The Curious Incidence of Monetary Policy Shocks Across the Income Distribution*

Tobias Broer† John Kramer‡ Kurt Mitman§

March 14, 2021

Abstract

We use high-frequency administrative data from Germany to study the effect of monetary policy shocks on incomes and employment prospects along the income distribution. We find that income growth at the bottom of the income distribution is substantially more affected by monetary policy shocks. Much of this heterogeneity comes from stronger effects of these shocks on the job-finding and separation rates of the poor.

1 Introduction

Do monetary policy interventions affect poor workers' incomes and employment prospects more than those of the rich? Answering this question is important to assess the welfare effects of monetary policy, and thus for policy design. It is also important for the transmission of monetary policy shocks to aggregate demand, as the consumption of poorer households is likely to react more strongly to fluctuations in their incomes (Patterson et al., 2019), and because substantial heterogeneous responses of labor market risk may change the transmission of monetary policy (Ravn and Sterk, 2017; Werning, 2015).

We use a long panel of detailed administrative data from Germany, containing individual labor market biographies including earnings. The high frequency nature of our data allows us to estimate responses of earnings and transitions in employment status to monetary policy shocks, which we identify using high-frequency changes in Overnight Indexed Swap rates.

*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

We find that monetary policy shocks disproportionately affects the poor. In particular, below the first quintile of the income distribution the effect of a contractionary shock to monetary policy is more than four times larger than above the top quintile.Much of this heterogeneous incidence on individual earnings arises because monetary policy has stronger effects on the labor market prospects of poor workers, who experience a substantially stronger rise in separation risk after a contractionary monetary policy shock. We also document substantial, and strongly heterogeneous, effects of monetary policy shocks on the consequences of unemployment.

Implications of our findings for policy and macro-modelling

[To be added]

Relation to the literature

A large literature empirically investigates the heterogeneous effects business cycles on individual income risk using administrative datasets, see Guvenen et al. (2015, 2017) (US), Halvorsen et al. (2020) (Norway), Hoffmann and Malacrino (2019) (Italy), De Nardi et al. (2019), (Netherlands and US). Our high-frequency dataset allows us to study the effect of monetary policy shocks on earnings and employment transitions, which we find to have significantly different effects than average business cycles.

We contribute more directly a small empirical literature on the effects of monetary policy on inequality (Coibion et al., 2012). Holm et al. (2020) show that contractionary shocks reduced nonfinancial incomes, but most so at the bottom of the liquid asset distribution. We perform our analysis at monthly frequency, and look at the dynamic effects of monetary policy shocks on both earnings and labor market transitions. Moreover, we do this for the largest European economy, Germany. This makes it crucial to identify exogenous changes in interest rates, which we do using high-frequency changes in Overnight Indexed Swap rates.[1]

Our findings contribute to an empirical foundation for the large literature on the effect of aggregate shocks in economies with heterogeneity in wealth and income. Auclert (2019) shows that, in a large family of macroeconomic models[2], the elasticity of individual earnings to aggregate earnings is a crucial statistic in evaluating the effectiveness of monetary policy. Along similar lines, Werning (2015) and Ravn and Sterk (2017), among others[3], point to the importance of cyclical earnings *risk* as a crucial factor that governs the macroeconomy's

[1]see e.g. Gertler and Karadi (2015), Almgren et al. (2019)

[2]see e.g. Kaplan et al. (2018), Bilbiie (2020), Hagedorn et al. (2019)

[3]see e.g. Gornemann et al. (2016), Challe (2020)

response to aggregate shocks. We provide an empirical foundation for these studies, as we estimate both the elasticity of earnings and labor market transition probabilities (i.e. risk) along the distribution.
[To be completed]

The next section presents the data and desribes the structure of income and employment transitions in our sample on average. Section 4 investigates the effects monetary policy shocks. Section 5 concludes.

2 Data

We use administrative social security data for about 1.7 million German individuals from the Sample of Integrated Employment Biographies (provided by the Research Data Center, FDZ). Our data covers the years between 1975 and 2014 (although most of our analysis starts in 1999), and excludes civil servants and self-employed individuals. Each observation in the dataset is a labor-market spell.[4] We convert these spells into monthly employment histories for each individual, resulting in about 70 million person-quarter observations. An appendix describes how we deal with top-coding, and earnings observations below the social-security threshold.

We most often focus on changes in earnings and employment status over a 12-month period between t and $t + 12$. We define as *employed* individuals who receive labor income, as *unemployed* individuals who are not employed and claimed unemployment benefits after the end of their most recent employment spell, and the rest as *non-employed*.

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. Our preferred proxy is average earnings over the five years preceding quarter t as in Guvenen et al. (2017), but we also study alternative definitions below.[5]

To understand how key variables evolve along the distribution of permanent incomes (henceforth simply the "income distribution"), Table ?? reports descriptive statistics within deciles of our permanent income measure in January 2000, the starting years of our main analysis below.[6]

[4] 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).

[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] 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).

Table 1: Descriptive statistics by decile, first quarter 2000

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	mean	mean	mean	mean	mean	mean	mean	mean	mean	mean
Female	0.64	0.64	0.67	0.63	0.51	0.39	0.31	0.26	0.23	0.12
Age	38.65	39.68	41.10	40.80	39.78	39.89	41.02	42.25	43.08	45.63
Education	2.25	2.30	2.28	2.30	2.29	2.24	2.26	2.43	2.67	3.54
Quarterly earn.	1391.40	2242.45	3150.31	4105.40	5209.87	6216.76	7187.36	8291.95	9740.35	12195.03
Empl next year	0.31	0.48	0.62	0.70	0.76	0.83	0.87	0.90	0.90	0.92
Fully Attached	0.38	0.56	0.71	0.79	0.85	0.91	0.94	0.95	0.95	0.97
Non-employed	0.56	0.38	0.25	0.19	0.13	0.08	0.05	0.04	0.03	0.02
Observations	45282	45282	45282	45282	45282	45282	45282	45283	45281	45282

Note: The tables show average values for selected variables across deciles of the permanent income distribution in the first quarters of 1980 and 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. We impute education following imputation procedure 1 in Fitzenberger et al. (2005) Quarterly earnings are computed as average nominal daily earnings (see text) multiplied by 90. Fully attached individuals are those liable to social security without special characteristics.

