

Macroeconomic Risk Chair
SCOR-PSE

The Curious Incidence of Monetary Policy Across the Income Distribution

Tobias Broer

PSE, IIES, Stockholm
University and CEPR

John Kramer

IIES, Stockholm
University

Kurt Mitman

IIES, Stockholm
University, CEPR and IZA

Executive Summary

How does monetary policy intervention affect the earnings and employment prospects of individuals across the income distribution? Does the unequal incidence of monetary policy across the distribution amplify or dampen the response of aggregate consumption to changes in interest rates or future consumption? The burgeoning heterogeneous-agent New Keynesian (HANK) literature has identified labor income as an important channel through which household heterogeneity impacts the transmission of monetary policy. Answering the foregoing questions is key to understand the transmission of monetary policy to the aggregate economy. However, the literature is still lacking direct empirical evidence on these transmission channels.

In this paper, Tobias Broer, John Kramer and Kurt Mitman first study empirically the heterogeneous effects of monetary policy surprises on labor earnings across the income distribution. Using high-frequency data on labor earnings and labor market status from Germany, they show that monetary policy has significantly larger effects on the earnings of low-income workers. This is mainly because their job-loss risk responds more strongly to changes in interest rates than that of high-income workers. This unequal incidence significantly reduces income inequality in response to monetary expansions and has long-lasting effects on employment rates of poor workers, which remain elevated even years after the initial shock. In particular, the authors find that an unexpected interest rate cut leads the Gini coefficient of labor earnings to fall significantly. In addition, monetary policy has significant effects on medium-run employment prospects: individuals who become unemployed in the month of a monetary policy expansion find jobs significantly faster, have significantly higher earnings, and remain employed significantly longer.

The authors then use a structural model to show how this heterogeneous incidence of monetary policy strongly amplifies its effect on aggregate demand. Relative to a model where unemployment risk is homogeneous across the distribution, heterogeneous incidence further amplifies the unemployment-risk channel because monetary policy affects more strongly the riskier workers who account for the bulk of precautionary savings. This positive association of level and cyclicity of risk in the cross-section makes aggregate precautionary savings more responsive to monetary policy. Their analysis suggests quantitatively important results: consumption increases by about a third after a monetary policy intervention.

The Curious Incidence of Monetary Policy Across the Income Distribution*

Tobias Broer[†] John Kramer[‡] Kurt Mitman[§]

August 24, 2022

Abstract

We use high-frequency administrative data from Germany to study the effects of monetary policy on income and employment across the earnings distribution. Earnings growth at the bottom of the distribution is substantially more elastic to policy shocks. This unequal incidence is driven by differences in the response of employment risk across the distribution: job loss is more countercyclical for lower-earnings households. Viewed through the lens of a standard incomplete-markets model, the heterogeneous incidence substantially amplifies the equilibrium response of aggregate consumption to shocks.

*We thank Per Krusell, Serdar Ozkan, Kjetil Storesletten and seminar participants at various institutions and conferences for helpful comments. Agneta Berge and Carolina Lindholm provided outstanding research assistance. Support from the European Research Council ERC Starting Grant 759482, the Jan Wallander och Tom Hedelius stiftelse, and the Ragnar Söderbergs stiftelse is gratefully acknowledged. 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.

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

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

[§]IIES, Stockholm University, CEPR and IZA, kurt.mitman@iies.su.se

1 Introduction

How do monetary policy interventions affect the earnings and employment prospects of individuals across the income distribution? Does the unequal incidence of monetary policy across the distribution amplify or dampen the response of aggregate consumption to changes in interest rates or future consumption? The burgeoning heterogeneous-agent New Keynesian (HANK) literature has identified labor income as an important channel through which household heterogeneity impacts the transmission of monetary policy (inter alia, [Auclert, 2019](#); [Bilbiie, 2018](#); [Hagedorn et al., 2019](#); [Kaplan et al., 2018](#)). Answers to the foregoing questions are key for understanding the transmission of monetary policy to the aggregate economy. However, there is little direct empirical evidence from large advanced economies on these transmission channels.¹

In this paper, we first study empirically the heterogeneous effects of monetary policy surprises on labor earnings across the income distribution. Our findings show that monetary policy has significantly larger effects on the earnings of low-income workers. This is mainly because their job-loss risk responds more strongly to changes in interest rates than that of workers with higher incomes. This unequal incidence significantly reduces income inequality in response to monetary expansions and has long-lasting effects on employment rates of poor workers, which remain elevated even years after the initial shock. Second, we use a structural model to show how this heterogeneous incidence of monetary policy on unemployment risk along the income distribution strongly amplifies its effect on aggregate demand. This holds relative to both a standard representative-agent model without unemployment risk and a model where unemployment risk is homogeneous across the distribution.

For our empirical analysis, we use a long panel of detailed administrative data from Germany, containing individual labor market biographies including earnings. Labor market status is observed at a daily frequency. The high-frequency nature of our data allows us to estimate responses of earnings and labor market transition probabilities to monetary policy shocks and high-frequency changes in aggregate earnings. This sets our paper apart from the literature that empirically investigates the heterogeneous effects of business cycles on individual income risk using administrative datasets beginning with [Guvenen et al. \(2015\)](#) for the US.² Our dataset allows us to understand the importance of changes in employment status for earnings changes. While the previous literature has speculated that the larger

¹An important exception is ([Coibion et al., 2012](#)), which studies the impact of monetary policy shocks using the CEX data in the US. Some recent studies ([Holm et al., 2020](#), e.g.) examine the transmission of monetary policy using administrative data similar to ours in small Nordic countries.

²See also, [Guvenen et al. \(2017\)](#); [Patterson et al. \(2019\)](#) for the US, and for other countries [Halvorsen et al. \(2020\)](#) (Norway), [Hoffmann and Malacrino \(2019\)](#) (Italy), [De Nardi et al. \(2019\)](#), (Netherlands and US)

sensitivity of earnings at the bottom of the distribution was due to non-employment risk, ours, to our knowledge, is the first to show that this is the case.

We identify monetary policy surprises using high-frequency changes in Overnight Indexed Swap (OIS) rates for the Eurozone.³ We then use the identified shock series for estimating the impact of shocks on labor earnings using local projections a la [Jordà \(2005\)](#). Monetary policy affects labor earnings most at the bottom of the permanent-income distribution (proxied by past labor earnings as in [?](#)). In response to an expansionary monetary policy surprise, earnings growth rises about three times as much in the bottom quintile as it does at the top. The differential growth is accounted for by a substantially stronger fall in separation rates into non-employment at the bottom of the distribution. In contrast, job-finding rates, while pro-cyclical, rise homogeneously across the distribution. Similarly, the earnings growth of workers who remain employed increases, but with mostly uniform effects across the income distribution.

These heterogeneous earnings responses across the distribution give rise to strong redistributive forces. An unexpected interest rate cut leads the Gini coefficient of labor earnings to fall significantly. In addition, monetary policy has significant effects on medium-run employment prospects: individuals who become unemployed in the month of a monetary policy expansion find jobs significantly faster, have significantly higher earnings, and remain employed significantly longer.

To understand the implications of our empirical findings for the aggregate economy, we extend the framework of [Werning \(2015\)](#) to account for heterogeneous unemployment risk—the key force behind the heterogeneous earnings responses we document. [Werning \(2015\)](#) shows how countercyclical unemployment risk amplifies the effect of interest-rate shocks on aggregate demand through precautionary savings, as workers who fear unemployment reduce their consumption in recessions by more than the fall in their permanent income.⁴ Heterogeneous incidence further amplifies this unemployment-risk channel because monetary policy affects more strongly the riskier workers who account for the bulk of precautionary savings. This positive association of level and cyclical risk in the cross-section makes aggregate precautionary savings more responsive to monetary policy. Our analysis suggests that this increase is quantitatively important, raising the consumption response to monetary policy interventions by about a third.

The rest of the paper is organized as follows. The next section presents the data and

³See e.g. [Gertler and Karadi \(2015\)](#); [Nakamura and Steinsson \(2018\)](#), and [Almgren et al. \(2019\)](#) for discussion of high-frequency identification of monetary policy shocks in the U.S. and Eurozone, respectively.

⁴Prior to Werning’s paper, this unemployment risk-precautionary savings channel was analyzed in [Ravn and Sterk \(2017, 2021\)](#). And related papers include [Acharya and Dogra \(2020\)](#); [Challe \(2020\)](#); [Gornemann et al. \(2016\)](#).

describes the structure of income and employment transitions in our sample on average. Section 2 describes how we identify monetary policy surprises, and how we use them to study their heterogeneous incidence across the earnings distribution. Section 5 investigates the implications of our findings for aggregate consumption responses to monetary policy shocks. Section 6 concludes.

2 Data

We use administrative social security data on a two-percent sample of all labor-market histories in Germany from the Sample of Integrated Employment Biographies (provided by the Research Data Center, FDZ).⁵ This dataset contains about 1.7 million individuals but excludes civil servants and self-employed individuals. For our analysis, we utilize data on the years between 1995 and 2013. Each observation in the original dataset is a labor-market spell (Ganzer et al., 2017).⁶ For our purposes, we convert these spells into monthly employment histories for each individual, resulting in about 300 million person-month observations. Each such observation includes an individual’s employment status and their average daily labor earnings, which we aggregate to the monthly level. Earnings are deflated using the Harmonized Index for Consumer Prices for Germany.⁷ For about ten percent of individuals in our sample, earnings are top coded; we exclude these observations. All non-employed workers are coded to have zero income.

