	0.12

Table 2: Welfare change, by gender and marital status, for switch to optimal system.

and the increase in in-work benefits would not be enough to compensate for the welfare cost of foregone leisure. As a result, the optimal welfare system under the more realistic NL process does not imply these major changes and is much closer to the system that was in place before 2016.

Figures 6 and 7 show how, under either wage process, the optimal policy mix results in higher part-time and lower full-time labor market participation by single women, with a significant increase in participation overall. The rise in overall participation is driven by the increase in the relative return to work, as measured by the difference between the level of in-work benefits (the circles) and total benefits (the corresponding lines) in Figure 5. As we have discussed, the increase is particularly large under the canonical wage process and accounts for the much larger rise in participation. As discussed above, the switch from full-time to part-time is due to the fact that the higher withdrawal rate for in-work benefits increases the effective tax rate and reduces the benefits of working full-time rather than part-time. The effect is, again, more pronounced under the canonical process which features an order of magnitude larger increases in the withdrawal rate for in-work benefits compared to the NL process.

Table 2 reports the welfare change associated with the switch to the optimal benefit system under either wage processes. The welfare change is expressed as the percentage change in consumption (constant across ages and states) that would make a 25 year old (the initial age in the model) agent in the benchmark economy indifferent to being born in the counterfactual economy. The first column reports the welfare change from the perspective of after the realization of gender and marital status, but before the draw of the initial number of children. Finally, the “overall” measure in the first row is under the




Figure 6: Labor force participation under NL process: optimal benefit system.

full veil of ignorance, including the realization of the gender and marital status draw.

Despite the substantial differences in the optimal benefit policies across the two wage processes, the overall welfare gains associated with moving from the benchmark to the optimal system are pretty similar under the two wage processes. In both cases, the switch to the optimal benefit system implies an increase in welfare of approximately 0.2 percentage points. Under both wage processes, the main beneficiaries of the reform are single women, whose welfare gain is more than twice the overall one, while couples are hardly affected. It is this difference that drives the differences in ex-ante welfare between the two economies. Among single women, though, the two wage processes imply a very different distribution of welfare gains and losses. The gains are distributed rather evenly under the NL process, but accrue to single women with children, at the expense of single women without, under the canonical process. The intuition is the following. As we discussed above, under the canonical process the benefit reform involves a substantial shift from income support to in-work benefits. In-work benefits entail an additional £2,010




Figure 7: Labor force participation under canonical process: optimal benefit system.

for households with children, which accounts for the substantially bigger gap between the yellow and the blue line in the top-right panel of Figure 5 compared to the top-left one. As women respond to the higher incentives to work, the child-related component of in-work benefits more than compensates, on average, single women with children who switch to working under the reform. On the other hand, for single women without children the lower utility of leisure associated with post-reform labor market participation is not compensated by the child-related transfer.

To sum up, an important message of the comparison of welfare maximizing benefit configurations is that imposing the misspecified canonical wage process implies policy prescriptions that are very different from the actual policies in place before the 2016 reform. Instead, the pre-reform system is close to the optimal one implied by the model when the latter is fed our richer wage process.

Parameters	
Base level (singles)	3886
Couple supplement	2210
Per-child (up to 2) supplement	1730
Earnings disregard	2304
Withdrawal rate	63%

Table 3: Universal Credit parameters




Figure 8: Implied total level of benefits, by pre-tax income level, comparing benchmark vs. Universal Credit. For singles, circles represent benefit entitlement for non-working individuals under the benchmark.

