ISSN 2282-6483

Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

A tail of labor supply and a tale of monetary policy

Cristiano Cantore
Haroon Mumtaz
Filippo Ferroni
Angeliki Theophilopoulou

Quaderni - Working Paper DSE N°1210

A tail of labor supply and a tale of monetary policy*

Cristiano Cantore
Sapienza University of Rome

Filippo Ferroni
University of Bologna

Haroon Mumtaz
Queen Mary, University of London

Angeliki Theophilopoulou
Brunel University London

September 2, 2025

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 U.S. 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.

Keywords: Monetary policy, Household Survey, FAVARs, HANK.

JEL classification: E52, E32, C10

*We would like to thank Valerio Pieroni for the lenghty discussions and suggestions about the HANK analysis we also thank Michele Andreolli, Guido Ascari, Saleem Bahaj, Christian Bayer, Gadi Barlevy, Florin O. Bilbiie, Davide Debortoli, Luca Fornaro, Luca Gambetti, Nezih Guner, Chris Huckfeldt, Mathias Klein, Guido Lorenzoni, Leonardo Melosi, Carlo Pizzinelli, Ricardo Reis, Kjetil Storesletten, Dan Sullivan, Paolo Surico, Gianluca Violante, three anonymous referees and participants at numerous conferences and seminars for comments and suggestions.

Non Technical Summary

The paper asks a simple question with a surprising answer: when the central bank raises interest rates, do people change how much they work, and does that change differ by income? Looking at U.S. survey data after unexpected rate hikes, the authors find that total hours and earnings in the economy fall, as standard stories predict. But among people at the bottom of the income distribution, hours worked actually rise. These workers are also less likely to separate from their jobs and their behavior is more sensitive to interest rate moves than that of middle- and higher-income households. Because low- and moderate-income workers account for a meaningful share of total hours, this extra effort partially offsets the overall decline in labor input.

Why would those with the least resources work more when the economy cools? The paper points to a straightforward intuition. Monetary tightening lowers real wages and raises the cost of servicing debt. Households with little savings and tight budgets have a harder time smoothing their spending by dipping into assets or borrowing. Instead, they smooth by supplying more labor. In other words, the "income effect" of feeling poorer dominates the usual "substitution effect" that would make people work less when wages soften. The authors also argue that this pattern reflects people's choices rather than firms' selective hiring or scheduling. The result holds in panel data, remains when focusing on full-time workers, and is accompanied by falling wages alongside rising hours at the bottom–signs that supply, not demand, is doing most of the work.

To test the interpretation, the paper builds a standard macro model that allows households to differ and lets each choose how much to work. Calibrated to match realistic variation in income, assets, and borrowing limits, the model reproduces the data's key feature: after a contractionary monetary shock, lower-income households increase their labor effort while higher-income households reduce it. In the model, the mechanism comes from strong income effects among constrained families, driven by lower real wages and higher debt payments.

These differences matter for the broader story of how monetary policy moves the economy. When poorer households can flex their hours, they rely less on cutting spending to adjust, so the hit to overall demand is smaller than in models that assume everyone behaves the same. The authors show that this cushions the output decline that typically accompanies disinflation. Measured by the sacrifice ratio–the cumulative output loss per percentage point reduction in inflation over the first year–the cost is about one-third lower in the version of the model where labor supply can vary across households (for example, roughly 0.67 rather than 1.02 under a calibration with strong income effects).

The upshot is that the "tail" of the income distribution changes the "tale" of monetary policy. If central banks ignore how low-income households adjust their work, they risk overstating the growth cost of bringing inflation down and misreading the trade-offs involved in raising rates.

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 (Bilbiie (2008), Auclert (2019), 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; i.e., 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 worked at a more granular, disaggregated level. To do this, we study the effects of unexpected shifts in the monetary policy stance on the amount of hours worked by households with different income levels using survey data for the U.S. 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

[1]E.g. Galí, Smets and Wouters (2012); Dyrda and Pedroni (2022); Wolf (2021); Auclert and Rognlie (2020); amongst 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 on 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. E.g. 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."

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; i.e. 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, constrained or hand-to-mouth (HTM) households face a tighter intertemporal budget constraint and a wedge in their Euler equation, making their marginal utility of consumption highly sensitive to income fluctuations. 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. This mechanism directly links the intertemporal and intratemporal choices of households, highlighting 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 elasticity of intertemporal substitution, 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.

1.1 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 (Cloyne et al. (2020), Auclert (2019)), 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, Lopez, Pastorino and Salgado (2020) and Amir-Ahmadi, Matthes and Wang (2021) document heterogeneity in the responses of hours worked and unemployment across U.S. 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 adjustments 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, Grigsby, Qian and Walsh (2025) also study the distributional effects of the US monetary shocks using monthly VARs and data from the 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, Jansson, Klein and Rogantini-Picco (2022), Andersen, Johannesen, Jørgensen and Peydró (2021) and Holm, Paul and Tischbirek (2021)) focus on labor income and capture combined effects on both the extensive and intensive margins. Coglianese, 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, unlike most of these studies (except Broer et al. (2022)), we use data at monthly frequency, 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 U.S., Das, Hambur, Hellwig and Spray (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. Most estimates come from idiosyncratic income shocks, such as lottery winnings. For example, Cesarini, Lindqvist, Notowidigdo and Östling (2017) use Swedish administrative data and find modest income effects, while Golosov, Graber, Mogstad and Novgorodsky (2023), using U.S. 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 et al. (2024), Bayer, Born and Luetticke (2024), Bilbiie (2024)). A notable exception is Athreya et al. (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 challengesespecially when introducing labor market frictions like sticky wages, which often rely on homogeneous labor supply to unions. A recent contribution by Gerke, Giesen, Lozej and Rottger (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 Current Population Survey (CPS), sponsored jointly by the U.S. Census Bureau and the U.S. Bureau of Labor Statistics (BLS).

The CPS survey is conducted at a monthly frequency on a sample of about 60,000 U.S. 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 earnings. 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 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 hourly earnings. Our panel includes $P_{\leq J}$ for $J = 5, \ldots 95$ capturing the cumulative distribution.

[5] Households are interviewed in the CPS for four months and then again for four months after an eight-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.

We also include quintile bins, i.e. five percentile groups of earnings. We refer to less than or equal to the 20^{th} percentile with P_{20}, greater than 20^{th} and less than equal to 40^{th} percentile with P_{20-40} and so on. We calculate average earnings 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]

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 labour 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:

$$Y_t = c + \sum_{j=1}^{P} \beta_j Y_{t-j} + u_t \tag{1}$$

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

[7] Unless otherwise specified, we apply a filter to the survey data. In particular, we drop respondents that lie in the top and bottom first 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 labour 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.

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:

$$X_t = c + b\tau + \Lambda F_t + \xi_t \tag{2}$$

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 times 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. Stock and Watson (2008) and Mertens and Ravn (2013)). The residuals u_t are related to structural shocks ε_t via:

$$u_t = A_0 \varepsilon_t \tag{3}$$

where $cov(u_t) = \Sigma = A_0 A_0'$. 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 exogeniety conditions: $cov(m_t, \varepsilon_{1t}) = \alpha \neq 0$ and $cov(m_t, \varepsilon_{-t}) = 0$. As discussed in 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 for e.g. Gertler and Karadi (2015)) can be predicted by high frequency changes in macroeconomic and financial. We use their orthogonalised 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 one-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 et al. (2021). The parameters of the VAR model in (1) are estimated using a Bayesian approach. The Markov chain Monte-Carlo algorithm is described in appendix B. We employ 21,000 iterations, retaining every 2^{nd} draw after a burn-in period of 1000.[10]

2.3 Response of aggregate and dis-aggregate 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 percent after one year; the lag and magnitude effects of the shock are roughly in line with what available in the literature on the empirical transmission of monetary policy shocks, see e.g. Bauer and Swanson (2023). The consumers price index falls on impact by 0.4 percent and remains persistently low thereafter with a peak effect of 0.5 percent after ten months. The monetary contraction is associated with a deterioration of labor market indicators. The unemployment rate increases peaking at 0.4 percent two and half years after the shock, similar to Miranda-Agrippino and Ricco (2021). Aggregate hours worked decline on impact by 0.4 percent and continue falling during the following ten month reaching a trough of 0.8 percent. Stock market prices react negatively, with a peak response of around 8 percent 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-a-is with the U.K. pound appreciates. In short, these results accord well with theory.

The top panel of figure 2 shows our main result, that 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 aggregate. Hours fall persistently in the range of 0.5 to 0.8 percent at their trough in the middle 60 percent of the earning distribution, i.e. earnings between the 20 and 80 percentile. Hours worked also decline for the high income individuals albeit the peak effect is smaller, i.e. 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 six 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

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

[10] The prior distributions for the VAR parameters are standard and described in appendix B.

Figure 1: Impulse responses to a monetary policy shock.

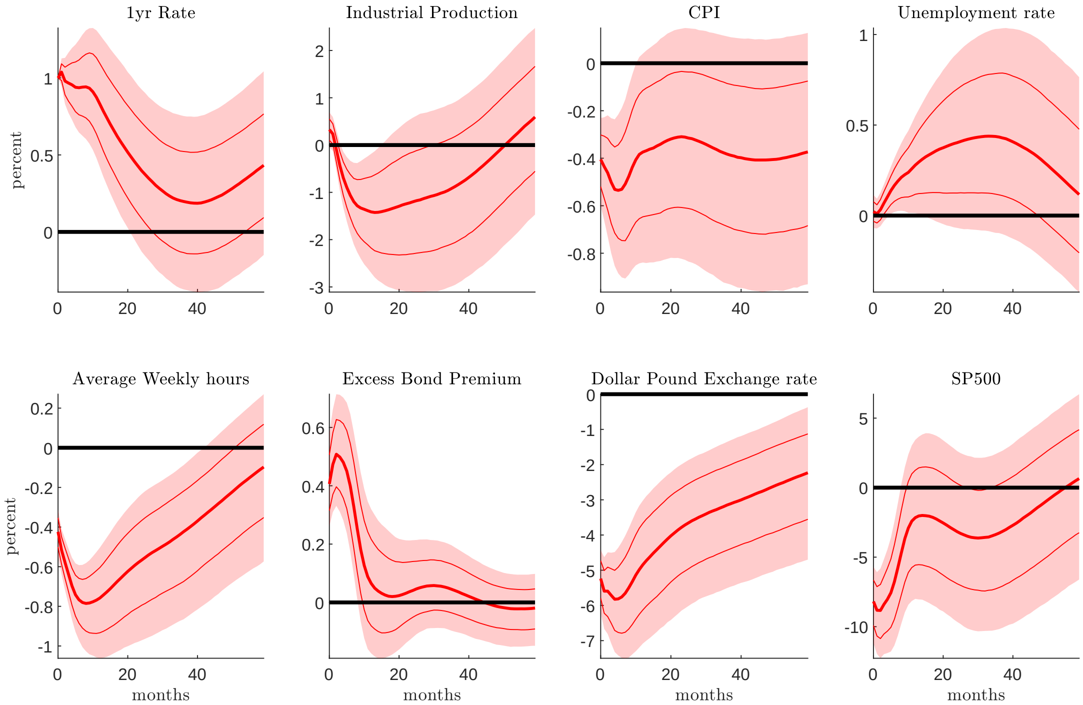

Notes: This figure depicts the impulse responses to a monetary policy tightening shock. From top left to bottom right: one year government bond yield; industrial production; consumer price index; unemployment rate; aggregate hours worked; excess bond premium; UK-US exchange rate; S&P500. Dark (light)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

particular, the horizontal axis reports different percentiles of the wage earning distribution and the vertical axis 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 two years (six 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 20th and up to the 30th percentile of the wage distribution. The former (20th percentile) explains 16 percent of aggregate hours worked and 27 percent of the growth rate of aggregate hours, respectively; the latter (30th percentile) explains 29 percent of aggregate hours and 44 percent 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 which excludes the bottom 20 percent of the wage distribution. Figure 3 reports the responses of hours worked for the first quintile of the earning distribution

Figure 2: Distribution of responses to a monetary policy shock.

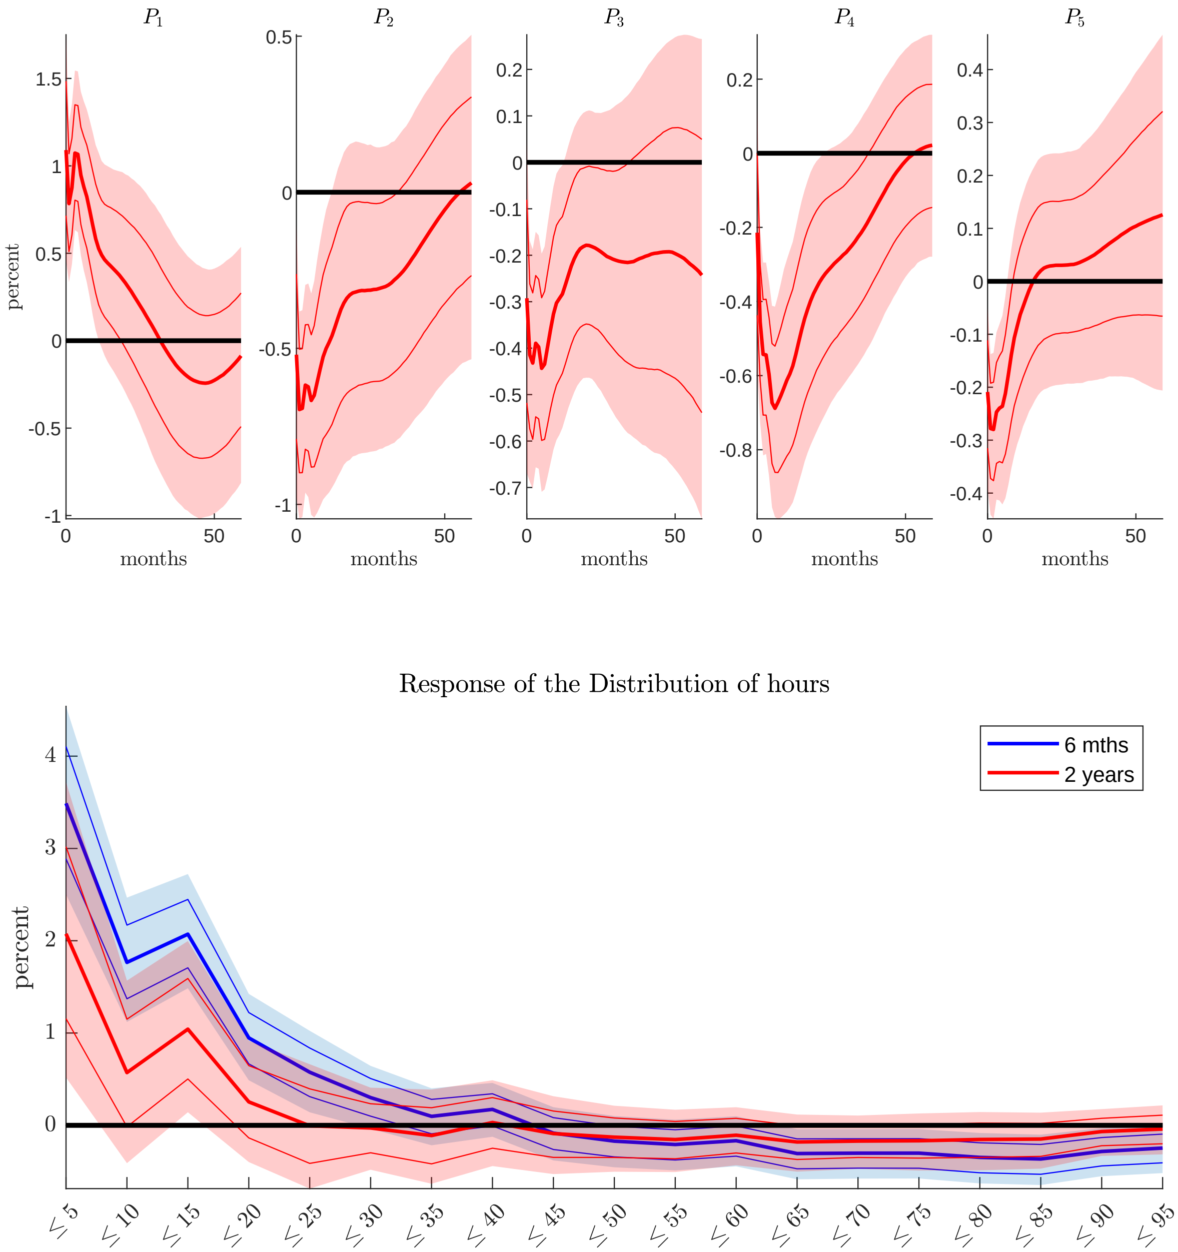

Notes: This figure depicts the impulse responses to a monetary policy tightening shock. Dark (light) red areas 68 (90)% confidence sets.

(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 first quintile of the earning distribution (third panel). The response of the counterfactual (third panel) measure is 50 percent 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.

Figure 3: Impulse responses to a monetary policy shock.

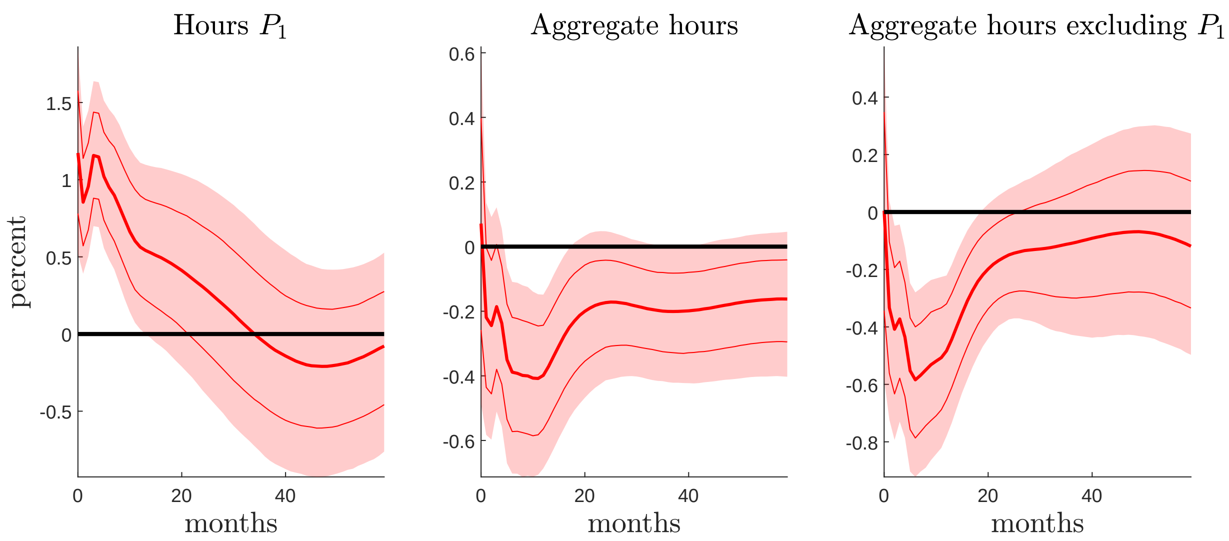

Notes: This figure depicts the impulse responses to a monetary policy tightening shock. The left panel reports the responses of hours worked for the first quintile of the earning distribution, the cental panel the aggregate measure of hours worked using the CPS data, and right panel the alternative aggregate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

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 an 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, 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

[11] For this exercise we use the Kansas Fed extract of the CPS by setting the variable lfdetail76 equal to either 1 or 2.

Figure 4: Responses of hours for full-time employees

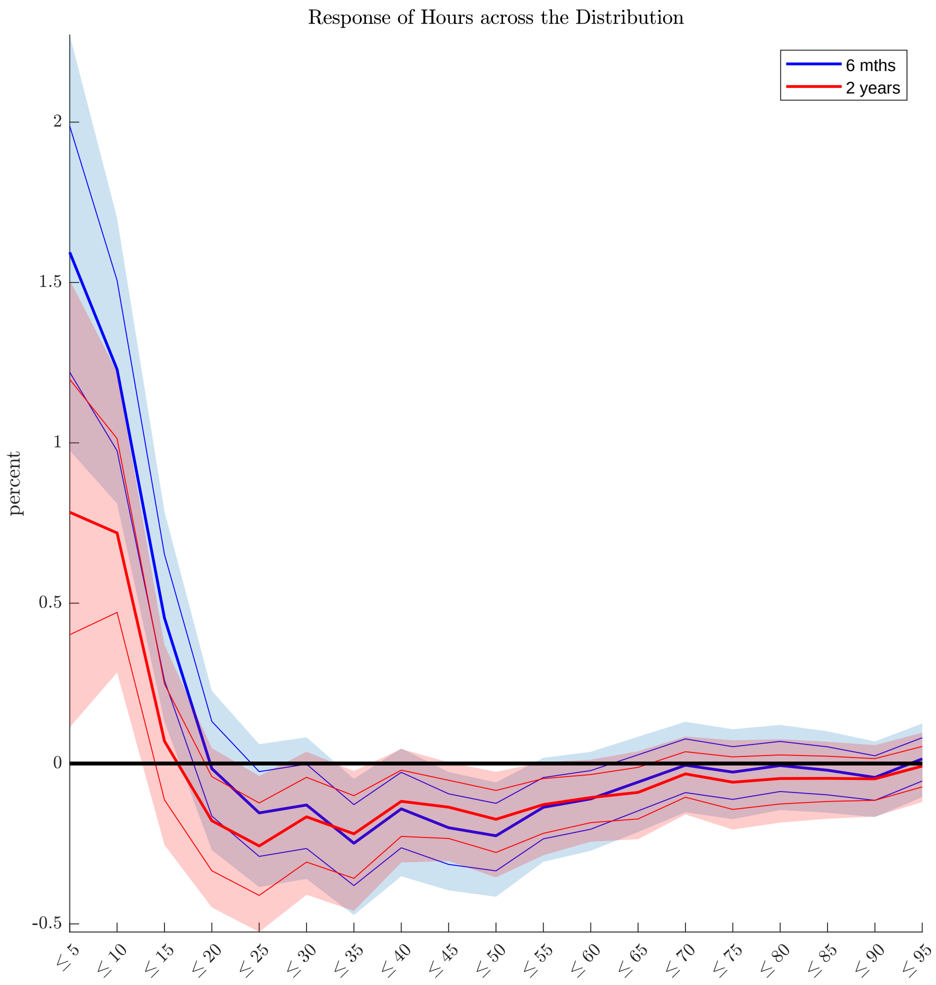

Notes: This figure depicts the impulse responses to a monetary policy tightening shock. Dark (light) red areas 68 (90)% confidence sets.