Because we are interested in the effect of monetary policy on labor earnings and employment status, we focus on individuals with a high degree of attachment to the labor market. In particular, we restrict our sample to employed individuals liable to social security without special characteristics, (thus excluding, for example, trainees and marginal part-time workers) and the unemployed, defined as individuals who received unemployment benefits (ALG I) at the beginning of their current non-employment spell.

To study the differences in the earnings responses to monetary policy across the income distribution, we rank individuals in a given period $t - 1$ according to a proxy measure of their *permanent* incomes. Our preferred proxy is average earnings over the five years preceding month $t - 1$, as in Guvenen et al. (2017).⁸ Using this measure, in every month $t - 1$, we sort

⁵We rely on the factually anonymous version of the Sample of Integrated Labour Market Biographies (SIAB-Regionalfile) – Version 7514. 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.

⁶Employment relationships longer than 12 months are split into multiple spells. We drop spells that are shorter than 1 month. Potentially missing spells are imputed according to Drews et al. (2007).

⁷Obtained from Eurostat, series `prc_hicp_midx`.

⁸Our estimation sample comprises the period between 2000M1 to 2012M12. However, we make use of

individuals into quantiles, conditional on gender and five-year age brackets. We restrict the sample to workers who have at least one earnings observation in the five years prior to period $t - 1$ in order to avoid bunching at zero.

To understand how key observables evolve along our permanent income distribution (henceforth simply the “income distribution”), Table 1 reports descriptive statistics across deciles for the month of January 2010.⁹¹⁰

Table 1: Averages within deciles of permanent income, first quarter 2010

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
	mean	mean	mean	mean	mean	mean	mean	mean	mean	mean
Education	2.47	2.42	2.36	2.36	2.35	2.38	2.43	2.51	2.69	3.13
Monthly earnings	1337.67	1588.82	1807.78	2018.00	2221.79	2449.59	2707.51	3012.21	3421.54	4261.82
Employed	0.81	0.87	0.92	0.95	0.97	0.98	0.98	0.99	0.99	0.98
Job finding	0.36	0.50	0.52	0.54	0.55	0.58	0.57	0.53	0.60	0.53
Job loss	0.07	0.06	0.04	0.03	0.02	0.02	0.01	0.01	0.01	0.02
Observations	28878	28871	28872	28870	28870	28874	28871	28871	28872	28866

Note: The table shows values of different variables averaged within deciles of the permanent income distribution in January 2010. Deciles are conditional on five-year age brackets and gender. Education takes a value of 1 for individuals without a degree, 2 for vocational training, 3 for high school, 4 for high school and vocational training, 5 for graduates of technical colleges and 6 for university graduates. We impute education following the imputation procedure in [Fitzenberger et al. \(2005\)](#). Monthly earnings are nominal values. Job-finding and job-loss refer to U to E and E to U transitions over twelve months, respectively. The deciles are computed conditional on age and gender. These variables are thus not reported.

In our dataset, education is measured by a categorical variable between 1 (no degree) and 5 (university graduate). Because 70 percent of all individuals in our sample indicate vocational training as their highest qualification, education levels are very similar across the first 8 deciles, but strongly rise across the top, where degrees from technical colleges and universities are more common. The gradient of nominal earnings across the distribution is substantial, with average earnings in the top decile more than 3 times higher than in the first. Employment rates are high in this sample of highly-attached individuals. They average 81 percent in the bottom decile, and rise steeply across the bottom half of the distribution to flatten out around 98 percent above the median. Job-finding rates (defined as 12-month transitions of the unemployed into employment) are equal between 50 and 60 percent in all deciles but the first, where they are substantially lower (about a third). Job-loss probabilities (similarly defined) fall monotonically, from 6 to 2 percent, across the distribution.

data from 1995 in order to compute our backward-looking permanent income measure, but only consider monetary policy surprises from 2000M1 to 2012M12.

⁹The deciles are computed conditional on age and gender. These variables are thus not reported in Table 1.

¹⁰Note that, with some abuse of language but hopefully no room for confusion, we call deciles both the 9 points of the distribution as well as the 10 groups they define (we proceed similarly for other quantiles).

3 Estimation strategy

3.1 Identifying monetary policy surprises

We focus on the period between January 2000 and December 2012, when European monetary policy was conducted by the ECB.¹¹ Since the German economy accounts for roughly one-quarter of Euro-area GDP it is likely that the ECB’s monetary policy was heavily influenced by German economic performance. Hence, when estimating the impact of interest rate changes on the German economy, endogeneity is an important concern.

To identify monetary policy surprises our approach follows [Almgren et al. \(2019\)](#). We rely on high-frequency changes in Overnight Indexed Swap (OIS) rates. We use these changes to instrument for unexpected changes in the interest rate which the ECB charges for its main refinancing operations (MROs), which we denote as Δi_t . Every six weeks, on Thursdays, the ECB governing council meets to decide on monetary policy actions. At 13:45 CET, a press release is posted, which concisely summarizes the decisions taken by the governing council. Subsequently, at 14:30 CET, the president of the ECB holds a press conference, first motivating the decisions taken in an introductory statement and later taking questions from the audience. Our instrument, Z_t , equals the change in 3-month EONIA OIS rates in response to these two events in a narrow time window around them. If this measure is large, in absolute terms, we conclude that the decisions taken by the ECB Governing Council were not expected by financial markets and vice versa. The identifying assumption underlying the approach is that no other news is released during the above-mentioned short time windows which have an impact on the effectiveness of monetary policy.¹²

Our main empirical specification to estimate the effects of monetary policy surprises on economic variables is the following regression:

$$x_{t+h} - x_{t-1} = \alpha_h + \beta_h \Delta i_t + \gamma_h X_{t-1} + \varepsilon_{t,h} \quad (1)$$

where x_{t-1} represents the value of the economic variable in question one period before the monetary policy surprise, and x_{t+h} represents its value h periods after the shock. We condition this growth rate on x_{t-1} , as opposed to x_t , because it is conceivable that monetary policy has contemporaneous effects on x_t , which would invalidate all growth rates going forward. The vector X_{t-1} represents a set of control variables consisting of three lags of the instrument Z_t and Δi_t , as well as calendar month dummies to control for seasonality.

¹¹The high-frequency identification approach outlined here cannot be implemented for earlier time periods, as the Bundesbank did not relay its policy decision on a precisely planned schedule on the announcement day.

¹²For more information, see [Almgren et al. \(2019\)](#).

3.2 Aggregate effects of monetary policy

Before moving to individual incomes, we investigate the effect of monetary policy shocks on the aggregate economy in Germany. To this end, we estimate the regression in Equation (1), replacing x with (i) the logarithm of the HICP price index for Germany, (ii) the logarithm of industrial production, (iii) the German unemployment rate, and (iv) the real interest rate, computed as the change in the logarithm of the German price index between months $t + 1$ and t subtracted from the ECB’s policy rate in period t :

Figure 1 shows the impulse responses to a surprise increase in the policy rate by one standard deviation (following Gertler and Karadi, 2015), estimated using Equation (1). The horizontal axis measures the time after the monetary policy shock in months, the vertical axis measures the percentage point change in the variable in question. The top left panel indicates that the inflation rate does not strongly react to the surprise in either direction. Although it initially increases significantly, most point estimates going forward are insignificantly different from zero. The response of industrial production is reported in the top right panel. According to the textbook theory of monetary policy, production should contract following a monetary tightening. The graph indicates that this is the case. The unemployment rate (bottom left panel) increases slowly but significantly so. The real interest rate increases after the monetary policy shock, but then returns to zero after about 1.5 years. Most of the subsequent estimates are insignificant.

In addition to standard macro aggregates, Figure 10 in the appendix plots the change in aggregate earnings (i.e., average earnings across all individuals) and the average probability of remaining employed for our full sample. The left panel shows that the response of earnings to a contractionary monetary policy surprise of one-standard-deviation builds up gradually, reaching a point estimate of about 0.5 percentage points after two years. This reduction in average earnings is accompanied by an increase in the probability of transitioning into unemployment (a fall in the probability of being employed). Twelve months after the shock, the probability of remaining employed decreases by 0.2 percentage points.

4 The impact of monetary policy across the distribution

4.1 Earnings growth

Our aim is to estimate the effect of monetary policy surprises on earnings growth rates, separately for individuals in different quantiles of the permanent income distribution. As described in Section 2, we sort individuals into quantiles based on their permanent income in period $t - 1$. We split the distribution into 20 quantiles, or ventiles, conditioning on age and

Figure 1: Aggregate responses to monetary policy surprises



Note: The figure shows the impulse responses of aggregate variables to a surprise one standard deviation increase in the policy interest rate, estimated using the LPIV outlined in Equation (1). The *Top Left Panel* shows the change in the inflation rate, calculated as the change in the logarithm of the HICP for Germany. The *Top Right Panel* shows the percentage change in industrial production, calculated as the log difference, and the *Bottom Left Panel* shows the change in the unemployment rate. The *Bottom Right Panel* shows the change in the real interest rate, calculated as the inflation rate subtracted from the policy rate. The sample period is from 2000 until 2013. The shaded areas indicate 68 percent confidence intervals.

gender. For each of these quantile groups, we first compute average earnings as

$$\overline{earn}_{t+h}^q = \frac{1}{N^q} \sum_{i=1}^{N^q} earn_{i,t+h} \quad \forall i \in q \text{ at } t-1$$

where $earn_{i,t+h}$ represents the labor earnings of an individual in month $t+h$ who was sorted into quantile q in month $t-1$. Since labor earnings $earn_{i,t+h}$ are zero in unemployment, *individual* earnings growth rates between two periods, especially in log-differences, are only defined for the continuously employed in our sample. However, by computing *average* earnings in $t+h$, we can subsume zero-earners into this aggregate and retain them when computing

earnings growth rates.

