The Impact of Aggregate Fluctuations Across the UK Income Distribution

Tomas Key* Jamie Lenney†

July 2024

Abstract

In this paper, we examine the response of earnings and employment to fluctuations in aggregate economic activity (GDP) across the income distribution. Using data from the UK's Labour Force Survey, we present evidence that aggregate fluctuations have economically significant but heterogeneous impacts across the income distribution. Sensitivity is greatest at the very bottom (first decile) of the income distribution and smallest in the upper middle (seventh and eight deciles) of the distribution. The transmission of GDP fluctuations also differs across the income distribution. Changes to hours worked and employment explain the majority of the labour earnings response in the bottom half of the distribution, whereas changes to the hourly wage are more important in the top half. In a further decomposition, we show that the changes to employment are largely due to fluctuations in the employment to unemployment transition rate. We also find that GDP fluctuations are positively correlated with job switching in the bottom half of the distribution.

*Bank of England: tomas.key@bankofengland.co.uk

†Bank of England: jamie.lenney@bankofengland.co.uk

1 Introduction

Understanding the incidence and impact of fluctuations in aggregate economic activity across the distribution of individuals and households is important when considering the welfare implications of business cycles and the macroeconomic policies that are enacted in response to these fluctuations.

Recent studies, notably Guvenen et al. (2017), have begun to explore these distributional questions. Our aim is to contribute to this literature by focusing on two areas that have received relatively less attention: a decomposition of changes in labor income over the business cycle and income distribution into changes in employment, hours worked, and hourly pay; and an examination of these phenomena within the context of the UK economy.

We do so using microdata from the UK's Labour Force Survey (LFS), a large household survey which underpins the UK's national labour market statistics. A particular advantage of the LFS for our purposes is its short panel element, whereby individuals remain in the survey for up to five quarters. That enables us to observe how changes to the labour income and employment status of individuals with different initial circumstances are related to changes in the aggregate economy.

The specific approach that we take is to estimate OLS regressions of four quarter changes to different individual outcome variables on four quarter changes to Gross Domestic Product (GDP) over the same period. We show that the relationships between changes to GDP and the outcome variables that we consider can be best interpreted as elasticities with respect to transitory fluctuations in GDP. The main outcome variables that we consider are labour earnings, hours worked and labour market status, including whether individuals change job. We estimate these regressions separately by initial pay decile in order to assess which parts of the income distribution are most affected by aggregate fluctuations and along what margins. Using the estimates from these regressions we are able to approximately decompose changes in labour earnings across the income distribution into intensive and extensive margin contributions.

We find that temporary fluctuations in GDP have economically significant but heterogeneous effects on real pay across the income distribution. While on average real pay responds by about 0.7 per cent to a 1 per cent movement in GDP, we find the largest effects are at the very bottom of the income distribution (first decile) and the smallest effects are in the upper middle of the distribution (eighth decile). The reasons why earnings change also differs across the income distribution. Those in the bottom half of the income distribution largely experience adjustment through changes to their working hours, and a significant share of that change is accounted for by changes to the likelihood of them becoming unemployed. Those in the upper middle and top of the distribution see adjustment to their pay mainly through changes to their hourly wage. We also find that the response of the job switching probability to movements in GDP is greater for those in the bottom half of the income distribution than those in the top half.

The fact that those in the bottom half of the income distribution are more likely to move into and out of employment is also true on average, or unconditionally. An important difference that we uncover is the margin of adjustment. In general, most movements from employment to non-employment are due to transitions into inactivity. However, conditional on GDP fluctuations, it is the unemployment margin that dominates, not the inactivity margin, which is estimated to be insignificant in our analysis.

Although, as mentioned above, we show that these results should be interpreted as responses to transitory fluctuations in GDP, we also consider the responses to unanticipated movements in GDP that are explicitly identified using either monetary policy shocks or GDP forecast errors from a VAR. We conduct these exercises both as a validation of our headline results and because responses to monetary policy shocks are of particular interest given the role of monetary policy in stabilising the economy. We find that the responses to these identified shocks are broadly similar to our headline results described above.

In further analysis, we find that when estimating our regressions at the household instead of the individual level, the earnings responses are flatter across the distribution with the exception of the first income decile. We also examine whether controlling for negative GDP growth periods, such as the Great Recession, changes our results, but again find that our broad conclusions are unchanged. Due to our focus on earnings changes, most of our analysis is restricted to individuals that are initially employed. When we analyse employment transition rates for those that are initially not employed, we find a greater sensitivity to aggregate fluctuations for those that are initially unemployed than for those that are initially employed or inactive in the labour market. We also analyse the variance of pay growth and estimate a negative relationship between changes to GDP and pay growth variance across the distribution. This follows from the correlation between GDP fluctuations and extensive margin labour market transitions, as increases in GDP reduce the probability of movements out of employment.

Our results will be of interest to fiscal and monetary policy makers in helping inform them of the potential distributional and welfare implications of any policy changes that may impact the aggregate business cycle.[1] The facts presented in this paper will also be noteworthy for economists seeking to calibrate macroeconomic models that incorporate meaningful heterogeneity at the business cycle frequency to analyse the implications of fiscal and monetary policy e.g. models in the developing HANK literature following from McKay & Reis (2016).[2]

1.1 Related Literature

This paper fits into a growing literature that seeks to understand the distributional implications for individuals and households of fluctuations in the aggregate business cycle. Guvenen et al. (2017) use a large administrative dataset for the US to document a U-shaped response to fluctuations in GDP across the income distribution, with the incomes of the poorest and richest individuals found to be most exposed to the US business cycle. Other recent studies such as Amberg et al. (2022), Andersen et al. (2022) and Holm et al. (2021) have focused on the distributional implications of monetary policy shocks using administrative data from Sweden, Denmark and Norway respectively.[3] The conclusions from these papers vary, but a common thread is that there is excess sensitivity to business cycle fluctuations induced by monetary policy at the bottom and top of the distribution, with the response of labour income most important at the bottom of the distribution, and changes to capital income more important at the top.

This literature has mostly focused on the response of incomes, but some papers such as Broer et al. (2022) and Hoffmann & Malacrino (2019) also consider extensive margin adjustment. Hoffmann & Malacrino (2019) analyse administrative data from Italy and conclude that employment changes and spells of unemployment contribute to the pro-cyclical skewness of income. As in this paper, Broer et al. (2022) find that the extensive margin is key to explaining the excess sensitivity of income to the business cycle at the bottom of the income distribution. We contribute to this aspect of the literature by focusing on adjustment along both the intensive and extensive margins, and a more granular decomposition of extensive margin adjustment into the transitions into inactivity and unemployment, and also to job mobility.

Our paper also contributes to a small UK literature on the incidence of aggregate shocks. Bell et al. (2022) use the UK's Annual Survey of Hours and Earnings (ASHE) - a sample of employees - to study earnings dynamics and inequality over a long sample (1975-2020). They

[1] Particularly UK policy makers, advisers and economic modellers.

[2] Including the descriptive statistics.

[3] While the findings of these papers can potentially be extrapolated to larger developed economies, contributions such as in this paper are important in establishing the evidence base for this.

find that the variance of earnings has increased over time and that earnings exhibit pro-cyclical skewness. They also document the exposure of individuals with different characteristics to aggregate shocks, finding that it is declining in age, the size of their employer, skill level and permanent earnings. Relative to their paper, by using the LFS we are able to observe individuals who are not employed, including by studying the extent to which job-loss contributes to the response of labour earnings to aggregate shocks. Our sample also includes employees that are not captured by ASHE, either because they work for a very small firm or because their pay is below the threshold that requires their employer to enter them into the PAYE tax system. Cantore et al. (2023) use US and UK data to study the effect of monetary policy on hours worked and unemployment across the income distribution, including in a pseudo-panel constructed from the LFS. These authors find an initial counter-cyclical response of hours worked at the very bottom of the income distribution conditional on a monetary policy shock, though this effect does not persist to the peak transmission period in the case of the UK. Compared to this study, we leverage the panel element of the LFS to trace individual outcomes and focus on the broader business cycle as well as monetary policy induced fluctuations.

In addition, we contribute to a wider UK literature on cyclical earnings dynamics and the role of labour market transitions in generating employment fluctuations. Schaefer & Singleton (2019) use the ASHE data to measure the job-level response of wages and hours worked to business cycle fluctuations. They find that firms reduced the real hourly wages of both new hires and job stayers within jobs following the financial crisis, but that only the hours of new hires were affected. Relative to this study, we consider the response of pay to aggregate fluctuations across the income distribution, and are able to examine the role of job-loss. Finally, Elsby et al. (2011), Gomes (2012), Razzu & Singleton (2016) and Singleton (2018) examine the contribution of labour market transitions to cyclical movements in the UK unemployment and inactivity rates. In common with this literature, we find that cyclical variation in the probability of an unemployment transition is the predominant driver of employment fluctuations.

The rest of this paper is structured as follows. Section 2 describes the LFS data in more detail and presents key statistics on income, hours and labour market transition rates across the income distribution. Section 3 describes our empirical approach. Section 4 reports our main results and Section 5 concludes.

2 Data and Descriptive Statistics

2.1 Data

This study focuses on data from the UK's Labour Force Survey (LFS) which allows us to track both changes in labour earnings and the labour market status of individuals. The LFS is the largest household survey in the UK and is used to construct headline labour market statistics, such as the unemployment rate, labour force participation rate and hours worked. The survey has been run on a quarterly basis from the spring of 1992 and is designed to achieve a sample of $36,000$ households in each quarter. Individuals from sampled households stay in the survey for five quarters (waves) to enable the analysis of labour market transitions i.e. movements between employment, unemployment and inactivity. Since the spring of 1997, participants have been asked to report their pre-tax labour earnings in both wave 1 and wave 5, enabling analysis of changes in labour income over a 1 year period for employed individuals. Participants are asked to report their pay for the main and, if applicable, 2nd job in which they were employed during a particular reference week. Alongside their current labour market status and recent employment history, they also report the number of hours they work per week in their main job. The LFS measure of labour earnings, while not the official statistic (AWE), tracks the official measure very closely (see Figure A.1.1) and is the preferred source for the pay of the

Table 1: Average Characteristics by Income Decile

(a) Income

Period	1	2	3	4	5	6	7	8	9	10
(1996,2003]	44	101	157	204	246	292	346	416	517	856
(2003,2007]	61	133	200	253	303	357	426	511	637	1048
(2007,2011]	69	146	220	279	335	400	476	573	719	1195
(2011,2015]	75	158	234	296	356	426	506	612	768	1273
(2015,2019]	89	185	267	332	394	465	549	662	828	1378

(b) Other Characteristics

Period	1	2	3	4	5	6	7	8	9	10
Avg. Hours	14.4	24.0	32.9	37.0	38.6	39.6	40.4	40.6	40.7	41.9
Part Time (PT)	0.92	0.78	0.36	0.17	0.11	0.08	0.06	0.04	0.03	0.02
PT Student	0.29	0.08	0.02	0.01	0.00	0.00	0.00	0.00	0.00	0.00
Female Shr.	0.73	0.76	0.66	0.55	0.48	0.43	0.38	0.35	0.32	0.22
Avg. Age	36.0	39.6	38.0	38.0	38.9	39.5	39.7	40.6	41.9	43.2
Dependents	0.75	0.77	0.58	0.50	0.50	0.51	0.54	0.59	0.66	0.78
Age>60 Shr.	0.13	0.10	0.07	0.06	0.06	0.05	0.05	0.04	0.04	0.04
High Skill Shr.	0.05	0.08	0.09	0.11	0.16	0.23	0.31	0.43	0.59	0.77

Note: Panel (a) reports nominal pre-tax average weekly earnings by income decile for the specified periods. Panel (b) reports the average of other characteristics by income decile across all periods. Dependents is the average numbers of dependents under 16 within the individuals household. High Skill is based on the ONS classification of occupations (SOC) into four skill levels: we report the share in the highest skill group (level 4).

low paid and part time workers (ONS 2015).

