

On the Effects of Monetary Policy Shocks on Earnings and Consumption Heterogeneity

Minsu Chang*

Georgetown University

Frank Schorfheide

University of Pennsylvania,

CEPR, PIER, NBER

This Version: February 12, 2022

Abstract

In this paper we use the functional vector autoregression (VAR) framework of Chang, Chen, and Schorfheide (2021) to study the effects of monetary policy shocks (conventional and informational) on the cross-sectional distribution of earnings and consumption. We find that an expansionary monetary policy shock reduces earnings inequality. The reduction is generated by what we call the employment channel. In the left tail of the earnings distribution, the expansion lifts individuals out of unemployment and thereby reduces the earnings dispersion. For consumption we obtain the opposite result: the expansionary policy shock raises the mass in the right tail of the consumption distribution, thereby increasing consumption inequality. (JEL C11, C32, C52, E32)

Key words: Earnings Distribution, Functional Vector Autoregressions, Heterogeneous Agent Models, Monetary Policy Shocks

* Correspondence: M. Chang: Department of Economics, Georgetown University, Washington, D.C., 20057-1036. Email: minsu.chang@georgetown.edu (Chang). F. Schorfheide: Department of Economics, University of Pennsylvania, Philadelphia, PA 19104-6297. Email: schorf@ssc.upenn.edu (Schorfheide). We thank participants at various seminars and conferences for helpful suggestions. Schorfheide gratefully acknowledges financial support from the National Science Foundation under Grant SES 1851634.

1 Introduction

Traditionally, the effects of monetary policy interventions have been predominantly studied through the lens of models that abstract from micro-level heterogeneity, such as structural vector autoregressions (VARs) specified in terms of macroeconomic aggregates or as representative agent New Keynesian (RANK) models. However, in view of concerns about rising inequalities in advanced economies in the aftermath of the global financial crisis, there is growing interest in the distributional impacts of conventional and unconventional monetary policies.

In this paper we use the functional vector autoregression (VAR) framework of Chang, Chen, and Schorfheide (2021) to study the effects of monetary policy shocks (conventional and informational) on the cross-sectional distribution of earnings and consumption. In addition, we examine whether the inclusion of distributional information changes inference about the effects of monetary policy shocks on aggregate variables. The importance of heterogeneity for the propagation of monetary policy shocks is also a central question in the recent literature on heterogeneous agent New Keynesian (HANK) models. The empirical findings generated by our functional VAR analysis can be used to assess the fit and empirical adequacy of HANK models.

The textbook effect of an expansionary monetary policy shock is a temporary fall in the real interest rate that stimulates demand and increases aggregate output and nominal prices. Kaplan, Moll, and Violante (2018) emphasize that a decomposition into direct and indirect effects is useful for the analysis of the propagation of monetary policy shocks in heterogeneous agent models.¹ The direct effect is generated through the consumption Euler equation: an expansionary monetary policy lowers the real rate and creates an incentive for households to consume in the current period rather than to save for future consumption. Indirect effects are generated through general equilibrium mechanisms that alter the income and wealth distribution. For instance, increased labor demand might raise wages and employment. A rising price level may generate income losses for recipients of nominal government transfers. Moreover, to the extent that debt contracts are nominal, inflation shifts wealth from lenders to borrowers.² The indirect effect on aggregate consumption crucially depends on the households' marginal propensity to consume (MPC).

¹An alternative but closely related decomposition is provided by Auclert (2019).

²This channel is due to Fisher (1933) and has been recently studied in Doepke and Schneider (2006).

To identify a monetary policy shock we use instrumental variables (interest rate and stock price surprises) proposed by Jarocinski and Karadi (2020) and Nakamura and Steinsson (2018). By measuring how a monetary policy shock affects earnings inequality we examine a specific dimension of the indirect channel. An expansionary monetary policy intervention lifts workers out of unemployment and thereby benefits individuals in the left tail of the earnings distribution. We standardize individual earnings by $2/3$ (approximately the labor share) of GDP per capita and decompose the earnings distribution into a continuous part and a pointmass at zero that corresponds to the unemployment rate. Our analysis of the response of the cross-sectional distribution of consumption captures on the one hand the indirect effect of rising earnings onto consumption, but also the direct effect of a real-interest change on household consumption.

