# The Curious Incidence of Shocks Along the Income Distribution*

Tobias Broer† John Kramer‡ Kurt Mitman§

April 24, 2020

**Abstract**

We use detailed labor market data from Germany to study the comovement of individual earnings as well as separation and job-finding probabilities with fluctuations in aggregate output. In addition, we estimate impulse responses of these variables to monetary policy shocks identified from high-frequency changes in Overnight Indexed Swap rates. We find that workers at the bottom of the income distribution are more exposed to aggregate earnings risk in general, and to monetary policy shocks specifically. Income risk for the poor is almost entirely 'extensive', or due to cyclical fluctuations in labor market transitions, through procyclical job-finding and countercyclical separation rates. At higher incomes, in contrast, 'intensive' risk, from fluctuating income growth in continued employment, accounts for half of cyclical movements in income risk.

# 1 Introduction

How does income and employment risk vary across the income distribution? Do business cycles in general, and monetary policy interventions in particular, affect poor workers' incomes more than those of the rich? And is this mainly due to fluctuations in the likelihood of losing or finding a job, or in wages and salaries? Answering these questions is important to assess the welfare costs of business cycles and for policy design. It is also likely to be important for the transmission of shocks to aggregate demand and production, as poorer households tend to react more strongly to fluctuations in their incomes (Patterson et al., 2019), and cyclical fluctuations in labor market risk may amplify other shocks (Ravn and Sterk, 2017).

To answer these questions we use a long high-frequency panel of detailed administrative data from Germany, containing individual labor market biographies including earnings. We first study the comovement of earnings, separation rates and job-finding rates with measures of aggregate earnings activity. We find strong heterogeneity in this comovement along the

*We thank Per Krusell, Kjetil Storesletten and seminar participants at various institutions for comments. Agneta Berge and Carolina Lindholm provided outstanding research assistance.

†Paris School of Economics, IIES, Stockholm University, and CEPR, tobias.broer@iies.su.se

‡IIES, Stockholm University, john.kramer@iies.su.se

§IIES, Stockholm University, Danmarks Nationalbank and CEPR, kurt.mitman@iies.su.se

distribution of permanent incomes. In particular, income and employment risk at the bottom of the distribution are substantially more affected by business cycles than those at the top.
The high frequency of our data allows us to also study the heterogeneous effects of policy interventions on earnings and employment risk along the income distribution. In particular, we estimate impulse responses to monetary policy shocks, which we identify using high-frequency changes in Overnight Indexed Swap rates. We find that workers at the bottom of the income distribution are not only more exposed to aggregate earnings risk in general, but also to monetary policy shocks specifically. In particular, a contractionary shock to monetary policy lowers earnings at the first quintile between two and three times more than at the top quintile. The relative incidence of the policy shock on the poor in terms of employment risk is even more pronounced: the probability of becoming unemployed responds about five times as much to monetary policy at the bottom compared to the top quintile, while job-finding probabilities respond twice as strongly.
Our dataset allows us to identify the sources of this heterogeneous incidence of aggregate fluctuations on individual earnings and employment risk. For this, we decompose business cycle fluctuations in individual earnings growth into an extensive component, due to cyclical transitions in and out of employment or job switching, and an intensive component, due to cyclical income changes associated with each transition. At a frequency of one year, cyclical income risk at the bottom of the distribution is almost entirely extensive, dominated by fluctuations in job-finding rates and procyclical job-switching that is associated with large income gains for the poor. At higher incomes, in contrast, where employment relationships last substantially longer, cyclical income risk is split equally into an intensive part (from cyclical movements in job stayers' income growth) and an extensive part. The latter captures countercyclical flows into non-employment, but also countercyclical job-switching, which in the top half of the distribution is associated with income losses on average.
Interestingly, we find that much of this heterogeneity is associated with recessions, where the decline in comovement of individual and aggregate earnings along the distribution is substantial. During times of above-average growth, in contrast, the pattern is hump-shaped. One interpretation of our findings is thus that poorer workers suffer in busts that feature less, and less attractive, opportunities for finding and switching jobs, and a higher likelihood of getting fired. Higher-income workers, in contrast, benefit in booms that bring wage rises for job-stayers.
Overall, our results suggest that to understand fluctuations in income risk, we need to understand both cyclical variations in employment transitions and in the associated income changes, and the substantial heterogeneity along the income distribution. This is important for the specification of models of individual income dynamics that are basic inputs to analyses of consumption dynamics, asset prices, or fluctuations in aggregate demand. Specifically, because job-stayers experience small and relatively stable income changes relative to those who switch or lose employment it seems important to condition income dynamics on employment status. Moreover, the observation that income changes of job-switchers become smaller on average (and eventually negative) along the income distribution and less procyclical (in fact eventually countercyclical) suggests that on-the-job search and firm-induced separations are both important sources of earnings dynamics, with different relevance at different points in the distribution. Finally, while the previous feature broadly resonates with job-ladder-type earnings dynamics, the high separation and low job-finding rates at the bottom of

the distribution suggest that there are other sources of heterogeneity in income dynamics correlated with our permanent income measure.
Our results also have implications for the costs of business cycles. In particular the fact that individuals with lower permanent incomes experience more procyclical income growth mainly due to extensive movements in job finding and separation rates acts to increase the cost of business cycles relative to standard calculations based on either a representative agent or homogeneous income risk. Finally, monetary policymakers concerned about the welfare consequences of their policies should be aware that these have disproportional effects on the poor.

Our paper contributes to a growing literature ...

The next section presents the data, and some key descriptive statistics. Section 3 desribes the structure of income and employment risk in the sample on average. Section 4 studies business cycle fluctuations in risk. Section 5 looks at impulse responses after a monetary policy shock.

# 2 Data

We work with the anonymised version of the Sample of Integrated Employment Biographies, obtained from the Research Data Center (FDZ) in Germany. It contains administrative social security data for about 1.7 million individuals, excluding civil servants and self-employed individuals, for the years between 1975 and 2014. Each observation in the dataset is a labor-market spell.[1] We convert these spells into monthly employment histories for each individual, resulting in about 300 million person-month observations.[2] Each of these contains information about labor market status, receipt of unemployment benefits and daily gross earnings, calculated as the cumulated gross earnings during an employment spell, divided by the length of the spell in days. If an individual has multiple employers, we sum her earnings from all jobs in a month. We count a month (or quarter) as belonging to an employment spell if the worker is employed for more than half of it. In addition, the data contains information about key individual characteristics, such as age, gender, education level, and employer location. The maximum (and median) employment spell length is 12 months, as employers are required to report total wage payments and days worked once a year, unless the employment relationship is terminated before the end of this period.[3] Time aggregation thus affects our measure of monthly individual earnings (which are calculated as an average over the employment spell), but not the information on employment status.
The data are censored from above at the upper earnings limit specified by an individual's pension insurance, which increases over time. This applies to around 6-8% of our earnings observations each month. Furthermore, until 1999, earnings below the social security reporting

[1] Employment relationships longer than 12 months are split into multiple spells. We drop spells that are shorter than 1 month. Potentially missing spells are imputed according to Drews et al. (2007).

[2] When we investigate unconditional risk, we use quarterly data. For our estimation of the effect of monetary policy surprises on the labor market, we use monthly data.

[3] Another reason for the end of a spell is an employee's change of health insurance provider.

threshold (marginal part-time work) were not recorded.[4]
We use two different definitions of employment. First, we define individuals as "fully attached" to the labor market if they are liable to social security without special characteristics. This definition excludes trainees, marginal part-time workers, employees in partial retirement, interns, student trainees and casual workers, all of which account for about 20% of total employment. Our second, broader definition also includes all individuals whose employment status falls under the aforementioned special characteristics, unless explicitly stated otherwise. We define individuals as non-employed if they are not employed.

# Descriptive Statistics

We study the differences in income- and employment-risk across the income distribution by ranking individuals in a given period according to a proxy measure of their *permanent* income, which is less prone to noise from temporary income movements and should capture welfare and consumption differences more accurately than current incomes. In what follows, we assign individuals to quantiles based on this permanent income measure and refer to the distribution simply as the "income distribution". Our prefered proxy for permanent income is average earnings over the five years preceding quarter $t$ as in Guvenen et al. (2017), but we also study alternative definitions below.[5] Importantly, the top ten percentiles are directly affected by the upper reporting limit for earnings (which grows discretely over time in nominal terms). For this reason, the comovement between individual and aggregate earnings for these percentiles is not informative and we do not report results for this group. Furthermore, we exclude individuals whose employer is located in counties which belonged to the German Democratic Republic before 1990, as their inclusion leads to discontinuities.
To understand how key variables evolve along this income distribution, Tables 1a and 1b report descriptive statistics for the whole sample, split into deciles of our permanent income measure in January 2000 and 2010.[6] Note that, with some abuse of language but hopefully no room for confusion, we call deciles both the 9 points of the distribution as well as the 10 groups they define (we proceed similarly for other quantiles).
Nominal quarterly earnings, defined as the sum of labor earnings and benefits, increase across the permanent income distribution Tables 1a and 1b, with the average income in the highest deciles between four and five times larger than the same figure in the first decile. The percentage of women decreases markedly across deciles, and the fraction of women in the upper deciles increases only mildly over time. Importantly for our analysis, which, for the most part, abstracts from life-cycle heterogeneity, the mean age increases only modestly across the earnings distribution, with individuals in the top quintile only around three years older than the average age in the sample in 2000. Education levels are rather flat from

[4] We impute data that is likely due to spell errors following (Böhm et al., 2019) and impute education where data are missing or inconsistent following Fitzenberger (2005).

[5] Due to the construction of permanent income, our sample is restricted to workers who have at least one earnings observation in the five years prior to period $t$.

[6] The number of individuals in each decile differs because daily earnings are rounded to whole Euros. Education takes a value of 1 for individuals without a degree, 2 for vocational training, 3 for high school, 4 for high school and vocational training, 5 for graduates of technical colleges and 6 for university graduates. We impute education following imputation procedure 1 in Fitzenberger et al. (2005).

Table 1: Descriptive statistics by Decile

(a) First Quarter 2000

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
|---|---|---|---|---|---|---|---|---|---|---|
| | mean | mean | mean | mean | mean | mean | mean | mean | mean | mean |
| Female | 0.70 | 0.62 | 0.64 | 0.63 | 0.52 | 0.40 | 0.31 | 0.27 | 0.23 | 0.12 |
| Age | 38.19 | 39.98 | 41.02 | 40.99 | 39.96 | 39.80 | 40.95 | 42.20 | 43.05 | 45.64 |
| Education | 1.67 | 1.77 | 1.83 | 2.13 | 2.17 | 2.16 | 2.18 | 2.33 | 2.57 | 3.42 |
| Daily wage | 84.65 | 128.50 | 184.97 | 231.55 | 267.18 | 298.23 | 319.52 | 335.54 | 347.87 | 370.08 |
| Empl next year | 0.27 | 0.40 | 0.55 | 0.66 | 0.74 | 0.81 | 0.86 | 0.89 | 0.90 | 0.91 |
| Fully Attached | 0.33 | 0.47 | 0.64 | 0.76 | 0.83 | 0.89 | 0.93 | 0.94 | 0.95 | 0.96 |
| Non-employed | 0.62 | 0.47 | 0.32 | 0.22 | 0.15 | 0.09 | 0.06 | 0.04 | 0.04 | 0.02 |
| Observations | 47203 | 47203 | 47203 | 47267 | 47138 | 47203 | 47202 | 47203 | 47203 | 47202 |

(b) First Quarter 2000

| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) | (10) |
|---|---|---|---|---|---|---|---|---|---|---|
| | mean | mean | mean | mean | mean | mean | mean | mean | mean | mean |
| Female | 0.65 | 0.64 | 0.66 | 0.64 | 0.52 | 0.42 | 0.36 | 0.31 | 0.25 | 0.14 |
| Age | 39.94 | 40.66 | 41.87 | 42.10 | 41.64 | 42.33 | 43.55 | 44.54 | 45.15 | 46.95 |
| Education | 1.84 | 2.04 | 2.18 | 2.28 | 2.35 | 2.34 | 2.39 | 2.54 | 2.95 | 4.01 |
| Daily wage | 106.01 | 189.53 | 249.26 | 292.57 | 323.57 | 350.31 | 370.72 | 387.08 | 405.40 | 433.30 |
| Empl next year | 0.31 | 0.54 | 0.67 | 0.75 | 0.80 | 0.84 | 0.87 | 0.89 | 0.91 | 0.93 |
| Fully Attached | 0.37 | 0.61 | 0.75 | 0.83 | 0.87 | 0.90 | 0.92 | 0.93 | 0.93 | 0.95 |
| Non-employed | 0.56 | 0.32 | 0.21 | 0.15 | 0.11 | 0.08 | 0.05 | 0.04 | 0.03 | 0.02 |
| Observations | 44741 | 44740 | 44741 | 44745 | 44735 | 44741 | 44740 | 44741 | 44740 | 44740 |

**Note:** The table shows average values for selected variables across deciles of the permanent income distribution in the first quarter of 2000. Education takes a value of 1 for individuals without a degree, 2 for vocational training, 3 for high school, 4 for high school and vocational training, 5 for graduates of technical colleges and 6 for university graduates. Quarterly earnings are computed as average daily earnings (see text) multiplied by 75. Fully attached individuals are those liable to social security without special characteristics.

deciles 4 to 7, but are substantially lower (higher) at the bottom (top) deciles, respectively. Attachment to the labor market in general, and employment rates in particular, are low in the bottom decile, where more than 50 percent of individuals are non-employed, but rise steeply across the bottom half of the distribution, flattening thereafter.

