The Cyclicality of Earnings Growth along the Distribution - Causes and Consequences

John Kramer*

First Draft: January 26, 2022
Current Draft: May 22, 2026

Abstract

Earnings growth is more procyclical at the bottom of the income distribution than at the top. Using administrative data from Germany, I decompose this heterogeneity by labor-market transition type. Job-finding accounts for 60% of earnings procyclicality in the lowest decile, declining monotonically to near zero above the median, while the contribution of job-separations is roughly constant across the distribution. Two forces drive this pattern: non-employment is concentrated at the bottom of the income distribution, and job-finding rates are highly procyclical. I build a heterogeneous-agent business cycle model with directed search that endogenously reproduces these findings. Workers choose where to search, trading off wages against job-finding probabilities. In response to positive productivity shocks, firms post more vacancies, disproportionately raising job-finding rates for low-income workers, among whom non-employment is most prevalent. The model replicates the heterogeneous procyclicality of earnings growth and the dominant role of job-finding along an endogenous earnings distribution, without targeting these moments in the calibration. I use the model to evaluate countercyclical hiring subsidies. The policy nearly eliminates the procyclicality of vacancy posting and reduces the welfare cost of recessions by 30%, with gains concentrated among low-productivity unemployed workers whose job-finding rates are most sensitive to the cycle.

*University of Copenhagen, Department of Economics, Øster Farimagsgade 5, Building 26, 1353 København K, Denmark; email address: john.kramer@econ.ku.dk.

1 Introduction

Recessions disproportionately affect the incomes of the poor. More generally, their earnings growth follows the business cycle more closely than does the earnings growth of the rich (Guvenen et al., 2017). I show that this heterogeneity in income growth over the business cycle is mainly driven by individuals transitioning from non-employment into employment. To investigate the consequences of this phenomenon for policy, I propose a heterogeneous-agent business cycle model that endogenously generates heterogeneous procyclicality in earnings growth, along an endogenous income distribution, driven by job-finding. I use this model to evaluate the effectiveness of countercyclical hiring subsidies in reducing the cost of business cycles.

Using data on individual yearly earnings in the US, Guvenen et al. (2017) show that the yearly earnings growth at the bottom of the permanent income distribution comoves more closely with GDP growth than earnings growth at the top. I show, using administrative microdata, that the same pattern holds in Germany, suggesting that it is not a US-specific phenomenon.[1] My main empirical contribution is the investigation of the sources of this heterogeneity, since it has implications for the distribution of business cycle risk along the income distribution.

To this end, I decompose earnings growth into the contributions of different labor-market transitions. Importantly, my data, unlike many other administrative datasets, allows me to observe employment spells at daily frequency, and distinguish precisely between job-finding and job-separation episodes. I find that the main driver of the heterogeneity in earnings growth over the business cycle are the earnings changes of workers transitioning from non-employment to employment. Job-finding is responsible for 60% of the procyclicality in the first decile, declining monotonically to near zero above the median. By contrast, separations raise the level of earnings procyclicality but do not generate heterogeneity in it. This distinction is important because it implies that the steeply declining pattern of earnings betas is almost entirely accounted for by job-finding, making it a natural target for policy intervention.

I show this by, first, restricting the sample to individuals who stay employed continuously. Within this subsample, earnings procyclicality is homogeneous, implying

[1] Relatedly, studies investigating the impact of monetary policy on earnings growth along the income distribution find almost identical patterns of strongly downward sloping incidence of expansionary monetary policy (Amberg et al., 2022; Andersen et al., 2023; Broer et al., 2020; Holm et al., 2021).

that the earnings changes of workers who transition along the extensive margin must be responsible for the heterogeneity.[2] In a second step, I isolate job-finding as the key source of heterogeneity by comparing two subsamples: the continuously employed sample, and another that includes both continuously employed workers and job-finders (but excludes job-separators). The fact that heterogeneity reappears when job-finders are added back demonstrates that transitions from non-employment to employment, rather than job separations, drive the procyclical pattern of earnings growth across the income distribution. This suggests that policies targeting transitions into employment could have a disproportionate impact on stabilizing earnings and reducing inequality over the business cycle, especially for the income-poor.

I uncover the driving forces behind the findings by looking beyond earnings, at the heterogeneous cyclicality of labor market transitions. At the bottom of the permanent income distribution, unemployment is much higher than at the top: individuals in the first decile transition from non-employment to employment with a probability of less than ten percent, while the same probability is 20 percent in the top decile. The second factor is the fact that the job-finding probability is highly procyclical: as aggregate earnings rise by one percentage point, the probability of finding a job rises by about one percentage point. Combined, these two forces lead earnings growth at the low end of the permanent income distribution to comove very strongly with the business cycle.

These empirical patterns motivate the development of a macroeconomic model capable of reproducing the heterogeneous procyclicality of earnings and evaluating policy interventions. The second part of the paper presents such a model, featuring idiosyncratic and aggregate risk, and a frictional labor market, similar to Birinci and See (2023). Crucially, the model endogenously produces the earnings dynamics documented in the data, unlike models that impose exogenous income processes to fit them (Guvenen et al., 2014), and both worker and firm choices are free to respond to policy interventions.

The model features heterogeneous, risk-averse workers who search for work across submarkets which offer different wages and job-finding probabilities. Workers differ along three dimensions: productivity, wealth and wages, and search for jobs while unemployed and employed. When deciding where to search for jobs, they endogenously trade off higher job-finding probabilities and higher wages. Hence, in

[2]This is in line with Hoffmann and Malacrino (2019) who show that procyclicality in higher order moments of the earnings growth distribution is driven by extensive margin transitions.

the model, both worker choices and firm choices (about vacancy posting) influence earnings growth. Agents can smooth consumption and insure against earnings fluctuations by saving in a risk-free asset.

Highly productive unemployed workers have strong incentives to find work quickly, for two reasons. First, their potential market wages substantially exceed their income in unemployment, implying that they are missing out on potential consumption. Second, these workers can insure against a future fall in their productivity by finding employment, as wages are fixed for the duration of the match. Conversely, low productivity workers have high reservation wages, relative to their productivity, because, for them, unemployment benefit income has a meaningful disincentive effect: their (exogenous) separation probabilities are high, and their productivity is likely to increase over time. Not wanting to get "trapped" in low paying jobs, as searching while employed is less efficient, this leads them to search in submarkets which offer high wages, relative to their productivities, but low job-finding probabilities. These forces endogenously create heterogeneous job-finding probabilities along the income distribution, and concentrate unemployment at the bottom – a key feature of the data.

An important insight provided by the model is that heterogeneous idiosyncratic productivity is key for matching the observed heterogeneity in the data. Conditional on productivity, wealthier workers have lower incentives to find work, as they can insure their consumption through their savings (see, e.g., Eeckhout and Sepahsalari, 2024; Repele, 2025). Hence, models that feature homogeneous productivity, or constant heterogeneous productivity, would predict job-finding probabilities that fall along the earnings distribution – as earnings are correlated with wealth. This, however, is counterfactual to the data.

I calibrate the model to match moments of the German labor market, including the heterogeneous labor-market transition rates I document in the data. I leave earnings growth and fluctuations over the business cycle untargeted. In steady state, the model is able to closely replicate earnings growth rates along the distribution. Further, the model very closely matches the distribution of liquid asset holdings in the data, which is similarly untargeted. Matching asset holdings is particularly important as workers' self-insurance must be taken into account when assessing the insurance value of policies.

I solve for the model's response to a sequence of aggregate productivity shocks. It successfully, endogenously, reproduces the strongly decreasing procyclicality of earnings growth along the income distribution, as well as the heterogeneous contribution of job-finding. The main drivers behind the result are the high unemployment incidence among low-productivity workers, combined with procyclical vacancy posting by firms. In a recession, as vacancy posting decreases, this has a disproportional impact on earnings growth at the bottom of the income distribution, where non-employment is most concentrated. The model closely matches the procyclicality of earnings growth in the lower deciles, capturing the key empirical patterns. At the top of the distribution, the model produces somewhat stronger procyclicality than observed in the data, since wage adjustments for continuously employed workers are absent in the model.

To evaluate the welfare consequences of the heterogeneous procyclicality in earnings growth, I introduce countercyclical hiring subsidies into the model, financed with a tax on labor income. The policy is aimed at reducing the risk of lower vacancy posting during recessions, which is disproportionately borne by workers at the bottom of the recent earnings distribution. These workers are more likely to be unemployed and hold fewer assets, and are thus particularly vulnerable to business cycle fluctuations. In the counterfactual, firms receive a one-time lump-sum transfer from the government when a match is formed, with the size of the subsidy indexed to aggregate productivity. This leads to a pronounced reduction in the volatility of vacancy posting: in the baseline model, a negative 1% productivity shock leads to a 6% reduction in vacancy posting; in the counterfactual economy, posting is almost unaffected.

The welfare effect of the policy is positive: the average worker would need to be given 0.06% of additional baseline consumption per period in order to be as well off during a recession as in the economy with a hiring subsidy. This effect is sizeable, when compared with the overall welfare effect of a 1% negative productivity shock in the baseline economy, the consumption equivalent variation of which is 0.2%. Thus, the hiring subsidy lowers the welfare impact of recessions by 30%.

The policy's positive effects are concentrated among low-productivity unemployed workers. For them, job-finding probabilities are most procyclical over the business cycle and the size of the subsidy is large relative to their wages. The policy experiment thus shows that reducing the procyclicality of vacancy creation has welfare benefits, because it reduces a risk faced by workers over the business cycle.

Related Literature The empirical part of this paper is related to a fast growing literature documenting the heterogeneity in earnings growth over the business cycle.[3] Guvenen et al. (2014) show that higher order moments such as skewness and kurtosis are cyclical, as well as heterogeneous across the labor income distribution. Similar results have been found for other countries, such as Germany, France and Sweden (Busch et al., 2022) and Denmark (Harmenberg and Sievertsen, 2017). The main motivation for the present paper is that Guvenen et al. (2017) and Dany-Knedlik et al. (2021) show that in the US, "earnings betas", i.e., the correlations between individual and aggregate earnings growth, are higher at the bottom of the income distribution. My paper adds to this literature by empirically decomposing the sources of the documented earnings growth heterogeneity. I isolate job-finding, as distinct from job-separation, as the primary source of heterogeneous earnings procyclicality across the income distribution. This distinction has direct policy implications, as job-finding is potentially responsive to vacancy-side interventions, motivating the policy analysis in the second part of the paper.

In contributions closely related to the empirical section of this paper, Hoffmann and Malacrino (2019) show that extensive margin transitions drive the cyclicality of the earnings distribution's higher order moments in Italy; and Broer et al. (2020) find that the incidence of monetary policy on earnings growth in Germany is driven by the extensive margin. However, neither paper attributes identifies procyclical job-finding as the main labor market flow responsible for their findings.

The model presented in this paper is a directed search framework. As such, it is related to Gregory et al. (2025), who study such a model with permanent worker heterogeneity to match empirical patterns in labor market transitions. Their framework provides a rich description of labor market dynamics, including endogenous separations, but abstracts from precautionary savings, limiting its suitability for policy analysis. By contrast, my model incorporates both incomplete markets and savings behavior, and is thus more closely related to that of Birinci and See (2023), who show that time-varying unemployment benefits can help smooth consumption in recessions. While their government budget clears in expectation with fixed tax and interest rates, my model allows these prices to adjust endogenously over the business cycle.

A key feature of the model is the inclusion of persistent, idiosyncratic worker-level productivity shocks, which are crucial for matching observed patterns of job search across the earnings distribution. In standard settings, wealth-poor workers have strong incentives to accept jobs quickly, which would lead to higher job-finding

[3]A related literature studies earnings growth heterogeneity while not focusing on cyclicality (see, e.g., De Nardi et al., 2021; Guvenen et al., 2014; Halvorsen et al., 2019).

rates at the bottom of the income distribution (Eeckhout and Sepahsalari, 2024; Repele, 2025). Introducing time variation in productivity reverses this relationship and generates job-finding behavior across the distribution in line with the data.

Moreover, the model replicates the empirical distribution of asset holdings across income groups, consistent with Chaumont and Shi (2022). This joint ability to match both labor market transitions and wealth heterogeneity is central to the model's quantitative performance and its suitability for policy analysis.

Finally, this paper relates to a long literature investigating policies which stabilize the macroeconomy during a recession (see, e.g., Costain and Reiter, 2005; McKay and Reis, 2016). Birinci and See (2023) study the implications of heterogeneity for unemployment insurance provision. I evaluate the effectiveness of hiring subsidies, similar to Kitao et al. (2011), who show that this instrument can be particularly powerful when labor markets are weak.

The rest of the paper is organized as follows. Section 2 presents the empirical analysis of the paper. Section 3 outlines the model, Section 4 describes the calibration. Section 5 presents the model results. Section 6 discusses the policy experiments and section 7 concludes.

2 Empirical Analysis

In this section, I compute "earnings betas" for Germany (Guvenen et al., 2017), i.e., the procyclicality of earnings growth, across the permanent income distribution. I decompose the estimates into contributions of different labor market transitions. For the full sample, earnings growth is significantly more procyclical at the bottom of the distribution. Restricting the sample to job-stayers, however, reveals no heterogeneity. Thus heterogeneous procyclicality must stem from extensive margin transitions. Decomposing these further, I show that the earnings growth of job-finders drives most of the heterogeneity in earnings betas. The outsized influence of job-finding on the procyclicality of earnings growth at the bottom of the distribution has two reasons: first, the incidence of unemployment is higher and, second, job-finding is procyclical. Combined, these two forces make earnings growth much more procyclical at the bottom of the distribution.

2.1 Data

For the empirical exercise, I utilize the Sample of Integrated Labor Market Biographies (SIAB), constructed by the Research Data Center (FDZ) of the German Federal Employment Agency (BA).[4] It contains administrative data on a representative two percent sample of German labor market histories between 1975 and 2019, with the exception of civil servants, students and self-employed individuals (Frodermann et al., 2021).[5]

The dataset is divided into labor market spells, which mark distinct episodes during an individual's labor market biography. For each such spell, the data provides information about its start- and end-date, as well as the average daily earnings throughout its duration. Since firms are required to notify the responsible social security agencies about their employees at least once per year (Frodermann et al., 2021), the maximum spell length is one year.

For employed individuals, the earnings measure is pre-tax earnings liable to social security contributions; for non-employed individuals, some unemployment benefit receipts are reported.[6] Periods of non-employment without benefit receipts are missing in the dataset.[7] To construct complete labor market biographies, I declare any missing values at any point during an individual's life as non-employment episodes, except for those before their first or after their last non-missing episode in the data.[8]

Around 5% of German employees are civil servants (Statistisches Bundesamt, 2026), but are missing from the SIAB sample. Thus, transitions into and out of civil service could potentially lead to misclassifications of employment periods as non-employment. However, since I drop missing values after the last non-missing

[4]I use the factually anonymous Version of the Sample of Integrated Labour Market Biographies (SIAB-Regionalfile) – Version 7519. Research Data Centre (FDZ) of the Federal Employment Agency (BA) at the Institute for Employment Research (IAB). Data access was provided via a Scientific Use File supplied by the FDZ of the BA at the IAB.

[5]The excluded individuals make up around 20% of the workforce (Busch et al., 2022).

[6] Three forms of unemployment benefits are reported: i) Unemployment Benefits I (*Arbeitslosengeld I*), which represent an earnings replacement payment immediately after job-loss, for a limited time, (ii) Unemployment Help (*Arbeitslosenhilfe*) are benefits claimable after the exhaustion of unemployment benefits, (iii) Unemployment Benefits II (*Arbeitslosengeld II*), the successor of the Unemployment Help program, implemented through the Agenda 2010 program in 2005. There is no data on Unemployment Benefits II between the years 2004 and 2007.

[7]For the latter years of the sample period, the data also contain information on job-search and marginal part-time employment. To keep the sample consistent over time, however, I assign individuals in these categories to non-employment.

[8]For individuals younger than 55 years, I also declare as non-employed all episodes after their last employment observation if that observation falls into years after 2014. The number of non-employed individuals would fall drastically towards the end of the sample, otherwise.

episode from the sample, this concern is minimal. German civil service is close to an absorbing state—in 2020, only 574 workers left the civil service (Deutscher Bundestag, 2025)—and thus entrants likely never reappear in the SIAB sample.