We include in our sample all individuals aged 16 and over and focus on the period starting in 1997 (when we can start tracking changes in labour income) and end in 2019 (prior to the pandemic). More detail on the construction of our data is provided in Appendix A.1.1.

2.1.1 Limitations

LFS respondents who are unemployed or inactive in the labour market do not report any income or receive an income weight. Therefore, the majority of our analysis is restricted to those employed in wave 1 who we then track through to wave 5. Furthermore, self-employed individuals are not asked the income questions in the LFS and so our analysis also abstracts from this section of the labour force. Since the onset of the COVID pandemic in 2020/2021, the LFS has struggled with response rates and uncertainty over the true underlying population has created difficulty in the construction of accurate survey weights. This underscores our decision to end our sample in 2019Q4. Finally, while the LFS is a large nationally representative survey, the sample is smaller than the administrative datasets used in similar analyses in other countries. For this reason we are somewhat constrained in how finely we are able to cut the data and generally conduct our analysis at the income decile level.

2.2 Descriptive Statistics

Our analysis focuses on the response of labour earnings, hours worked and labour market transitions to fluctuations in economic activity across the income distribution. Like Broer et al. (2022), we first group households by initial income decile and then take averages over each decile in current and future periods (quarters):

$$y_{g,t+h} = \frac{1}{\sum_i \mathbf{1}_{(i,t)\in g} w_{i,t}} \sum_i \mathbf{1}_{(i,t)\in g} y_{i,t+h} w_{i,t} \tag{1}$$

where i is an individual in the LFS in quarter $t+h$, $h \in \{0,4\}$, g is an income group (defined at time t), $w_{i,t}$ is the LFS income weight and y is a variable of interest, e.g. labour earnings or an employment indicator.[4] [5]

Average weekly earnings, hours worked and other characteristics are reported by income decile in Table 1. The LFS captures the inequality in UK labour earnings reasonably well, with the top decile reporting incomes 15 times the level of the bottom decile, and the 9th decile reporting incomes at 4.5 times the level of the 2nd decile. However, due to income censoring and a lack of coverage at the very top of the distribution, these ratios are lower than those reported by the UK's tax authority (HMRC), which are closer to 25 and 5.5, respectively.[6] In terms of hours worked, we see that hours are increasing in weekly income, though the differences in hours worked are small in the top half of the distribution, where there are few part time workers. The bottom two deciles are notable for the share of part time workers, older workers and the female share relative to the other income buckets. And within those two deciles, the share of students stands out in the bottom income bucket. Individuals at the bottom and top of the income distribution also tend to live in households with more dependents under the age of 16.

This study is also focused on individual labour market transitions. Figure 1 panel (a) reports average transition rates between quarter t and $t+4$ by income decile in quarter t.[7] We see that average labour market transition rates are close to monotonically decreasing in income. The probability of transitioning out of employment (blue line) is highest at the bottom of the income distribution, and almost four times higher than at the top of the distribution. On average, most (around $\frac{3}{4}$) of these transitions out of employment are into inactivity (orange line), with the remainder explained by transitions into unemployment (cyan line). Those at the bottom of the income distribution are also more likely to switch jobs (purple line) than those at the top.

The propensity for more frequent labour market transitions at the bottom of the distribution is further reflected in higher real pay growth variance for these groups. Panel (b) of Figure 1 plots average pay growth variance by initial income decile which takes the form of an incomplete U shape, with variance highest at the bottom and lowest in the upper-middle deciles 7 and 8. By comparing to the orange dashed line we can see that the differences in pay growth variance are nearly completely explained by extensive margin transitions, with the variance for those that remain employed in both periods nearly flat across the income distribution.

Finally, while average transition rates vary across the distribution it's notable that even in the middle and upper deciles, transitions between jobs or into non-employment are quite common. This underscores the value of being able to follow individuals in a panel as opposed to conducting analysis on a pseudo-panel, where this margin is either omitted or imputed. In terms of movements along the income distribution between waves, in further analysis (see Appendix A.1.3) we find that the large majority of individuals either remain in the same income decile or move into adjacent deciles.

[4] In practice, our results are insensitive to the use of income weights as demonstrated in Figure A.1.4. This is due to the fact that these weights are highly correlated with income decile which we condition on in our analysis.

[5] In Broer et al. (2022) households are grouped by permanent income decile based on income history and observables. In the LFS we only observe incomes twice in wave 1 and 4 quarters later in wave 5, and so bin based on regular income reported in wave 1.

[6] HMRC Survey of Personal Incomes.

[7] Note that unlike our main results, these averages are not conditional on changes in GDP and so represent a mix of steady-state and business cycle induced transitions.

Figure 1: Unconditional Moments

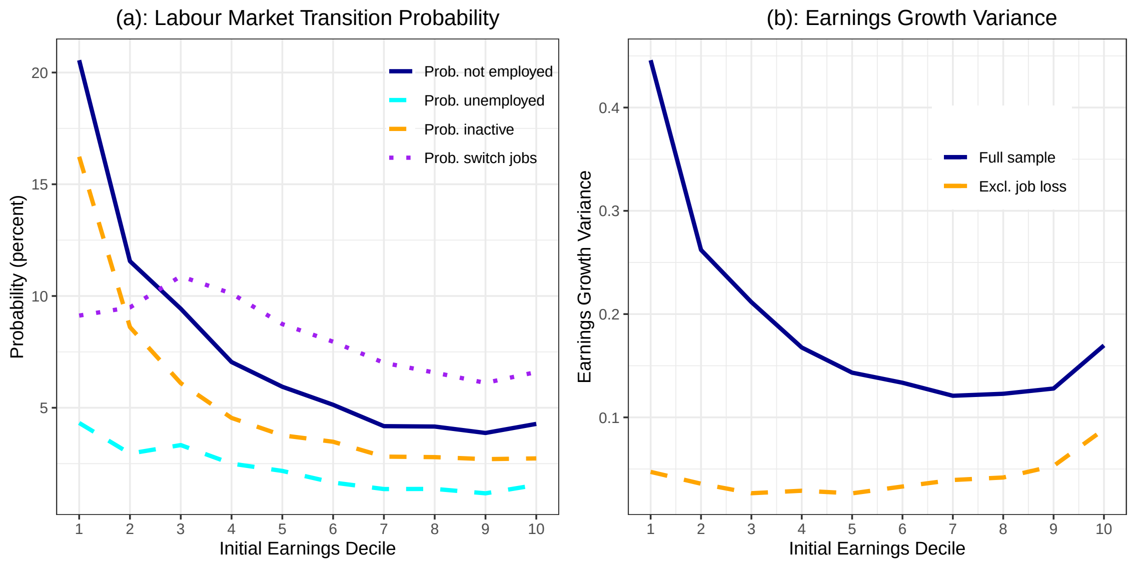

Note: Panel (a) shows average labour market transition probabilities between t and $t+4$. The dark blue line shows the probability of leaving employment; the cyan line shows the probability of becoming unemployed; the orange line shows the probability of becoming inactive; and the purple line shows the probability of changing employer. Panel (b) plots the variance of earnings growth between period t and $t+4$ across our sample period. Individuals are sorted into deciles in each quarter based on their earnings in t. The sample period is 1997Q2-2019Q4.

3 Empirical Framework

Our empirical approach follows other similar work, such as Guvenen et al. (2017) or Broer et al. (2022), by running simple regressions of changes in different outcome variables on changes in aggregate economic conditions (GDP) by income decile. We refer to the elasticity, β_g, resulting from these regression as a GDP beta. The outcomes we focus on are: real labour income growth (log difference); hours growth (log difference); the linear probability of a transition to unemployment or inactivity; and the linear probability of changing jobs.[8] Our main GDP beta regressions are specified as follows:

$$y_{g,t+4} - y_{g,t} = \alpha_g + \beta_g \Delta GDP_{t+4} + \epsilon_{g,t} \tag{2}$$

where β_g is an unconditional GDP elasticity for GDP growth over the same four quarter period. This elasticity can be interpreted as a variance weighted average of elasticities with respect to a slow moving growth trend and the elasticity to transitory business-cycle frequency fluctuations.

To see this, consider the case in which the data generating process for individual wage growth is as follows:

$$dy_{i,t} = dL_{i,t} + dz_{i,t} + \gamma_1 dg_t + \gamma_2 da_t \tag{3}$$

where the change in GDP is $dgdp_t = dg_t + da_t$ is composed of a slow moving growth trend dg and faster moving business cycle component da. Individual pay growth is also determined by returns to age and experience, dL, and by idiosyncratic shocks dz. If we were to then estimate the regression:

$$dy_{i,t} = \alpha + \beta dgdp_t + \epsilon_{i,t} \tag{4}$$

[8] We use real GDP growth and real income growth in our regressions, where the latter is nominal earnings deflated using the GDP deflator. More detail on the construction of our data is provided in Appendix A.1.1.

we would obtain the following coefficient estimates for α and β:

$$\alpha = E[dL] + E[dz] + \frac{E[dgdp]}{Var(dgdp)}(\gamma_1 - \gamma_2)E[da^2] \tag{5}$$

$$\beta = \frac{Var(dg)}{Var(dgdp)}\gamma_1 + \frac{Var(da)}{Var(dgdp)}\gamma_2 \tag{6}$$

where we have assumed $E[da] = 0$ and $E[dadg] = 0$, i.e. transitory aggregate shocks average zero and are uncorrelated with long-run trend growth. From Equation (6), we can see the unconditional GDP beta (β) is a variance share weighted average of the elasticities with respect to the trend and the business cycle growth components of GDP. If most of the variance of UK GDP growth is determined by shorter run business cycle movements, as suggested by Melolinna & Tóth (2016) when estimating a GDP filter under similar assumptions, then this unconditional elasticity will largely reflect the elasticity with respect to those transitory business cycle frequency movements in GDP. From Equation (5), we can also see that the constant α_g largely picks up the income group specific trends, such as mean reversion of individual idiosyncratic transitory shocks and growth related to the life-cycle earnings profile.

Based on the insights above and in addition to our main approach, we also estimate GDP betas after isolating more specific transitory business cycle frequency movements (da) in GDP. We do so because it may be the case that the responses to transitory shocks are sufficiently different to unexpected changes in trend growth i.e. $\gamma_1 \neq \gamma_2$. The response to explicitly identified shocks such as monetary policy shocks is also of interest in and of itself and a focus of the related literature (e.g. Amberg et al. (2022)). To do so, we estimate Equation (2) via two stage least squares, instrumenting for changes in GDP with either accumulated monetary policy shocks prior to time t or with the unexpected component of GDP growth between period t and $t+4$.[9] Specifically, we estimate the effects of 1 per cent changes to GDP associated with either surprise movements in past interest rates or with changes to GDP that would have been difficult to predict given the information set in the quarter preceding wave 1. Both methods draw upon a six variable VAR outlined in Appendix A.2, with the first stage explaining a significant share of movements in GDP as indicated in Table A.2.1.[10]

4 Results

4.1 GDP Betas

Our main estimates of β_g from Equation (2) for real earnings and hours are displayed in panels (a) and (b) of Figure 2. The impact of GDP movements on real earnings (panel (a)) is economically and statistically significant across the income distribution and on average a movement in GDP of 1 per cent is associated with a change in real labour earnings of around 0.7 per cent.[11] Despite some volatility at the bottom, the overall picture is of a mildly downward sloping GDP beta estimate for earnings as we move across the distribution from the lowest to the highest income groups, with some reversal as we move past the lowest variance income groups (deciles 7 and 8) to the top decile. Panel (b) of Figure 2 focuses on the response of hours worked to a 1 per cent change in GDP. Here we see a much more notably downward sloping GDP beta,

[9] We use high frequency shocks identified in 30 minute windows around MPC instruments and focus on movements in the three month short sterling futures contract expiring 3-6 months after the announcement (see Cesa-Bianchi et al. (2020) for more details). To instrument for changes in GDP we use accumulated shocks from one and two years prior in line with the significance of the transmission dynamics captured in the VAR illustrated in Figure A.2.1.