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 six months (red) and two years (blue) post-shock, across the income distribution. Despite wider confidence 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 outflows[12] 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,

[12]See footnote 8 and Appendix A.1 for details.

Figure 5: Distribution of responses of the growth rate of hours worked and of employment after monetary policy tightening.

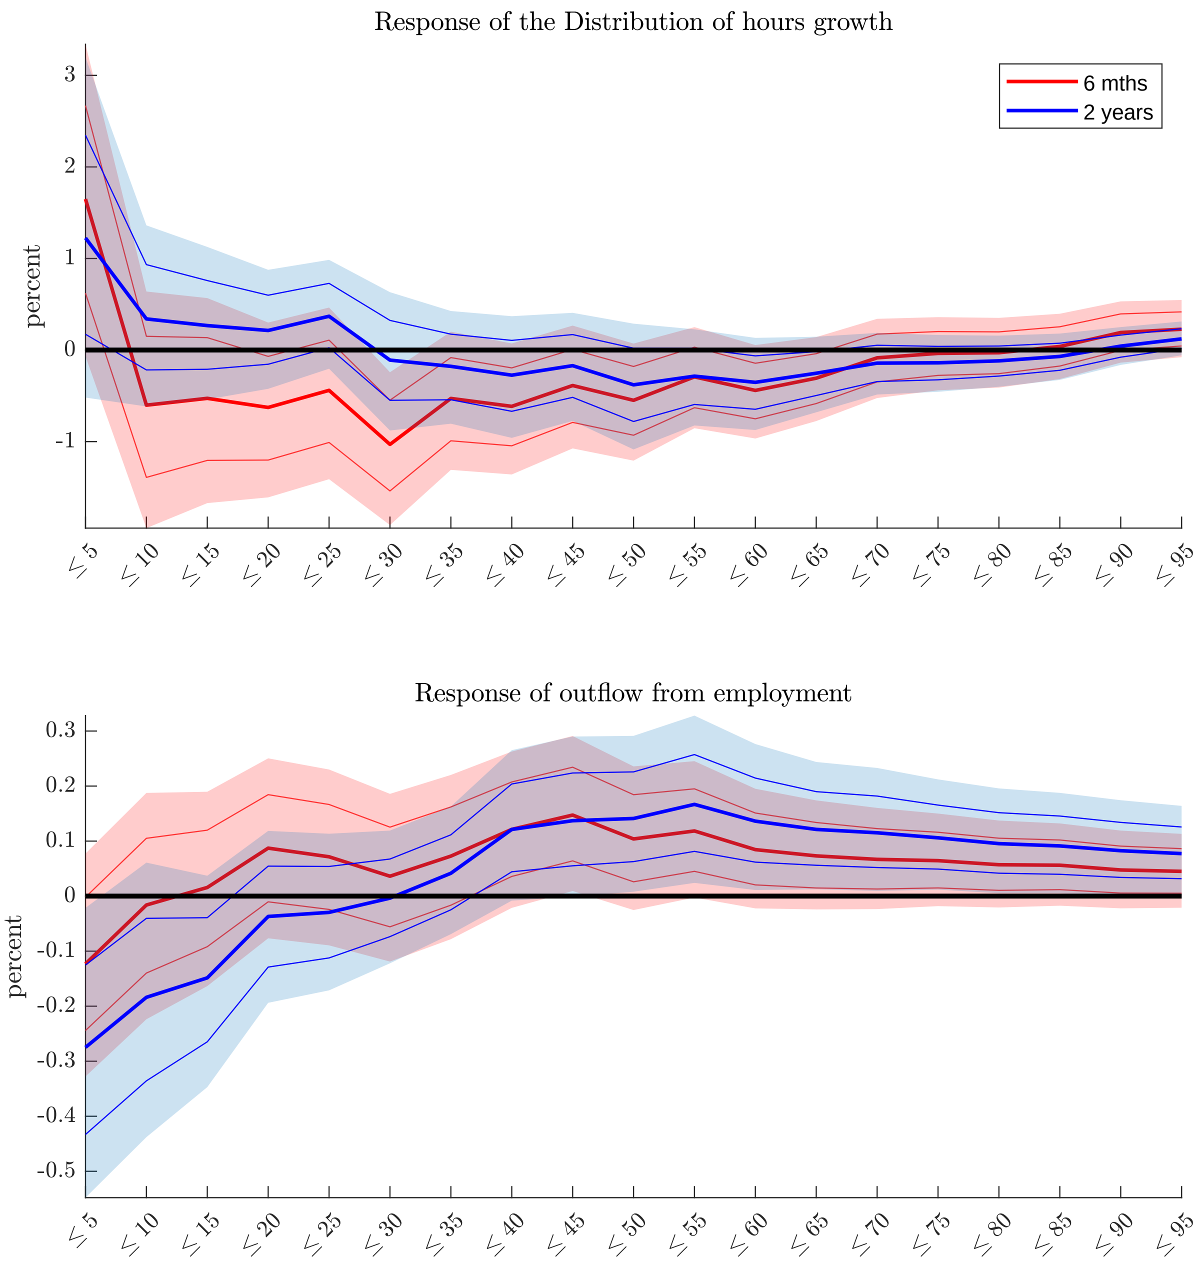

Notes: The top panel depicts response of the growth rate of hours worked growth six months (red) and two (blue) years after the shock using the panel version of the CPS. The bottom panel depicts response of the outflows from employment six months (red) and two (blue) years after the shock using the panel version of the CPS. Shaded areas (solid lines) 68 (90)% confidence sets.

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 labor demand curve in response to shifts in labor supply. The responses of real wages at different income percentiles are reported in the Appendix D.3 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 (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 (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 (Appendix Figure D.7).

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

[13] For more details see appendix D.2.

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:

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

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):

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

where r_t is the real interest rate, affected directly 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.

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 elasticity of intertemporal substitution 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 income and interest rates. 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

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

[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.

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 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.

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, Bardóczy, Rognlie and Straub (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.

3.2.1 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

[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.

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, $\mathcal{E} = \{e_1, e_2, \ldots, e_{n_e}\}$, 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:

$$\log e_{it} = \rho_e \log e_{it-1} + \sigma_e \epsilon_{it},$$

with normal innovations $\epsilon_{it} \sim \mathcal{N}(0, 1)$, and 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, characterized by the following Bellman equation:

$$\begin{aligned} V_t\left(e_{it}, a_{it-1}\right) = \max_{c_{it}, h_{it}, a_{it}} & \left\{ \frac{c_{it}^{1-\sigma^{-1}}}{1-\sigma^{-1}} - \varphi \frac{h_{it}^{1+\nu}}{1+\nu} + \beta \mathbb{E}_t V_{t+1}\left(e_{it+1}, a_{it}\right) \right\} \\ \text{s.t.} \quad & c_{it} + a_{it} = (1 + r_t) a_{it-1} + w_t e_{it} h_{it} - \tau_t \bar{\tau}(e_{it}) + d_t \bar{d}(e_{it}) \\ & a_{it} \geq \underline{a}. \end{aligned}$$

Households receive an hourly wage w_t, pay taxes, and receive dividends from firm ownership according to incidence rules $\bar{\tau}(e)$ and $\bar{d}(e)$.[17]

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):

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

Here, Z_t denotes aggregate productivity that may be time-varying. In a symmetric equilibrium, gross inflation $1 + \pi_t = \frac{P_t}{P_{t-1}}$ 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:

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

[17] 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.

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:

$$\tau_t = r_t B.$$

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

- Asset market:
$$B = \int_0^1 a_{it}\, di,$$
- Labor market:
$$Y_t = H_t = \int_0^1 e_{it} h_{it}\, di,$$
- 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.[18]

Calibration The model is solved using the *sequence-space Jacobians* approach pioneered by Auclert et al. (2021).[19] For technical details, we refer the reader to their paper. The calibration is also mostly taken from Auclert et al. (2021) and summarised in Table 1. The main difference from their calibration is that here we also choose B to target a percentage of HTM agents in 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 ($\underline{a} = -0.5$).[20] This calibration implies an income-weighted aggregate MPC in steady state equal to 12.5%.

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

[19] See Bayer and Luetticke (2020) for an alternative solution method.

[20] 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 MPC $= 11\%$ and HTM $= 17\%$. All the results presented in this section do not change substantially if we use their original calibration.

Parameter		Value	Target
Households			
β	Discount factor	0.98	$r = 0.005$
φ	Disutility of labor	0.78	$H = 1$
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	
Firms			
μ	Steady-state markup	1.2	
κ	Slope of Phillips curve	0.1	
Policy			
B	Bond supply	3.84	HTM = 0.2
ϕ	Taylor rule coefficient on inflation	1.5	
Monetary Policy Shock			
ρ_ϵ	Persistence	0.61	
σ_ϵ	Standard Deviation	0.0025	
Discretization			
n_e	Points in Markov chain for e	7	
n_a	Points on asset grid	500	
$\bar{a}$	Upper limit on asset grid	150	
Untargeted Steady State			
MPC	% aggregate Marginal Propensity to Consume	12.5%	

Table 1: One-asset HANK calibration

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.[21]

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, 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

[21] See appendix G.1 for details.

Figure 6: Impulse responses to a 25 basis points Monetary Policy Shock

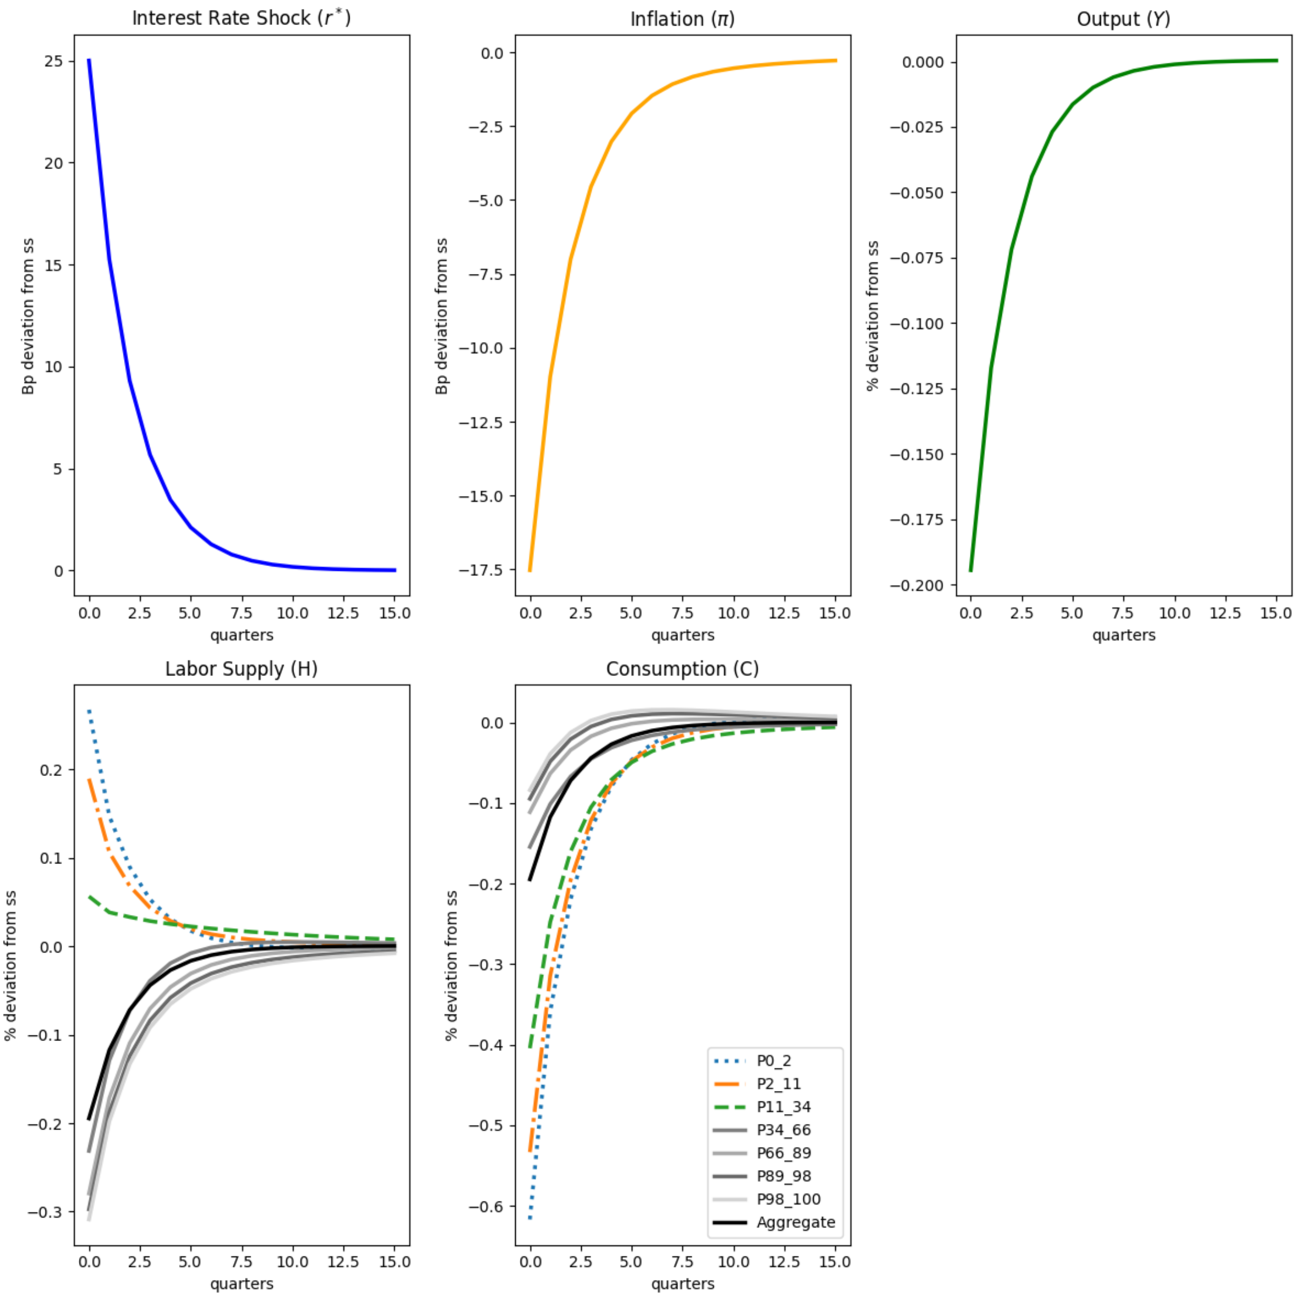

Figure 7: Impact Impulse Responses of Labor Supply Across the Income Distribution

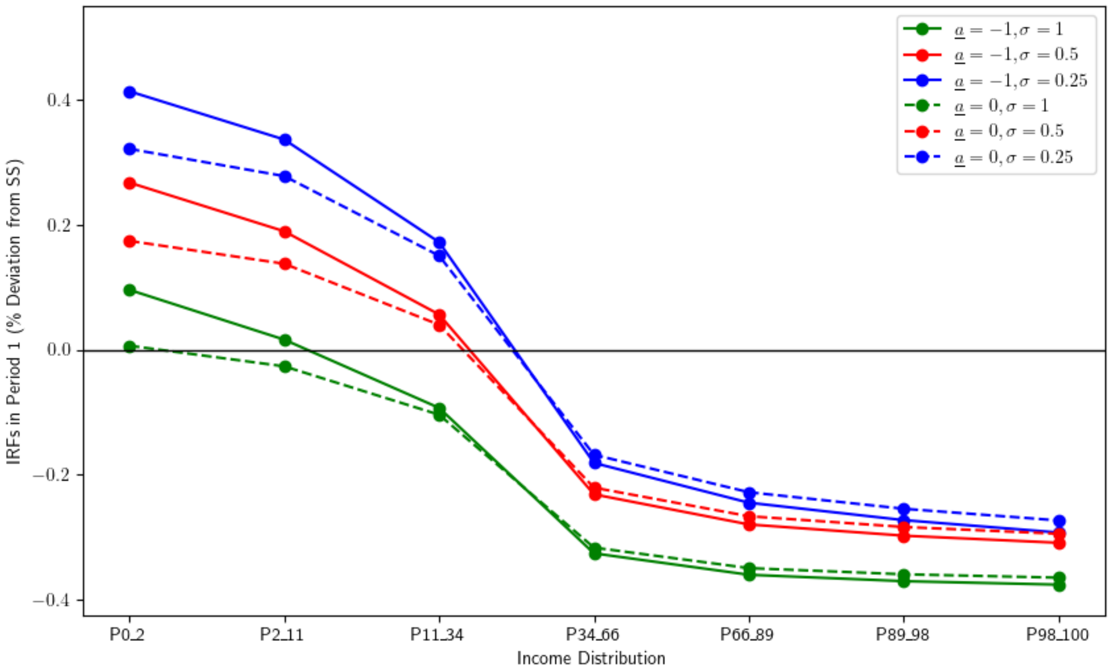

model is able to capture our empirical evidence on the heterogeneous labor supply behavior across the income distribution in response to a monetary policy shock. It is important to stress that, since the model features a representative firm, labor demand is homogeneous across agents. However, while the 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. 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.[22]

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 P0_2 to P98_100), under the six different model calibrations. Solid lines correspond to the baseline borrowing constraint

[22]Appendix G.2 reports policy functions and IRFs for all calibrations.

calibration $\underline{a} = -0.5$, allowing limited debt, while dashed lines represent a tighter borrowing constraint $\underline{a} = 0$ (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 $\underline{a} = 0$ & $\sigma = 1$, 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 ($\underline{a} = -0.5$) 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 elasticity of intertemporal substitution (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 ($\underline{a} < 0$) 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. 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 term of these main channels. The top left panel of this figure reports the decomposition of the aggregate hours worked while the subsequent panels 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. This is mainly due to the standard intertemporal substitution effect-less consumption today (more labor supply) versus more tomorrow. But there is also another channel through which higher rates increase labor supply today: an increase in the interest repayment of existing debt induces agents to work more. 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

Figure 8: Decomposition of labor supply responses in HANK.

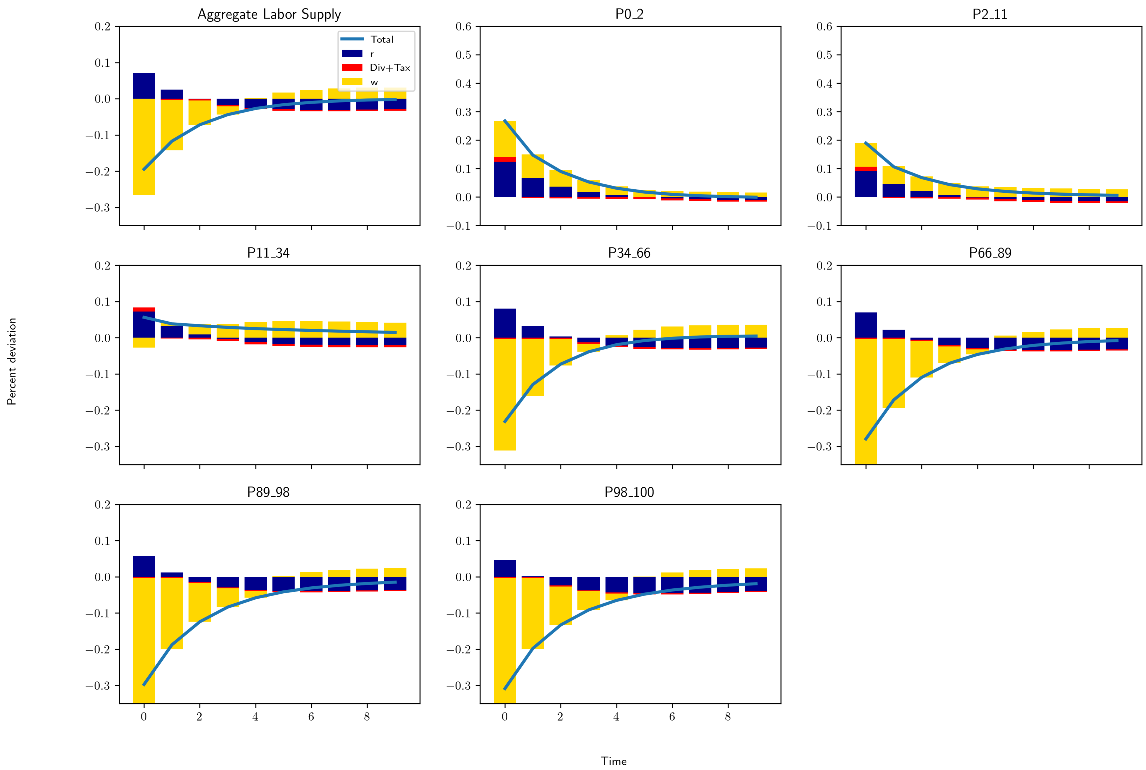

income state. First, we notice that the real rate channel on impact is stronger at the bottom of the income distribution, and it remains positive for more periods compared to agents in the middle to the top of the distribution. This is in line with the fact that the income effect from debt repayments is larger for poorer individuals. Moreover, the yellow bars flip sign at the bottom, pointing towards 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 are responsible for matching our empirical evidence on the response of hours worked at the bottom of the income distribution.[23]

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 et al. (2024), total labor hours are evenly allocated across all households by a representative union,

[23] Appendix G.3 shows the same figures for consumption.

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.[24] Figure G.21 in the 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.[25]

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: One-asset HANK steady-state comparison

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 marginal propensity to consume 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 adjust-

[24] 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 Appendix shows the IRFs match for the baseline σ case.

[25] 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.

ment 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. 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 the cumulative output loss and cumulative inflation decline over the first four quarters after the shock (i.e., over one 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.

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

Table 3: Sacrifice ratios across model calibrations.

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 elasticity of intertemporal substitution 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 act 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, N., Jansson, T., Klein, M. and Rogantini-Picco, A.: 2022, Five Facts about the Distributional Income Effects of Monetary Policy Shocks, *American Economic Review: Insights* .

Amir-Ahmadi, P., Matthes, C. and Wang, M.-C.: 2021, What Does Monetary Policy Do To Different People?, *Mimeo.*

Andersen, A. L., Johannesen, N., Jørgensen, M. and Peydró, J.-L.: 2021, Monetary policy and inequality, *Economics Working Papers 1761*, Department of Economics and Business, Universitat Pompeu Fabra.