I deflate all earnings using the quarterly consumer price index.[9] For the subset of non-employed individuals who receive unemployment benefits, I also observe their benefit income. However, in the baseline estimation, I set the earnings for all non-employed individuals to zero for two reasons. First, the duration and generosity of unemployment benefits changed throughout the sample, which implies that the the unemployment benefit series is not consistent over time. Second, focusing only on employment income is in line with the previous literature studying earnings volatility (see, e.g., Busch et al., 2022). In appendix A.6, I perform a robustness exercise including unemployment benefits, which yields qualitatively similar results to the baseline.

I convert average daily earnings during labor market spells to quarterly earnings. To obtain a binary indicator for quarterly employment status, I assign individuals who are employed for more than half of a quarter into employment and all others into non-employment. Employed individuals are assigned the average earnings value of the longest employment spell during the quarter.[10]

Because social security contributions are capped at the assessment ceiling for pension insurance, the earnings data are censored from above, affecting around 6 % of all observations. In most of the empirical analysis that follows, I focus on the lower 90% of the earnings distribution but impute top-earnings following Card et al. (2013) and Dustmann et al. (2009). Their approach relies on predicting the censored observations using a Tobit estimation.[11] Note that because the imputed earnings are randomly drawn from estimated distributions, the earnings growth they produce for each individual mainly reflect movements in aggregate earnings and the earnings growth rates among the top 10 percent should be interpreted with this in mind.

To make the data consistent over time, I first exclude individuals employed in the states of the former German Democratic Republic. Second, marginal part-time

[9]The CPI index (base year 2015) is obtained from the Federal Reserve Bank of Saint Louis' FRED database, and is not seasonally adjusted. The variable ID is DEUCPIALLQINMEI.

[10]Employment spells that are longer than one quarter are split into multiple quarters. The average daily earnings for each quarter are those of the overarching spell.

[11]A second approach fits a Pareto distribution to the upper tail of the earnings distribution and draws replacements for the censored earnings observations using this method. Drawing on a separate dataset, Dustmann et al. (2009) conclude, that the Tobit approach is preferable. In the context of my exercise, results are very similar.

workers are declared as non-employed since they are only registered in the dataset after 1999 and are previously unobserved. Third, due to large fluctuations in the number of individuals whose employment status is coded as "Other employment status" (Frodermann et al., 2021), I label all of these as non-employed in all years of the sample. For my analysis, I restrict the sample period to 1980-2019 and to individuals older than 25 years and younger than 60, leaving close to 71 million person-quarter observations.

2.2 Descriptive Statistics

For each quarter, I generate an income distribution by sorting individuals according to their recent earnings, following Guvenen et al. (2014). Recent earnings are constructed by averaging individual quarterly labor earnings (including zeros) over five years prior to quarter t.[12] This measure is intended as a proxy for permanent income (the results are similar using lifetime earnings as a sorting variable) and all individuals are assigned to one of 20 quantiles along its distribution, trading off noise reduction (larger quantiles) and granularity (smaller quantiles).[13] I exclude individuals who do not receive any labor earnings during the five-year period. Furthermore, because age and gender are influential predictors of recent earnings, I construct the quantiles outlined above conditional on these two characteristics, by assigning individuals to 5-year age brackets.

The left-hand panel of Figure 1 shows the median of real quarterly earnings, as well as the 25th and 75th percentile, conditional on employment, within each quantile. Median labor earnings are about three times larger at the top, compared to the bottom. Interestingly, the slope of the earnings curve is relatively flat along the distribution, steepening only towards the top.

[12] As a robustness exercise, I construct a recent earnings distribution excluding zeros in the appendix and perform the same estimations as in the baseline. The results are very similar.

[13] Unless otherwise stated, I use “quantile” to describe the mass of individuals between two percentiles. Individuals in the second decile, e.g., are those between percentiles 10 and 20.

Figure 1: Earnings and Employment across the distribution

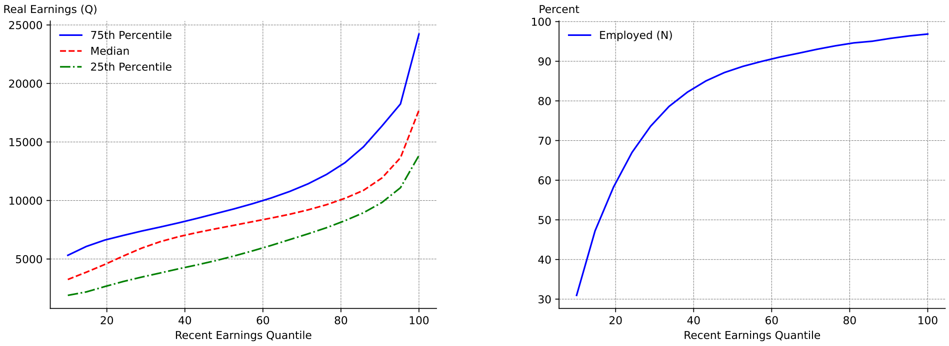

(a) Real Earnings

(b) Employment status

Note: The *Left Panel* shows the median, 25th and 75th percentile of real gross quarterly labor earnings by quantile, deflated using the CPI with base year 2015. The *Right Panel* shows the average employment share by quantile. Individuals are sorted into 20 quantiles each quarter based on their five-year earnings history. The sample period is 1980-2019.

The right-hand panel of Figure 1 plots the share of employed individuals across the distribution. While only close to 30 % of individuals in the first quantile are employed according to the definition outlined above, the fraction rises to 80 % at the 4th decile and to above 90 % beyond the median.

Next, I compute growth of average earnings by quantile, which is the key statistic I use to measure the procyclicality in earnings. Here, due to the high-frequency nature of my data, I am able to make an important modification to the approach used by Guvenen et al. (2017). Because their work relies on yearly earnings observations, it is not possible to identify whether an individual's earnings fluctuations are driven by non-employment during part of the calendar year, or by changes in labor earnings on the job. Because my dataset allows me to evaluate labor market status at higher frequency, I can distinguish between these two forces and explicitly identify the contribution of extensive margin transition to changes in labor earnings. However, *individual* real earnings growth rates are not always well defined when zero-earnings observations are involved. To circumvent this problem, I average real earnings within 20 quantiles (including zeros) and compute the growth of this measure over time (Krueger et al., 2016):

$$\Delta \overline{y}^q_{t,k} = \log(\overline{y}^q_{t+k}) - \log(\overline{y}^q_t) \tag{1}$$

$$\text{where } \overline{y}^q_{t+k} = \frac{1}{N_q} \sum_i y_{i,t+k} \quad \forall i \in q \text{ at } t$$

where $\overline{y}^q_{t+k}$ represents average earnings in quarter $t + k$, for all individuals who are in quantile q in period t and, consequently, $\Delta\overline{y}^q_t$ is the log-growth of average earnings. This approach has the advantage of being able to include zeros in a log-growth rate, as they are subsumed in the aggregate. When using individual growth rates, this is not possible.

Furthermore, my dataset allows me to compute earnings growth rates separately by labor market transitions. I consider five such transitions: job-stayers are individuals who remain with the same employer between periods t and $t + k$, job-switchers are individuals who are employed in periods t and $t + k$, but with different employers, job losers are individuals who are employed in quarter t but non-employed in period $t+k$, job finders are non-employed in period t and employed in period $t + k$ and finally, all other individuals are non-employed in both periods t and $t + k$. Note that I make no restrictions on labor market status for periods $t < p < t + k$, except for stayers. In what follows, I decompose average earnings growth over time along these dimensions.

The solid blue line in the left-hand panel of Figure 2 displays average quarterly earnings growth (k=1) across the earnings distribution. While earnings growth is as high as 10 % in the first quantile, it mean reverts around the median. Next, I exclude all individuals who transition from employment to non-employment between quarters t and $t + 1$ (black, dotted line). This removes all downside risk for the employed, but retains the upside potential for the non-employed. Unsurprisingly, earnings growth for this restricted sample is considerably higher, especially at the bottom, thus increasing growth by more than 0.1 log-points. At the top of the distribution, growth only increases by around 0.02 log-points when separators are excluded.

Figure 2: Earnings growth and Transitions across the distribution

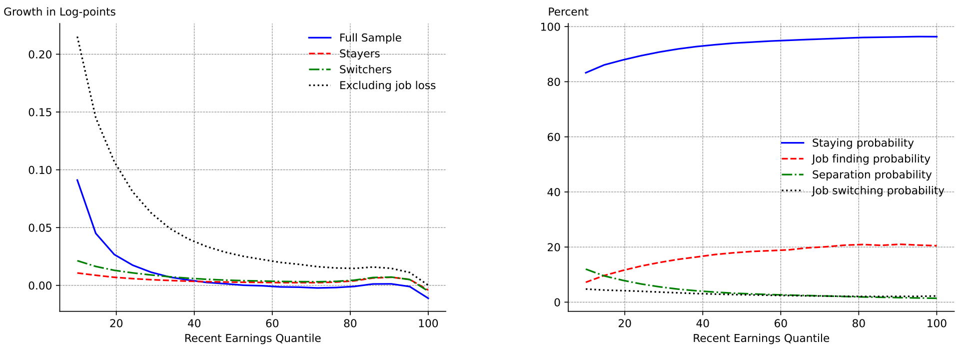

(a) Earnings growth

(b) Transition probabilities

Note: The *Left Panel* shows average earnings growth between quarters t and $t+1$ for various sub-samples, calculated according to Equation (1), within quantile. The blue line represents average quarterly earnings growth for the full sample, the black line excludes job-losers, the green line limits the sample to those who are employed in periods t and $t+1$ and the red line only includes job-stayers. The *Right Panel* shows the average transition probabilities between employment states. The blue line represents the probability of staying employed, the red line represents the job-finding probability, the green line the probability of separation and the black line the probability of switching employers. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The sample period is 1980-2019.

Next, I remove all extensive-margin transitions, by restricting the sample to those individuals who are employed in periods t and $t+1$ (green line). This almost perfectly equalizes average earnings growth across the distribution, with now only slightly higher growth towards the bottom end of the distribution. Restricting the sample even further, to only include individuals who stay with the same employer, shows that the slightly higher growth rate at the bottom seems to be mainly driven by job-switchers.

Focusing more explicitly on labor market transitions, the right-hand panel of Figure 2 shows quarterly transition probabilities across the distribution. There is a pronounced difference in separation probabilities between the bottom and the top of the earnings distribution: while the employed in the first quantiles face a 15 % chance of transitioning into non-employment by the next quarter, the probability is almost zero at the very top. Furthermore, while generally low, the job-switching probability is higher at the bottom as well. Shifting to the non-employed, the figure shows a striking upward slope in quarterly job-finding probabilities across the distribution, implying longer unemployment duration at the bottom. Together, these differences in transition probabilities are responsible for the strong heterogeneity in employment status documented in the right-hand panel of Figure 1.

2.3 Earnings procyclicality across the distribution

Next, I document heterogeneity in the procyclicality of earnings by quantile, similar to Guvenen et al. (2017). I regress the measure of earnings growth introduced in Equation (1) on aggregate earnings growth, by quantile. As a measure for aggregate growth, I choose growth in average earnings, constructed analogously to the quantile-specific growth measure:

$$\Delta Y_{t,k} = \log(\overline{y}_{t+k}) - \log(\overline{y}_t)$$
$$\text{where } \overline{y}_{t+k} = \frac{1}{N}\sum_i y_{i,t+k}$$

Using the growth rates constructed in this way, the procyclicality of earnings along the income distribution can be estimated using the following regression:

$$\Delta\overline{y}^q_{t,k} = \alpha + \beta^q_{Y,k}\Delta Y_{t,k} + X_t + \epsilon_{q,t} \tag{2}$$

where $\beta^q_{Y,k}$ measures the change in quantile-specific growth in response to a 1 percentage point increase in aggregate growth and X_t contains calendar-quarter dummies to account for seasonality. For this analysis, I focus on yearly quarter-on-quarter growth rates, implying $k = 4$.

The left-hand panel of Figure 3 plots the values of $\beta^q_{Y,4}$ obtained from Equation (2) across the earnings distribution. Individuals at the bottom of the distribution see their earnings growth rise by about 3 percentage points, on average, for every additional percentage point in aggregate earnings growth. The same relationship is slightly less than 1-for-1 beyond the 40th percentile, but increasing slightly again between deciles seven and eight. Beyond the 80th percentile, the graph shows a decrease in the earnings betas, likely driven by the imputation of top incomes.

These results are strikingly similar to the ones documented by Guvenen et al. (2021) for the case of the US. The black dashed line in the top-left panel of Figure 3 reports the earnings betas for males, aged 36-45, as reported in the online appendix Table B1 in Guvenen et al. (2021).[14] In both countries, earnings growth is about three times as procyclical at the bottom of the income distribution as it is at the top. The slope is steeper in Germany, implying that around the median, earnings

[14]The data are only available for deciles.

growth is more procyclical in the US.

Figure 3: Cyclicality of Earnings Growth

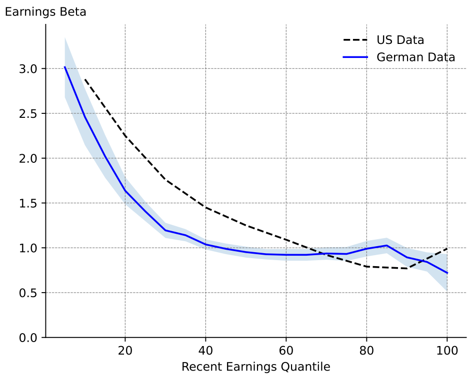

(a) Earnings Beta

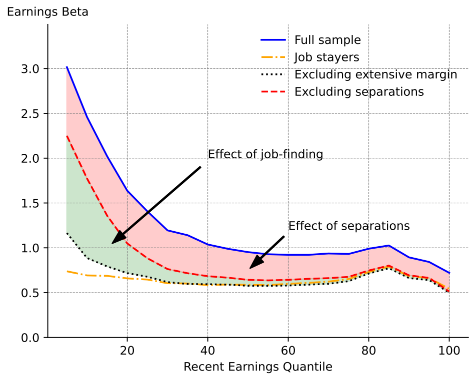

(b) Earnings beta, subsamples

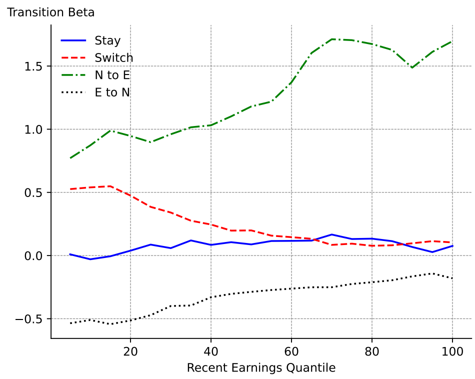

(c) Cyclicality of transition probabilities

Note: The *Left Panel* plots the coefficients $\beta^q_{Y,4}$ from Equation (2) (solid blue), by quantile. The shaded area represents 95 % confidence intervals. The dashed black line reproduces the GDP earnings beta for men aged 36-45, as reported by Guvenen et al. (2017). The *Right Panel* shows the estimates for the coefficient $\beta^q_{Y,4}$ from Equation (2) for two additional subsamples. The solid line utilizes the full sample, the red, dashed line excludes individuals who are employed in quarter t and non-employed in quarter $t+4$, the black, dotted line restricts the sample to those employed or unemployed in both periods t and $t+4$, the dash-dotted line restricts the sample to those employed by the same employer in both periods t and $t+4$. The area between the solid and dashed lines represents the change in the beta-coefficients upon the exclusion of individuals who separate. The green-shaded region marks the contribution of job-finding to the beta-coefficients. The *Bottom Panel* plots the coefficients $\beta^q_{Y,4,trans}$ in Equation (3). The solid blue line restricts the sample to job-stayers, the dashed red line to job switchers, the dash-dotted green line to job-finders and the dotted black line to separators. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The sample period is 1980-2019.

In order to investigate possible reasons for the stark heterogeneity in the procyclicality of earnings growth rates across the distribution, I proceed analogously to Figure 2, by restricting the sample to subgroups and investigating the effect on the regression coefficients in Equation (2). The right-hand panel in Figure 3 shows the results of this exercise. When restricting the sample to individuals who stay in the same job between quarters t and $t + 4$, the coefficients are close to homogeneous along the earnings distribution (yellow line). Earnings are procyclical for job-stayers, but their earnings grow only half as much as aggregate earnings.