Specifically, in Equation (1), we set $x_{t+h} = \log(\overline{earn}_{t+h})$ and, for each quantile, estimate

$$\Delta \log(\overline{earn}_{t+h}^q) = \alpha_h + \beta_h^q \Delta i_t + \theta X_t + \epsilon_{t+h}^q \quad (2)$$

where $\Delta x_{t+1} = x_{t+h} - x_{t-1}$. The coefficient β_h^q captures the effect of a 100 basis point change in interest rates Δi_t , in period t , (instrumented by Z_t , as described in Section 3, following Stock and Watson (2018)) on earnings growth in quantile q between periods $t - 1$ and $t + h$.

We scale the size of the exogenous interest rate change, Δi , such that it causes a 1 percentage point increase in the growth rate of *aggregate* earnings, twelve months after the shock.¹³ This allows us to compare the change in earnings growth rates across quantiles associated with an *unconditional* one-percent change in aggregate earnings (as in Guvenen et al. (2017)) to that of a *conditional* change in aggregate earnings of equal size, caused by a monetary policy innovation:


$$\Delta \log(\overline{earn}_{t+h}^q) = \alpha_{Y,h} + \beta_{Y,h}^q \Delta \log(\overline{earn}_{t+h}) + \theta X_t + \epsilon_{Y,t+h}^q. \quad (3)$$

Here, the coefficient $\beta_{Y,h}^q$ represents the change in the quantile-specific earnings average in response to an unconditional change in the overall earnings average.

The blue line in the left panel of Figure 2 reports the quantile-specific earnings changes between months $t - 1$ and $t + 12$, induced by an exogenous interest rate change which raises aggregate earnings by one percentage point over the same period. Recall that, since the maximum length of an employment spell in our dataset is twelve months, earnings growth between $t - 1$ and $t + 12$ is always computed using earnings observation drawn from two different employment spells.

¹³For reference, the left panel of Figure 10 plots the response of aggregate earnings in our sample to an exogenous *one-standard-deviation* rise in the policy rate. Aggregate earnings fall by roughly 0.3 percentage points. Hence, to induce a 1 percentage point rise in aggregate earnings, the policy rate must fall by three times as much.

Figure 2: Regression coefficients β_{12}^q across the income distribution



Note: The *Left Panel* plots the coefficients β_{12}^q in Equation (2) (scaled by an expansionary monetary policy surprise consistent with a one-percent increase in aggregate earnings) and $\beta_{Y,12}^q$ in Equation (3), separately for individuals who shared the same ventile of the permanent income distribution in period $t - 1$. Income growth is computed as the log-change in the average income of individuals who were in the same ventile at time $t - 1$. The *Right Panel* compares the coefficients β_{12}^q for the full sample in a ventile (gray dashed line) to $\beta_{12}^{q,E}$ and $\beta_{Y,12}^{q,E}$, estimated on a smaller sample of individuals who are employed both in period $t - 1$ and $t + 12$ (the blue and red lines, respectively). Ventiles are constructed based on average earnings during the five years prior to $t - 1$, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

Earnings changes in response to expansionary monetary policy exhibit a pronounced U-shape across the permanent-income distribution. In particular, the earnings of the poorest individuals, in the bottom ventile, respond almost three times as much as earnings at the median. Moving up the income distribution, this response declines strongly in magnitude, to about two-thirds of the median effect, in ventiles 15 to 19. Earnings of the income-rich, in the top ventile, respond more, about twice as strong as median earnings. This up-tick is somewhat surprising, since we drop top-coded observations, but is likely driven by the income fluctuations of individuals whose earnings are close to, but below the top-code limit.

The red line in the left panel of Figure 2 depicts the point estimates $\beta_{Y,12}^q$, summarising the comovement of individual and aggregate earnings growth rates without conditioning on monetary policy surprises. As documented in Guvenen et al. (2017) for the US economy, this comovement also has a U-shaped relationship with the level of individual permanent incomes, very similar to that of earnings changes due to monetary policy (although with a somewhat less pronounced increase in the extreme ventiles).

The estimates of β_{12}^q , depicted in the left panel of Figure 2, conflate the effect monetary policy has on labor earnings with the effect it has on employment probabilities, as they are based on the changes in average labor earnings of all individuals in a given quantile (including

the unemployed who have zero labor earnings). However, because average earnings, \overline{earn}_t , equal the product of the average labor earnings of the employed, \overline{earn}_t^E times the employment rate, we can compute the following decomposition

$$\log(\overline{earn}_{t+h}^q) = \log(\overline{earn}_{t+h}^{q,E}) + \log\left(\frac{N_{t+h}^{q,E}}{N^q}\right) \quad (4)$$

where $\overline{earn}_{t+h}^{q,E}$ represents the average earnings of employed individuals in month $t + h$, who were sorted into quantile q in period $t - 1$. In the second expression, $N_{t+h}^{q,E}$ represents the number of employed individuals in period $t + h$ who were sorted into quantile q in period $t - 1$. Thus, Equation (4) implies that changes in average labor earnings across quantiles are the sum of two separate effects: the changes in the labor earnings of the employed (which we denote the *intensive*-margin effect), and changes in the employment rate (*extensive*-margin effect).

To isolate the heterogeneity in the intensive-margin effect, we substitute the change in average earnings of the employed, $\overline{earn}_{t+h}^{q,E}$, in place of its full-sample counterpart \overline{earn}_{t+h}^q in Equation (2). The resulting coefficients, which we refer to as $\beta_{12}^{q,E}$, are displayed in the blue line in the right panel of Figure 2. As before, we scale the point estimates such that the initial exogenous interest rate change Δi causes aggregate earnings growth to rise by one percentage point over the subsequent twelve months.

Earnings of the employed appear to be much less affected by monetary policy surprises than earnings in the full sample. The estimates are less heterogeneous across the distribution and substantially smaller in magnitude. In response to an exogenous change to the policy rate, the earnings growth of the employed rises by about 0.7 percentage points in the first quantile. Across the first five quantiles, this effect declines, but is essentially flat between ventiles 9 and 19, before rising substantially in the top ventile. The difference between the estimates of β_{12}^q (dashed black line) and $\beta_{12}^{q,E}$ is most pronounced in the bottom ventile, where the extensive margin of employment accounts for two thirds of monetary policy's effect on average labor earnings. This role of the extensive margin declines across the income distribution, to about a quarter of the overall effect.

In the next subsection, we investigate how monetary policy affects labor market transition probabilities.

4.2 Labor market transitions

We observe each individual in our sample either as employed or as unemployed. Let $s_{i,1}$ be an individual's labor market status in period $t - 1$ and $s_{i,2}$ be the labor market status of the

same individual in some future period. Then, there are four different transitions between $s_{i,1} \in \{E, U\}$ and $s_{i,2} \in \{E, U\}$. In addition, we also identify a subset of “switchers” who are observed as employed in both periods, but with different employers ($s_2 = \text{switch}$).

For each quantile along the permanent-income distribution, we aggregate the individual transitions into transition probabilities:

$$TR_{t+h}^{q,s_1,s_2} = \frac{1}{N_{q,s_1}} \sum_{i \in q,s_1} \mathbb{I}_{s_1,s_2}.$$

According to this definition, TR_{t+h}^{q,s_1,s_2} is the fraction of all individuals who are sorted into quantile q in period $t - 1$ and observed in state s_1 at $t - 1$ (N^{1,s_1}), who have transitioned to state s_2 by period $t + h$.

Similarly to Equation (2), we then estimate the following regression separately for each quantile-subsample:

$$TR_{t+h}^{q,s_1,s_2} = \alpha + \gamma_h^{q,s_1,s_2} \Delta i_t + \theta X_t + \epsilon_{t+h}^{q,s_1,s_2} \quad (5)$$

where the coefficient $\gamma_h^{s_1,s_2}$ measures the percentage point change in the share of individuals in state s_1 that make a particular labor market transition in response to a monetary policy surprise, for a given quantile q . Again, the vector X_t contains calendar-month dummies and three lagged values of Δi_t and Z_t .