# 3 Income and employment risk along the income distribution

This section presents some key facts about average individual income growth and labor market transition frequencies along the income distribution in our data that are important to understand the cyclical fluctuations of income risk. First, less stable employment and lower job-finding rates in the lower part of the income distribution explain higher unemployment rates there. Conditional on staying employed, however, income growth of poorer workers is higher than that of the rich, particularly for job-switchers (who experience strong income growth at the bottom, but income declines in the upper half of the distribution). This implies overall mean reversion in incomes in the full sample.

Figure 1: Average earnings growth

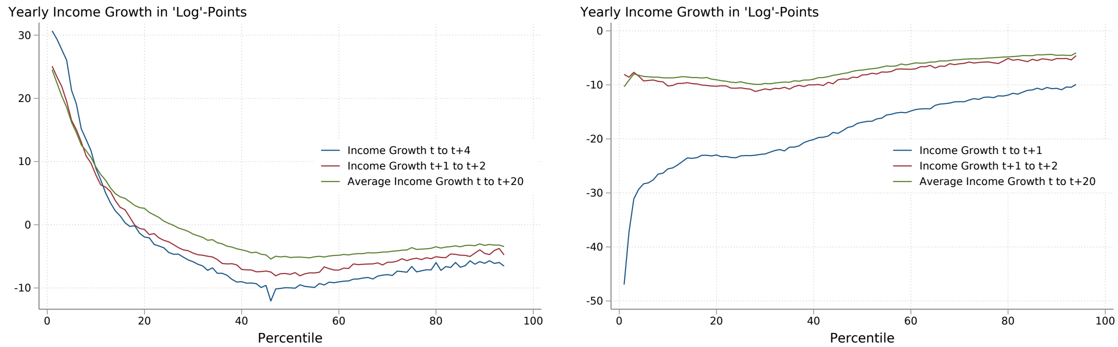


a) Full sample

b) Fully attached

**Note:** The *Left Panel* plots the average income growth for percentiles 1 to 95 for the full sample, for three different horizons. Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. Income growth is computed as difference in the hyperbolic sine transform of post-benefit income between quarters $t$ and $t+4$ (blue), $t+5$ and $t+9$ (red) and the average yearly growth rates between $t$ and $t+20$ (green). The sample period is from 1980 until 2014. The *Right Panel* plots average income growth for percentiles 1 to 95 of the subsample of individuals that are fully attached in period $t$.

Figure 1 reports total average earnings growth rates for three different horizons across the distribution. In order to account for transitions into and out of employment, which may imply 0 earnings, we use the hyberbolic sine transformation to calculate income growth

$$\Delta earn_{j,t} = sinh^{-1}(earn_{j,t+4}) - sinh^{-1}(earn_{j,t}) \tag{1}$$

where $sinh^{-1}(x) = ln(x + \sqrt{x^2+1})$. Away from zero, this measure is equivalent to using the logarithm.

The first pattern apparent from the left panel of Figure 1 is one of mean reversion, as the average five-year growth rate is strongly positive at the bottom of the distribution, but negative above the median (if less strongly so). This finding is consistent with similar patterns documented for the US (Guvenen et al., 2015).

Apart from overall mean reversion, the left panel of Figure 1 also shows a mild U-shape, as average income growth is moderately increasing in the upper half of the income distribution. This fact is more apparent when focusing only on individuals who are fully attached in period $t$, as done in the right panel of Figure 1. Their average income growth increases strongly across the distribution, particularly at the 1 year horizon.

Figure 2 shows average probabilities of labor market transitions in the full sample.[7] Job-finding rates (N to E) are substantially lower for the income-poor, while at the same time,

[7] Conditional on labor market transitions the patterns for the fully attached are very similar to those of the full sample in Figure 2 and thus not shown

their employment is less stable. In particular, transitions to non-employment (E to N) are twice as likely in the bottom decile than on average. Together, these two forces explain the strongly falling incidence of nonemployment along the income distribution in Table 1b.

The right panel of Figure 2 shows that labor market transitions are strongly related to income changes, with gains and losses for those that find and lose jobs, respectively (depicted along the right axis), an order of magnitude larger than income changes for the continually employed.[8] For the latter, there is substantial heterogeneity in income growth along the distribution: Poor job-switchers see strong income gains (of more than 25 percent on average in the bottom quintile), while job-switchers in the upper half of the distribution experience income losses of between 5 and 10 percent on average. Job-stayers also experience substantial average income growth in the bottom quintile while their income is unchanged on average in the rest of the distribution.

The mean reversion in Figure 1 thus results from substantially higher income growth of job-switchers and stayers at the bottom of the distribution. The U-shape in the left panel, and the strongly increasing nature of average income growth among the fully attached, in contrast are due to the falling incidence of job-switching and job loss as permanent income levels rise, both of which are, in that part of the distribution, associated with income losses.

Figure 2: Average transition rates and earnings growth by labor market transition

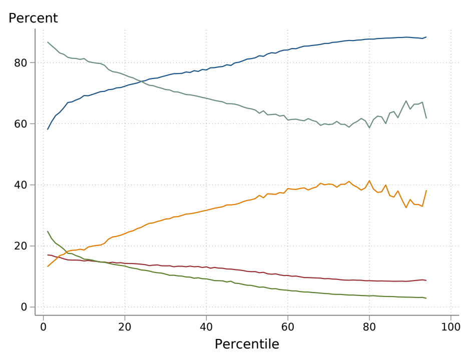


a) Average transition rates

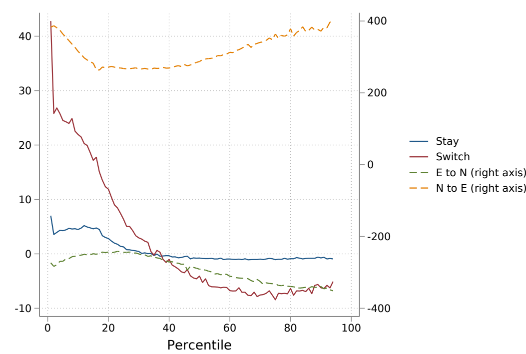


b) Average earnings growth by labor market transition

**Note:** The *Left Panel* plots the average 4-quarter frequencies of transition between labor market states for percentiles 1 to 95 for the full sample. Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. Labor market transitions are defined as having the same job in period $t + 4$ as in $t$ (Stay), having alternative employment (Switch), moving from employment to non-employment (E-to-N), from non-employment to employment (N-to-N), or remaining non-employed (N-to-N). The sample period is from 1980 until 2014. The *Right Panel* plots average income growth within groups that make a given labor market transition for percentiles 1 to 95 for the full sample. Income growth is computed as difference in the hyperbolic sine transform of post-benefit income between quarters $t$ and $t + 4$.

[8] Note that the hyperbolic sine cannot be interpreted as percentage changes in cases of growth to or from 0. The sine-difference may be considerably larger than 100, which is impossible for percentage changes, as individuals cannot lose more than their earnings in case of non-employment.

# 4 The cyclicality of income and employment risk

This section studies business cycle fluctuations in income and employment risk across the income distribution. For this we compute, within each percentile, the comovement of individual earnings growth with measures of aggregate activity. The procedure is close to that in Guvenen et al. (2017). Importantly, our data set also allows us to study how frequencies of labor market transitions comove with aggregate activity at different income levels. Hence, we can decompose income risk into an intensive margin, associated with cyclical income changes within a transition; and an extensive margin, due to cyclical changes in transition probabilities. We find that income growth is substantially more procyclical for the income poor. Moreover, cyclical income risk at the bottom of the distribution is almost entirely extensive, dominated by fluctuations in job-finding rates and procyclical job-switching that is associated with large income gains for the poor. Further up the income distribution, income growth is less procyclical, accounted for in equal parts by intensive risk (from procyclical wages for job-stayers) and extensive risk (from countercyclical separation rates).

To quantify the relative size of cyclical variation in earnings risk across the income distribution, we follow Guvenen et al. (2017) and estimate the comovement between individual and aggregate activity using the following regression[9]:

$$\Delta earn^p_{j,t} = \alpha + \beta^p_Y \Delta Y + \epsilon_{j,t} \tag{2}$$

where $\Delta earn^p_{j,t}$ represents the growth in real earnings, defined as the sum of labor earnings and benefits, of an individual $j$ in percentile $p$. Earnings growth is measured as the hyperbolic sine-difference between quarters $t$ and $t+4$. $\Delta Y$ represents the growth rate in real aggregate activity between quarters $t$ and $t+4$. Our preferred measure of aggregate activity is aggregate earnings ($Y = EARN$), calculated as the sum of earnings for all employed, fully-attached individuals in our dataset. We denote the regression coefficient in equation (2) for this measure $\beta^p_{EARN}$. As an alternative measure of aggregate activity, we also consider real GDP ($Y = GDP$), in which case we denote the regression coefficient $\beta^p_{GDP}$.[10]

## 4.1 Income and employment risk for the employed

We first look at income and employment risk for the employed. Figure 3 plots the regression coefficients $\beta^p_{EARN}$ and $\beta^p_{GDP}$ in Equation (2) by percentile for individuals who are fully-attached in $t$ and employed in $t+4$ (but may have been unemployed or switched jobs between these periods). On average, a rise in aggregate earnings growth of 1 percentage point increases individual incomes by half as much (in the left panel). The degree of comovement is strongly heterogenous across the income distribution: the coefficient $\beta^p_{EARN}$ declines as income rises in the bottom quartile, then remains flat at about 50 percent up until the 80th percentile

[9] We focus on quarterly earnings because one of our measures of aggregate activity, GDP, is reported quarterly.

[10] We obtain quarterly data for deflated, seasonally adjusted GDP from the German Federal Statistical Office. Because no continuous series exists for our sample period, we merge data from 1970-1991 with data from 1991-2018. After normalizing both series by their 1991Q1 values, we append the latter to the former for quarters after 1991Q1.

and declines strongly throughout the top quintile. We attribute part of this decline at the top of the distribution to an increasing incidence of censoring. The comovement of individual income growth with GDP growth, as measured by the coefficient $\beta^p_{GDP}$ (in the right panel) is on average only about half as large, and less precisely estimated. Its shape is similar to that for aggregate earnings, although somewhat more U-shaped around a local minimum at around the 20th percentile.

Figure 3: Regression coefficients $\beta^p_{EARN}$ and $\beta^p_{GDP}$ - employed


a) Regression coefficients $\beta^p_{EARN}$ b) Regression coefficients $\beta^p_{GDP}$

**Note:** The *Left Panel* plots the coefficient $\beta^p_{EARN}$ in Equation (2) for percentiles 1 to 95 for individuals who are categorized as fully attached to the labor force in period $t$ and as employed in period $t+4$. Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. Income growth is computed as difference in the hyperbolic sine transform of post-benefit income between quarters $t$ and $t+4$. The sample period is from 1980 until 2014. The shaded area indicates 95 percent confidence bands. The *Right Panel* plots the coefficients $\beta^p_{gdp}$ across percentiles with the same sampling restrictions.

Figure 4 plots the regression coefficients $\beta^p_{EARN}$ and $\beta^p_{GDP}$ in Equation (2) by percentile for all individuals who are fully-attached in $t$, but leaves their status unrestricted in period $t+4$. Hence, it allows for an extensive margin of earnings cyclicality.

The average comovement of aggregate and individual earnings growth is more than doubled, and equals about 1 for the three upper quartiles. The bottom quartile, in contrast, shows substantially higher sensitivity of income growth to aggregate fluctuations. A one percentage point increase in aggregate earnings growth thus increases aggregate earnings in the bottom decile by more than one and a half percentage points.

Figure 4: Regression coefficients $\beta^p_{EARN}$ and $\beta^p_{GDP}$ - initially employed

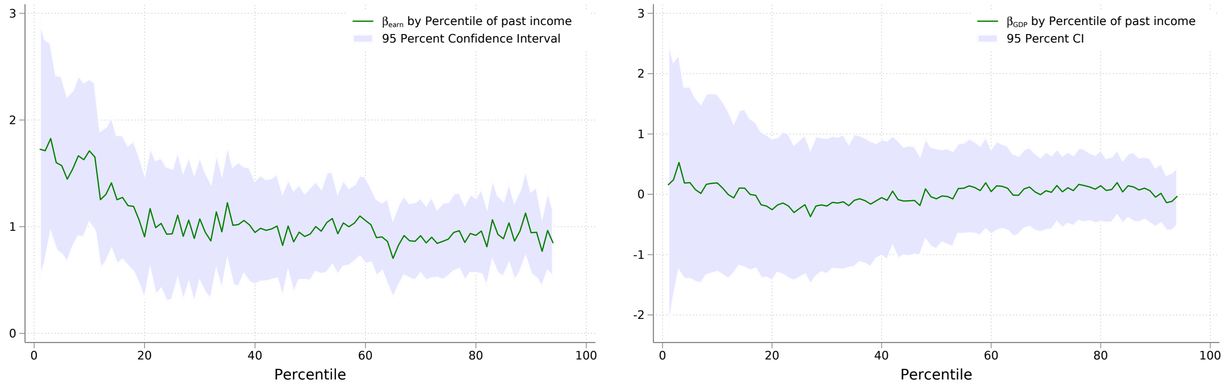


a) Regression coefficients $\beta^p_{EARN}$ b) Regression coefficients $\beta^p_{GDP}$

**Note:** The *Left Panel* plots the coefficient $\beta^p_{EARN}$ in Equation (2) for percentiles 1 to 95 for individuals who are categorized as fully attached to the labor force in period $t$. Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. The sample period is from 1980 until 2014. The shaded area indicates 95 percent confidence bands. The *Right Panel* plots the coefficients $\beta^p_{gdp}$ across percentiles with the same sampling restrictions.