Athreya, K., Owens, A. and Schwartzman, F.: 2017, Does redistribution increase output? the centrality of labor supply, *Quantitative Economics* **8**(3), 761–808.

Auclert, A.: 2019, Monetary Policy and the Redistribution Channel, *American Economic Review* **109**(6), 2333–2367.

Auclert, A., Bardóczy, B., Rognlie, M. and Straub, L.: 2021, Using the sequence-space jacobian to solve and estimate heterogeneous-agent models, *Econometrica* **89**(5), 2375–2408.
URL: *https://onlinelibrary.wiley.com/doi/abs/10.3982/ECTA17434*

Auclert, A. and Rognlie, M.: 2020, Inequality and Aggregate Demand, *Mimeo.*

Auclert, A., Rognlie, M. and Straub, L.: 2024, The intertemporal keynesian cross, *Journal of Political Economy* **132**(12), 4068–4121.
URL: *https://doi.org/10.1086/732531*

Bai, J. and Ng, S.: 2002, Determining the Number of Factors in Approximate Factor Models, *Econometrica* **70**(1), 191–221.

Banbura, M., Giannone, D. and Reichlin, L.: 2007, Bayesian vars with large panels, *CEPR Discussion Papers* **6326**.

Barigozzi, M., Lippi, M. and Luciani, M.: 2021, Large-dimensional Dynamic Factor Models: Estimation of Impulse-Response Functions with I(1) cointegrated factors, *Journal of Econometrics* **221**(2), 455–482.

Bauer, M. D. and Swanson, E. T.: 2023, A Reassessment of Monetary Policy Surprises and High-Frequency Identification, *NBER Macroeconomics Annual* **37**(1), 87–155.
URL: *https://ideas.repec.org/a/ucp/macann/doi10.1086-723574.html*

Bayer, C., Born, B. and Luetticke, R.: 2024, Shocks, frictions, and inequality in us business cycles, *American Economic Review* **114**(5), 1211â47.
URL: *https://www.aeaweb.org/articles?id=10.1257/aer.20201875*

Bayer, C. and Luetticke, R.: 2020, Solving heterogeneous agent models in discrete time with many idiosyncratic states by perturbation methods, *Quantitative Economics* **11**, 1253–1288.

Bilbiie, F.: 2008, Limited asset markets participation, monetary policy and (inverted) aggregate demand logic, *Journal of Economic Theory* **140**(1), 162–196.

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

Bilbiie, F. O.: 2024, Monetary policy and heterogeneity: An analytical framework*, *The Review of Economic Studies* p. rdae066.
URL: *https://doi.org/10.1093/restud/rdae066*

Bilbiie, F. O., Monacelli, T. and Perotti, R.: 2013, Public debt and redistribution with borrowing constraints, *The Economic Journal* **123**(566), F64–F98.

Broer, T., Kramer, J. and Mitman, K.: 2022, The Curious Incidence of Monetary Policy Shocks Across the Income Distribution, *Working Paper Series 416*, Sveriges Riksbank (Central Bank of Sweden).

Canova, F. and Ferroni, F.: 2022, Mind the gap! Stylized dynamic facts and structural models, *American Economic Journal: Macroeconomics* **forthcoming**.

Cesarini, D., Lindqvist, E., Notowidigdo, M. J. and Östling, R.: 2017, The effect of wealth on individual and household labor supply: Evidence from swedish lotteries, *American Economic Review* **107**(12), 3917–46.

Cloyne, Ferreira, C. and Surico, P.: 2020, Monetary Policy when Households have Debt: New Evidence on the Transmission Mechanism, *Review of Economic Studies* **87**(1), 102–129.

Coglianese, J., Olsson, M. and Patterson, C.: 2025, Monetary policy and the labor market: A quasi-experiment in sweden, *American Economic Review* .

Das, M., Hambur, J., Hellwig, K. and Spray, J.: 2025, Labor supply effects of monetary policy: Evidence from australian mortgage holders, *mimeo.*

De Giorgi, G. and Gambetti, L.: 2017, Business cycle fluctuations and the distribution of consumption, *Review of Economic Dynamics* **23**, 19 – 41.

Del Canto, F., Grigsby, J., Qian, E. and Walsh, C.: 2025, Are inflationary shocks regressive? a feasible set approach, *The Quarterly Journal of Economics* .

Dyrda, S. and Pedroni, M.: 2022, Optimal fiscal policy in a model with uninsurable idiosyncratic shocks, *Review of Economic Studies* .

Forni, M. and Gambetti, L.: 2014, Sufficient information in structural VARs, *Journal of Monetary Economics* **66**(C), 124–136.

Galí, J., Smets, F. and Wouters, R.: 2012, Unemployment in an estimated new keynesian model, *NBER Macroeconomics Annual* **26**, 329–360.

Gerke, R., Giesen, S., Lozej, M. and Rottger, J.: 2024, On household labour supply in sticky-wage HANK models, *Technical report*.

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

Golosov, M., Graber, M., Mogstad, M. and Novgorodsky, D.: 2023, How americans respond to idiosyncratic and exogenous changes in household wealth and unearned income*, *The Quarterly Journal of Economics* **139**(2), 1321–1395.
URL: *https://doi.org/10.1093/qje/qjad053*

Graves, S., Huckfeldt, C. and Swanson, E. T.: 2023, The labor demand and labor supply channels of monetary policy, *mimeo*.

Guerrieri, V. and Lorenzoni, G.: 2017, Credit Crises, Precautionary Savings, and the Liquidity Trap, *The Quarterly Journal of Economics* **132**(3), 1427–1467.

Holm, M. B., Paul, P. and Tischbirek, A.: 2021, The transmission of monetary policy under the microscope, *Journal of Political Economy* **129**(10), 2861–2904.

Hubert, P. and Savignac, F.: 2024, The heterogeneous effects of monetary policy on labor income: Disentangling the extensive and intensive margins, *Journal of Political Economy Macroeconomics* **2**(4), 687–719.
URL: *https://doi.org/10.1086/732469*

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

Katz, L. F. and Murphy, K. M.: 1992, Changes in relative wages, 1963â1987: Supply and demand factors, *The Quarterly Journal of Economics* **107**(1), 35–78.

Kehoe, P. J., Lopez, P., Pastorino, E. and Salgado, S.: 2020, On the Dynamic Effects of Monetary Policy With Heterogeneous Agents, *Mimeo*.

McKay, A., Nakamura, E. and Steinsson, J.: 2016, The power of forward guidance revisited, *American Economic Review* **106**(10), 3133–58.

Mertens, K. and Ravn, M. O.: 2013, The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States, *American Economic Review* **103**(4), 1212–47.

Miranda-Agrippino, S. and Ricco, G.: 2021, The transmission of monetary policy shocks, *American Economic Journal: Macroeconomics* **13**(3), 74–107.

Pijoan-Mas, J.: 2006, Precautionary savings or working longer hours?, *Review of Economic Dynamics* **9**(2), 326–352.

Rotemberg, J. J.: 1982, Sticky prices in the united states, *Journal of Political Economy* **90**(6), 1187–1211.

Stock, J. H. and Watson, M. W.: 2008, What's new in econometrics - time series, *Lecture 7*, National Bureau of Economic Research, Inc.

Wolf, C.: 2021, The Missing Intercept: A Demand Equivalence Approach, *National Bureau of Economic Research WP29558*.

Online Appendix

A Household level data

A.1 US: CPS Survey

For the US, individual level data is obtained from the current population survey (CPS) for each month from 1985 to 2019. In the benchmark case we use the uniform CPS outgoing rotation group extracts created by CEPR data. Our benchmark measure of hours is the variable hourslw (Hours last week, all jobs). We use the variable rw as our measure of real hourly wage. This is the recommended consistent real wage variable across the sample we employ. For details see the CEPR FAQ.

Figure A.1: Characteristics along the wage distribution. Average across the sample.

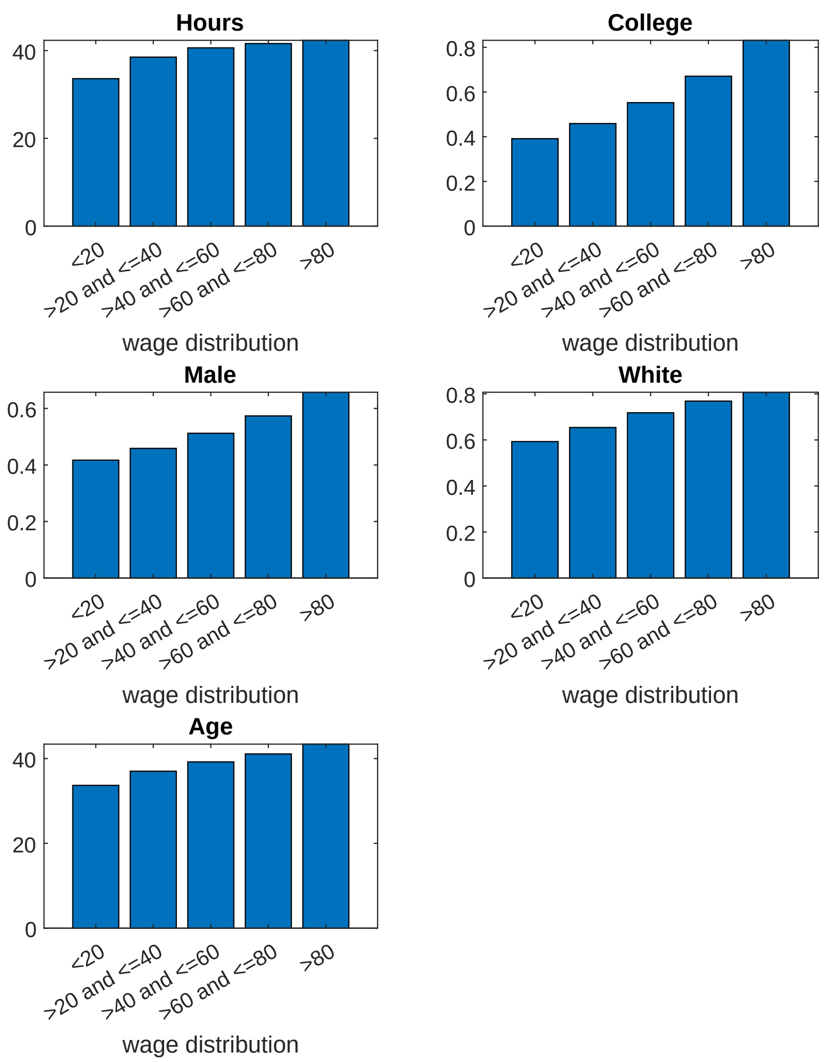

Figure A.2: Proportion of individuals in different industries.

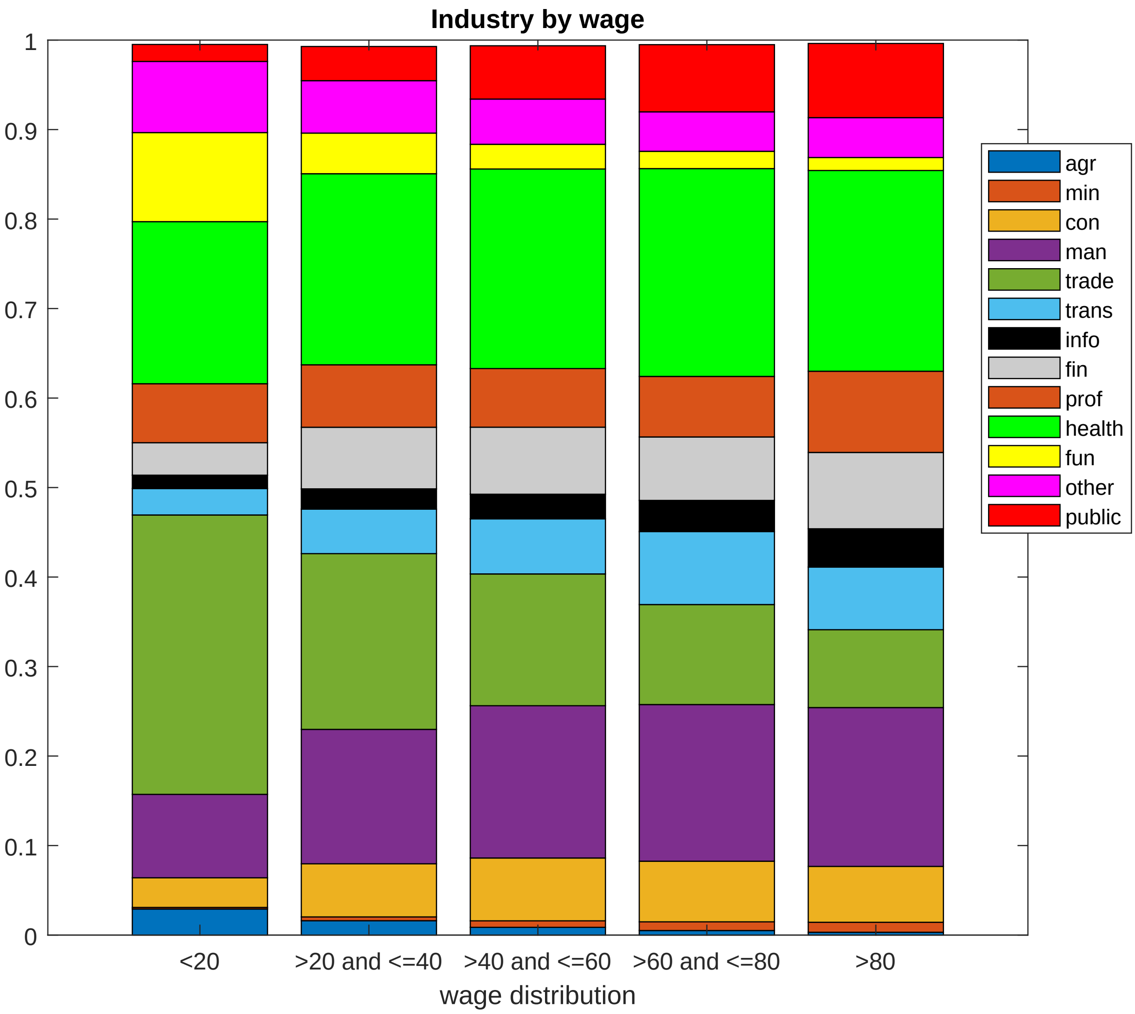

Notes: The proportions are averaged over the sample. agr is Agriculture, forestry, fishing and hunting. min is mining, con is construction, man is manufacturing, trade is wholesale and retail trade, trans is transport and utilities, info is information, fin is financial activities, prof is professional and business services, health is education and health services, fun is leisure and hospitality, other is other services and public denotes public administration.

Figure A.1 provides information regarding the characteristics of the earnings distribution. Respondents in the right tail of the distribution tend to be older, better educated, are likely to work longer hours and be white and Male. Figure A.2 shows industry of employment across the wage distribution. At the left tail, industries such as wholesale and retail trade, health, leisure and manufacturing are important.

As discussed above, we also employ the longitudinally matched version of CPS provided by the Kansas Fed. We employ the variable wageperhrclean82 as our measure of hourly earnings. This series is deflated by CPI. The change in hours is constructed as the difference in the log of hours82 and hours82_tm12 where the former denotes a consistent series constructed using actual or usual hours and the latter is the 12 month lag of this variable. We construct the rate of transition from employment using the labour market status variable mlr76 and the 12 month lag of this variable mlr76_tm12. As discussed in the text, we impute earnings for individuals that are not employed or do not report earnings data. To do this we regress earnings on experience and individual characteristics including gender, occupation, industry of employment, geographical characteristics (US state, metropolitan indicator) and race. The fitted values from this regression are used to obtain imputed earnings.The coefficients of the regression are shocked and used to produce predicted values. These are assigned to individuals with missing earnings data using predicted mean matching. We produce 5 replicates, with the final imputed data taken to be the mean across these replications.

A.2 Comparison with aggregate data

The top panel of Figure A.3 compares aggregate actual hours from the CPS to a measure of monthly hours from the Bureau of Labour statistics (average weekly hours of production and non-supervisory employees). The CPS data captures the main movements in the aggregate data fairly well.

Figure A.3: Comparison of survey based total hours (blue) with aggregate (orange).

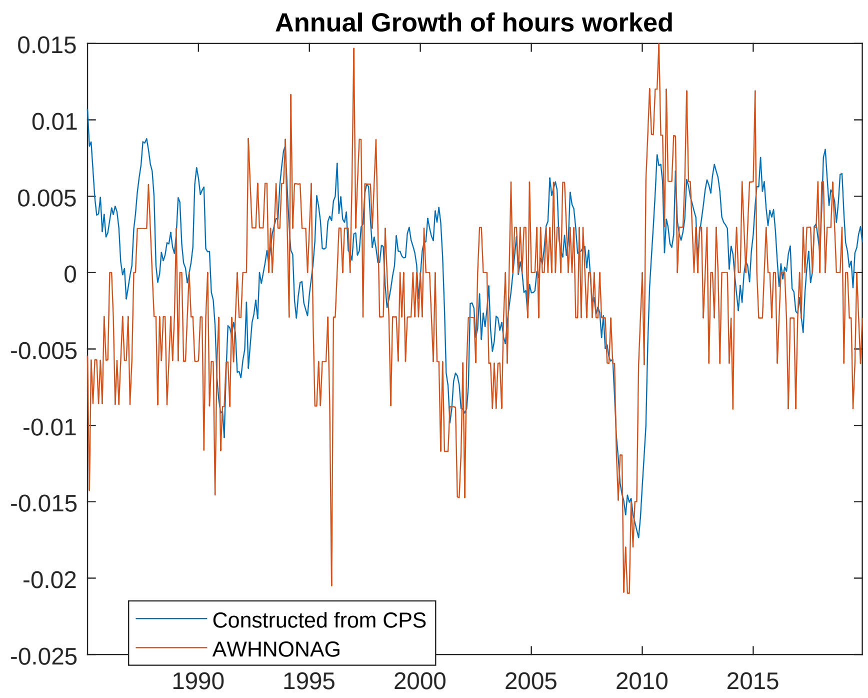

B Priors and estimation of the FAVAR

The FAVAR model is defined by the following equations:

$$X_{it} = c_i + b_i\tau + \Lambda_i F_t + \xi_{it} \tag{B.1}$$

$$Y_t = c + \sum_{j=1}^{P} \beta_j Y_{t-j} + u_t \tag{B.2}$$

$$cov(u_t) = \Sigma = A_0 A_0' \tag{B.3}$$

Where $Y_t = \underbrace{\begin{pmatrix} R_t \\ F_t \end{pmatrix}}_{N\times 1}$, R_t denotes the 1 year interest rate, $i = 1, 2, .., M$ denotes the cross-sectional dimension of the panel data-set X_{it} while $t = 1, 2, .., T$ is the dimension. As described in Barigozzi et al. (2021), the factors can be consistently estimated using a principal components (PC) estimator. In particular, the factor loadings are estimated via PC analysis of the first differenced data ΔX_{it}. With these in hand, the factors are estimated as $\hat{F}_t = \frac{1}{M}\left(\hat{\Lambda}'\tilde{X}_t\right)$. Here,$\Lambda$ is the matrix of factor loadings, $\tilde{X}_t$ is given by $(x_{1t}, x_{2t}, ..., x_{Mt})$ where $x_{it} = X_{it} - \hat{c}_i - \hat{b}_i\tau$ Note that Barigozzi et al. (2021) describe a procedure to check if the ith series contains a linear trend and that $\hat{b}_i$ is different from zero.

Given the estimated factors, the VAR in equations B.2 is estimated using a Bayesian methods.

B.1 Priors

Denote the var coefficients as $B = vec([\beta_1, \beta_2, .., \beta_P, c])$. We follow Banbura, Giannone and Reichlin (2007) and use a Natural Conjugate prior implemented via dummy observations. The priors are implemented by the dummy observations y_D and x_D that are defined as:

$$y_D = \begin{bmatrix} \frac{diag(\gamma_1 s_1 ... \gamma_n s_n)}{\kappa} \\ 0_{N\times(P-1)\times N} \\ diag(s_1 ... s_n) \\ .... \\ 0_{EX\times N} \end{bmatrix}, \quad x_D = \begin{bmatrix} \frac{J_P \otimes diag(s_1...s_n)}{\kappa} & 0_{NP\times EX} \\ .... & \\ 0_{N\times(NP)+EX} & \\ .... & \\ 0_{EX\times NP} & I_{EX}\times 1/c \end{bmatrix} \tag{B.4}$$

where $J_P = diag(1, 2, ..., P)$, γ_1 to γ_n denote the prior mean for the parameters on the first lag obtained by estimating individual AR(1) regressions, s_1 to s_n is an estimate of the variance of the endogenous variables obtained individual AR(1) regressions, κ measures the tightness of the prior on the autoregressive VAR coefficients, and c is the tightness of the prior on the remaining regressors. We set $\kappa = 0.2$ and $c = 1000$. We also implement priors on the sum of coefficients (see Banbura et al. (2007)). The dummy observations for this prior are defined as:

$$\tilde{y}_D = \frac{diag(\gamma_1\mu_1...\gamma_n\mu_n)}{\tau}, \tilde{x}_D = \begin{pmatrix} (1_{1\times P}) \otimes \frac{diag(\gamma_1\mu_1...\gamma_n\mu_n)}{\tau} & 0_{N\times EX} \end{pmatrix} \tag{B.5}$$

where μ_i is the sample average of the ith variable. As in Banbura et al. (2007) we set $\tau = 10\kappa$. The total number of dummy observations is T_D.

B.2 MCMC algorithm

Banbura et al. (2007) show that posterior distribution can be written as:

$$g(\Sigma|Y) \sim iW\left(\bar{\Sigma}, T_D + 2 + T - K\right) \tag{B.6}$$

$$g(B|\Sigma, Y) \sim N\left(\bar{B}, \Sigma \otimes \left(X_*'X_*\right)^{-1}\right) \tag{B.7}$$

where iW denotes the inverse Wishart distribution, K denotes the number of regressors in each equation of the VAR model. Note that $Y_* = \begin{pmatrix} Y \\ y_D \\ \tilde{y}_D \end{pmatrix}$ and $X_* = \begin{pmatrix} X \\ x_D \\ \tilde{x}_D \end{pmatrix}$, X collects the regressors, and

$$\begin{aligned} \tilde{B} &= \left(X_*'X_*\right)^{-1}\left(X_*'Y_*\right) \\ \bar{B} &= vec\left(\tilde{B}\right) \\ \bar{\Sigma} &= \left(Y_* - X_*\tilde{B}\right)'\left(Y_* - X_*\tilde{B}\right) \end{aligned}$$

Posterior draws can be easily generated by drawing Σ from the marginal distribution in B.6 and then b from the conditional distribution in equation B.7. We set the number of draws to 21,000 with a burn-in of 1,000. We retain every second draw after the burn-in period.