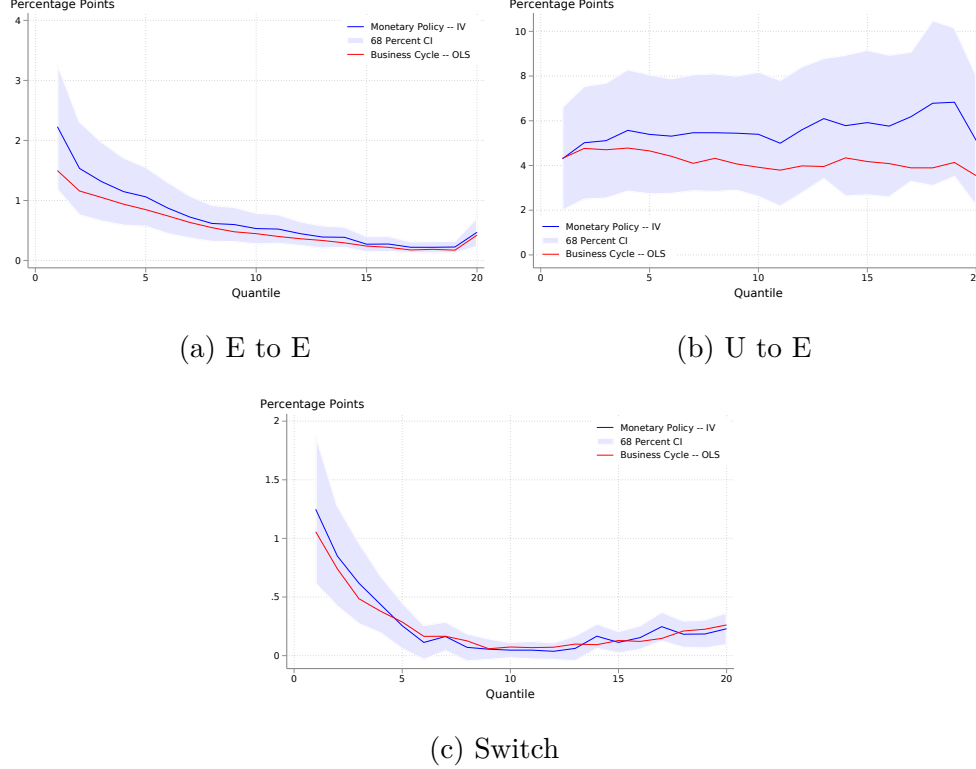
The blue line in the top left panel of Figure 3 shows the point estimates for $\gamma_{12}^{q,EE}$ (again scaled by an expansionary monetary policy surprise consistent with a one-percent increase in aggregate earnings), summarising the effect of a monetary policy surprise on transitions from employment to employment. As with earnings, we document strong heterogeneity in the incidence of monetary policy surprises along the income distribution. For the poorest individuals in the sample, the interest rate change decreases the probability of moving to unemployment by on average two percentage points. Moving up the income distribution this effect declines monotonically to less than 0.5 percentage points. The top ventile is again affected somewhat more strongly.

Analogous to section 4.1, we can compare the estimates conditional on monetary policy with those of unconditional comovement between transition probabilities and aggregate earnings changes.¹⁴ The resulting coefficients are displayed as the red line in the top left panel of Figure 3. Interestingly, the reduction in transitions into unemployment is somewhat more pronounced for the expansionary monetary policy shock than for an unconditional

¹⁴The regression is of the same form as Equation 5, with Δi substituted for with changes in aggregate income ΔY . We label the resulting coefficient $\gamma_{Y,h}^{q,s_1,s_2}$.

increase in aggregate earnings. The difference is largest at the low end of the permanent income distribution.

Figure 3: Regression coefficients γ_{12}^q across the income distribution



Note: The *Top Left Panel* plots the coefficients $\gamma_{12}^{q,EE}$ in Equation (5) (scaled by an expansionary monetary policy surprise consistent with a one-percent increase in aggregate earnings, blue line) and $\gamma_{Y,12}^{q,EE}$ (red line), from a version of Equation (5) which quantifies unconditional comovement (see text). Both quantify the change in transition probabilities for the employed in $t - 1$ to employment in period $t + 12$ (E to E). The *Top Right Panel* plots the scaled coefficients $\gamma_{12}^{q,UE}$, and $\gamma_{Y,12}^{q,UE}$, for the share of unemployed transiting to employment (U to E). The *Bottom Panel* plots the scaled coefficient $\gamma_{Y,12}^{q,switch}$ and $\gamma_{Y,12}^{q,switch}$ for the share of the employed who change employment relation. Ventiles are constructed based on average earnings during the five years prior to $t - 1$, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

The top right panel of Figure 3 shows the scaled point estimates for γ_{12}^{UE} , summarising the effect of an expansionary monetary policy surprise on the probability of unemployed individuals transitioning to employment. This effect is on average more than 5 percentage points. Contrary to the stronger effect on the likelihood of E-to-E transitions, U-to-E transitions respond slightly less to monetary policy at the bottom of the distribution. In particular, while monetary policy shocks affect the transition probabilities of the income-poor similarly to average fluctuations (as summarised by their comovement with average earnings, in the red line), a gap between the two opens up along the income distribution.

The results in the top panels of Figure 3 thus show that the substantially stronger extensive-margin effect of monetary policy on employment shares of the poor is largely accounted for by their more responsive employment-to-employment transitions. The bottom panel of Figure 3 further investigates the source of this heterogeneity. It shows the scaled point estimates for $\beta_{TR,12}^{switch}$, summarising the effect of monetary policy surprises on the frequency of transitions between two different employment relationships. An expansionary monetary policy surprise makes job-switching more likely in the bottom quartile, but has little effect in the rest of the distribution. A similar pattern holds for the effect on job-switching of unconditional fluctuations in average earnings.


4.3 Inequality

The previous results beg the question how inequality in labor earnings develops in response to changes in monetary policy. To investigate this, we substitute values of the aggregate monthly Gini coefficient, $gini_{t+h}$, for x in Equation (1). Importantly, we include unemployed individuals in our calculations, with their labor earnings set to zero, as above.

The left panel in Figure 4 plots the change in the Gini coefficient in response to an expansionary monetary policy shock over time.¹⁵ Inequality falls significantly, for two years after the shock, then reverts back. Throughout our sample period, the average value of the Gini coefficient is close to 0.3, implying that monetary policy has economically significant effects on this measure, decreasing it by close to five percent at the trough of the impulse response function in Figure 4.

¹⁵As before, the monetary policy surprise is scaled to cause aggregate earnings to rise by one percentage point over twelve months.

Figure 4: Gini coefficient Impulse Response



Note: The *Left Panel* shows the change in the Gini coefficient of labor earnings (including zeros), $gini$, in response to an expansionary monetary policy surprise, consistent with a one-percent increase in aggregate earnings, over time. The *Right Panel* shows the change in the Gini coefficient of labor earnings, including unemployment benefit receipts, $gini^{UI}$, in response to an expansionary monetary policy surprise, consistent with a one-percent increase in aggregate earnings, over time. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

Because our dataset also includes some information about unemployment benefit receipts, we can calculate the Gini coefficient taking these benefits into account. As before, we substitute $gini_{t+h}^{UI}$ into Equation (1) and compute the impulse response of this statistic to an expansionary monetary policy surprise. The implied change in inequality, plotted in the right panel of Figure 4, is substantially smaller, compared to the case when the unemployed’s earnings are set to zero. Although $gini^{UI}$ decreases significantly, after about two years, the change is economically small: two percent relative to its average value. The unemployment benefit system, therefore, appears to attenuate the effect that monetary policy has on inequality.

4.4 Earnings and labor market prospects after unemployment

Figure 3 shows that much of the effect of monetary policy on average earnings, and most of its heterogeneous incidence, is due to the response of labor market transitions between employment and unemployment. Because the costs of unemployment are strongly affected by its duration and effect on future earnings, this section investigates the effect of monetary policy shocks on re-employment probabilities and earnings after unemployment. We focus on two groups of individuals: employed who become non-employed in the period of the surprised (t) and those who retain their jobs in t . We then investigate how earnings of the second group evolve relative to the first, and how monetary policy affects the difference. For $k = -6, -5, \dots, 36$ we run the following regression for both groups and for three terciles of


the permanent income distribution:

$$x_{t+k} = \alpha_{x,h} + \gamma_{x,h}\Delta i_t + \theta_{x,h}X_t + \epsilon_{x,t}. \quad (6)$$

where, $x \in \{earn, emp\}$ corresponds to monthly individual earnings ($earn_{i,t}$) in levels or an indicator variable $emp_{i,t}$ that takes the value 1 when an individual is employed, and 0 otherwise. Again, Δi_t represents the interest rate change in period t , instrumented using Z_t as before, and X_t contains calendar-month dummies and three lags of the interest rate change as well as the instrument. In Equation (6), $\alpha_{x,h}$ equals the average earnings or employment h months after an unemployment shock in the absence of monetary policy surprises, for both groups. In turn, $\gamma_{x,h}$ quantifies the impact of monetary policy on these variables. The regressions are similar in spirit to that in [Davis and Von Wachter \(2011\)](#), who also investigate earnings paths of the unemployed relative to those who remain employed. We focus on individuals who become unemployed after an employment spell that lasted at least 6 months. Because this substantially reduces the sample size, we report results for terciles, rather than ventiles, of the permanent income distribution.

Figure 5 shows the results of the exercise for employment probabilities for the first tercile. We include results for the second and third terciles in Appendix Figure 14. The red line in the right panel shows the probability of being employed for individuals who transitioned to nonemployment in period 0 *relative to those who did not*. Similarly, the green line in the same graphs shows how an expansionary monetary policy surprise, which causes aggregate earnings to rise by one percent over twelve months, affects this probability. An expansionary monetary policy shock increases employment probabilities by about four percent. This is in line with the results reported in Figure 3, for more granular sorting. The top right in Figure 5 isolates the effect of a monetary policy shock on employment probabilities in the first tercile of permanent income, of those becoming unemployed in period 0 *relative to those who don't*. For the first six months, the effect is relatively small and insignificant. After one year, however, it grows to four percent and stays significant for an extended period of time. The results for the other terciles are similar.