The large increase in the $\beta^p_{EARN}$ coefficients in Figure 4 suggests that countercyclical job separations are a key source of cyclicality. To quantify the cyclicality of labor market transitions for the employed more generally, we estimate the following regression

$$TR^j_{s_1,s_2} = \beta^p_{TR}\Delta EARN_t + \varepsilon_{j,t} \tag{3}$$

Here $TR^j_{s_1,s_2}$ is a binary variable indicating whether an individual $j$ transitioned from state $s_1$ to $s_2$. In this section, $s_1$ corresponds to fully-attached employment, while $s_2$ captures three possible transitions: staying at the same employer (labelled "stay" below), switching employers ("switch"), or moving to non-employment ("E to N"). Intuitively, the coefficient $\beta^p_{TR}$ measures the percentage point change in the transition probability in response to a one percentage point increase in GDP growth.

The left panel of Figure 5 depicts the regression coefficients $\beta^p_{TR}$ across the income distribution. As expected, the probability of moving to non-employment is strongly countercyclical at the bottom of the income distribution: for an individual in the 10th percentile, an increase in aggregate earnings growth by one percentage point is associated with a decline in the probability of moving to non-employment by about half a percentage point. This effect declines strongly throughout the distribution, however, to about a a tenth of a percentage point at the 80th percentile. The probability of staying in the same job is procyclical throughout the distribution, but, again, more so in the bottom quintile. Interestingly, the probability of switching jobs is procyclial below the median, but countercyclical above. This is consistent with an interpretation where for the income poor, job-switching is mainly an opportunity (that materialises more often in times of high aggregate activity), but a hazard for the rich (whose likelihood declines in good times).

Because different labor market transitions are associated with vastly different income growth on average (as shown in Figure 2 in Section 2), the left panel of Figure 5 captures an 'extensive' source of income risk from transitions between labor market states. The right panel of Figure 5, in contrast, captures 'intensive risk' from fluctuations of income growth conditional on a given transition. It plots the regression coefficients $\beta^p_{EARN}$ in Equation (2) estimated separately for individuals in a percentile who share a given transition across labor market states. Income growth for job stayers comoves positively with aggregate earnings, and becomes moderately more procyclical throughout most of the distribution. Income growth of job switchers is approximately acyclical, and constant across the distribution. The dominant pattern in the right panel of Figure 5, however, is the strong decline along the distribution in the comovement between aggregate and individual income growth of job losers, which is strongly countercyclical at the bottom of the distribution but approximately acyclical for the income-rich.

Figure 5: Regression coefficients $\beta^p_{TR}$ and $\beta^p_{EARN}$ for different labor market transitions - initially employed

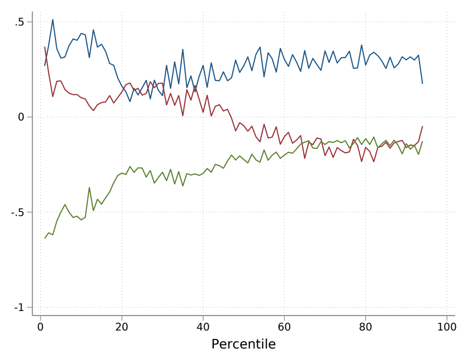


a) Regression coefficients $\beta^p_{TR}$

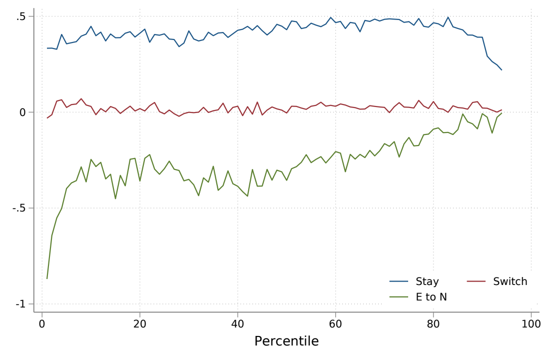


b) Regression coefficients $\beta^p_{EARN}$ for groups that share labor market transition

**Note:** For individuals who are fully attached to the labor force in period $t$, the *Left Panel* plots the coefficient $\beta^p_{TR}$ in Equation (3) from separate regressions for those individuals who in period $t+4$ have the same job (Stay), have alternative employment (Switch), or are non-employed (E-to-N). Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. The sample period is from 1980 until 2014. The *Right Panel* plots the coefficients $\beta^p_{EARN}$ in Equation 2 estimated separately for groups that make the aforementioned transitions.

In order to decompose the comovement between individual and aggregate earnings in Figure 4 into an intensive and an extensive margin, we can combine the sample averages presented in Figures 1 and 2 with the cyclicality measures presented in Figure 5 as follows:

$$\beta_{EARN} \approx \beta_{ext} + \beta_{int} \tag{4}$$

$$\beta_{int} = \sum_{k=1}^{J} \left[\bar{\mu}_k \beta_{int,k}\right] \tag{5}$$

$$\beta_{ext} = \sum_{k=1}^{J} \left[(\bar{\Delta} earn_k - \bar{\Delta} earn)\beta_{tr,k}\right] \tag{6}$$

where we suppress the percentile subscript and $\beta_{ext}$ and $\beta_{int}$ are the measures of comovement due to, respectively, extensive and intensive income risk. The index $k$ represents the three different groups of stayers, switchers, and job losers. $\Delta earn_{kt}$ and $\mu_{kt}$ are, respectively, group $k$'s average income growth and its share in the total number of individuals in percentile $p$ in period $t$. Bars over variables denote averages over time. The share of comovement due to intensive risk $\beta_{int}$ in Equation (23) equals the sum of the group-specific regression coefficients $\beta_{EARN,k}$ weighted by the groups' share in the sample. Equation (24) expresses the share of comovement that arises from extensive risk as the sum of regression coefficients $\beta_{tr,k}$ weighted by the groups' average income growth $\bar{\Delta} earn_k$ (as a deviation from the mean $\bar{\Delta} earn$). As shown in A.6, the decomposition is exact up to a third order term (that equals less than 0.06 percent on average).

Figure 6 shows the decomposition. As suggested by Figure 5, cyclical movements in extensive risk, in the form of countercyclical separations and procyclical job-finding, are the dominant source of cyclical income risk at the bottom of the distribution. Their share falls, however, across the distribution, while that accounted for by intensive risk rises. In the top quintile the two account for approximately equal shares of cyclical income risk. This is because both job loss and job-switching are unlikely for the income-rich. So intensive risk (from procyclical incomes of stayers), and extensive risk (from procyclical probabilities of staying in the same job) are the main sources of cyclical income risk, approximately accounting for equal shares of the overall comovement between individual and aggregate earnings.

Figure 6: Decomposition of $\beta^p_{EARN}$ - initially employed

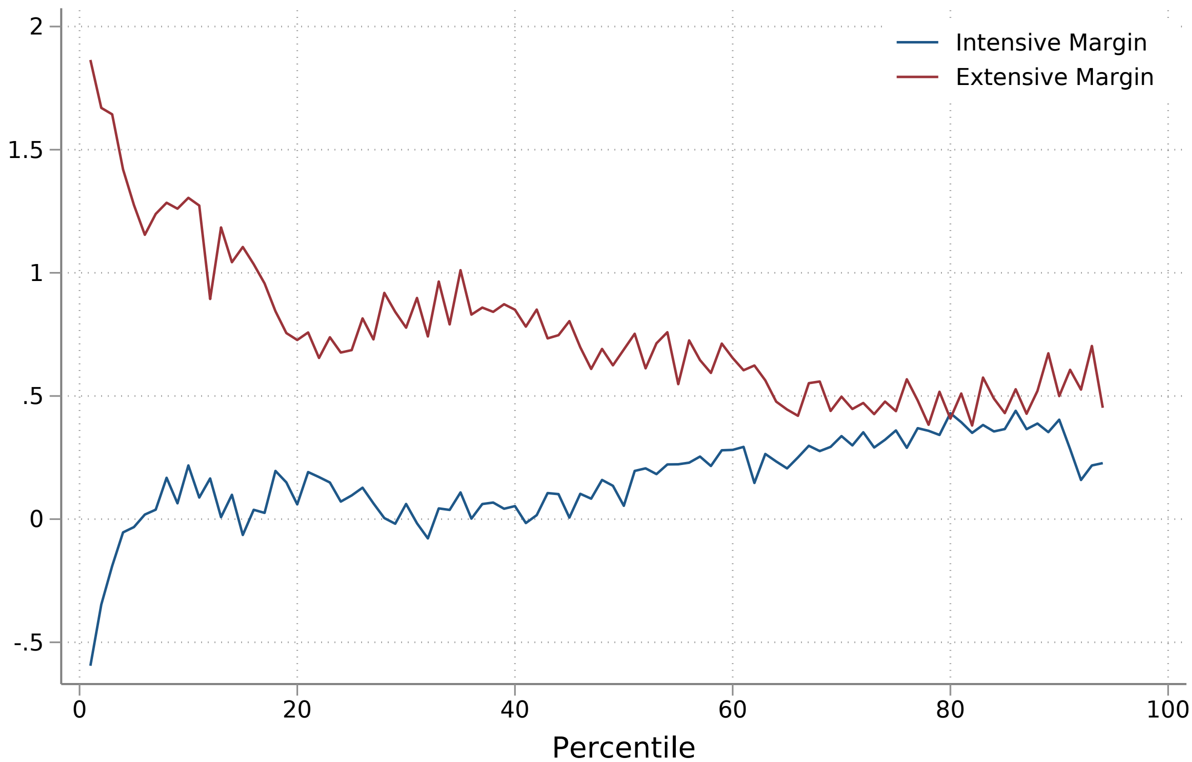


**Note:** For individuals who are fully attached to the labor force in period $t$, the Figure presents the two sources of cyclical income risk $\beta_{ext}$ (red line) and $\beta_{int}$ (blue line) according to the decomposition (22). Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. The sample period is from 1980 until 2014.

## 4.2 Income and employment risk in the whole sample

This section studies cyclical income and employment risk across the income distribution in the full sample of individuals, where non-employment and partial attachment to the labor market are common, particularly in the lower quantiles. We show how income risk is, again, substantially more procyclical below median income, mainly because of procyclical job-finding and countercyclical separation rates. Extensive labor market risk, from cyclical transition rates, thus accounts for more than three quarters of the cyclicality below the median. The importance of intensive risk, from cyclical income growth within given labor market transitions, increases along the distribution, and again accounts for approximately half of the cyclical income risk in the top quintile.

To interpret the results, it is important to bear in mind the heterogeneity within percentiles of the full sample. In particular, as Tables 1a and 1b show, individuals who are non-employed or employed but not fully attached account for more than half the sample in the bottom percentiles. In line with this, job-finding rates are lower and separation rates higher there, making labor market transitions the dominant source of cyclical income risk. In the upper half of the distribution, in contrast, the fully attached account for almost 95 percent of the sample. There, the results for the full sample therefore mirror those for the (initially) fully attached individuals considered in Section 4.1.

Figure 7 depicts the coefficients $\beta^p_{EARN}$ (left panel) and $\beta^p_{GDP}$ (right panel) in Equation (2) for the whole sample. Overall, individual income risk in the whole sample is substantially more cyclical than for the fully attached in Section 4.1 (although the coefficients are also less precisely estimated). The pattern of strongly declining cyclicality is, however, similar. In particular, a one percentage point increase in the growth of aggregate earnings increasing individual earnings growth by between 1.5 and 2 percentage points in the bottom quartile of the distribution. Above the median, individual income growth moves about one-for-one with aggregate earnings.

Figure 7: Regression coefficients $\beta^p_{EARN}$ and $\beta^p_{GDP}$ – Full sample


b) Regression coefficients $\beta^p_{EARN}$

b) Regression coefficients $\beta^p_{GDP}$

**Note:** The *Left Panel* plots the coefficient $\beta^p_{EARN}$ in Equation (2) for percentiles 1 to 95 for all individuals. Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. Income growth is computed as difference in the hyperbolic sine transform of post-benefit income between quarters $t$ and $t + 4$. The sample period is from 1980 until 2014. The shaded area indicates 95 percent confidence bands. The *Right Panel* plots the coefficients $\beta^p_{gdp}$ across percentiles with the same sampling restrictions.