The empirical analysis generates the following findings. First, adding the earnings distribution to the VAR does not change inference about the effect of a monetary policy shock on the aggregate variables. Second, an expansionary monetary policy shock reduces earning inequality because the unemployment rate falls and individuals with previously no earnings receive positive earnings (employment channel). The estimated effects are broadly consistent with the HANK model with indivisible labor studied by Ma (2021). If we focus solely on the continuous part of the earnings distribution, then the effect on inequality is small and short-lived. Thus, the employment channel dominates. Third, the expansionary policy shock raises the mass in the right tail of the consumption distribution, thereby increasing consumption inequality. Thus, measures of earnings and consumption inequality move in opposite directions. This observation would be consistent with (i) wealthy households disproportionately benefiting from an increase in unearned income caused by expansionary monetary policy, (ii) a substantial fraction of wealthy “hand-to-mouth” consumers with a high MPC responding strongly to income increases, (iii) and heterogeneity in the direct effect of monetary policy whereby high consumption/income households’ consumption decisions are more interest rate sensitive.

This paper is connected to several strands of the literature. First, there is an empirical literature that studies the effect of policy shocks on measures of inequality using VARs or local projections, e.g., Anderson, Inoue, and Rossi (2016), Coibion, Gorodnichenko, Kueng, and Silvia (2017), Furceri, Loungani, and Zdzienicka (2018), or Guerello (2018). Rather than including the entire distribution, papers in this literature include a low-dimensional vector of distributional statistics, e.g., percentiles or measures of inequality, into the empirical model.

Our functional approach is more comprehensive than existing approaches that rely on

specific statistics, because by modeling the dynamics of the entire distribution, we can compute responses of any distributional statistic from our model. Thus, our method generates more detailed information about distributional dynamics which may be valuable in its own right or it can be used for a more thorough evaluation of structural models. In our framework, there is no need to consider different empirical specifications, e.g., one for percentiles, and another one for the Gini coefficient or related measures of inequality. Moreover, unlike, say, a VAR that includes quantiles of the earnings distribution which may cross in a forward simulation, our model is theoretically coherent, in that the cross-sectional densities are always non-negative and integrate to a finite value (either one or the employment rate).

Coibion, Gorodnichenko, Kueng, and Silvia (2017) study the effects of monetary policy shocks on consumption and income inequality in the US since 1980 using local projections. They consider various measures of economic inequality such as the Gini coefficient, the cross-sectional standard deviation, and the log difference between the 90th and 10th percentiles. Their key finding is that contractionary monetary policy systematically increases inequality in labor earnings, total income, consumption and total expenditures. While our earnings results are qualitatively similar, we obtain the opposite finding for consumption: a rise instead of a fall in inequality. Furceri, Loungani, and Zdzienicka (2018) study the effect of conventional monetary policy shocks on income inequality using a panel of 32 advanced and emerging economies over the period from 1990 to 2013 using local projections. They also find an increase in earnings inequality in response to a contractionary monetary policy shock.

Second, our paper is related to the literature on HANK models which also quantifies the effect of aggregate shocks on cross-sectional distributions. For instance, Kaplan and Violante (2018a) provide an analysis on how a contractionary monetary shock affects the distribution of consumption with the two-asset HANK model. Most closely connected to our empirical analysis is the work by Ma (2021) who calibrates a HANK model with indivisible labor supply based on Chang and Kim (2006). His model emphasizes the employment channel studied in our paper and generates effects that are quantitatively in line with our VAR estimates.

Bayer, Born, and Luetticke (2020) estimate a HANK model using Bayesian techniques. To study the importance of inequality for the business cycle, they conduct the estimation with and without data on inequality. They find that the addition of distributional data does not change what we infer about the aggregate shocks and frictions driving the US business cycle. Their finding is consistent with our results, obtained by comparing IRFs from an

SVAR with only aggregate variables to an SVAR that includes cross-sectional information in addition.

Part of the HANK literature focuses on comparing the differences in the propagation of monetary policy shocks in heterogeneous versus representative agent environments, often emphasizing amplification or dampening mechanisms, e.g., Acharya, Chen, Del Negro, Dogra, Matlin, and Safarti (2021), Auclert (2019), Kaplan and Violante (2018b), Kaplan, Moll, and Violante (2018). Our structural functional VAR analysis does not speak to this issue, because it is difficult to create a compelling RANK counterfactual in a VAR that has been estimated on data from an economy with heterogeneous households and firms. What we can do in our framework is to examine whether cross-sectional information is helpful in predicting aggregate outcomes (Granger causality).