[10] F statistic of 12.85 and 40.1 for the monetary policy shocks and forecast error approach respectively.

[11] Bell et al. (2022) estimate 0.38pp using a longer sample of UK earnings data from ASHE.

with an estimated GDP beta for hours similar to that of earnings in the bottom half of the distribution but significantly smaller in the top half of the distribution.

The different shape of the GDP beta profile for earnings and hours displayed in panels (a) and (b) of Figure 2 suggests that the margin of adjustment to fluctuations in aggregate conditions varies across the income distribution, with hours playing the dominant role in the bottom half of the distribution whereas changes in hourly earnings would seem to be more important at the top. In panels (c) and (d) of Figure 2, we therefore present some indicative decompositions of the GDP beta estimates by comparing our headline estimates to point estimates from different sub-samples. Panel (c) compares the point estimate from panel (a) in dark blue to the results from two sub-samples. The orange dashed line shows the point estimate only for those individuals that remain employed in period $t + 4$. The estimated GDP betas for those that remain employed are uniformly below the whole sample estimates, and by a large margin for some income deciles in the bottom half of the distribution. This points to an important role for the extensive margin in accounting for the sensitivity of earnings to GDP fluctuations across the income distribution, particularly at the bottom. The dashed cyan line in panel (c) further restricts the sample to those that stay in the same job. It is not materially different from the orange dashed line, suggesting that job switching is not a significant determinant of the overall GDP beta. In a final step, panel (d) focuses on those that remain employed in both periods (the orange line) and approximately decomposes the GDP beta point estimate into changes in hourly pay and changes in hours.[12] The purple dashed line plots the GDP beta for hours for those that remain employed and shows it to be downward sloping across the income distribution. This contrasts with hourly pay (green dashed line) which, though more volatile, is upward sloping on average; it explains a smaller share than hours at the bottom of the distribution and almost the entire adjustment from the middle to the top of the distribution.

The results in Figure 2 point to an important role for the extensive margin in explaining the sensitivity of incomes across the UK income distribution. In Figure 3, we therefore further investigate this margin of adjustment by using labour market transitions as the dependent variable in Equation (2). Panel (a) plots the GDP beta for the linear probability of moving from employment in period t to non-employment in period $t+4$. A 1 per cent increase (decrease) in GDP is associated with around a 0.2 percentage point reduction (increase) in the probability of moving to non-employment on average across the distribution. With the exception of the first and second deciles, the point estimate of the GDP beta for non-employment is upwards sloping towards the upper middle portion of the income distribution where it is not significantly different from zero.

Focussing on the bottom two deciles, it's perhaps surprising relative to the patterns across the rest of the distribution to see reversals in the point estimates between deciles at the bottom of the distribution, as we do in panel (a) of Figures 2 and 3. From Table 1, we know that the bottom two deciles are mostly part-time workers and when we remove those that are initially part-time employed from our sample, the results are smoother and more downward sloping without significant reversals across the distribution (see Figure A.3.4). However, one of the advantages of our approach using the LFS is that we capture these low paid and part-time employees where they are often missed in other analyses.[13] It's likely that workers in this part of the distribution may exhibit different and more heterogeneous labour supply characteristics. For example, aside from being part-time employed, we have shown that those in the first decile are more likely to be a student or above 60 whereas those in the second decile are part of

[12]Whilst close in practice, the product of the arithmetic weighted mean of the change in hours and change in hourly pay does not equal the arithmetic weighted mean of the change in pay (hours x hourly pay).

[13]By deliberately excluding part-time workers from the sample or by using data that fails to capture them for administrative reasons. For example, the ASHE dataset misses employees from non-VAT registered firms and employees below the national insurance earnings threshold.

Figure 2: GDP Betas for Earnings and Hours

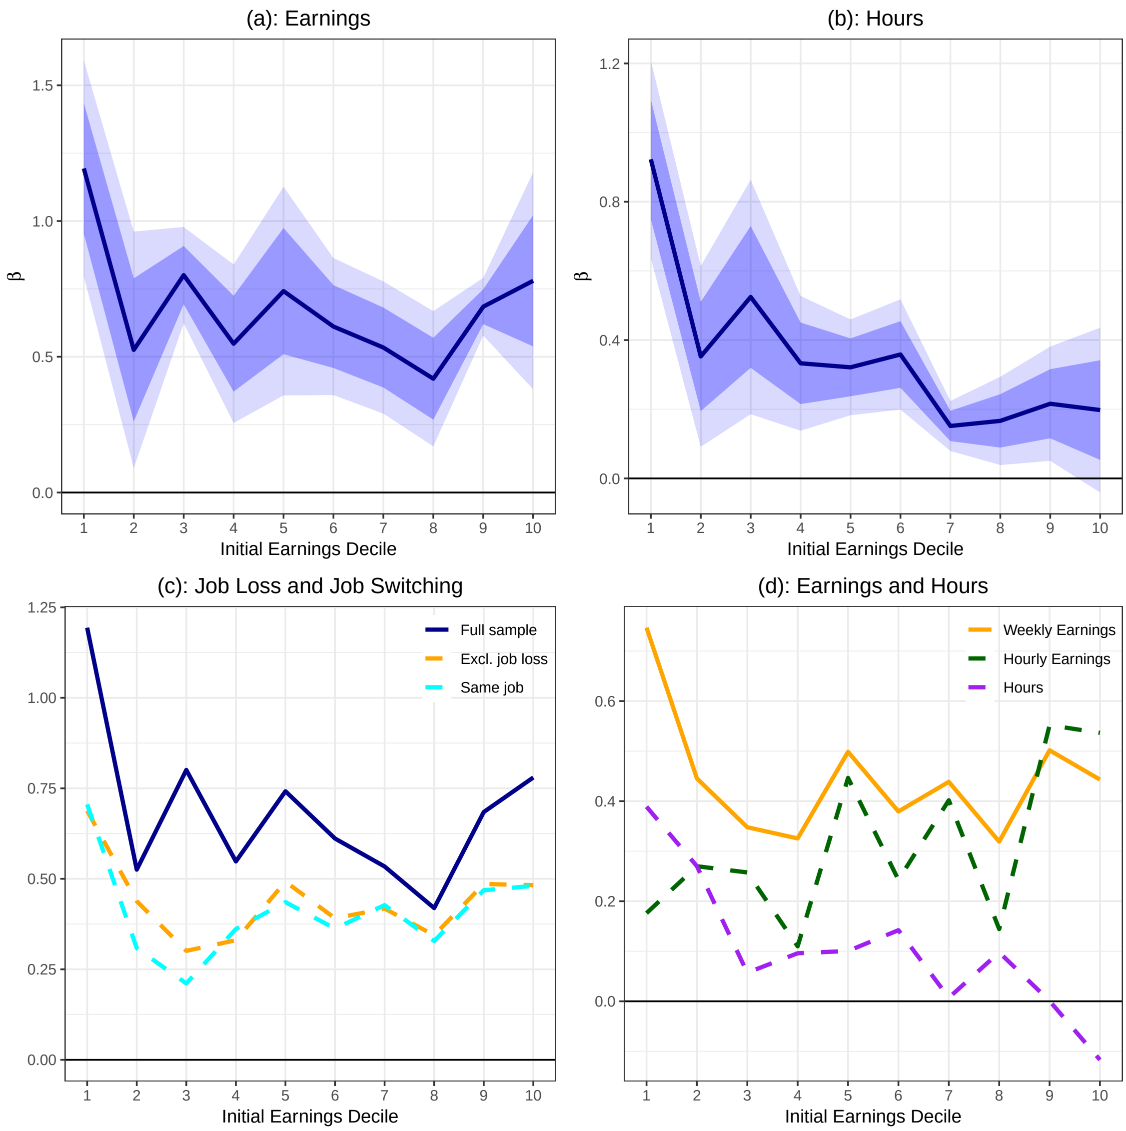

Note: Panel (a) plots the coefficients β_g from Equation (2) for changes in labour income and panel (b) for changes in hours worked in their main job. The darker (lighter) shaded area represents the 68% (90%) confidence intervals calculated using HAC standard errors. Panel (c) plots the coefficients β_g from Equation (2) for different samples: the blue line plots the coefficients for the full sample, as in panel (a); the orange dashed line plots the coefficients for those individuals that remain employed in $t+4$; and the cyan line plots the coefficients for those individuals that remain with the same employer in $t+4$. Panel (d) plots the coefficients β_g from Equation (2) with different dependent variables, for the sample of individuals that are employed in both periods: the orange line plots the coefficients using weekly pay as the dependent variable, as in the orange line in panel (c); the green line plots the coefficients using hourly pay as the dependent variable; the purple line plots the coefficients using weekly hours worked as the dependent variable. Individuals are sorted into deciles in each quarter based on their earnings in t. The sample period is 1997Q2-2019Q4 in each panel.

Figure 3: GDP Betas for Labour Market Transition Probabilities

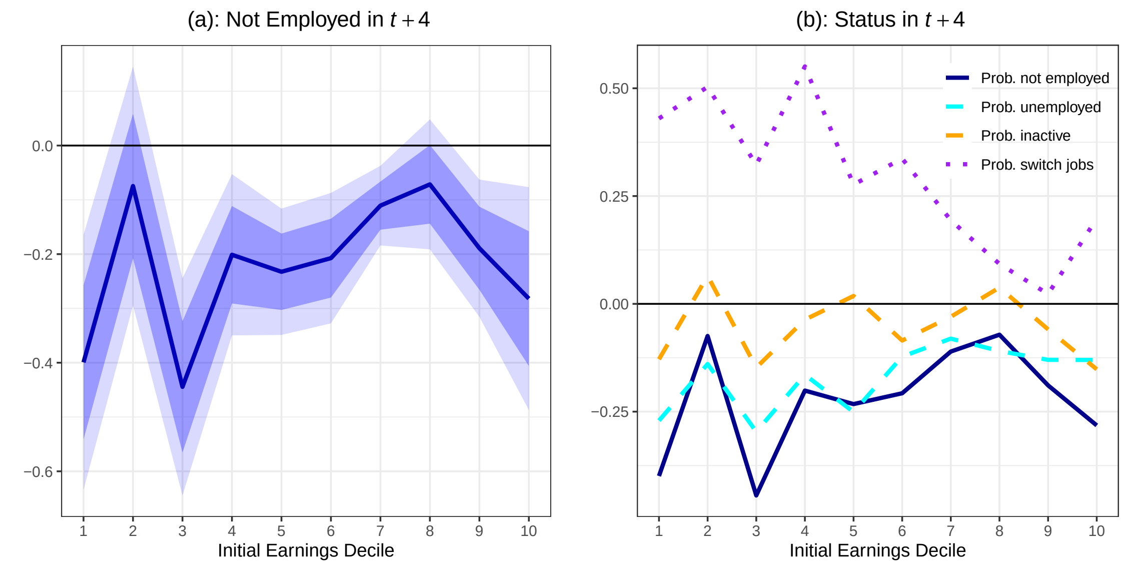

Note: Panel (a) plots the coefficients β_g from Equation (2), with the probability of non-employment in $t+4$ as the dependent variable. The darker (lighter) shaded area represents the 68% (90%) confidence intervals calculated using HAC standard errors. Panel (b) plots the coefficients β_g from Equation (2) with different dependent variables: the blue line plots the coefficients using the probability of non-employment as the dependent variable, as in panel (a); the cyan line plots the coefficients using the probability of unemployment as the dependent variable; the orange line plots the coefficients using the probability of inactivity as the dependent variable; and the purple line plots the coefficients using the probability of having a new employer as the dependent variable. Individuals are sorted into deciles in each quarter based on their earnings in t. The sample period is 1997Q2-2019Q4.

households with higher numbers of dependants but are otherwise more similar to the other deciles. It's possible that part-time students and older workers may, for example, be more demand determined, taking hours as they come. Whereas those in the second decile may exhibit more insensitive labour supply as welfare program incentives (e.g. low income support) and caring responsibilities may dictate their labour supply to a greater degree than other parts of the distribution (Beffy et al. 2019).