Figure 5: Effect of monetary policy shock on re-employment probabilities

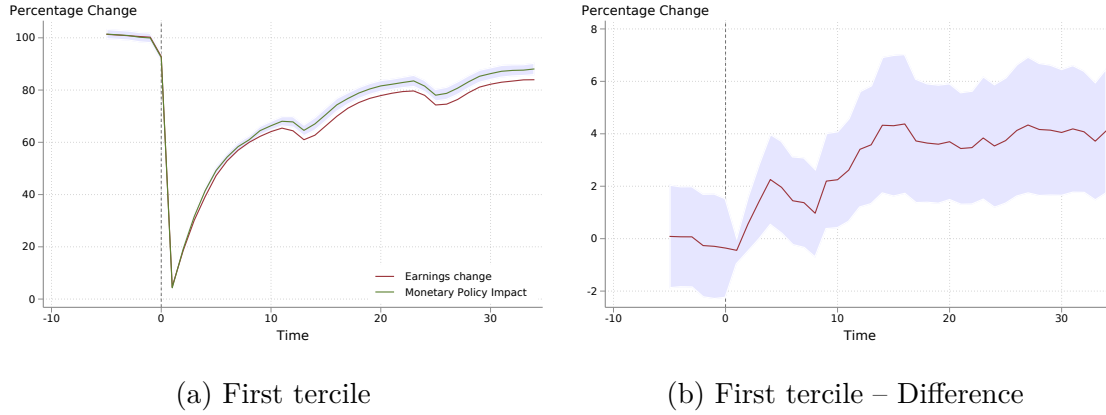


Note: The panels show the employment probability of individuals who transition into unemployment in month $t = 0$, *relative to those who don't*, with and without a monetary policy surprise (scaled to be consistent with a one-percent increase in aggregate earnings) in the blue and red lines, respectively, over time. The *Top Left, and Bottom Panels* show results for three subsamples comprising individuals in the lower, middle and upper tercile of the permanent income distribution, respectively. The *Top Right Panel* shows the difference between the red and blue lines in the top left panel, corresponding to the effect of a monetary policy surprise on re-employment earnings of individuals in the lowest tercile. Terciles are constructed based on average earnings during the five years prior to the unemployment spell in period 0, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

Figure 6 shows the results for the same exercise using earnings as the dependent variable. Again, the results in the right panel represent earnings of individuals who become unemployed in period 0 *relative to those who don't*. Naturally, earnings approach zero in the first month after the transition into unemployment. However, some individuals find new jobs in the same month, implying positive average earnings.¹⁶ An accommodative monetary policy surprise steepens the slope of this recovery. Here, as shown in the right panel, such a shock increases earnings by about 4 percent, 20 months after unemployment. These effects are stronger than for the other two terciles, plotted in Appendix Figure 15.

¹⁶Note that an individual is counted as employed if they were employed for at least half of the month. In the figure, the noticeable drop in earnings at period 0 is explained by individuals transitioning to unemployment towards the end of month 0, but still counting as employed in said month.


Figure 6: Effect of monetary policy on average earnings after unemployment



Note: The panels show the average earnings of individuals who transition into unemployment in month 0 with and without monetary policy surprise (scaled to be consistent with a one-percent increase in aggregate earnings), in the blue and red lines, respectively, over time. The *Top Left, and Bottom Panels* show results for three subsamples comprising individuals in the lower, middle and upper tercile of the permanent income distribution, respectively. The *Top Right Panel* shows the difference between the red and blue lines in the top left panel, corresponding to the effect of a monetary policy surprise on re-employment earnings of individuals in the lowest tercile. Terciles are constructed based on average earnings during the five years prior to the unemployment spell in period 0, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

Since the sample contains some information on unemployment benefits, we can investigate how these benefits affect earnings changes after unemployment. Specifically, we substitute the benefit income for the zero-earnings assumption used in the previous analysis. Figure 7 shows the results of this exercise. The fall in earnings upon unemployment is much less pronounced. Whereas before, earnings fell close to zero upon unemployment, the fall is now closer to 40 percent. However, as the left panel shows, monetary policy still has a significant impact on earnings. After a year, earnings of individuals who lose their employment during periods of accommodating monetary policy have recovered on average about two percent more than those who become unemployed in normal times. These effects are similar to the other two terciles, plotted in Appendix Figure 16.

Figure 7: Effect of monetary policy on post-benefit earnings after unemployment



Note: The panels show the average earnings, including unemployment benefit receipts, of individuals who transition into unemployment in month $t = 0$, with and without monetary policy surprise (scaled to be consistent with a one-percent increase in aggregate earnings), in the blue and red lines, respectively, over time. The *Left Panel* shows results for the subsample comprising individuals in the lower tercile of the permanent income distribution. The *Right Panel* shows the difference between the red and blue line in the left panel, corresponding to the effect of a monetary policy surprise on re-employment earnings of individuals in the lowest tercile. Terciles are constructed based on average earnings during the five years prior to the unemployment spell in period 0, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

Figure 8 summarizes the heterogeneous incidence of monetary-policy shocks (again normalized to imply a 1-percentage-point increase in aggregate earnings) on earnings and employment probabilities following an unemployment event. In line with our previous results, the estimated responses are strongest for income-poor workers, in the first tercile, in all three panels of Figure 8. The point estimate of the 12-month response of pre-employment earnings for the first tercile, for example, is 3 percentage points larger than for the second. In line with the U-shape reported in Figure 2, the estimated responses for the third tercile are on average larger than for the second, but less consistent over time and imprecisely estimated as they aggregate the strong responses at the very top of the distribution and the weaker responses just below in Figure 2.

Figure 8: Effect of monetary policy on the recently unemployed



Note: The *Left Panel* show the effect of a monetary policy surprise (scaled to be consistent with a one-percent increase in aggregate earnings), on employment probabilities of individuals who become unemployed in period 0 across terciles. The *Right Panel* shows the effect of the same monetary policy surprise on earnings. The solid line represents estimates for the first tercile, the dashed line those for the second and the dash-dotted line those for the third. The shaded area represents 68 percent confidence bands on the estimates of the second tercile. Terciles are constructed based on average earnings during the five years prior to the unemployment spell in period 0, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

5 Implications for Aggregate Demand

Recent literature has pointed out that cyclical variations in employment risk may act as an amplifying mechanism for business cycles (Broer et al., 2021; Graves, 2020; Ravn and Sterk, 2021). If workers reduce consumption and build up precautionary savings when separation risk rises (in recessions), the resulting contraction in demand could deepen the downturn. In this section, we explore how the heterogeneity we document in Section 4.1 affects the dynamics of aggregate demand in response to monetary policy shocks. We differ relative to previous analyses that abstract from heterogeneous incidence of unemployment risk (Acharya and Dogra, 2020; Auclert, 2019; Patterson et al., 2019). In our analysis, we follow Werning (2015) and focus on the household “demand block” without explicitly specifying the supply side of the economy. We consider an extension of his framework to account for job-finding and separation risk, and heterogeneity in risk across the distribution of earnings.

5.1 A Model with heterogeneous employment risk

The economy is populated by a unit measure of households. There is a finite set of household types indexed by $i \in I$, with measures $\mu^i > 0$, such that $\sum_{i \in I} \mu^i = 1$. Households have

identical preferences over consumption:

$$u = \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t U(c_t(s^t)) \quad (7)$$

where c_t is consumption, $s_t \in S^i$ represents an idiosyncratic shock that follows a stochastic process that is identical and independent across households. s^t is the history of such shocks from 0 to t . In particular, we assume $U(c) = \frac{c^{1-\sigma}}{1-\sigma}$ (CRRA preferences). Household income $y_t^i(s_t)$ can depend on realizations of the idiosyncratic shock s_t and aggregate income Y_t according to:

$$y_t^i(s_t) = \gamma_t^i(s_t, Y_t). \quad (8)$$

Households can save and borrow in one-period risk-free bonds, b_t , subject to a type-specific borrowing constraint $-B_t^i(s_t, Y_t)$. The household budget constraint is thus given by:

$$c_t + b_t \leq y_t^i(s_t) + R_{t-1}b_{t-1} \quad (9)$$

where R_t is the real interest rate on risk-free bonds.

The cross-sectional distribution of households across idiosyncratic states and bond holdings is denoted $\Omega(s, b)$. The equilibrium definition is standard, such that households optimize, markets clear and the evolution of the cross-sectional distribution is generated from the optimal policy choices of households and the stochastic process for s_t .

