

A TAIL OF LABOR SUPPLY AND A TALE OF MONETARY POLICY

Cristiano Cantore

Sapienza University of Rome, Italy

Filippo Ferroni

University of Bologna, Italy

Haroon Mumtaz

Queen Mary, University of London, UK

Angeliki Theophilopoulou

Brunel University of London, UK

Abstract

We study the interaction between monetary policy and labor supply decisions at the household level. We uncover evidence of heterogeneous responses and a strong countercyclicality of hours worked in the left tail of the income distribution following a monetary policy shock in the US. Specifically, while aggregate hours and labor earnings decline after a monetary tightening, individuals at the bottom of the income distribution increase their hours worked. Moreover, this positive labor supply response is quantitatively significant, substantially dampening the decline in aggregate hours worked. We show that the empirical patterns are consistent with a standard one-asset HANK model featuring endogenous labor supply. The model reveals that strong income effects at the bottom of the distribution can account for the observed countercyclical labor responses, highlighting how labor supply adjustments act as an additional margin through which households smooth consumption. Comparing this specification to a model with a homogeneous labor supply, we find that labor supply heterogeneity reduces the aggregate MPC and attenuates the transmission of monetary policy through aggregate demand. As a result, the output cost of disinflation is lower in economies where poorer households can flexibly adjust their labor effort, easing the trade-off faced by the central bank. (JEL: E52, E32, C10)

The editor in charge of this paper was Guido Lorenzoni.

Acknowledgments: We thank the Editor Guido Lorenzoni and three anonymous referees. We also thank Valerio Pieroni for the lengthy discussions and suggestions about the HANK analysis. We also thank Michele Andreoli, Guido Ascari, Saleem Bahaj, Christian Bayer, Gadi Barlevy, Florin O. Bilbiie, Emanuele Brancati, Davide Debortoli, Luca Fornaro, Francesco Furlanetto, Luca Gambetti, Nezh Guner, Chris Huckfeldt, Mathias Klein, Leonardo Melosi, Carlo Pizzinelli, Ricardo Reis, Kjetil Storesletten, Dan

Journal of the European Economic Association 2026 00(0):1–30

<https://doi.org/10.1093/jeea/jvag030>

© The Author(s) 2026. Published by Oxford University Press on behalf of European Economic Association. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.

1. Introduction

Do people adjust how much they want to work when the central bank's monetary policy stance shifts? More specifically, does an interest rate hike induce individuals to work more or fewer hours? And does this effect differ across households with different levels of income (or earnings)?

The vast literature on the heterogeneous effects of monetary policy has focused on the inter-temporal channel that affects the consumption and savings plans of households (Auclert 2019 (Auclert 2019); Bilbiie 2008; Cloyne, Ferreira, and Surico 2020; Kaplan, Moll, and Violante 2018). However, changes in consumption plans induced by variation in rates also influence the intratemporal allocation between consumption and leisure; that is, a household's desired supply of labor depends on how each individual can substitute consumption with working time and/or compensate with different sources of income. In standard models, the lower wage rates induced by a contractionary monetary policy have two effects on households' labor supply: a substitution effect that reduces how much households would prefer to work and an income effect that increases it.

The majority of the theoretical macro literature assumes no or negligible income effects on the labor supply.¹ It is often thought that income effects are small because—being short-lived—monetary policy shocks do not have large effects on lifetime income, which is what matters for an optimizing worker–consumer.² Moreover, monetary policy is traditionally viewed as affecting labor demand through the extensive margin and having little effect on labor supply.³

The scope of this paper is to revisit this channel and study the transmission mechanism of monetary policy to the labor supply decisions at the household level.⁴ First, we offer novel empirical evidence on the effect of monetary policy on hours

Sullivan, Paolo Surico, Gianluca Violante, and participants at numerous conferences and seminars for comments and suggestions.

E-mail: cristiano.cantore@uniroma1.it (Cantore); filippo.ferroni2@unibo.it (Ferroni); h.mumtaz@qmul.ac.uk (Mumtaz); Angeliki.Theophilopoulou@brunel.ac.uk (Theophilopoulou)

1. For example, Galí, Smets, and Wouters (2012), Dyrda and Pedroni (2022), Wolf (2021), and Auclert and Rognlie (2020), among others.
2. However, most of the empirical evidence used to support this assumption focuses on other shocks and not on monetary policy shocks. See the literature review section for more details.
3. For example, quoting the Federal Reserve Chairman Jerome Powell in a speech on November 30, 2022, “Policies to support labor supply are not the domain of the Fed: Our tools work principally on demand.”
4. The consequences of monetary policy actions on the labor market dynamics are not only of interest in academic cycles. Policymakers have expressed considerable interest in labor market outcomes across the whole spectrum of the population and in particular in low- and moderate-income communities. For example, in a [Jackson Hole speech on August 27, 2020](#), J. Powell said in unveiling the new Fed strategy that “our revised statement emphasizes that maximum employment is a broad-based and inclusive goal. This change reflects our appreciation for the benefits of a strong labor market, particularly for many in low- and moderate-income communities.”

worked at a more granular, disaggregated level. To do this, we study the effects of unexpected shifts in the monetary policy stance on hours worked by households with different income levels using survey data for the US. We find that individuals at the bottom of the income distribution increase their hours worked following a monetary policy tightening, in contrast with conventional macroeconomic theory. At the same time, aggregate hours and wages across the whole distribution decline. This adjustment occurs through both the intensive and extensive margins of labor supply, but with important heterogeneity across the income distribution. In particular, low-income individuals tend to increase their hours worked and exhibit lower separation rates following a monetary contraction. Moreover, their response is more sensitive to interest rate variations compared with other percentiles of income in the population. As the labor supplied by low- and moderate-income households (*the tail of labor supply*) represents both a non-negligible share of the volatility and a relevant proportion of hours worked in the aggregate, this response is also quantitatively relevant from a macro perspective.

The countercyclicality of hours worked in the left tail of the income distribution observed in the data is an equilibrium outcome resulting from the interaction of households' labor supply forces and firms' labor demand factors and consistent with multiple explanations. For example, as the recession induced by a contractionary monetary policy increases the probability of becoming unemployed, households with limited income sources have incentives to work more. Similarly, individuals who are close to their borrowing limits may need to work more hours to meet their debt obligations when interest rates rise. Supply-side explanations suggest that when lacking buffer savings or non-labor income sources, households with low- and moderate-incomes have less room to maneuver during tough economic times, and by varying their labor supply, they can smooth consumption along the business cycle. Alternatively, on the demand side, firms may lay off temporary or part-time workers and adjust the labor's input by utilizing more of their existing labor force, inducing selection in the sample. While it is very difficult to isolate the dominant channel responsible for our empirical findings and a combination of these stories is more likely, several pieces of evidence suggest that selection is not the dominant force in this context; that is, the results carry over when using the panel dimensions of our survey data and/or when isolating the response of only full-time employed workers. Finally, the evidence of falling wages alongside rising hours worked at the bottom of the income distribution suggests that labor supply forces—rather than labor demand shifts—play a dominant role in driving this pattern.

It is therefore natural to ask whether the labor supply behavior observed among low-income individuals can be theoretically rationalized. Another way to assess whether the empirical patterns are primarily driven by labor supply rather than labor demand forces is to study them in a structural model. If a standard heterogeneous-agent framework with endogenous labor supply—abstracting from heterogeneity in labor demand—can reproduce the countercyclical responses of hours worked at the bottom of the income distribution, this would provide strong support for a supply-side interpretation of the data. The second contribution of the paper is thus to explore these

mechanisms in theory and assess their implications for the transmission (*the tale*) of monetary policy.

We start by providing a simple intuition of the mechanism at work. Borrowing constraints and limited consumption smoothing (Athreya, Owens, and Schwartzman 2017) are likely to drive stronger income effects on labor supply decisions for households with low income and limited assets. In particular, for households at or near the borrowing limit, intertemporal smoothing is limited, so monetary policy affects their labor supply primarily through contemporaneous cash-flow and wealth effects (e.g., wages/transfers and the real return or real burden on existing nominal positions), which raise marginal utility of consumption and strengthen income effects in the intratemporal condition. By analyzing how borrowing constraints and the curvature of the utility function with respect to consumption affect the optimal choices of consumption and leisure, we show that monetary policy shocks can generate an increase in labor supply among constrained households, driven by income effects rather than standard intertemporal substitution. This mechanism highlights the importance of heterogeneity in borrowing constraints and marginal propensities to work.

We then turn to a quantitative analysis using a standard one-asset heterogeneous-agent New Keynesian (HANK) model with endogenous labor supply and nominal price rigidities. We show that this “off-the-shelf” HANK model, calibrated to match plausible features of household heterogeneity, is able to reproduce the key labor supply patterns we uncover in the data: following a contractionary monetary policy shock, households in the lower part of the income distribution increase labor supply, while labor supply declines among higher-income households. With this model, we can also decompose the labor supply response into its underlying channels. This reveals that the countercyclical labor supply at the bottom of the income distribution is primarily driven by income effects—specifically, the decline in real wages and the increase in debt repayment burdens following a monetary tightening. We then systematically study how the strength of these heterogeneous labor supply responses varies across different model calibrations by altering the elasticity of intertemporal substitution (EIS) and the borrowing limit. To further quantify the macroeconomic implications of heterogeneous labor supply, we compare this baseline model to a similar HANK economy where labor supply is homogeneous across households, as in Auclert, Rognlie, and Straub (2024). To ensure comparability, we calibrate both models to match the labor supply response of the median agent type, and examine differences in the steady state and in the responses to monetary policy shocks.

We find that allowing for heterogeneous labor supply has quantitatively significant implications for monetary transmission. In particular, the steady-state aggregate marginal propensity to consume (MPC) is systematically lower in models with endogenous, heterogeneous labor supply than in comparable models where labor supply is homogeneous. This difference is especially pronounced at low values of the EIS, where income effects are stronger and constrained households rely more on labor effort to buffer shocks. This additional adjustment margin dampens the aggregate consumption response to monetary policy and, crucially, reduces the real

cost of disinflation for the monetary authority. We quantify this effect by computing the sacrifice ratio, defined as the cumulative percentage output loss per cumulative percentage point reduction in inflation over the first year following a contractionary monetary shock. Across different calibrations of the EIS, we find that the sacrifice ratio is systematically lower in the model with heterogeneous labor supply. For instance, under a low EIS (high income effect), the sacrifice ratio falls from 1.02 in the homogeneous labor supply model to 0.67 in the heterogeneous labor supply one model, a 35% reduction in the output cost of disinflation. This result arises because low-income households increase labor effort in response to the shock, partially offsetting the decline in consumption and mitigating the contraction in aggregate demand. From a policy perspective, this implies that failing to account for heterogeneity in labor supply may lead central banks to overestimate the output costs of achieving disinflation and misjudge the trade-offs involved in monetary tightening.