The gradient of nominal earnings across the distribution is substantial, with average earnings in the top decile about 9 times higher than in the first. Employment rates are less than 50 percent in the bottom decile, but rise steeply across the bottom half of the distribution and average around 95 percent above the median. Job-finding rates (as defined by 12-month transitions of those without employment into employment) are below one third in the bottom decile, and follow a similar pattern across the distribution. Education levels are surprisingly similar within the bottom three quartiles, but substantially higher in the top quartile. 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. Finally, the gradient of gender composition is substantial (with only 12 percent of women in the top permanent income decile in 2000).

3 Earnings responses to monetary policy surprises

This section investigates the effects of the monetary policy shocks on individual earnings and labor market transitions in different parts of the income distribution. For this, we divide the distribution into 20 equally sized bins (ventiles).

3.1 Identifying monetary policy surprises

We focus on the period between January 1999 and December 2012, when European monetary policy was conducted by the ECB.[7] 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, Z_t, as in (Almgren et al., 2019). 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.[8]

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 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.

3.2 Earnings and labor market transitions after a monetary policy surprise

This section documents the effect of monetary policy shocks on earnings and labor market transitions in our sample. For this, we estimate two regressions. First, we estimate:

$$earn_{t+h} - earn_{t-1} = \alpha + \beta_{earn,h}\Delta i_t + \gamma X_t + \epsilon_t \tag{1}$$

[7] The high-frequency identification approach outlined here cannot be implemented for earlier time periods, as the Bundesbank did not relay it's policy decision on a precisely planned schedule on the announcement day.

[8] 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).

where the LHS of the equation captures the growth in average real monthly earnings in our dataset between periods $t-1$, i.e. one period before the shock, and period $t+h$. The vector X_t contains calendar month dummies and lagged values of Δi_t and Z_t. Earnings growth is deflated using the Harmonized Index for Consumer Prices for Germany.[9] Note that, 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 are always drawn from two different spells.

To study the effect of monetary policy on transitions in the labor market, we assign individuals every month to one of three labor market states: we define as *employed* [John to fill in]. We define as *unemployed* all those [John to fill in]. The remainder we define as *non-employed*. Similarly to (1), we then estimate the following regression

$$TR^{s_1,s_2}_{t+h} = \alpha + \beta_{TR,h}\Delta i_t + X_t + \varepsilon_h \tag{2}$$

where $TR^{s_1,s_2}_{v,t+h}$, indicates the share of individuals transitioning from state s_1 to s_2 by period $t+h$, conditional on being in state s_1 at $t-1$. s_1 captures two possible states, namely employed E and non-employed N, while s_2 now captures two possible labor market states: moving from employment to employment ("E to E"), or moving to non-employment ("E to N"). Hence, $\beta^v_{TR,h}$ measures the percentage point change in the probability of a particular labor market transition in response to a monetary policy shock. Again, the vector X_t contains calendar month dummies and lagged values of Δi_t and Z_t.

We compare the dynamic effect of this particular shock, to monetary policy, captured by $\beta_{earn,h}$ and $\beta_{TR,h}$, to an alternative measure that quantifies the comovement of individual earnings and transition probabilities with a measure of the business cycle more generally (as resulting from all shocks to the economy). Specifically, in the spirit of Guvenen et al. (2017), we use average earnings in our sample as a proxy of the business cycle and estimate the following regression:

$$x_{t+h} = \alpha_q + \beta^Y_{x,h}\Delta_h Y_t + \gamma X_t + \epsilon_t \tag{3}$$

where the dependent variable is either earnings growth or the share of transitions between two labor market states ($x_{t+h} \in \{earn_{t+h} - earn_{t-1}, TR^{s_1,s_2}_{t+h}\}$), as respectively in (1) and (2), and $\Delta_h Y_t$ denotes the (log-) change in average earnings between period $t+h$ and $t-1$.

[9]Obtained from Eurostat, series prc_hicp_midx.

3.3 Monetary policy effects on earnings across the income distribution

This section shows that the response of individual earnings and labor market transition probabilities to monetary policy shocks differs substantially along the income distribution. For this, we estimate $\beta_{Earn,12}$, the one-year response in (1), separately for the earnings growth of individuals who are in period t in the same ventile of the distribution of our permanent-income proxy.

Figure 1: Regression coefficients $\beta_{EARN,12}$ for ventiles of the income distribution

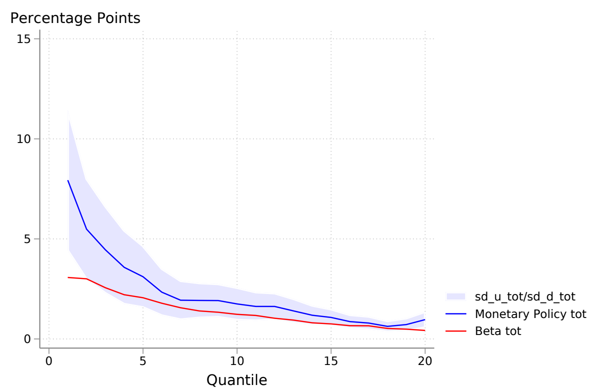

a) Regression coefficients β^p_{EARN} for the Full sample

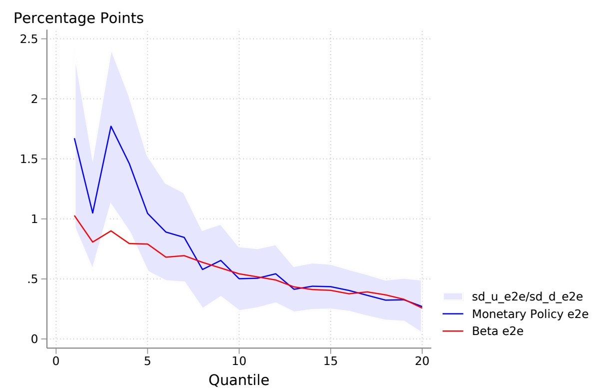

b) Regression coefficients β^p_{EARN} for E to E

Note: The *First Panel* plots the coefficient $\beta_{EARN,12}$ in Equation (1) estimated separately for individiuals who shared the same ventile of the permanent-income distribution in period t, for the full sample. Income growth is computed as average (log) growth in income of individuals who were in the same ventile at time t. The sample period is from 2000 until 2014. The shaded area indicates 95 percent confidence bands. The *Second Panel* plots the coefficients $\beta_{EARN,12}$, only including individuals who transition form employment to employment.