Figure 8: Regression coefficients $\beta^p_{TR}$ and $\beta^p_{EARN}$ for different labor market transitions - full sample

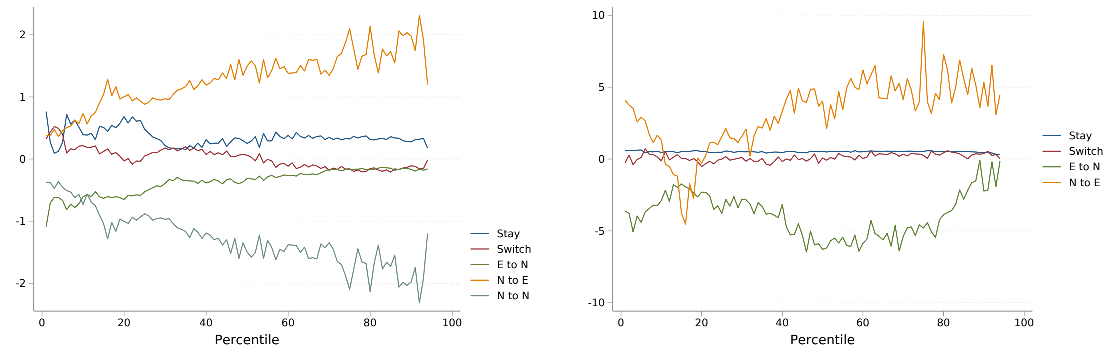


a) Regression coefficients $\beta^p_{TR}$

b) Regression coefficients $\beta^p_{EARN}$

**Note:** For the whole sample, the *Left Panel* plots the coefficient $\beta^p_{TR}$ in (3) from separate regressions for those individuals who in period $t + 4$ have the same job (Stay) as in $t$, have alternative employment (Switch), move form employment to non-employment (E-to-N), move from non-employment to employment (N-to-N), or remain non-employed (N-to-N). Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. The sample period is from 1980 until 2014. The *Right Panel* plots the coefficients $\beta^p_{EARN}$ estimated separately for groups that make the aforementioned transitions.

Figure 8 depicts the cyclicality of labor market transitions (left panel) and of average income growth within groups that share labor market transitions (right panel). Job-finding rates of the non-employed are strongly procyclical, and become strongly more so when non-employment becomes less common along the income distribution. Transitions from employment exhibit similar cyclicality as those considered in Section 4.1: all are more cyclical in the two bottom quintiles. Job-switching is again procyclical there, but eventually turns mildly countercyclical in the upper half of the distribution. Separation rates are strongly countercyclical for the bottom two quintiles, but - as with the fully attached - their cyclicality declines with higher incomes.

Figure 9: Decomposition of $\beta^p_{EARN}$ - full sample

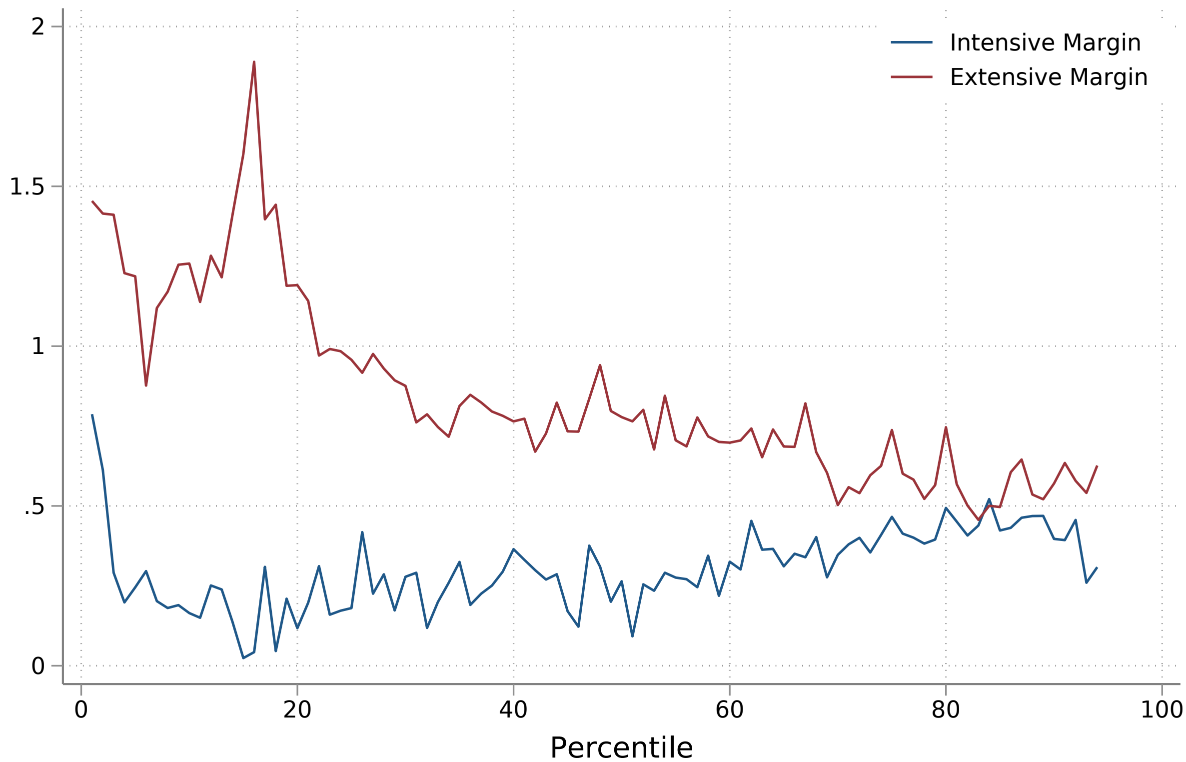


**Note:** The Figure presents the two sources of cyclical income risk $\beta_{ext}$ (red line) and $\beta_{int}$ (blue line) according to the decomposition (22).

Figure 9 shows how in the whole sample, the regression coefficients $\beta^p_{EARN}$ in Figure 8 sum to shares of intensive and extensive risk that have a similar shape across the distribution as those of the fully attached in Section 4.1. Cyclical income risk at the bottom of the distribution is thus strongly dominated by the cyclicality of extensive risk from labor market transitions, particularly in job finding and separation rates. Moving up the income distribution this extensive source of risk becomes substantially less important as job-staying becomes the dominant transition. With the probability of staying in the same job and the associated income growth both procyclical, extensive and intensive sources of cyclical income risk are again of approximately equal importance for cyclical income risk in the upper quintile.

## 4.3 Asymmetries over the business cycle

Are the effects we find symmetric over the business cycle or is the correlation we find dominated by booms, or busts? To answer this question, we divide our sample into periods with above and below average aggregate earnings growth and run the above analysis for each period separately.[11]

Figure 10 reports the results from this analysis. The red line traces out the regression coefficients $\beta^p_{earn}$ for periods of recessions, the blue line for booms. The difference between the two is striking and significant. The pattern generated in recessions is similar to the one uncovered in previous graphs. The regression coefficients $\beta^p_{EARN}$ estimated in booms, in contrast, follow a hump shape across the distribution. Again, this pattern sets apart the bottom quintile (where income growth responds strongly to fluctuations in aggregate earnings growth during periods of below-average activity) from much of the rest of the distribution (where income growth is more responsive in booms, at least until the top quintile). The pattern again reverses in the top quintile (where bonus payments might make incomes more responsive to aggregate conditions in booms).

Figure 10: Asymmetry of regression coefficients $\beta^p_{EARN}$ - full sample

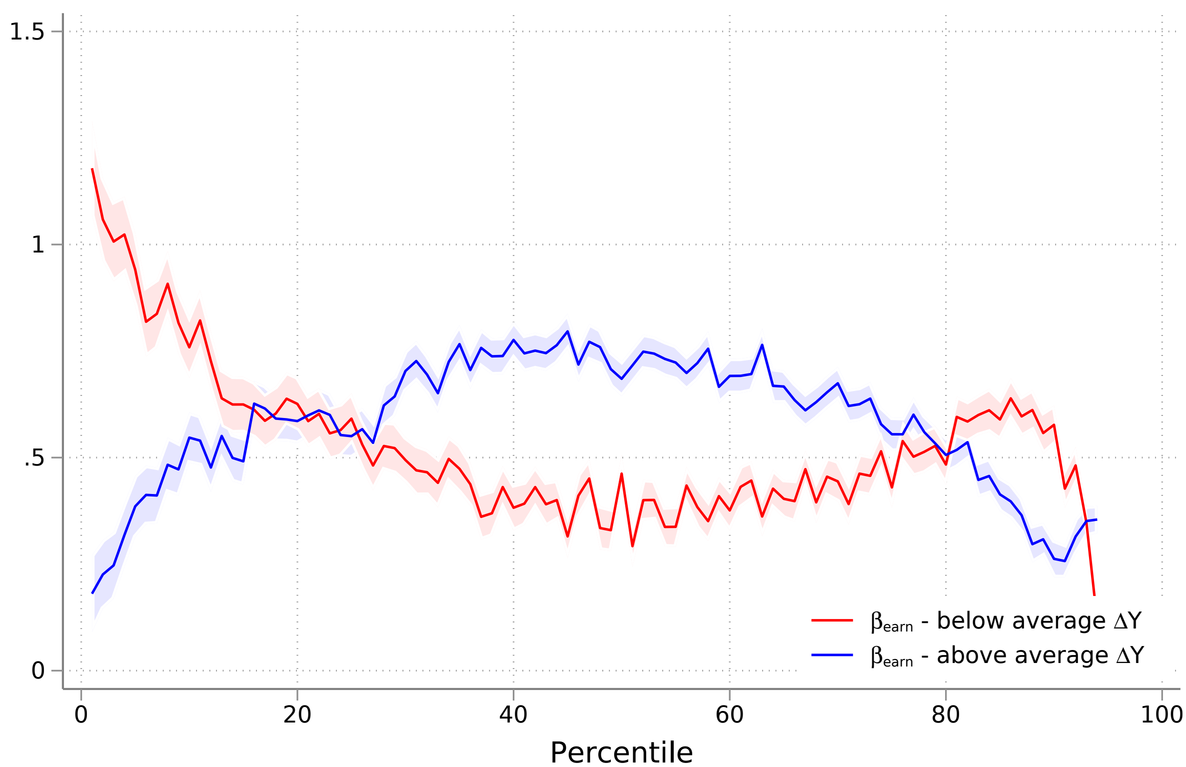


**Note:** The Figure plots the coefficient $\beta^p_{EARN}$ in Equation (2) for percentiles 1 to 95 for all individuals. Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. Income growth is computed as the hyperbolic sine transform of post-benefit income. The regressions are estimated for periods with above and below average aggregate earnings growth. The sample period is from 1980 until 2014. The shaded area indicates 95 percent confidence bands.

[11] For our sample since 1980, there is no authority in Germany which declares periods to be recessionary as the NBER does in the US. Consequently, declaring below average year-on-year quarterly aggregate income growth as such an indicator seems reasonable.

This is important because previous research conducted on US data finds that recessions lead to stronger earnings decreases for the bottom of the income distribution, while barely affecting the top (Guvenen et al., 2015). The figure above indicates that similar results hold for Germany.

## 4.4 Robustness

Our benchmark results so far were based on the definition of permanent income used in Guvenen et al. (2017), based on average earnings during the 5 years prior to the base period $t$. Our findings are robust to changing this definition of permanent income to that of an individual fixed effect in earning. To obtain it, we estimate the following regression

$$earn_{j,t} = \beta \widetilde{X}_{j,t} + \delta_j + \varepsilon_{j,t} \tag{7}$$

where $\widetilde{X}_{j,t}$ is a vector of individual-specific control variables containing an age polynomial, and dummies for education, industry and occupation; and $\delta_j$ is an individual fixed effect. While the previous sorting into percentiles was based on values calculated each quarter, the fixed effect is constant across time for each individual.
The results using percentiles based on the fixed effect described above are reported in Figure 11. The $\beta^p_{EARN}$ coefficients estimated across the permanent income distribution are very similar in size to those reported in Figure 4. Regardless of the construction of the proxy for permanent income, those at the bottom of the permanent income distribution are more exposed to aggregate income movements than those at the top.

Figure 11: Regression coefficients $\beta^p_{EARN}$ - alternative measures of permanent income

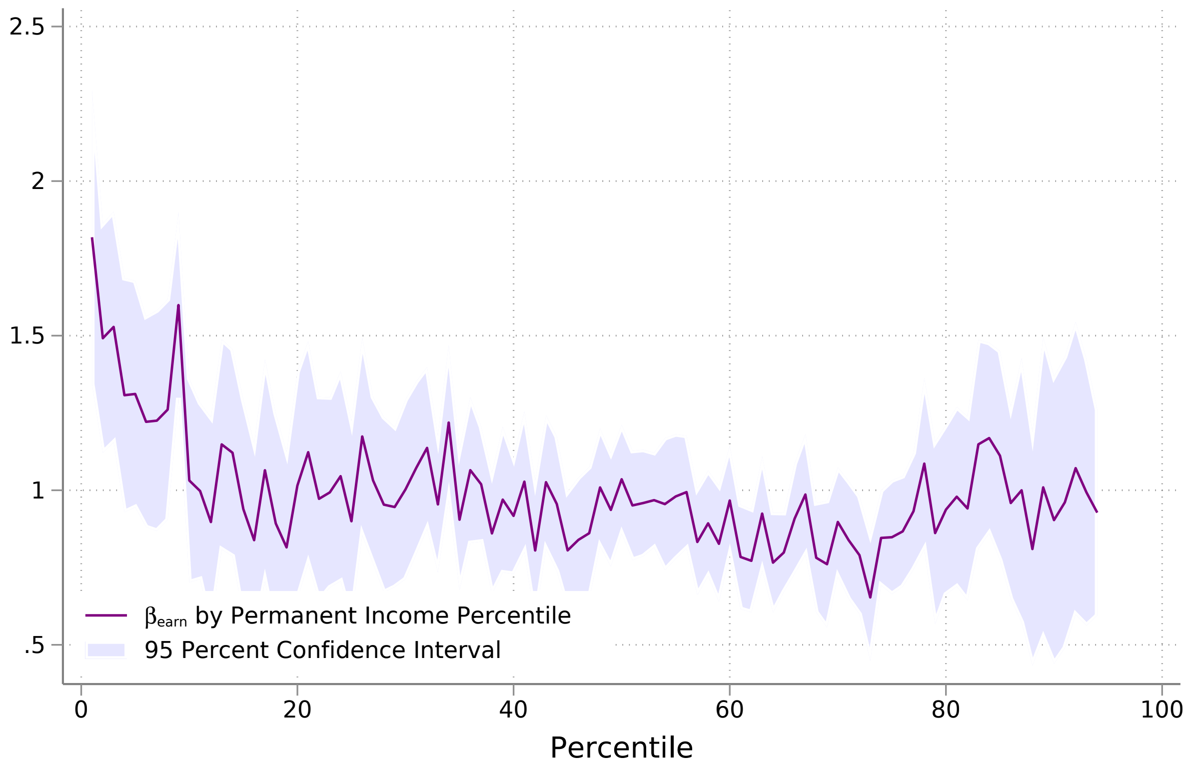


**Note:** This figure plots the coefficient $\beta^p_{EARN}$ in Equation (2) for percentiles 1 to 95. The sample period is from 1980 until 2014. The shaded area indicates 95 percent confidence bands. Individuals are sorted into percentiles fixed effect $\delta_j$ from the regression in Equation (7).