C IV Identification

For a given draw of B, Σ and u_t, we obtain the first column of $A_0$by using the procedure proposed by Mertens and Ravn (2013). We assume that the instrument is relevant and exogenous:

$$\begin{aligned} cov\left(m_t, \varepsilon_{1t}\right) &= \alpha \\ cov\left(m_t, \varepsilon_t^-\right) &= 0 \end{aligned}$$

where ε_{1t} denotes the structural shock of interest that is ordered first for convenience, while ε_t^-represent all remaining shocks and $\varepsilon_t = \begin{pmatrix} \varepsilon_{1t} & \varepsilon_t^- \end{pmatrix}$. Re-writing the relevance and exogeneity conditions in vector form:

$$E(m_t\varepsilon_t') = [\alpha \quad 0] \tag{C.1}$$

$$E(m_t\varepsilon_t'A_0') = [\alpha \quad 0]A_0' \tag{C.2}$$

$$E(m_t u_t') = \alpha a_0 \tag{C.3}$$

where a_0 is a $(1 \times R)$ vector corresponding to the first row of A_0' (hence first column of A_0). An estimate of $E(m_t u_t') = \begin{pmatrix} E(m_t u_{1t}') \\ E(m_t u_{2t}') \\ . \\ . \\ E(m_t u_{Nt}') \end{pmatrix}'$ can be easily obtained by using a linear regression. However, α on the RHS of equation C.3 is unknown. This parameter can be eliminated by normalising the left and the right hand side by dividing by the first element of $E(m_t u_t')$ and a_0, respectively. Therefore the impulse vector to a unit shock is given by

$$\tilde{a}_0 = \begin{pmatrix} 1 \\ \frac{E(m_t u_{2t}')}{E(m_t u_{1t}')} \\ . \\ . \\ \frac{E(m_t u_{Nt}')}{E(m_t u_{1t}')} \end{pmatrix}' .$$

D Robustness

D.1 Identification

We employ the instrument of Miranda-Agrippino and Ricco (2021). Miranda-Agrippino and Ricco (2021) argue that high frequency instruments such as those used in Gertler and Karadi (2015) contain information about the policy shock and a signal regarding central bank's information about the state of the economy. Miranda-Agrippino and Ricco (2021) construct their proxy as the high frequency changes in federal funds futures that are orthogonal to Greenbook forecasts and data revisions.

Figure D.1 shows the response of the distribution of hours to a contractionary policy shock obtained from this model. As in the benchmark case, hours increase towards the left tail of the earnings distribution.

As an alternative identification scheme, we employ sign restrictions. We assume that a contractionary monetary policy shock increases the one year rate,the Federal Funds rate, the unemployment rate and the excess bond premium and leads to a decrease in CPI and industrial production. The restrictions are imposed on the first 6 periods after the shock. Figure D.2 shows the response of hours to one standard deviation shock. While the response are less precisely estimated, the point to the conclusions obtained from the benchmark model. A contractionary monetary policy shock leads to an increase in hours at the left tail of the distribution.

Figure D.1: Response of hours worked

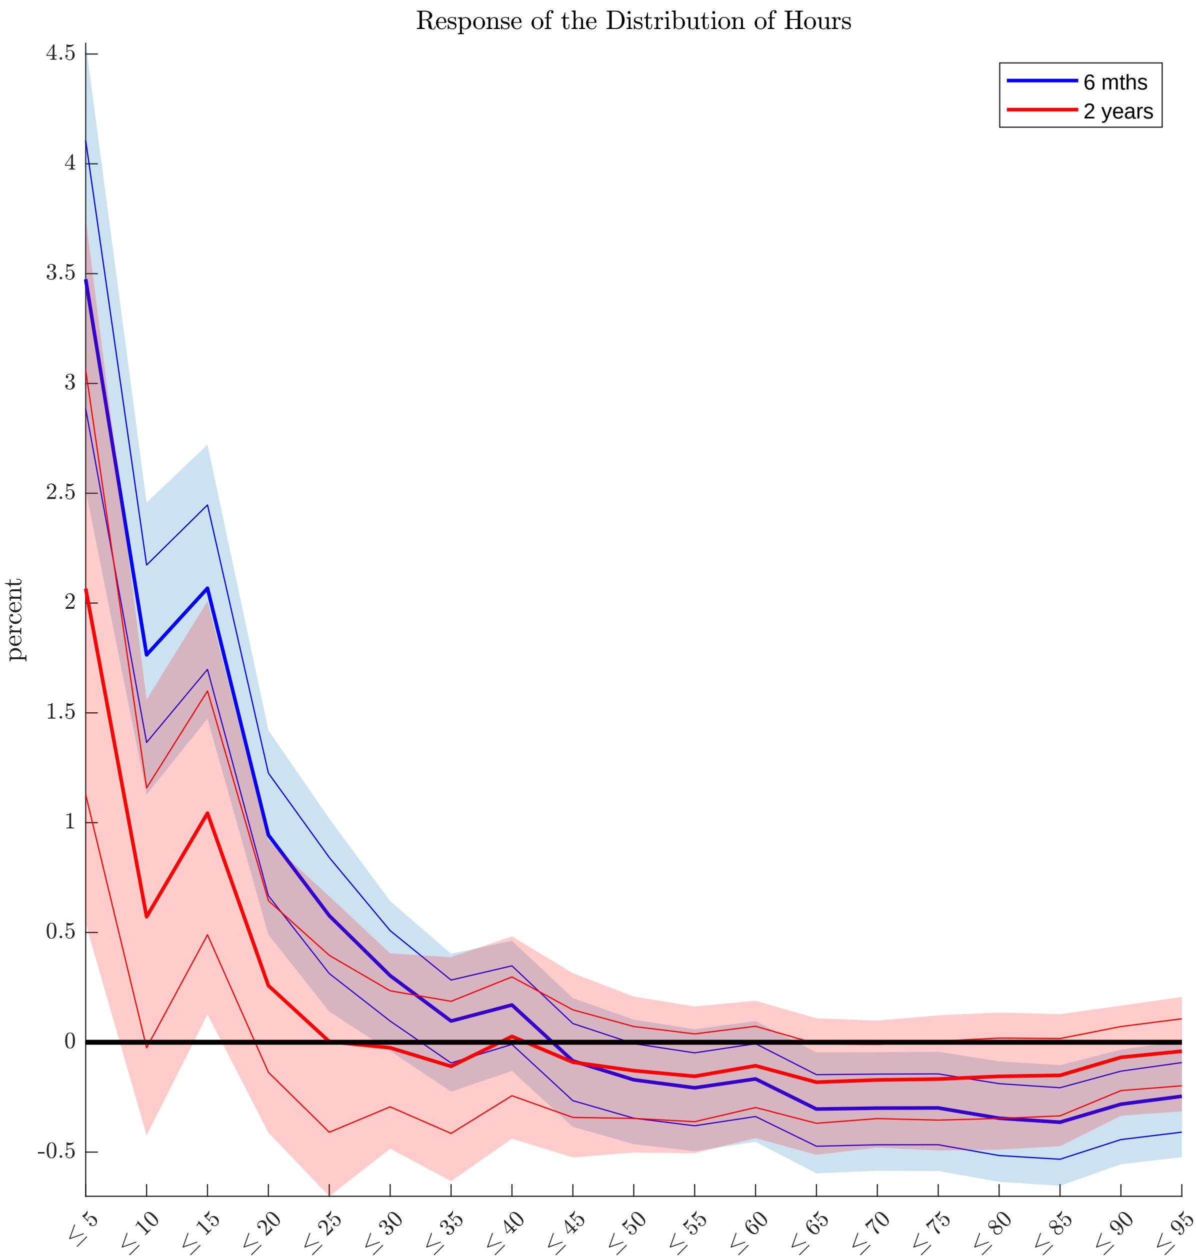

Notes: Identification using the instrument of Miranda-Agrippino and Ricco (2021). The blue line and shaded area represent the response at a 6-month horizon, while the red line and shaded area illustrate the response at a 2-year horizon.

Figure D.2: Response of hours worked

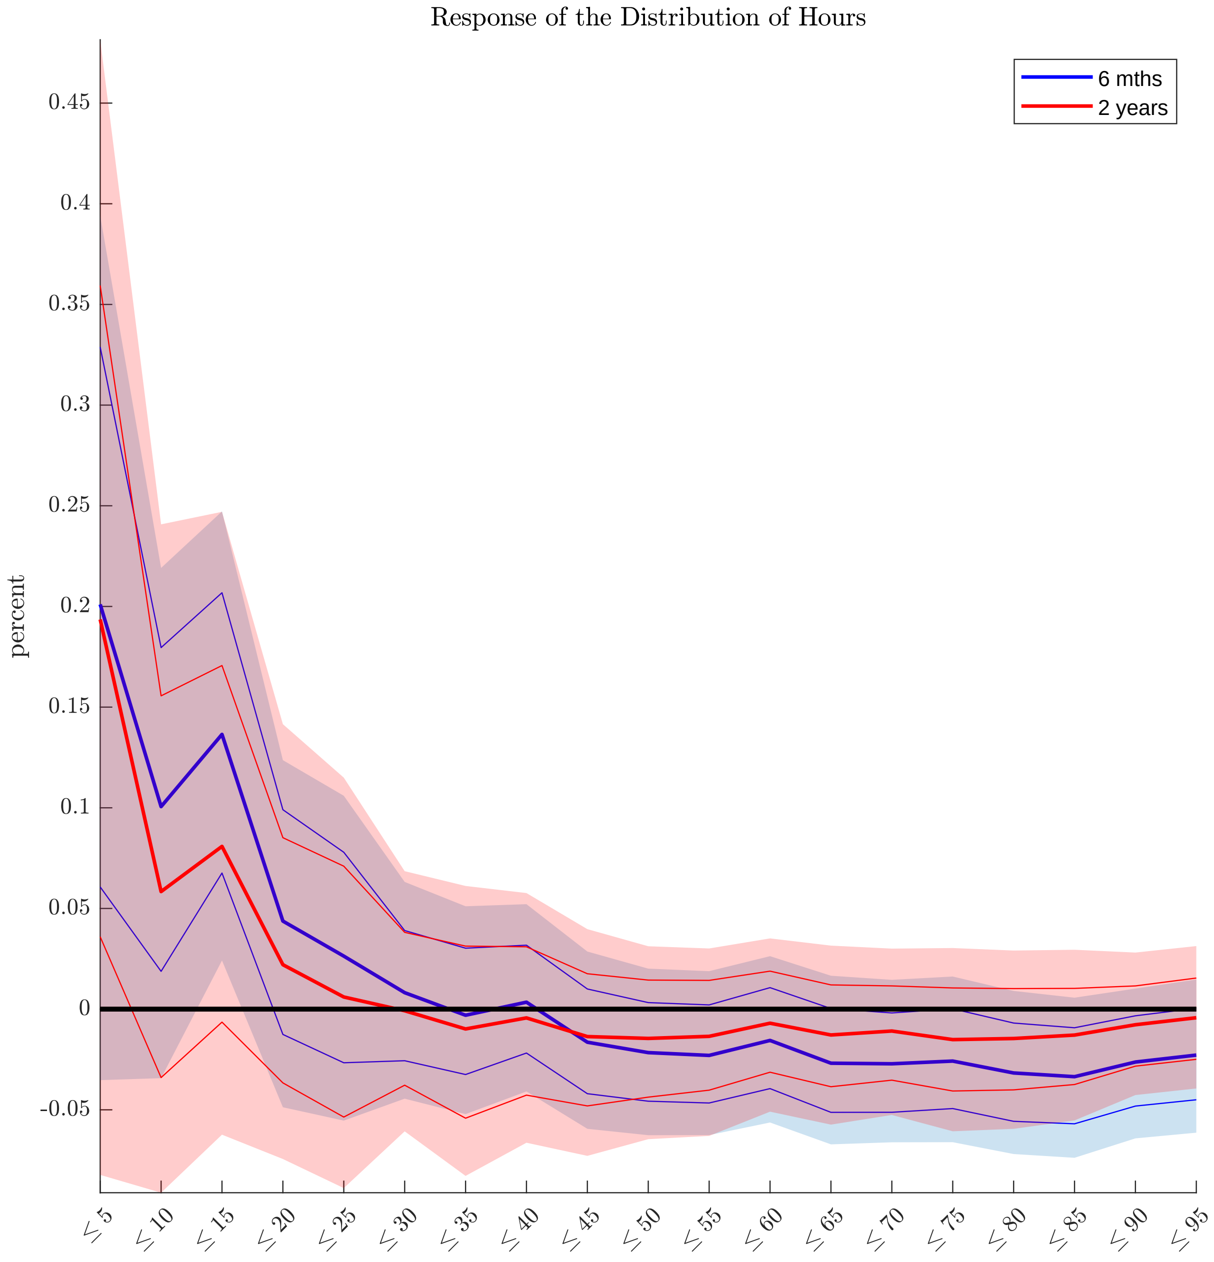

Notes: Identification using sign restrictions: The blue line and shaded area represent the response at a 6-month horizon, while the red line and shaded area illustrate the response at a 2-year horizon.

D.2 Hours by industry and by education attainments

Figure D.3 shows the proportion of workers in 13 key industries across the overall wage distribution. It is clear the left tail features workers in a variety of industries with the bulk employed in wholesale and retail trade, leisure and hospitality and education and health services. Construction constitutes a relatively small proportion of the total.

We re-estimate our benchmark model, splitting hours in each wage percentile group by the industry. In other words, for each group defined by the wage distribution (i.e. less than the 5th percentile, less than 10th percentile and so on until the 95th percentile of wage), we split hours for individuals in to the 13 industry groups discussed above. [26] Figure D.4 presents the impulse response of hours in each industry along the wage distribution. The responses at the left tail vary substantially. Sectors such as manufacturing, wholesale and retail trade and Education and health services show an increase in hours at the left tail after a monetary contraction. Some sectors such as construction, leisure and hospitality are more cyclical with hours declining at the left tail.

Figure D.5 shows that the left tail of the wage distribution is dominated by individuals with school education and some college.

Figure D.6 zooms in on employment by industry/education level for the left tail of the wage distribution. As noted above, Wholesale and Retail trade, Health and Education are important sectors of low wage employees. This figure shows that in these industries, most employees have high school education, followed by those with some college education.

We re-estimate our benchmark model, splitting hours in each wage percentile group by the level of education. In other words, for each group defined by the wage distribution (i.e. less than the 5th percentile, less than 10th percentile and so on until the 95th percentile of wage), we split hours for individuals with different level of education. We include average hours for each wage group in the FAVAR model.

Figure D.7 shows the response of hours in wage percentile for the different education groups. While the response of LTHS and somecollege displays strong cyclical behavior, hours of low wage individuals with high school or college education increases substantially.

[26] The model has a substantially larger cross-section and we include to 25 factors to account for the aggregate and disaggregated variables adequately. To keep the model computationally tractable, we reduce the number of lags to 6.

Figure D.3: Proportion of individuals in different industries.

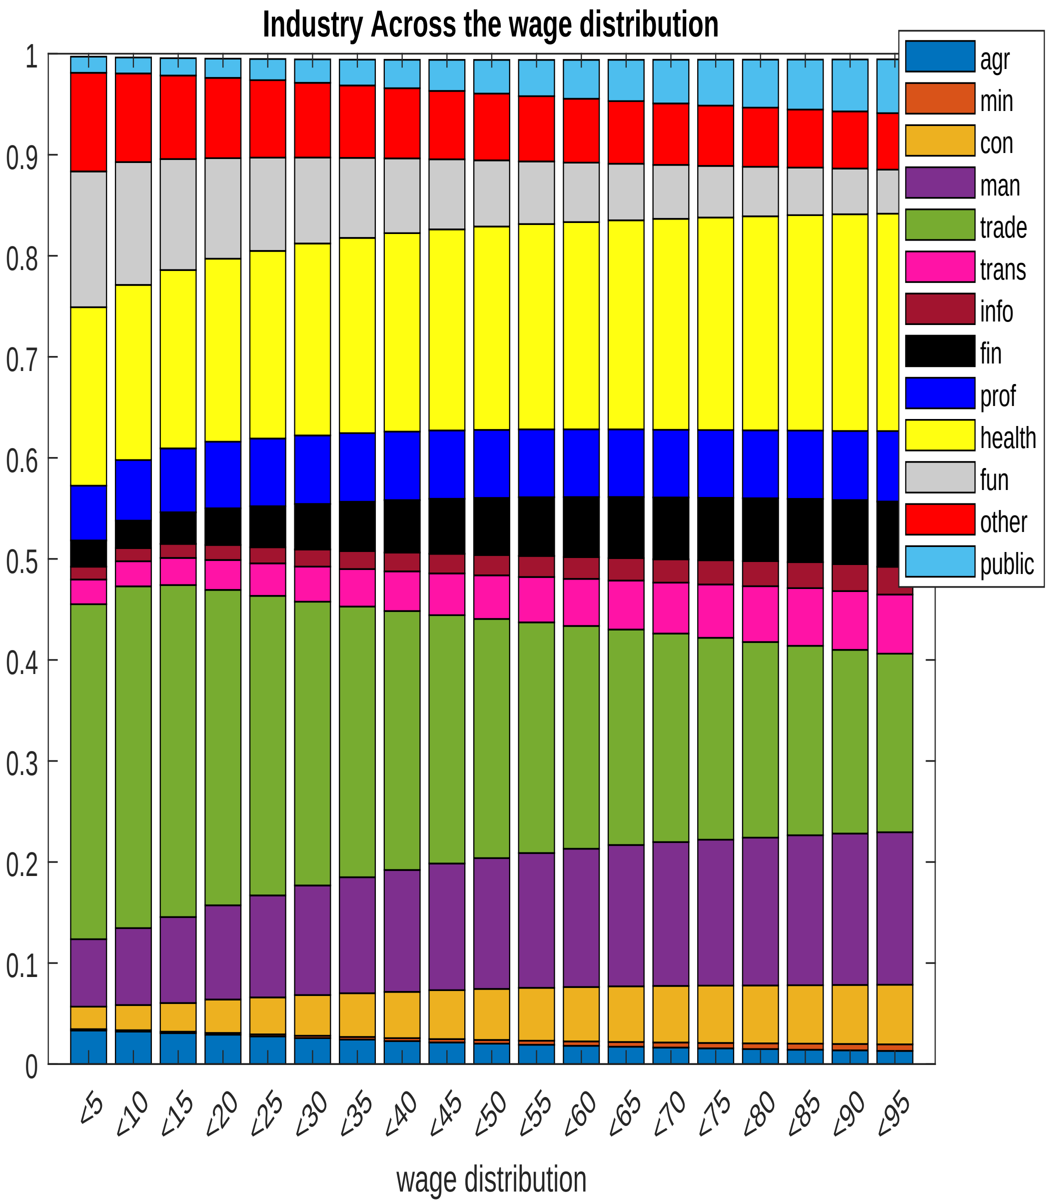

Notes: The proportions are averaged over the sample. agr is Agriculture, forestry, fishing and hunting. min is mining, con is construction, man is manufacturing, trade is Wholesale and retail trade, trans is transport and utilities, info is information, fin is financial activities, prof is professional and business services, health is education and health services, fun is leisure and hospitality, other is other services and public denotes public administration.

Figure D.4: IRF of hours at the 6 mth horizon to a monetary contraction.

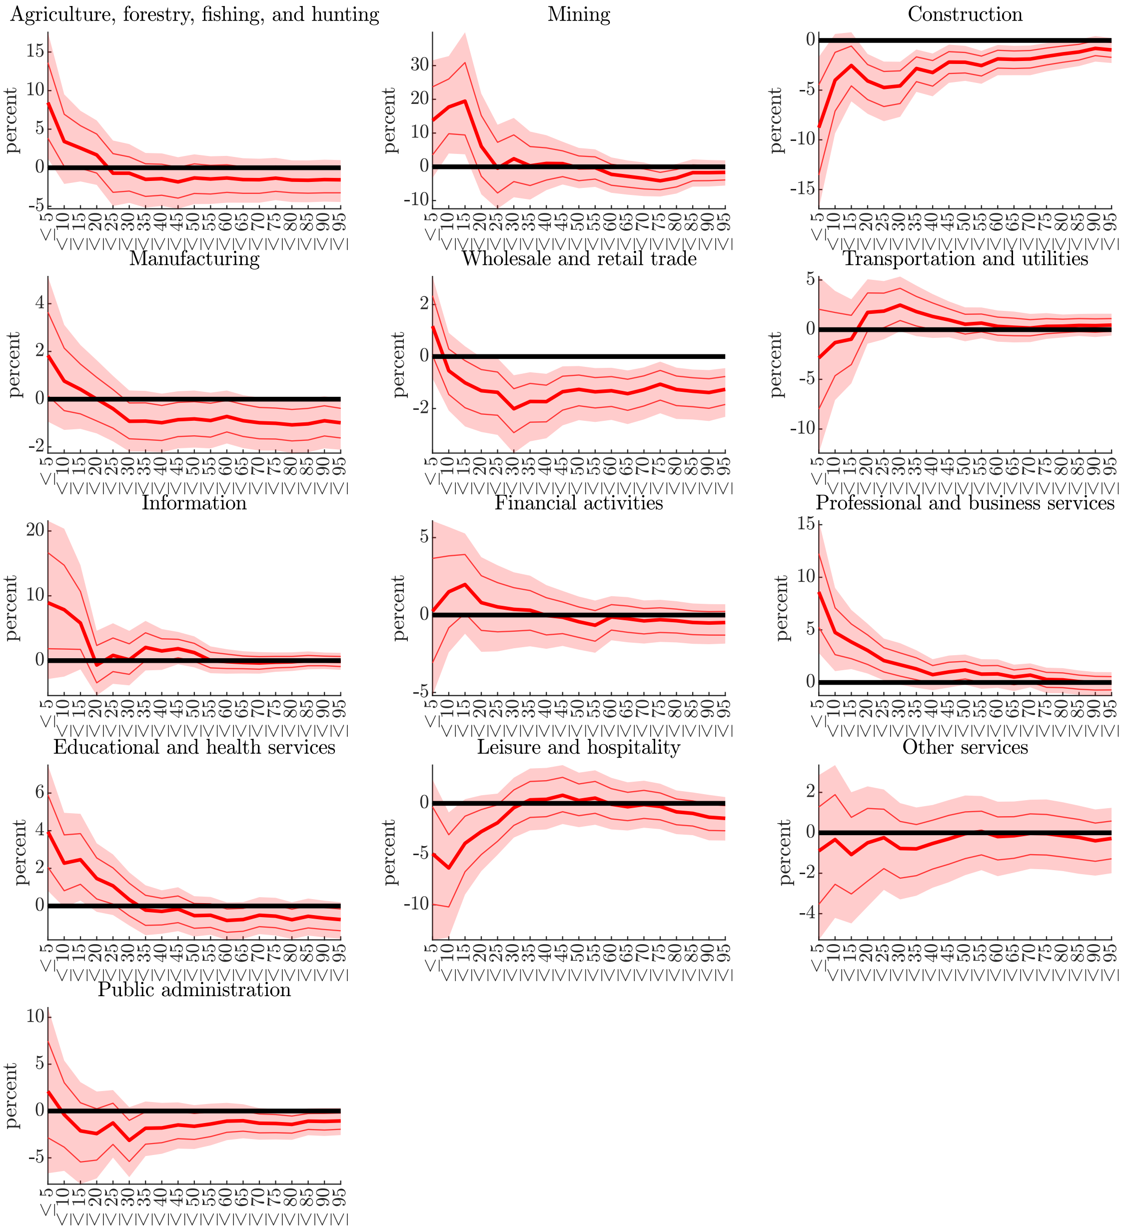

Notes: The X-axis of each figure denotes the percentiles of the earnings distribution.