In panel (b) of Figure 3, we decompose the probability of moving out of employment by showing the point estimates for the linear probability of becoming unemployed (cyan line) and the probability of becoming inactive in the labour market. We also plot the probability of switching jobs (purple line). The close proximity of the unemployment line to the overall blue non-employment line suggests that it is the unemployment margin that drives the response of extensive margins transitions to GDP fluctuations. Movements to inactivity are statistically insignificant. This latter result contrasts with the steady-state or average transition probabilities shown in Figure 1, where the majority of movements into non-employment are explained by movements into inactivity. But it is in line with previous findings that cyclical variation in the probability of an unemployment transition is the predominant driver of employment fluctuations in the UK (Elsby et al. 2011, Gomes 2012, Razzu & Singleton 2016, Singleton 2018). This underscores the importance of conditioning on GDP, as otherwise a policymaker could erroneously infer the margin through which their policies are likely to transmit. Finally, in Figure 3 we note that the GDP betas for job switching are downward sloping across the income distribution, more so than for the unconditional transition rate (Figure 1).

4.2 Transitory Shocks

As discussed in Section 3, the GDP beta estimates reported in the previous sub-section reflect a variance-weighted average of trend growth (dg) and more transitory business cycle movements in GDP around trend (da). In this sub-section, we attempt to isolate those more transitory, business cycle frequency fluctuations by using GDP forecast errors and monetary policy shocks as first stage instruments for GDP growth, as outlined in Table A.2.1. The key results are shown in Figure 4. The first row shows the results for the forecast errors approach. The pattern and magnitude of the estimates are quite similar to that of our main results. We again find significant effects across the income distribution and the GDP beta for earnings (panel (a)) is larger in the bottom half of the distribution and lowest in the upper middle. It does, however, tick up more noticeably at the top of the distribution. The extensive margin again plays an important role in determining this pattern (panel (b)), with the point estimates similar to our main result (Figure 3).

The bottom row of Figure 4 shows the results for monetary policy shocks, i.e. movements in earnings and employment probabilities conditional on a 1 per cent movement in GDP that is due to past interest rate changes. Earnings are again significantly impacted across the income distribution. While still elevated at the bottom of the income distribution, the point estimates of the GDP beta for earnings in panel (c) are noticeably flatter than for the other estimates, though the confidence intervals are larger. The pattern for the linear probability of non-employment (panel (d)) is more similar to our other results, with heightened sensitivity in the bottom half of the distribution. Taken together with additional results reported in Appendix A.2.1, the overall conclusions that we draw from the instrumental variable approach are similar to those from the unconditional approach, as hypothesised in Section 3.

4.3 Other Exercises

We have also conducted a number of additional exercises which we summarise in this section and illustrate in Appendix A.3. As a household survey, the LFS interviews each member of sampled households which enables us to aggregate our dataset to the household level. It is possible that GDP beta for earnings could look different when grouping and sorting households instead of individuals.[14] For example, an individual's labour supply decision may reflect the labour market status of other members of their household (Lundberg 1985). For most income deciles, we find that the GDP beta estimates are not significantly different from the individual-level results reported in Section 4.1. One difference is a greater sensitivity in the first decile. Taken together, these results do not suggest that within-household insurance cushions the response of incomes to GDP fluctuations.

While we do not observe an individual's income when not employed in wave 1 (period t), we do observe their labour market status and subsequent status 4 quarters later. We can therefore estimate the GDP betas for the probability of being not employed, unemployed and inactive for those initially not in employment. When doing so, we again find that in response to GDP fluctuations it is the unemployment to unemployment margin that dominates. Furthermore, we find the GDP beta for the probability of non-employment is larger (more negative or sensitive) for the initially unemployed than the initially employed (see Figure 3).

We also considered whether controlling for (discounting) periods of negative growth – and in particular the Great Recession – affects our GDP beta estimates. We do this by including indicators of negative growth between periods t and $t + 4$ as an interaction term with GDP growth in Equation (2). The estimates produced by this specification are similar to our main results, particularly for the extensive margin transitions. We do find some statistically signifi-

[14]The majority of the related literature has conducted its analysis at the individual level.

Figure 4: GDP Betas for Transitory Shocks

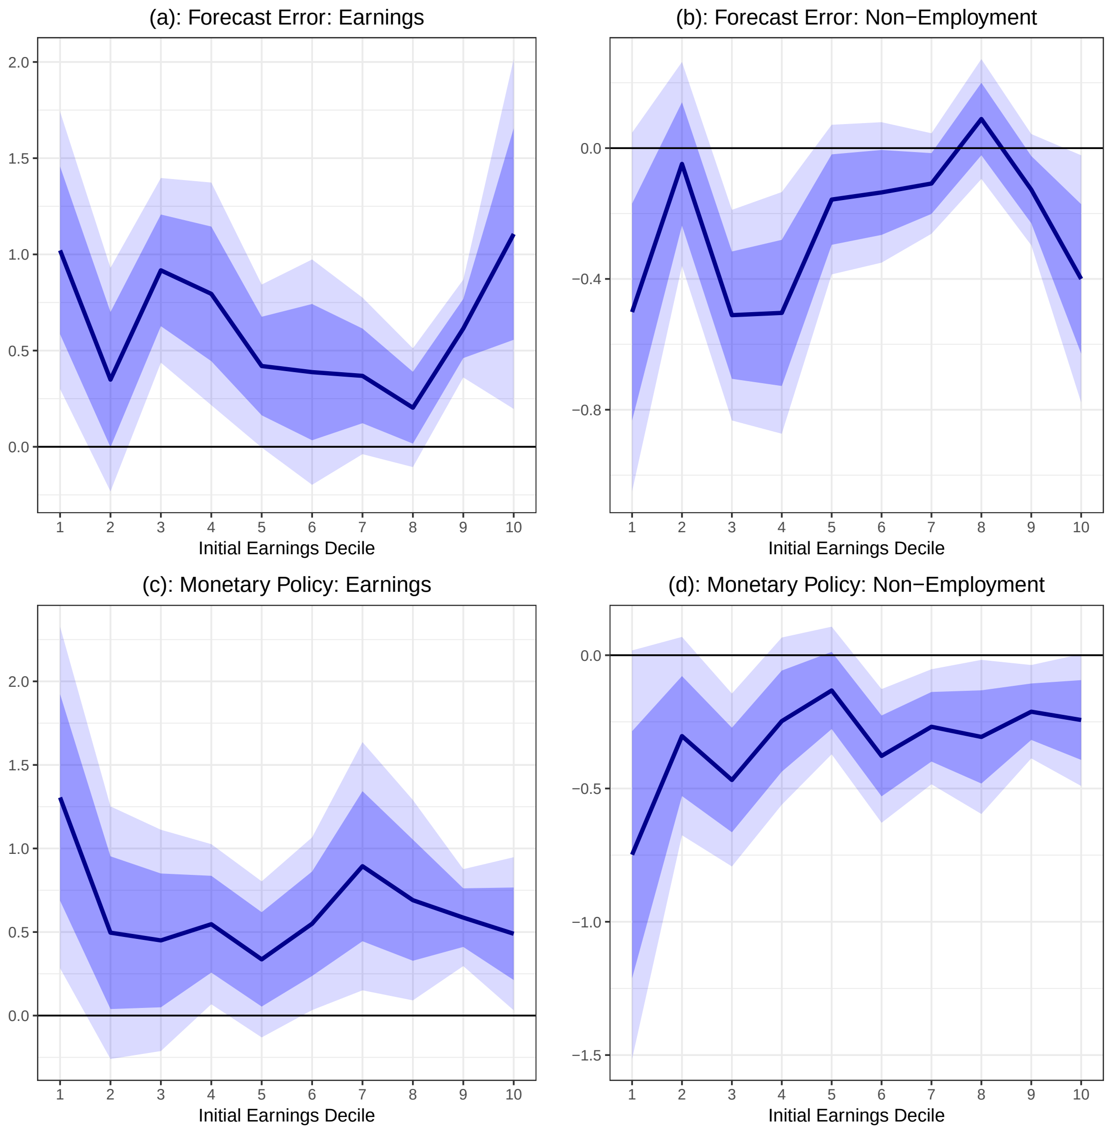

Note: Panel (a) and panel (b) plot the coefficients β_g from Equation (2) under an instrumental variable approach using GDP forecast errors from the VAR discussed in Appendix A.2 as an instrument for GDP growth. Panel (c) and (d) follow the same procedure but use accumulated high frequency monetary policy shocks as the instrument. The darker (lighter) shaded area represents the 68% (90%) confidence intervals calculated using HAC standard errors. Individuals are sorted into deciles in each quarter based on their earnings in t. The sample period is 1997Q2-2019Q4.

cant differences in the estimated GDP betas for earnings for some deciles, with estimated GDP betas higher on average in response to positive GDP movements, consistent with findings of downward real wage rigidity (Grigsby et al. 2021).

The change in the variance of income growth across the distribution by income decile is a further object of interest and a theoretically important determinant of household consumption. We find that the GDP betas for earnings variance are negative across the distribution, and slightly more so in the bottom half, though this relationship is flatter than for actual earnings (the mean). The fact that the GDP betas for pay growth variance are negative can be explained by the negative correlation between extensive margin labour market transitions and GDP movements. A positive economic shock reduces the probability of becoming unemployed (Figure 3), which reduces variance in pay growth for the initially employed. This is consistent with findings in the literature. For example, Storesletten et al. (2001) estimate that the variance of persistent shocks to disposable household income more than doubles in US recessions.

5 Conclusion

Using data from the UK's Labour Force Survey, this paper provides empirical evidence on the impact of GDP fluctuations on real earnings and labour market outcomes across the income distribution. We find that aggregate fluctuations have statistically and economically significant heterogeneous impacts across the income distribution. In particular, we find the greatest sensitivity at the very bottom of the income distribution and the least sensitivity in the upper middle portion of the income distribution. When decomposing changes in real earnings, we find that changes in earnings in the bottom half of the distribution are explained more through changes to hours worked and transitions into unemployment. In contrast, adjustments in the top half of the distribution are better explained by changes to hourly wages. These findings also hold on average across identification strategies e.g. conditioning specifically on monetary policy shocks.

Future research could build on this work by rationalising these GDP betas and steady-state labour market transition rates in a heterogeneous agent economic model of UK business cycles.