Our analysis abstracts from additional sources of heterogeneity, for example along the life-cycle dimension, partly motivated by the low correlation between age and income in the sample in Tables 1a and 1b. Those tables did, however, show a strong decline in the share of female individuals along the income distribution. Does this affect our results? Figure 12 shows the coefficients $\beta^p_{EARN}$ estimated on the subsample of men only. The patterns are very similar to the benchmark estimates in Figures 4 and 7, but the decline in comovement along the distribution is even more pronounced. In particular, the regression coefficient $\beta^p_{EARN}$ is about a third higher for men at the bottom of the distribution than in the sample of both sexes in Figure 4. There is, however, one exception: incomes in the bottom quintile of the full sample of men (in the left panel of Figure12), where men account for only about a third of the population, is about as procyclical as in the full sample with both sexes, leading to a hump shape in $\beta^p_{EARN}$ at the bottom of the distribution.

Figure 12: Regression coefficients $\beta^p_{EARN}$ - Only men

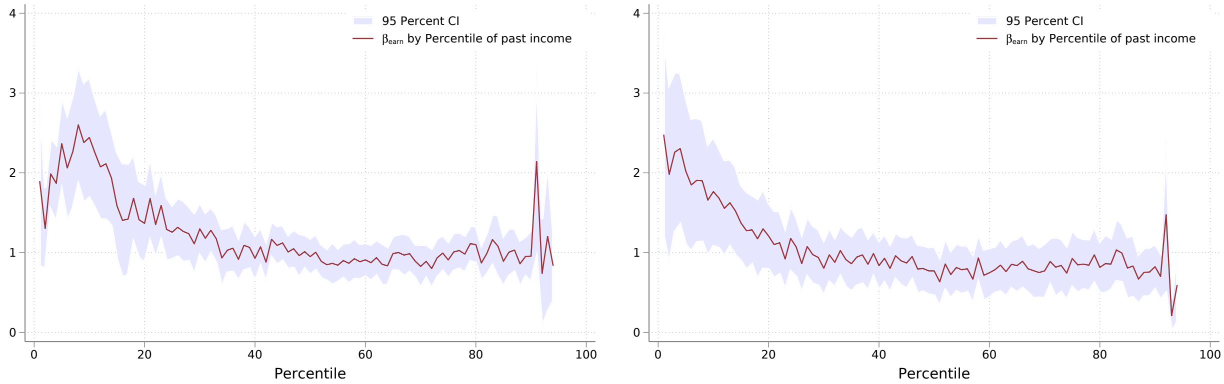


a) Full sample

b) Fully attached

**Note:** The *Left Panel* plots the coefficient $\beta^p_{EARN}$ in Equation (2) for percentiles 1 to 95 for the group of men in the full sample. Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t$. Income growth is computed as difference in the hyperbolic sine transform of post-benefit income between quarters $t$ and $t+4$. The sample period is from 1980 until 2014. The shaded area indicates 95 percent confidence bands. The *Right Panel* plots the coefficients $\beta^p_{gdp}$ across percentiles for men that were fully attached in period $t$.

# 5 Impulse responses to a monetary policy surprise

The previous section studied unconditional correlations between aggregate and individual earnings. This section investigates the effects of a particular shock, to monetary policy, on individual earnings in different parts of the income distribution.

We focus on the period since January 1999, when European monetary policy has been conducted by the ECB. Since the German economy accounts for roughly a quarter of Euro area GDP, however, it is likely that ECB monetary policy is heavily influenced by German economic performance. Hence, when estimating the impact of interest rate changes on the German economy, endogeneity is an important concern.
To identify monetary policy surprises, we use high-frequency data on Overnight Index Swap (OIS) rates as an instrument, as in (Jarocinski and Karadi, 2018). In an OIS agreement, two parties exchange a floating overnight interest rate, in our case the European Overnight Index Average (EONIA), for a fixed rate over a prespecified period of time and on a notional principal. At the end of the swap-period, the contract is cash settled. Furthermore, as the contracts are highly collateralized, counterparty risk is minimal.[12]
Every six weeks, on Thursdays, the ECB Governing Council meets to decide on monetary policy actions. On such days, at 13:45 CET, a press release is posted which concisely summarizes the decisions taken by the Governing council. Subsequently, at 14:30 CET, the president of the ECB holds a press conference, first motivating the decisions taken in an introductory statement and later taking questions from the audience. Our instrument

[12]For a more detailed description of OIS swaps and their similarities to Federal Funds Futures, which have been used to identify monetary policy surprises in the US (Gertler and Karadi, 2015), see (Lloyd, 2018).

measures the change in 3 month EONIA OIS rates in response to these two events in a narrow time window around them. We calculate the average rates in windows 15 minutes before and 30 minutes after the press release and the press conference. We take the difference between the pre-and post- window in each case and sum the two. If this measure is large, in absolute terms, we conclude that the decisions taken by the ECB Governing Council were not expected by financial markets.

## 5.1 Aggregate responses to monetary policy surprises

To test the validity of our monetary policy shocks, we run the following local projection regression:

$$x_{t+h} - x_{t-1} = \alpha + \beta_h \Delta i_t + +\gamma_h X_{t-1} + \varepsilon_{t,h} \tag{8}$$

where $x$ represents (i) the inflation rate as measured by the logarithm of German HICP, (ii) the logarithm of industrial production and (iii) the German unemployment rate. The vector $X_{t-1}$ represents a set of control variables consisting of one lag of the instrument, $\Delta i_t$, the inflation rate and $x$.

Figure 13 shows the impulse responses to a 100 Basis Point shock to the interest rate, estimated using Equation (8). The horizontal axis measures time after the monetary policy surprise in months, the vertical axis measures the percentage point change in the aggregate in question. The top left graph indicates that the inflation rate does not strongly react to the surprise in either direction. Although it initially increases significantly, most point estimates going forward are insignificant. The response of Industrial production is reported in the top right graph. According to the textbook theory of monetary policy, production should contract following a monetary tightening. The graph indicates that this is the case, however very imprecisely estimated.

Lastly, the bottom panel in Figure 13 shows the estimated response of the unemployment rate following a surprising tightening of monetary policy. Again, the response is very imprecisely estimated but points in the expected direction, with the unemployment rate increasing, and the response reaching its maximum after around 16 months.

## 5.2 Earnings after a monetary policy surprise

We use our measure of unexpected interest rate changes as an instrument to estimate the following regression, separately for every *decile* of the individual earnings distribution (as the computationally intensive calculation of the standard errors makes estimation with too many quantiles infeasible):

$$w^d_{j,t+h} - w^d_{j,t-1} = \alpha + \beta^d_h \Delta i_t + \varepsilon_{j,h} \tag{9}$$

where $w^d_{j,t}$ represents the hyperbolic sine of deflated post-benefit earnings for individual $j$ in decile $d$ at time $t$. To account for seasonality, we control for calendar-month fixed effects. The left-hand side of the regression measures the percentage change in daily earnings, deflated using the Harmonized Index for Consumer Prices for Germany[13], between the month before the monetary policy surprise $t-1$ and period $t+h$.

[13]Obtained from Eurostat, series prc_hicp_midx.

Figure 13: Aggregate responses to monetary policy surprises

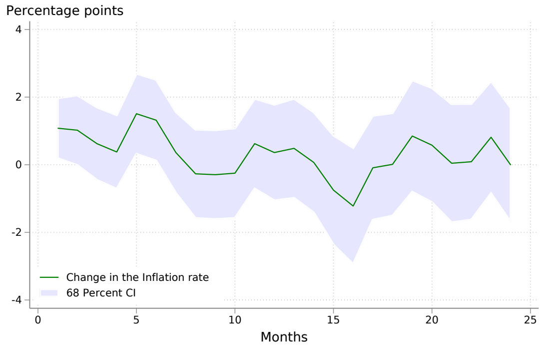


(a) Impulse response of the inflation rate

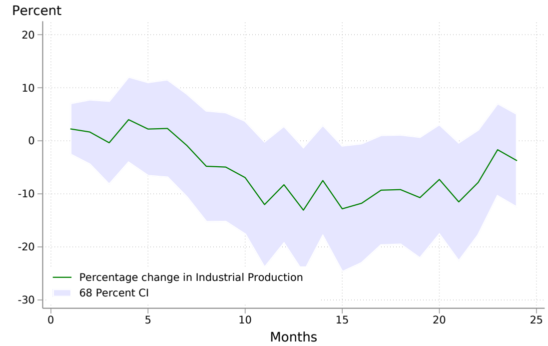


(b) Impulse response of Industrial Production

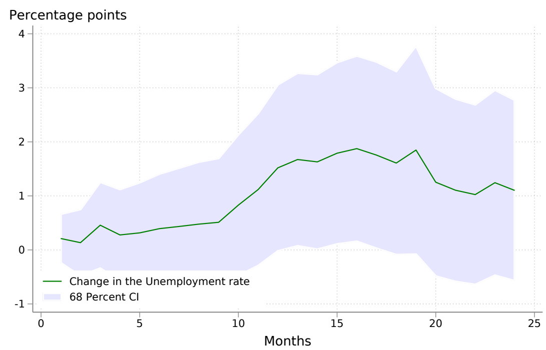


(c) Impulse response of the unemployment rate

**Note:** The Figure shows the impulse responses of aggregate variables to a 100 BP monetary policy surprise, estimated using the Local Projection outlined in Equation (8). The *Top Left Panel* shows the change in the inflation rate, calculated as the change in the logarithm of the HICP for Germany. The *Top Right Panel* shows the percentage change in industrial production, calculated as the log difference, and the *Bottom Panel* shows the change in the unemployment rate. The sample period is from 2000 until 2014. The shaded areas indicate 68 percent confidence intervals.

The left panel in Figure 15 plots the resulting point estimates of the impulse responses at $h = 12$ for a contractionary monetary policy surprise which results in a 100BP reduction in the policy rate. Since the maximum length of an employment spell is twelve months, and we measure the percentage change in average earnings between $t-1$ and $t+12$, the earnings observation must come from two different spells.

Individuals in the bottom deciles experience much larger negative changes in their real earnings than those in the top-deciles in response to a contraction in monetary policy. All estimates except the coefficient for the top decile are significantly different from zero.

Figure 15: Wage responses to monetary policy surprises

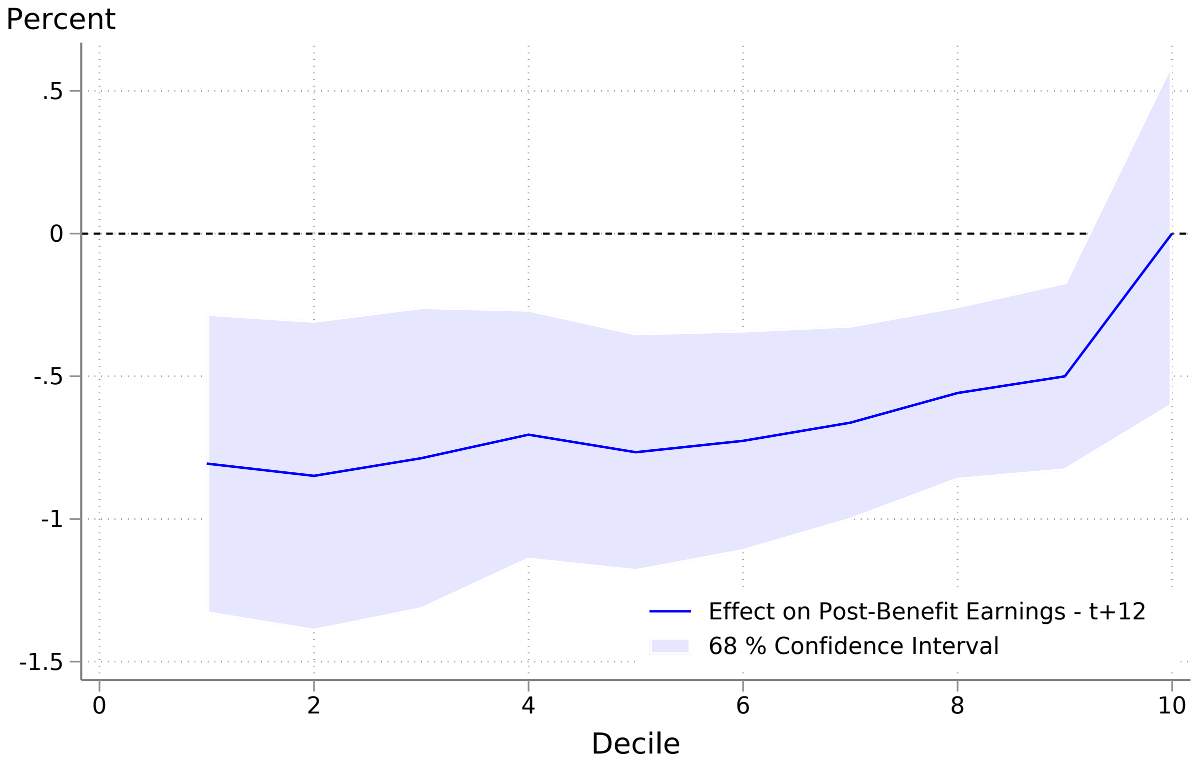


**Note:** The Figure plots the coefficients $\beta_{12}$ in Equation (9) across deciles. Deciles are computed based on a five year average of earnings prior to the monetary policy surprise. The shaded area indicates Discroll-Kraay 68 percent confidence bands. The sample is restricted to individuals who are employed in the period before the surprise. The sample period is from 2000 until 2014.