Next, I reintroduce individuals who change employers or who are non-employed in both periods t and $t+4$. This exercise allows some labor market transitions, e.g. job switching, and non-employment in any of the periods between, but excluding t and $t + 4$. This raises the coefficients slightly towards the bottom of the income distribution. I interpret the difference between the beta coefficients estimated for the full sample (solid line) on the one hand and the ones estimated for this restricted sample on the other (dotted line) as the contribution of the extensive margin. The graph suggests that this contribution is considerably larger towards the bottom of the distribution than it is at the top.

The above findings are in line with Hoffmann and Malacrino (2019), who, using annual Italian data, split earnings growth into the contributions of employment time fluctuations and wage changes. They find that the cyclicality of both skewness and variance of the income growth distribution are almost entirely driven by extensive margin transitions. My findings also mirror those by Broer et al. (2020), who, using the same German data, show that the impact of monetary policy on monthly earnings growth is strongest at the low end of the income distribution, mainly due to extensive margin transitions.

Going beyond these established regularities, I decompose the extensive margin into the contributions by job-finders and job-separators presents a challenge. It is not possible to compute earnings growth rates for either group individually, as all job-finders, by definition, have zero earnings in period t, and all separators earn zero in period $t+4$. In order to still pin down which margin is responsible for the strong heterogeneity in earnings growth rates across the distribution, I restrict the sample in steps. First, I focus on all individuals who do **not** transition from employment in quarter t to non-employment in quarter $t + 4$ (i.e., I exclude separators) and re-estimate the regression in Equation (2). The resulting coefficients are plotted as the red, dashed line in the right-hand panel of Figure 3b. Across the distribution, the earnings betas for this subsample are lower than they are for the full sample. The shifts are similar, with coefficients decreasing by between 0.25 and 0.75 points. Therefore, the contribution of separations to the overall beta-estimates appears to be fairly homogeneous between the bottom and the top of the recent earnings distribution. Importantly, the earnings betas estimated for the described subsample are still steeply decreasing from close to 2.5 in the first decile to 0.5 around the median of the income distribution.

Crucially, the beta-coefficients estimated for the above subsample also identify the contribution of job-finders to the earnings beta. In the right-hand panel of Figure 3b, this is the shaded green area between the dotted black line, the beta-coefficients for the subsample of job-stayers, and the dashed red line, the coefficients when the estimation sample includes job-stayers and job-finders (but excludes job-separators). The difference between the two lines must be due to the addition of job-finders. In the first decile of the recent earnings distribution, the addition of job-finders increases the beta coefficient by more than one full percentage point, implying that there, job-finding drives a large share of the procyclicality in earnings. The effect decreases towards the median and is close to zero approaching the seventh decile.

From this decomposition, I conclude that the main reason for the observed heterogeneity in the procyclicality of earnings growth rates along the income distribution is earnings growth due to job-finding. This is in line with Figure 1, which shows that non-employment is considerably more prevalent towards the lower end of the income distribution. Next, I show that job-finding is procyclical at the bottom of the recent earnings distribution.

Similarly to Equation (2), I estimate a regression of transition probabilities on aggregate earnings:

$$\overline{trans}^{q}_{t,k} = \gamma + \beta^{q}_{Y,k,trans} \Delta Y_t + X_t + \varepsilon_t \tag{3}$$
$$\text{where } \overline{trans}^{q}_{t,k} = \frac{1}{N_q} \sum_i trans_{i,t+k} \forall i \in q$$

where $\overline{trans}^{q}_{t}$ is the average conditional transition probability between two states, for all individuals in quantile q at quarter t. I focus on four transition paths ($trans$): (i) staying with the same employer, (ii) switching to a new employer, (iii) employed to non-employed and (iv) non-employed to employed.

The bottom panel of Figure 3 plots the values of $\beta^{q}_{Y,4,trans}$ from Equation (3) across the recent earnings distribution. While job-staying is relatively acyclical, job-switching is moderately more procyclical at the low end of the distribution, compared to the top. Job-separations are countercyclical everywhere, but more so at the bottom. Crucially, a one percentage point increase in aggregate earnings growth leads to a one percentage point increase in the job-finding probability at the bottom of the distribution, but procyclicality increases along the distribution. Note that, despite this increasing procyclicality, changes in job-finding probabilities have heterogeneous impacts on earnings growth, because the stock of non-employed workers varies across the distribution (see Figure 1b).[15] Similarly, despite the countercyclical and heterogeneous nature of transitions into non-employment, their effect on the earnings betas is relatively homogeneous due to cross-quantile differences in employment incidence and earnings levels.

Together, these results explain the heterogeneous incidence of job-finding in driving earnings procyclicality. Non-employment is considerably more common at the low end of the recent earnings distribution. This is driven by higher separation and lower job-finding rates (right-hand panel in Figure 2). When aggregate earnings rise, job-finding increases across the distribution, driving large income gains at the bottom, but almost none at the top, as there are very few non-employed workers there.

In order to assess the consequences of these findings for welfare and policy, in the next section, I build a model that can reproduce the effects observed in the data, while accounting for important margins of insurance like savings and unemployment benefits.

3 Model

In order to conduct welfare analyses of policy proposals, I develop a model that mirrors the empirical earnings beta and its decomposition by generating an endogenous earnings distribution and endogenous, heterogeneous labor-market transitions across this distribution. The environment is a discrete-time, infinite-horizon directed-search framework in the spirit of Chaumont and Shi (2022), bringing together a unit mass of workers and a continuum of firms. I extend their model by allowing for individual and aggregate productivity shocks. Hence, the model presented here features ex-ante homogeneous agents, who differ along three dimensions ex-post: productivity, wealth and wages. In the model, job-finding probabilities and wages are endogenous outcomes of agents' decisions and can therefore be affected by labor market policies. Due to search frictions, the model features involuntary unemployment, endogenous to the aggregate state of the economy; and

[15] Appendix A.2 reports similar results in percent, as opposed to percentage points. At the bottom of the income distribution, job-finding is again more procyclical than separations.

due to borrowing constraints, households are unable to perfectly insure against this risk.

3.1 Environment

The consumers in the model are risk averse, receive utility from consumption and discount the future at rate β. Hence, their expected utilities are

$$U_t = \sum_{s=t}^{\infty} \beta^{s-t} u(c_s).$$

The utility function $u : \mathbb{R}_+ \to \mathbb{R}$ is twice differentiable, strictly increasing, strictly concave and $u'(0) = \infty$. All consumers can save in a risk-free asset a, which pays interest rate r. Borrowing is not possible.

Households can either be employed or unemployed, with the employed supplying one unit of labor inelastically and earning $(1-\tau)w$. The variable w represents the constant wage associated with their firm-worker match and τ is a labor income tax imposed by the government. In equilibrium, wages differ across workers, giving rise to an endogenous earnings distribution as in the data. The unemployed receive unemployment benefits $b(z)$, indexed by individual productivity, and enjoy a value of leisure l. Consequently, workers differ according to their asset holdings, their employment status, their productivity and their wages.

Each individual is endowed with productivity z_t, which evolves according to the Markov process

$$\log(z_t) = \rho_z \log(z_{t-1}) + \varepsilon_t$$
$$\varepsilon_t \sim N(0, \sigma_z)$$

Labor market search in the model is directed. Households, employed and unemployed, can direct their search at wage submarkets indexed by productivity and wealth. In each submarket, workers only meet firms willing to employ them at the submarket-specific wage rate. Consequently, if a worker meets a (new) firm, there is no need for wage bargaining, and they start their relationship at the beginning of the subsequent period. Wages are fixed for the duration of the match. The unemployed can search every period, while the employed can only search with probability $0 < \Lambda < 1$, capturing the fact that most of their time is spent working.

Firms are owned by risk-neutral investors, who maximize profits and discount the future at the risk-free interest rate. They post vacancies across wage submarkets.

The cost of posting a vacancy is κ and there is free entry. A firm-worker match produces zA according to the worker's productivity z and aggregate productivity A. Existing matches separate for two possible reasons: (i) exogenously with probability $\delta(z)$, or (ii) endogenously, if the worker finds a new job.

Submarkets are indexed by the worker's wealth in order to keep the problem tractable and less costly to solve computationally.[16] The indexation allows firms to predict the workers' job-finding probabilities on-the-job.[17] These affect the firm's expected match profits, which in turn influence which wage-submarket a firm enters. If worker wealth was unobservable to the firm, they would need to form expectations over the worker's wealth level given her productivity. This, in turn, would require them to know the distribution of wealth conditional on productivity. Furthermore, upon meeting, firms and workers would need to bargain over the wage, since match surpluses would differ depending on the workers' assets.[18] Both features would make the problem computationally more difficult.

Within each submarket, searchers and firms meet according to a matching function $M(S, V)$, where S and V represent the mass of searchers and vacancies within each submarket, respectively. The function is strictly increasing, concave in both arguments and exhibits constant returns to scale. If a vacant firm meets a worker, they form a match. The probability of a vacancy being filled can be expressed as $q(\theta) = M(S, V)/V = M(\frac{1}{\theta}, 1)$, where $\theta = \frac{V}{S}$ represents the submarket tightness. Likewise, the probability of a worker finding a match can be written as $\eta(\theta) = M(1, \theta)$. The vacancy filling probability $q(\theta)$ is decreasing in θ, while the job finding probability $\eta(\theta)$ is increasing in θ. If a vacant firm meets more than one worker, it randomly decides which one to hire. Submarkets that are not active because no worker or no firm visits them have $\theta = 0$.

The government finances unemployment benefits and bond issuance through the aforementioned tax on labor income, τ. It follows a budget rule, trading off

[16] If wealth was unobservable to the firm, workers in the model would have an incentive to signal their wealth to the firm (Chaumont and Shi, 2022). In equilibrium, asset-rich workers will choose to search for jobs in submarkets that offer high wages and low job-finding probabilities. Hence, they are less likely to leave their current match for a new one than wealth-poor workers. This, all else equal, is attractive to the firm, since it implies a longer match duration. Chaumont and Shi (2022) argue that workers who can signal to the firm any level of wealth less than or equal to their actual holdings, will choose to signal the true value. This, in turn, implies that submarkets would endogenously separate by the worker's wealth level.

[17] If there was no on-the-job search, implying that the probability of an endogenous quit was zero, workers' assets would not enter the firm's problem. I choose to allow workers to search in order for the model to be able to match the labor market flows I observe in the data.

[18] This problem is similar to Krusell et al. (2010), who solve a random search model with workers who can hold assets.

fluctuations in bond issuance with tax rate changes. Further, the government taxes all profits lump sum.[19]

A worker's state vector can be written as (L, z, x, w), where $L \in \{U, E\}$ represents her employment status, $x \in \mathcal{X} \equiv [0, \bar{x}] \subseteq \mathbb{R}_+$ represents her beginning of period wealth, $z \in Z \subseteq \mathbb{R}_+$ represents her idiosyncratic productivity realisation, and $w \in \mathcal{W} \subseteq [0, \overline{w}]$ is her wage, if she is employed.

The aggregate state of the economy can be expressed as $\psi = (A, \Omega)$, with $A \in \mathbb{R}_+$ representing the aggregate productivity state of the economy and $\Omega : \{U, E\} \times \mathcal{X} \times \mathcal{Z} \times \mathcal{W} \to [0, 1]$ denoting the distribution of agents across states. The aggregate state follows the law of motion $\psi' = Q(\psi)$.

Model timing: At the beginning of each period, the aggregate and individual productivity innovations are revealed. Subsequently, individuals who matched with a vacant firm in the previous period start their new jobs. In the next phase, matches produce and households consume and make their savings choices. In the last stage, firms post vacancies across submarkets and consumers choose the submarket in which to search. Lastly, separations and job-matching take place. In the notation that follows, the value functions of the workers are evaluated after production has taken place and wages were paid, but before the consumption/savings decision.

3.2 Firms

An unmatched firm can post a vacancy into a submarket indexed by the worker's productivity z, her cash-on-hand x and the posted wage w'. The value of a vacancy is

$$V_F(z, x, w'; \psi) = -\kappa + \frac{1}{1+r}\mathbb{E}\left[q(z, x, w'; \psi)J(z', x', w'; \psi') + (1 - q(z, x, w'; \psi))V_F'\right] \tag{4}$$

The vacancy posting cost is κ, $q(z, x, w'; \psi)$ represents the probability that the firm will meet a worker and r is the interest rate for investments between period t and $t + 1$. If the vacancy is not filled, the firm can post a new vacancy in the subsequent period. Because the firm's owners are risk neutral, they discount the future at the risk-free interest rate.

[19] An alternative specification follows Kaplan et al. (2014), where each agent receives a fraction of total profits each period, in relation to their share of total risk-free investment. In my model, using such an approach would require guessing a path for aggregate profits along the transition path to aggregate productivity shocks, which would make the problem computationally much more costly to solve.

A firm that meets a worker in a submarket starts producing output in the subsequent period. Its value, then, is

$$
\begin{aligned}
J(z,x,w;\psi) = Az - w + \frac{1}{1+r}\mathbb{E}\,[\delta(z')V_{new} \\
&+ (1-\delta(z'))\Lambda\eta(\theta(z,x,w';\psi))V_{new} \\
&+ (1-\delta(z'))(1-\Lambda\eta(\theta(z,x,w';\psi))J(z',x',w;\psi')]
\end{aligned}
\tag{5}
$$

The firm produces according to aggregate productivity A and the worker's productivity z. With probability $\delta(z')$, the match dissolves exogenously before the production phase in the next period. If this does not occur, the worker may find a new match on-the-job with probability $\eta(\theta(z,x,w';\psi))$, which depends on the worker's submarket choice w'. In both cases, the firm has the opportunity to post a new vacancy V_{new} in the following period. If the match does not separate, it continues into the next period at the same wage w.

Because entry into all submarkets is free, in equilibrium, the value of posting a vacancy is driven to zero, $V_F = 0$. Utilizing this, it is possible to rewrite Equation (4) as

$$
\frac{1}{1+r}\mathbb{E}\,[J(z',x',w';\psi')] = \frac{\kappa}{q(\theta(z,x,w';\psi))}.
$$

Intuitively, the expected value of a match must be equal to the vacancy posting cost, adjusted for the probability of meeting a worker, in each submarket. Since submarkets are separated along the worker's state variables z and x, the equation pins down the relationship between the wage and tightness **within** each submarket:

$$
\theta(z,x,w';\psi) = \begin{cases} q^{-1}\left(\frac{\kappa(1+r)}{\mathbb{E}[J(z',x',w';\psi')]}\right) & \text{if } \mathbb{E}\,[J(z',x',w';\psi')] \geq \kappa(1+r) \\ 0 & \text{otherwise} \end{cases}
\tag{6}
$$

The inequality constraint implies that firms will not post vacancies in submarkets where the present discounted value of forming a match is too low to cover the vacancy posting cost, even if a match was guaranteed ($q(\theta) = 1$). In these markets, tightness is zero and no matches are formed. To save space, in what follows, I suppress the dependencies of θ on current productivity z, current worker wealth x and the posted wage w'; instead of $\eta(\theta(z,x,w';\psi'))$, I use $\eta(w')$ to indicate that, conditional on the workers state variables, wage submarket choices dictate job-finding probabilities.

3.3 Worker Problems

3.3.1 Unemployed

The unemployed enter each period with state variables (i) productivity, z, and (ii) cash-on-hand, x, and maximize their expected discounted utilities. In recursive notation, their problem is

$$
\begin{aligned}
V^u(z,x;\psi) = \max_{a'} \quad & u(c) + l + \beta \max_{w'} \mathbb{E}\Bigg[\underbrace{\eta(w')V^e(z',x_n',w';\psi')}_{\text{Find Work}} \\
& \qquad + \underbrace{(1-\eta(w'))V^u(z',x_u';\psi')}_{\text{Stay Unemployed}}\Bigg] \qquad (7) \\
\text{subject to} \quad & x_u' = (1+r)a' + b(z') \\
& x_n' = (1+r)a' + (1-\tau)w' \\
& c + a' \leq x
\end{aligned}
$$

Using backward induction, households first solve which submarket w' to search in, taking as given the choice for consumption c and savings a'. As discussed in section 3.2, the free entry condition implies that, conditional on worker productivity z and wealth x, each submarket wage w' implies a unique submarket tightness, which pins down a submarket specific job-finding probability $\eta(w')$. When choosing in which submarket to search, workers trade off higher wages w', which promise higher continuation values V^w_{t+1} in case a match is formed, with higher job-finding probabilities. Note that unemployed households who successfully match to a firm cannot immediately separate exogenously. If an unemployed worker matches with a vacant firm, their cash-on-hand in the following period is x_n'. Otherwise, they remain unemployed, receiving benefits $b(z)$ and leisure utility l.