5.1.1 Cyclical Employment Risk

Our empirical findings in Section 4.1 illustrate that:

1. Conditional on staying employed, the cyclical growth of earnings is (approximately) uniform across the income distribution, and earnings growth is less cyclical than aggregate earnings growth
2. Employment risk is countercyclical, and substantially more so at the bottom of the distribution

We can capture these empirical findings in the model with the following specification, extending the model proposed by [Werning \(2015\)](#) with type-specific countercyclical employment risk. Households can be either employed ($\epsilon_t^i = 1$) or unemployed ($\epsilon_t^i = 0$). Define the type-specific employment rate as $e_t^i = \mathbb{E}\epsilon_t^i$. Aggregate earnings can thus be characterized by:

$$Y_t = \sum_{i \in I} \mu^i \left[e_t^i \bar{y}_i Y_t^{\psi^i} + (1 - e_t^i) \underline{y}_i Y_t^{\psi^i} \right] \quad (10)$$

where we assume that household income follows aggregate income according to $y_t^i(s_t) = \bar{y}_i Y_t^{\psi^i}$ for the employed and according to $y_t^i(s_t) = \underline{y}_i Y_t^{\psi^i}$ for the unemployed. The parameter ψ^i is the elasticity of income with respect to the aggregate, conditional on not changing employment states. Employment by type evolves according to a type-specific separation rate $\delta^i(Y_t)$ and job-finding rate $f^i(Y_t)$:

$$e_t^i = \left(1 - \delta^i(Y_t) \left(1 - f^i(Y_t)\right)\right) e_{t-1}^i + f^i(Y_t) \left(1 - e_{t-1}^i\right) \quad (11)$$

where the timing is such that we allow for households that lose their job to find one immediately within the period.

5.1.2 The dynamics of aggregate demand

In traditional macroeconomic models with a representative household or complete markets, the dynamics of aggregate demand are characterized by the dynamic optimality condition, or Euler equation, for aggregate consumption. With incomplete markets and idiosyncratic risk, there is typically no such condition. To characterize the effect of heterogeneous incidence on aggregate demand in a transparent way, we follow [Werning \(2015\)](#) and consider the zero-liquidity limit of the economy (where there is no asset trade in equilibrium), which allows the individual optimality conditions to be aggregated to a condition for aggregate consumption demand that is similar to the representative household's Euler equation in the absence of heterogeneity. In particular, as originally discussed in [Krusell et al. \(2011\)](#), the [Huggett \(1993\)](#) economy with the tightest borrowing limit ($B_t^i(s_t, Y_t) = 0 \forall t$) generates a well defined stochastic discount factor for the economy. The agent with the strongest incentive to save is the one who prices the risk-free bond.¹⁷ In our setting, it is clear that one of the employed agents will be this ‘marginal saver’ (since unemployed agents have positive expected income growth and would thus like to borrow). The first-order condition for employed agents is given by:

$$U'(\bar{y}^i Y_t^{\psi^i}) = \beta R_t \mathbb{E}_t \left[\underbrace{\left(1 - \delta^i(Y_{t+1}) \left(1 - f^i(Y_{t+1})\right)\right) U'(\bar{y}^i Y_{t+1}^{\psi^i})}_{\text{remains employed}} + \underbrace{\delta^i(Y_{t+1}) \left(1 - f^i(Y_{t+1})\right) U'(\underline{y}^i Y_{t+1}^{\psi^i})}_{\text{becomes unemployed}} \right] \quad (12)$$

¹⁷In principle, this provides only a lower bound for the equilibrium bond price, as at any higher price all agents would simply be constrained by the zero-borrowing limit. The equilibrium bond price is unique and equal to this bound, however, when there is an arbitrarily small supply of bonds.

where the right-hand side equals the expected marginal utility across employment and unemployment states. The bond price is thus determined as:

$$1 = \beta R_t \max_{i \in I} \mathbb{E}_t \left[\frac{(1 - \delta^i(Y_{t+1}) (1 - f^i(Y_{t+1}))) U'(\bar{y}^i Y_{t+1}^{\psi^i})}{U'(\bar{y}^i Y_t^{\psi^i})} + \frac{\delta^i(Y_{t+1}) (1 - f^i(Y_{t+1})) U'(\underline{y}^i Y_{t+1}^{\psi^i})}{U'(\bar{y}^i Y_t^{\psi^i})} \right] \quad (13)$$

Using the fact that $U'(c) = c^{-\sigma}$, and the in the aggregate $Y = C$, we can characterize the equilibrium.

Proposition 1. *In the economy we have the aggregate Euler relation:*

$$\hat{U}'(C) = \hat{\beta}(C') R \mathbb{E} \hat{U}'(C') \quad (14)$$

where $\hat{U}'(C) \equiv C^{-\sigma\psi_j}$, where $j = \arg \max$ of the discount rate function $\hat{\beta}$, which is decreasing and given by:

$$\hat{\beta} \equiv \beta \max_{i \in I} \left((1 - \delta^i(C') + \delta^i(C') f^i(C')) + \delta^i(C') (1 - f^i(C')) U'(\underline{y}^i / \bar{y}^i) \right)$$

Proof. Straightforward extension of [Werning \(2015\)](#) Proposition 4. □

Equation (14) characterises the dynamic response of aggregate consumption demand to shocks. In particular, two elements determine the response of C to any shock in the economy. First, as in standard representative-agent models, the elasticity of intertemporal substitution governs the aggregate response of consumption to an unexpected and purely temporary change in the real interest rate R . However, in contrast to the representative agent case, in our framework with (heterogeneous) unemployment risk, that elasticity depends on the responsiveness of individual incomes to aggregate income governed by ψ^j . If $\psi^j < 1$, the Euler relation implies that consumption is more elastic to the interest rate under incomplete markets than when markets are complete, as in [Werning \(2015\)](#). Effectively, the elasticity of intertemporal substitution rises from $\frac{1}{\sigma}$ to $\frac{1}{\sigma\psi^j}$. Intuitively, if earnings respond little to changes in aggregate output ($\psi \rightarrow 0$), changes in output must feed into earnings through the extensive margin, increasing risk.

The second element in equation (14) which determines the dynamics of aggregate demand is the discount factor $\hat{\beta}$ in (14), equal to the equilibrium change in current aggregate consumption C in response to a given change in future consumption C' . In our framework, $\hat{\beta}$ depends on the time-varying probability of staying employed. If this probability is pro-cyclical (as we

show in Section 4.1), then the discount factor $\beta(C)$ is decreasing in aggregate consumption. As a result, contemporaneous consumption responds more than one for one to changes in future consumption, implying that changes to *future* interest rates have a greater effect on today's consumption than changes in the current interest rate.

5.2 Quantification

Equation (14) allows us to quantify the extent of amplification implied by incomplete markets using our empirical evidence for Germany. To do so requires estimates for 1) the elasticity of earnings conditional on staying employed by type ψ^i ; 2) the ratio of the earnings of employed and unemployed by type $\underline{y}^i/\bar{y}^i$; 3) functions for the separation rate as a function of aggregate earnings by type $\delta^i(Y)$; and 4) functions for the job-finding rate as a function of aggregate earnings by type $f^i(Y)$. We focus on two types of households, $i = L, H$, corresponding to the top and bottom halves of the permanent income distribution.

We can measure the ψ^i directly as the response of individual earnings to aggregate income β_q for individuals employed in $t - 1$ and $t + 12$ in Figure 2. The low earnings types have $\psi^L = 0.39$ and the high earnings types $\psi^H = 0.47$ (see Table 2). As discussed in the previous section, the heterogeneity in sensitivity of earnings conditional on being employed is small.

The ratio of earnings by type $\underline{y}^i/\bar{y}^i$ is set to be equal to 0.95 and represents the consumption fall upon unemployment, both for low- and high-earnings individuals Kolsrud et al. (2018).

Next, for $\delta^i(Y)$ and $f^i(Y)$ we log-linearize the functions, yielding

$$\delta^i(Y) = d_0^i + d_1^i (\log(Y) - \log(Y_{ss})) \quad (15)$$

$$f^i(Y) = f_0^i + f_1^i (\log(Y) - \log(Y_{ss})) \quad (16)$$

The estimates for d_0^i and f_0^i are based on the average separation rate into non-employment and job-finding rate from non-employment over the sample. The regression coefficients $\gamma_{12}^{q,E,U}$ and $\gamma_{12}^{q,U,E}$ in Figure 3—the effect of a one percent change in aggregate earnings induced by a monetary surprise on the probability of going from E to U and E to E—identify d_1^i and f_1^i , respectively. The estimates are summarized in Table 2.

We study the effect of small shocks, implying small fluctuations in employment risk and earnings. Incentives to save for the employed are thus mainly governed by the average (or ‘steady-state’) probability to move to unemployment d_0^i , which is highest for low-income individuals (see Table 1). L type households are thus the marginal savers that determine the equilibrium dynamics of aggregate demand via equation (14). The effective elasticity of intertemporal substitution in the Euler relation is given by $\frac{1}{\psi^L \sigma} \approx 2.6 \times \frac{1}{\sigma}$, implying that consumption is more than twice as elastic to interest rates, as compared to a framework with

Table 2: Parameter values

	(1)	(2)	(3)
	Pooled	Low Type	High Type
MP effect on earnings of E to E (ψ^i)	0.44	0.39	0.47
MP effect on separations (d_1^i)	-0.59	-0.93	-0.32
Steady-state separations (d_0^i)	3.20%	4.86%	1.73%
MP effect on job-finding (f_1^i)	5.07	4.94	5.69
Steady-state job-finding (f_0^i)	33.82%	32.66%	41.08%

complete markets or a representative agent. Thus, incomplete markets generate substantial amplification of the response of consumption to contemporaneous interest rate changes.