The left panel of Figure 1 shows the point estimates of β_{EARN} in Equation (1) (the blue line), together with 95-percent confidence bands (the shaded area). [ADD DISCUSSION OF RESULTS]

The right panel of Figure 1 shows the point estimates of β_{EARN} when we restrict our sample to those individuals who [remain/ are] employed [between/in] periods t and $t-1$. [ADD DISCUSSION OF RESULTS]

Figure 1 also shows the $\beta^Y_{earn,12}$ coefficients in Equation (3), as a measure of comovement between individual and average annual earnings growth. [ADD DISCUSSION OF RESULTS]

3.4 Monetary policy effects on labor market transitions across the income distribution

The substantially smaller effect of monetary policy shocks on earnings of the employed in Figure 1 suggests that much of the policy effect comes from changes in the probability of job-finding and separation. In this section, we therefore look at transitions between different labor market states. Similar to the previous section, we therefore estimate $\beta_{TR,12}$, the one-year response in (2), separately for every ventile of the income distribution, and plot the result in Figure 2.

Figure 2: Regression coefficients β^q_{TR}

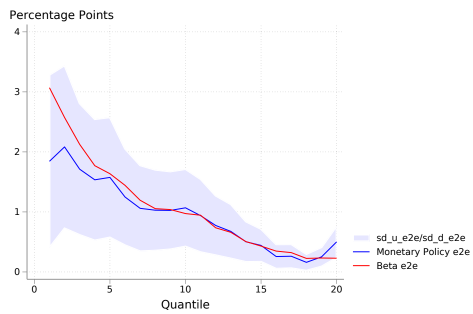

a) E to E

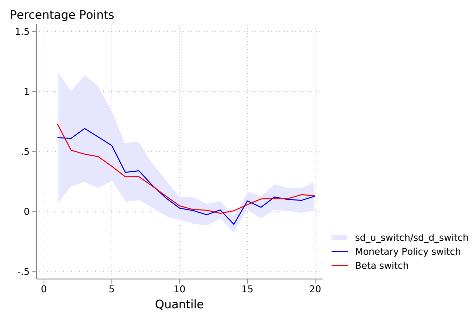

b) Switch

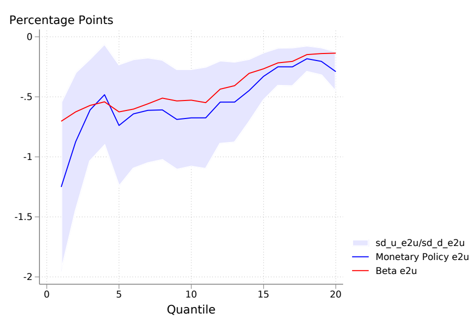

c) E to U

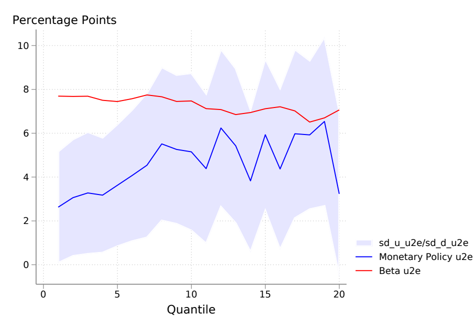

d) U to E

Note: The *first Panel* plots the coefficient β^q_{TR} in (**??**) from separate regressions for those individuals who remain employed between periods t and $t+4$ (E to E). The *second Panel* restricts the sample to those employed in period t and $t+4$ that transition through non-employment (Switch). The *third panel* restricts the sample to those moving from employment to unemployment (E to U). The *fourth panel* restricts the sample to those moving from unemployment to employment (U to E). Individuals are sorted into quantiles based on their average income throughout the five years prior to quarter t. The sample period is from 2000 until 2014.

The top-left panel of Figure 2 shows ... [ADD DISCUSSION OF RESULTS]

3.5 Monetary policy effects on labor market prospects and wages after unemployment

THe analysis so far suggests that much of the effect of monetary policy, and a substantal part of its heterogeneous incidence, arises from labor market transitions. In this section, we focus on one particular transition, from employment to unemployment, and investigate the effect of monetary policy shocks on re-hiring wages and re-remployment probabilities. For this, we first trace the average employment status (either 0 or 1) and the logarithm of real monthly earnings $earn_t$ of all individuals who become non-employed a period t, from $k = 6$ months before their non-employment until $k = 36$ months afterwards. Then, for each k we run the following regression:

$$y_{t+k} = \alpha_t + \beta X_t + \beta_{y,k} \Delta i_t + \epsilon_t$$

where y_t is monthly earnings or employment status for all individuals who become non-employed in period $t = 1$, X_t contains calendar-month dummies and $\Delta i_t + k$ represents the interest rate change in period t, instrumented using Z_t. In this way, we identify (i) the earnings and employment responses upon non-employment, and (ii) the impact of monetary policy on these time paths. The regression is similar to **?**, but they investigate earnings paths relative to those who remain employed. We, instead, focus on the subsample of those who become non-employed and report their earnings paths in different monetary policy regimes. For our analysis, we restrict the sample to those individuals who are employed for six months before becoming non-employed.
Figures 3 and 4 show the results.

4 Conclusion

References

Almgren, M., J. E. Gallegos, J. Kramer, and R. Lima (2019). Monetary policy and liquidity constraints: Evidence from the euro area. *Available at SSRN 3422687*.

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

Bilbiie, F. O. (2020). The new keynesian cross. *Journal of Monetary Economics 114*, 90–108.