2.2.2 Universal Credit

The combination of key benefits (Income-Based JSA, Income-Related Employment and Support Allowance, Income Support, Working Tax Credit, Child Tax Credit and Housing Benefits) that we have modelled in the benchmark economy described was replaced, with the exception of Child Benefits, by a unified benefit system: Universal Credit. Universal Credit was first piloted in 2013 in a few regions and then gradually rolled out to all of Great Britain from May 2016 to December 2018. The aim of this section is to analyse the changes in labour market participation and welfare under Universal Credit relative to the pre-reform system in the benchmark economy. For comparability with the




Figure 9: Labor force participation under NL process: Universal Credit vs optimal benefits.

analysis in Section 2.2.1, we impose that the introduction of Universal Credit is budget neutral by scaling all statutory allowances proportionally by 0.86. Table 3 reports the corresponding parameters.

Two features of Universal Credit are worth pointing out. First, it features a £2304 earnings disregard for families with children. Second, benefits are withdrawn as a function of *after-tax* income, rather than pre-tax income in the pre-reform system.

Figure 8 reports benefits levels as a fraction of pre-tax income in our benchmark economy and under Universal Credit. The main takeaway is that, compared to our benchmark, Universal Credit entails lower benefits for households without children and for very low-income couples with children, and higher benefits for the rest of households with children.

Figure 9 compares labor force participation under Universal Credit and in the benchmark, pre-reform, benefit system under the richer wage processes. Universal Credit substantially reduces benefits for single women with no children. This results in higher

Group	Average	0 kids	1 kid	2 kids	3+ kids
Overall	0.45				
Single men	-0.07				
Single women	-0.15	-1.65	2.44	3.98	3.44
Couples	0.69	0.44	0.75	0.97	1.02

Table 4: Welfare change for switch to Universal Credit. NL process.

(part-time) labor force participation from middle age onward, when children have left the household. As a result, part-time labor force participation at older ages is higher than in the optimal formulation of the previous system. In contrast, participation is lower for single women between age 35 and 45, compared to the optimal system, because a significant number of them have children and receive higher benefits independently from labor force participation under Universal Credit than in our benchmark system.

Table 4 reports the steady-state changes in welfare associated with switching from the benchmark pre-reform benefit configuration to Universal Credit. The switch to Universal Credit entails an increase in average (overall) welfare of 0.45 percentage points. Looking into the distribution of the overall gains, though, reveals that the main beneficiaries are households with children, while singles without children are worse off. In particular, while the welfare loss for single men is relatively low, single women with children experience a very sizeable welfare loss of 1.65 percentage points. The implications under the canonical processes are similar.

It may seem surprising that, unlike singles, couples with no children benefit from the reform. In fact, one has to realize that welfare in Table 4 is from the perspective of households at age 25. So the difference between couples and single women with no children is that the former have a much higher probability than the latter of having children, and therefore benefit from the earnings disregard, later in life. This drives their higher welfare in ex ante, expected, terms.

It is instructive to compare these welfare changes with those associated with a switch to the optimal pre-UC benefit system⁷ in Table 2. Relative to the benchmark, the optimal pre-UC system redistributes more towards singles, particularly single women, while Universal Credit mostly benefits couples. Since couples constitute the majority of the population, Universal Credit entails a two-to-three times larger welfare increase under

⁷It is worth point out that, as we have shown in the previous section, that the benchmark benefit system that was in place pre-Universal Credit is also close to the optimal one that we compute by optimizing over the same set of policy instruments.

the utilitarian (i.e. unweighted) social welfare criterion.

Contrary to the optimal benefit reform, though, Universal Credit *reduces* overall benefits for singles without children at every level of income (top left panel in Figure 8) relative to the benchmark. This implies a negative income effect which induces higher labor force participation, and lower welfare, for single working women on low income.

The reason why Universal Credit can achieve higher average welfare than the optimal benefit reform in Section 2.2.1 is optimal within the class of linear in-work and income-floor benefit functions. The non-linear Universal Credit benefit function is therefore not nested in that class because it also adopts income disregards and asset testing.

3 Conclusions

A growing body of empirical work that takes advantage of large, administrative datasets and new statistical techniques provides evidence that households' labour income dynamics are substantially richer than those implied by the *canonical* income process – with constant variance and persistence – that are typically used in studies that evaluate welfare policies.