Figure D.5: Proportion of individuals with education less than high school (LTHS), high school (HS), some college (somecollege), college and advanced degrees (advanced).

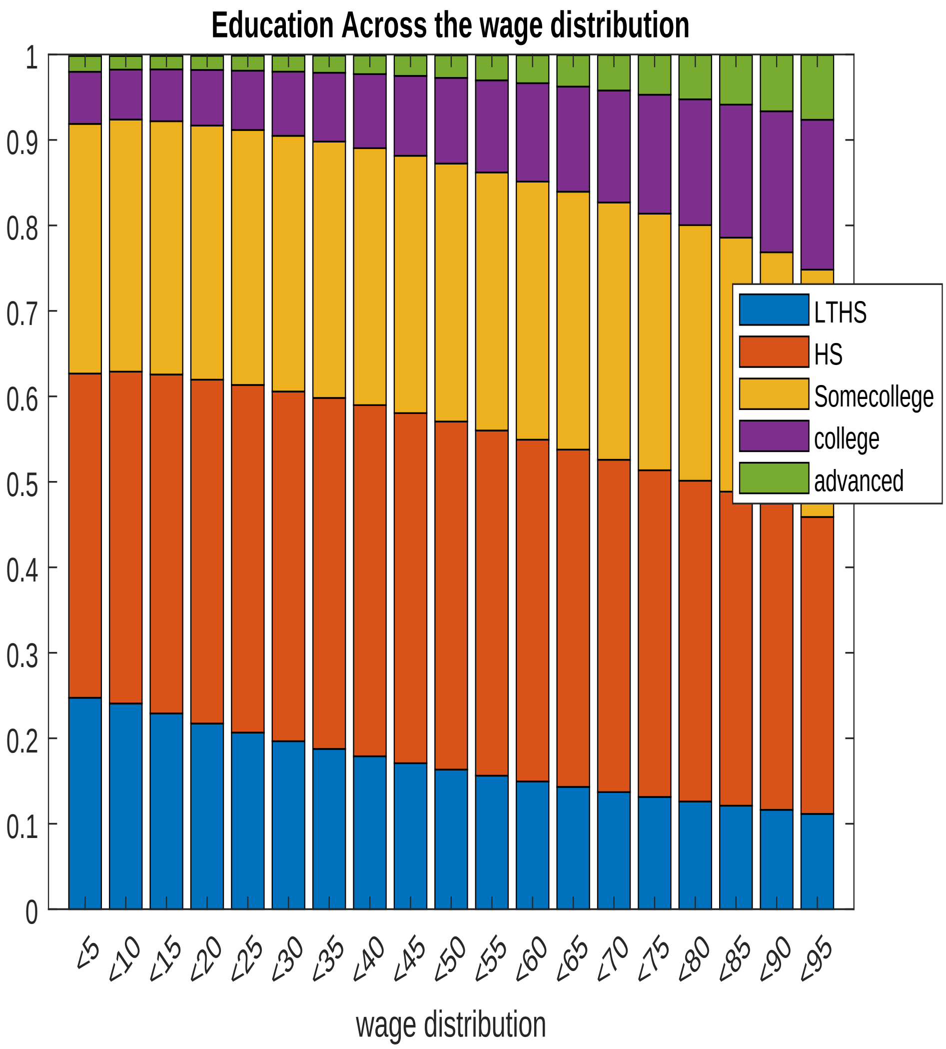

Notes: The proportions are averaged over the sample.

Figure D.6: Employees by level of education in each industry for individuals earning below the 20th percentile of wage.

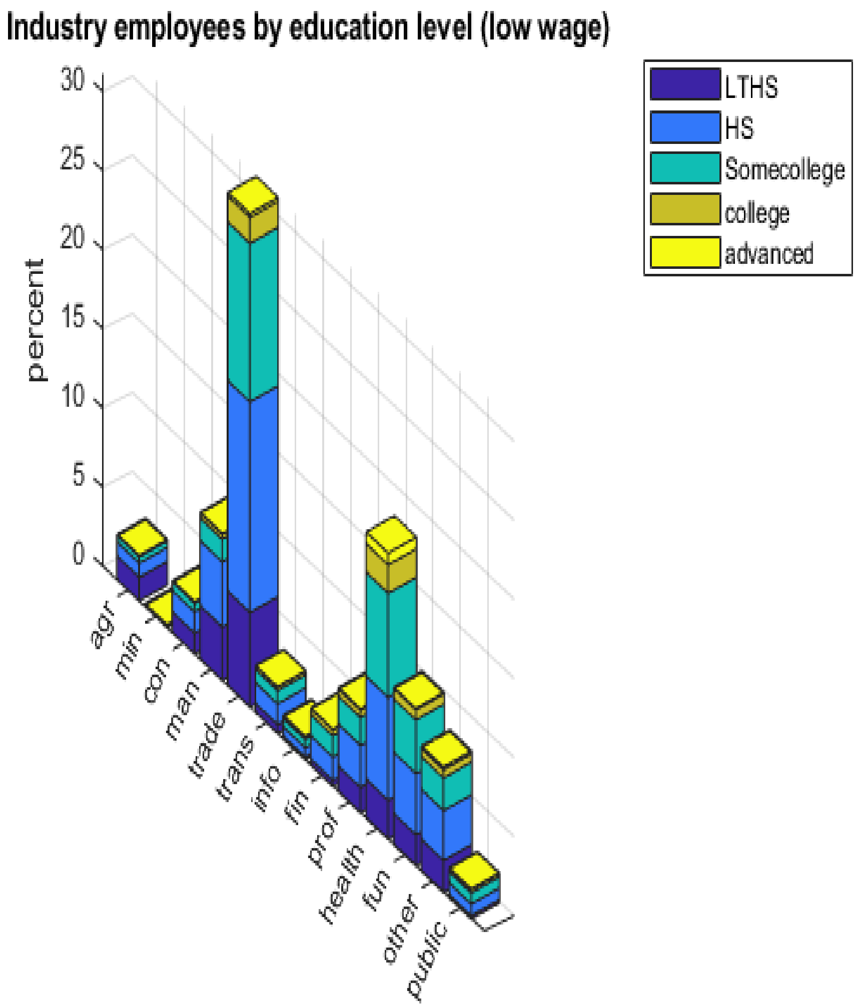

Notes: Proportion of individuals with education less than high school (LTHS), high school (HS), some college (somecollege), college and advanced degrees (advanced). The proportions are averaged over the sample.

Figure D.7: IRF of hours at the 6mth horizon to a monetary contraction

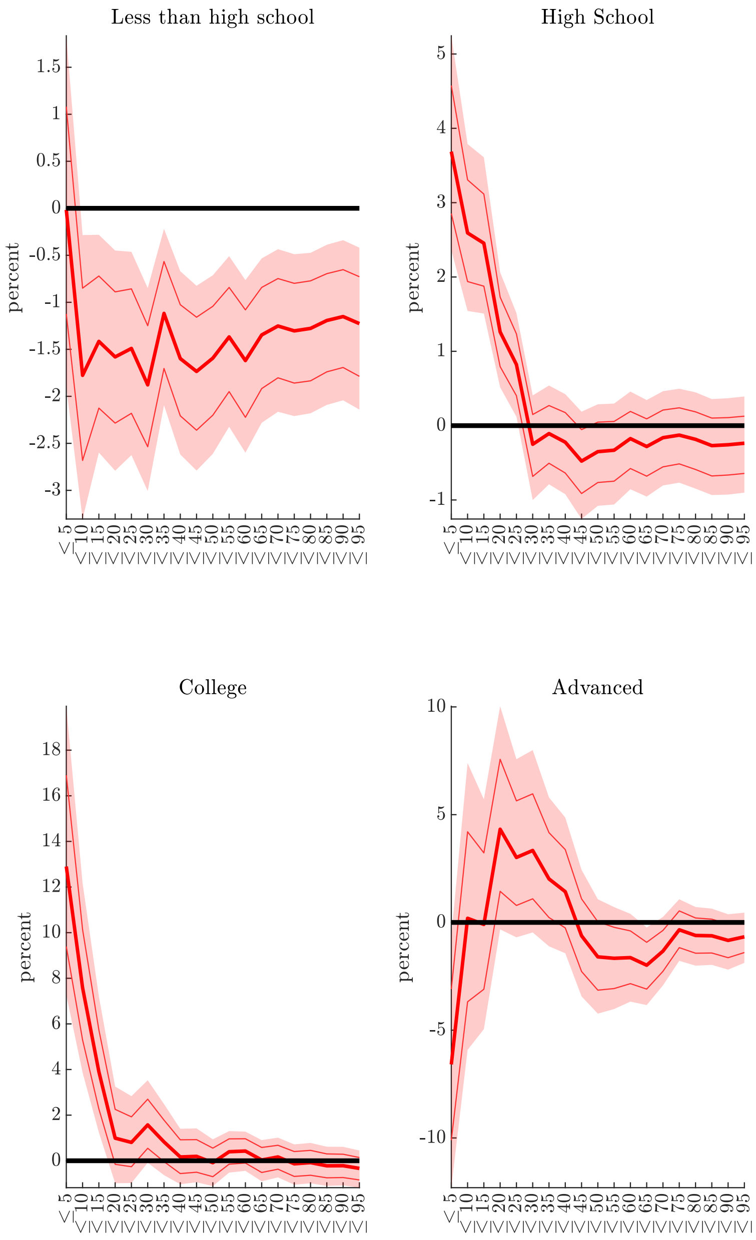

D.3 Response of real wages

Figure D.8 reports the responses of real wages both in the aggregate and across different percentiles of the income distribution. We find that for individuals in the middle and upper parts of the distribution, real wages and hours worked both decline following a monetary tightening. This pattern is consistent with a reduction in labor demand.

In contrast, for individuals in the lower end of the distribution, real wages also decline, but hours worked increase. This divergence suggests that, in this group, an increase in labor supply is the dominant force driving the response. The overall picture thus points to heterogeneous labor market adjustment mechanisms across the income distribution, with labor supply playing a more prominent role for lower-income workers.

Figure D.8: Responses of real wages after monetary policy tightening.

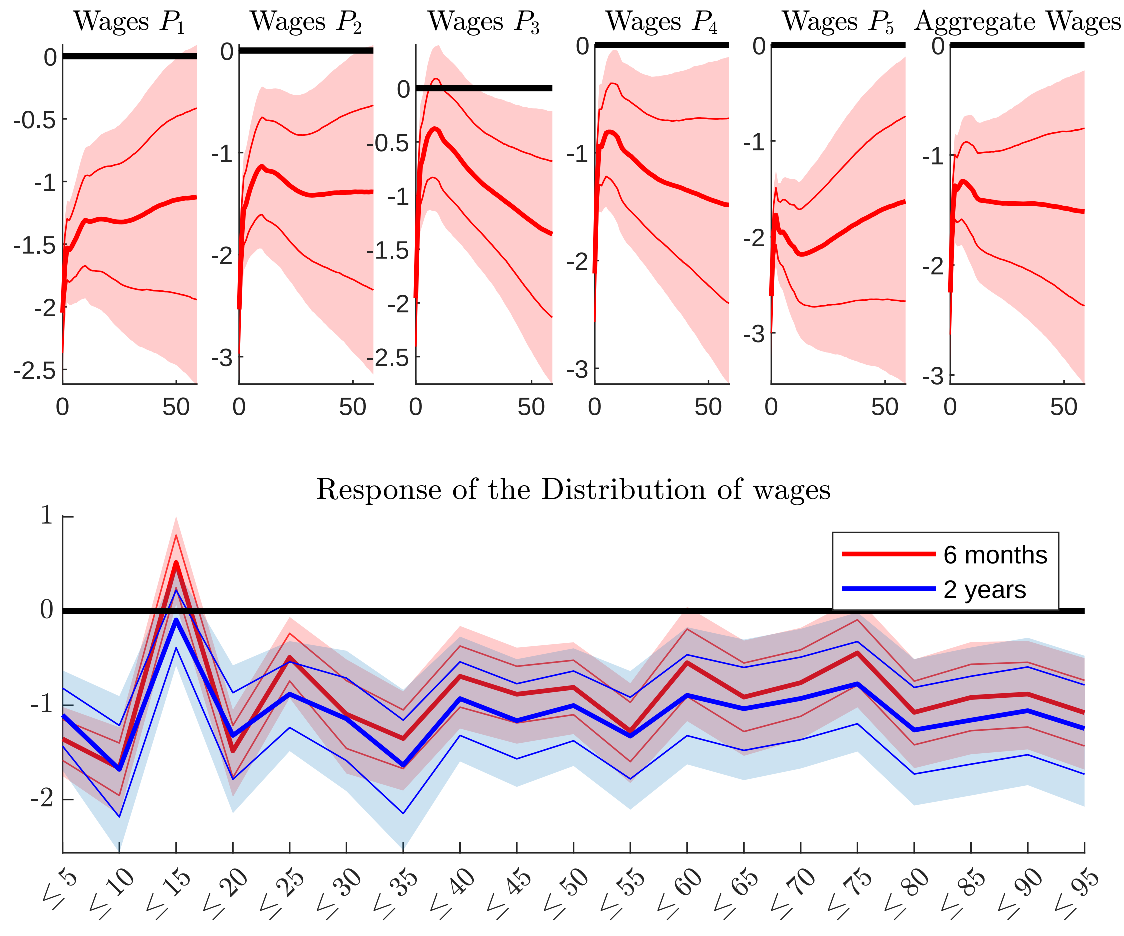

Notes: The top panels show the response of the real wages at different quintile of the wage distribution and the aggregate real wages (rightmost panel). The bottom panel displays the distribution of the responses of real wages six months (red) and two (blue) years after the shock. Shaded areas (solid lines) 68 (90)% confidence sets.

E Sample of Individuals Employed Continuously for Three Months

Figure E.1: IRF of hours for Individuals Employed Continuously for Three Months

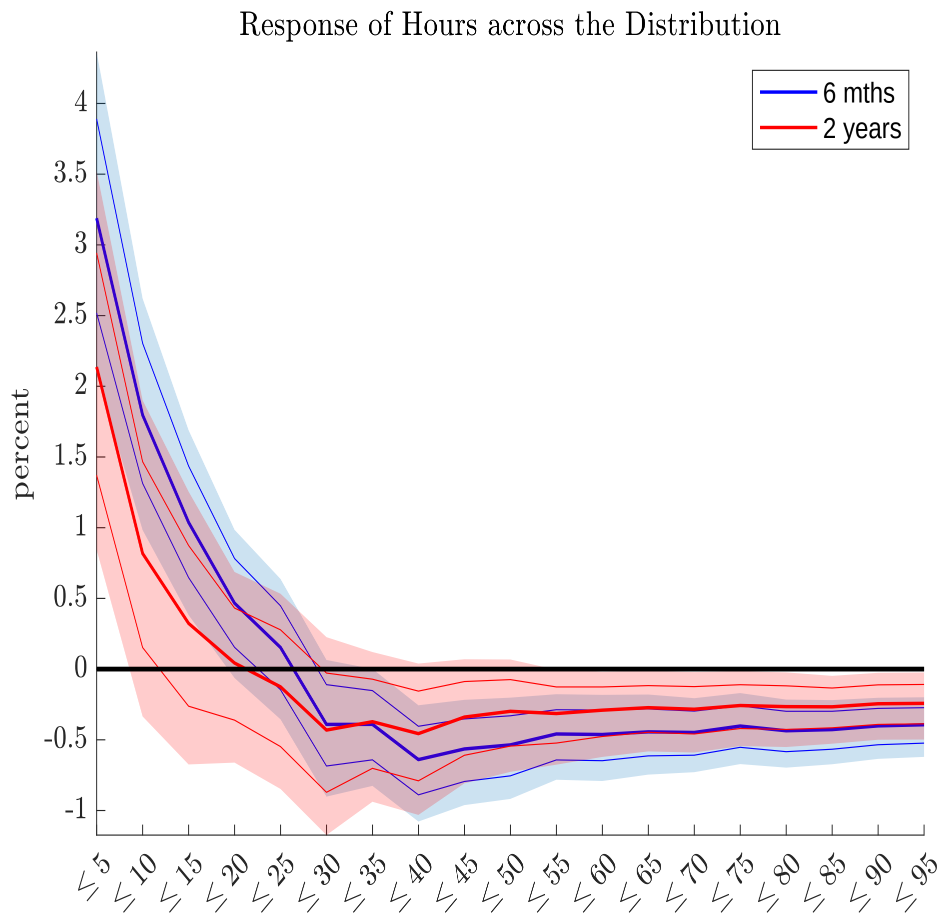

We restrict the sample to workers employed for three months by setting the variables `mlr76`, `mlr76_tm1`, and `mlr76_tm2` equal to 1. For this exercise we use the Kansas Fed extract. Figure E.1 shows that Hours for this cohort go up at the left tail of the Wage distribution after a contractionary monetary policy shock.

F TANK model with Borrowers and Savers

A simple way to characterize the household heterogeneity and retain analytical tractability is to consider a two agents model where one type of agents are net *borrower* (indexed with B) and the other net *saver* (indexed with S), as in Bilbiie, Monacelli and Perotti (2013).[27] The key features of this class of models are that borrowers are more impatient than savers, have no access to government bonds, and can borrow up to a limit. Following Bilbiie (2020) and McKay et al. (2016) we consider an economy where the monetary authority can effectively choose the real rate which makes the algebra simpler.

The key equations of the log-linearized model are reported in table F.1.

	Log-linearized Conditions	
1:	Labor Supply S	$\nu \hat{H}_t^S = \hat{w}_t - \sigma^{-1}\hat{c}_t^S$
2:	Euler S	$\hat{c}_t^S = \hat{c}_{t+1\|t}^S - \sigma\left(\hat{R}_t - \hat{\Pi}_{t+1\|t}\right)$
3:	Labor Supply B	$\nu \hat{H}_t^B = \hat{w}_t - \sigma^{-1}\hat{c}_t^B$
4:	Budget constraint B	$\hat{c}_t^B \gamma + \bar{D}(\hat{R}_{t-1} - \hat{\Pi}_t) = \left(\hat{w}_t + \hat{H}_t^B\right)$
5:	Phillips Curve	$\hat{\Pi}_t = \beta E_t \hat{\Pi}_{t+1} + \kappa \hat{w}_t$
6:	Aggregate C	$\hat{c}_t = \lambda\gamma\hat{c}_t^B + (1-\lambda\gamma)\hat{c}_t^S$
7:	Aggregate B	$\hat{c}_t = \hat{H}_t = \lambda \hat{H}_t^B + (1-\lambda)\hat{H}_t^S$
8:	Taylor Rule	$\hat{R}_t = \hat{\Pi}_{t+1\|t} + \epsilon_t^m$

Table F.1: Log-linearized Conditions of Savers/Borrowers model

This section follows a different notational convention that the one used for the HANK model in the main text. Small case letters represent real variables while capital letters represent nominal variables or hours worked. A 'hat' on top of the variable denotes percentage deviations from the steady state. The assumption that borrowers discount more future consumption, $\beta^B < \beta^S = \beta$, implies that they become net borrower in equilibrium with the borrowing limit ($\bar{D}$) always binding.[28] γ is a steady state parameter which captures the consumption inequality between borrowers and savers, i.e. $\gamma = c^B/c = 1 + \bar{D}(\beta - 1) < 1$. Notice that when $\bar{D} = 0 \rightarrow \gamma = 1$ and the model is identical to the one with Savers and HTM consumers.[29] In this model, a fraction of agents are not on the Euler equation and cannot optimize intertemporally. Importantly, and differently from a model with HTM that are not net borrowers, a change in the nominal rate will have an impact not only on the

[27] See their paper for details on the model derivation. Relative to them we simplify it further by abstracting from government debt, expenditure, and redistribution concerns. We follow Bilbiie (2020) and assume that there is a production subsidy that induces marginal cost pricing which implies that the steady state of marginal costs is 1 which simplifies substantially the steady state and the log-linearized conditions.

[28] Note that in equation (4) in table F.1 $\bar{D}$ is effectively divided by the steady state of total income/consumption. But this is =1 one in this simple set up.

[29] In that set up, Bilbiie (2008) showed that $\sigma < 1$ is the condition for the hours of hand-to-mouth agents to increase following a decline in their income.

time t consumption and labor supply decision but also on the $t+1$ decisions because the debt repayments at $t+1$ depend on the time t interest rates.

Under mild conditions, borrowers have an incentive to increase their labor supply after an interest rate hike; this is formalized in the following proposition.