Figure 3: Effect of monetary policy shock on re-employment probabilities

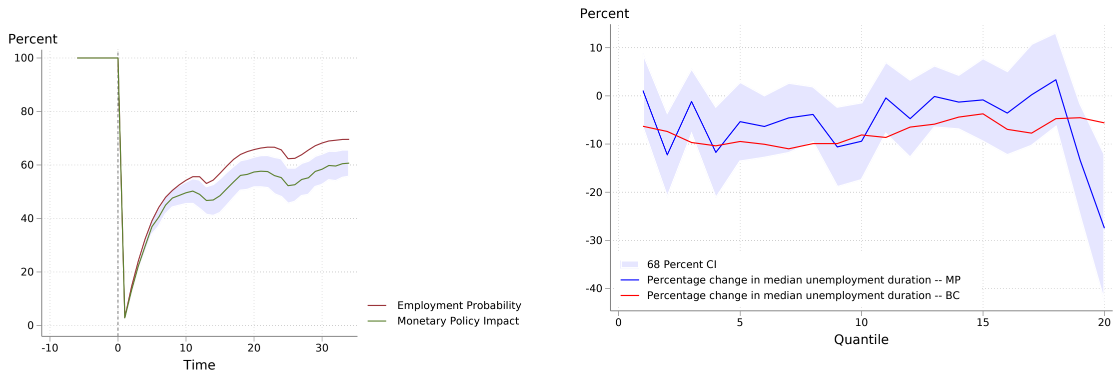

(a) Average employment probability

(b) Unemployment duration across the income distribution

Note: The *Left Panel* shows the employment probability of individuals who transition into non-employment in month $t = 0$ with and without a 25 basis point monetary policy surprise, over time. The *right panel* plots the coefficient $\beta^q_{Y,post}$ in Equation (**??**) by quantile q. Individuals are sorted into quantiles by their average income throughout the five years prior to quarter t. The sample period is from 2000 until 2014. The shaded area indicates 68 percent confidence bands.

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

Challe, E. (2020). Uninsured unemployment risk and optimal monetary policy in a zero-liquidity economy. *American Economic Journal: Macroeconomics 12*(2), 241–83.

Christiano, L. J., M. Eichenbaum, and C. L. Evans (1999). Monetary policy shocks: What have we learned and to what end? *Handbook of macroeconomics 1*, 65–148.

Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2012). Innocent bystanders? monetary policy and inequality in the us. Technical report, National Bureau of Economic Research.

De Nardi, M., G. Fella, M. Knoef, G. Paz-Pardo, and R. Van Ooijen (2019). Family and government insurance: Wage, earnings, and income risks in the netherlands and the us. *Journal of Public Economics 193*, 104327.

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).

Figure 4: Effect of monetary policy on re-employment earnings

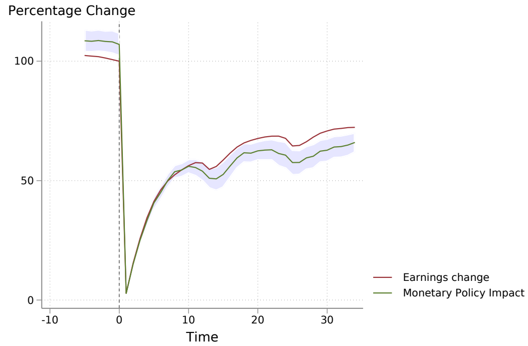

(a) Earnings level following unemployment

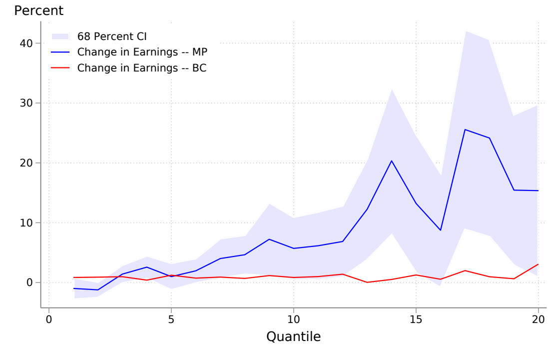

(b) Earnings change across the income distribution

Note: The *Left Panel* shows the earnings change for individuals who transition into non-employment in month $t = 0$ with and without a 25 basis point monetary policy surprise, over time. The *Right Panel* shows the earnings change for individuals who initially claim unemployment benefits. The shaded areas represent 68 percent confidence intervals. The sample period is from 2000 until 2014. We restrict the sample to those individuals who are employed for six months before becoming non-employed.

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

Gornemann, N., K. Kuester, and M. Nakajima (2016). Doves for the rich, hawks for the poor? distributional consequences of monetary policy.

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.

Hagedorn, M., J. Luo, I. Manovskii, and K. Mitman (2019). Forward guidance. *Journal of Monetary Economics 102*, 1–23.

Halvorsen, E., H. A. Holter, S. Ozkan, and K. Storesletten (2020). Dissecting idiosyncratic earnings risk. *Available at SSRN*.

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

Holm, M. B., P. Paul, and A. Tischbirek (2020). The transmission of monetary policy under the microscope. Federal Reserve Bank of San Francisco.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review 95*(1), 161–182.

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

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.

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

Werning, I. (2015). Incomplete markets and aggregate demand. Technical report, National Bureau of Economic Research.

Appendix

A Data

The data are censored from above at the upper earnings limit specified by an individual's pension insurance, which increases over time. We include as much information about the top of the distribution as possible, and therefore exclude top-coded observations only when they do not contain information about annual earnings growth, i.e. when an individual's earnings are top coded both at the beginning *and* the end of a period over which earnings growth is calculated.[10] This applies to around 6-8% of our earnings observations each quarter.[11] Furthermore, until 1999, earnings below the social security reporting threshold (marginal part-time work) were not recorded.

[10] In this way, we partly preserve the income growth of individuals who, e.g. have a non-censured earnings observation in period t, but not in period $t+4$.

[11] 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).