The paper is organized as follows: the next subsection discusses the existing literature. Section 2 describes the data and the empirical strategy and presents our empirical evidence. Section 3 presents a structural model that accounts for this evidence and investigates the implication for the transmission of monetary policy. Finally, Section 4 provides some concluding remarks.

Related Literature

This paper contributes to the literature on monetary policy and household heterogeneity. While most empirical work has focused on balance sheet composition and the heterogeneity in MPCs following monetary shocks (Auclert 2019; Cloyne, Ferreira, and Surico 2020), we instead study how such shocks affect labor supply decisions across households. By examining heterogeneous labor supply responses in HANK models, we also highlight their implications for aggregate MPCs.

Kehoe et al. (2020) and Amir-Ahmadi, Matthes, and Wang (2021) document heterogeneity in the responses of hours worked and unemployment across US demographic groups. The former finds that labor supply is less cyclical for older and college-educated workers, while the latter shows large variation in unemployment responses. We complement these studies by sorting households by income bins rather than demographic traits and focusing on the intensive margin of labor supply.

Graves, Huckfeldt, and Swanson (2023) study the effect of monetary policy on the labor market flows and find that a monetary policy tightening induces an increase in the fraction of labor force non-participants reporting that they want a job and an increase in the number of distinct job search methods by unemployed individuals. Both these margins of adjustment are consistent with an increase in the labor supply of non-employed individuals. These results are in line and complementary with our findings on the increase of hours worked of workers with low or moderate income (both currently employed or coming from non-employment) following a monetary policy tightening.

Del Canto et al. (2025) also study the distributional effects of the US monetary shocks using monthly VARs and data from the Current Population Survey (CPS), but their focus is normative rather than positive.

Several papers use administrative data to study the heterogeneous effects of monetary policy on labor market outcomes. For Scandinavian countries, Amberg et al. (2022), Andersen et al. (2021) and Holm, Paul, and Tischbirek (2021) focus on labor income and capture combined effects on both the extensive and intensive margins. Coglianesi, Olsson, and Patterson (2025) analyze administrative data from Sweden and show that unemployment responses to monetary shocks vary across the earnings distribution, focusing on labor market transitions. Hubert and Savignac (2024) find that in France, most of the variation in labor income for the bottom half of the distribution stems from the extensive margin, while Broer, Kramer, and Mitman (2022) document heterogeneous unemployment risk in Germany, with low-income workers facing more pro-cyclical separation rates. However, none of these studies can disentangle hours worked from wages, as we do here. Our contribution is to identify a distinct transmission channel: the heterogeneous response of hours worked to a monetary policy shock. Moreover, we use monthly outcomes, as in Broer, Kramer, and Mitman (2022), which allows us to exploit a longer time-series dimension to identify the transmission of monetary policy shocks.

Motivated by our empirical findings for the US, Das et al. (2025) study the effects of monetary policy on hours worked using administrative data from Australia. Leveraging high-frequency identification and individual-level income and hours data, they find that labor supply responses are stronger among low-income and low-liquidity individuals. Their results confirm that income effects play a key role in shaping labor supply reactions to interest rate shocks.

As discussed in the introduction, macroeconomic models often assume negligible income effects on labor supply. However, the empirical evidence supporting this view rarely focuses on business cycle shocks, eventually delivering mixed conclusions. Most estimates come from idiosyncratic income shocks, such as lottery winnings. For example, Cesarini et al. (2017) use Swedish administrative data and find modest income effects, while Golosov et al. (2023), using US data, argue that labor supply responses to lottery winnings are sizable and not negligible.

From a theoretical perspective, we contribute to the literature on micro-level heterogeneity in New Keynesian models. Most existing work focuses on the consumption channel of monetary policy while abstracting from labor supply heterogeneity (e.g., Auclert 2019; Auclert, Rognlie, and Straub 2024; Bayer, Born, and Luettticke 2024; Bilbiie 2024). A notable exception is Athreya, Owens, and Schwartzman (2017), who emphasize the role of labor supply decisions and marginal propensities to work in shaping the effects of fiscal transfers. To our knowledge, we are the first to study this channel in the context of monetary policy. Similarly, Guerrieri and Lorenzoni (2017) explore how different utility calibrations affect labor supply responses to credit shocks in a heterogeneous-agent model with incomplete markets.

Importantly, while our empirical and theoretical results highlight the relevance of heterogeneous labor supply, incorporating this feature into HANK models presents

challenges—especially when introducing labor market frictions like sticky wages, which often rely on homogeneous labor supply to unions. A recent contribution by Gerke et al. (2024), motivated by our work, addresses this by allowing unions to internalize household-specific labor supply—either through heterogeneous hours or constraints ensuring households do not work beyond their individual optimum. Their results, consistent with ours, show that accounting for labor supply heterogeneity dampens the effects of monetary policy on output, wages, and inflation even in the presence of wage rigidities.

2. Monetary Policy and Labor Market Outcomes along the Income Distribution

In this section, we describe the data sources and construction and the empirical strategy to identify monetary policy shocks, and we present our empirical evidence about the transmission of these shocks to household-level variables. Our main empirical evidence is constructed using the information on US labor earnings and hours worked at the individual level.

We find evidence that the individuals at the left tail of the income distribution typically increase the weekly amount of hours worked after a monetary policy tightening. The response of these individuals contributes to a non-negligible fraction of the response of aggregate hours worked. Moreover, we find that hours worked are more sensitive on the left tail of the income distribution.

2.1. Household-Level Data

Our source of individual-level data is the CPS, sponsored jointly by the US Census Bureau and the US Bureau of Labor Statistics (BLS).

The CPS survey is conducted at a monthly frequency on a sample of about 60,000 US households; it contains detailed information about the demographic characteristics of households, labor market attitudes, and labor earnings. We employ the uniform CPS outgoing rotation group extracts created by CEPR data for our benchmark analysis.⁵ In each month of the sample that runs from 1985 to 2019, we extract individual-level data on hours worked and hourly real wages, with individuals sorted on labor earnings.⁶ Our measure of hours corresponds to hours worked in the previous week in all jobs. We use the consistent series for hourly wages in 2019 dollars created by CEPR as our measure of wage. We construct a pseudo panel using earning percentile groups. We split the population into two groups, $P_{\leq J}$ and $P_{> J}$ where J ranges from 5 to 95 with

5. Households are interviewed in the CPS for 4 months and then again for 4 months after an 8-month break. In the fourth and eighth months of interviews, when households are about to rotate out of the interviews, they are asked additional questions about earnings. For more details on the outgoing rotation group, see [CPS notes](#).

6. Gross income data are not available at the monthly frequency in the CPS.

increments of 5. For example, when $J = 5$, $P_{\leq 5}$ refers to the group that consists of respondents who fall below that 5th percentile of labor earnings. Our panel includes $P_{\leq J}$ for $J = 5, \dots, 95$ capturing the cumulative distribution. We also include quintile bins, that is, five percentile groups of earnings. We refer to less than or equal to the 20th percentile with P_{20} , greater than 20th and less than equal to 40th percentile with P_{20-40} and so on. We calculate average wage and average weekly hours for each group using the survey weights. Repeating this across all months in the sample provides a time series of average earnings and average weekly hours by percentile group.⁷

As discussed in the next section, the dataset used for estimation combines aggregate and income-group-level data, considering one percentile split at a time, that is, $P_{\leq 5}$ and $P_{>5}$. We therefore estimate a separate empirical model for each percentile split.

The demographic characteristics vary substantially across bins. As we move from the left to the right of the earnings distribution, respondents tend to be older and better educated; they are more likely to work longer hours and be white and male. On the left tail of the wage distribution, industries such as wholesale and retail trade, health and education, leisure, and manufacturing are important in terms of employment.

While the CEPR extracts provide consistent data to construct a pseudo panel, we are unable to follow individuals over time. In the empirical analysis below, we also consider the impact of monetary policy shocks on transitions from employment. To construct this variable, we use the longitudinally matched version of CPS provided by the [Kansas Fed](#). This allows to track the employment status of individuals across a year and compute the rate at which individuals in the different wage percentile groups move from employment to being unemployed or exiting the labor force.⁸

2.2. Empirical Model

To estimate the impact of monetary policy shocks on the labor supply for different slices of the population, we use a factor-augmented VAR (FAVAR) model. The model is defined by the VAR:

$$Y_t = c + \sum_{j=1}^P \beta_j Y_{t-j} + u_t, \quad (1)$$

7. Unless otherwise specified, we apply a filter to the survey data. In particular, we drop respondents that lie in the top and bottom 1st percentile of the earnings distribution or are aged less than 18 or more than 66.

8. In principal, this data can be constructed by simply counting the transitions across wage groups. In practice, the number of individuals that report labor market status in two consecutive years and have earnings data available is low, leading to an implausibly low estimated transition rate, especially at the left tail of the wage distribution. In order to deal with this problem, we impute earnings for individuals with missing earnings data using mincer-type regressions. In particular, we regress earnings on level of education, a measure of experience (age minus year of schooling minus six), and individual characteristics, including race, sex, and industry of occupation. The fitted values from this regression are used to obtain imputed earnings for individuals that do not report this data.

with

$$Y_t = \begin{pmatrix} R_t \\ \hat{F}_t \end{pmatrix},$$

and where R_t denotes a policy interest rate and \hat{F}_t represents factors that summarize information in a panel of macroeconomic and financial series and the survey-based data on income and hours, described above. The factors are estimated using the non-stationary factor model of Barigozzi, Lippi, and Luciani (2021). A key advantage of this approach is that it allows us to use the data in levels. This is convenient as we are primarily interested in the impact of policy on the level of wages/labor earnings and hours in the percentile groups. Denote X_t as the $(M \times 1)$ data matrix that contains the panel of macroeconomic and financial series that summarize information about the economy, and also includes the average hours and average real earnings in the earnings percentile groups described above. The observation equation of the FAVAR is defined as

$$X_t = c + b\tau + \Lambda F_t + \xi_t \quad (2)$$

where c is an intercept, τ denotes a time-trend, F_t are the K non-stationary factors, Λ is a $M \times K$ matrix of factor loadings, and ξ_t are idiosyncratic components that are allowed to be $I(1)$ or $I(0)$. Note that the idiosyncratic components corresponding to the survey-based data can be interpreted as shocks that are specific to those groups and also capture possible measurement errors. In contrast, the shocks to equation (1) represent macroeconomic or common shocks. The response to these common shocks is relevant to our investigation. This ability to estimate the impact of macroeconomic shocks while accounting for idiosyncratic disturbances is a key advantage of the FAVAR over a VAR, where these two sources of fluctuations may be conflated (see De Giorgi and Gambetti 2017). Moreover, expanding the cross-sectional dimension of the VAR with factors is important also for identification purposes as it reduces the problem of information deficiency (see, e.g., Forni and Gambetti 2014) and shock deformation (see, e.g., Canova and Ferroni 2022).