This paper is the first to establish that the rich dynamics of labour income documented for other countries also applies to the UK. Rather than being constant, the variance and persistence of labour earnings display substantial differences by age and labour income history. These rich dynamics are a feature not only of earnings, but also of wages. Hence, they reflect genuine labour income risk rather than being merely the byproduct of the adjustment of hours to wage shocks.

We show that ignoring such richer dynamics when estimating stochastic labour income processes implies biased estimates of important moments on the data. In particular, the canonical model underestimates the persistence of shocks to female wages and overestimates the persistence of shocks to male earnings relative to a richer, flexible earnings process which does not impose the constancy of variance and persistence.

Correctly estimating the persistence of labour income shocks is important to capture labour income risk, because persistence crucially affects agents' ability to insulate consumption from income shocks through dissaving (self-insurance). This is why we investigate how allowing for a richer labour income dynamics affects the evaluation of

welfare policies compared to the standard, canonical income process. To do so, we build and estimate a structural life-cycle model in the vein of those typically used in the literature on the evaluation of tax and welfare policies. Our model features households with different family structure to capture the following important channels. First, both the need for resources and the level of welfare benefits in the UK depend on family structure; that is, the presence of a spouse and the number of dependent children. Second, allowing for both single and married households is crucial because labour income pooling within families and the possibility of adjusting the labour supply of the secondary earner are important margins of insurance at the household level.

We use our model to evaluate to evaluate alternative benefit reforms under richer and canonical labour income processes. Our findings confirm that correctly capturing the dynamics of labour income is important to evaluate the costs and benefits of welfare policies. In particular, we analyse a hypothetical reform that chooses the structure of the two main benefits – income support and in-work benefits – to maximize (utilitarian) welfare in the economy. Namely, for either type of benefit, the maximum benefit level for single individuals and the rate at which benefits are withdrawn as income increases are chosen to maximize welfare, while leaving unchanged the spouse- and children-related component of benefits. This reform entails small welfare gains compared to the pre-2016, benchmark UK benefit configuration. Although the welfare results are robust to the way the earnings process is modelled, the optimal benefit configuration is very different between the canonical and flexible risks. Under the flexible earnings process, the optimal benefit configuration is very similar to the pre-reform one. In contrast, if one were to ignore the rich wage dynamics that we estimate from the data and simply assume a canonical wage process, one would conclude that optimal benefits during the same period should have been very different. In particular, the optimal policy would incorrectly prescribe a trebling of in-work benefits and a much faster (from 40 to 100 per cent) withdrawal rate for benefits. As a result, the canonical wage implies a much larger increase in female labour force participation when moving to the optimal policy. The intuition is that the canonical wage process underestimates the average persistence of shocks to female wages, relative to the richer process. Since more transitory shocks are easier to self-insure through borrowing, the optimal policy under the canonical process is skewed towards providing incentives to work rather than insurance against low labour

income realisations.

We also consider a reform that mimics the switch to Universal Credit which was introduced in 2016 and completed in 2018. Universal Credit includes an earnings disregard for households with children and thus does not belong to the class of linear benefit functions that we consider for optimality. We find that, thanks to this additional margin the move to Universal Credit implies overall welfare gains which are larger than those under the optimal benefit system. However, this average improvement masks heterogeneous effects. The main beneficiaries are households with children, who benefit from the earnings disregard, while singles without children lose out.