References

Amberg, N., Jansson, T., Klein, M. & Picco, A. R. (2022), 'Five facts about the distributional income effects of monetary policy shocks', *American Economic Review: Insights* **4**(3), 289–304.

Andersen, A. L., Johannesen, N., Jørgensen, M. & Peydró, J.-L. (2022), 'Monetary policy and inequality', *Univ. of Copenhagen Dept. of Economics Discussion Paper, CEBI Working Paper* **9**, 22.

Bauer, M. D. & Swanson, E. T. (2023), 'A reassessment of monetary policy surprises and high-frequency identification', *NBER Macroeconomics Annual* **37**(1), 87–155.

Beffy, M., Blundell, R., Bozio, A., Laroque, G. & To, M. (2019), 'Labour supply and taxation with restricted choices', *Journal of Econometrics* **211**(1), 16–46.

Bell, B., Bloom, N. & Blundell, J. (2022), 'Income dynamics in the United Kingdom and the impact of the Covid-19 recession', *Quantitative Economics* **13**(4), 1849–1878.

Broer, T., Kramer, J. V. & Mitman, K. (2022), The curious incidence of monetary policy across the income distribution, Discussion paper, Centre for Economic Policy Research.

Cantore, C., Ferroni, F., Mumtaz, H. & Theophilopoulou, A. (2023), A tail of labor supply and a tale of monetary policy, Discussion paper, Centre for Macroeconomics.

Cesa-Bianchi, A., Thwaites, G. & Vicondoa, A. (2020), 'Monetary policy transmission in the United Kingdom: A high frequency identification approach', *European Economic Review* **123**, 103375.

Elsby, M. W., Smith, J. C. & Wadsworth, J. (2011), 'The role of worker flows in the dynamics and distribution of UK unemployment', *Oxford Review of Economic Policy* **27**(2), 338–363.

Gomes, P. (2012), 'Labour market flows: Facts from the United Kingdom', *Labour Economics* **19**(2), 165–175.

Grigsby, J., Hurst, E. & Yildirmaz, A. (2021), 'Aggregate nominal wage adjustments: New evidence from administrative payroll data', *American Economic Review* **111**(2), 428–471.

Guvenen, F., Schulhofer-Wohl, S., Song, J. & Yogo, M. (2017), 'Worker betas: Five facts about systematic earnings risk', *American Economic Review* **107**(5), 398–403.

Hoffmann, E. B. & Malacrino, D. (2019), 'Employment time and the cyclicality of earnings growth', *Journal of Public Economics* **169**, 160–171.

Holm, M. B., Paul, P. & Tischbirek, A. (2021), 'The transmission of monetary policy under the microscope', *Journal of Political Economy* **129**(10), 2861–2904.

Jentsch, C. & Lunsford, K. G. (2022), 'Asymptotically valid bootstrap inference for proxy SVARs', *Journal of Business & Economic Statistics* **40**(4), 1876–1891.

Lundberg, S. (1985), 'The added worker effect', *Journal of Labor Economics* **3**(1, Part 1), 11–37.

McKay, A. & Reis, R. (2016), 'The role of automatic stabilizers in the US business cycle', *Econometrica* **84**(1), 141–194.

Melolinna, M. & Tóth, M. (2016), 'Output gaps, inflation and financial cycles in the United Kingdom', *Bank of England, Staff Working Paper 585* .

ONS (2015), 'Labour Force Survey (LFS) QMI'.

Razzu, G. & Singleton, C. (2016), 'Gender and the business cycle: An analysis of labour markets in the US and UK', *Journal of Macroeconomics* **47**, 131–146.

Schaefer, D. & Singleton, C. (2019), 'Cyclical labor costs within jobs', *European Economic Review* **120**, 103317.

Singleton, C. (2018), 'Long-term unemployment and the great recession: Evidence from UK stocks and flows', *Scottish Journal of Political Economy* **65**(2), 105–126.

Stock, J. H. & Watson, M. W. (2018), 'Identification and estimation of dynamic causal effects in macroeconomics using external instruments', *The Economic Journal* **128**(610), 917–948.

Storesletten, K., Telmer, C. I. & Yaron, A. (2001), 'The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk', *European Economic Review* **45**(7), 1311–1339.

A Appendix

A.1 Data

In this section, we provide more detail on the LFS variables and other data that we use in our analysis.

A.1.1 LFS data

As discussed in Section 3, we construct the following outcome variables using the LFS data:

- **Sample**

  We include in our sample those aged 16 and over from LFS cross sectional data waves 1997Q2 to 2019Q4.

- **Labour income**

  The main outcome that we consider is the annual change in real weekly labour income: nominal weekly labour income deflated using the GDP deflator (see next section). We use the sum of gross weekly pay in the main job (*GRSSWK*) and, if applicable, second job (*GRSSWK2*).

  In the decomposition that we report, we also consider the change in real hourly pay in an individual's main job. This is constructed by dividing nominal weekly pay in an individual's main job (as described above) by the hours that the individual usually works in that job (described below).

  We group individuals reporting income in wave 1 into initial income deciles in each LFS wave. If individuals report that the amount that they were paid for their main job was different to usual, we use the usual pay that they report (*USUGPAY*) converted to a weekly amount using the period covered by that pay (*GRSPRD*).

  For those employed in wave 1, we remove those earning less than £37.5 in real terms, approximately half the real rate of job seekers allowance. We also windsorize calculated real income growth, removing the top 1 per cent as suspected outliers.

- **Hours worked**

  The LFS only records information on the usual hours that individuals work in their main job. Depending on an individual's circumstances, we construct this using information from the following variables: total usual hours worked excluding lunch breaks (*TOTUS1*); usual hours worked excluding overtime (*USUHR*); and usual hours of paid overtime (*POTHR*).

- **Labour market transitions**

  Labour market transitions are constructed using the labour force status (*ILODEFR*) reported by individuals in wave 1 and wave 5 of the survey.

- **Job switching**

  Individuals who have changed employer are identified using information on their tenure (*EMPLEN*).

- **Income Weights**

  Our analysis is focused on income changes and so we use the cross sectional data wave 1 income weights to weight individuals when aggregating the data as in equation 1. As

Figure A.1.1: LFS Microdata Compared to Official Series.

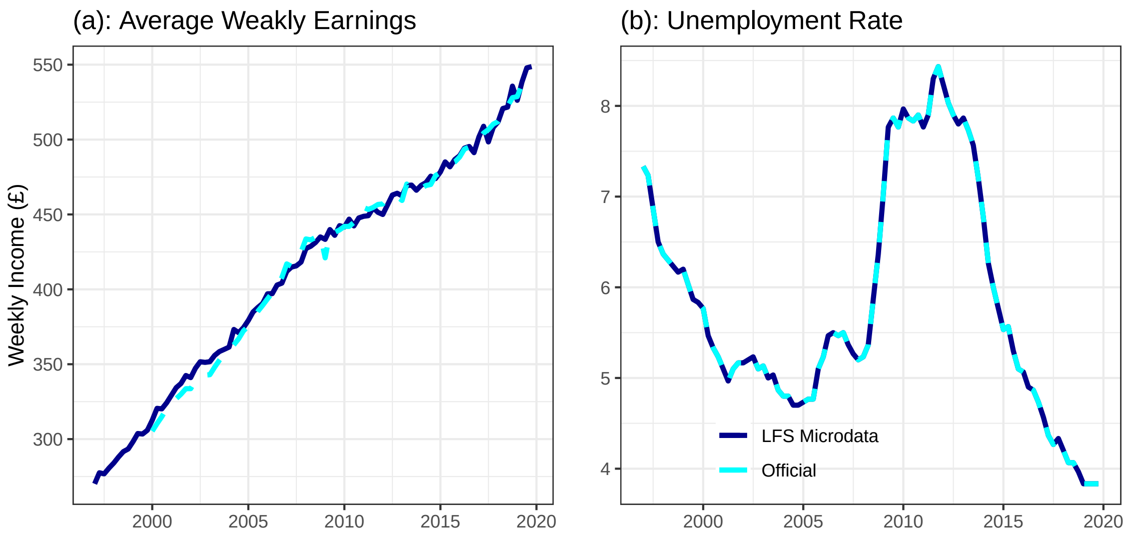

Note: Figure compares aggregate series constructed using the LFS microdata used in this study to the headline figures reported by ONS. The LFS data aggregates are calculated using the provided survey weights, and then seasonally adjusted using the X13-ARIMA procedure. The official AWE measure from ONS is available from January 2000.

some employed in wave 1 will become not employed in wave 5 we use the income weights from wave 1 in the wave 1 and wave 5 aggregations. Our results are not sensitive to the use of weights, however, as indicated in Figure A.1.4.

A.1.2 Other data

- **GDP**

  Latest ONS estimate of real GDP ($ABMI$), the annual log change of which is used as an explanatory variable in Equation (2).

- **GDP deflator**

  Latest ONS estimate of the GDP deflator ($YBGB$), used to delate labour income constructed using the LFS data.

A.2 Transitory shocks

The SVAR described below supports our analysis by demonstrating the effect of and timing of our chosen monetary policy shocks on GDP, and by providing a framework for forecasting GDP from which we can extract four quarter ahead forecast errors. The VAR is composed of six variables:

1. 1 year Rate: Monthly average of 12 month spot rate on UK government debt (Bank of England);

2. Exchange Rate: UK effective exchange rate index (Bank of England). A geometric weighted average of selected bilateral exchange rates;

3. FTSE: FTSE All share index (Refinitiv);

Figure A.1.2: Earnings Growth Distribution

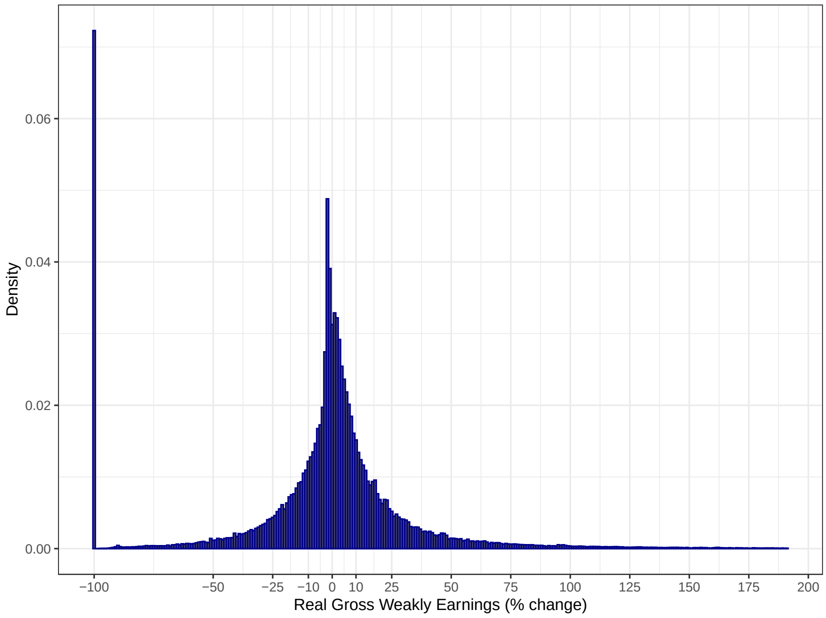

Note: Figure reports the distribution of growth in real gross weakly earnings for indivdiuals accross the entire sample 1997-2019.