The macro and financial data in X_t is obtained from the [FRED-MD](#) database. This monthly database contains 149 time series covering real activity, employment, inflation, money, credit, spreads, and asset prices. The sample starts in 1985m1, which is the first observation of hours worked constructed in the CPS, while the last observation is 2019m12.

To identify a monetary policy shock, we use an external instrument approach (see, e.g., Mertens and Ravn 2013 (Stock and Watson 2008); Stock and Watson 2008). The residuals u_t are related to structural shocks ε_t via:

$$u_t = A_0 \varepsilon_t \quad (3)$$

where $\text{cov}(u_t) = \Sigma = A_0 A_0'$. We denote the shock of interest as ε_{1t} and the remaining disturbances as ε_{-t} . Identification of ε_{1t} is based on the instrument m_t that satisfies the relevance and exogeneity conditions: $\text{cov}(m_t, \varepsilon_{1t}) = \alpha \neq 0$ and $\text{cov}(m_t, \varepsilon_{-t}) = 0$. As discussed in [Online Appendix C](#), these conditions can be combined with

the covariance restrictions to obtain an estimate of the relevant column of the contemporaneous impact matrix A_0 .

Our benchmark instrument used to identify the monetary policy shock is taken from Bauer and Swanson (2023). Bauer and Swanson (2023) show that instruments for monetary policy shocks based on high-frequency yield curve movements around FOMC meetings (see, e.g., Gertler and Karadi 2015) can be predicted by high-frequency changes in macroeconomic and financial variables. We use their orthogonalized version of the [instrument](#) that is exogenous to these fluctuations in our benchmark model. Following Gertler and Karadi (2015), R_t is assumed to be the 1-year government bond yield. The number of factors in the FAVAR model is set to 9 on the basis of the information criteria of Bai and Ng (2002) and the lag length is set at 12.⁹

The unobserved factors in (2) are estimated using the principal component estimator described in Barigozzi, Lippi, and Luciani (2021). The parameters of the VAR model in (1) are estimated using a Bayesian approach. The Markov chain Monte-Carlo algorithm is described in [Online Appendix B](#). We employ 21,000 iterations, retaining every 2nd draw after a burn-in period of 1,000.¹⁰

2.3. Response of Aggregate and Dis-ggregate Variables

Figure 1 shows the response of some key aggregate variables to a contractionary monetary policy shock in the US. The results are obtained using the monthly FAVAR, which includes data on the distribution of hours from the CPS. The size of the shock is normalized to generate an increase in the one government bond yield of one percent.

Industrial production contracts and the peak decline is about 1.4% after 1 year; the lag and magnitude effects of the shock are roughly in line with what is available in the literature on the empirical transmission of monetary policy shocks, see, for example, Bauer and Swanson (2023). The consumers price index falls on impact by 0.4% and remains persistently low thereafter with a peak effect of 0.5% after 10 months. The monetary contraction is associated with a deterioration of labor market indicators. The unemployment rate increases, peaking at 0.4% two and a half years after the shock, similar to Miranda-Agrippino and Ricco (2021). Aggregate hours worked decline on impact by 0.4% and continue falling during the following 10 months, reaching a trough of 0.8 percent. Stock market prices react negatively, with a peak response of around 8% on impact. Financial conditions, measured by the excess bond premium, deteriorate on impact and remain tight for a year. The exchange rate of the dollar vis-à-vis with the UK pound appreciates. In short, these results accord well with theory.

The top panel of Figure 2 shows our main result, which is the response of ‘actual’ hours worked for individuals at different slices of the earning distribution after a monetary policy tightening. For individuals above the 20th percentile of the earning distribution, the impulse response of hours worked is qualitatively similar to the

9. Our main results are not sensitive to this choice.

10. The prior distributions for the VAR parameters are standard and described in [Online Appendix B](#).

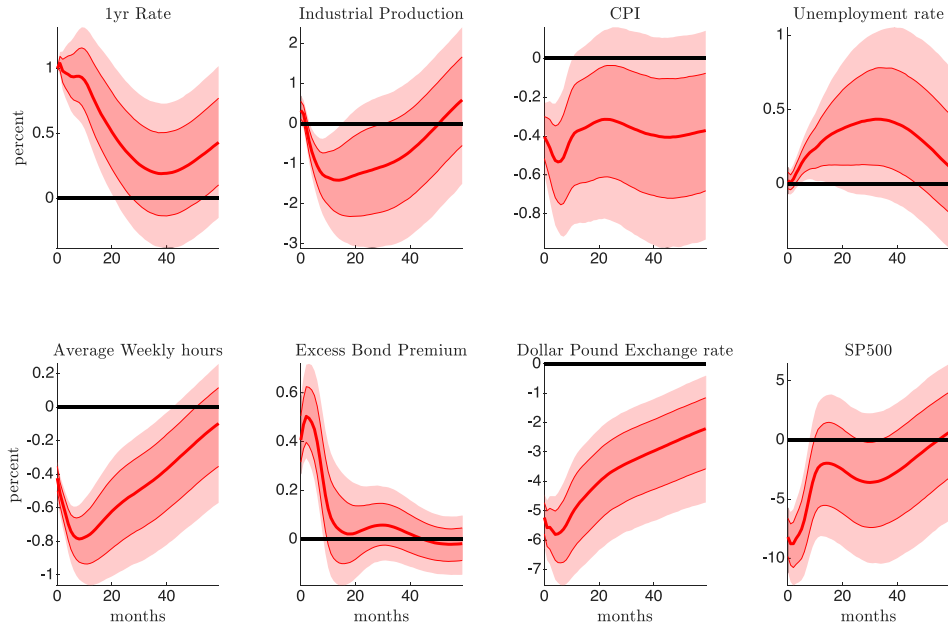


FIGURE 1. Impulse responses to a monetary policy shock. This figure depicts the impulse responses to a monetary policy tightening shock. From top left to bottom right: 1-year government bond yield; industrial production; consumer price index; unemployment rate; aggregate hours worked; excess bond premium; UK-US exchange rate; S&P500. Dark (light) red areas: 68 (90)% confidence sets.

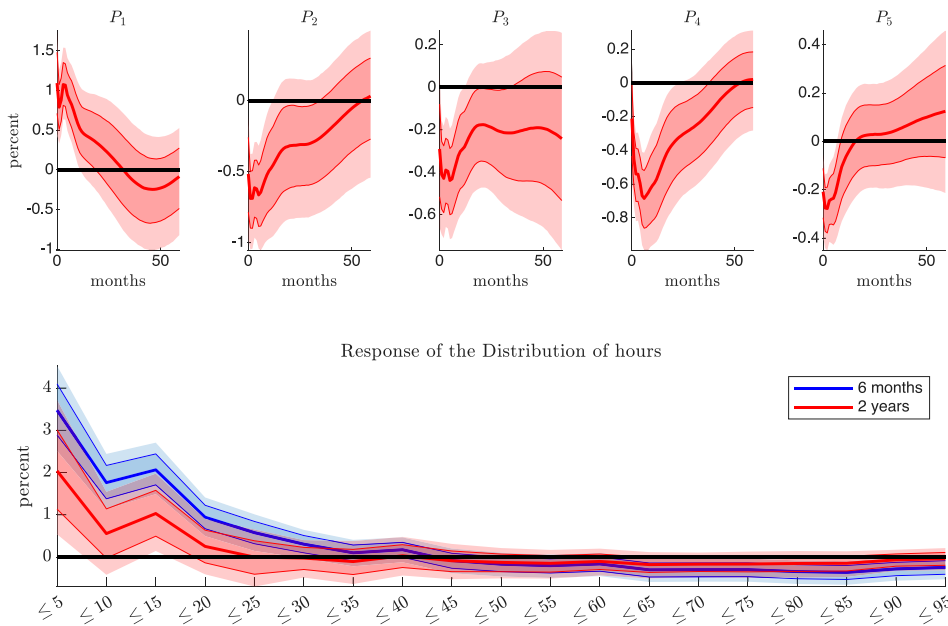


FIGURE 2. Distribution of responses to a monetary policy shock. This figure depicts the impulse responses to a monetary policy tightening shock. Dark (light) red areas: 68 (90)% confidence sets.

aggregate. Hours fall persistently in the range of 0.5%–0.8% at their trough in the middle 60% of the earning distribution, that is, earnings between the 20th and 80th percentile. Hours worked also decline for the high-income individuals, albeit the peak effect is smaller, that is, about 0.3 percent, and more short-lived. In contrast, hours display a persistent increase on the left tail of the wage distribution. The peak response of hours occurs at about a 6-month horizon and is estimated to increase by 1 percent. Interestingly, the responses are more volatile on the left side of the earning distribution. The bottom panel of Figure 2 reports the response of hours worked for households at the left tail of the wage distribution by varying the cut-off threshold. In particular, the horizontal axis reports different percentiles of the wage-earning distribution, and the vertical axis reports the responses of hours worked after the shock for an income group below a certain percentile. Red (blue) lines and areas report the point and confidence sets 2 years (6 months) after the shock. From this figure, we can conclude that in the left tail of the income distribution, a monetary policy tightening causes individuals to work more hours.