In the first stage, knowing their choice $w'(a')$, workers choose the optimal consumption in each period, subject to their beginning of period wealth x, their productivity z, and the borrowing constraint.

3.3.2 Employed

Like for the unemployed, an employed worker's problem consists of a consumption-savings choice and a choice of which wage submarket to search in. An employed worker enters the period with her individual productivity level, z, cash-on-hand, x, and wage, w. Recursively, her problem can be formulated as

$$\begin{aligned}
V^e(z,x,w;\psi) = \max_{a',c} \quad & u(c) + \\
& \beta \max_{w'} \mathbb{E}\Bigg[\underbrace{(1-\delta(z'))(1-\Lambda\eta(w'))V^e(z',x'_e,w;\psi')}_{\text{Stay in same job}} \\
& \quad + \underbrace{(1-\delta(z'))\Lambda\eta(w')V^e(z',x'_n,w';\psi')}_{\text{Switch jobs}} \\
& \quad + \underbrace{\delta(z')V^u(z',x'_u;\psi')}_{\text{Separations}}\Bigg] \qquad (8) \\
\text{subject to} \quad x'_u &= (1+r)a' + b(z') \\
x'_e &= (1+r)a' + (1-\tau)w \\
x'_n &= (1+r)a' + (1-\tau)w' \\
c + a' &\leq x
\end{aligned}$$

Note that the separation probability $\delta(z')$ governs separation from the current match, since newly found matches cannot immediately terminate.

Analogous to unemployed's problem, an employed worker's problem can be solved by backward induction. The employed first choose the wage submarket in which they want to search on-the-job, taking savings and consumption decision as given. Each period, the employed are only able to search with probability Λ. Further, with exogenous probability $\delta(z')$, an employed worker becomes unemployed at the beginning of the next period, irrespective of their search outcome.

The employed search in the same submarkets as the unemployed, conditional on their respective savings decisions and productivity. Hence, with probability $\eta(w')$, their search is successful and they move to a new wage w' starting in the next period (unless they get separated exogenously). If their job-search is unsuccessful, or they are unable to search, they continue working with their current employer.

Heterogeneous productivity levels affect job-search as follows: in equilibrium, workers who enjoy high levels of idiosyncratic productivity search in submarkets with higher job-finding probabilities and lower wages, relative to their productivities. The reasons are two-fold. First, low-productivity workers are subject to exogenously higher separation rates. This makes job search relatively less attractive for them, compared to highly productive workers. Second, because individual level productivity z changes over time, highly productive workers want to insure themselves by matching themselves to an employer. Since wages are fixed for the duration of a match, these workers stand to gain more from being employed at a wage that is low, *relative to their productivity*, than from waiting. These channels are not present in Chaumont and Shi (2022), but are crucial for matching the heterogeneous labor market transition probabilities discussed in Section 4.1.

Due to risk aversion, all households will try to smooth their consumption. Unemployed households with little wealth, whose income is likely to increase due to a positive job-finding probability and mean reversion in individual productivity, would like to borrow, but are prevented from doing so by the borrowing constraint. To move away from the constraint as quickly as possible, they search in submarkets that offer relatively low wages but high job-finding probabilities. At the same time, rich unemployed households search in markets with lower job-finding probabilities but higher wages. Due to their asset holdings, they are insured against the risk of not finding a job that is associated with higher wage submarkets. Both Eeckhout and Sepahsalari (2024) and Repele (2025) focus on this channel, but it is of second order here, as the impact of productivity on labor market search choices are much larger than those of wealth. The wage policy functions for the calibrated model are discussed in Section 6.

3.4 Government budget

The government collects income taxes and all firm profits, with which it finances a stock of risk free bonds:

$$\begin{aligned} B(1+r) + UI &= B' + \tau W + \Pi \qquad (9) \\ \text{where } W &= \int_{i\in\Omega} w_i \quad d\Omega \\ UI &= \int_{i\in\Omega} b_i I_i^U \quad d\Omega \end{aligned}$$

where B represents the stock of bonds at the beginning of period t, B' represents newly issued bonds, W are total earnings in the economy and UI is the total outlays of the unemployment insurance system.

The government follows the following budget rule, similar to Auclert and Mitman (2018) and Bayer et al. (2024):

$$\tau - \tau^* = \phi_\tau(B' - B^*) \qquad (10)$$

where τ^* and B^* are the steady state tax rate and bond issuance, respectively. The parameter $\phi_\tau > 0$ governs how elastic bond issuance is relative to tax changes. For $\phi_\tau = 0$, taxes are constant and the budget is balanced through bond issuance. As the parameter rises, taxes become more elastic.

3.5 Equilibrium

Recursive Equilibrium definition: A recursive equilibrium in this economy is a path for the interest rate r and the labor income tax τ, a set of household savings policy functions $g_a^E(z, x, w; \psi)$ and $g_a^U(z, x; \psi)$, a set of household wage choice functions $g_w^E(z, x, w; \psi)$ and $g_w^U(z, x; \psi)$ and submarket tightnesses $\theta(z, x, w', \psi)$, such that

- the household's asset and wage policy functions are consistent with the problems in Equation (7) and (8),
- the government budget in Equation (9) is balanced,
- the free entry condition implies submarket tightness according to Equation (6),
- asset markets clear, such that

$$B' = \int_{i \in \Omega} g_a^E(z, x, w; \psi) + g_a^U(z, x; \psi) \quad d\Omega$$

- the aggregate state of the economy ψ evolves according to $\psi' = Q(\psi)$

The equilibrium, formulated in this way, is difficult to compute, because the distribution of agents Ω is an infinite-dimensional object. However, the model permits a block recursive formulation, as discussed in Menzio and Shi (2010) and Menzio and Shi (2011), which decouples the labor search decision rules and submarket tightnesses from the distribution of agents Ω. The proposition and the corresponding proof follow Karahan and Rhee (2019), Herkenhoff (2019) and Birinci and See (2023). Importantly, block recursivity only materializes when r and τ are given, as these still depend on the distribution of agents across states, Ω.

Block Recursive Equilibrium (BRE) definition: A block recursive equilibrium is an equilibrium in which, given a path for the interest rate r and the labor income tax τ, the households' policy functions and submarket tightnesses only depend on the aggregate productivity state A, but not on the distribution of agents Ω.

Proposition *If i) utility function $u(\cdot)$ is strictly increasing, strictly concave, and satisfies the Inada conditions; ii) choice sets $\mathcal{W}$ and $\mathcal{A}$, and sets of exogenous productivity processes z and A are bounded; iii) the matching function M exhibits constant returns to scale; and iv) All policies are restricted to depend on the aggregate state only through aggregate match productivity, then there exists a unique BRE for this economy, given paths for r and τ.*

Proof See online appendix C

4 Calibration

I calibrate the steady state of the model to match a set of labor market moments I observe in my dataset. For this exercise I set aggregate productivity to $A = 1$. I use the Generalized Method of Moments to minimize the distance between the targeted moments in the model and the data. Tables 1-2 and Figure 4 summarize the model calibration.

The period length in the model is set to one quarter. I set the annual risk-free interest rate to 4%, resulting in a quarterly rate of $r = (1.04)^{\frac{1}{4}} - 1 = 0.00985$. To calibrate the bond supply B, I target the average liquid-asset to income ratio in Germany, which I obtain from the Household Finance and Consumption Survey. I utilize the survey's second wave, which was mainly conducted in 2014. Household asset holdings are categorized into liquid and illiquid following the approach by Kaplan et al. (2014), with slight modifications to the definitions used in their exercise: as household income, I classify the sum of self-reported labor earnings and social transfers, excluding private transfers and pension income, as this comes closest to the income measure I observe in the SIAB panel.[20] Further, I classify savings accounts as liquid savings, as my model is quarterly. The resulting average liquid-asset to income ratio for Germany is 2.67.

The utility function for workers is given by

$$u(c) = \frac{c^{1-\sigma}}{1-\sigma}$$

[20]To establish comparability between the HFCS and the SIAB panel, and reliability of the income variable in the HFCS, I calculate the median quarterly income in 2014 in both datasets. In the SIAB, the individual median income is €5700, while in the HFCS, quarterly household income is €4160. Note that the SIAB values are gross earnings, while the HFCS are net of tax.

where σ represents the coefficient of relative risk aversion. I set this parameter to 2, in line with standard calibrations in the literature.

Labor Market Following Den Haan et al. (2000), labor market matches in the model are formed according to the matching function

$$M(S,V) = \frac{SV}{(S^x + V^x)^{\frac{1}{x}}}$$

where S is a mass of searchers and V a mass of vacancies. This function guarantees that the job-finding and vacancy-filling probabilities lie within the interval $[0,1]$. I set the vacancy cost $\kappa = 0.53$, following Hagedorn and Manovskii (2008). For the exogenous quarterly separation probabilities, I assume the following functional form:

$$\delta(z) = \max\{0.01, \delta_0 * e^{\delta_1 z}\} \tag{11}$$

This allows for steadily decaying separations across the productivity distribution, with a minimum of one percent. The latter is necessary to ensure sufficient turnover in very high productivity states. I calibrate the parameters such that the model matches the average separation probability in the data, as well as the slope of separations across the distribution, resulting in $\delta_0 = 3.594$ and $\delta_1 = -5.21$.

Job-finding across the distribution is governed by the parameter of the matching function x and the value of leisure in unemployment, h. In the model, a higher value of x increases job-finding probabilities across the distribution, and thus I set $x = 0.721$ to match the average job-finding probability in the data. Conversely, leisure in unemployment is particularly relevant for the relative probability of job-finding of low- and high-earning households. I calibrate it to match the slope of the job-finding probability along the income distribution, resulting in $l = 0.067$.

Government In Germany, the replacement rate is 60 % of the last net-wage. In order to save on state variables (past wages), I model this benefit by indexing unemployment benefits to productivity such that $b(z) = \phi_B z$ with $\phi_B = 0.6$, similar to McKay and Reis (2016). I set the parameter governing the elasticity of the tax rate $\phi_\tau = 2$. Both Auclert and Mitman (2018) and Bayer et al. (2024) estimate this parameter for the US but find very different values: 0.024 and 2.796, respectively. The results reported below are not sensitive to this parameter.

Productivity process The household's productivity when employed is a combination of aggregate productivity in the economy and the consumer's individual productivity. Both evolve according to AR(1) processes. Individual productivity follows

$$\log(z_{t+1}) = \rho_z \log(z_t) + \epsilon_{t+1} \tag{12}$$

with persistence ρ_z and innovations $\varepsilon \sim N(0, \sigma_\epsilon)$. In the model, I discretize the process to a grid $\{z_1, ..., z_s\}$ with $s = 7$ using the Rouwenhorst method. To calibrate the parameter ρ_z, I first estimate the following earnings process for *the continuously employed* in both the model and in the data:

$$\log(y_{i,t+1}) = \rho_y \log(y_{i,t}) + e_{t+1}. \tag{13}$$

When calibrating, I exclude extensive margin transitions, which will be an endogenous outcome of the model. Then, I choose the parameter ρ_z in the productivity process such that the two match. In the data, I residualize earnings in the spirit of Heathcote et al. (2010) by regressing labor earnings on observable characteristics:

$$\log(earn_{i,t}) = \alpha + \gamma_1 X_{i,t} + \gamma_2 T_t + y_{i,t}$$

where $earn_{i,t}$ is an individual's quarterly earnings and $X_{i,t}$ contains an age polynomial, dummies for gender and education. I remove these aspects of the earnings process, because the model does not allow for heterogeneity along these dimensions. Using the variance/covariance structure between earnings observations of periods t and $t + s$, I back out the parameter $\rho_z = (\text{Cov}(y_t, y_{t+s})/\sigma_\epsilon^2)^{(1/(s-t))}$. As $s \to \infty$ (around $s = 50$), the latter equation converges to $\rho_{y,data} = 0.988$ (around $s = 20$) in the data. To calibrate the volatility of the earnings process, σ_ϵ, I target the ratio of the average gross earnings between the 25th and the 75th percentile, which is 2.885. In the model, I set $\rho_z = 0.978$ and $\sigma_\epsilon^2 = 0.082$.

Table 1: Externally Calibrated Parameters

Parameter	Description	Value
r	Interest Rate	0.00985
σ	Coefficient of relative risk aversion	2
ϕ_B	Replacement Rate	0.60
κ	Vacancy posting cost	0.53
ϕ_τ	Parameter in budget rule	2

Table 2: Internally Calibrated Parameters

Parameter	Description	Value	Target Moment	Data	Model
B	Bond supply	2.514	Liq. wealth to income	2.670	2.673
Λ	OTJ search Probability	0.374	Average switching prob.	0.027	0.028
x	Matching Funct parameter	0.703	Average job finding	0.134	0.142
l	Leisure in unemployment	0.068	Slope of job-finding	0.133	0.113
δ_0	Intercept of separation fctn	3.682	Average separations	0.040	0.037
δ_1	Slope of separation fctn	-5.405	Slope of separations	0.538	0.552
ρ_z	Persistence of prod. process	0.978	Cont. empl. earnings	0.988	0.982
σ_ϵ	Volatility of prod. process	0.082	Rel. earnings $P75/P25$	2.885	2.579

Note: Parameters calibrated via GMM to match moments from the SIAB (1980–2019) and HFCS (2014). "Slope" moments are regression coefficients of transition probabilities on recent earnings percentile. The separation function is $\delta(z) = \max 0.01, e^{\delta_0+\delta_1 z}$. Earnings persistence is estimated from the autocovariance structure for continuously employed workers. $P75/P25$ is the ratio of median earnings at the 75th to the 25th percentile.

Figure 4: Transition Probabilities - Model and Data

Note: This figure compares the transition probabilities in the data to their analogs in the model. Dashed lines represent probabilities obtained in the data, solid lines are obtained from the model. The blue lines represent job-finding probabilities, the yellow lines represent separation probabilities and the red lines represent job-switching probabilities. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The sample period is 1980-2019.

The calibrated model is able to replicate the transition probabilities in the data fairly well. Figure 4 shows the fit between the targeted labor market transition probabilities in the data (dashed) and their model counterparts (solid). The model lines are, perhaps surprisingly, more jagged than the empirical estimates. The reason for this is that while the model is almost entirely simulated non-stochastically, the mapping from each possible individual state (z, x, w) to a percentile of average earnings over the past 5-years can only be produced stochastically. Assigning an exact recent earnings percentile distribution to each state would require one to keep track of all possible paths leading to said specific state, which is impossible with large state-spaces. Instead, I simulate the economy for 21 quarters with 15,000,000 individuals to find the probability distribution across recent earnings percentiles for each point in the state space. Still, as the space of possible paths is vast ($14105^{20} \sim 10^{83}$), results are somewhat jagged.

The separation probabilities (yellow) in the model (solid) come close to those in the data (dashed). The model slightly undershoots the probability of moving from employment to unemployment in the first quintile. Job switching (red) was only targeted using one parameter, and still the model matches the observed switching patterns reasonably well, although it overshoots by about one percentage point. The inverted U-shape of the job switching probability is due to the fact that low productivity individuals have little incentive to search on-the-job, as they can only secure markedly higher wages if they experience a positive productivity shock. Job-finding probabilities (blue) are closely aligned between the data and the model, despite only being targeted by two parameters. In the next sub-section I describe how the model achieves this.

The model also performs well on several untargeted, yet important dimensions. First, the left panel in Figure 5 shows the earnings growth rates along the recent earnings distribution that result from the model, for the full sample, as well as excluding job-loss. The model results (solid) are very similar to the data (dashed), albeit more jagged towards the top of the distribution. The calibration only targets aggregate moments of the earnings process, but matching the heterogeneity across the distribution is a success of the model.