Figure A.1.3: Earnings Decile Transition Matrix

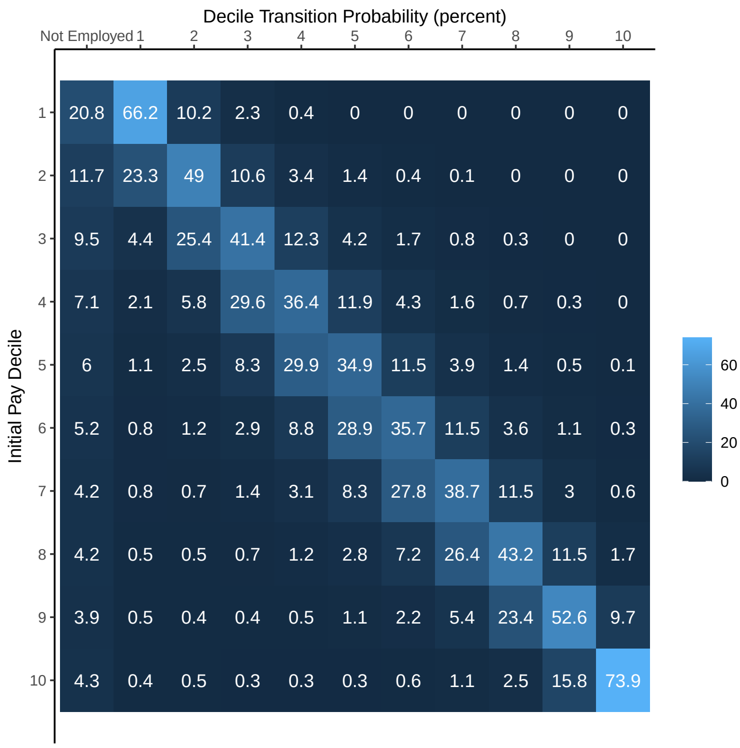

Note: Figure reports the average probability over our sample of transitioning to each pay decile or non-employment over 4 quarters conditional on an individual's pay decile in period t.

Figure A.1.4: Impact of Person Income Weights on GDP Betas

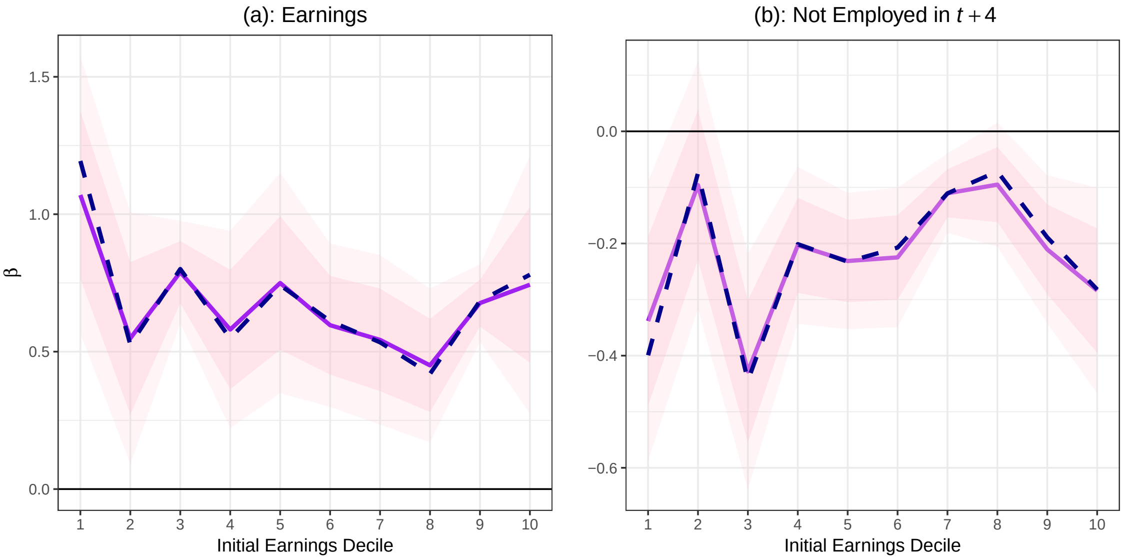

Note: Figure reports the GDP betas for earnings and labour market transition probabilities from Figure 2 panel (a) and Figure 3 panel (a) constructed without using ONS person income weights. The blue dashed lines are the corresponding results using the income weights, as reported in Section 4.

4. Spreads: Difference between ICE Bank Of America Sterling Corporate Index and 5 year UK government debt yield;

5. GDP: Monthly Gross Domestic Product (ONS);

6. Core CPI: Consumer price index excluding energy, food, alcohol and tobacco (ONS). Seasonally adjusted.

The VAR is estimated using monthly observations from 1996 to 2019 and 12 lags. The effect of monetary policy is identified using the external (proxy) instruments approach (Stock & Watson 2018). The instrument for monetary policy shocks is taken as the change in the price of sterling futures contracts in the 30 minutes around UK Monetary Policy Committee announcements. Specifically, we take the difference in the sterling futures contract that settles in the quarter following the announcement based on the 3-Month London Interbank Offered Rate (LIBOR). Following Bauer & Swanson (2023), we further clean our instrument $\hat{z}_t$ for information effects by orthogonalising with respect to information x_{t-} available in the 30 days priors to each announcement window:

$$\hat{z}_t = \beta x_{t-} + z_t \tag{A.2.1}$$

where x_{t-} includes the change in the FTSE all share in the 30 days prior to the annoucnment, the change in the effective exchange rate, the change in the 1 year government borrowing rate, change in corporate spreads and the change in GDP in the month prior to the announcement.

Figure A.2.1 shows the effect of our identified monetary policy shocks in the monthly SVAR. We can see the effects of the monetary policy shock on GDP take some time to build and only become significant at the 68 per cent confidence interval beyond 18 months and 90 per cent level beyond 24 months.

When looking to assess the impact of identified transitory shocks in our analysis we do so by estimating Equation (2) via two stage least squares. We instrument for four quarter changes in GDP in two ways. In one approach we sum monetary policy shocks from one and two years prior given the significance of the transmission dynamics captured in Figure A.2.1.

Figure A.2.1: Impulse Response to a 1 per cent Monetary Policy Innovation

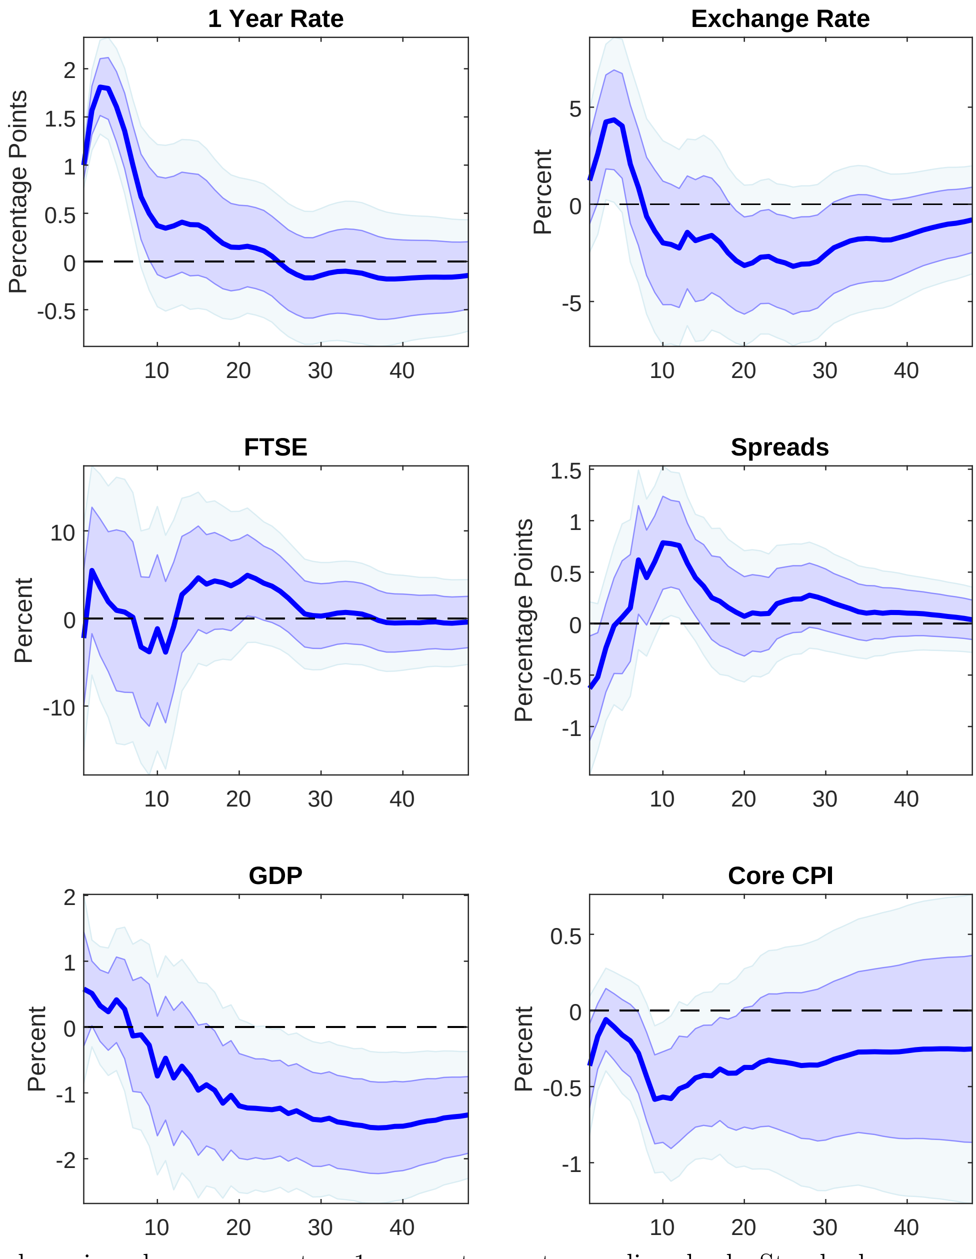

Note: Figure shows impulse responses to a 1 per cent monetary policy shock. Standard errors are derived using a moving block bootstrap (Jentsch & Lunsford 2022) with centered 68 and 90 per cent confidence intervals reported in blue.

Table A.2.1: First Stage Regression on GDP

Instrument:	Monetary Policy Shock	Forecast Error
$z_{1,t}$	-6.32*	1.11**
	(3.60)	(0.456)
$z_{2,t}$	-4.06	
	(2.689)	
constant	0.020***	0.019***
	(0.003)	(0.004)
F-stat	12.85***	40.1***
R^2	0.22	0.32

Note: Table shows results from first stage regression $\Delta GDP_{t+4} = \delta z_t + \epsilon_t$. In the monetary policy shock implementation z is accumulated high frequency shocks from one (quarters $t-1$ to $t-4$) and two years prior (quarters $t-5$ to $t-8$). In the forecast error version z is the 1 year ahead GDP forecast error at time t from the VAR detailed in Appendix A.2. HAC standard errors reported: *p<0.1; **p<0.05; ***p<0.01.

In the second approach we use the VAR to forecast GDP growth over the next four quarters and take the residual of that forecast and actual GDP growth as our instrument. This purpose of this instrument is to isolate the unexpected or surprise component of GDP growth between waves. The results from the first stage regression is outlined in Table A.2.1 where we can see that both approaches are able to explain a significant amount of the variation in GDP growth.

A.2.1 Further Results with Transitory Shocks

In this subsection we further decompose the results in Figure 4, following the approach that we followed in constructing Figures 2 and 3. Figures A.2.2 and A.2.3 decompose the GDP betas for earnings in the forecast error and monetary policy shock specifications, respectively, into extensive and intensive margin components. Figure A.2.4 decomposes the extensive margin for each identification approach into movements to unemployment, inactivity and between jobs.