Importance of the Tail. To have a sense of the importance of the response of hours worked of low- and moderate-income individuals for aggregate quantities, we looked at different statistics. First, we computed the proportion of the variance of hours explained by the left tail, up to the 20th and up to the 30th percentile of the wage distribution. The former (20th percentile) explains 16% of aggregate hours worked and 27% of the growth rate of aggregate hours, respectively; the latter (30th percentile) explains 29% of aggregate hours and 44% of the growth rate of aggregate hours. Second, we constructed an aggregate measure of hours worked based on CPS data and an alternative aggregate measure of hours worked that excludes the bottom 20% of the wage distribution. Figure 3 reports the responses of hours worked for the first quintile of the earning distribution (first panel), the aggregate measure of hours worked constructed using the CPS data (second panel), and the alternative aggregate measure of hours worked, which excludes the first quintile of the earning distribution (third panel). The response of the counterfactual (third panel) measure is 50% larger than the actual aggregate one (second panel). This suggests that the response of low-income individuals plays a non-trivial role in shaping the aggregate labor market outcomes and in particular in dampening the response of aggregate hours to a monetary policy shock.

The extent to which this milder contraction in labor supply translates into less amplification of other variables (especially inflation) is less clear. For answering the latter, we need to construct a hypothetical counterfactual economy without the left tail of labor supply. The empirical model does not allow to run such counterfactuals. The structural model presented in Section 3 can shed some light on this point.

Composition Effects. A potential concern about the empirical evidence presented earlier is that the observed increase in hours worked among low-income individuals following a monetary tightening may reflect composition effects. For instance, if part-time or low-hour workers are more likely to exit employment during downturns,

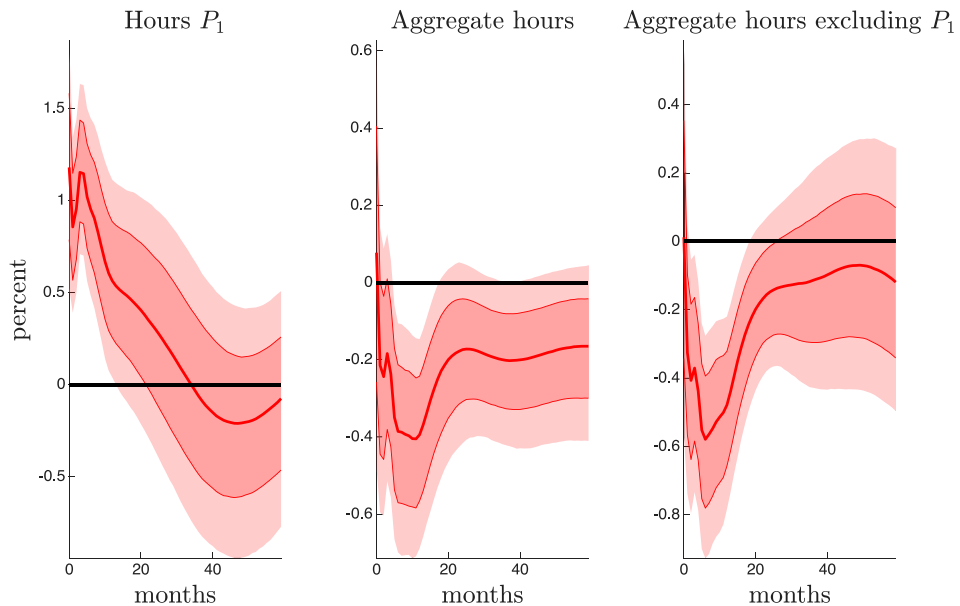


FIGURE 3. Impulse responses to a monetary policy shock. This figure depicts the impulse responses to a monetary policy tightening shock. The left panel reports the responses of hours worked for the first quintile of the earning distribution, the central panel the aggregate measure of hours worked using the CPS data, and the right panel the alternative aggregate measure of hours worked, which excludes the first quintile of the earning distribution. Dark (light) red areas: 68 (90)% confidence sets.

average hours could rise mechanically even if individual labor supply remains unchanged. To address this concern, we first restrict the sample of our analysis to full-time workers.¹¹ Figure 4 displays the response of hours worked for different income levels. Removing part-time workers does not invalidate our main findings and hours increase at the left tail of the earnings distribution after a monetary contraction.

Up to now our empirical analysis is constructed using a repeated cross-section, which even controlling for part- and full-time workers might still be prone to composition effects. To rule those out, we leverage the panel dimension of the CPS constructed by the Kansas City Fed (<https://cps.kansascityfed.org/>) and track changes in hours worked at the individual level. Specifically, we compute the change in hours between month t and $t + 12$ and use its average within income groups as the dependent variable in a FAVAR framework. This approach mitigates composition concerns inherent in cross-sectional averages.

The top panel of Figure 5 shows the growth in hours worked 6 months (red) and 2 years (blue) post-shock across the income distribution. Despite wider confidence

11. For this exercise, we use the [Kansas Fed](https://cps.kansascityfed.org/) extract of the CPS by setting the variable `lfdetail76` equal to either 1 or 2.

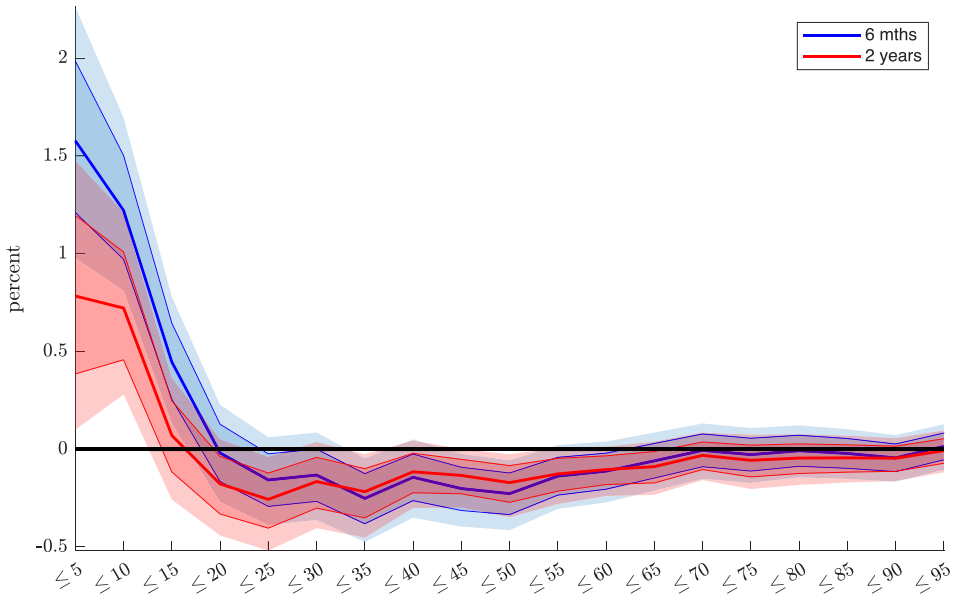


FIGURE 4. Responses of hours for full-time employees. This figure depicts the impulse responses to a monetary policy tightening shock. Dark (light) red areas: 68 (90)% confidence sets.

intervals, the results support our main finding: low-income individuals increase their hours worked in response to contractionary policy. The bottom panel of Figure 5 reports employment outflows¹² over the same horizon. While outflows rise for middle- and high-income groups, they decline for low-income individuals, indicating stronger job attachment in the lower tail of the distribution. These findings confirm that both intensive and extensive labor supply margins respond to monetary shocks, but in income-dependent ways. In the next sections, we interpret these patterns through the lens of theory.

Supply and Demand Factors. The countercyclical increase in labor supply among low-income workers reflects the interplay between household labor supply and firm-level labor demand. For middle- and high-income individuals, both real wages and hours worked decline following a monetary tightening, consistent with a leftward shift in labor demand. In contrast, among low-income workers, we observe rising hours despite falling real wages—indicating that income effects on labor supply dominate in this group. This pattern of negative comovement between wages and hours is suggestive of supply-side forces at play, as discussed in Katz and Murphy (1992), where such comovement typically reflects movements along a downward-sloping

12. See footnote 8 and [Online Appendix A.1](#) for details.




FIGURE 5. Distribution of responses of the growth rate of hours worked and of employment after monetary policy tightening. The top panel depicts the response of the growth rate of hours worked growth 6 months (red) and 2 years (blue) after the shock using the panel version of the CPS. The bottom panel depicts response of the outflows from employment 6 months (red) and 2 years (blue) after the shock using the panel version of the CPS. Shaded areas (solid lines): 68 (90)% confidence sets.

labor demand curve in response to shifts in labor supply. The responses of real wages at different income percentiles are reported in the [Online Appendix D.3](#) and [Figure D.8](#).

Sectors and Education. We further explore the heterogeneity behind this pattern by disaggregating responses by industry and education.¹³ Low-wage workers are largely employed in sectors such as wholesale and retail trade, leisure and hospitality, and education and health services ([Online Appendix Figure D.3](#)). Within these sectors, we observe clear positive responses of hours worked after a monetary contraction, especially at the bottom of the wage distribution ([Online Appendix Figure D.4](#)).

Educational attainment also plays a role. Low-income, non-college-educated individuals exhibit a larger increase in labor supply after monetary tightening than their college-educated peers suggesting that education moderates the strength of income effects ([Online Appendix Figure D.7](#)).

13. For more details, see [Online Appendix D.2](#).

3. Labor Supply and Heterogeneity

We now turn to a theoretical model to rationalize our empirical findings and assess their implications for the monetary policy transmission mechanism. We start from a general and stylized framework that clarifies the novel link between heterogeneous labor supply and monetary policy. Specifically, constrained agents facing tighter financial conditions tend to sacrifice leisure and increase labor supply to sustain their consumption in response to a decline in income.

Our empirical evidence, demonstrating that hours worked increase among households in the lower part of the income distribution after a monetary policy tightening, aligns with several potential theories. While the observed equilibrium outcome reflects both supply- and demand-side factors, the inverse movement of wages and hours worked—wages falling while hours increase—points predominantly toward a labor supply response driven by income effects. For this reason, in this section, we focus the theoretical analysis on labor supply heterogeneity while assuming a standard labor demand side, modeled through a representative firm with homogeneous labor demand.