Proposition 1 *Under SADL ($\lambda < \frac{1}{1+\nu(1+\bar{D}\kappa)}$) and $\sigma < \frac{1+\bar{D}\kappa}{\gamma}$, a rate hike at time t induces an increases in the borrowers labor supply both at time t and $t+1$.*

The proof is on the next page. At the core of this result, we have that consumption and the labor supply of the borrowers move in opposite directions. This can be appreciated when combining the time $t+1$ optimal response of borrowers in terms of consumption and labor supply after a monetary policy shock which are[30]

$$\hat{H}^B_{t+1} = \frac{\bar{D}(\nu\lambda\gamma\sigma - 1 + \lambda)}{(\nu\lambda - 1 + \lambda)(1+\gamma\nu\sigma)}\epsilon^m_t,$$
$$\hat{c}^B_{t+1} = \frac{\nu\sigma\bar{D}}{(\nu\lambda - 1 + \lambda)(1+\gamma\nu\sigma)}\epsilon^m_t.$$

Combining the latter two equations we have that

$$\hat{c}^B_{t+1} = \frac{\nu\sigma}{\nu\lambda\gamma\sigma - 1 + \lambda}\hat{H}^B_{t+1}.$$

The numerator is positive. Notice that if the conditions of Proposition 1 hold, we have $\lambda < \frac{1}{1+\nu(1+\bar{D}\kappa)} < \frac{1}{1+\nu\sigma\gamma}$, which implies that $(1-\lambda-\nu\sigma\lambda\gamma) > 0$; hence the denominator is negative.

Derivation of Proposition 1 Recall that $\epsilon^m_{t+j} = 0$ for $j > 0$ and $\epsilon^m_t \neq 0$. This implies that from $t+2$ onward the economy is back to steady states and all quantities are zero. This means also that $\hat{R}_{t+j} = \hat{\Pi}_{t+j+1|t+j}$ for $j > 0$, which implies that

$$\hat{c}^S_{t+1} = 0$$

The saver labor supply becomes

$$\nu\hat{H}^S_{t+1} = \hat{w}_{t+1}$$

Using the borrowers BC we have

$$\hat{c}^B_{t+1}\gamma + \bar{D}(\hat{R}_t - \hat{\Pi}_{t+1}) = \left(\hat{w}_{t+1} + \hat{H}^B_{t+1}\right)$$
$$\hat{c}^B_{t+1}\gamma + \bar{D}(\hat{\Pi}_{t+1|t} + \epsilon^m_t - \hat{\Pi}_{t+1}) = (1+\nu)\hat{H}^B_{t+1} + 1/\sigma\hat{c}^B_{t+1}$$
$$1/\sigma\hat{c}^B_{t+1} = \frac{1+\nu}{\gamma\sigma-1}\hat{H}^B_{t+1} - \frac{\bar{D}}{\gamma\sigma-1}\epsilon^m_t$$

[30] The time t optimal responses are analogous but more involved. To ease the notation we only discuss the time $t+1$ decisions.

notice that in absence of shocks in $t+1$ $\hat{\Pi}_{t+1} = \hat{\Pi}_{t+1|t}$. Combining the to labor supply conditions we have

$$
\begin{aligned}
\nu \hat{H}^S_{t+1} + \sigma^{-1}\hat{c}^S_{t+1} &= \nu \hat{H}^B_{t+1} + \sigma^{-1}\hat{c}^B_{t+1} \\
\nu \hat{H}^S_{t+1} &= \nu \hat{H}^B_{t+1} + \frac{1+\nu}{\gamma\sigma - 1}\hat{H}^B_{t+1} - \frac{\bar{D}}{\gamma\sigma - 1}\epsilon^m_t \\
\nu \hat{H}^S_{t+1} &= \frac{1+\gamma\nu\sigma}{\gamma\sigma - 1}\hat{H}^B_{t+1} - \frac{\bar{D}}{\gamma\sigma - 1}\epsilon^m_t
\end{aligned}
$$

Combining the aggregate conditions we have

$$
\begin{aligned}
\lambda\gamma\hat{c}^B_{t+1} + (1-\lambda\gamma)\hat{c}^S_{t+1} =& \lambda\hat{H}^B_{t+1} + (1-\lambda)\hat{H}^S_{t+1} \\
\lambda\gamma\sigma\left[\frac{1+\nu}{\gamma\sigma-1}\hat{H}^B_{t+1} - \frac{\bar{D}}{\gamma\sigma-1}\epsilon^m_t\right] =& \lambda\hat{H}^B_{t+1} + (1-\lambda)\left[\frac{1+\gamma\nu\sigma}{\nu(\gamma\sigma-1)}\hat{H}^B_{t+1} - \frac{\bar{D}}{\nu(\gamma\sigma-1)}\epsilon^m_t\right] \\
\hat{H}^B_{t+1}\left[\frac{(1+\nu)\lambda\gamma\sigma}{\gamma\sigma-1} - \lambda - (1-\lambda)\frac{1+\gamma\nu\sigma}{\nu(\gamma\sigma-1)}\right] =& \left[\frac{\bar{D}\lambda\gamma\sigma}{\gamma\sigma-1} - (1-\lambda)\frac{\bar{D}}{\nu(\gamma\sigma-1)}\right]\epsilon^m_t \\
\hat{H}^B_{t+1}\left[\frac{(1+\nu)\nu\lambda\gamma\sigma - \lambda\nu(\gamma\sigma-1) - (1-\lambda)(1+\gamma\nu\sigma)}{\nu(\gamma\sigma-1)}\right] =& \frac{\nu\lambda\gamma\sigma - (1-\lambda)}{\nu(\gamma\sigma-1)}\bar{D}\epsilon^m_t \\
\hat{H}^B_{t+1}\left[\frac{\nu\lambda(\gamma\sigma+\nu\gamma\sigma-\gamma\sigma+1) - (1-\lambda)(1+\gamma\nu\sigma)}{\nu(\gamma\sigma-1)}\right] =& \frac{\nu\lambda\gamma\sigma - (1-\lambda)}{\nu(\gamma\sigma-1)}\bar{D}\epsilon^m_t \\
\hat{H}^B_{t+1}\frac{(\nu\lambda-1+\lambda)(1+\gamma\nu\sigma)}{\nu(\gamma\sigma-1)} =& \frac{\nu\lambda\gamma\sigma - (1-\lambda)}{\nu(\gamma\sigma-1)}\bar{D}\epsilon^m_t
\end{aligned}
$$

which yield to

$$
\hat{H}^B_{t+1} = \frac{\nu\lambda\gamma\sigma - 1 + \lambda}{(\nu\lambda-1+\lambda)(1+\gamma\nu\sigma)}\bar{D}\ \epsilon^m_t \tag{F.1}
$$

This implies that borrower consumption at time $t+1$ is

$$
\begin{aligned}
\hat{c}^B_{t+1} =& \frac{\sigma(1+\nu)}{\gamma\sigma-1}\hat{H}^B_{t+1} - \frac{\sigma\bar{D}}{\gamma\sigma-1}\epsilon^m_t \\
=& \frac{\sigma(1+\nu)}{\gamma\sigma-1}\frac{\nu\lambda\gamma\sigma-1+\lambda}{(\nu\lambda-1+\lambda)(1+\gamma\nu\sigma)}\bar{D}\ \epsilon^m_t - \frac{\sigma\bar{D}}{\gamma\sigma-1}\epsilon^m_t \\
=& \epsilon^m_t\frac{\sigma\bar{D}}{\gamma\sigma-1}\left[\frac{(1+\nu)(-1+\lambda(1+\nu\gamma\sigma)) - (\nu\lambda-1+\lambda)(1+\gamma\nu\sigma)}{(\nu\lambda-1+\lambda)(1+\gamma\nu\sigma)}\right] \\
=& \epsilon^m_t\frac{\sigma\bar{D}}{\gamma\sigma-1}\left[\frac{-(1+\nu) + (1+\nu)\lambda(1+\nu\gamma\sigma) - \nu\lambda(1+\gamma\nu\sigma) + (1-\lambda)(1+\gamma\nu\sigma)}{(\nu\lambda-1+\lambda)(1+\gamma\nu\sigma)}\right] \\
=& \epsilon^m_t\frac{\sigma\bar{D}}{\gamma\sigma-1}\left[\frac{-1-\nu+1+\gamma\nu\sigma}{(\nu\lambda-1+\lambda)(1+\gamma\nu\sigma)}\right] \\
=& \epsilon^m_t\frac{\nu\sigma\bar{D}}{(\nu\lambda-1+\lambda)(1+\gamma\nu\sigma)}
\end{aligned}
$$

and wages

$$
\begin{aligned}
\hat{w}_{t+1} = & \nu \hat{H}^B_{t+1} + 1/\sigma \hat{c}^B_{t+1} \\
= & \nu \frac{\nu\lambda\gamma\sigma - 1 + \lambda}{(\nu\lambda - 1 + \lambda)(1+\gamma\nu\sigma)} \bar{D}\ \epsilon^m_t + \epsilon^m_t \frac{\nu\bar{D}}{(\nu\lambda - 1 + \lambda)(1+\gamma\nu\sigma)} \\
= & \epsilon^m_t \bar{D}\nu \frac{\nu\lambda\gamma\sigma - 1 + \lambda + 1}{(\nu\lambda - 1 + \lambda)(1+\gamma\nu\sigma)} \\
= & \epsilon^m_t \frac{\bar{D}\nu\lambda}{\nu\lambda - 1 + \lambda}
\end{aligned}
$$

and inflation

$$
\hat{\Pi}_{t+1} = \beta \hat{\Pi}_{t+2|t+1} + \kappa \hat{w}_{t+1} = \epsilon^m_t \frac{\bar{D}\nu\lambda\kappa}{\nu\lambda - 1 + \lambda}
$$

Now, we are in a position to solve for time t. Solving the Euler equation forward we have

$$
\hat{c}^S_t = -\sigma \epsilon^m_t
$$

From the NKP we have an expression for today inflation

$$
\hat{\Pi}_t = \beta \hat{\Pi}_{t+1|t} + \kappa \hat{w}_t = \epsilon^m_t \frac{\bar{D}\nu\lambda\kappa\beta}{\nu\lambda - 1 + \lambda} + \kappa \hat{w}_t
$$

Using the borrowers BC we have

$$
\gamma \hat{c}^B_t + \bar{D}(\hat{R}_{t-1} - \hat{\Pi}_t) = \hat{w}_t + \hat{H}^B_t
$$

$$
\gamma \hat{c}^B_t - \epsilon^m_t \bar{D} \frac{\bar{D}\nu\lambda\kappa\beta}{\nu\lambda - 1 + \lambda} - \bar{D}\kappa\nu \hat{H}^B_t - \bar{D}\kappa/\sigma \hat{c}^B_t = \nu \hat{H}^B_t + 1/\sigma \hat{c}^B_t + \hat{H}^B_t
$$

which yields to

$$
\hat{c}^B_t = \frac{\sigma(1 + \nu(1 + \bar{D}\kappa))}{\gamma\sigma - 1 - \bar{D}\kappa} \hat{H}^B_t + \frac{\bar{D}^2\sigma\nu\lambda\kappa\beta}{e_1 e_0} \epsilon^m_t
$$

where $e_1 = \gamma\sigma - 1 - \bar{D}\kappa$ and $e_0 = \nu\lambda - 1 + \lambda$. Combining the labor supply decision we have

$$
\begin{aligned}
\nu \hat{H}^S_t + \sigma^{-1}\hat{c}^S_t = & \nu \hat{H}^B_t + \sigma^{-1}\hat{c}^B_t \\
\nu \hat{H}^S_t - \epsilon^m_t = & \nu \hat{H}^B_t + \frac{1 + \nu + \nu\bar{D}\kappa}{\gamma\sigma - 1 - \bar{D}\kappa} \hat{H}^B_t + \frac{\bar{D}^2\nu\lambda\kappa\beta}{(\gamma\sigma - 1 - \bar{D}\kappa)(\nu\lambda - 1 + \lambda)} \epsilon^m_t
\end{aligned}
$$

which yields to

$$
\hat{H}^S_t = \frac{1 + \nu\gamma\sigma}{\nu(\gamma\sigma - 1 - \bar{D}\kappa)} \hat{H}^B_t + \frac{\bar{D}^2\nu\lambda\kappa\beta + e_0 e_1}{\nu e_0 e_1} \epsilon^m_t
$$

where $e_1 = \gamma\sigma - 1 - \bar{D}\kappa$ and $e_0 = \nu\lambda - 1 + \lambda$. Combining the aggregate conditions we have

$$
\lambda\gamma \hat{c}^B_{t+1} + (1 - \lambda\gamma)\hat{c}^S_{t+1} = \lambda \hat{H}^B_{t+1} + (1 - \lambda)\hat{H}^S_{t+1}
$$

$$
\begin{aligned}
& \lambda\gamma \left[\frac{\sigma(1 + \nu(1 + \bar{D}\kappa))}{\gamma\sigma - 1 - \bar{D}\kappa} \hat{H}^B_t + \frac{\bar{D}^2\sigma\nu\lambda\kappa\beta}{e_1 e_0} \epsilon^m_t \right] + (1 - \lambda\gamma)[-\sigma\epsilon^m_t] \\
& = \lambda \hat{H}^B_{t+1} + (1 - \lambda) \left[\frac{1 + \nu\gamma\sigma}{\nu(\gamma\sigma - 1 - \bar{D}\kappa)} \hat{H}^B_t + \frac{\bar{D}^2\nu\lambda\kappa\beta + e_0 e_1}{\nu e_0 e_1} \epsilon^m_t \right]
\end{aligned}
$$

$$\hat{H}_t^B \left[\frac{\lambda\gamma\nu\sigma(1+\nu(1+\bar{D}\kappa))}{\nu(\gamma\sigma-1-\bar{D}\kappa)} - \frac{\lambda\nu(\gamma\sigma-1-\bar{D}\kappa)}{\nu(\gamma\sigma-1-\bar{D}\kappa)} - \frac{(1-\lambda)(1+\nu\gamma\sigma)}{\nu(\gamma\sigma-1-\bar{D}\kappa)}\right]$$
$$= \epsilon_t^m \left[\sigma(1-\lambda\gamma) - \gamma\lambda\frac{\bar{D}^2\sigma\nu\lambda\kappa\beta}{e_1e_0} + (1-\lambda)\frac{\bar{D}^2\nu\lambda\kappa\beta + e_0e_1}{\nu e_0 e_1}\right]$$

Focusing on terms inside the left hand side square bracket we have

$$\begin{aligned}
&\nu^{-1}e_1^{-1}\left[\nu\lambda(\sigma\gamma+\sigma\gamma\nu(1+\bar{D}\kappa)) - \lambda\nu(\gamma\sigma-1-\bar{D}\kappa) - (1-\lambda)(1+\nu\gamma\sigma)\right]\\
&= \nu^{-1}e_1^{-1}\left[\nu\lambda(1+\sigma\gamma\nu)(1+\bar{D}\kappa) - (1-\lambda)(1+\nu\gamma\sigma)\right]\\
&= \nu^{-1}e_1^{-1}(1+\sigma\gamma\nu)(\nu\lambda(1+\bar{D}\kappa)-1+\lambda)
\end{aligned}$$

Focusing on terms inside the right hand side square bracket we have

$$\begin{aligned}
&(\nu e_0 e_1)^{-1}\left[\nu\sigma(1-\lambda\gamma)e_0e_1 - \nu\gamma\lambda(\bar{D}^2\sigma\nu\lambda\kappa\beta) + (1-\lambda)(\bar{D}^2\nu\lambda\kappa\beta + e_0e_1)\right]\\
&= (\nu e_0 e_1)^{-1}\left[\nu\sigma(1-\lambda\gamma)e_0e_1 - \nu\gamma\lambda(\bar{D}^2\sigma\nu\lambda\kappa\beta) + (1-\lambda)\bar{D}^2\nu\lambda\kappa\beta + (1-\lambda)e_0e_1\right]\\
&= (\nu e_0 e_1)^{-1}\left[e_0e_1(\nu\sigma(1-\lambda\gamma)+1-\lambda) - \nu\gamma\lambda\sigma\bar{D}^2\nu\lambda\kappa\beta + (1-\lambda)\bar{D}^2\nu\lambda\kappa\beta\right]\\
&= (\nu e_0 e_1)^{-1}\left[\nu\sigma e_0e_1 + e_0e_1(1-\lambda-\nu\sigma\lambda\gamma) + \bar{D}^2\nu\lambda\kappa\beta(1-\lambda-\nu\sigma\lambda\gamma)\right]\\
&= (\nu e_0 e_1)^{-1}\left[\nu\sigma e_0e_1 + (e_0e_1+\bar{D}^2\nu\lambda\kappa\beta)(1-\lambda-\nu\sigma\lambda\gamma)\right]
\end{aligned}$$

Rearranging terms we have

$$\hat{H}_t^B = \frac{\nu\sigma e_0e_1 + (e_0e_1+\bar{D}^2\nu\lambda\kappa\beta)(1-\lambda-\nu\sigma\lambda\gamma)}{(1+\sigma\gamma\nu)(\nu\lambda(1+\bar{D}\kappa)-1+\lambda)(\nu\lambda-1+\lambda)}\ \epsilon_t^m \tag{F.2}$$

where $e_1 = \gamma\sigma-1-\bar{D}\kappa$ and $e_0 = \nu\lambda-1+\lambda$. While still not very tractable we can derive a set of sufficient conditions such that the latter expression becomes positive. These conditions are

1. $\lambda < \frac{1}{1+\nu(1+\bar{D}\kappa)}$
2. $\sigma < \frac{1+\bar{D}\kappa}{\gamma}$

Condition 1 implies that $e_1 = [\nu\lambda(1+\bar{D}\kappa)-1+\lambda] < 0$; this implies also that $(\nu\lambda-1+\lambda) < 0$. Thus, if condition 1 holds, the denominator is positive. Condition 2 implies that $e_0 = (\sigma\gamma-1-\bar{D}\kappa) < 0$; this implies also that $e_0e_1 > 0$. Moreover, if condition 2 hold, it is the case that

$$\lambda < \frac{1}{1+\nu(1+\bar{D}\kappa)} < \frac{1}{1+\nu\sigma\gamma}$$

The latter implies that $(1-\lambda-\nu\sigma\lambda\gamma) > 0$. Therefore also the numerator is positive. These conditions imply also that the coefficient in (F.1) is positive.

G One Asset HANK Appendix

G.1 Steady State

Table G.1 summarizes the mapping between the discrete income states in the model and their corresponding positions in the income distribution. It shows that, in line with the data, both median consumption and labor supply increase monotonically with income. Figure G.1 plots the labor supply policy as a function of assets for each income group. It shows that, for most of the asset distribution and away from the borrowing constraint, more productive households (higher income states) supply more labor at any given asset level. Within a given productivity group, poorer households (with fewer assets) tend to work more. However, as asset holdings approach the borrowing constraint (i.e., become small or negative), among the three income states that include HTM households (20% of the population), the steepness of the labor supply policy function more than offsets productivity differences. In this region, poorer households may supply more labor than more productive ones with similar asset levels. Figure G.2 plots the consumption policy function for each income group. This shows standard results for HANK models where consumption is strictly increasing in assets for all groups, with higher-income households consuming more than lower-income households at any given asset level. Near the borrowing constraint, consumption increases more steeply with assets, consistent with stronger precautionary motives for households close to hand-to-mouth status.

Table G.1: Mapping of income states to percentiles, income levels, median consumption, and median hours worked in the steady state.

State	**Tag**	**Percentile**	**Income**	**Consumption**	**Hours**
1	P0_2	1.6%	0.26	0.58	0.91
2	P2_11	10.9%	0.39	0.69	0.93
3	P11_34	34.4%	0.59	0.84	0.95
4	P34_66	65.6%	0.88	1.00	0.97
5	P66_89	89.1%	1.33	1.20	0.99
6	P89_98	98.4%	2.00	1.42	1.03
7	P98_100	100.0%	3.01	1.67	1.07

Notes: Here hours are not weighted by productivity and by the asset distribution in each income state. Doing so and considering mean hours per income state would result in $H = \sum_e \int h(e, a)\, dD(e, a) = 1$ consistent with the steady state normalization.

Figure G.1: Labor Supply Policy Functions

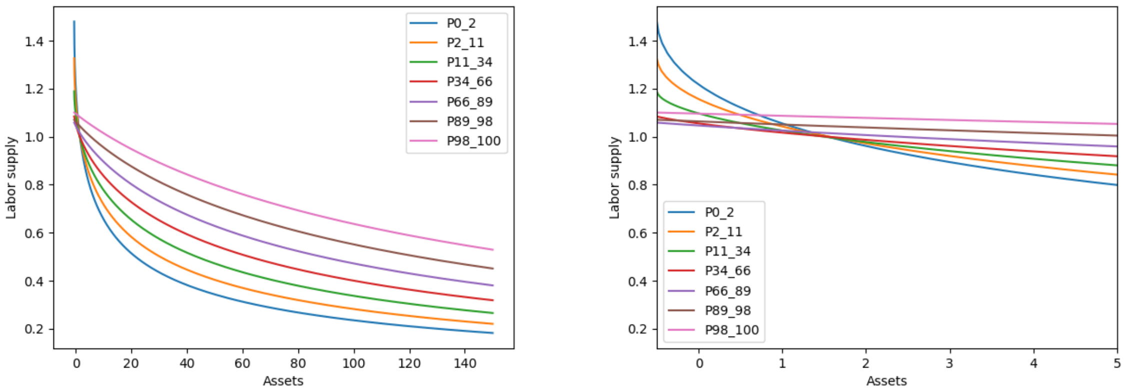

Notes: Note: The left-hand side panel shows it for the full asset grid, while the right-hand side zooms in on the lower tail of the wealth distribution, where borrowing constraints bind.