B Additional Results

B.1 Aggregate responses to monetary policy surprises

Before moving to individual incomes, we investigate the effect of monetary policy shocks on the aggregate economy. We run the following local projection regression following Jordà (2005)

$$x_{t+h} - x_{t-1} = \alpha + \beta_h \Delta i_t + \gamma_h X_{t-1} + \varepsilon_{t,h} \tag{4}$$

where Δi_t captures the change in the ECB's policy rate, x represents (i) the monthly inflation rate as measured by the logarithm of German HICP, (ii) the logarithm of industrial production or (iii) the German unemployment rate. The vector X_{t-1} represents a set of control variables consisting of one lag of the instrument Z_t, Δi_t and of x; lastly, it contains calendar month dummies.

Figure 5 shows the impulse responses to a 25 basis point shock to the interest rate, estimated using Equation (4). 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 unemployment rate (bottom panel) increases and industrial production (top right panel) declines in response to the surprise increase in the policy rate. Both respond with a lag of about 6 to 8 months, and then follow a hump-shaped pattern. The responses are, however, rather imprecisely estimated. This is particularly true for the response of inflation (top left panel).

For consistency with the previous literature, we report impulse responses for a 25 basis point shock to the nominal interest rate. The magnitudes of the responses are large, especially when compared to what is usually found in the literature when responses are estimated using standard VARs, i.e. without the use of external instruments (e.g., Christiano et al., 1999). This difference in magnitudes also arises when comparing single equation approaches to externally identified VARs, i.e. SVAR-IV (Coibion et al., 2012; Stock and Watson, 2018). It is important to note, however, that the standard deviation of our shock series is 2 basis points. Because the coefficient in our first stage regression takes a value of 1.4, this implies that a large monthly surprises in our sample affects the nominal interest rate by 3 basis points. Consequently, a 25 basis point surprise should be, in reference to our sample, considered to be a very large outlier, leading to very strong movements in the outcome variables.

Figure 5: Aggregate responses to monetary policy surprises

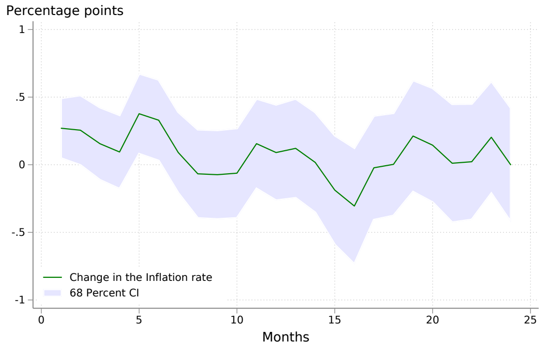

(a) Impulse response of the inflation rate

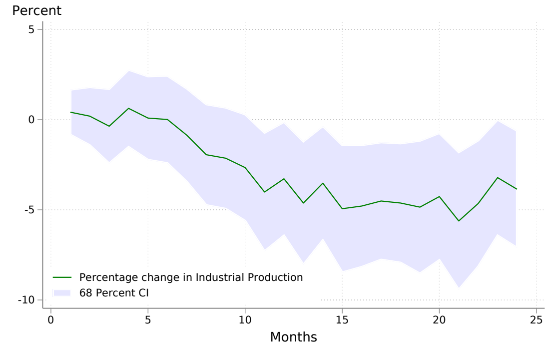

(b) Impulse response of Industrial Production

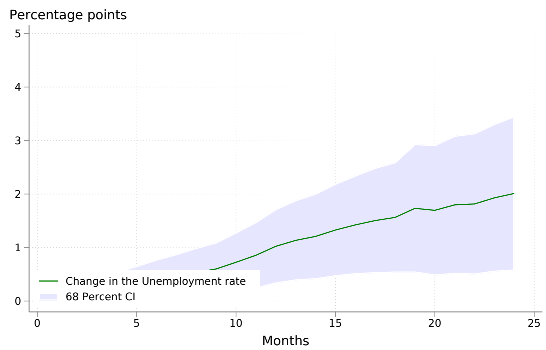

(c) Impulse response of the unemployment rate

Note: The Figure shows the impulse responses of aggregate variables to a 100 bp surprise increase in the the policy interest rate, estimated using the Local Projection outlined in Equation (4). 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.

B.2 Unemployment rate

The left panel of Figure B.2 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.[12] 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 B.2 shows that the unemployment rates calculated using different

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

Figure 6: German unemployment rates

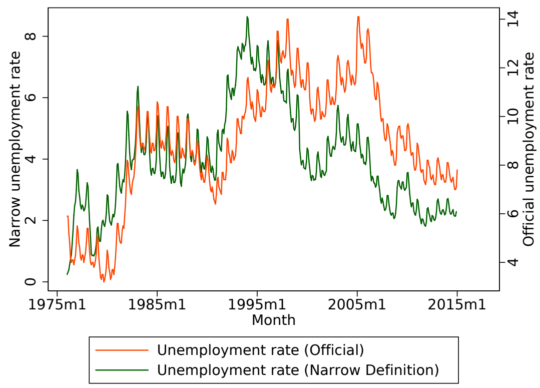

(a) Official rate and narrow definition (only west)

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

definitions of employment behave very similarly.

B.3 Decomposition

The findings in the previous section suggest that the extensive margin of earnings is more important for the heterogeneity of cyclicality across the distribution than the intensive margin. However, the approach used above does not allow us to ascertain whether the heterogeneity is driven by cyclical changes in the transition probabilities or changes in the earnings growth given a certain labor market transition. Here, we estimate the same quantities as above, however, using individual income growth rates as opposed to the growth of average income. This allows us to find, up to a first order approximation, the contributions of the cyclicality of labor market transitions and the cyclicality of earnings changes, given a transition, respectively.