3.1. Income Effects

We begin by analyzing the household's problem in partial equilibrium, taking labor income as given, to highlight how the strength of the income effect on individual labor supply is shaped by two key forces: the curvature of the utility function and the tightness of the borrowing constraint.¹⁴

A key feature in macroeconomic models determining household labor supply is the intratemporal optimality condition governing the trade-off between consumption and leisure. Let H_t denote hours worked at time t , C_t consumption, and $U(C_t, H_t)$ the household's utility function. With w_t as the real wage rate, this optimality condition is

$$-\frac{U_h(C_t, H_t)}{U_c(C_t, H_t)} = w_t, \quad (4)$$

where U_h and U_c are partial derivatives with respect to hours and consumption, respectively.

Households also face an intertemporal consumption decision summarized by the Euler equation (abstracting from uncertainty for now):

$$\frac{U_c(C_t, H_t)}{U_c(C_{t+1}, H_{t+1})} = \beta(1 + r_t)(1 + \omega_t), \quad (5)$$

where r_t is the real interest rate, affected by monetary policy, and ω_t represents a wedge arising from borrowing constraints or other financial frictions. Typically, constrained households face a positive wedge ($\omega_t > 0$), reflecting a higher marginal utility of current consumption. When the borrowing constraint binds, the Euler

14. We thank an anonymous referee for suggesting how to structure this section.

equation is replaced by a complementary slackness condition; in that case, the relevant link from monetary policy to constrained households operates through contemporaneous disposable-income and wealth (cash-flow) effects rather than intertemporal substitution.

Following a contractionary monetary policy shock that raises r_t , the household's intratemporal optimality condition determines how hours worked adjust in response to changes in income and wages. The strength of this labor supply adjustment crucially depends on the curvature of the utility function over consumption. We assume preferences of the form:

$$U(C_t, H_t) = \frac{C_t^{1-1/\sigma}}{1-1/\sigma} - \varphi \frac{H_t^{1+\nu}}{1+\nu},$$

where σ denotes the EIS and ν the inverse of the Frisch elasticity of labor supply.

A lower σ (corresponding to a more concave utility in consumption) implies that marginal utility reacts more strongly to income changes, amplifying the labor supply response when consumption declines. Households facing a negative income shock will then supply more labor to stabilize their utility. Conversely, a higher σ flattens the utility curve, making households less sensitive to fluctuations in income, and thus dampening the adjustment of hours worked.¹⁵

Borrowing constraints (via ω_t in the Euler equation) further magnify this effect by forcing some households to adjust labor supply more aggressively in response to changes in current disposable income and the real return on existing asset/debt positions. The borrowing limit directly affects the strength of this labor supply channel. A tighter borrowing constraint (or higher equilibrium debt levels) amplifies the sensitivity of labor supply to monetary policy shocks: when the real interest rate rises, higher debt repayments reduce disposable income, strengthening the income effect. Constrained households, unable to smooth consumption through borrowing, respond by increasing their labor supply to offset the higher financial burden.

In [Online Appendix F](#), we illustrate how these forces interact in a simple two-agent (borrower-saver) economy where borrowers face tighter borrowing limits and higher impatience relative to savers. We show analytically how borrowers' labor supply response after an interest rate increase depends directly on the EIS (decreasing in σ) and their borrowing limit (increasing in the net debt position).

Thus, both the EIS and borrowing constraints crucially shape the extent to which the labor supply of constrained households becomes countercyclical following monetary policy shocks, offering a clear theoretical interpretation in terms of labor supply of our empirical findings.

In the next section, we move to a general equilibrium analysis and consider a standard one-asset heterogeneous agents New Keynesian model with staggered price setting.

15. See Bilbiie (2008), who shows, in a two-agent NK model, how σ affects the sign of the response of the labor supply of hand-to-mouth agents.

3.2. Heterogeneous Labor Supply in HANK

The purpose of this section is to demonstrate that our empirical evidence aligns with the implications for heterogeneous labor supply responses in a standard HANK model. To do so, we use the one-asset HANK model with endogenous labor supply, as in Auclert et al. (2021).¹⁶ The goal is to examine how different calibrations of the model affect the behavior of labor supply across the income and wealth distribution, and to compare the aggregate implications for the monetary transmission mechanism with and without heterogeneous labor supply responses in the model. The main results are as follows: (i) An off-the-shelf HANK model with heterogeneous labor supply and homogeneous labor demand can replicate our empirical findings on the labor supply response across the income distribution. (ii) The main driver of the response at the lower end of the distribution is an income effect, stemming from falling wages and rising debt repayments. Moreover, the strength of this response is amplified by the curvature of the utility function in consumption and the tightness of the borrowing limit, further highlighting the importance of income effects. (iii) The presence of labor supply heterogeneity lowers the real cost of disinflation, making it easier for the central bank to achieve its inflation target.

One-Asset HANK. The model features a heterogeneous agents sector similar to McKay, Nakamura, and Steinsson (2016), coupled with a standard New Keynesian supply-side block.

Households. There is a unit mass of ex-ante identical households who differ ex-post by their labor productivity e (“skill”) and asset holdings a . For notational simplicity, we use the subscript i to denote household-level outcomes, instead of writing them explicitly as functions of state variables.

Skill e follows a time-invariant discrete Markov chain with n_e states, $\mathcal{E} = \{e_1, e_2, \dots, e_{n_e}\}$, and exogenous transition probabilities $P(e', e)$. This introduces cyclical income risk in the model. The stationary distribution of P is denoted by π_e . Average labor productivity $\int_0^1 e_{i,t} di$ is invariant and normalized to 1.

We assume that $P(e', e)$ discretizes a log AR(1) process:

$$\log e_{it} = \rho_e \log e_{it-1} + \sigma_{\epsilon,e} \epsilon_{it},$$

with normal innovations $\epsilon_{it} \sim \mathcal{N}(0, 1)$.¹⁷ We use the Rouwenhorst method for discretization.

Households can freely choose the number of hours h they work, subject to an additively separable utility cost of working. Consumption, savings, and labor choices c , a , and h are the solution to the household’s utility maximization problem,

16. The only difference with Auclert et al. (2021) is that we abstract from government spending, given the focus on monetary policy here. Results are not affected by this choice.

17. In Table 1, we report the stationary cross-sectional dispersion $\text{Std}(\log e) = \sigma_e$; given ρ_e , the corresponding innovation standard deviation is $\sigma_{\epsilon,e} = \sigma_e \sqrt{1 - \rho_e^2}$ (for the baseline calibration, $\sigma_{\epsilon,e} \approx 0.13$).

characterized by the following Bellman equation:

$$V_t(e_{it}, a_{it-1}) = \max_{c_{it}, h_{it}, a_{it}} \left\{ \frac{c_{it}^{1-\sigma^{-1}}}{1-\sigma^{-1}} - \varphi \frac{h_{it}^{1+\nu}}{1+\nu} + \beta \mathbb{E}_t V_{t+1}(e_{it+1}, a_{it}) \right\}$$

subj. to: $c_{it} + a_{it} = (1+r_t)a_{it-1} + w_t e_{it} h_{it} - \tau_t \bar{\tau}(e_{it}) + d_t \bar{d}(e_{it})$

$a_{it} \geq \underline{a}$.

Households receive an hourly wage w_t , pay taxes, and receive dividends from firm ownership according to incidence rules $\bar{\tau}(e)$ and $\bar{d}(e)$.¹⁸

Firms. A competitive final goods firm assembles its output using a CES production function $Y_t = \left(\int_0^1 y_{jt}^{\frac{1}{\mu}} dj \right)^\mu$, giving rise to a standard CES demand system for the continuum of intermediate goods. These intermediates, in turn, are supplied by monopolists who produce using only labor, such that $y_{jt} = Z_t h_{jt}$, and are subject to quadratic price adjustment costs a la Rotemberg (1982):

$$\psi_t(p_{jt}, p_{jt-1}) = \frac{\mu}{\mu-1} \frac{1}{2\kappa} \left[\log \left(\frac{p_{jt}}{p_{jt-1}} \right) \right]^2 Y_t.$$

Here, Z_t denotes aggregate productivity that may be time-varying. In a symmetric equilibrium, gross inflation $1 + \pi_t = P_t/P_{t-1}$ evolves according to the Phillips curve:

$$\log(1 + \pi_t) = \kappa \left(\frac{w_t}{Z_t} - \frac{1}{\mu} \right) + \frac{1}{1+r_{t+1}} \frac{Y_{t+1}}{Y_t} \log(1 + \pi_{t+1}),$$

and dividends equal output net of labor and price adjustment costs: $d_t = Y_t - w_t H_t - \psi_t$.

Policy. Monetary policy sets the nominal rate on bonds according to a standard Taylor rule:

$$i_t = r^* + \phi \pi_t + \epsilon_t,$$

where r_t^* is the economy's long-run "natural rate" and ϵ_t the monetary policy shock. The real interest rate in period t is determined by the previously set nominal rate and inflation so that

$$1 + r_t = \frac{1 + i_{t-1}}{1 + \pi_t}.$$

The fiscal authority, in turn, issues a constant amount of government bonds B each period and collects the lump-sum taxes τ_t already mentioned above. Since there are no other spending items (abstracting from government spending), it chooses the tax rate to cover its interest rate payments every period:

$$\tau_t = r_t B.$$

18. This implies that skill e determines not only a household's income per hour worked, but also the amount of lump-sum taxes she must pay and the dividends she receives, both of which are distributed proportionally to e .

Market clearing. In an equilibrium, the following market clearing conditions must hold:

- Asset market:

$$B = \int_0^1 a_{it} di,$$

- Labor market:

$$\frac{Y_t}{Z_t} = N_t = \int_0^1 e_{it} h_{it} di,$$

- Goods market:

$$Y_t = \int_0^1 c_{it} di + \psi_t.$$

These in turn imply that aggregate household savings equal government-provided liquidity, labor demand equals supply in efficiency units, and the final good is used for consumption and price adjustment costs.¹⁹

Calibration. The model is solved using the *sequence-space Jacobians* approach pioneered by Auclert et al. (2021).²⁰ For technical details, we refer the reader to their paper. The calibration is also mostly taken from Auclert et al. (2021) and summarized in Table 1. The main difference from their calibration is that here we also choose B to target a percentage of HTM agents in a steady state of 20%. This is to ensure a comparison across different models and/or calibrations, keeping fixed the steady-state proportion of constrained agents. The other difference is that, in the baseline calibration, we allow for some borrowing in equilibrium ($\bar{a} = -0.5$).²¹ This calibration implies an income-weighted aggregate MPC in steady state equal to 12.5%.