Figure G.2: Consumption policy function with baseline calibration

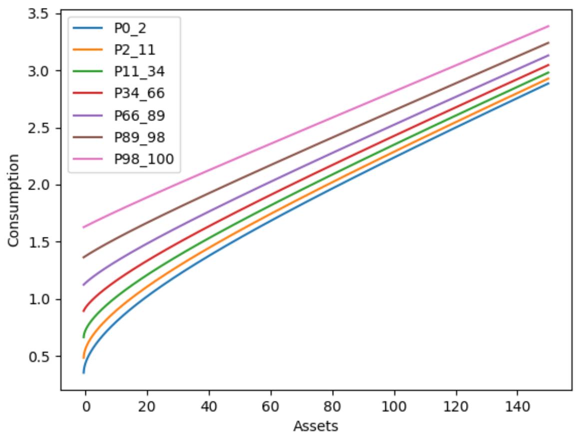

G.2 Policy functions and IRFs across alternative calibrations of the HANK model

Figure G.3: Labor Supply Policy Functions with $\underline{a} = -0.5$ and $\sigma = 1$

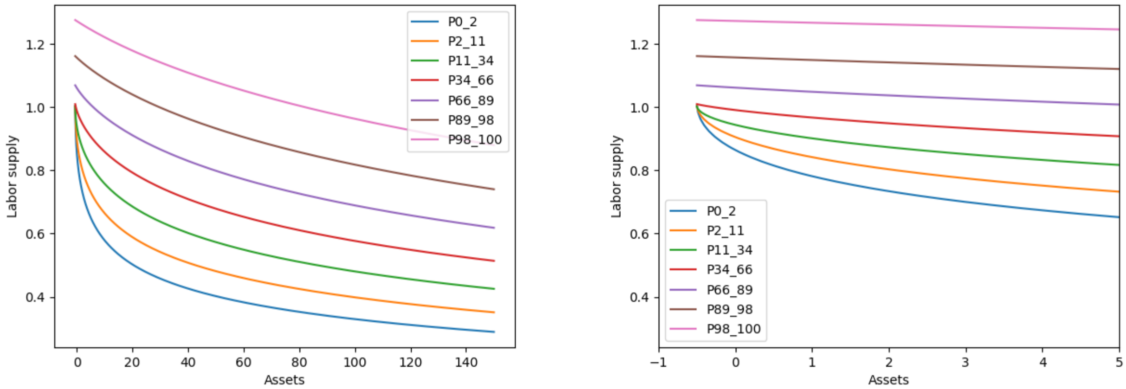

Figure G.4: Consumption policy function with $\underline{a} = -0.5$ and $\sigma = 1$

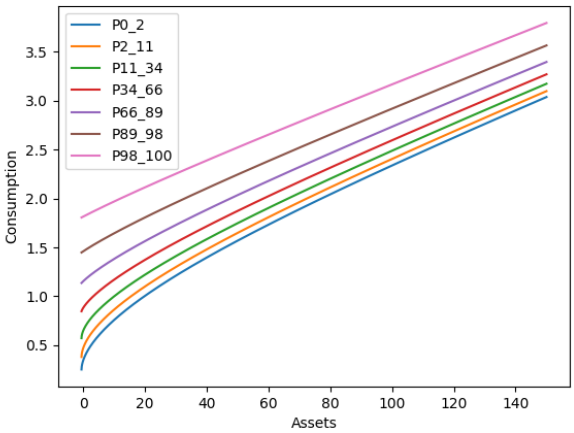

Figure G.5: Impulse responses to a 25 basis points Monetary Policy Shock with $\underline{a} = -0.5$ and $\sigma = 1$

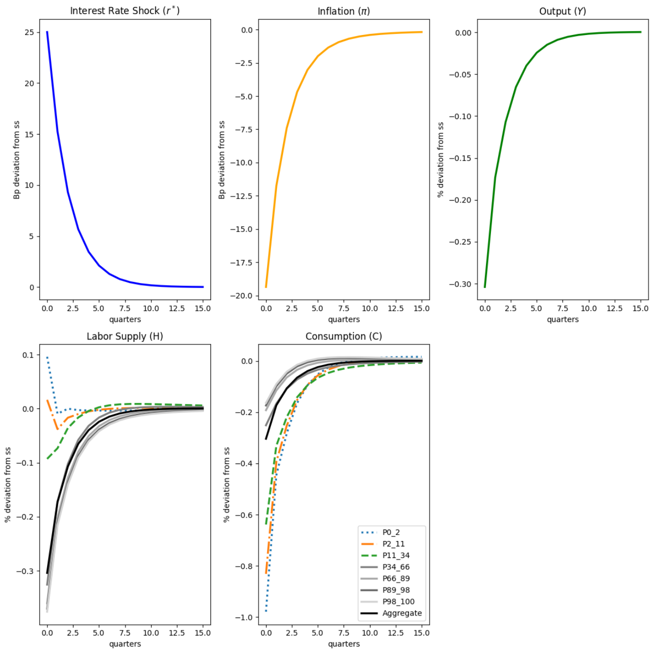

Figure G.6: Labor Supply Policy Functions with $\underline{a} = -0.5$ and $\sigma = 0.25$

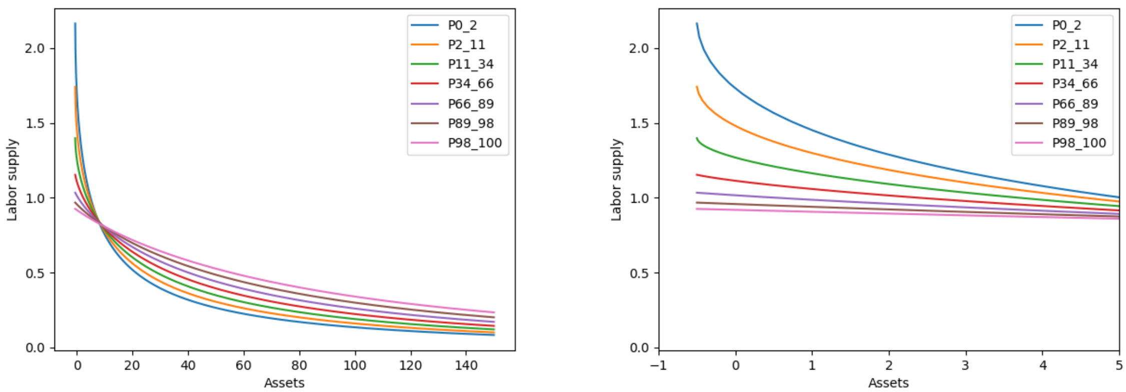

Figure G.7: Consumption policy function with $\underline{a} = -0.5$ and $\sigma = 0.25$

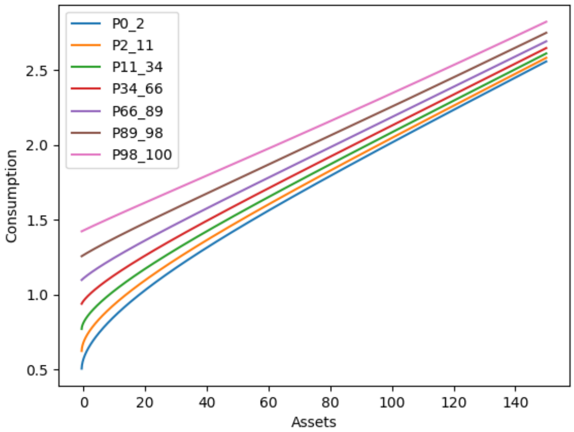

Figure G.8: Impulse responses to a 25 basis points Monetary Policy Shock with $\underline{a} = -0.5$ and $\sigma = 0.25$

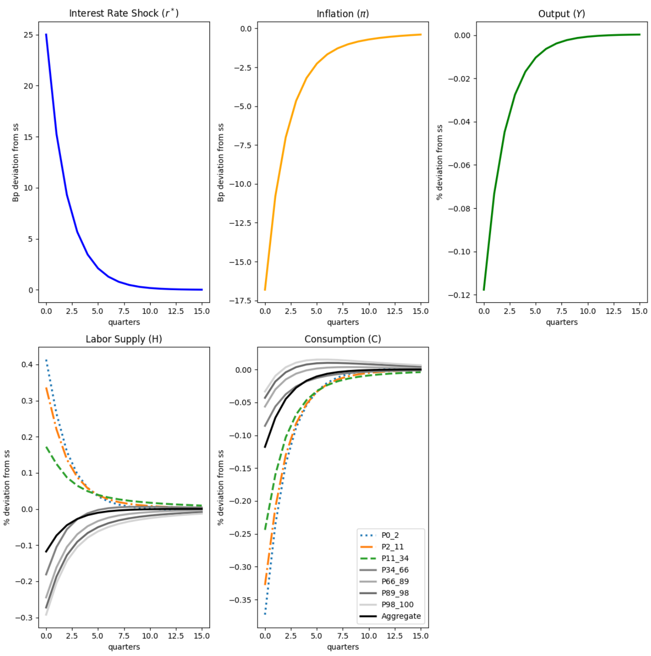

Figure G.9: Labor Supply Policy Functions with $\underline{a} = 0$ and $\sigma = 1$

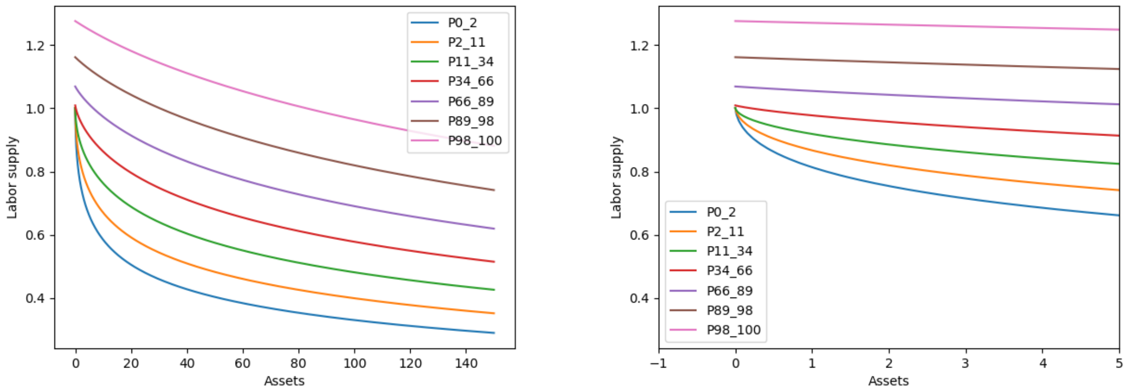

Figure G.10: Consumption policy function with $\underline{a} = 0$ and $\sigma = 1$

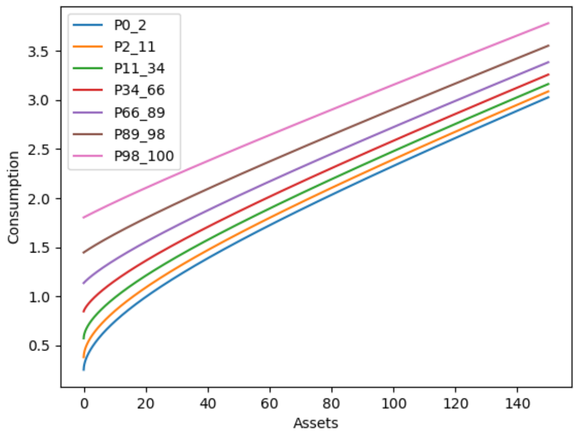

Figure G.11: Impulse responses to a 25 basis points Monetary Policy Shock with $\underline{a} = 0$ and $\sigma = 1$

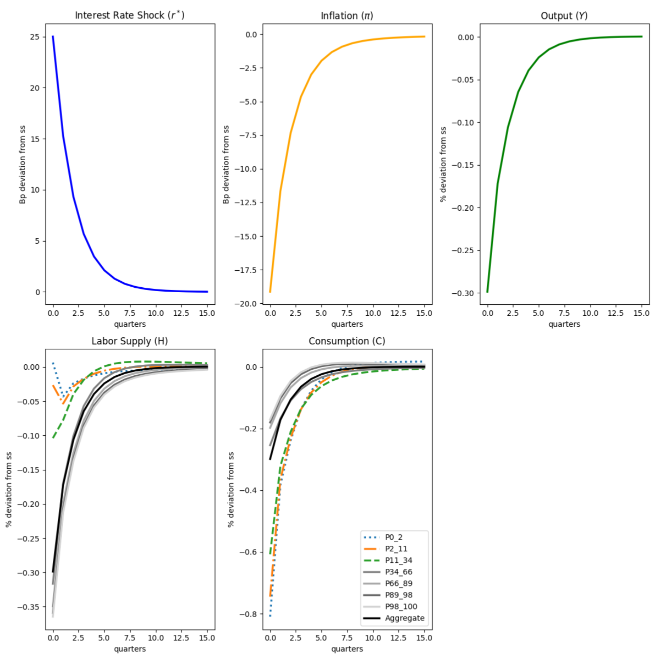

Figure G.12: Labor Supply Policy Functions with $\underline{a} = 0$ and $\sigma = 0.5$

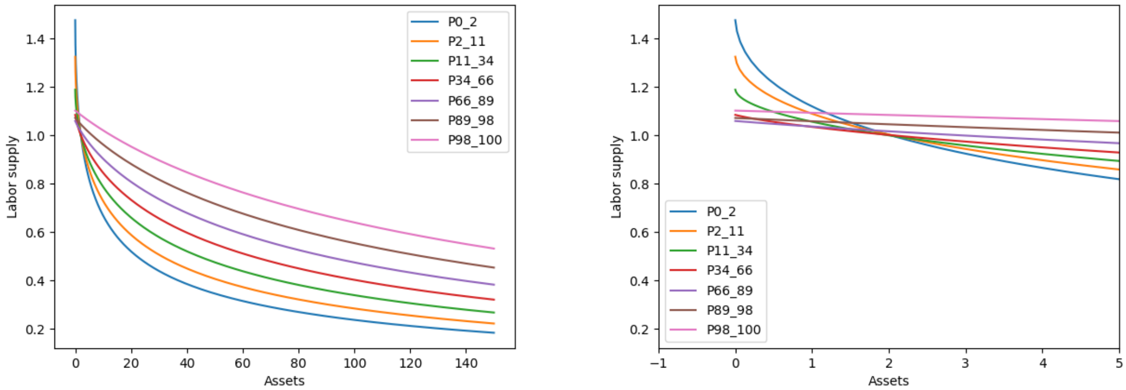

Figure G.13: Consumption policy function with $\underline{a} = 0$ and $\sigma = 0.5$

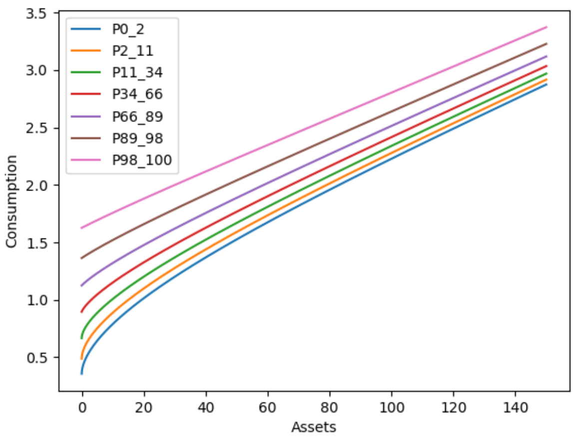

Figure G.14: Impulse responses to a 25 basis points Monetary Policy Shock with $\underline{a} = 0$ and $\sigma = 0.5$

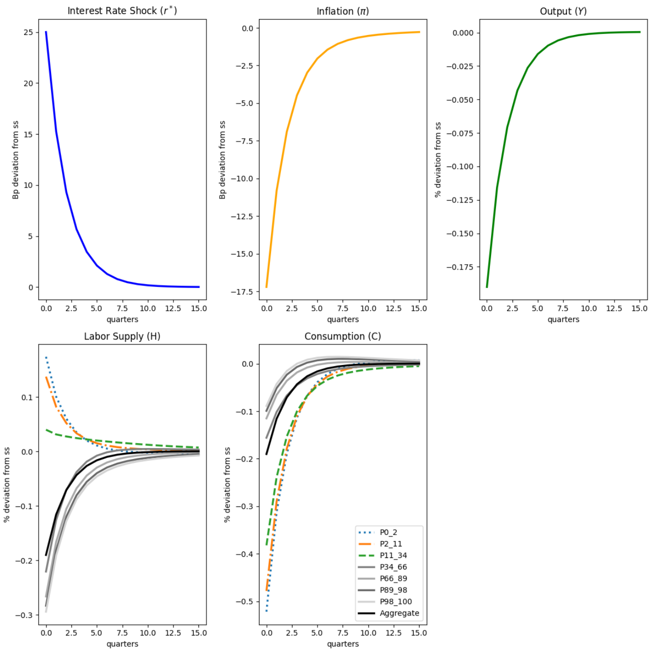

Figure G.15: Labor Supply Policy Functions with $\underline{a} = 0$ and $\sigma = 0.25$

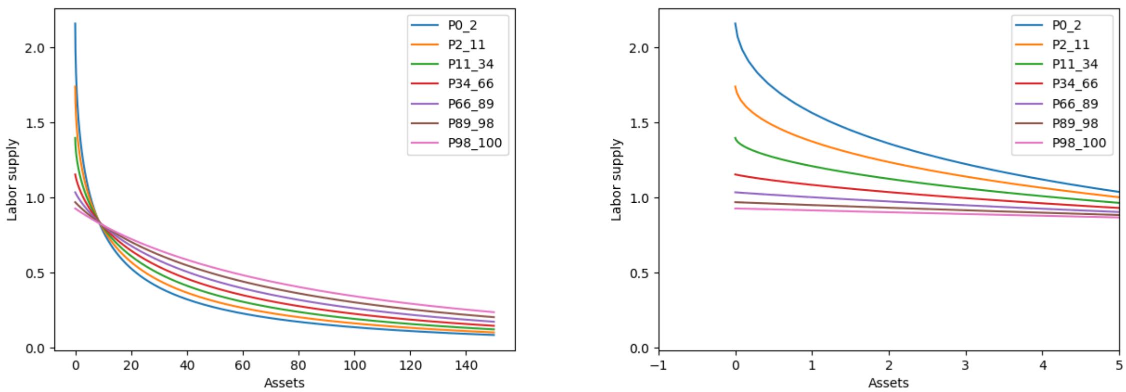

Figure G.16: Consumption policy function with $\underline{a} = 0$ and $\sigma = 0.25$

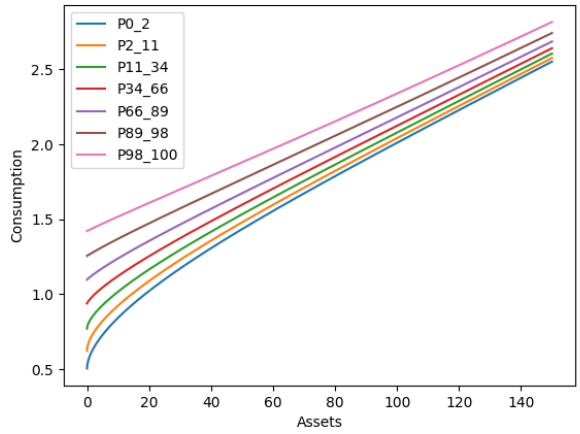

Figure G.17: Impulse responses to a 25 basis points Monetary Policy Shock with $\underline{a} = 0$ and $\sigma = 0.25$

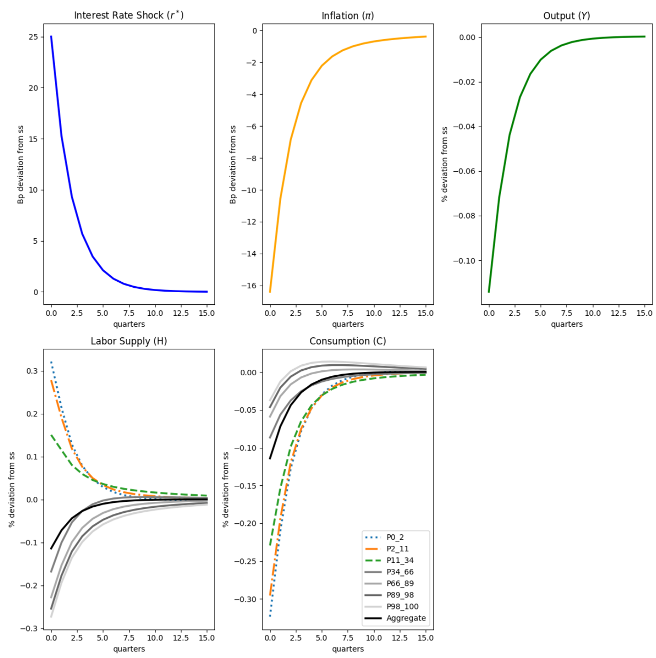

G.3 IRFs decompositions HANK

Figure G.18: Decomposition of consumption Response in HANK.

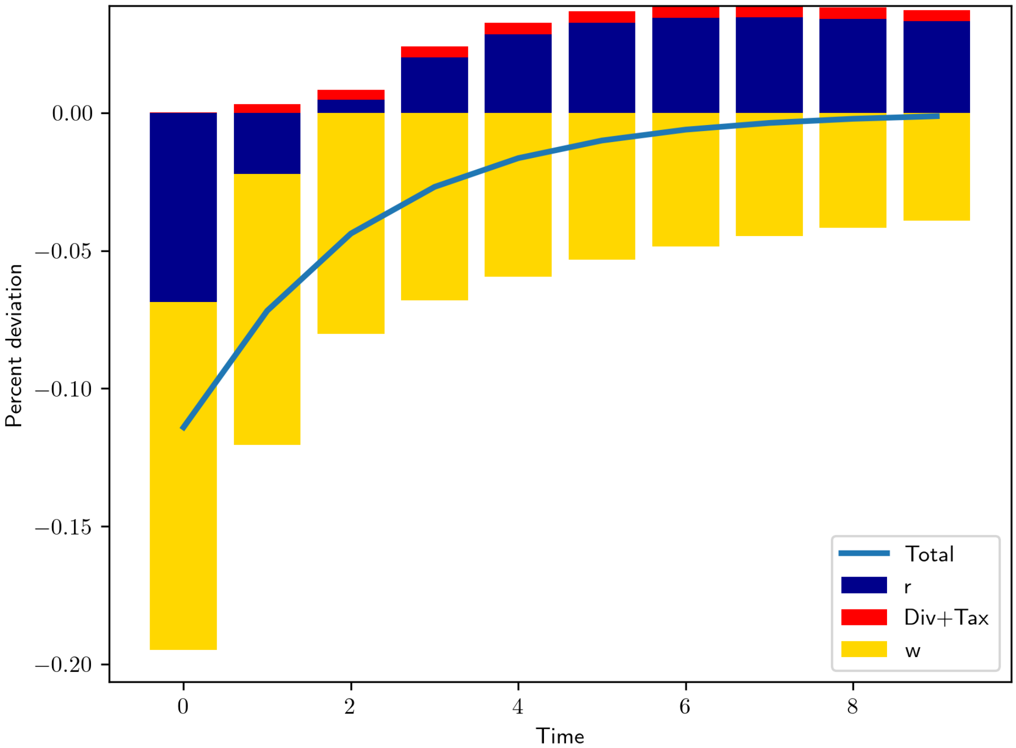

Notes: Note: Percent deviation in aggregate labor hours from steady state, decomposed into marginal channels: interest rate, dividends, taxes, and wages.