Figure 5: Steady state earnings growth rates

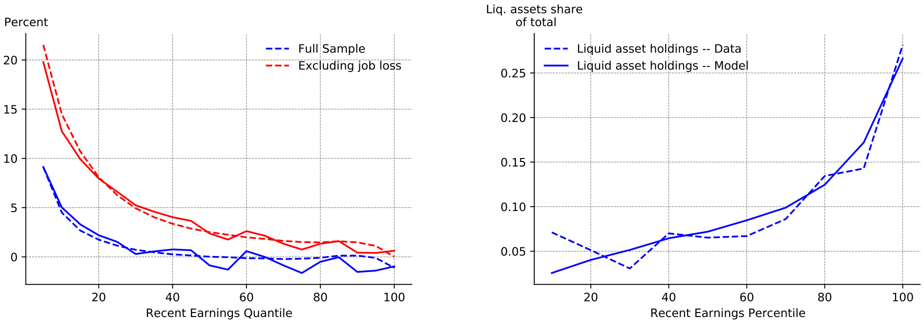

(a) Earnings growth (b) Liquid asset holdings

Note: The *Left Panel* shows a comparison between quarterly growth-rates implied by the model and the same statistic measured in the data. Dashed lines represent the data, solid lines represent the model's output. The blue line represents earnings growth for the all workers, the red line excludes job-loss. The *Right Panel* shows liquid asset holdings in the model (solid) and the data (dashed). Liquid asset holdings are calculated using data from the Households Finance and Consumption Survey (HFCS), for details, see the text. Quantile specific holdings are displayed relative to total holdings across quantiles. Note that recent earnings in the HFCS are calculated as yearly earnings reported in the survey, as the data cannot be matched to German administrative labor market data.

The right panel of Figure 5 shows liquid asset holdings in the model (solid) and compares it to liquid assets reported in the HFCS. It is not possible to sort households in the HFCS into 5-year past earnings bins, as the data is limited to the time period of the survey. Instead, I sort individuals according to their current, self-reported earnings. Barring these limitations, data and model line up remarkably well in terms of liquid asset holdings. The calibration only targets the aggregate liquid asset-to-income ratio, implying that matching asset holdings across the earnings distribution is an achievement of the model. Appendix B compares the earnings growth rates along the intensive margin as well as the distribution of marginal propensities to consume between the model and the data. The model is able to reproduce these relatively closely, as well.

4.1 Policy functions

As shown in Figure 4, job-finding probabilities are strongly increasing along the income distribution, in both the model and the data. The reason for the model's success along this dimension is illustrated in Figure 6. The left-hand panel shows the wage search policy functions of the unemployed for different levels of cash-on-hand (x-axis), as well as idiosyncratic productivity levels (colors) *relative to* *their productivity realization*. The right panel shows the corresponding job-finding probabilities. Unemployed individuals with high productivity realizations (green) search in submarkets that offer lower wages, relative to their current productivities, which leads them to higher job-finding probabilities. These workers have a strong incentive to become matched: if they find employment, their wages are fixed, even if their productivity realizations fall in the near future. Note that, for highly productive workers, while wages are low relative to their productivity realizations, wage levels are still higher than wages for unproductive workers.

At the lower end of the productivity spectrum (blue), the unemployed have less of an incentive to find employment. Their outside option (unemployment benefits) is attractive, relative to (i) the wage they could secure in the market and (ii) the high separation probabilities they face if they find employment. Furthermore, their productivity may rise in the future, in which case being tied to a job is disadvantageous, since matched individuals can only search for new jobs with probability $\Lambda < 1$. Together, these forces imply a strongly increasing job-finding probability along the recent earnings distribution seen in Figure 4.

Wealth (or cash-on-hand) has the opposite effect to productivity: wealthier workers search for higher wages and have lower job-finding probabilities, conditional on productivity. This is because individuals with relatively little wealth, wish to find work, so they can increase their consumption. In order to achieve this, they trade off high wages in favor of higher job-finding probabilities. As they get richer, the unemployed are more able to insure against the risk of not finding a job and search for high wages. This effect features prominently in Chaumont and Shi (2022) and Repele (2025). Additionally, here, the marginal consumption benefit of finding a job, $u'(c) * w$, falls as wealth increases, but the marginal loss of leisure (h) is constant. Hence, conditional on productivity, a higher level of wealth decreases the desire to find work even beyond what the pure precautionary motive would imply.

Figure 6 thus shows that the inclusion of individual productivity shocks is crucial for enabling the model to produce increasing job-finding rates along the income distribution. A model without such productivity differences would predict lower job-finding probabilities for higher earners, due to the positive effect of earnings on wealth. This, however, is counterfactual to the data.

Figure 6: Submarket choice of the unemployed

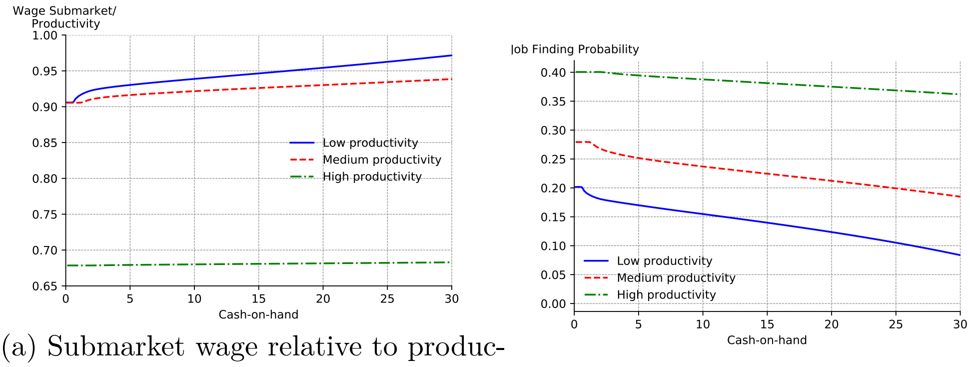

(a) Submarket wage relative to productivity

(b) Job-finding probability

Note: The *Left Panel* shows the wage submarket choices for workers with high, medium and low-productivity realizations. The x-axis represents their cash-on-hand state and the y-axis represents their wage choice relative to their current productivity realization. The *Right Panel* shows the associated job-finding probabilities.

The mechanism for employed workers who search on the job is analogous. Higher productivity realizations make workers search for jobs at wages that are low, relative to their productivities, but offer high job-finding probabilities. Wealth, on the other hand, makes individuals search for jobs that pay higher wages but offer lower job-finding probabilities. As for the unemployed, the productivity channel dominates the wealth effect. Additionally, employed workers take their current wages into account, never searching in wage submarkets that offer lower pay.

4.2 Aggregate shocks in the model

To estimate earnings betas in the model, I re-introduce aggregate shocks and solve for the economy's response along a series of such shocks over time. The dynamic programming problem in an economy with aggregate risk is difficult to solve, as the aggregate state of the economy $\psi = (\Omega, A)$ is an infinite dimensional object. I rely on the solution method proposed by Boppart et al. (2018) (BKM) to solve this problem. The approach employs the impulse response to a single unexpected shock in aggregate productivity as a first-order approximation to the global solution of a model in which aggregate productivity fluctuates constantly. In order for this method to provide accurate results, the model's policy functions, and hence aggregate outcomes, need to be approximately linear in the aggregate productivity shock. I show that this is the case in Appendix B.1.

Along the transition path, I assume that the government balances its budget according to the fiscal rule in Equation (10) by varying both the labor-income tax rate τ_t and the bond issuance B_t. Each period, the real interest rate r_t adjusts to clear the asset market. I solve for the path of the tax rate τ_t and the interest rate r_t using the sequence space Jacobian method proposed by Auclert et al. (2021).[21]

Importantly, in addition to economic aggregates, the BKM approach allows me to characterize the behavior of sub-aggregates, such as the earnings growth rates across different quantiles.

5 Results

5.1 Earnings procyclicality in the model

Using impulse responses obtained as described above, I linearly approximate the model's response to any aggregate productivity shock of the same sign. The aggregate shocks are drawn from a normal distribution with volatility $\sigma_{agg} = 1\%$ and persistence $\rho_{agg} = 0.9$. The impulses are obtained non-stochastically, except for the assignment into recent earnings quantiles. As described in Section 4, obtaining the recent earnings distribution within each state in the model's distribution Ω over time is computationally very costly. I stochastically simulate the paths of 15 million individuals such that in each period, I obtain the distribution $\Omega : \{U, E\} \times \mathcal{X} \times \mathcal{Z} \times \mathcal{W} \times T \times q \to [0, 1]$, where q represents the quantile in the recent earnings distribution. From the simulation, I construct the model-generated analog to the dataset discussed in Section 2 and use it to estimate Equation (2), the gross-earnings betas, in the model. To account for simulation uncertainty, I perform this step three times and plot the range of estimates in the figures.

Figure 7 shows the result of this exercise.[22] The solid lines represent the model generated earnings betas, which were not targeted in the calibration. The coefficients imply that earnings at the bottom of the recent earnings distribution are about four times as procyclical as those around the median, which is in line with the data. The model replicates the steeply decreasing earnings betas, observed in the data, implying that earnings growth is more procyclical for individuals at the bottom of the income distribution. In contrast to the pattern in Figure 3b, the model does not replicate the flattening out of the empirical earnings betas beyond the median. This is because, in the model, wages are fixed for the duration

[21] For a discussion of the computation, see online appendix D.

[22] Appendix B.2 reports the procyclicality estimates for the transition probabilities.

of their matches. Thus, while earnings of individuals employed in t and $t+4$ are slightly procyclical, they undershoot the data. As job-staying is the most common transition path for individuals at the top of the earnings distribution, and wages in continuous matches are fixed, this drives the correlation between aggregate and quantile specific earnings growth to zero. Hence, the model implies that earnings procyclicality decrease towards zero at the top of the income distribution.

Figure 7: Model implied earnings betas

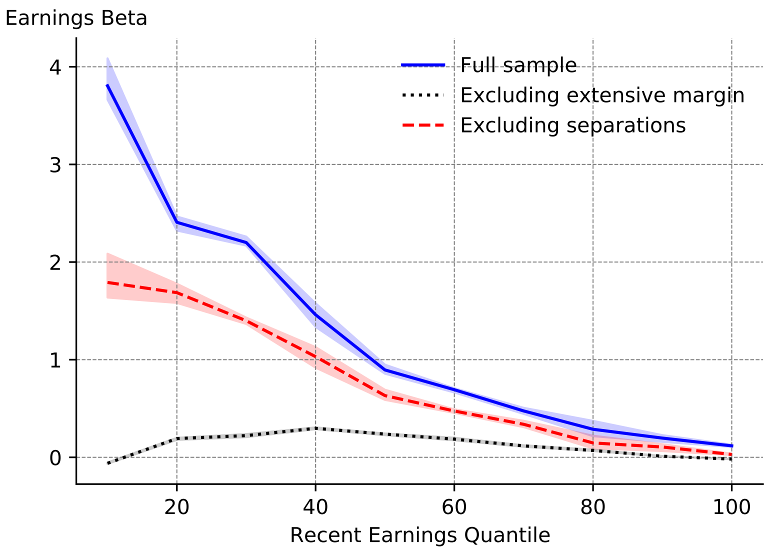

Note: The figure shows the results for the regression coefficients in Equation (2) in the model. The blue lines represent results for the full sample, the red lines represent results for the subsample excluding job-separators, the black line plots the coefficients for individuals in employment in quarters t and $t+4$. The results are obtained through 10 simulations with 3 million individuals, the shaded areas show the range of estimates across simulations, the point estimates are the averages across simulations.

When I restrict the sample by excluding individuals who are employed in quarter t and non-employed in quarter $t+4$ the result is the red dashed line. As before, I interpret the difference between the blue and red lines as the contribution of separations to the procyclicality of earnings growth. While this difference is fairly constant in the data, the model implies that it decreases towards the top of the income distribution. Still, along the entire distribution, job-finding is the main driver behind heterogeneity in earnings betas.

The model is able to successfully generate the patterns observed in the data through the mechanism outlined in Figure 8. In the top left panel the blue solid line represents the steady state submarket wage choices of highly productive unemployed workers. The top right panel shows the resulting job-finding probabilities in the same color. In response to a negative productivity shock, firm value decreases, all else equal. This leads to fewer firms entering into all submarkets. Consequently, workers would face lower job-finding probabilities if they followed their steady-state wage submarket choices. In the top right panel, the black dotted line shows the job-finding probabilities in a partial equilibrium exercise, holding worker policies constant at their steady state values. In this hypothetical, workers not adjusting their submarket choices results in a two percentage point decrease in job-finding probabilities. To counteract this decrease, after the shock hits, productive workers choose to search in submarkets that offer lower wages but job-finding probabilities close to the steady state. This is exemplified by the red dashed line, which plots the equilibrium quantities for wage and job-finding when workers and firms make optimal choices.

The bottom right panel in Figure 8 shows that unemployed low-productivity workers do less to counteract the fall in their job-finding probabilities. The red and black lines lie much more closely together than they do for high-productivity individuals. Thus, job-finding probabilities in the model are much more procyclical at the bottom.

Figure 8: Model mechanism

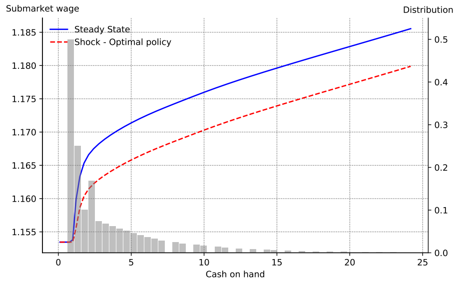

(a) Wage policy, high productivity

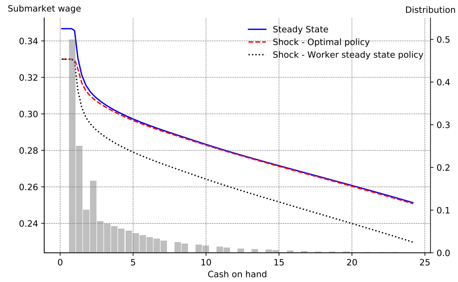

(b) Job-finding, high productivity

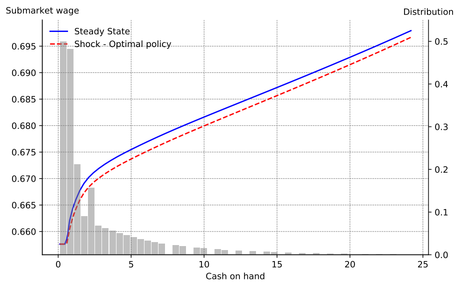

(c) Wage policy, low productivity

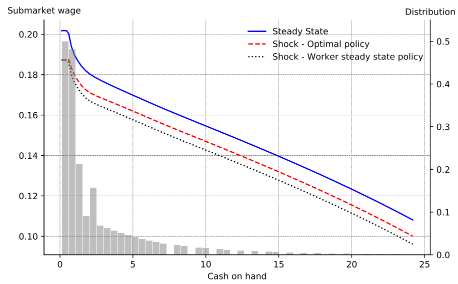

(d) Job-finding, low productivity

Note: The *Top Panels* show the wage choices (left) and the job-finding probabilities (right) of high-productivity workers. The solid lines represent steady state outcomes, the dashed lines represent outcomes in the first period after the realization of a negative productivity shock, the dotted line represents the job-finding probabilities after a negative productivity shock when worker choices are held fixed at their steady state values. The *Bottom Panels* contain the same information for low-productivity individuals.

As productivity differences translate into differences in recent earnings, the choices outlined above lead to job-finding, and hence average earnings, to decrease most at the bottom of the distribution. As mentioned in Section 2, procyclical job-finding probabilities are a key ingredient in the heterogeneous procyclicality of earnings growth. Figure 20 shows that the model actually overstates the strength of this relationship relative to the data, likely leading to the overestimation of the earnings beta in Figure 7.

The second contributing factor for the heterogeneous procyclicality of earnings growth in the data is the heterogeneous incidence of unemployment across the distribution. As Figure 17 in the appendix shows, the model reproduces this relationship very well.

6 Countercyclical hiring subsidies

Using the calibrated model outlined in section 3, I conduct a policy experiment: the introduction of countercyclical hiring subsidies as a business cycle stabilizer. As described above, one of the driving forces behind the heterogeneous procyclicality in earnings are procyclical job-finding probabilities. These concentrate the labor market effects of recessions at the bottom of the distribution – where unemployment incidence is highest. Hiring subsidies directly address this channel by incentivizing firms to keep posting vacancies, allowing job-finding to remain closer to its steady state level. Thus, such a policy insures the poorest, high non-employment-risk workers by lowering the risk caused by a decrease in vacancy posting, which they cannot counteract by adjusting their search decisions.