## 5.3 Labor market transitions after a monetary policy surprise

As before, we can also investigate the responses of transition probabilities to monetary policy surprises. Again, we instrument the following regression, separately for every *decile* of the permanent earnings distribution:

$$emp^d_{j,t+h} = \alpha + \beta_h \Delta i_t + \varepsilon_{j,h} \tag{10}$$

where $emp^d_{t+h}$ is a binary variable measuring the employment status of individual $j$ at time $t+h$. We also control for calendar month fixed effects. The coefficient of interest is $\beta_h$, which measures the impact of a monetary policy surprise at time $t$ on employment $h$ months after. As before, deciles are based on a five-year average of past earnings before the surprise. Using Equation (10), we can compute the change in the probability of (i) remaining employed or (ii) finding a job by conditioning on employment or non-employed at $t-1$, respectively.
The left panel in Figure 16 plots the change in the probability of remaining employed, one year ($h = 12$) after a contractionary monetary policy shock that leads to a 100BP reduction in

the policy rate.[14] In the first permanent income decile, such a shock leads to a 10% reduction in the probability of remaining employed, while the top of the distribution is hardly affected.

Figure 16: Labor market transitions in response to monetary policy surprises

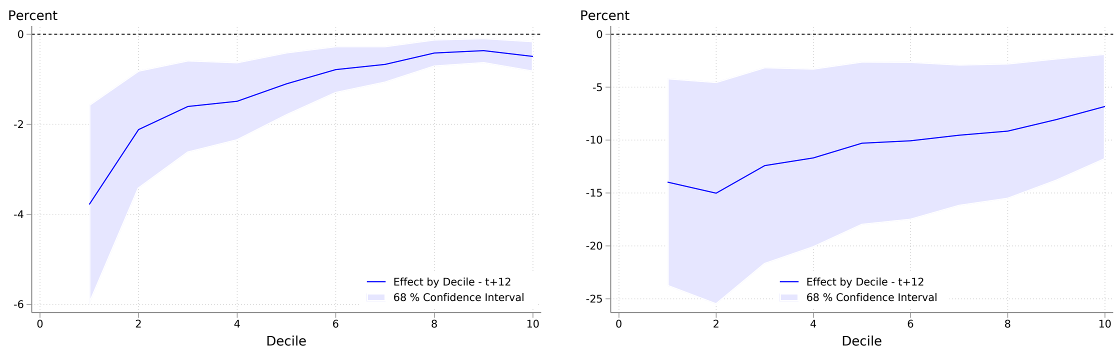


(a) 12 month responses – Employed (b) 12 month responses – Non-Employed

**Note:** The *Left Panel* shows the change in the probability of staying employed 12 months after a contractionary monetary policy shock, conditional on being employed in the period before the shock, by decile. The *Right Panel* shows the change in the probability of finding employment 12 months after a contractionary monetary policy shock, conditional on being non-employed in the period before the shock, by decile. The shaded areas represent 68 percent Discroll-Kraay confidence intervals. The sample period is from 2000 until 2014.

The right panel of figure 16 shows how the transition from non-employment to employment is affected by a monetary policy surprise. The probability of finding employment declines across deciles, with the bottom of the distribution entirely unaffected. The magnitudes are striking: individuals with high permanent income see their job finding probablity fall by 25%. Furthermore, it is unintuitive to see the top deciles most affected, while these individuals were almost unaffected in all previously reported results.
To provide this intuition, we plot the impulse responses relative to their steady state values in Figure 17. The steady state is calculated as the average of the calendar month dummies in Equation (10). The left panel shows that the probability of remaining employed is steeply increasing with permanent income: individuals in the first decile who are employed in month $t$ have less than an 80% chance of being employed 12 months later. The same probability is higher than 95% for the tenth decile. A monetary policy shock steepens the slope even further.
The pattern is similar for the job-finding probability. It, too is steeply increasing across the permanent income distribution. Non-employed individuals in the tenth decile are four times more likely to be employed in period $t + 12$ than those in the first. However, a contractionary monetary policy shock leads to a significant flattening of this slope. Hence, the unintuitive finding that the job-finding probability is most affected for the top of the distribution is explained by the fact that in the first deciles, the probability is already very low, implying that it cannot fall much further.[15]

[14] We report the full impulse responses in the appendix.

[15] To properly tackle the issue of probabilities, we repeat the estimation using a probit regression reported in the appendix. The results are very similar.

Figure 17: Transition to employment

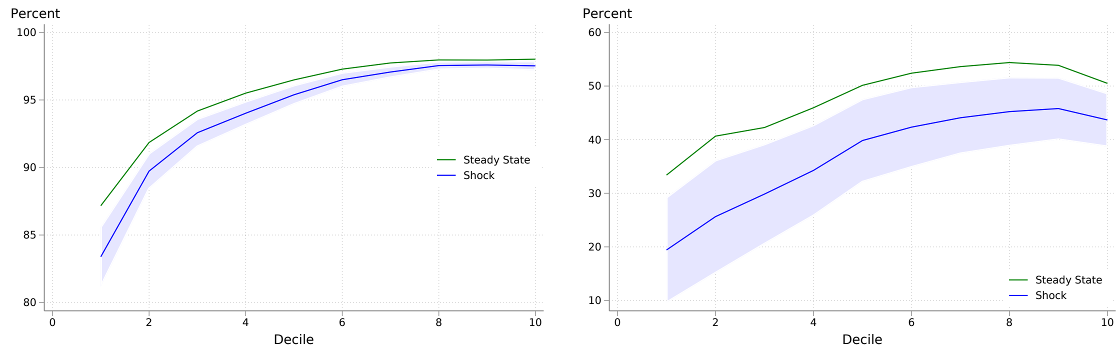


(a) 12 month responses – Employed (b) 12 month responses – Non-Employed

**Note:** The *Left Panel* shows the steady state probability of remaining employed in $t+12$, conditional on being employed at $t-1$, in green. Furthermore, it shows the conditional probability of remaining employed after a monetary policy shock. The *Right Panel* shows the steady state probability of being employed in $t+12$, conditional on being non-employed at $t-1$, in green. Furthermore, is shows the conditional probability of finding a job after a contractionary monetary policy surprise. The shaded areas represent 68 percent Discroll-Kraay confidence intervals. The sample period is from 2000 until 2014.

# 6 Inspecting the mechanism: Transmission of monetary policy

In previous sections, we have provided evidence for (i) aggregate earnings growth being correlated more or less with individual earnings growth across the permanent income distribution and (ii) monetary policy affecting income growth differentionally across said distrubtion. This poses the natural question whether the two are related to each other.

An important question in the literature on monetary policy transmission is whether monetary policy itself affects only aggregate GDP, which then has effects on individuals (indirect effect), or whether there is an effect through which individuals are affected directly. We can take these questions to our dataset and estimate the following regression to find tentative answers:

$$\Delta earn^p_{j,t} = \alpha_1 + \beta^p_{earn}\Delta earn_{agg} + \epsilon_{j,t} \tag{11}$$

where we instrument aggregate earnings growth $\Delta earn_{agg}$ with the instrument presented above. Intuitively, this captures the effect of aggregate income growth on individual income growth *due to* monetary policy. Naturally, our approach cannot speak to direct effects which might result from financial positions an individual holds (see e.g. Auclert, 2019), unless these affect their labor market outcomes.

The left panel in Figure 18 reports the results from the OLS estimation (blue) and the IV (red). The former is equivalent to the results reported in Figure 3, with percentiles collapsed to deciles. We find no significant difference between the individual/aggregate correlation, whether the growth in the aggregate is driven by monetary policy or not.

Similarly, the right panel reports results from an analogous regression, substituting with the employment state as the dependent variable. Again, we find no significant difference between

Figure 18: Non-employment to employment

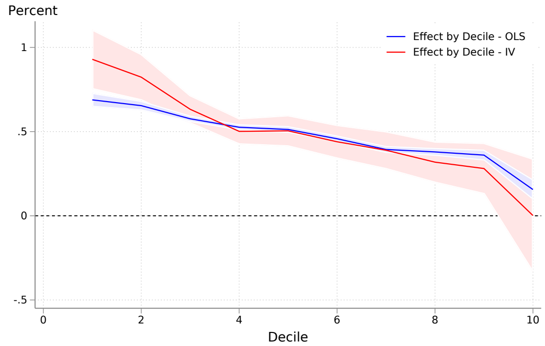


(a) 12 month responses – Wages

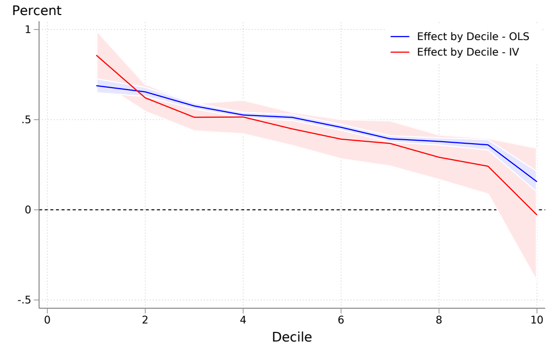


(b) 12 month responses – Wages FA

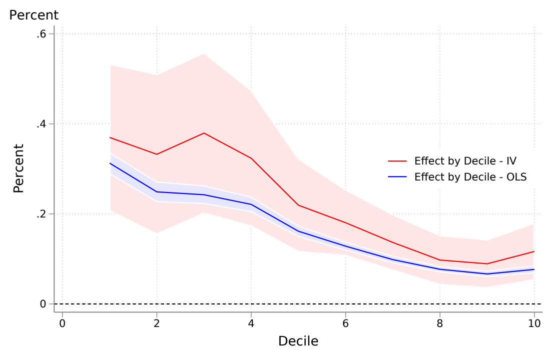


(c) 12 month responses – Employment

**Note:** This figure plots $\beta_{12}$ in Equation (2) across deciles. Deciles are computed based on a five year average of earnings prior to the shock. The standard errors are Discroll-Kraay.

the OLS and IV estimates. These estimates imply that while it is conceivable that monetary policy directly affects individuals, we find no evidence for a direct effect on labor earnings.

# 7 Conclusion

Our answers to the opening questions of our paper are as follows. In Germany, income risk from job loss and low job-finding rates, as well as the effect of business cycles on individual income growth, all decline strongly along the income distribution. Importantly we find similar patterns also in response to monetary policy surprises, which affect income growth as well as separation and job finding rates substantially more at the bottom of the income distribution. Interestingly, we found that the sources of this comovement also changed along the distribution: fluctuations in extensive risk, due to transitions between labor market states that are associated with different income levels, dominate for the poor. But in line with the rising share of job-stayers, the role of extensive risk falls, and that of intensive risk (in particular from procyclical income growth in continued employment relationships) rises along the distribution, with the two accounting for about equal shares at the top.

Our analysis focused on heterogeneity in individual life-time resources, as summarized by our measures of permanent income. The strong heterogeneity we find along its distribution, together with the small correlation of incomes with age and the similar results we document for men, make us confident that this is a particular relevant dimension to study. It would however be fruitful to look at additional heterogeneity in future work, for example along the life-cycle dimension, potentially interacted with gender differences.

Our results have implications for policy and research alike. To monetary policy makers, our message is that their actions may have substantial effects on economic inequality because interest rate changes affect incomes and employment prospects at the bottom of the income distribution substantially more than at the top. For those concerned with economic inequality more generally, our results highlight the dominant role of cyclical separations and job finding at the bottom of the distribution, contrasted with more important wage and salary fluctuations at the top. An interesting open question is the role of German labor market institutions for these patterns, such as the publicly funded short time work (*Kurzarbeit*) when firms face adverse conditions. Institutions themselves may have different effects along the distribution. For example, one may expect short time work to be more relevant for fully attached workers than for those in less standard employment relationships.

Finally, we hope that the stylised facts we document provide some valuable inputs for future research. For example, following Guvenen et al. (2015), many studies have investigated the role of countercyclical skewness and excess kurtosis of income innovations for consumption dynamics, asset prices, etc. Our work suggests that the link that some of these features have with employment dynamics, and their heterogeneity across the income distribution, should be incorporated explicitly in such analyses. Finally, the literature on heterogeneous-agent New Keynesian models has highlighted the potential role of heterogeneous incidence of income risk for transmission of shocks via aggregate demand when marginal propensities to consume are also heterogeneous. Our results suggest that studying heterogeneous incidence in particular of employment risk in such a framework may provide important insights. We aim to carry out such an analysis in the future.