Rationalizing the Empirical Evidence. Mapping the discrete income states to their positions in the steady-state distribution reveals that both consumption and labor supply, on average, increase with income—consistent with empirical evidence. However, poorer households near the borrowing constraint tend to work more than richer ones due to stronger income effects and tighter liquidity.²²

Figure 6 shows the impulse responses to a monetary policy tightening, assuming that ϵ_t follows an AR(1) process with persistence $\rho_\epsilon = 0.61$ and standard deviation $\sigma_\epsilon = 0.0025$. The upper panels display the dynamics of the interest rate shock,

19. As is well known, price adjustment costs don't matter in linearized models.

20. See Bayer and Lueticke (2020) for an alternative solution method.

21. The implied β and φ are almost identical to theirs, while our calibration requires a lower amount of liquidity compared to their model, where they set $B = 5.6$ and obtain $\text{MPC} = 11\%$ and $\text{HTM} = 17\%$. All the results presented in this section do not change substantially if we use their original calibration.

22. See Online Appendix G.1 for details.

TABLE 1. One-asset HANK calibration.

Parameter		Value	Target
<i>Households</i>			
β	Discount factor	0.98	Asset market clearing
φ	Disutility of labor	0.78	Labor market clearing
σ	EIS	.5	
ν	Inverse Frisch	2	
\underline{a}	Borrowing constraint	-0.5	
ρ_e	Autocorrelation of earnings	0.966	
σ_e	Cross-sectional std of log earnings	0.5	
<i>Firms</i>			
μ	Steady-state markup	1.2	
κ	Slope of Phillips curve	0.1	
<i>Policy</i>			
B	Bond supply	3.84	HTM = 0.2
ϕ	Taylor rule coefficient on inflation	1.5	
<i>Monetary policy shock</i>			
ρ_ϵ	Persistence	0.61	
σ_ϵ	Standard deviation	0.0025	
<i>Discretization</i>			
n_e	Points in Markov chain for e	7	
n_a	Points on asset grid	500	
\bar{a}	Upper limit on asset grid	150	
<i>Untargeted steady state</i>			
MPC	% aggregate marginal propensity to consume	12.5%	
r	Quarterly real rate	0.5%	
H	Effective aggregate labor supply	1	

inflation, and aggregate output. The lower panels highlight heterogeneous labor supply and consumption responses across income groups. Poorer households (e.g., P0_2, P2_11, and P11_34) increase labor supply following the shock, consistent with a dominant income effect driven by tight borrowing constraints. In contrast, higher-income households reduce labor supply, reflecting standard substitution effects. Consumption falls across all groups but declines more sharply for lower-income households, as expected. These results demonstrate that an off-the-shelf HANK model is able to capture our empirical evidence on the heterogeneous labor supply behavior across the income distribution in response to a monetary policy shock. It is important to stress that, since the model features a representative firm, labor demand is homogeneous across agents. However, while the physical unit of labor is identical across households, effective labor supply differs due to heterogeneity in individual productivity.

Next, in order to assess how labor supply responses to monetary policy shocks vary across the income distribution, and to understand the role of key structural parameters, we perform a comparative exercise across different model calibrations.

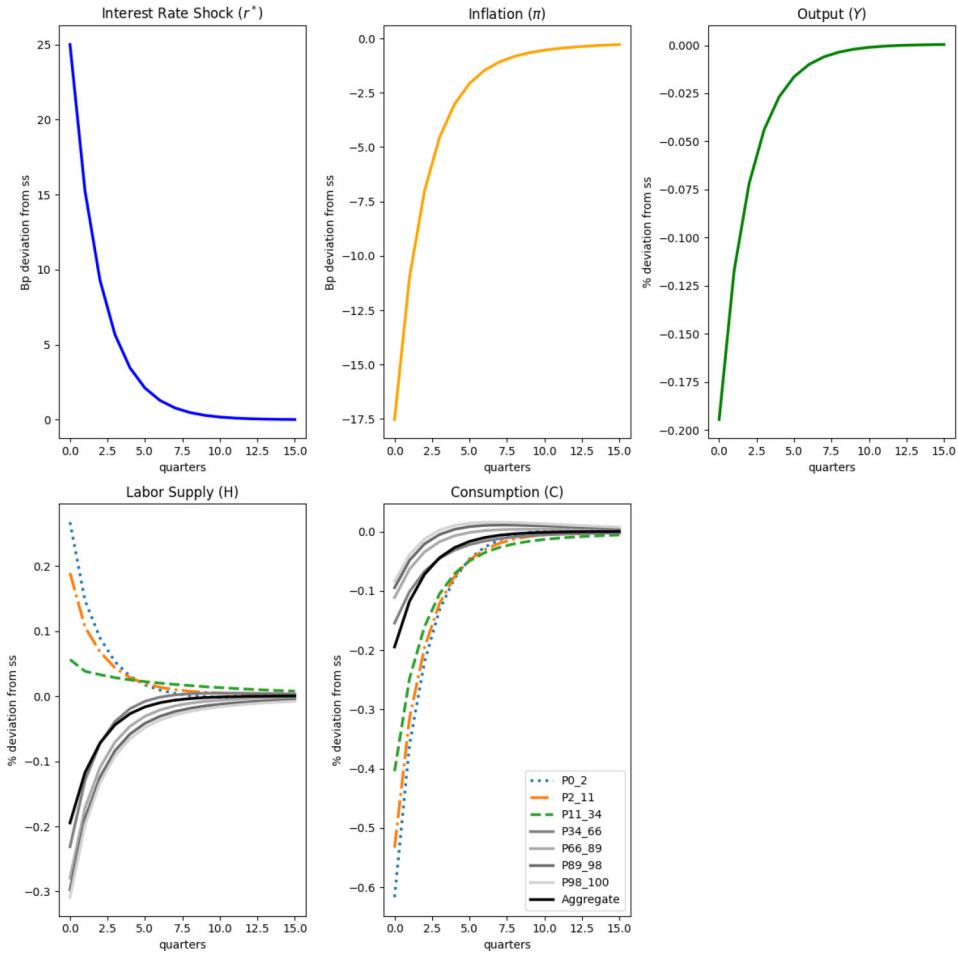


FIGURE 6. Impulse responses to a 25 basis points Monetary Policy Shock.

Specifically, we vary the two critical parameters identified in the previous section: the borrowing constraint (comparing $\underline{a} = 0$ —no borrowing allowed—with the baseline calibration $\underline{a} = -0.5$) and the EIS, which governs households’ willingness to smooth consumption over time and the relative strength of the income and substitution effects on their labor supply.²³

By analyzing the impact impulse responses (IRFs) of labor supply across income groups under these alternative specifications, we aim to isolate how tighter borrowing constraints and the EIS interact with income heterogeneity to shape aggregate and disaggregate labor market dynamics following monetary policy shocks. Figure 7 plots the impact responses of labor supply for households across seven income groups (from

23. [Online Appendix G.2](#) reports policy functions and IRFs for all calibrations.




FIGURE 7. Impact impulse responses of labor supply across the income distribution.

P0_2 to P98_100), under the six different model calibrations. Solid lines correspond to the baseline borrowing constraint calibration $a = -0.5$, allowing limited debt, while dashed lines represent a tighter borrowing constraint $a = 0$ (no debt allowed). Different colors denote different values of the EIS: blue for $\sigma = 0.25$ (high income effect), red for $\sigma = 0.5$ (baseline), and green for $\sigma = 1.0$ (low income effect).

Across all models, except for the case $a = 0$ & $\sigma = 1$, we observe that the bottom income groups (P0_2, P2_11) increase their labor supply after a monetary tightening, while upper income groups (P34_66 and above) reduce labor supply. Allowing for borrowing ($a = -0.5$) amplifies the need for immediate labor adjustment, leading to higher increases in labor supply at the bottom compared to the no-borrowing case. Similarly, lower EIS (lower σ) amplifies labor supply responses at the bottom, as households are more willing to reallocate labor to smooth consumption. In contrast, when EIS is high ($\sigma = 1$), the labor supply responses are significantly muted across the distribution, especially among the lower-income groups. Overall, the interaction between borrowing constraints and EIS shapes the heterogeneous labor market adjustments following monetary policy shocks. Importantly, this figure shows that while a low EIS is sufficient to induce a change in the sign of the labor supply IRFs for households in the left tail of the income distribution, allowing for a negative borrowing limit ($a < 0$) is necessary to generate a larger magnitude of the (absolute value of) labor supply responses at the bottom of the distribution, in line with our empirical results.

IRFs Decomposition. Contrary to the FAVAR, this model allows us to check what is driving the left tail of the labor supply response. To understand this, we decompose the total impulse responses into marginal contributions from individual channels.

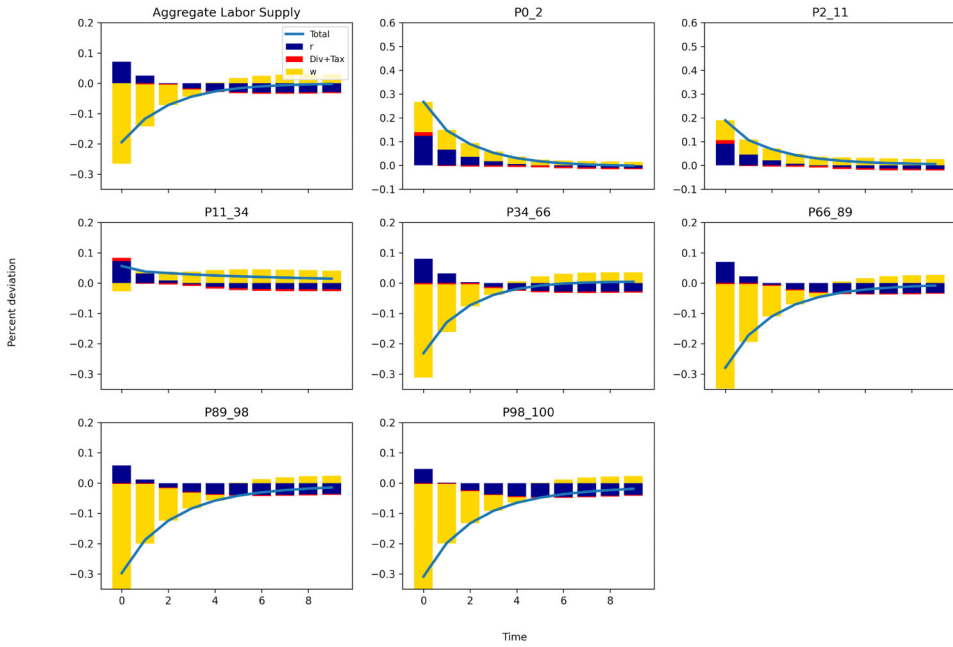


FIGURE 8. Decomposition of labor supply responses in HANK.