References

- Abowd, John and Card, David (1989), ‘On the covariance structure of earnings and hours changes’, *Econometrica* **57**(2), 411–45.
- Arellano, Manuel, Blundell, Richard and Bonhomme, Stéphane (2017), ‘Earnings and consumption dynamics: A non-linear panel data framework’, *Econometrica* **85**(3), 693–734.
- Bick, Alexander and Fuchs-Schündeln, Nicola (2017), ‘Taxation and Labour Supply of Married Couples across Countries: A Macroeconomic Analysis’, *The Review of Economic Studies* **85**(3), 1543–1576.
- Blundell, Richard, Costa Dias, Monica, Meghir, Costas and Shaw, Jonathan (2016), ‘Female labor supply, human capital, and welfare reform’, *Econometrica* **84**(5), 1705–1753.
- Borella, Margherita, De Nardi, Mariacristina and Yang, Fang (2016), ‘Marriage-related policies in an estimated life-cycle model of households’ labor supply and savings for two cohorts’, Working paper.
- De Nardi, Mariacristina, Fella, Giulio and Paz-Pardo, Gonzalo (2020), ‘Nonlinear household earnings dynamics, self-insurance, and welfare’, *Journal of the European Economic Association* **18**(2), 890–926.
- De Nardi, Mariacristina, Fella, Giulio, Knoef, Marike, Paz-Pardo, Gonzalo and Van Ooijen, Raun (2021), ‘Family and government insurance: Wage, earnings, and income risks in the netherlands and the us’, *Journal of Public Economics* **193**, 104327.
- Guner, Nezih, Kaygusuz, Remzi and Ventura, Gustavo (2012), ‘Taxation and Household Labour Supply’, *The Review of Economic Studies* **79**(3), 1113–1149.
- Guvenen, Fatih, Karahan, Fatih, Ozkan, Serdar and Song, Jae (2016), What do data on millions of U.S. workers reveal about life-cycle earnings risk? Working paper, University of Minnesota.

Keane, Michael P. and Wolpin, Kenneth I. (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’, *International Economic Review* **51**(3), 851–892.

Ozkan, Serdar, Storesletten, Kjetil, Holter, Hans and Halvorsen, Elin (2017), Dissecting Idiosyncratic Income Risk, 2017 meeting papers, Society for Economic Dynamics.

A Conference Programme

Earnings, Risk, and Insurance

London, April 26th 2019

Venue: Charter Suite CILIP (IFS building), 7 Ridgmount Street, London WC1E 7AE

Organizers: Mariacristina De Nardi (Minneapolis Fed and UCL) and Giulio Fella (QMUL)

9.00-9.30 **Registration and Coffee**

9.30-10.10 Mariacristina De Nardi (Minneapolis Fed), **Jul Fella (MALE)** and Gonzalo Pardo (UCL) “Household Earnings Risk, Government Policy, and Welfare.”

10.10-10.25 **Hamish Low** (Oxford) Discussion

10.30-11.10 Richard Blundell (UCL), Margherita Borella (Torino), **Jeanne Commault** (Sciences Po), Mariacristina De Nardi (Minneapolis Fed) “Old Age Risks and Responses.”

11.10-11.25 **Eric French** (UCL) Discussion

11.30-12.00 **Coffee break**

12.00-12.40 **Luigi Pistaferri** (Stanford) and Hamish Low (Oxford) “Efficiency Aspects of Disability Insurance: Evidence from Merged Survey-Administrative Data.”

12.40-12.55 **Amanda Michaud** (UWO) Discussion

13.00-14.30 **Lunch**

14.30-15.10 Tetsuya Kaji, **Elena Manresa** (NYU), and Guillaume Pouliot “Artificial Intelligence for Structural Estimation.”

15.10-15.25 **Aureo de Paula** (UCL) Discussion

- 15.30-16.10 Michael Keane (UNSW), **Zvi Eckstein (IDC)**, and Osnat Lifshitz (IDC) “Career and Family Decisions: Cohorts Born 1935–1975.”
- 16.10-16.25 **Margherita Borella** (Torino) Discussion
- 16.30-17.00 **Coffee break**
- 17.00-17.40 Gregory Jolivet (Bristol) and **Fabien Postel-Vinay** (UCL) “A Structural Analysis of Health and Labor Market Trajectories.”
- 17.40-17.55 **Joseph Mullins (UMN)** Discussion
- 19.00 **Dinner**