In addition to the higher effective EIS, the presence of cyclical separation risk implies that the effective discount factor $\hat{\beta}$ in the Euler relation will be decreasing in aggregate earnings. Differentiating the log-linearized expression in Proposition 1 and using $C = Y$ yields:

$$\frac{\partial \hat{\beta}}{\partial C} = \left(\left[\frac{y^L}{\bar{y}^L} \right]^{-\sigma \psi^L} - 1 \right) (d_1^L - d_0^L f_1^L - d_1^L f_0^L - 2d_1^L f_1^L dC). \quad (17)$$

Assuming a baseline EIS of $1/2$ ($\sigma = 2$) yields $\frac{\partial \hat{\beta}}{\partial C} = -0.084$. Future increases in aggregate consumption induce more than a one-for-one movement in aggregate consumption today. To first order, a one percent increase in consumption in the next period would lead to a 1.1% increase in consumption today. This provides a sufficient statistic for the extent of demand amplification as a result of market incompleteness and heterogeneous incidence of cyclical earnings risk.

5.2.1 Pooled risk counterfactual

Past research (Auclert, 2019; Patterson et al., 2019) has focused mostly on the unequal incidence of level earnings changes, but not on the unequal incidence of *risk* across the distribution. Our empirical findings highlight that the incidence of earnings changes conditional on being employed is homogeneous across the earnings distribution (0.39 for the bottom half, 0.47 for the top half, and a pooled estimate of 0.44). The unequal incidence in earnings is driven by the unequal incidence of the risk of moving out of employment. The first fact implies that heterogeneous incidence does not affect the elasticity of substitution much beyond the effect of cyclical individual earnings per se. The dynamics of the discount factor $\hat{\beta}$, in contrast, are substantially changed by heterogeneous incidence. In particular, the more cyclical unemployment risk of L types increases the cyclicity of demand through

this channel.

One way to highlight the importance of this unequal incidence of risk is to compute the Euler relation from Proposition 1 based on pooled data (i.e., without heterogeneous types based on past earnings). The coefficients for that experiment are in Column 1 of Table 2.


First, the effective increase in the EIS is lower in the pooled version (increasing by a factor of 2.3 vs 2.6 in the heterogeneous case). Second, the derivative of the discount factor with respect to aggregate consumption is about 40% smaller ($\frac{\partial \hat{\beta}}{\partial C} = -0.053$ as opposed to -0.084). Thus, ignoring the unequal incidence of risk would imply significantly less demand amplification to interest rates and future consumption.

5.2.2 Impulse responses

We illustrate the amplification of responses to monetary policy shocks by computing the impulse response to a persistent monetary policy shock. We decompose the response of consumption into channels driven by cyclical unemployment risk and its heterogeneous incidence. We compare three versions of the model, one with heterogeneous risk as in the data, one with homogeneous risk, and one with perfectly insured idiosyncratic risk. First, we set the real interest rate in Equation (14) equal to its average value in the data (2.4%). Then we compute the steady-state discount factors β necessary to solve the Euler equation under heterogeneous and pooled risk. We obtain $\beta_{het} = 0.972$ and $\beta_{pool} = 0.973$. Into each of the two economies, we introduce a monetary policy shock of 25 basis points that decays following an AR(1) process with persistence of 0.7.

Figure 9 shows the impulse responses to this shock in heterogeneous and pooled risk cases, as well as the representative agent case (with perfectly insured idiosyncratic risk). The representative-agent economy experiences the smallest initial drop in aggregate output: less than 50 basis points. In the pooled-risk economy, output drops by slightly more than one percentage point on impact, as the rise in unemployment risk contracts consumption demand by more than in the representative-agent case. The impulse response in the economy with the heterogeneous incidence of risk further amplifies the drop in consumption on impact up to 1.5 percentage points. Thus, accounting for the heterogeneous incidence of monetary policy on unemployment risk yields significantly larger output responses.

Figure 9: Model impulse responses



Note: The figure shows the impulse responses of output in the representative agent economy (dash-dotted line), the economy with the homogeneous incidence of unemployment risk (dashed line), and the one with the heterogeneous incidence to a monetary policy shock (solid line). The shock raises the real interest rate by 25 basis points and then follows an AR(1) process with persistence of 0.7. For more details on the respective models and their calibration, see the text.

6 Conclusion

Monetary policy surprises affect income growth substantially more at the bottom of the earnings distribution. This heterogeneous incidence is mainly driven by the response of separation rates for the poor. Job-finding rates and earnings growth of the employed are both procyclical, but with little differences across the distribution. While our findings are for Germany, we believe that they are most likely applicable more broadly. First, the heterogeneous incidence of risk is apparent for all changes in aggregate earnings (not just monetary surprises). Second, a larger elasticity of earnings of the poor has been documented across a large number of countries. Thus, it is reasonable to conclude that the mechanism that drives that larger elasticity in Germany is likely at play in other industrialized economies.

Using a general incomplete-markets setting with type-specific cyclical unemployment risk, we showed how the concentration of movements in separation risk among high-risk workers amplifies aggregate-demand responses to interest rate changes by making precautionary savings more volatile. Based on our estimates, this makes consumption more than twice as responsive to interest rates as would prevail under complete markets, and substantially more responsive than in a model with homogeneous risk.

Our work suggests that the burgeoning HANK literature needs to take seriously the documented heterogeneity in employment dynamics across the income distribution and that it should incorporate it explicitly into its analyses. Our findings also suggest that studying policies that aim to reduce this heterogeneous income risk can significantly reduce aggregate fluctuations. We leave the study and design of such policies for future work.

References

- Acharya, S. and K. Dogra (2020). Understanding hank: Insights from a prank. *Econometrica* 88(3), 1113–1158.
- Almgren, M., J. E. Gallegos, J. Kramer, and R. Lima (2019). Monetary policy and liquidity constraints: Evidence from the euro area. *Available at SSRN 3422687*.
- Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review* 109(6), 2333–67.
- Bilbiie, F. O. (2018). Monetary policy and heterogeneity: An analytical framework. *Available at SSRN 3106805*.
- Broer, T., J. Druedahl, K. Harmenberg, and E. Öberg (2021). The unemployment-risk channel in business-cycle fluctuations.
- Challe, E. (2020). Uninsured unemployment risk and optimal monetary policy in a zero-liquidity economy. *American Economic Journal: Macroeconomics* 12(2), 241–83.
- Coibion, O., Y. Gorodnichenko, L. Kueng, and J. Silvia (2012). Innocent bystanders? monetary policy and inequality in the us. Technical report, National Bureau of Economic Research.
- Davis, S. J. and T. Von Wachter (2011). Recessions and the costs of job loss. *Brookings papers on economic activity* 2011(2), 1.
- De Nardi, M., G. Fella, M. Knoef, G. Paz-Pardo, and R. Van Ooijen (2019). Family and government insurance: Wage, earnings, and income risks in the netherlands and the us. *Journal of Public Economics* 193, 104327.
- Drews, N., D. Groll, and P. Jacobebbinghaus (2007). Programmierbeispiele zur aufbereitung von fdz personendaten in stata.
- Fitzenberger, B., A. Osikominu, and R. Völter (2005). Imputation rules to improve the education variable in the iab employment subsample. *ZEW-Centre for European Economic Research Discussion Paper* (05-010).
- Ganzer, A., A. Schmucker, P. Vom Berge, and A. Wurdack (2017). Sample of integrated labour market biographies-regional file 1975-2014. *FDZ-Datenreport 1*, 2017.

- Gertler, M. and P. Karadi (2015). Monetary policy surprises, credit costs, and economic activity. *American Economic Journal: Macroeconomics* 7(1), 44–76.
- Gornemann, N., K. Kuester, and M. Nakajima (2016). Doves for the rich, hawks for the poor? distributional consequences of monetary policy.
- Graves, S. (2020). Does unemployment risk affect business cycle dynamics? *International Finance Discussion Paper* (1298).
- Guenen, F., F. Karahan, S. Ozkan, and J. Song (2015). What do data on millions of us workers reveal about life-cycle earnings risk? Technical report, National Bureau of Economic Research.
- Guenen, F., S. Schulhofer-Wohl, J. Song, and M. Yogo (2017). Worker betas: Five facts about systematic earnings risk. *American Economic Review* 107(5), 398–403.
- Hagedorn, M., J. Luo, I. Manovskii, and K. Mitman (2019). Forward guidance. *Journal of Monetary Economics* 102, 1–23.
- Halvorsen, E., H. A. Holter, S. Ozkan, and K. Storesletten (2020). Dissecting idiosyncratic earnings risk. *Available at SSRN*.
- Hoffmann, E. B. and D. Malacrino (2019). Employment time and the cyclicity of earnings growth. *Journal of Public Economics* 169, 160–171.
- Holm, M. B., P. Paul, and A. Tischbirek (2020). The transmission of monetary policy under the microscope. Federal Reserve Bank of San Francisco.
- Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control* 17(5-6), 953–969.
- Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review* 95(1), 161–182.
- Kaplan, G., B. Moll, and G. L. Violante (2018). Monetary policy according to hank. *American Economic Review* 108(3), 697–743.
- Kolsrud, J., C. Landais, P. Nilsson, and J. Spinnewijn (2018). The optimal timing of unemployment benefits: Theory and evidence from sweden. *American Economic Review* 108(4-5), 985–1033.
- Krusell, P., T. Mukoyama, and A. A. Smith Jr (2011). Asset prices in a huggett economy. *Journal of Economic Theory* 146(3), 812–844.

- Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: the information effect. *The Quarterly Journal of Economics* 133(3), 1283–1330.
- Patterson, C. et al. (2019). The matching multiplier and the amplification of recessions. In *2019 Meeting Papers*, Number 95. Society for Economic Dynamics.
- Ravn, M. O. and V. Sterk (2017, oct). Job uncertainty and deep recessions. *Journal of Monetary Economics* 90, 125–141.
- Ravn, M. O. and V. Sterk (2021). Macroeconomic fluctuations with hank & sam: An analytical approach. *Journal of the European Economic Association* 19(2), 1162–1202.
- Stock, J. H. and M. W. Watson (2018). Identification and estimation of dynamic causal effects in macroeconomics using external instruments. *The Economic Journal* 128(610), 917–948.
- Werning, I. (2015). Incomplete markets and aggregate demand. Technical report, National Bureau of Economic Research.


Appendix

A Micro Data

We use the Sample of Integrated Labor Market Biographies (SIAB) data. The SIAB data is provided in the form of labor market spells, each at most one year in duration, reporting the average daily wage during the spell. We convert these spells into monthly observations and multiply the daily wages by 30 in order to ascertain monthly earnings. If an individual reports multiple simultaneous spells during a month, we keep the spell that is classified as “Subject to social security without special characteristics” (as classified in Table A4 of [Ganzer et al. \(2017\)](#)). If one of the simultaneous spells implies non-employment, we keep that spell and classify the individual as non-employed. We classify individuals who earn less than the lower social security contribution limit as non-employed. All non-employed workers are coded to have zero income.

We classify as unemployed those individuals who receive unemployment benefits (ALG). Because the definition and eligibility of these benefits changed over time, we declare any individuals who are non-employed but started their non-employment spell in unemployment as unemployed for the whole duration of the non-employment spell. This addresses in particular

Figure 10: Aggregate Impulse responses



(a) Regression coefficients β_h for the full sample (b) Regression coefficients γ_h for the full sample

Note: Panel a) plots the coefficient β_h in Equation (2), scaled by a one-standard-error contractionary monetary policy surprise, estimated on the whole sample. Panel b) plots the coefficients $\gamma_h^{E,E}$ for individuals who transition from employment to employment ($s_1 = s_2 = E$), again for the whole sample. The shaded area indicates 68 percent confidence bands.

the shortening of unemployment benefit eligibility around 2005. All earnings are deflated into real earnings using the monthly CPI index obtained from the OECD.

B Additional Results

B.1 Aggregate Impulse responses

Figure 10 plots the coefficients β_h in Equation (2) and $\gamma_h^{E,E}$, scaled by a one-standard-error contractionary monetary policy surprise, for the whole sample.

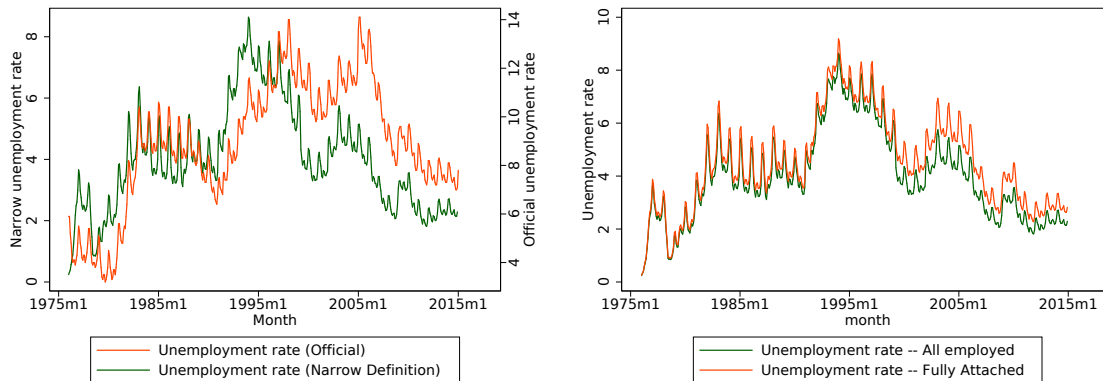
B.2 Unemployment rate

The left panel of Figure 11 compares the unemployment rate for Germany resulting from our first definition of employment to the official rate reported by the German statistics office, computed using survey data. Importantly, we include only individuals whose place of work lies in the counties that were part of the Federal Republic of Germany before 1990.¹⁸ The two rates move closely together, especially before the reunification in 1990. After 1995, the narrow unemployment rate is, however, systematically lower than the officially reported one.

The right panel of Figure 11 shows that the unemployment rates calculated using different definitions of employment behave very similarly.

¹⁸For the non-employed, location information is not available. We use the last employer's location.

Figure 11: German unemployment rates




(a) Official rate and narrow definition (only west)
(b) Narrow unemployment and fully attached (only west)

B.3 Men and Women

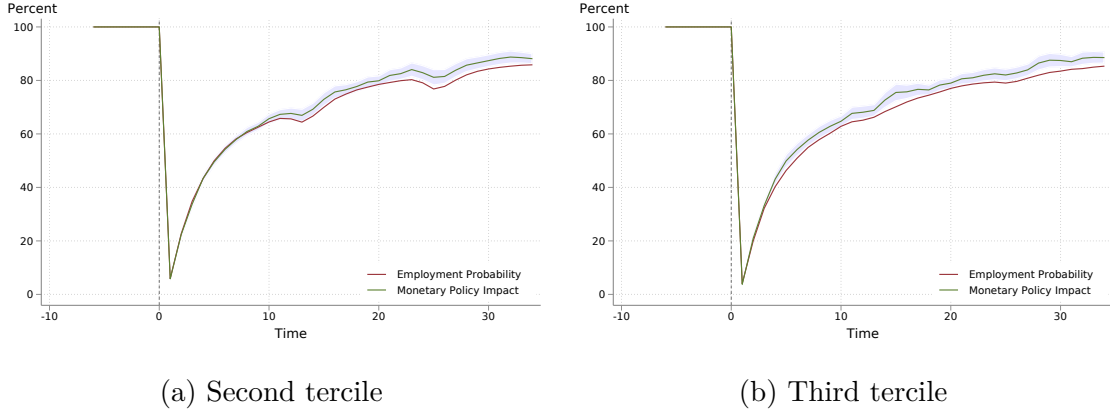
In the main text, we sort individuals into quantiles based on their average earnings between months t and $t - 60$ while conditioning on gender and age. Here, we report the procyclicality of earnings for men and women separately. Figures 12 and 13 indicate that results are not different for the two genders.

Figure 12: Regression coefficients β_{12} for prime age men and women




Note: The figure plots the coefficients β_{12}^q in Equation (2) (scaled by an expansionary monetary policy surprise consistent with a one-percent increase in aggregate earnings), separately for prime-age men and women who shared the same ventile of the permanent income distribution in period $t - 1$. Income growth is computed as the log-change in the average income of individuals who were in the same ventile at time $t - 1$. Ventiles within gender are constructed based on average earnings during the five years prior to $t - 1$. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

Figure 14: Effect of monetary policy shock on re-employment probabilities



Note: The panels show the employment probability of individuals who transition into unemployment in month $t = 0$, with and without monetary policy surprise (scaled to be consistent with a one-percent increase in aggregate earnings), in the blue and red lines, respectively, over time. The *Left, and Right Panels* show results for two subsamples comprising individuals in the middle and upper tercile of the permanent income distribution, respectively. Terciles are constructed based on average earnings during the five years prior to the unemployment spell in period 0, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.


Figure 13: Regression coefficients γ^q - Prime age men and women



Note: The *Left Panel* plots the coefficients $\gamma_{12}^{q,E,E}$ in Equation (5) (scaled by an expansionary monetary policy surprise consistent with a one-percent increase in aggregate earnings, blue line), separately for prime-age men and women. It quantifies the change in transition probabilities for the employed in $t - 1$ to employment in period $t + 12$ (E to E). The *Right Panel* plots the scaled coefficients $\gamma_{12}^{q,U,E}$, separately for prime-age men and women. Ventiles, by gender, are constructed based on average earnings during the five years prior to $t - 1$. The sample period is 2000-2013.


B.4 Earnings and labor market prospects after unemployment

Figure 15: Effect of monetary policy on average earnings after unemployment



Note: The panels show the average earnings of individuals who transition into unemployment in month $t = 0$, with and without monetary policy surprise (scaled to be consistent with a one-percent increase in aggregate earnings), in the blue and red lines, respectively, over time. The *Left, and Right Panels* show results for two subsamples comprising individuals in the middle and upper tercile of the permanent income distribution, respectively. Terciles are constructed based on average earnings during the five years prior to the unemployment spell in period 0, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.

Figure 16: Effect of monetary policy on post-benefit earnings after unemployment



Note: The panels show the average earnings, including unemployment benefit receipts, of individuals who transition into unemployment in month $t = 0$, with and without monetary policy surprise (scaled to be consistent with a one-percent increase in aggregate earnings), in the blue and red lines, respectively, over time. The *Left, and Right Panels* show results for two subsamples comprising individuals in the middle and upper tercile of the permanent income distribution, respectively. Terciles are constructed based on average earnings during the five years prior to the unemployment spell in period 0, conditional on gender and five-year age brackets. The shaded area indicates 68 percent confidence bands. The sample period is 2000-2013.