Specifically, we compute the Jacobians of household consumption with respect to the following inputs: the real interest rate (r), labor income (w), and transfers, defined as dividends (Div) minus taxes (Tax). Figure 8 shows the decomposition of total effective labor supply (top left panel) in terms of these main channels. The top left panel of this figure reports the decomposition of the aggregate hours worked, while the subsequent panels report the same decomposition for different percentiles of the income distribution. Labor income is the main driver of the decline in aggregate hours worked (yellow bars). A lower real hourly wage reduces the aggregate incentive to work. The real rate (blue bars), instead, pushes aggregate labor supply up during the first two quarters. In the Sequence Space Jacobian (SSJ) decomposition, the r input isolates the marginal effect of the real interest rate holding wages and fiscal flows (Div and Tax) fixed. This contribution reflects intertemporal-substitution effects for unconstrained households and contemporaneous cash-flow/wealth effects on existing asset and debt positions; for low-income/near-HTM households, the debt-service (cash-flow) channel is the key driver of the positive r contribution. Transfers, defined as the difference between dividends and taxes, marginally affect the dynamics (red bars).

Now the question is, how does this picture change when we zoom in at the individual level? And what is the force behind the left tail of the labor supply? The other panels in Figure 8 answer these questions by plotting the decomposition of labor supply by each income state. First, we notice that the real rate channel on impact is

stronger at the bottom of the income distribution, and it remains positive for more periods compared to agents in the middle to the top of the distribution. This is in line with the fact that the income effect from debt repayments is larger for poorer individuals. Moreover, the yellow bars flip sign at the bottom, pointing toward another income effect coming from labor income. For poor households, the income effect of a lower hourly wage dominates the substitution effect. Hence, in line with the simple analytical discussion presented at the beginning of this section, these results show that the combination of borrowing constraints and strong income effects on labor supply is responsible for matching our empirical evidence on the response of hours worked at the bottom of the income distribution.²⁴

Quantitative Implications. To further quantify the macroeconomic implications of heterogeneous labor supply behavior, we conduct an exercise comparing our baseline HANK model with endogenous labor supply to a HANK model (HANK-HomL) featuring homogeneous labor supply. In the model with homogeneous labor supply, following Auclert, Rognlie, and Straub (2024), total labor hours are evenly allocated across all households by a representative union, such that each individual's labor supply is rationed to be the same ($h_{it} = H_t$) at all times. While they introduce this assumption to facilitate the incorporation of sticky wages, here we maintain sticky prices and flexible wages to ensure comparability with the baseline models featuring heterogeneous labor supply. Moreover, as discussed in Section 1.1, Gerke et al. (2024) show similar results to ours in the presence of wage rigidities.

To ensure a meaningful comparison between models and isolate the implications of the left tail of the labor supply, we need a calibration that ensures that the labor supply response of the middle- and top-income agents is similar across the models while the one at the bottom is substantially different. To do so, we calibrate the EIS in the HANK-HomL in three cases (high, baseline, and low EIS) to match the impact response of the median agent across the two models.²⁵ Figure G.21 in the Online Appendix confirms that this calibration of σ across the two models does not affect the relative size of direct versus indirect effects of monetary policy. Moreover, even in the HANK-HomL, we keep calibrating B to match 20% fraction of HTM households in steady state. We then examine the sensitivity of the aggregate MPC to variations in the EIS, which, as we have shown, governs the strength of labor supply responses at the bottom of the income distribution.²⁶

24. Online Appendix G.3 shows the same figures for consumption.

25. The values of low, baseline, and high σ for the HANK model are 0.25, 0.5, and 1, respectively. The corresponding values in HANK-HomL that match the impact response of hours worked for the median agent in HANK are 0.445, 0.625, and 1.1. Figure G.20 in the Online Appendix shows the IRFs match for the baseline σ case. As a robustness check, Online Appendix G.4 instead recalibrates HANK-HomL to match the impact response (period 0) of aggregate labor L and reports the implied sacrifice ratios in Table G.2. The baseline sacrifice ratios under this median-match calibration are presented later in the paper (Table 3).

26. For this exercise, we use the baseline calibration of the borrowing limit $\underline{a} = -0.5$. Using $\underline{a} = 0$, we get virtually the same results as the borrowing limit has little impact on the steady state.

TABLE 2. One-asset HANK steady-state comparison.

Parameter	Low σ		Baseline σ		High σ	
	HANK-HomL	HANK	HANK-HomL	HANK	HANK-HomL	HANK
β	0.97	0.97	0.98	0.98	0.99	0.99
φ	0.83	0.71	0.83	0.78	0.83	0.85
B	1.98	3.80	2.52	3.84	3.38	4.61
MPC	15.3%	11.7%	14.8%	12.5%	14.2%	13.2%

Table 2 presents the results from this exercise. Across different values of the EIS parameter (σ), we find notable differences between the HANK-HomL and the baseline HANK model. In particular, the aggregate MPC is systematically lower in the model featuring heterogeneous labor supply compared to the HANK-HomL counterpart. Interestingly, this difference widens as the value of the EIS declines. This is because the MPC increases in σ in the HANK model while it decreases in HANK-HomL. The underlying mechanism is that, in the HANK model, households can use labor supply as a buffer to smooth consumption when faced with adverse shocks, thereby reducing their MPC out of transitory income (Pijoan-Mas 2006). In contrast, in the HANK-HomL model, labor supply is fixed and cannot be used as a margin of adjustment, forcing households to respond to shocks primarily through changes in consumption. This fundamental difference in adjustment channels explains the diverging MPC responses as σ varies. This is also evident by looking at the disutility of labor parameter (φ) that clears the labor market in HANK: a lower σ (higher income effect) induces more work effort, which is captured by a lower implied φ in equilibrium.

Accounting for heterogeneity in labor supply substantially dampens the aggregate consumption response to monetary policy shocks, as low-income households increase labor effort to buffer income losses—thereby lowering the aggregate MPC. This additional margin of adjustment weakens the transmission of monetary policy through aggregate demand. [Online Appendix G.4](#) provides full model comparisons, distributional impulse responses, and contribution analyses across the income distribution.

Policy Implications. Finally, we discuss the policy implications of the left tail of labor supply. More precisely, we ask: How does heterogeneity in labor supply responses affect the conduct of monetary policy? For instance, does it make central bank targets easier or harder to achieve? In recent years, the Federal Reserve had to raise interest rates significantly to bring inflation under control. This raises the question: Does the real activity cost of disinflation depend on the presence—and strength—of the heterogeneous labor supply channel we have highlighted?

In particular, we aim to quantify the output costs associated with a monetary-policy-induced reduction in inflation. To do so, we compute the sacrifice ratios implied by each model calibration. The sacrifice ratio measures the cumulative loss in output, expressed as a percentage deviation from steady state, per cumulative percentage point reduction in inflation following a monetary policy tightening. Specifically, we compute

TABLE 3. Sacrifice ratios across model calibrations.

	HANK-HomL	HANK
Low σ	1.02	0.67
Baseline σ	1.23	1.07
High σ	1.58	1.50

the cumulative output loss and cumulative inflation decline over the first four quarters after the shock (i.e., over 1 year), and define the sacrifice ratio as the ratio of these two quantities. Formally, the sacrifice ratio is given by

$$\text{Sacrifice Ratio} = \frac{\sum_{t=1}^4 \Delta Y_t}{\sum_{t=1}^4 \Delta \pi_t}$$

where ΔY_t denotes the percentage deviation of output from steady state in quarter t , and $\Delta \pi_t$ denotes the deviation of inflation from steady state in quarter t .

Comparing sacrifice ratios across different models allows us to assess how labor supply heterogeneity and borrowing constraints influence the real cost of disinflation.

Table 3 reports the sacrifice ratios computed for both models across different calibrations of the EIS, as in the previous exercises. Two important findings emerge. First, the sacrifice ratio increases with σ , indicating that a higher EIS makes output more sensitive to monetary policy shocks. When households can more easily shift consumption across periods (i.e., when intertemporal substitution is higher), contractionary monetary policy induces larger output declines, thus requiring a greater cumulative sacrifice to reduce inflation. Second, heterogeneous labor supply behavior implies a substantially lower cost of disinflation for the monetary authority due to the countercyclical labor supply response at the bottom of the income distribution. As shown in previous sections, when σ increases, this countercyclicality diminishes, and the gap between the two models narrows accordingly.

4. Conclusion

In this paper, we study the interaction between monetary policy and labor supply decisions at the household level. Our first contribution is to establish new empirical facts. Using individual-level survey data from the United States and a factor-augmented VAR methodology, we find that the response of hours worked to monetary policy shocks is heterogeneous across the income distribution. While aggregate hours decline after a contractionary monetary policy shock, employed individuals at the bottom of the income distribution increase their labor supply. Moreover, the response of low- and moderate-income households is not only opposite in sign, but also larger in magnitude compared to that of higher-income groups.

Given that the labor supplied by low- and moderate-income households accounts for a non-negligible share of aggregate labor supply and its volatility, these heterogeneous responses are quantitatively relevant for the aggregate effects

of monetary policy. While multiple mechanisms may contribute to the observed heterogeneity, several pieces of evidence—such as the robustness of results in panel data, among full-time employed workers, and across different sectors—point away from selection effects and support a dominant role for labor supply forces. In particular, the combination of falling wages and rising hours at the bottom of the distribution is consistent with strong income effects among constrained households.

The second contribution of the paper is to explore the theoretical implications of our empirical findings for the transmission of monetary policy. Another way to assess whether the empirical patterns are primarily driven by labor supply rather than labor demand forces is to study them in a structural model. We begin by illustrating the mechanism in a simple framework, showing how borrowing constraints and limited asset holdings amplify income effects and induce constrained households to increase labor supply in response to tighter monetary conditions. We then embed this mechanism into a one-asset HANK model with nominal rigidities and endogenous labor supply.