To estimate the cyclicality of conditional earnings changes by labor market transition, in particular at the extensive margin, we need to accomodate zero-earnings observations. To achieve this, we calculate individual earnings growth as

$$\Delta earn_{i,t}^{p} = sinh^{-1}\left(earn_{i,t+4}\right) - sinh^{-1}\left(earn_{i,t}\right) \tag{5}$$

and run the following regression:

$$\Delta earn_{i,t}^{p} = \alpha + \beta_{Y}^{p} \Delta Y + \gamma X_{t} + \epsilon_{j,t} \tag{6}$$

where $\Delta earn_{i,t}^{p}$ represents the earnings change of individual i who belongs to percentile p, between periods t and $t+4$. The vector X_t contains calendar quarter dummies and a dummy controlling for the period after 2005, in which marginal part-time workers enter the sample. As before, this regression yields a coefficient β^p_{earn}, quantifying the correlation between individual incomes and aggregate income across percentiles. Note, however, that the construction of the left-hand side in equation (6) differs from that in equation (3) as follows: the former is a measure of individual earnings changes, while the latter uses the change in average earnings in percentile p between periods t and $t+4$. Figure 7 plots the coefficients of β^p_{earn} obtained with the two different methodologies. The two lines are different in levels, but almost parallel to each other. The reason for the level difference is that when computing earnings changes involving zeros using hyperbolic sine, the level of the non-zero observation matters. If, e.g., an individual transitions from employment to unemployment, i.e., to zero earnings, the hyperbolic sine is (close to) the logarithm of the last earnings. This feature of the approach leads the β_{earn} estimates to be about twice those obtained using the initial estimation.

Crucially, using individual earnings growth allows us to decompose β^p_{earn} into two components:

$$\beta_{EARN} \approx \beta_{trans} + \beta_{\Delta earn|trans} \tag{7}$$

$$\beta_{trans} = \sum_{k=1}^{J} \left[\bar{\mu}_k \beta_{int,k}\right] \tag{8}$$

$$\beta_{\Delta earn|trans} = \sum_{k=1}^{J} \left[(\bar{\Delta} earn_k - \bar{\Delta} earn)\beta_{tr,k}\right] \tag{9}$$

The first term captures the elasticity of transition probabilities, while the latter captures the elasticity of conditional earnings changes, respectively, with respect to aggregate earnings changes. This decomposition was not possible with our initial estimation, as there we take logs of averages, while here we rely on the averages of logs. The index k represents the five different transition paths between E and N, accounting for job-stayers and job-switchers separately. $\Delta earn_{kt}$ and μ_{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 decomposition is exact up to a third order term (that equals less than 0.06 percent on average)

Intuitively, β_{trans} estimates the cyclicality of individual earnings in response to aggregate earnings movements, holding conditional earnings changes fixed. Similarly, $\beta_{\Delta earn|trans}$ represents the correlation of individual earnings with aggregate earnings, if transition probabilities remained constant.

Figure 7: Decomposition of β^{p}_{EARN}

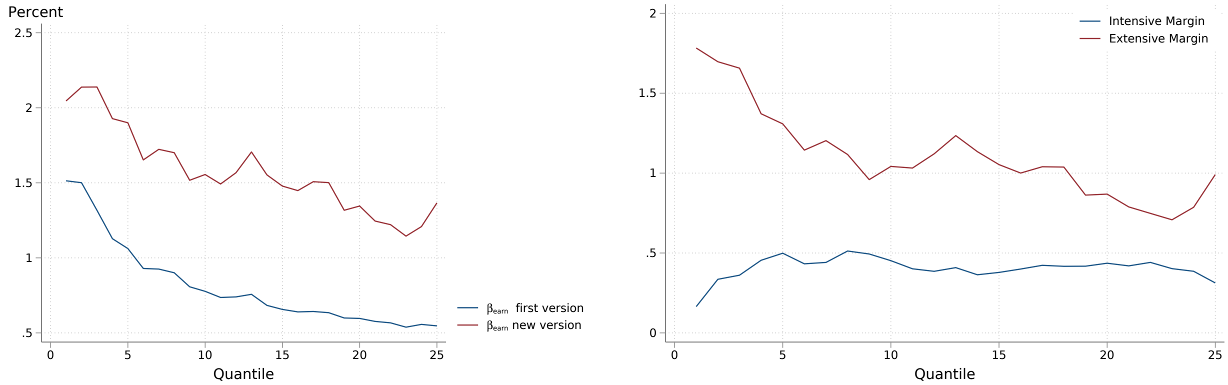

(a) New estimation

(b) Decomposition

Note: The *Left Panel* plots the coefficients β^{p}_{earn} in Equations (3) and (6). The blue line relies on the growth of average earnings, whereas the red line uses individual earnings growth on the LHS. The sample period is from 1980 until 2014. The *Right Panel* plots the coefficients β^{p}_{trans} and $\beta^{p}_{\Delta earn|trans}$ as outlined in equation (20)

B.4 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 8 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 8: Full Impulse Responses

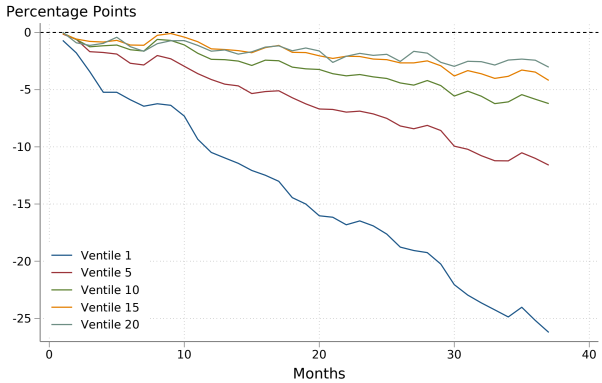

(a) Probability of remaining employed

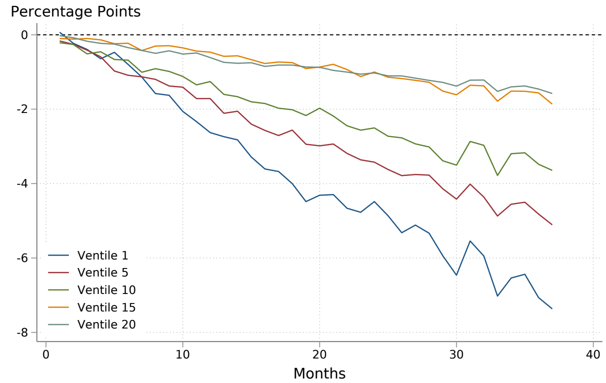

(b) Probability of finding employment

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.