I implement this policy by letting the government award (or tax) firms the lump-sum transfer $v(A)$, indexed by aggregate productivity A, with the following functional form:

$$v(A) = \vartheta(A - \overline{A})$$

where the multiplier ϑ is calibrated such that it minimizes the variance of vacancy posting over the business cycle. This requires $v(0.99) = 0.0821$, implying a hiring subsidy of close to 660 euros upon impact of a negative 1% productivity shock. As aggregate match productivity returns to its steady state value, the hiring subsidy decreases back to zero. The process is analogous upon the arrival of positive productivity shocks.

I calculate the welfare effects by estimating the percentage increase in consumption necessary in order to make agents indifferent between facing a recession with and without countercyclical hiring subsidies and, following Krusell et al. (2010):

$$E_0 \left[\sum_{t=0}^{\infty} \beta^{\mathrm{t}} u\left((1+\lambda)c_{i,t}\right) \right] = E_0 \left[\sum_{t=0}^{\infty} \beta^{t} u\left(\bar{c}_{i,t}\right) \right] \tag{14}$$

where $c_{i,t}$ is each agent's consumption in the baseline economy and $\bar{c}_{i,t}$ is their consumption in the counterfactual economy. I evaluate Equation (14) in period 0 after an unexpected, one-time negative productivity shock, in both economies. The value of λ indicates how much additional consumption each agent would require to be indifferent between the baseline and the counterfactual economy, as opposed to the baseline economy. Finally, I average the values of λ over all agents in the model to arrive at the utilitarian consumption equivalent variation of the counterfactual economy.

Figure 9 shows the transition paths for economies with and without hiring subsidies in response to a negative productivity shock of 1%. The left-hand panel displays the percentage change in the number of vacancies created. In the baseline economy, vacancy creation contracts by close to 6 percent. With the subsidy, this fall is almost entirely eliminated: vacancy creation rises marginally on impact, before falling by less than 0.2 percent.

Figure 9: Effect of hiring subsidies

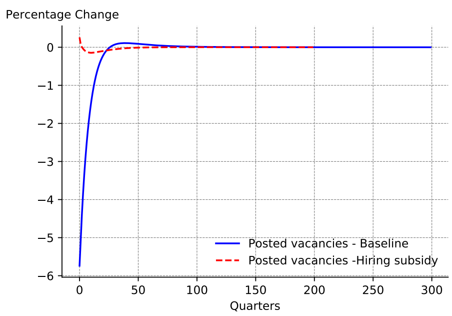

(a) Posted Vacancies

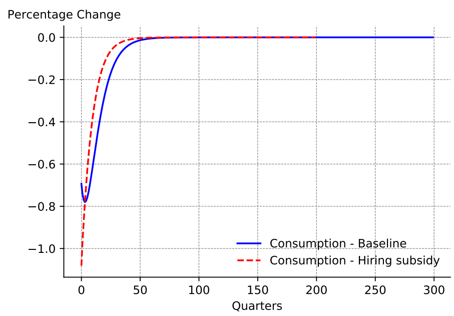

(b) Consumption response to -1% shock

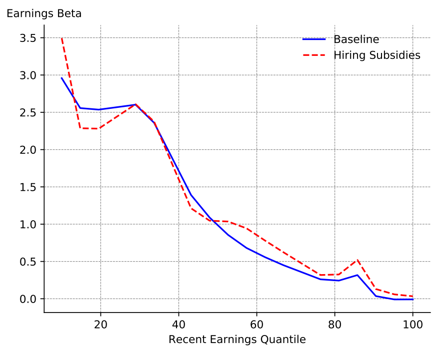

(c) Earnings betas with hiring subsidy

Note: The *Left Panel* shows the response of vacancy posting from firms entering low-productivity submarkets to a negative aggregate productivity shock of 1 percent in the baseline economy (solid) and the economy with hiring subsidies (dashed). The *Right Panel* shows the response of aggregate consumption to a negative productivity shock in the two economies.

The right-hand panel of Figure 9 shows the response of aggregate consumption. The initial impact of a negative productivity shock is now larger than in the baseline economy. However, the economy with hiring subsidies rebounds faster. The stronger initial decline is driven by the finance a larger deficit with the hiring subsidy. This leads to a decrease in overall consumption, but especially at the top.

Welfare, as measured by Equation (14), immediately after the realization of a negative, unexpected, transitory shock to aggregate match productivity, increases. To reach the same level of aggregate welfare in the baseline economy, each worker would need to be given an increase of 0.06% in their consumption, forever. To put this number into perspective, I calculate the welfare cost of a recession in both the baseline economy and the economy with a hiring subsidy, using the consumption equivalent variation between the steady state economy and the first period after a negative productivity shock. If given a choice, workers who are faced with a recession, on average, would need to be given 0.2% more consumption in the baseline economy each period in order to be equally well off as those in the steady state. The same number in the economy with a hiring subsidy is 0.12%. The hiring subsidy thus attenuates the welfare costs of recessions by about 30 percent. The relatively small welfare impact of transitory recessions is driven by the fact that, for employed workers, who make up the vast majority of the economy, separation rates are exogenous and wages are fixed; hence, any recession has little impact on them except through aggregate variables like the tax and the interest rate.

The aggregate figure hides significant heterogeneity, however. As shown in Table 3, for low productivity workers, $\lambda = 0.17\%$, while, for high productivity individuals, the welfare effect is negative. Low productivity workers' job-finding probabilities are most affected by the subsidies, as the size of the subsidy relative to wages is largest for them. Hence, the largest welfare gains accrue to the workers at the low end of the productivity distribution. Furthermore, the benefits of the hiring subsidy are concentrated among the unemployed, as it considerably reduces the negative consumption effects of recessions.

Table 3: Welfare change by Type

Type		Overall	Unemployed	Employed
Low productivity	CE (%)	0.179	0.205	0.114
	Share		71.7%	28.3%
Medium productivity	CE (%)	0.052	0.179	0.042
	Share		7.8%	92.2%
High productivity	CE (%)	-0.014	0.101	-0.018
	Share		3.1%	96.9%

7 Conclusion

In this paper, I show that earnings growth is more procyclical at the low end of the income distribution due to extensive margin transitions over the business cycle. Using administrative microdata from Germany, I find that as aggregate earnings growth rises by 1 percentage point, earnings growth in the first decile of the income distribution rises by three times as much. At the top of the distribution, the comovement is less than one-for-one. Decomposing the effect, I show that at the bottom of the distribution, transitions from unemployment to employment account for half of the procyclicality, whereas in the last decile, they account for almost nothing.

In order to speak to the welfare consequences of these findings, I build a heterogeneous agent macroeconomic model with idiosyncratic and aggregate risk. The model features a frictional labor market in which workers direct their search at different submarkets, based on their characteristics. They endogenously trade off higher job-finding probabilities with higher wages. Highly productive workers find work faster by choosing submarkets which offer low wages, relative to productivity. The opposite is true for less productive workers. In this way, the model reproduces heterogeneous job-finding probabilities along the income distribution. I calibrate the model to match key features of the German labor market, such as heterogeneous labor market transition rates. Importantly, I leave earnings growth rates untargeted.

The model produces heterogeneity in the procyclicality of earnings growth similar to the data. The reason for these results are the lower job-finding probabilities for low productivity workers in steady state, combined with procyclical vacancy posting by firms. The (untargeted) slope of earnings procyclicality along the income distribution produced by the model is somewhat steeper than what I find in the data, which leads to an underprediction of procyclicality at the top of the distribution.

I introduce a countercyclical hiring subsidy into the model, to gauge the welfare effects of addressing the strong procyclicality of job-finding. Countercyclical hiring subsidies, which rise in recessions and fall in booms, reduce the cyclicality of vacancy posting by firms. This, in turn, reduces the impact of aggregate productivity fluctuations on the unemployed. After the realization of a negative productivity shock, workers in the baseline economy would require an increase of 0.06 percent in their consumption in order for aggregate welfare to be the same as in an economy with such hiring subsidies, implying an increase in overall welfare.

To assess the magnitude of this increase, I calculate the welfare impact of a single productivity shock in the baseline economy, computing how much consumption workers would be willing to forego each period to avoid the recessions. According to this measure, a negative productivity shock reduces welfare by 0.2 percent, implying that hiring subsidies lessen the welfare impact of recessions by 30 percent.

Importantly, my results hinge on the endogenous nature of the earnings distribution, and the endogenous responses of workers' job finding choices to aggregate shocks. Models that don't feature these channels are bound to struggle to arrive at the correct policy conclusion when investigating policies which affect workers labor market choices. Therefore, I argue that a model that matches the heterogeneous procyclicality of earnings growth, is crucial for understanding the impact of such labor market policies.

Acknowledgements

I am indebted to Kurt Mitman and Per Krusell for many discussions and much advice. I also wish to thank Mattias Almgren, Gualtiero Azzalini, Agneta Berge, Tobias Broer, Mitch Downey, John Hassler, Stefan Hinkelmann, Alexandre Kohlhas, Markus Kondziella, Kieran Larkin, Arash Nekoei, Josef Sigurdsson, Xueping Sun, and participants in the IIES Macro Group and the IIES Brownbag series. This paper uses data from the Eurosystem Household Finance and Consumption Survey. The results published and the related observations and analysis may not correspond to results or analysis of the data producers.

References

Amberg, N., T. Jansson, M. Klein, and A. R. Picco (2022). Five facts about the distributional income effects of monetary policy shocks. *American Economic Review: Insights 4*(3), 289–304.

Andersen, A. L., N. Johannesen, M. Jørgensen, and J.-L. Peydró (2023). Monetary policy and inequality. *The Journal of Finance 78*(5), 2945–2989.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2021). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica 89*(5), 2375–2408.

Auclert, A. and K. Mitman (2018). Consumer bankruptcy as aggregate demand management. In *Society for Economic Dynamics Meeting Papers*.

Bayer, C., B. Born, and R. Luetticke (2024). Shocks, frictions, and inequality in us business cycles. *American Economic Review 114*(5), 1211–1247.

Birinci, S. and K. See (2023). Labor market responses to unemployment insurance: The role of heterogeneity. *American Economic Journal: Macroeconomics 15*(3), 388–430.

Boppart, T., P. Krusell, and K. Mitman (2018). Exploiting mit shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control 89*, 68–92.

Broer, T., J. Kramer, and K. Mitman (2020). The curious incidence of shocks along the income distribution. Technical report, Discussion paper, IIES, Stockholm.

Busch, C., D. Domeij, F. Guvenen, and R. Madera (2022). Skewed idiosyncratic income risk over the business cycle: Sources and insurance. *American Economic Journal: Macroeconomics 14*(2), 207–242.

Card, D., J. Heining, and P. Kline (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics 128*(3), 967–1015.

Carroll, C. D. (2006). The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics letters 91*(3), 312–320.

Chaumont, G. and S. Shi (2022). Wealth accumulation, on-the-job search and inequality. *Journal of Monetary Economics 128*, 51–71.

Costain, J. S. and M. Reiter (2005). Stabilization versus insurance: welfare effects of procyclical taxation under incomplete markets. *Universitat Pompeu Fabra Economics Working Paper* (890).

Dany-Knedlik, G., A. Kriwoluzky, and S. Pasch (2021). Income business cycles.

De Nardi, M., G. Fella, M. Knoef, G. Paz-Pardo, and R. Van Ooijen (2021). Family and government insurance: Wage, earnings, and income risks in the netherlands and the us. *Journal of Public Economics 193*, 104327.

Den Haan, W. J., G. Ramey, and J. Watson (2000). Job destruction and propagation of shocks. *American economic review 90*(3), 482–498.

Deutscher Bundestag (2025). Schriftliche fragen mit den in der woche vom 27. januar 2025 eingegangenen antworten der bundesregierung: Frage 34 des abgeordneten konstantin kuhle (fdp). Bundestagsdrucksache 20/14810. Antwort durch den Parlamentarischen Staatssekretär Johann Saathoff.

Dustmann, C., J. Ludsteck, and U. Schönberg (2009). Revisiting the german wage structure. *The Quarterly journal of economics 124*(2), 843–881.

Eeckhout, J. and A. Sepahsalari (2024). The effect of wealth on worker productivity. *Review of Economic Studies 91*(3), 1584–1633.

Frodermann, C., A. Ganzer, A. Schmucker, P. Vom Berge, et al. (2021). Sample of integrated labour market biographies regional file (siab-r) 1975-2019. Technical report, Institut für Arbeitsmarkt-und Berufsforschung (IAB), Nürnberg.

Gregory, V., G. Menzio, and D. Wiczer (2025). The alpha beta gamma of the labor market. *Journal of Monetary Economics 150*, 103695.

Guvenen, F., F. Karahan, S. Ozkan, and J. Song (2021). What do data on millions of us workers reveal about lifecycle earnings dynamics? *Econometrica 89*(5), 2303–2339.

Guvenen, F., S. Ozkan, and J. Song (2014). The nature of countercyclical income risk. *Journal of Political Economy 122*(3), 621–660.

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. and I. Manovskii (2008). The cyclical behavior of equilibrium unemployment and vacancies revisited. *American Economic Review 98*(4), 1692–1706.

Halvorsen, E., H. Holter, K. Storesletten, S. Ozkan, et al. (2019). Dissecting idiosyncratic income risk. In *2019 Meeting Papers*, Number 1337. Society for Economic Dynamics.

Harmenberg, K. and H. H. Sievertsen (2017). The labor-market origins of cyclical income risk.

Heathcote, J., F. Perri, and G. L. Violante (2010). Unequal we stand: An empirical analysis of economic inequality in the united states, 1967–2006. *Review of Economic dynamics 13*(1), 15–51.

Herkenhoff, K. F. (2019). The impact of consumer credit access on unemployment. *The Review of Economic Studies 86*(6), 2605–2642.

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 (2021). The transmission of monetary policy under the microscope. *Journal of Political Economy 129*(10), 2861–2904.

Kaplan, G., G. L. Violante, and J. Weidner (2014). The wealthy hand-to-mouth. Technical report, National Bureau of Economic Research.

Karahan, F. and S. Rhee (2019). Geographic reallocation and unemployment during the great recession: The role of the housing bust. *Journal of Economic Dynamics and Control 100*, 47–69.

Kitao, S., A. Şahin, and J. Song (2011). Hiring subsidies, job creation and job destruction. *Economics Letters 113*(3), 248–251.

Krueger, D., K. Mitman, and F. Perri (2016). Macroeconomics and household heterogeneity. In *Handbook of Macroeconomics*, Volume 2, pp. 843–921. Elsevier.

Krusell, P., T. Mukoyama, and A. Şahin (2010). Labour-market matching with precautionary savings and aggregate fluctuations. *The Review of Economic Studies 77*(4), 1477–1507.

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

Menzio, G. and S. Shi (2010). Block recursive equilibria for stochastic models of search on the job. *Journal of Economic Theory 145*(4), 1453–1494.

Menzio, G. and S. Shi (2011). Efficient search on the job and the business cycle. *Journal of Political Economy 119*(3), 468–510.

Repele, A. (2025). Wealth sorting and cyclical employment risk.

Statistisches Bundesamt (2026). Beschäftigte im öffentlichen dienst nach geschlecht. Accessed: 2026-04-20.

Online Appendix

A Robustness

A.1 GDP

For robustness, I construct a second aggregate earnings growth measure using GDP. Growth is measured as the quarter-on-quarter log difference in real GDP between periods t and $t+4$:[23]

$$\Delta Y_{t,k}^{GDP} = \log(GDP_{t+k}) - \log(GDP_t)$$

I then substitute $\Delta Y_{t,k}^{GDP}$ into equation (2) and estimate the beta coefficients. Figure 10 reports the results from this exercise. While the beta-coefficients are shifted downwards, implying a lower cyclicality of quantile-specific earnings relative to GDP, the shape of the curve remains similar. Earnings at the bottom of the income distribution are still about three-times as procyclical as those at the top.