A.3 Additional results

A.3.1 Household Pay

As a household survey, the LFS interviews each member of selected households which enables us to aggregate our dataset to the household level. It is possible that GDP betas for pay could look different when grouping and sorting by households instead of individuals.[15] For example, individual labour supply decisions are likely to reflect household circumstances e.g. lower labour supply by one member may be compensated for by higher labour supply by another. Figure A.3.1 compares the overall GDP beta for pay for individuals (blue dashed line) to a similar estimate conducted at the household level (purple line and shading). For most pay deciles, the estimated GDP betas are not significantly different from each other. In fact, the principle difference between the two point estimates indicates greater sensitivity at the very bottom when sorting and aggregating by household. Outside of the bottom decile, the overall profile is flatter and even upward sloping between deciles four and seven suggesting that some of the hypothesised within household insurance may be present.

[15] The majority of the related literature has conducted its analysis at the individual level.

Figure A.2.2: Decomposition of GDP Betas - Forecast Errors

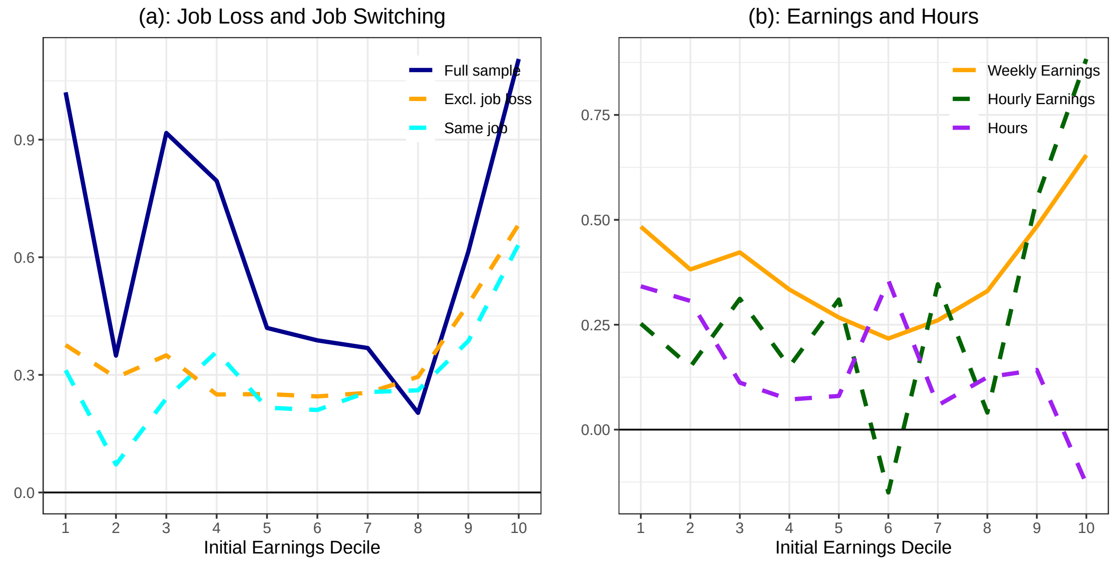

Notes: See notes for Figure 2. The coefficient estimates are the result of an instrumental variable strategy which uses GDP forecast errors from the VAR discussed in Appendix A.2 as an instrument for GDP growth.

Figure A.2.3: Decomposition of GDP Betas - MP shocks

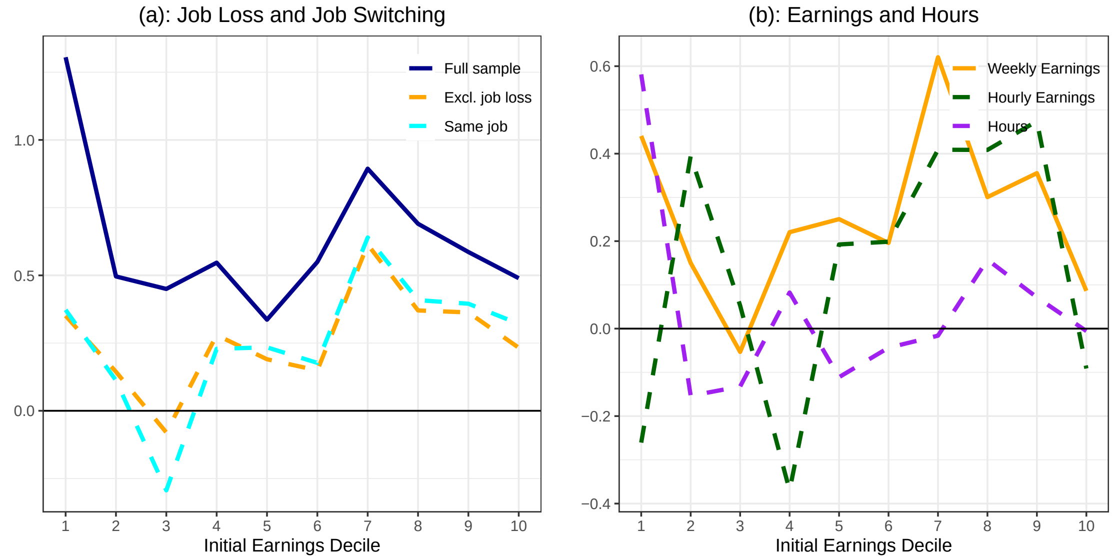

Notes: See notes for Figure 2. The coefficient estimates are the result of an instrumental variable strategy which uses accumulated high frequency monetary policy shocks as an instrument for GDP growth.

Figure A.2.4: GDP Betas for Labour Market Transition Probabilities – Transitory Shocks

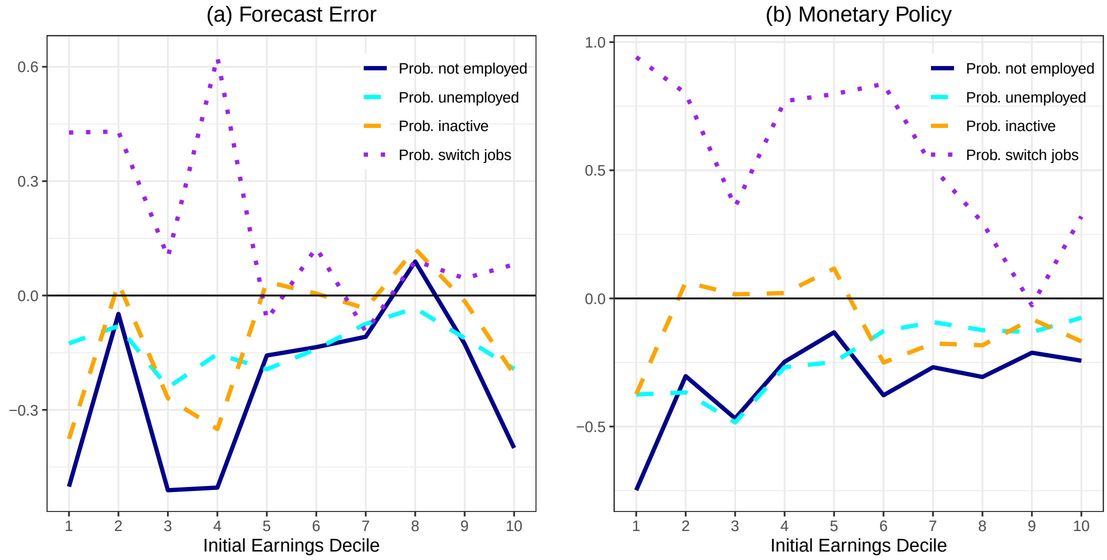

Notes: See notes for Figure 3. Panel (a) plots coefficient estimates that are the result of an instrumental variable strategy which uses GDP forecast errors from the VAR discussed in Appendix A.2 as an instrument for GDP growth. Panel (b) plots coefficient estimates that are the result of an instrumental variable strategy which uses accumulated high frequency monetary policy shocks as an instrument for GDP growth.

Figure A.3.1: GDP Betas for Household Pay

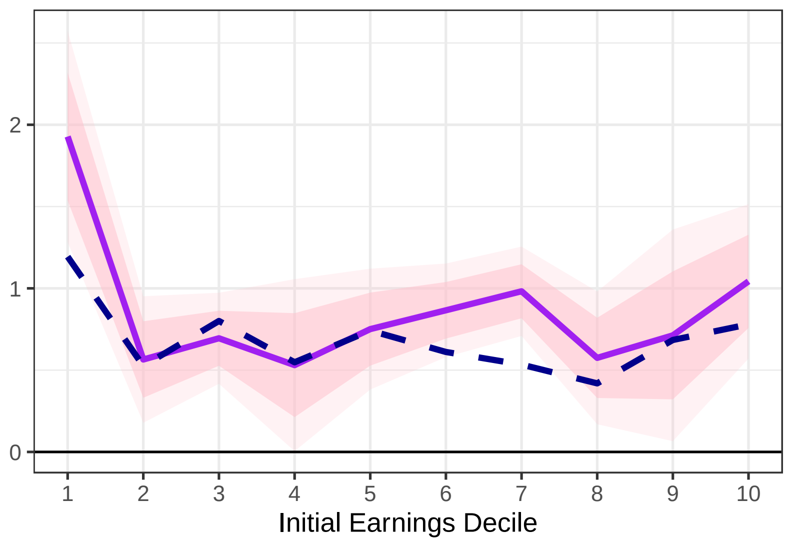

Note: The purple line plots the coefficients β_g from Equation (2) but aggregating and binning at the household level rather than the individual level. The darker (lighter) shaded area represents the 68% (90%) confidence intervals. Households are sorted into deciles in each quarter based on household earnings in t. The blue dashed line repeats the estimate from Figure 2 panel (a). The sample period is 1997Q2-2019Q4. HAC standard errors reported.

Figure A.3.2: Transitions from Non-Employment

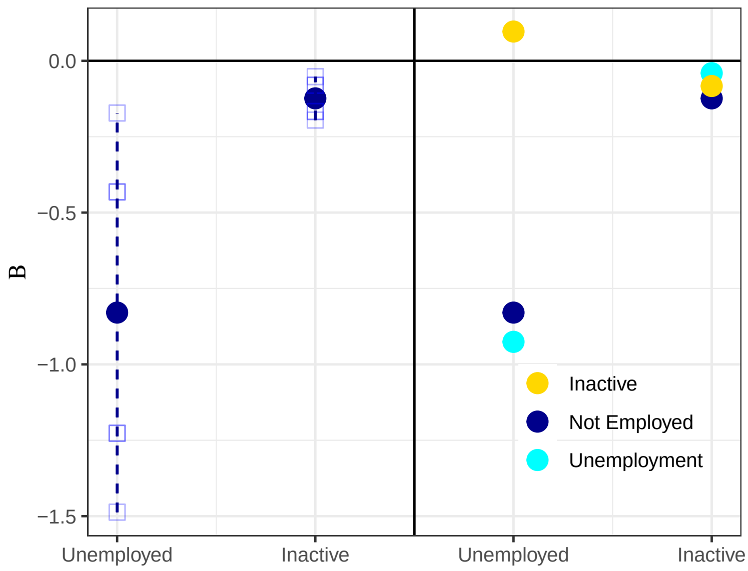

Note: Chart plots the GDP betas for the probability of being not employed/unemployed/inactive at wave $t+4$ conditional on being non-employed (inactive or unemployed) at wave t. The LHS of the chart includes the 68 and 90 per cent confidence interval for the probability of being not employed at $t+4$.