Third, there is an extensive literature on the statistical analysis of functional data. General treatments are provided in the books by Bosq (2000), Ramsey and Silverman (2005), and Horvath and Kokoszka (2012). The connection between this literature and the empirical methodology used in our paper is discussed in more detail in Chang, Chen, and Schorfheide (2021).

The remainder of this paper is organized as follows. Section 2 describes the methodological framework. Section 3 discusses the VAR identification of the monetary policy shock. The data sets used in the empirical analysis are described in Section 4. The results from the functional VAR with cross-sectional earnings data are presented in Section 5. The analysis with cross-sectional consumption data is conducted in Section 6. Finally, Section 7 concludes. An Online Appendix contains supplemental information on the methodology and additional results from the empirical analysis.

2 A Functional VAR

The econometric framework is based on Chang, Chen, and Schorfheide (2021). To make this paper self-contained, we provide a short summary.

The variables in the functional model comprise an $n_y \times 1$ vector of macroeconomic aggregates Y_t and a cross-sectional density $p_t(x)$. Throughout this paper, we will work with log densities defined as $\ell_t(x) = \ln p_t(x)$. We decompose Y_t and ℓ_t into a deterministic component $(Y_*, \ell_*(x))$ and fluctuations around the deterministic component. Let

$$Y_t = Y_* + \tilde{Y}_t, \quad \ell_t = \ell_* + \tilde{\ell}_t. \quad (1)$$

For notational convenience we assumed that the deterministic component is time-invariant and could be interpreted as a steady state. This assumption could be easily relaxed by letting (Y_*, ℓ_*) depend on t . We assume that the deviations from the deterministic component $(Y_t, \ell_t(x))$ evolve jointly according to the following linear functional vector autoregressive (fVAR) law of motion:

$$\begin{aligned}\tilde{Y}_t &= B_{yy}\tilde{Y}_{t-1} + \int B_{yl}(\tilde{x})\tilde{\ell}_{t-1}(\tilde{x})d\tilde{x} + u_{y,t} \\ \tilde{\ell}_t(x) &= B_{ly}(x)\tilde{Y}_{t-1} + \int B_{ll}(x, \tilde{x})\tilde{\ell}_{t-1}(\tilde{x})d\tilde{x} + u_{l,t}(x).\end{aligned}\tag{2}$$

Here $u_{y,t}$ is mean-zero random vector with covariance Ω_{yy} and $u_{l,t}(x)$ is a random element in a Hilbert space with covariance function $\Omega_{ll}(x, \tilde{x})$. We denote the covariance function for $u_{y,t}$ and $u_{l,t}(x)$ by $\Omega_{yl}(x)$. For now, (2) should be interpreted as a reduced-form fVAR in which $u_{y,t}$ and $u_{l,t}(x)$ are one-step-ahead forecast errors. One can easily add more lags to the system. (2) will subsequently serve as the state-transition equation in a functional state-space model.

2.1 Sampling and Measurement

We assume that in every period $t = 1, \dots, T$ an econometrician observes the macroeconomic aggregates Y_t as well as a sample of N_t draws x_{it} , $i = 1, \dots, N_t$ from the cross-sectional density $p_t(x)$. In practice, N_t is likely to vary from period to period, but for the subsequent exposition it will be more convenient to assume that $N_t = N$ for all t . We collect the time t cross-sectional observations in the vector $X_t = [x_{1t}, \dots, x_{Nt}]'$. We also assume that the draws x_{it} are independently and identically distributed (iid) over the cross-section and independent over time. The measurement equation for the cross-section observations takes the form

$$x_{it} \sim \text{iid } p_t(x) = \frac{\exp\{\ell_t(x)\}}{\int \exp\{\ell_t(x)\}dx}, \quad i = 1, \dots, N, \quad t = 1, \dots, T.\tag{3}$$

The assumption of x_{it} being *iid* across i and t is consistent with data sets that comprise repeated cross sections.³ The functional modeling approach does not require the econometrician to make assumptions about the evolution of x_{it} at the level of an individual, a household, or a firm.

³If the data exhibit spatial correlation, then our estimation approach below essentially replaces the likelihood function for x_{1t}, \dots, x_{Nt} by a composite likelihood function that ignores the spatial correlation; see Varin, Reid, and Firth (2011).