B.5 Probit estimation of transition probabilities

In the main body, we present results estimated using linear regression analysis. However, since Equation (**??**) features a binary dependent variable, and we are dealing with transition *probabilities*, we repeat the analysis presented above using a probit regression. Figure 9 shows that the results are almost equivalent to the linear model used before.

B.6 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

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

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

Figure 9: Non-employment to employment

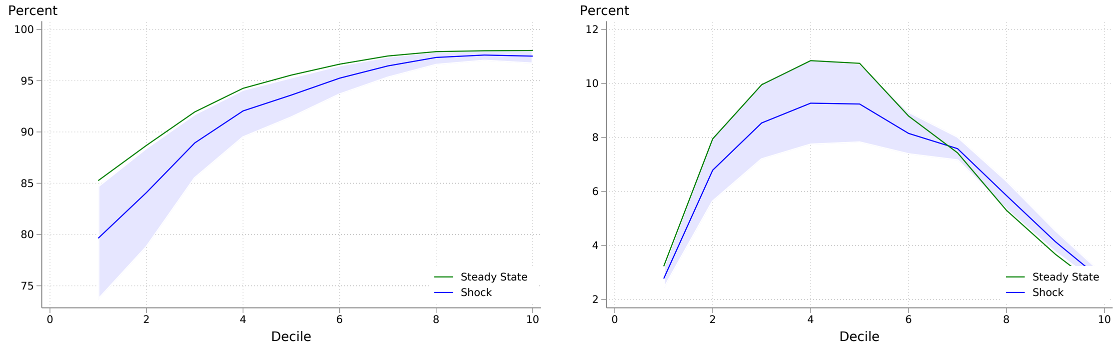

(a) 12 month responses – employed (b) 12 month responses – non-employed

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.

days of the year.

Figure 10 presents the average difference between the number of hirings (left panel) and separations (right panel) around monetary policy announcements, relative to days further 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 10: The labor market around monetary policy announcements

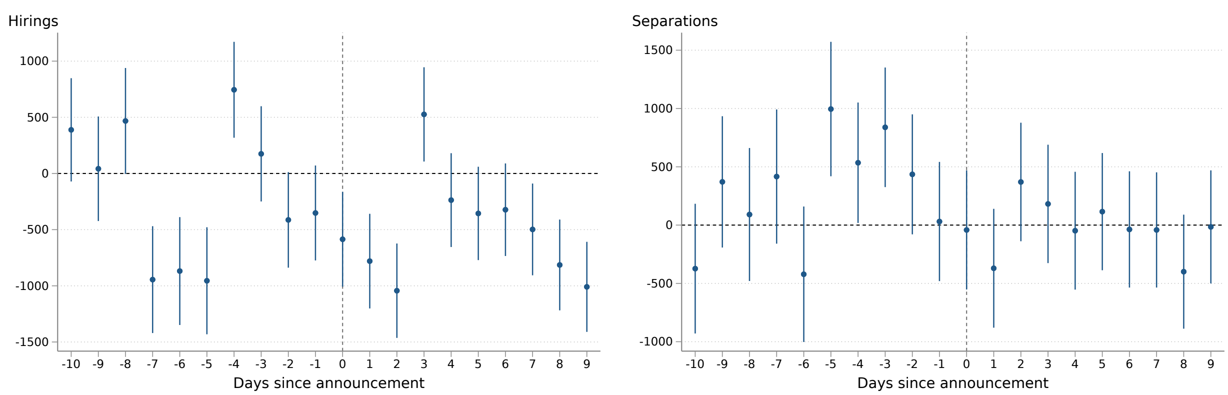

Figure 11: Contribution by flow – Full sample (left) and Fully attached (right)