Figure 10: Earnings beta with GDP

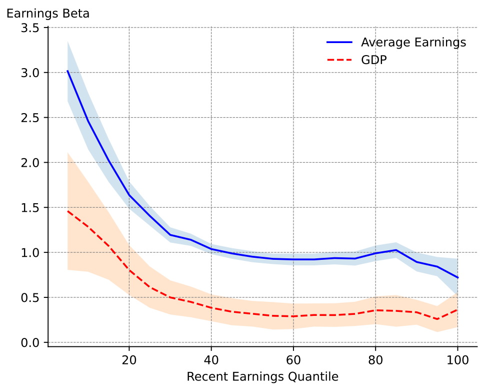

Note: This figure shows the coefficient $\beta_{Y,k}^{q}$ from Equation (2) by quantile, estimated using growth in aggregate average earnings (blue) GDP (red), respectively. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The sample period is 1980-2014.

[23] Deflated, seasonally and calendar adjusted GDP is obtained from the German Statistical Office. Due to German reunification in 1990, there only exist two separate time series for (i) West Germany and (ii) today's Germany, but both contain values for 1991. Consequently, I normalize both series by the GDP values for the first quarter of 1991 and append them.

A.2 Transition betas in logs

Figure 3c reports the percentage point changes in job-finding probabilities in response to aggregate earnings changes. However, since the unconditional transition probabilities are heterogeneous along the earnings distribution, as shown in Figure 2b, this appendix reports the percentage changes in transition probabilities in response to aggregate fluctuations.

I estimate these percentage transition betas using equation 3, but instead of the raw transition probabilities, I substitute their logged counterparts. In this exercise, the coefficient $\beta^{q}_{Y,k,trans}$ measures the elasticity of each labor market transition probability with respect to aggregate earnings.

Figure 11 displays the results. The elasticity of staying w.r.t. aggregate earnings changes is close to zero across the distribution. In contrast, both job-finding (green) and separations (black) are very strongly related to aggregate earnings, albeit in opposite directions. A one percentage point increase in aggregate earnings growth leads to a 4 percent increase in job-finding across the distribution. Separations have the opposite sign and are similarly homogeneous, in percentage terms.

Figure 11: Transition betas in logs

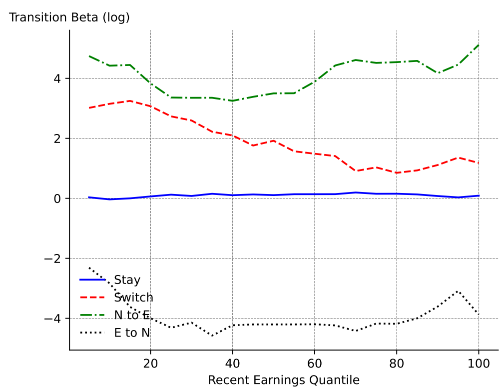

Note: This figure shows the coefficient $\beta^{q}_{Y,4,trans}$ from Equation (3) by quantile, for the log of the transition probability. The solid blue line restricts the sample to job-stayers, the dashed red line to job switchers, the dash-dotted green line to job-finders and the dotted black line to separators. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The sample period is 1980-2019.

Job-switching displays a decreasing pattern: in response to a one percentage point increase in aggregate earnings growth, the switching probability increases by around three percent at the low end of the permanent income distribution, but by around one percent at the top.

A.3 Age and Gender

In the graphs presented so far, all quantiles are constructed within age-gender bins. However, it is of course of interest how the earnings betas differ along those dimensions, as well. Figure 12, therefore, shows the estimates of β_Y^q in Equation (2) for quantiles based on recent earnings separately for each age group (left panel) and each gender (right panel).

Figure 12: Cyclicality of earnings – Age and Gender

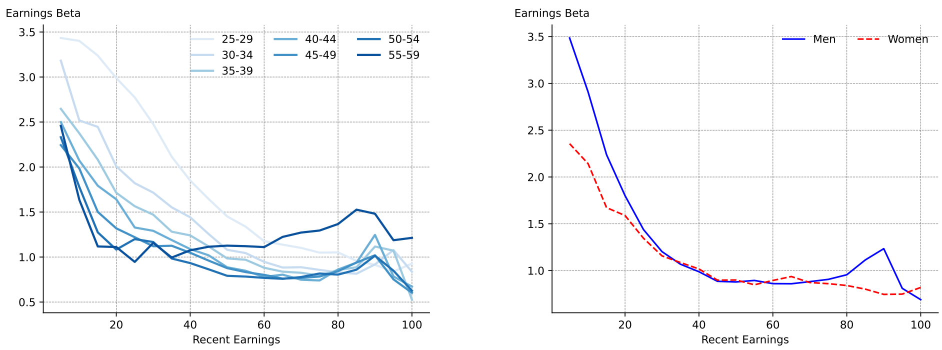

(a) Earnings Beta – Age

(b) Earnings beta – Gender

Note: The *Left Panel* shows the coefficient $\beta_{Y,4}^q$ from Equation (2) by quantile, estimated for subsamples of different five-year age groups. The *Right Panel* plots the coefficient $\beta_{Y,4}^q$ from Equation (2) by quantile, for men (blue) and women (red). Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The sample period is 1980-2014.

Similar to results presented in Guvenen et al. (2017), the size of the β_Y^q, as well as their heterogeneity, diminish with age (with the exception of the top earners of the last age group, where growth rates are imputed). While earnings growth of individuals between the ages of 25 and 34, at the bottom of the earnings distribution, moves with aggregate earnings by a factor of 3, this coefficient diminishes to 2 along the life-cycle. Towards to top of the distribution, procyclicality is weaker at all ages.

Earnings growth at the bottom is considerably more procyclical for men than it is for women as their β_Y^q is almost 50% larger. However, towards the median of the recent earnings distribution, the two graphs converge. At the very top, the two groups diverge again, potentially due to the fact that considerably more male earnings observations are censored and hence imputed. Hence, the difference is to be taken with a grain of salt.

A.4 Occupation and Industry

Another potential driver of the heterogeneous cyclicality in earnings growth rates are an individuals occupation or industry. Hence, again, I reestimate Equation (2), splitting the sample into ten separate occupation bins. Figure 13 presents the results. The familiar picture of a higher correlation between individual and aggregate earnings growth at the bottom of the distribution is present in all occupations, but to varying degrees. Among the most cyclical occupations related to construction and manufacturing, moving 3:1 with average earnings in the first quantile. In the same quantile, occupations related to health care and science only co-move slightly 2:1 with the aggregate. Towards the median quantile, all graphs converge towards a value between 0.5 and 1.

Figure 13: Cyclicality of earnings – Occupation

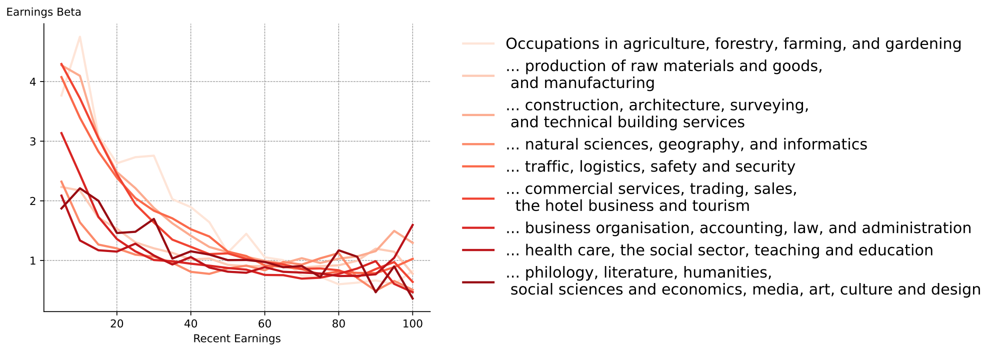

Note: The figure shows the coefficient $\beta^q_{Y,4}$ from Equation (2) by quantile, estimated for subsamples of different occupational groups. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The sample period is 1980-2014.

Figure 14 shows the results of performing a similar analysis separately by industry. For all industries, there is a strong downward trend in procyclicality as one moves up across quantiles. The least cyclical industry, for most quantiles, is public administration. Among the most cyclical are the automotive industry and agriculture.

Figure 14: Cyclicality of earnings – Industry

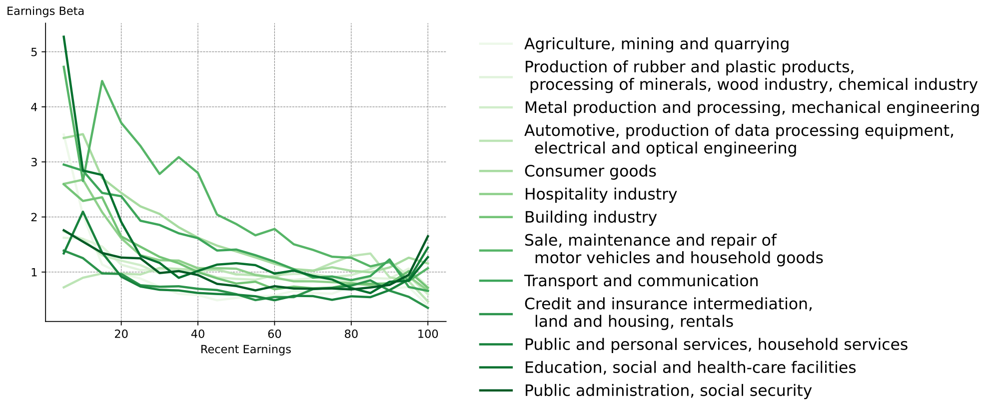

Note: The figure shows the coefficient $\beta^q_{Y,4}$ from Equation (2) by quantile, estimated for subsamples of different industry groups. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The sample period is 1980-2014.

A.5 Quantile ordering with Conditional earnings

In the baseline estimation of the earnings betas, I sort individuals into quantiles based on their earnings history over the previous five years, including zero-earnings spells, i.e., unconditional earnings. Alternatively, individuals can be sorted according to their conditional earnings history, i.e., excluding zeros. This measure is likely closer to individual level productivity, holding constant age and gender.

Figure 15 shows the result of this alternative sorting method on the earnings beta estimations (blue) and the baseline estimation discussed above (red). Sorting individuals based on conditional earnings considerably lowers the earnings beta for the bottom quantiles, but does not alter the general shape of the graph, which still indicates that those at the lower end of the (now conditional) earnings distribution receive earnings which are much more correlated with the aggregate business cycle than those further up in the same distribution.

Figure 15: Alternative Sorting – Conditional Earnings

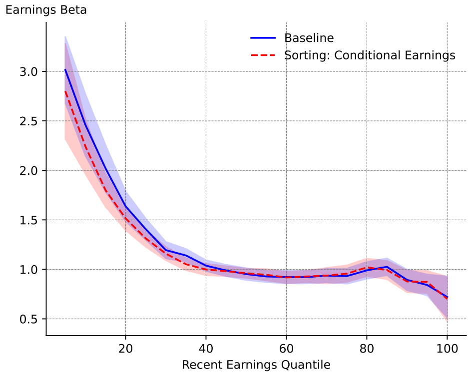

Note: This figure shows the coefficients $\beta^q_{Y,4}$, from equation (2), by quantile. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The blue line plots the coefficient when zero-earnings observations are included in this earnings history, the red line conditions on employment and excludes zeros from the earnings history. The sample period is 1980-2014.

A.6 Unemployment benefits

As outlined in footnote 6 the only consistent measure of unemployment benefits in the data are Unemployment Benefits I. They are paid out for a limited period of time after job-loss with a replacement rate of 65%. The left panel of Figure 16 compares the official German Unemployment rate[24] (blue) to the one obtained using the SIAB sample (red). To make the series comparable, I seasonally adjust the SIAB unemployment rate. From 1980 to 1998, the SIAB sample overestimates the unemployment rate by close to two percentage points. After 2000, the pattern reverses and it underestimates unemployment, especially towards the end of the sample. However, the dynamics of the two unemployment rates are very similar over time. Note that the duration of unemployment benefit receipts was shortened to 12 months in 2005, potentially leading to fewer observed unemployed individuals after that period.

[24] I obtain the Quarterly Registered Seasonally Adjusted Unemployment Rate for Germany from Fred (LMUNRRTTDEQ156S) .

Figure 16: Cyclicality of earnings – Unemployment benefits

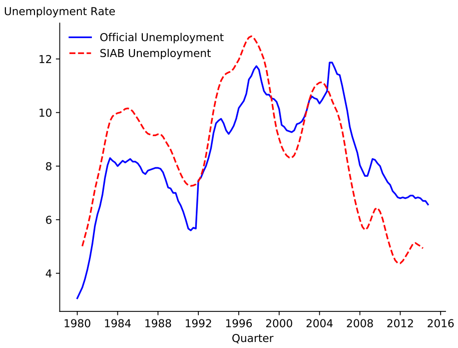

(a) Put unemployment rate here

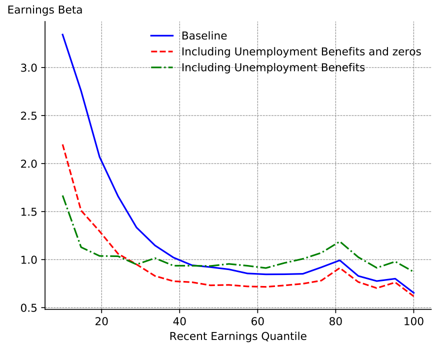

(b) Earnings betas – Unemployment benefits

Note: The *Left Panel* shows the official German unemployment rate (blue) and the unemployment rate implied by the SIAB sample (red). The official unemployment rate is constructed for West-Germany during the years before 1990. The *Right Panel* shows the coefficients $\beta^q_{Y,k}$, from equation (2), by quantile. The blue line represents the baseline model, which sets earnings to zero during non-employment spells. The red line shows the coefficients when unemployment benefits are included as earnings for the non-employed. The green line restricts the sample of the non-employed only to those individuals who receive unemployment benefits.

Here, I investigate the consequences of including non-employed earnings in the baseline estimation by re-estimating Equation (2) while including the observed benefits as earnings for the unemployed. I perform two robustness tests: (i) keeping earnings at zero for all non-employed individuals who are not unemployed (ii) excluding all non-employed individuals who are not unemployed. The former approach uses the same sample as the baseline, only changing earnings for some non-employed, while the second approach produces a smaller sample. Both approaches require new quantiles to be calculated.

The right panel of Figure 16 compares the baseline estimation, previously reported in Figure 3 (blue) to the two alternative approaches, including zero earnings for the non-unemployed non-employed (red) and excluding the non-unemployed non-employed (green). The Figure shows that the earnings betas at the bottom get progressively smaller with each step. When moving from the baseline to including non-employment earnings, this is intuitive, as unemployment benefits introduce an insurance mechanism that smooths the strongest earnings fluctuations. Not also that there is almost no difference between the two lines at the very top, implying that unemployment transitions are less important here.

Excluding zeros for the non-unemployed non-employed leads the betas at the bottom to fall further. This is likely due to a reshuffling of individuals across percentiles, dampening the effects seen in the other two graphs. Beyond the 6th quantile, the quantile specific earnings betas are slightly higher than using the other two approaches.

B Additional results

In addition to the calibration results reported in the main body, Figure 17 shows the average employment share across quantiles of the distribution. The model can match the steep incline in this metric that is present in the data, with employment shares rising from slightly over 30% to more than 95% at the top of the distribution. As pointed out in the empirical section, the relatively high presence of unemployment at the bottom of the distribution is one of the driving forces behind the heterogeneous procyclicality in earnings growth. Hence, the model matching this slope is an important ingredient in its performance. While the slope was not explicitly targeted, it is a result of the heterogeneous job-finding and separation probabilities across the recent earnings distribution.

Figure 17: Employment shares across the distribution

Note: This figure shows the share of employed individuals, by quantile. In both the model (solid) and the data (dashed), individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history.

Figure 18 shows earnings growth along the intensive margin (left panel) and the distribution of marginal propensities to consume (right panel) in the model and the data. In the model, earnings growth rates for job-stayers are zero by construction: wages, and hence earnings, are fixed for the duration of the match. The only way for individuals to increase their earnings is by switching jobs. The model considerably underpredicts the earnings growth for job-switchers (solid green line). Productivity shocks large enough to cause large jumps in earnings after job-switching are very rare in the model. This limits the size of earnings growth that job-switchers can achieve. Still, beyond the third decile, the model produces smaller-and-smaller earnings growth for job-switchers, which is in line with the data up to the eigth decile. The reason is that positive productivity shocks become less likely for high productivity individuals—once the highest productivity grid-point is reached, there is nowhere to go but down. This leads to (i) fewer switchers and (ii) small earnings changes conditional on switching. The model is unable to replicate the increase in earnings growth in the top part of the earnings distribution.