A.3.2 Transitions from Non-Employment

While we do not observe an individual's income when not employed in wave 1 (period t), we do observe their labour market status and subsequent status 4 quarters later. Figures A.3.2 and A.3.3 therefore plot the GDP betas for the probability of being not employed, unemployed and inactive. The figures show that in response to GDP fluctuations it is the unemployment to unemployment margin that dominates. In particular, the GDP beta for the probability of non-employment is higher (more negative) for the initially unemployed than the initially employed (see Figure 3).

A.3.3 Part Time Employees

Most of our results have exhibited some volatility and reversals within the first 3 deciles. And as shown in Table 1, there are a large number of part time employees in these deciles. In this sub-section, in order to better understand the role that part time employees may be playing in exaggerating or distorting our conclusions, we remove part-time employees from the initial period of our sample (period t) and re-estimate our key exercises.[16] The main results without part time employees are are shown in Figure A.3.4. Outside of the first three deciles, the responses are the same as there are very few part time employees in these deciles. The reversal between deciles two and three when using this sample is negligible. The estimated GDP beta for the first decile is now significantly larger, indicating that full time workers are more sensitive to fluctuations than part-time workers at the bottom of the income distribution.

A.3.4 Recessions

In this sub-section we assess whether controlling for (discounting) periods of negative growth, and in particular the Great Recession, affects our GDP beta estimates. We do this by including

[16] Note that employees transitioning to part-time work are retained in the sample.

Figure A.3.3: Transitions from Non-Employment (by shock)

Note: Chart plots the GDP betas for the probability of being not employed at wave $t + 4$ conditional on being non-employed (inactive or unemployed) at wave t. The chart includes the 68 and 90 per cent confidence interval which span the dashed blue line.

Figure A.3.4: GDP Betas (excluding part-time workers in t)

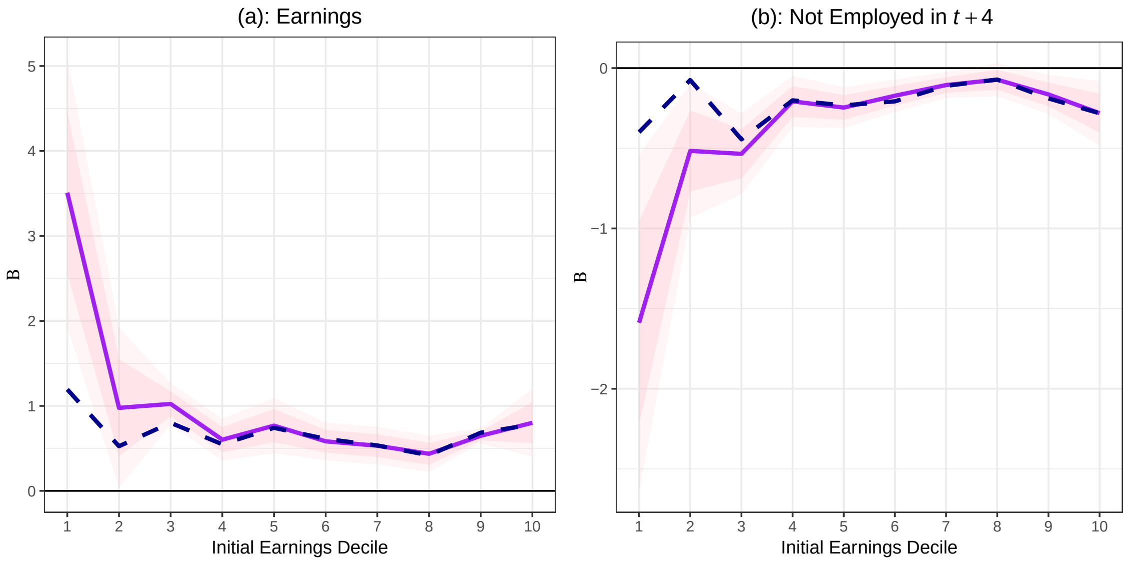

Note: The purple line plots the coefficients β_g from Equation (2) but excluding part time workers in the initial period. Income deciles are not recomputed. The darker (lighter) shaded area represents the 68% (90%) confidence intervals. Individuals are sorted into deciles in each quarter based on usual individual earnings in t. The blue dashed line repeats the estimate from Figure 2 panel (a). The sample period is 1997Q2-2019Q4. HAC standard errors reported.

Figure A.3.5: Impact of Recessions on GDP betas

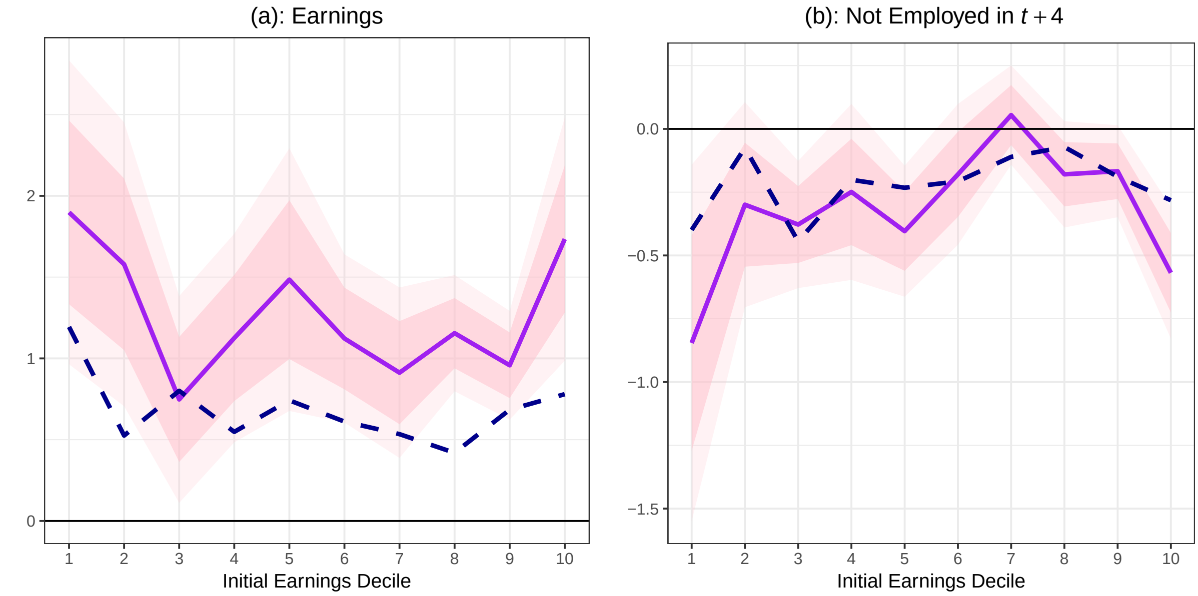

Note: Panel (a) and panel (b) plot GDP beta estimates β_g that control for the impact of recessions by including a negative growth indicator as an interaction term: $y_{g,t,4} - y_{g,t,0} = \alpha_g + \gamma_g \mathbf{1}_{\Delta GDP_{t+4}<0} + \beta_g \Delta GDP_{t+4} + \beta_{g_-} \Delta GDP_{t+4} \mathbf{1}_{\Delta GDP_{t+4}<0} + \epsilon_{g,t}$. The purple lines plot the estimate that controls for periods of negative growth which are compared to the full sample estimates in the blue dashed lines. The darker (lighter) shaded area represents the 68% (90%) confidence intervals. Individuals are sorted into deciles in each quarter based on their earnings in t. The sample period is 1997Q2-2019Q4. HAC standard errors reported.

indicators of negative growth between periods t and $t+4$ as an interaction term with GDP growth in Equation (2). Figure A.3.5 plots the estimated GDP betas from specifications which include this interaction term in purple. The point estimates from Figure 2 panel (a) and Figure 3 panel (a) are included as the blue dashed lines for comparison. Overall we see that the point estimates of the GDP betas for pay excluding the impact of recessions are largest at the bottom and smallest in the upper middle. The estimated average GDP beta across the distribution is larger, as also reported by Bell et al. (2022). This suggests that real labour incomes are more sensitive to positive than negative shocks, which is consistent with studies that find evidence of downward real and nominal rigidity.[17] Focusing on the extensive margin (panel (b)), the estimated profile is close to that from our baseline specification.

A.3.5 Income Variance

The change in the variance of income growth across the distribution is a further object of interest and a theoretically important determinant of household consumption. Figure A.3.6 plots the GDP betas for the change in the variance of pay growth by income decile for initially employed people as well as the GDP beta for the probability of changing income decile. We see that GDP betas for pay variance are negative across the distribution, and slightly more so in the bottom half though this relationship is flatter than for actual pay. The reduction in variance, while statically significant is relatively small when compared to the average variances reported in most income deciles in Figure 1. The fact that the GDP betas for pay growth variance are negative can be explained by the negative correlation between extensive margin labour market transitions and GDP movements. A positive economic shock reduces the probability of becoming unemployed (Figure 3), which reduces variance in pay growth for the initially

[17] See Grigsby et al. (2021) and references therein.

Figure A.3.6: GDP Betas for Earnings Variance

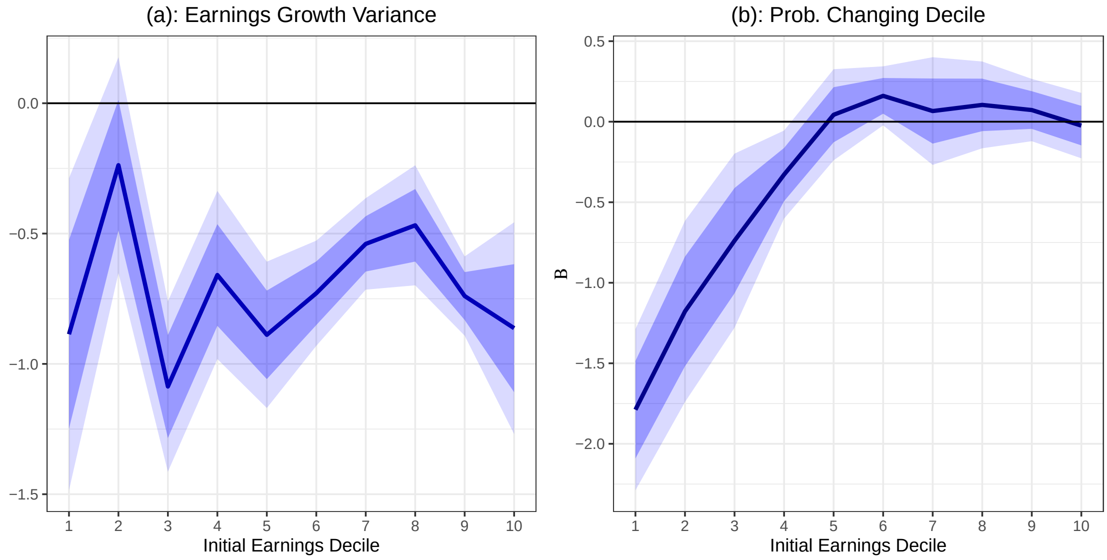

Note: Panel (a) plots the coefficients β_g from Equation (2) with within initial income decile pay variance $Var(dy_{t+4})$ on the LHS. The darker (lighter) shaded area represents the 68% (90%) confidence intervals. Panel (b) plots the GDP beta for the probability of changing income decile or changing labour market status (i.e. going to decile 0). Individuals are sorted into deciles in each quarter based on their earnings in t. The sample period is 1997Q2-2019Q4. HAC standard errors reported.

employed. This is more readily apparent in panel (b) which plots the probability of changing income deciles or becoming not employed i.e. going to a decile zero. Here there is a pronounced difference between the bottom and top half. The fact that the variance shift is quite flat and small across the distribution may be partially explained by the positive correlation between GDP movements and job switching (Figure 3), which pushes in the other direction for pay variance at the bottom of the income distribution.