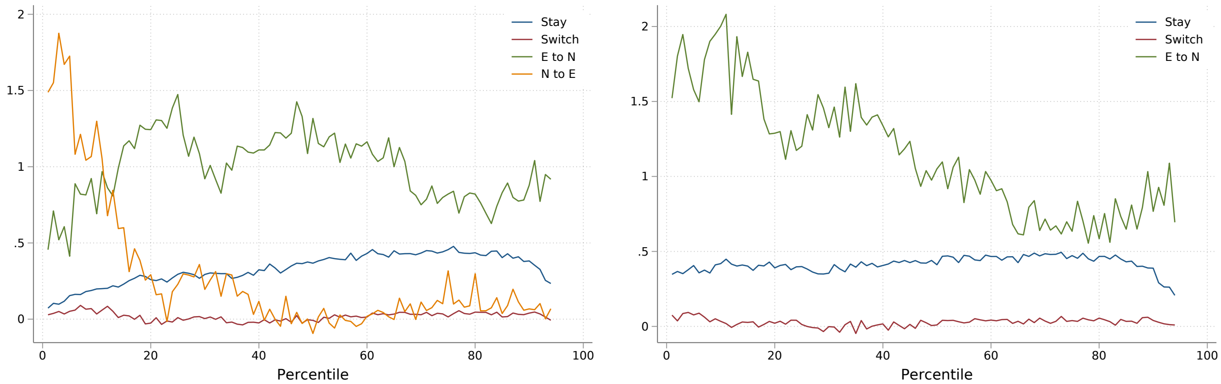

Figure 12: Extensive vs Intensive margin by flow – Full sample

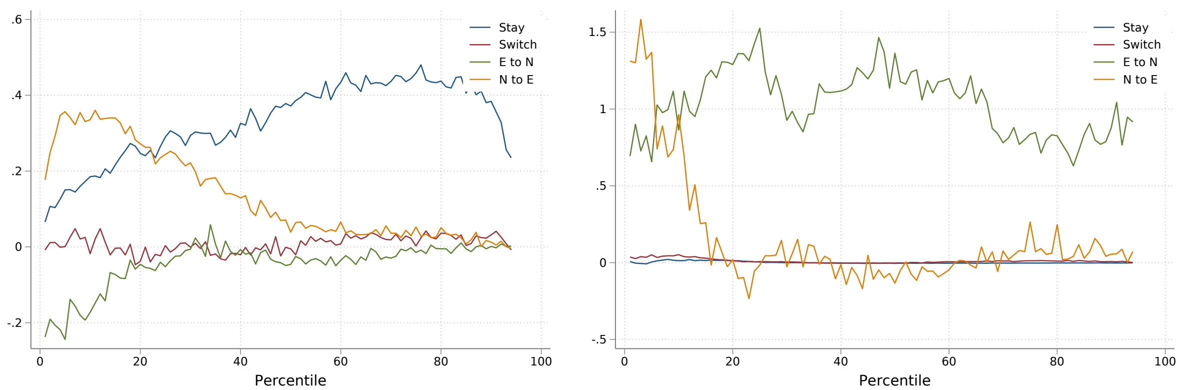

Figure 13: Extensive vs Intensive margin by flow – Fully attached

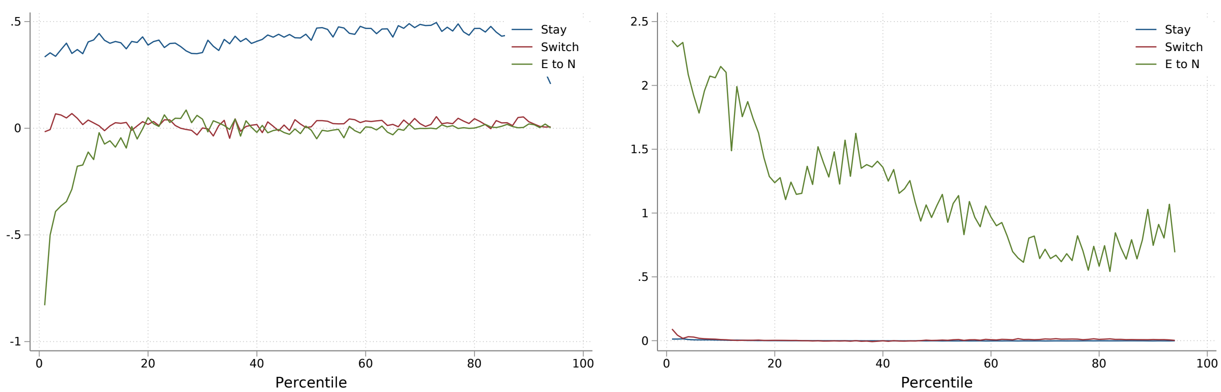

B.7 Decomposition of earnings betas

This section shows how to decompose the overall regression coefficient β_Y^p into a weighted average of β_{TR}^p in (**??**) and coefficients β_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 $\mu_j = \frac{\bar{N}_j}{\bar{N}}$.

$$\mathrm{Cov}(dy_{it}, dY_t) = \frac{1}{TOT} \sum_t \sum_{i=1}^{N_t} (dy_{it} - \bar{dy})(dY_t - \bar{dY}) \tag{11}$$

$$= \frac{1}{TOT} \sum_t \sum_{j=1}^{J} \sum_{i=1}^{N_{jt}} (dy_{ijt} - \bar{dy}_{jt} + \bar{dy}_{jt} - \bar{dy}_j + \bar{dy}_j - \bar{dy})(dY_t - \bar{dY}) \tag{12}$$

$$= \frac{1}{TOT} \sum_t \sum_{j=1}^{J} N_{jt} \left[0 + (\bar{dy}_{jt} - \bar{dy}_j) + (\bar{dy}_j - \bar{dy})\right] (dY_t - \bar{dY}) \tag{13}$$

$$= \frac{1}{TOT} \sum_t \sum_{j=1}^{J} \Big[\bar{N}_j (\bar{dy}_{jt} - \bar{dy}_j)(dY_t - \bar{dY}) \tag{14}$$

$$+ (N_{jt} - \bar{N}_j)(\bar{dy}_j - \bar{dy})(dY_t - \bar{dY}) + (N_{jt} - \bar{N}_j)(\bar{dy}_{jt} - \bar{dy}_j)(dY_t - \bar{dY})\Big] \tag{15}$$

$$= \frac{1}{\bar{N}} \sum_{j=1}^{J} \Bigg[\bar{N}_j \frac{1}{T} \sum_t (\bar{dy}_{jt} - \bar{dy}_j)(dY_t - \bar{dY}) \tag{16}$$

$$+ \frac{1}{T} \sum_t (N_{jt} - \bar{N}_j)(\bar{dy}_j - \bar{dy})(dY_t - \bar{dY}) + \frac{1}{T} \sum_t (N_{jt} - \bar{N}_j)(\bar{dy}_{jt} - \bar{dy}_j)(dY_t - \bar{dY})\Bigg] \tag{17}$$

$$= \sum_{j=1}^{J} \left[\bar{\mu}_j \mathrm{Cov}(\bar{dy}_{jt}, dY_t) + (\bar{dy}_j - \bar{dy}) \mathrm{Cov}(\mu_{jt}, dY_t) + \frac{1}{T} \sum_t (\mu_{jt} - \bar{\mu}_j)(\bar{dy}_{jt} - \bar{dy}_j)(dY_t - \bar{dY})\right] \tag{18}$$

$$\approx \sum_{j=1}^{J} \left[\bar{\mu}_j \mathrm{Cov}(\bar{dy}_{jt}, dY_t) + (\bar{dy}_j - \bar{dy}) \mathrm{Cov}(\mu_{jt}, dY_t)\right] \tag{19}$$

This yields

$$\beta_{EARN} = \frac{\text{Cov}(dy_{it}, dY_t)}{Var(dY_t)} \approx \beta_{int} + \beta_{ext} \tag{20}$$

$$\beta_{int} = \sum_{j=1}^{J} \bar{\mu}_j \beta_{int,j} = \sum_{j=1}^{J} \bar{\mu}_j \frac{\text{Cov}(\bar{d}y_{jt}, dY_t)}{Var(dY_t)} \tag{21}$$

$$\beta_{ext} = \sum_{j=1}^{J} (\bar{d}y_j - \bar{d}y) \beta_{TR,j} = \sum_{j=1}^{J} (\bar{d}y_j - \bar{d}y) \frac{\text{Cov}(\mu_{jt}, dY_t)}{Var(dY_t)} \tag{22}$$