Figure 18: Steady State earnings growth rates and liquidity

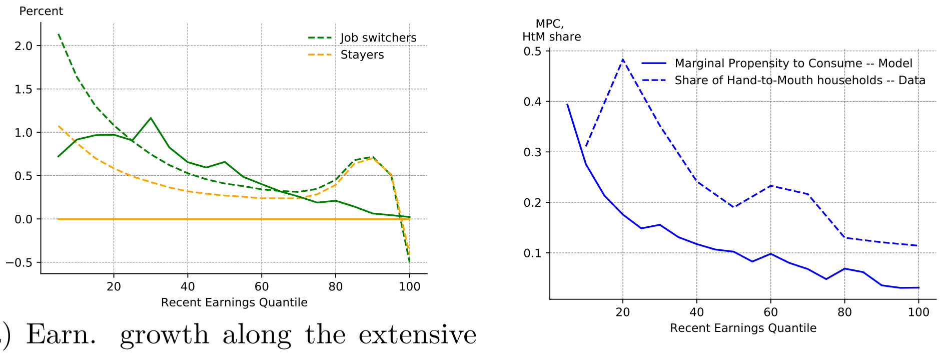

(a) Earn. growth along the extensive margin

(b) Liquid asset holdings

Note: The *Left Panel* shows a comparison between quarterly growth-rates implied by the model and the same statistic measured in the data. Dashed lines represent the data, solid lines represent the model's output. The yellow line represents earnings growth for job-stayers, the green line plots earnings growth for job-switchers. The *Right Panel* shows average marginal propensity to consume in the model (solid) and the share of hand-to-mouth households in the data, calculated according to Kaplan et al. (2014) (dashed). Hand-to-mouth shares are calculated using data from the Households Finance and Consumption Survey (HFCS), for details, see the text. Note that recent earnings in the HFCS are calculated using yearly earnings reported in the survey, as the data cannot be matched to German administrative labor market data.

The right panel of Figure 18 shows the distribution of marginal propensities to consume in the model and compares it to the share of households classified as hand-to-mouth, constructed along the lines of Kaplan et al. (2014), with the modifications mentioned in Section 4. This statistic identifies households which are likely to use all their income in a given month, implying little buffer against income shocks. These households are likely to have high marginal propensities to consume. Both the model and the data, the latter imperfectly measured, imply a steep negative gradient of marginal propensities to consume along the recent earnings distribution. This dimension was not targeted at all, implying that reproducing it is a success of the model.

B.1 Symmetry across aggregate shocks

I estimate two impulse responses to aggregate productivity shocks of $\pm 1\%$ and persistence $\rho = 0.9$. Figure 19 shows the responses of several aggregate variables to these shocks. The impulse responses to the negative productivity shocks are almost perfectly symmetrical to those after a positive shock. In response to the positive shock, output increases by close to 1 percent upon impact and then converges back to its steady state equilibrium. Wages also increase, but with a lag. This is due to the structure of the labor market: upon impact, job finding increases as firms want to hire workers who will temporarily be more productive. Workers trade off some of the increase in the job-finding probability for higher wages. Over time, this raises the average wage in the economy. As aggregate productivity converges back to its steady state value, these incentives wane away. The aggregate converges back to steady state as the workers who were hired at higher wages separate from their matches.

The labor income tax falls in response to a positive shock, for two reasons. Unemployment decreases, which leads to less government expenditures. Additionally, wages rise, implying a larger tax base. In period $t = 0$, as agents want to save some of their additional labor income for the future, interest rates fall. This is because, at the same time, as the government needs to finance lower outlays for the unemployment insurance scheme, it reduces the asset supply to keep the budget balanced, following its budget rule.

Figure 19: Impulse response to Productivity shock

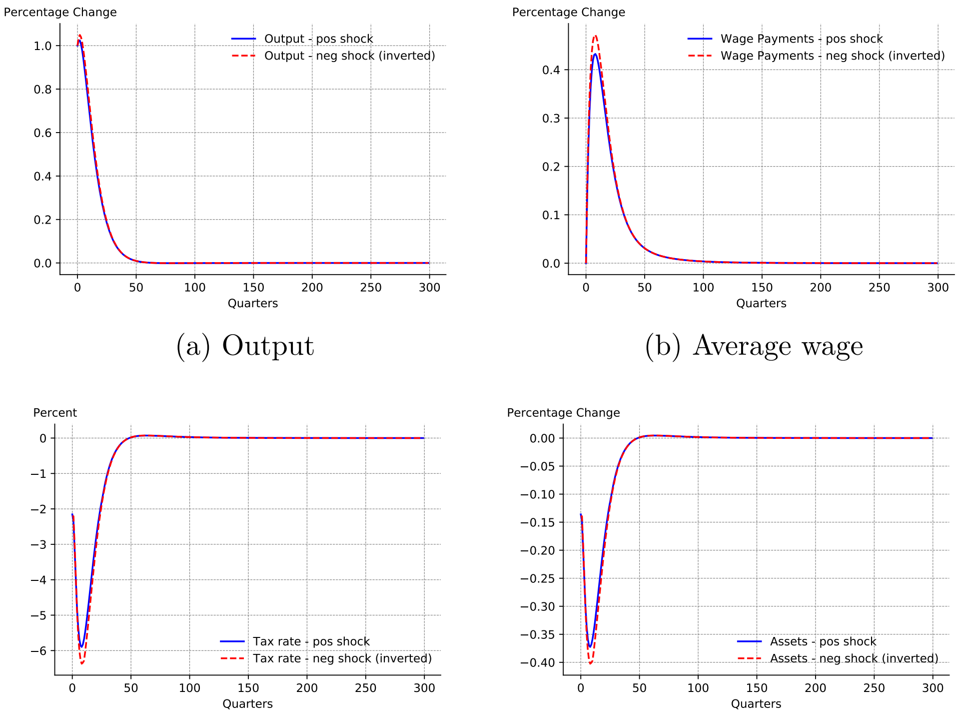

(a) Output

(b) Average wage

(c) Labor income tax

(d) Assets

Note: The *Top Left Panel* shows the response of output, in percent, relative to steady state, to a positive (solid) and negative (dashed) productivity shock which reduces aggregate match productivity by 1%, with persistence 0.9. The response to the negative shock is inverted. The *Top Right Panel* shows the responses of benefit payments, in percent, relative to the steady state; the negative response is inverted. The *Bottom Left Panel* shows the responses of the labor income tax rate to the two shocks. The *Bottom Right Panel* shows the response of asset demand, in percent, relative to the steady state, with the response to the negative shock inverted. The horizontal axis shows quarters since the shock.

B.2 Procyclicality of transition rates in the model

I compute the procyclicality of labor market transition probabilities in the model the same way as in the data, using Equation 3. The coefficient $\beta^q_{Y,trans}$ measures the percentage point change in the transition probability in question, in response to a one percentage point increase in aggregate earnings. Figure 20 displays the results. The left panel compares the procyclicality of employment-to-employment transitions between the model (solid) and the data (dashed). Both are downwards sloping along the recent earnings distribution. While the model closely replicates the data beyond the second decile, it severely overshoots the first quantile.

Figure 20: Procyclicality of transition probabilities

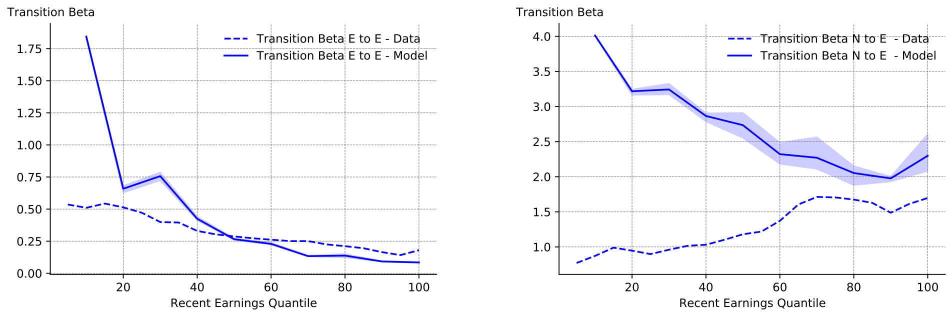

(a) Empl.-to-empl. transitions

(b) Job-finding

Note: The figure plots the coefficients $\beta^q_{Y,4,trans}$ in Equation (3) for the model and the data. The *Left Panel* displays the procyclicality of employment-to-employment transition probabilities along the recent earnings distribution. The *Right Panel* displays the procyclicality of job-finding probabilities along the recent earnings distribution. In both figures, the solid lines represent the model and the dashed lines represent the data. Individuals are sorted into 20 quantiles each quarter based on their most recent five-year earnings history. The model's results are obtained through 10 simulations with 3 million individuals, the shaded areas show the range of estimates across simulations, the point estimates are averages across these simulations. The sample period for the data is 1980-2019.

The right panel of the same figure compares the procyclicality of job-finding between the model and the data. In both, higher aggregate growth is associated with more job-finding across the distribution. Procyclicality is considerably higher in the model, which likely leads it to somewhat overestimate the contribution of job-finder's earnings growth to overall earnings procyclicality.

C Block recursive equilibrium

Block recursive equilibrium (BRE) definition: A block recursive equilibrium is an equilibrium in which, given a path for the interest rate r and the labor income tax τ, the households' policy functions and submarket tightnesses only depend on the aggregate productivity state A, but not on the distribution of agents Ω.

Proposition *If i) utility function $u(\cdot)$ is strictly increasing, strictly concave, and satisfies the Inada conditions; ii) choice sets $\mathcal{W}$ and $\mathcal{A}$, and sets of exogenous productivity processes z and A are bounded; iii) matching function M exhibits constant returns to scale; and iv) All policies are restricted to depend on the aggregate state only through aggregate match productivity, then there exists a unique BRE for this economy, given a path for r and τ.*

Proof As mentioned in the body of the paper, the proposition and the proof closely follow Karahan and Rhee (2019), Herkenhoff (2019) and Birinci and See (2023).

I prove the existence of a block recursive equilibrium in two steps. The first step is showing that the firms' value functions and the resulting market tightnesses only depend on the aggregate state ψ through aggregate productivity A, **given a tax rate** τ **and an interest rate** r. The second step shows that the households' policy and value functions are similarly independent of the distribution of agents Ω, **given** τ **and** r. Consequently, there is a solution to the households' problem which, together with the solution to the firms' problems and the resulting market tightnesses constitutes a BRE, as long as τ and r are given. Note that I condition on the tax rate τ and the interest rate r, two objects which, in order to clear the government budget and the asset market, will depend on the aggregate distribution of workers across states.

Let $\mathcal{J}(\mathcal{Z}, \mathcal{X}, \mathcal{W}, \mathcal{A}|\tau, r)$ be the set of continuous and bounded functions which map $\mathcal{J} : \mathcal{Z} \times \mathcal{X} \times \mathcal{W} \times \mathcal{A} \to \mathbb{R}$, given a tax rate τ and an interest rate r. Further, let $\mathbf{T}_{\mathcal{J}}$ be the operator associated with the firm's value function, Equation (5). One can verify, using Blackwell's sufficiency conditions, that $\mathbf{T}_{\mathcal{J}} : \mathcal{J} \to \mathcal{J}$ is a contraction, the unique fixed point of which I denote as $J^* \in \mathcal{J}$. From this, it follows that the firm's value function only depends on the aggregate state ψ through the state of aggregate productivity A. In turn, this implies that the wage posting choices by firms, conditional on worker wealth $\mathcal{X}$ and worker productivity $\mathcal{Z}$, are also only affected by the aggregate state ψ through aggregate productivity

A. Upon substituting J^* into Equation 6, I obtain

$$\theta^*(z,x,w';A) = \begin{cases} q^{-1}\left(\frac{\kappa(1+r)}{\mathbb{E}[J^*(z',x',w';A'|\tau,r)]}\right) & \text{if } w' \in \mathcal{W}(z,x,w;A) \\ 0 & \text{otherwise} \end{cases} \tag{15}$$

This condition shows that market tightness does not depend on the distribution of agents across states, Γ, apart from the interest rate r, which I am conditioning on.

Next, I move to the workers' problems. I combine all value functions into a single functional equation and show that this equation is a contraction. It can be shown that this function maps the set of functions which depend on the aggregate state ψ only through aggregate productivity A, conditional on r and τ. This function V is of the form $V : \{U, E\} \times \mathcal{X} \times \mathcal{Z} \times \mathcal{W} \times \Omega \to \mathbb{V}$, where Ω defines all possible realizations of the aggregate state. In this formulation,

$$V(U,z,x;\psi) = V^U(z,x;\psi|\tau,r)$$
$$V(E,z,x,w;\psi) = V^E(z,x,w;\psi|\tau,r)$$

I now define a set of functions $\mathcal{V} : \{U, E\} \times \mathcal{X} \times \mathcal{Z} \times \mathcal{W} \times \mathcal{A} \to \mathbb{V}$ and let $\mathbf{T}_V$ be an operator such that

$$\begin{aligned}
(T_V V) = &\mathbb{I}_U \Big\{\max_{c_U} \quad u(c_U) + l + \beta \Big[\max_{w'} \mathbb{E}\left[\eta(\theta(z,x,w';A|\tau,r))V(E,z',x'_n,w';A'|\tau,r)\right.\\
&+ \ (1-\eta(\theta(z,x,w';A|\tau,r)))V(U,z',x'_u;A'|\tau,r)]]\Big\}\\
&\mathbb{I}_E \Big\{\max_{a',c} \quad u(c) + \beta \max_{w'} \mathbb{E}\,[\\
&(1-\delta(z'))(1-\Lambda\eta(\theta(z,x,w';A|\tau,r)))V(E,z',x'_e,w;A'|\tau,r)\\
&+ \ (1-\delta(z'))\Lambda\eta(\theta(z,x,w';A|\tau,r))V(E,z',x'_n,w';A'|\tau,r)]\\
&+ \delta(z')V(U,z',x'_u;A'|\tau,r)\Big\}
\end{aligned}$$

subject to

$$\begin{aligned}
c + a' &\le x\\
x'_u &= (1+r)a' + b(z')\\
x'_e &= (1+r)a' + (1-\tau)w\\
x'_n &= (1+r)a' + (1-\tau)w'\\
A' &= F_A(A)\\
z' &= F_Z(z).
\end{aligned}$$

In this formula, I use the fact that submarket tightness θ does not depend on the aggregate distribution of workers across states. Further, x_u is the cash-on-hand of unemployed workers, x_n is the cash-on-hand value for workers who find a new job at w' and x_e is the cash-on-hand value for workers who stay employed at wage w. The first two lines represent the problem of an unemployed worker, the last three lines represent the problem of an employed worker.

If we assume that the utility function is bounded and continuous, then $\mathcal{V}$ is a set of bounded and continuous functions. It can be shown that the operator T_V maps $\mathcal{V} \rightarrow \mathcal{V}$. Using Blackwell's sufficiency conditions for a contraction and the assumptions on the boundedness of the sets defining the exogenous processes $\mathcal{Z}$ and $\mathcal{A}$, as well as the choice sets $\mathcal{W}$ and $\mathcal{X}$, one can show that T_V is a contraction with a fixed point $V^* \in \mathcal{V}$. Thus, the solution to the worker's problem does not depend on the distribution of workers across states Ω. This, in combination with the firm's problem above, constitutes a block recursive equilibrium. $\square$

D Computational appendix

To solve the model in steady state, I employ the following algorithm:

- Guess values for τ and r
- Solve the firms' and the workers' problems
  - Guess wage and consumption policy functions for the workers
  - Given these policies together with the free entry condition, solve the firms' problem
  - Given the resulting job-finding probabilities and taking wage choices as given, update the consumption policy function using an endogenous grid point method (Carroll, 2006)
  - Given the job-finding probabilities and the consumption policy functions, update the wage policy functions
- Using the policy functions computed in the previous step, compute the distribution of agents across states, Ω.
- Check that asset markets clear and the government budget holds
- If not, update the guess and repeat

Along the transition path, when solving for the economy's response to a single unexpected aggregate productivity shock, I employ the sequence space Jacobian method proposed by Auclert et al. (2021) to find a sequence of τ_t and r_t. Using this approach speeds up the solving process by a factor of 20, compared to traditional methods of gradient descent.

I calculate the earnings procyclicality measures using a combination of a non-stochastic simulation approach for the evolution of incomes and employment across states, and a stochastic simulation which sorts individuals into quantiles.