# References


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

Böhm, M. J., H.-M. v. Gaudecker, and F. Schran (2019). Occupation, growth, skill prices, and wage inequality.

Drews, N., D. Groll, and P. Jacobebbinghaus (2007). Programmierbeispiele zur aufbereitung von fdz personendaten in stata.

Fitzenberger, B., A. Osikominu, and R. Völter (2005). Imputation rules to improve the education variable in the iab employment subsample. *ZEW-Centre for European Economic Research Discussion Paper* (05-010).

Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics 7*(1), 44–76.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2015). What do data on millions of us workers reveal about life-cycle earnings risk? Technical report, National Bureau of Economic Research.

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

Jarocinski, M. and P. Karadi (2018). Deconstructing monetary policy surprises: the role of information shocks.

Lloyd, S. (2018). Overnight index swap market-based measures of monetary policy expectations.

Patterson, C. et al. (2019). The matching multiplier and the amplification of recessions. In *2019 Meeting Papers*, Number 95. Society for Economic Dynamics.

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

# Appendix

# A Additional Results

## A.1 Unemployment rate

The left panel of Figure A.1 compares the unemployment rate for Germany resulting from our first definition of employment to the official rate reported by the German statistics office, computed using survey data. Importantly, we include only individuals whose place of work lies in the counties that were part of the Federal Republic of Germany before 1990.[16] The two rates move closely together, especially before the reunification in 1990. After 1995, the narrow unemployment rate is, however, systematically lower than the officially reported one. The right panel of Figure A.1 shows that the unemployment rates calculated using different definitions of employment behave very similarly.

## A.2 Income and employment risk in the whole sample

In the analysis above, we group together unemployment (U) and nonparticipation (N). However, our results are similar if we keep flows disaggregated, as Figure 20 shows.

[16] For the non-employed, location information is not available. We use the last employer's location.

Figure 19: German unemployment rates

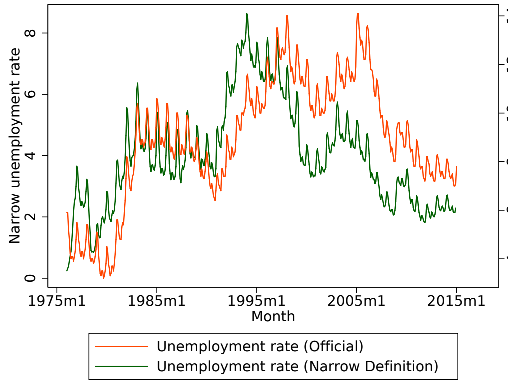


(a) Official rate and narrow definition (only west)


(b) Narrow unemployment and fully attached (only west)

Figure 20: Regression coefficients $\beta^p_{EARN}$ for more transitions - full sample

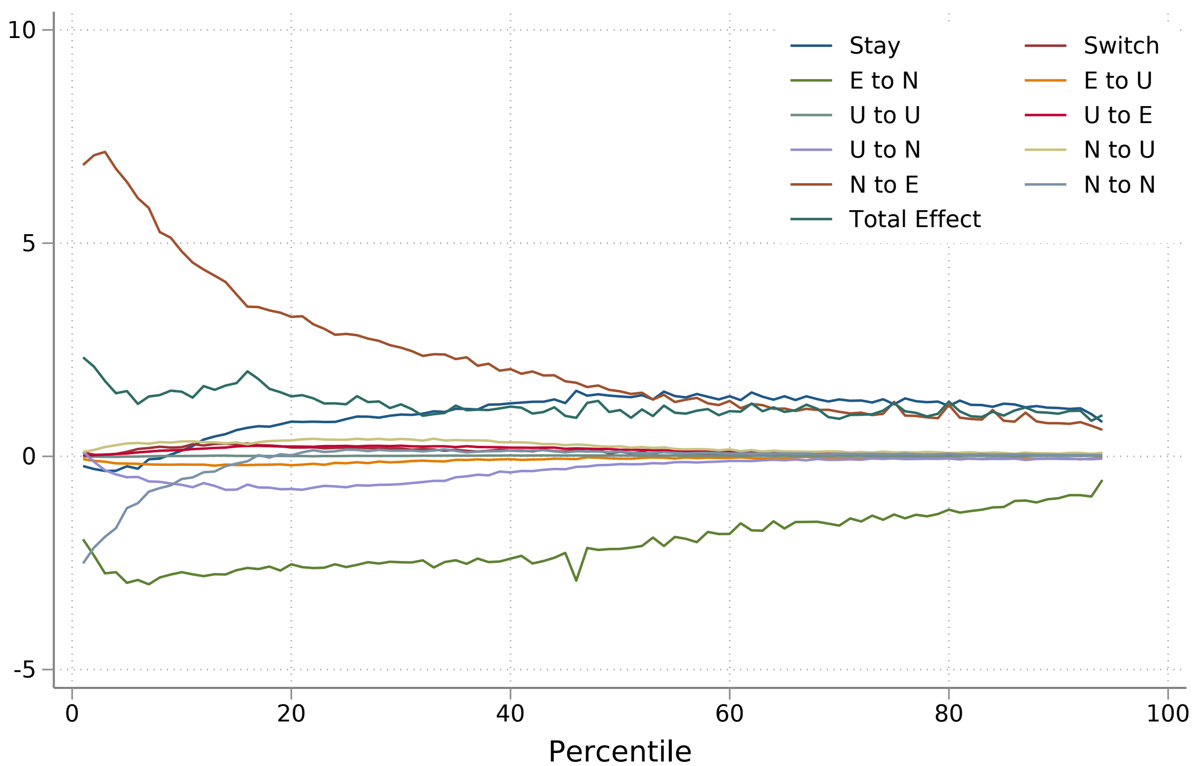


**Note:** The *Left Panel* plots the coefficient $\beta^p_{EARN}$ in Equation (2) for percentiles 1 to 95 for individuals who are categorized as fully attached to the labor force in period $t-4$ and as employed in period $t$. Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t-4$. Income growth is computed as difference in the hyperbolic sine transform of post-benefit income between quarters $t-4$ and $t$. The sample period is from 1980 until 2014. The shaded area indicates 95 percent confidence bands. The *Right Panel* plots the coefficients $\beta^p_{gdp}$ across percentiles with the same sampling restrictions.

We repeat the analysis from Figure 7 using longer time periods. The regression analysis is the same as before, but using longer time periods to calculate growth rates. Figure 21 reports the correlation between average 2 year (left) and 5 year (right) individual and aggregate income growth. The results are qualitatively very similar to the ones reported before, although the correlations are quantitatively smaller.

Figure 21: Regression coefficients $\beta^p_{EARN}$ with longer-term growth rates - full sample


**Note:** The *Left Panel* plots the coefficient $\beta^p_{EARN}$ in Equation (2) for percentiles 1 to 95 for the whole sample when growth rates of individual and aggregate incomes are calculated over eight quarters ($t-8$ to $t$). Individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t-8$. The *Right Panel* plots the coefficients $\beta^p_{EARN}$ when growth rates of individual and aggregate incomes are calculated over 20 quarters ($t-20$ to $t$) and individuals are sorted into percentiles by their average income throughout the five years prior to quarter $t-20$. Income growth is computed as difference in the hyperbolic sine transform of post-benefit income between quarters $t-4$ and $t$.

## A.3 Full impulse responses

In the main text, we only show results for the 12th month after a monetary policy shock, displaying the responses at this point across deciles. Here, we report the full impulse responses. The left panel of 22 shows the change in the probability of remaining employed after a monetary policy shock; only for five deciles to not clutter the graph too much. The right panel shows the probability of finding employment, given non-employment at $t-1$.

Figure 22: Full Impulse Responses

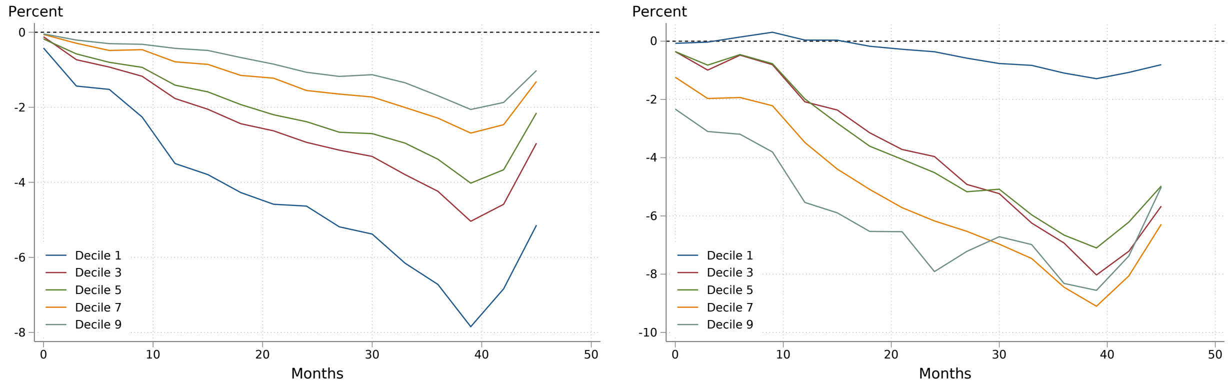


(a) Probability of remaining employed (b) Probability of finding employment

**Note:** The *Left Panel* shows the change in the probability of staying employed $t$ months after a contractionary monetary policy shock, for $t = 1, 2, ..., 50$ conditional on being employed in the period before the shock, for deciles $1, 3, 5, 7$ and 9. The *Right Panel* shows the change in the probability of finding employment 12 months after a contractionary monetary policy shock, conditional on being non-employed in the period before the shock, by decile. The shaded areas represent 68 percent Discroll-Kraay confidence intervals. The sample period is from 2000 until 2014.

## A.4 Probit estimation of transition probabilities

In the main body, we present results estimated using linear regression analysis. However, since Equation (10) features a binary dependent variable, and we are dealing with transition *probabilities*, we repeat the analysis presented above using a probit regression. Figure 23 shows that the results are almost equivalent to the linear model used before.

## A.5 Hiring and firing around monetary policy announcement dates

Because we have daily data on the beginning and end of each employment spell, we are able to conduct 'high-frequency' analysis around monetary policy meetings. One hypothesis is that labor market decisions are made immediately after monetary policy announcements, leading to spikes in hirings or firings in the data.
To investigate this further, we estimate the following regression

$$sep_t = \sum_{i=-10}^{10} I_i + \varepsilon_t \tag{12}$$

and an analogous regression for hirings. Because there are strong beginning- and end-of-month and -year effects, we include calendar-day dummies as well as dummies for the first and last days of the year.
Figure 24 presents the average difference between the number of hirings (left panel) and separations (right panel) around monetary policy announcements, relative to days further

Figure 23: Non-employment to employment

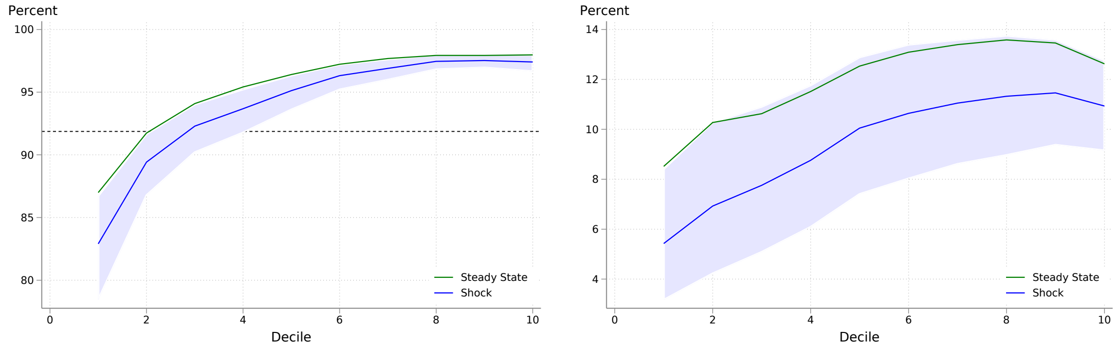


(a) 12 month responses – employed (b) 12 month responses – non-employed

**Note:** *Left Panel:* The figure shows the steady state probability of remaining employed in $t+12$, conditional on being employed at $t$. Furthermore, is shows the conditional probability of remaining employed after a monetary policy shock using a probit specification. *Right Panel:* The figure shows the steady state probability of being employed in $t+12$, conditional on being non-employed at $t$. Furthermore, is shows the conditional probability of finding a job after a monetary policy shock using a probit specification. The standard errors are clustered onat the month level, but **not** using the method developed by Discroll and Kraay. The deciles are computed using five year average earnings before the surprise.

away from such meetings. We find that there are more separations before meetings, with no significant effect afterwards. The effect on hirings appears to be ambiguous.