We show that, even if this has been overlooked thus far, an off-the-shelf HANK model featuring incomplete markets and borrowing constraints replicates the labor supply patterns observed in the data: Following an interest rate hike, poorer households increase their hours worked, while wealthier households reduce them. Importantly, the model allows us to decompose the labor supply responses into underlying channels, and we find that the countercyclical behavior at the bottom of the income distribution is primarily driven by income effects. We systematically study how these heterogeneous labor supply responses vary across different calibrations, changing the EIS and the borrowing limit, and compare them to a HANK model with homogeneous labor supply.

Our quantitative analysis shows that heterogeneity in labor supply behavior alters the aggregate MPC implied by the model. In particular, stronger countercyclical labor supply responses at the bottom of the distribution imply a lower MPC relative to the HANK model, where everybody supplies the same amount of hours, highlighting how labor supply acts as an additional margin through which households can smooth consumption. This additional adjustment margin also has substantial effects on the dynamics of the model. Importantly, it reduces the sacrifice ratio—the cumulative output loss per point of inflation reduction—faced by the monetary authority. Across model calibrations, we find that incorporating heterogeneous labor supply leads to a systematically lower sacrifice ratio, with reductions of up to 35%, implying that disinflationary policies may entail smaller output costs than previously estimated if this margin is accounted for.

Therefore, ignoring labor supply heterogeneity risks overstating households' vulnerability to monetary policy shocks and mischaracterizing both the distributional and aggregate effects of policy interventions.

References

- Amberg, Niklas**, Thomas Jansson, Mathias Klein, and Anna Rogantini-Picco (2022). “Five Facts about the Distributional Income Effects of Monetary Policy Shocks.” *American Economic Review: Insights* 4, 289–304.
- Amir-Ahmadi, Pooyan**, Christian Matthes, and Mu-Chun Wang (2021). “What Does Monetary Policy Do to Different People?” Working paper.
- Andersen, Asger Lau**, Niels Johannesen, Mia Jørgensen, and José-Luis Peydró (2021). “Monetary Policy and Inequality.” Economics Working Paper No. 1761, Dept. of Economics and Business, Universitat Pompeu Fabra.
- Athreya, Kartik**, Andrew Owens, and Felipe Schwartzman (2017). “Does Redistribution Increase Output? The Centrality of Labor Supply.” *Quantitative Economics*, 8, 761–808.
- Auclert, Adrien** (2019). “Monetary Policy and the Redistribution Channel.” *American Economic Review*, 109(6), 2333–2367.
- Auclert, Adrien**, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub (2021). “Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models.” *Econometrica*, 89, 2375–2408.
- Auclert, Adrien** and Matthew Rognlie (2020). “Inequality and Aggregate Demand.” Working paper.
- Auclert, Adrien**, Matthew Rognlie, and Ludwig Straub (2024). “The Intertemporal Keynesian Cross.” *Journal of Political Economy*, 132, 4068–4121.
- Bai, Jushan** and Serena Ng (2002). “Determining the Number of Factors in Approximate Factor Models.” *Econometrica*, 70, 191–221.
- Barigozzi, Matteo**, Marco Lippi, and Matteo Luciani (2021). “Large-dimensional Dynamic Factor Models: Estimation of Impulse-Response Functions with I(1) Cointegrated Factors.” *Journal of Econometrics*, 221, 455–482.
- Bauer, Michael D.** and Eric T. Swanson (2023). “A Reassessment of Monetary Policy Surprises and High-Frequency Identification.” *NBER Macroeconomics Annual*, 37, 87–155. <https://ideas.repec.org/a/ucp/macann/doi10.1086-723574.html>
- Bayer, Christian**, Benjamin Born, and Ralph Luetticke (2024). “Shocks, Frictions, and Inequality in US Business Cycles.” *American Economic Review*, 114(5), 1211–47. <https://www.aeaweb.org/articles?id=10.1257/aer.20201875>
- Bayer, Christian** and Ralph Luetticke (2020). “Solving Heterogeneous Agent Models in Discrete Time with Many Idiosyncratic States by Perturbation Methods.” *Quantitative Economics*, 11, 1253–1288.
- Bilbiie, Florin** (2008). “Limited Asset Markets Participation, Monetary Policy and (Inverted) Aggregate Demand Logic.” *Journal of Economic Theory*, 140, 162–196.
- Bilbiie, Florin O.** (2024). “Monetary Policy and Heterogeneity: An Analytical Framework.” *The Review of Economic Studies*, 92, 2398–2436. <https://doi.org/10.1093/restud/rdae066>
- Broer, T.**, J. Kramer, and K. Mitman (2022). “The Curious Incidence of Monetary Policy Shocks Across the Income Distribution.” Working Paper Series No. 416, Sveriges Riksbank, Central Bank of Sweden.
- Canova, Fabio** and Filippo Ferroni (2022). “Mind the Gap! Stylized Dynamic Facts and Structural Models.” *American Economic Journal: Macroeconomics*, 14(4), 104–135.
- Cesarini, David**, Erik Lindqvist, Matthew J. Notowidigdo, and Robert Östling (2017). “The Effect of Wealth on Individual and Household Labor Supply: Evidence from Swedish Lotteries.” *American Economic Review*, 107(12), 3917–46.
- Clodomiro, Ferreira Cloyne** and Paolo Surico (2020). “Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism.” *Review of Economic Studies*, 87, 102–129.
- Coglianesi, John**, Maria Olsson, and Christina Patterson (2025). “Monetary Policy and the Labor Market: A Quasi-Experiment in Sweden.” *American Economic Review*, 115(10), 3451–3486.
- Das, M.**, J. Hambur, K. Hellwig, and J. Spray (2025). “Labor Supply Effects of Monetary Policy: Evidence from Australian Mortgage Holders.” Working paper.
- De Giorgi, Giacomo** and Luca Gambetti (2017). “Business Cycle Fluctuations and the Distribution of Consumption.” *Review of Economic Dynamics*, 23, 19–41.

- Del Canto, Felipe, John Grigsby, Eric Qian, and Conor Walsh (2025). “Are Inflationary Shocks Regressive? A Feasible Set Approach.” *The Quarterly Journal of Economics*, 140(4), 2685–2747.
- Dyrda, Sebastian and Marcelo Pedroni (2022). “Optimal Fiscal Policy in a Model with Uninsurable Idiosyncratic Shocks.” *Review of Economic Studies*, 90, 744–780.
- Forni, Mario and Luca Gambetti (2014). “Sufficient Information in Structural VARs.” *Journal of Monetary Economics*, 66, 124–136.
- Gali, Jordi, Frank Smets, and Rafael Wouters (2012). “Unemployment in an Estimated New Keynesian Model.” *NBER Macroeconomics Annual*, 26, 329–360.
- Gerke, Rafael, Sebastian Giesen, Matija Lozej, and Joost Rottger (2024). “On Household Labour Supply in Sticky-Wage HANK Models.” Discussion Papers No. 01/2024, Deutsche Bundesbank, <https://ideas.repec.org/p/zbw/bubdps/282992.html>.
- Gertler, Mark and Peter Karadi (2015). “Monetary Policy Surprises, Credit Costs, and Economic Activity.” *American Economic Journal: Macroeconomics*, 7, 44–76.
- Golosov, Mikhail, Michael Graber, Magne Mogstad, and David Novgorodsky (2023). “How Americans Respond to Idiosyncratic and Exogenous Changes in Household Wealth and Unearned Income.” *The Quarterly Journal of Economics*, 139, 1321–1395. <https://doi.org/10.1093/qje/qjad053>
- Graves, S., C. Huckfeldt, and E. T. Swanson (2023). “The Labor Demand and Labor Supply Channels of Monetary Policy.” Working paper.
- Guerrieri, Veronica and Guido Lorenzoni (2017). “Credit Crises, Precautionary Savings, and the Liquidity Trap.” *The Quarterly Journal of Economics*, 132, 1427–1467.
- Holm, Martin Blomhoff, Pascal Paul, and Andreas Tischbirek (2021). “The Transmission of Monetary Policy under the Microscope.” *Journal of Political Economy*, 129, 2861–2904.
- Hubert, Paul and Frédérique Savignac (2024). “The Heterogeneous Effects of Monetary Policy on Labor Income: Disentangling the Extensive and Intensive Margins.” *Journal of Political Economy Macroeconomics*, 2, 687–719. <https://doi.org/10.1086/732469>
- Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018). “Monetary Policy According to HANK.” *American Economic Review*, 108(3), 697–743.
- Katz, Lawrence F. and Kevin M. Murphy (1992). “Changes in Relative Wages, 1963–1987: Supply and Demand Factors.” *The Quarterly Journal of Economics*, 107, 35–78.
- Kehoe, Patrick J., Pierlauro Lopez, Elena Pastorino, and Sergio Salgado (2020). “On the Dynamic Effects of Monetary Policy with Heterogeneous Agents.” Working paper.
- McKay, Alisdair, Emi Nakamura, and Jon Steinsson (2016). “The Power of Forward Guidance Revisited.” *American Economic Review*, 106(10), 3133–58.
- Mertens, Karel and Morten O. Ravn (2013). “The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States.” *American Economic Review*, 103(4), 1212–47.
- Miranda-Agrippino, Silvia and Giovanni Ricco (2021). “The Transmission of Monetary Policy Shocks.” *American Economic Journal: Macroeconomics*, 13(3), 74–107.
- Pi Joan-Mas, Josep (2006). “Precautionary Savings or Working Longer Hours?” *Review of Economic Dynamics*, 9, 326–352.
- Rotemberg, Julio J. (1982). “Sticky Prices in the United States.” *Journal of Political Economy*, 90, 1187–1211.
- Stock, James H. and Mark W. Watson (2008). “What’s New in Econometrics - Time Series.” Lecture 7, National Bureau of Economic Research, Inc.
- Wolf, Christian (2021). “The Missing Intercept: A Demand Equivalence Approach.” NBER Working Paper No. WP29558.

Supplementary Data

Supplementary data are available at [JEEA](https://academic.oup.com/jeea/advance-article/doi/10.1093/jeea/jvag030/6687749) online.

Journal of the European Economic Association 2026 00(0):1–30

<https://doi.org/10.1093/jeea/jvag030>

© The Author(s) 2026. Published by Oxford University Press on behalf of European Economic Association. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<https://creativecommons.org/licenses/by/4.0/>), which permits unrestricted reuse, distribution, and reproduction in any medium, provided the original work is properly cited.