Figure G.19: Decomposition of consumption Responses by Income Bin — HANK.

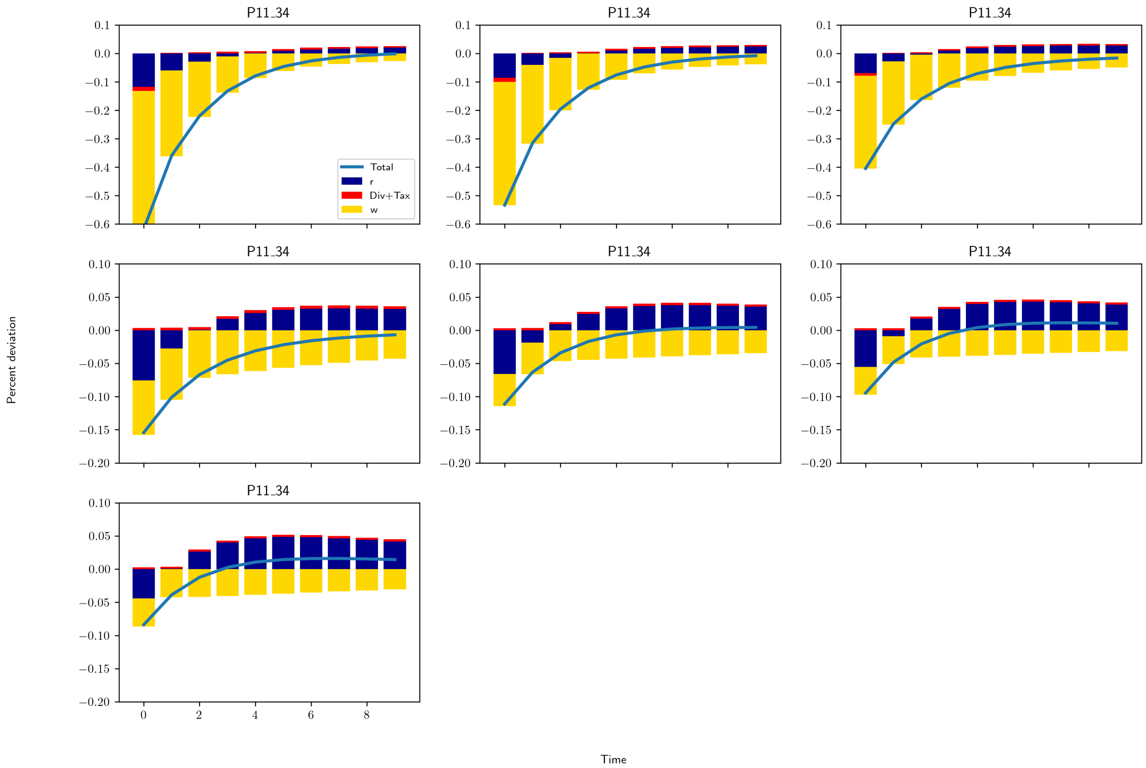

Notes: Notes: Each panel shows percent deviation in labor hours from steady state, decomposed into marginal channels: interest rate, dividends, taxes, and wages.

G.4 HANK vs HANK-HomL

Figure G.20: Impact response of labor supply across income states in HANK vs the aggregate impact response in HANK-HomL

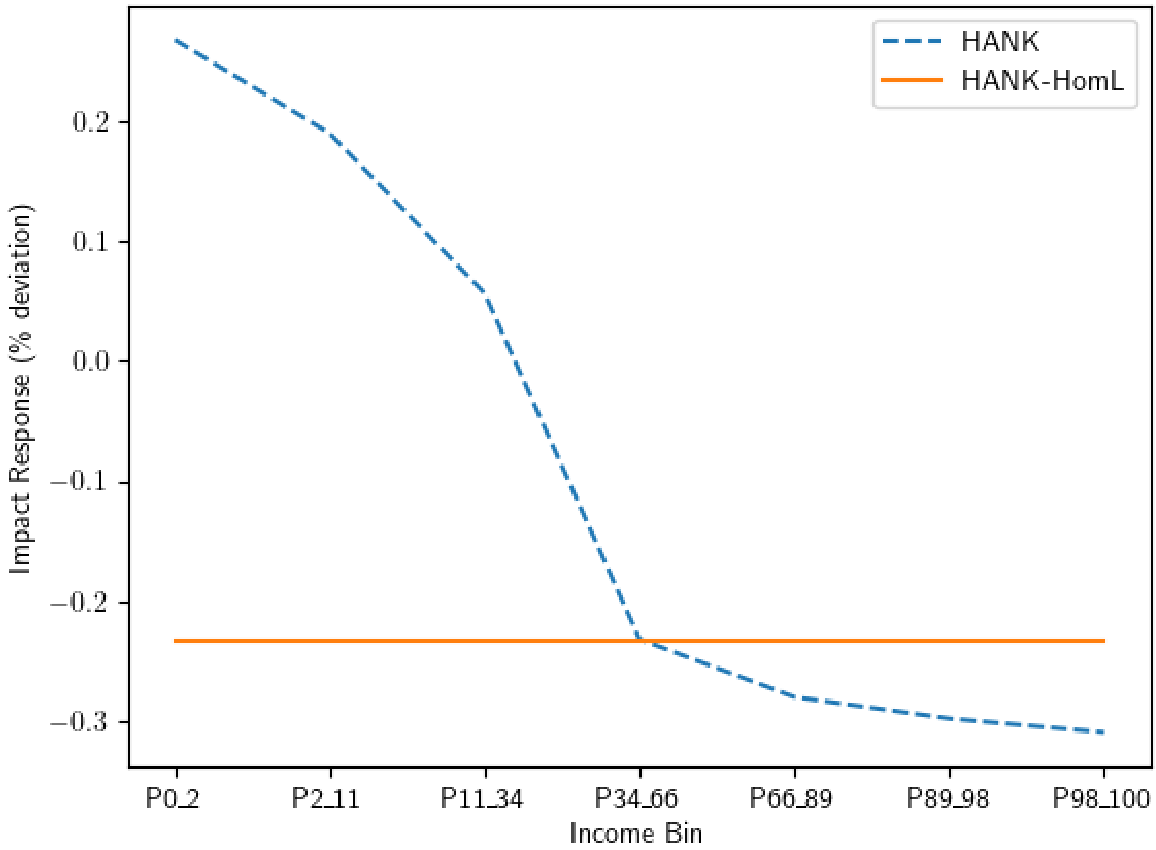

Figure G.21: Direct vs indirect effects of monetary policy on aggregate consumption in HANK vs HANK-HomL for baseline σ.

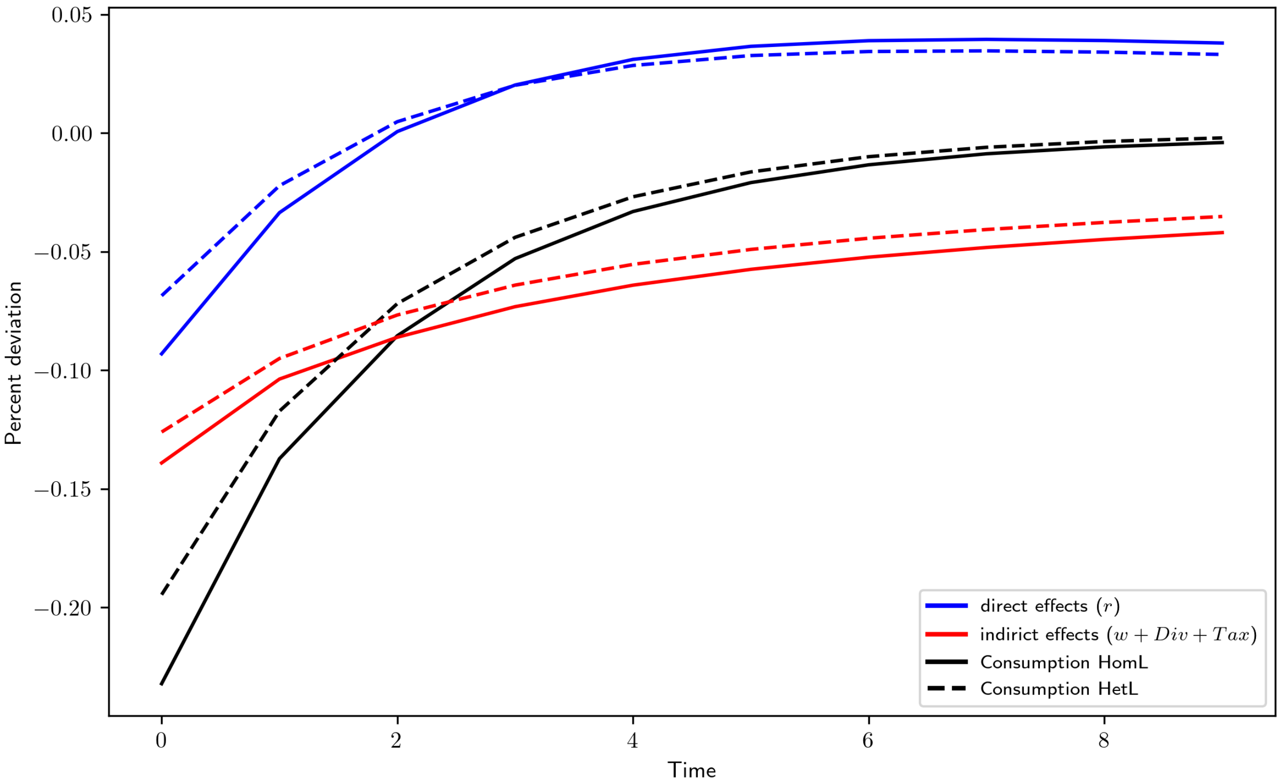

Table 2 in the main text highlights that accounting for heterogeneity in labor supply has quantitatively significant implications, notably lowering the economy's average MPC. This, in turn, affects the transmission of monetary policy. To illustrate this we compare here the dynamic responses to a contractionary monetary policy shock in the HANK and HANK-HomL models. Figure G.22 reports aggregate impulse responses for interest rates, inflation, output, wages, dividends, and the share of hand-to-mouth households. Aggregate price and dividend dynamics are broadly similar across models, as expected given their identical firm and policy structure. However, two notable differences emerge. First, the HTM share rises more and remains elevated longer in HANK, suggesting that flexible labor supply amplifies inequality in downturns.Second, aggregate consumption declines by 16% less on impact in HANK compared to HANK-HomL. This muted response reflects the additional margin of adjustment provided by heterogeneous labor supply, which allows low-income households to buffer shocks through work effort. We explore the distributional responses that drive this result next.

Figure G.23 shows that, by construction, labor supply responses are almost identical for the median agent type (34–66%). At the bottom of the distribution, agents increase their labor supply, consistent with our empirical evidence. In contrast, upper-middle and top-income households reduce their labor supply more in the HANK model compared to HANK-HomL. This highlights that the left tail of the labor supply distribution is responsible for the muted amplification of monetary policy on aggregate demand. Figure G.24 shows how these labor responses translate into individual consumption dynamics. The consumption decline is most severe at the bottom of the distribution in both models, but the drop is smaller and less persistent in the HANK model. At the top of the distribution, consumption responses are very similar across models.

Another way to quantify the role of different households across the income distribution in shaping the aggregate effects of monetary policy is to compute their respective contributions to the overall impact on consumption and labor supply, measured in basis point deviations. Figure G.25 presents the results.

Figure G.22: Impulse Responses of Core Macroeconomic Variables: HANK-HomL vs HANK

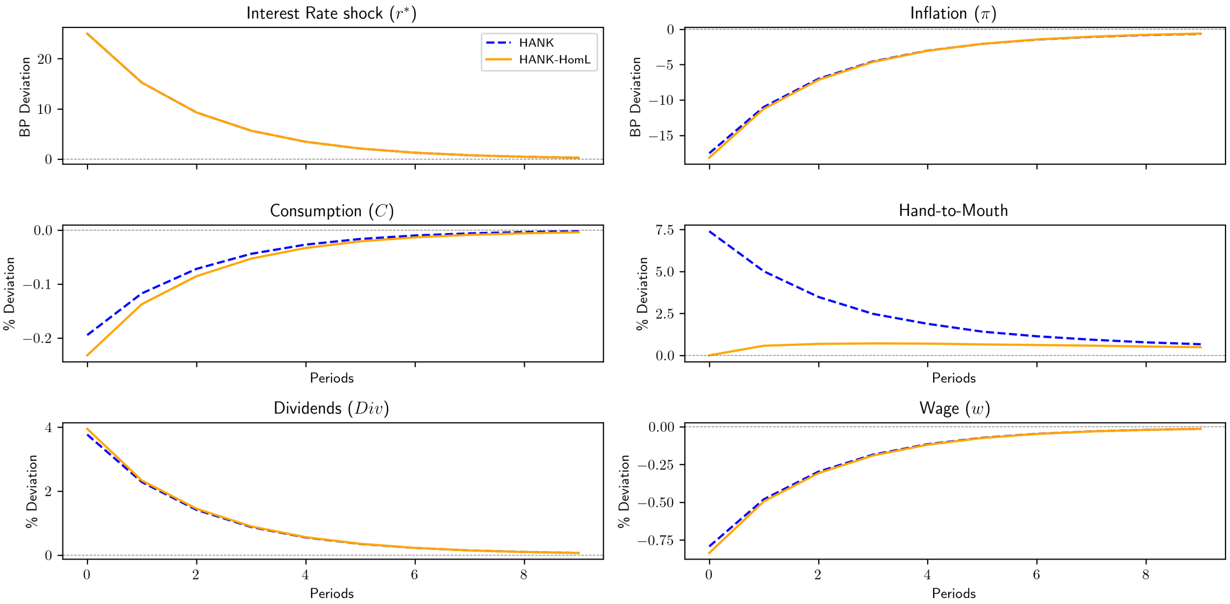

Figure G.23: Labor Supply Responses Across the Income Distribution

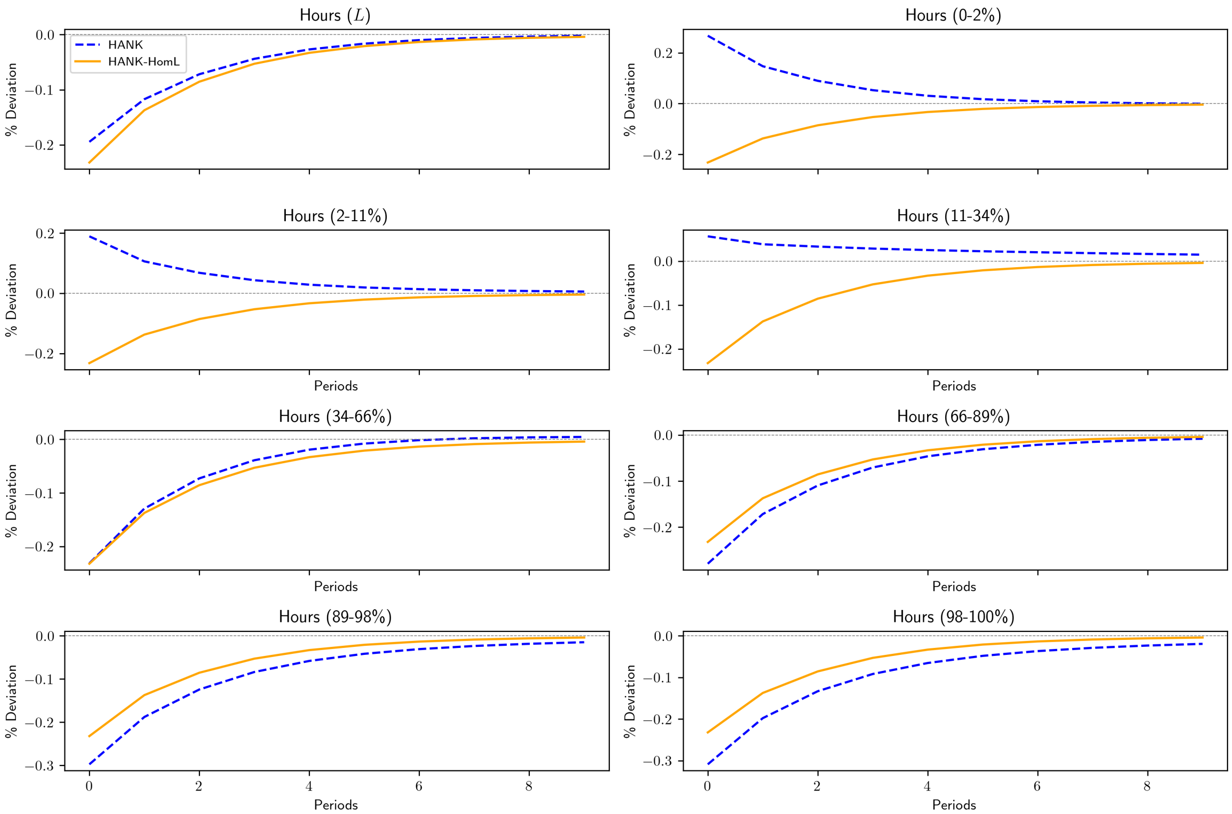

Figure G.24: Consumption Responses Across the Income Distribution

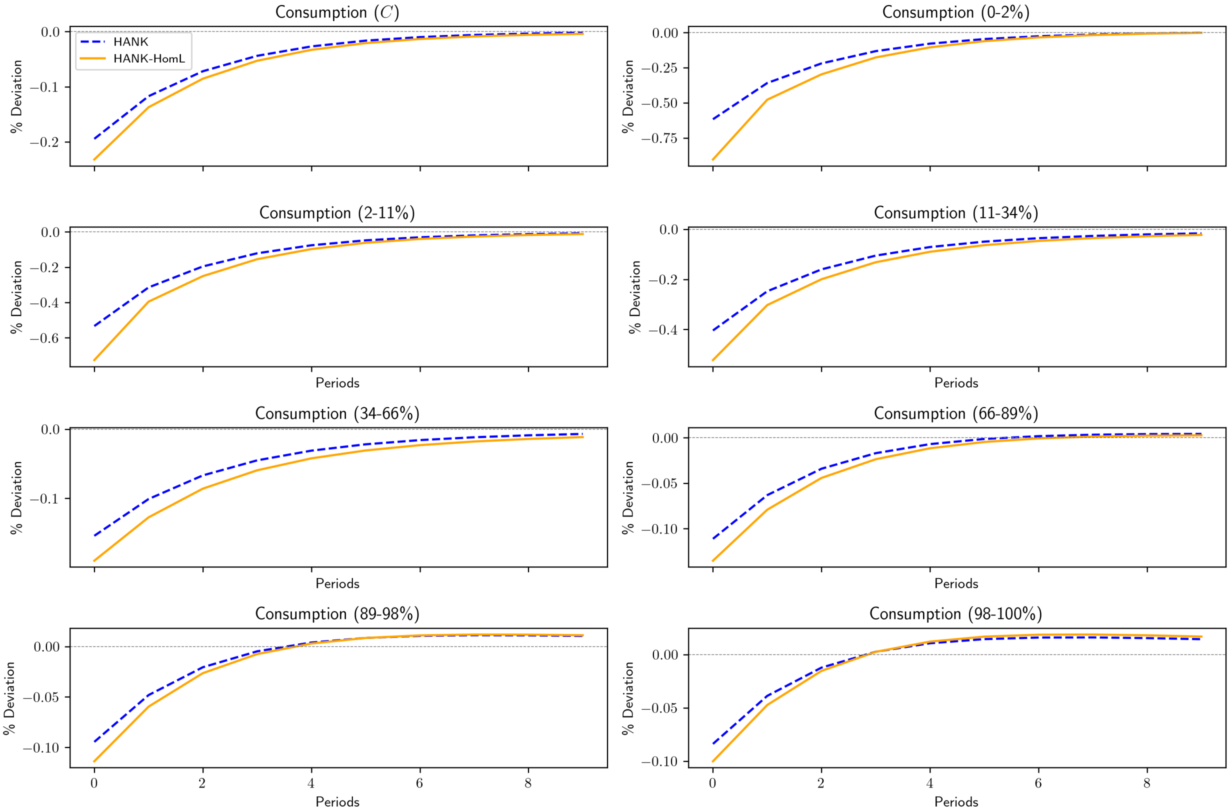

Figure G.25: Contribution of each income bin to the aggregate response.

(b) Contribution of each income bin to the aggregate response.

(a) Hours

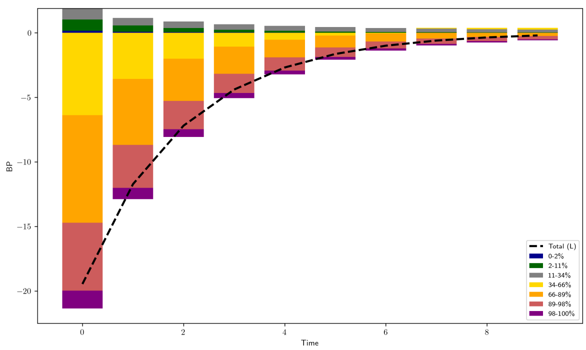

(c) Consumption

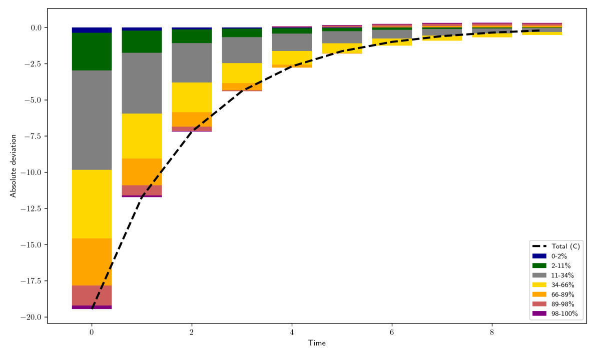

Quest'opera è soggetta alla licenza Creative Commons

CC BY-NC 4.0 DEED

Attribuzione - Non commerciale 4.0 Internazionale

Alma Mater Studiorum - Università di Bologna
DEPARTMENT OF ECONOMICS

Strada Maggiore 45
40125 Bologna - Italy
Tel. +39 051 2092604
Fax +39 051 2092664
http://www.dse.unibo.it