Figure 24: The labor market around monetary policy announcements

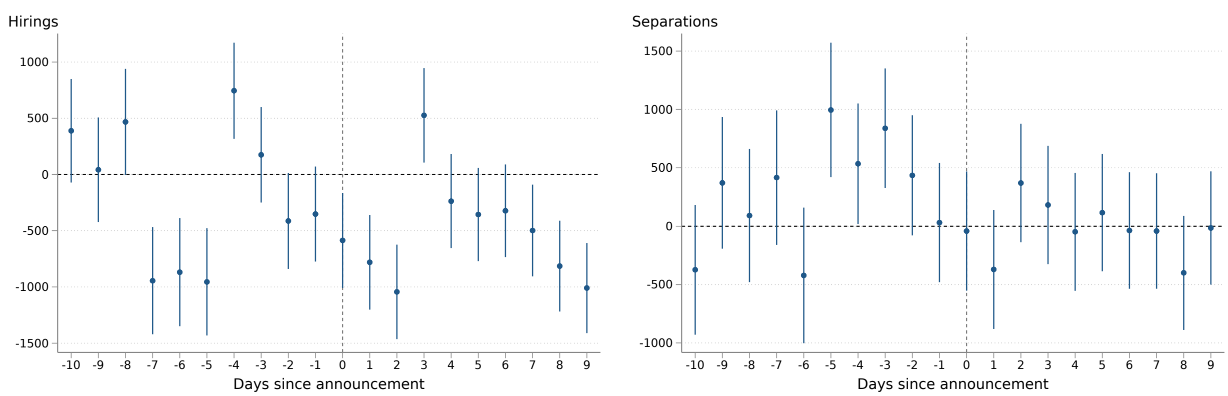

Figure 25: Contribution by flow – Full sample (left) and Fully attached (right)

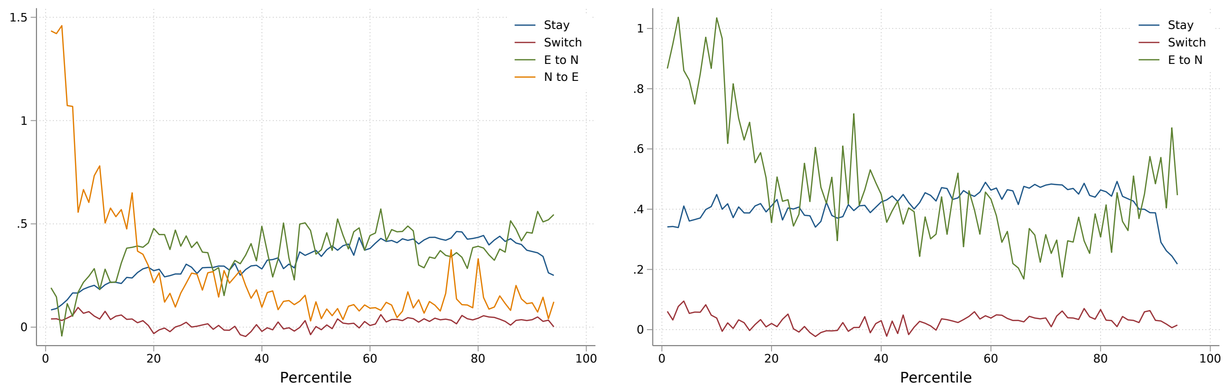


Figure 26: Extensive vs Intensive margin by flow – Full sample

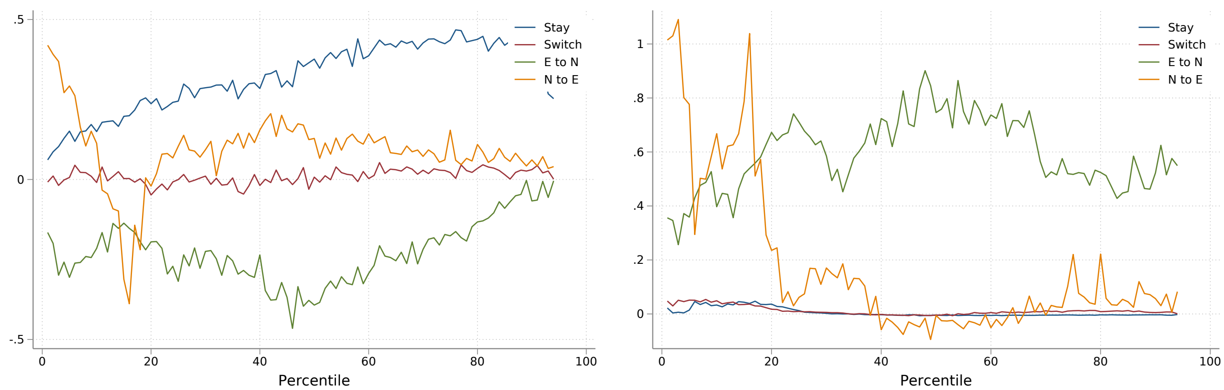


Figure 27: Extensive vs Intensive margin by flow – Fully attached

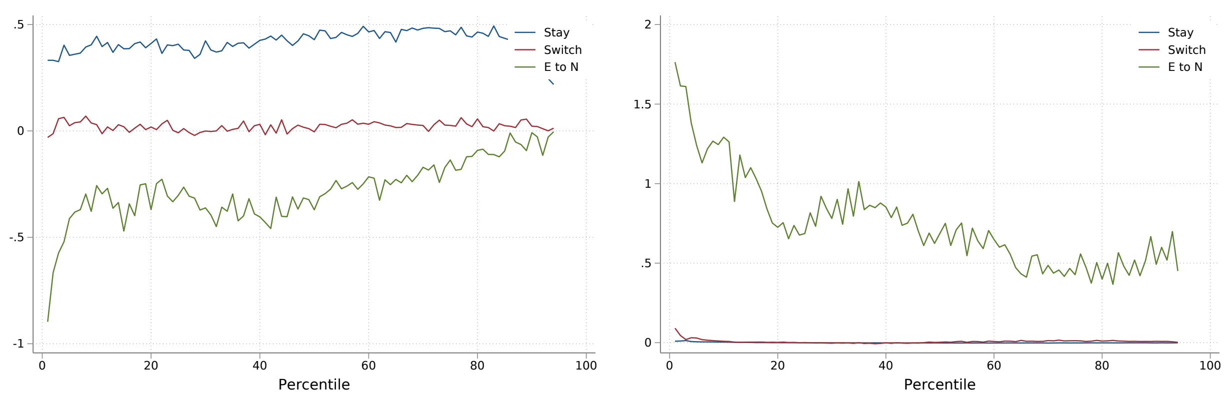

Figure 28: Comovement between individual and aggregate income – Only men

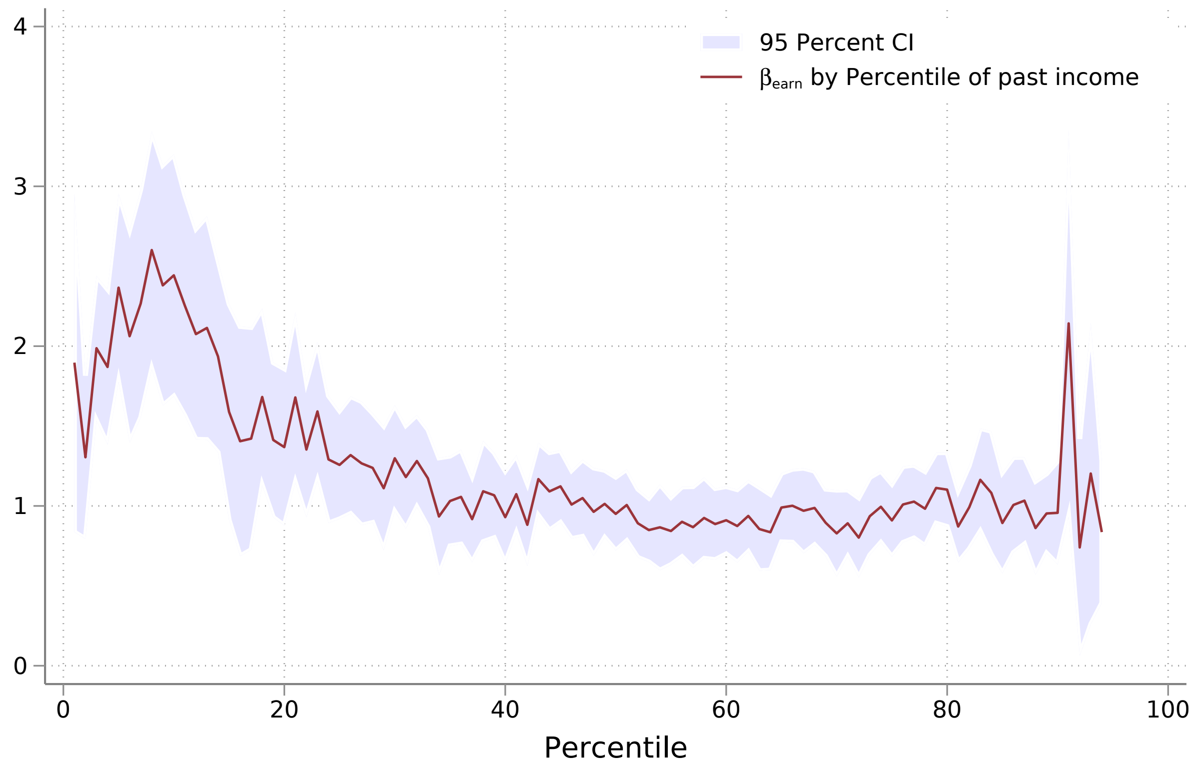


## A.6 Decomposition of earnings betas

This section shows how to decompose the overall regression coefficient $\beta_Y^p$ into a weighted average of $\beta_{TR}^p$ in (3) and coefficients $\beta_Y^p$ estimated for individual who share the same labor market transition. Specifically, we decompose the covariance between individual earnings growth (denoted $dy_{it}$ here to save space) and growth in aggregate activity ($dY_t$). A bar over a variable denotes averages, and variables without a subscript are averaged across the subscript dimension (so, e.g., $\bar{dy} = \sum_t \sum_{i=1}^{N_t} dy_{it}$). Index $i$ represents individuals; $j$ represents groups that share labor market transitions; $N_t$ and $N_{jt}$ denote, respectively, numbers of observations in period $t$ in total and in group j; $N_t\ TOT = \sum_{i,t} 1 = \sum_t N_t = \sum_t \sum_j N_{jt}$ is the total number of observations; $\bar{N} \equiv \frac{TOT}{T}$ and $\mu_j = \frac{\bar{N}_j}{\bar{N}}$.

$$\begin{aligned}
\text{Cov}(dy_{it}, dY_t) =& \frac{1}{TOT} \sum_t \sum_{i=1}^{N_t} (dy_{it} - \bar{d}y)(dY_t - \bar{d}Y) && (13)\\
=& \frac{1}{TOT} \sum_t \sum_{j=1}^{J} \sum_{i=1}^{N_{jt}} (dy_{ijt} - \bar{d}y_{jt} + \bar{d}y_{jt} - \bar{d}y_j + \bar{d}y_j - \bar{d}y)(dY_t - \bar{d}Y) && (14)\\
=& \frac{1}{TOT} \sum_t \sum_{j=1}^{J} N_{jt} \left[ 0 + (\bar{d}y_{jt} - \bar{d}y_j) + (\bar{d}y_j - \bar{d}y) \right] (dY_t - \bar{d}Y) && (15)\\
=& \frac{1}{TOT} \sum_t \sum_{j=1}^{J} \Big[ \bar{N}_j (\bar{d}y_{jt} - \bar{d}y_j)(dY_t - \bar{d}Y) && (16)\\
& + (N_{jt} - \bar{N}_j)(\bar{d}y_j - \bar{d}y)(dY_t - \bar{d}Y) + (N_{jt} - \bar{N}_j)(\bar{d}y_{jt} - \bar{d}y_j)(dY_t - \bar{d}Y) \Big] && (17)\\
=& \frac{1}{\bar{N}} \sum_{j=1}^{J} \Bigg[ \bar{N}_j \frac{1}{T} \sum_t (\bar{d}y_{jt} - \bar{d}y_j)(dY_t - \bar{d}Y) && (18)\\
& + \frac{1}{T} \sum_t (N_{jt} - \bar{N}_j)(\bar{d}y_j - \bar{d}y)(dY_t - \bar{d}Y) + \frac{1}{T} \sum_t (N_{jt} - \bar{N}_j)(\bar{d}y_{jt} - \bar{d}y_j)(dY_t - \bar{d}Y) \Bigg] && (19)\\
=& \sum_{j=1}^{J} \left[ \bar{\mu}_j \text{Cov}(\bar{d}y_{jt}, dY_t) + (\bar{d}y_j - \bar{d}y) \text{Cov}(\mu_{jt}, dY_t) + \frac{1}{T} \sum_t (\mu_{jt} - \bar{\mu}_j)(\bar{d}y_{jt} - \bar{d}y_j)(dY_t - \bar{d}Y) \right] && (20)\\
\approx& \sum_{j=1}^{J} \left[ \bar{\mu}_j \text{Cov}(\bar{d}y_{jt}, dY_t) + (\bar{d}y_j - \bar{d}y) \text{Cov}(\mu_{jt}, dY_t) \right] && (21)
\end{aligned}$$

This yields

$$\begin{aligned}
\beta_{EARN} &= \frac{\text{Cov}(dy_{it}, dY_t)}{Var(dY_t)} \approx \beta_{int} + \beta_{ext} && (22)\\
\beta_{int} &= \sum_{j=1}^{J} \bar{\mu}_j \beta_{int,j} = \sum_{j=1}^{J} \bar{\mu}_j \frac{\text{Cov}(\bar{d}y_{jt}, dY_t)}{Var(dY_t)} && (23)\\
\beta_{ext} &= \sum_{j=1}^{J} (\bar{d}y_j - \bar{d}y) \beta_{TR,j} = \sum_{j=1}^{J} (\bar{d}y_j - \bar{d}y) \frac{\text{Cov}(\mu_{jt}, dY_t)}{Var(dY_t)} && (24)
\end{aligned}$$