2.2 A Collection of Finite-Dimensional Models

Equations (1), (2) and (3) define a state-space model for the observables $\{Y_t, X_t\}_{t=1}^T$. The log density $\ell_t(\cdot)$ is the state variable. To implement the estimation of the functional model we use a collection of finite-dimensional representations, indexed by a superscript (K). Let

$$\ell_t^{(K)}(x) = \sum_{k=1}^K \alpha_{k,t} \zeta_k(x) = [\zeta_1(x), \dots, \zeta_K(x)] \cdot \begin{bmatrix} \alpha_{1,t} \\ \vdots \\ \alpha_{K,t} \end{bmatrix} = \zeta'(x) \alpha_t \quad (4)$$

and

$$\ell_*^{(K)}(x) = \zeta'(x) \alpha_*.$$

To simplify the notation a bit, we do not use (K) superscripts for the vectors $\zeta(x)$, α_t , and α_* . Here $\zeta_1(x), \zeta_2(x), \dots$ is a sequence of basis functions. We define $\tilde{\alpha}_t = \alpha_t - \alpha_*$ such that $\tilde{\ell}^{(K)}(x) = \ell_t^{(K)}(x) - \ell_*^{(K)}(x)$. For theoretical considerations it is convenient to demean the vector of basis functions and assume that $\int \zeta(x) dx = 0$. For applications this normalization is not important.

To construct the measurement equation of the cross-sectional observations in (3), we define the K -dimensional vector of sufficient statistics

$$\bar{\zeta}(X_t) = \frac{1}{N} \sum_{i=1}^N \zeta(x_{it}).$$

This allows us to write a K 'th order representation of the density of X_t :

$$\begin{aligned} p^{(K)}(X_t | \alpha_t) &= \exp \{ N \mathcal{L}^{(K)}(\alpha_t | X_t) \}, \\ \mathcal{L}^{(K)}(\alpha_t | X_t) &= \bar{\zeta}'(X_t) \alpha_t - \ln \int \exp \{ \zeta'(x) \alpha_t \} dx. \end{aligned} \quad (5)$$

We represent the kernels $B_{ll}(x, \tilde{x})$ and $B_{yl}(\tilde{x})$, the function $B_{ly}(x)$, and the functional innovation $u_{l,t}(x)$ that appear in the state-transition equation (2) as follows:

$$\begin{aligned} B_{ll}^{(K)}(x, \tilde{x}) &= \zeta'(x) B_{ll} \xi(\tilde{x}), & B_{yl}^{(K)}(x) &= B_{yl} \xi(\tilde{x}) \\ B_{ly}^{(K)}(x) &= \zeta(x)' B_{ly}, & u_{l,t}^{(K)}(x) &= \zeta'(x) u_{\alpha,t}, \end{aligned} \quad (6)$$

where $\xi(x)$ is a second $K \times 1$ vector of basis functions and $u_{\alpha,t}$ is a $K \times 1$ vector of innovations. The matrix B_{ll} is of dimension $K \times K$, B_{yl} is of dimension $n_y \times K$, and B_{ly} is of dimension

$K \times n_y$. Combining (1), (2), and (6) yields the following vector autoregressive system for the macroeconomic aggregates and the sieve coefficients (omitting K superscripts):

$$\begin{bmatrix} Y_t - Y_* \\ \alpha_t - \alpha_* \end{bmatrix} = \begin{bmatrix} B_{yy} & B_{yl}C_\alpha \\ B_{ly} & B_{ll}C_\alpha \end{bmatrix} \begin{bmatrix} Y_{t-1} - Y_* \\ \alpha_{t-1} - \alpha_* \end{bmatrix} + \begin{bmatrix} u_{y,t} \\ u_{\alpha,t} \end{bmatrix}, \quad (7)$$

where $C_\alpha = \int \xi(\tilde{x})\zeta'(\tilde{x})d\tilde{x}$. Let $u'_t = [u'_{y,t}, u'_{\alpha,t}]$. We will subsequently assume that the innovations are Gaussian:

$$u_t \sim N(0, \Sigma). \quad (8)$$

2.3 Further Implementation Details

As explained in detail in Chang, Chen, and Schorfheide (2021), in practice a few additional steps are required for the implementation. First, if the cross-sectional data are top-coded, the likelihood function in (5) needs to be adjusted accordingly.

Second, there might be linear dependencies in the α_t vector. Thus, after subtracting α_* (or $\alpha_{*,t}$ in case of the seasonal adjustment), we compress the α_t and $\hat{\alpha}_t$ into lower-dimensional vectors a_t and \hat{a}_t . Even if there are no exact linear dependencies, we normalize the log density coefficients to have unit variance. Using a_t instead of α_t and absorbing the matrix C_α into the matrices of regression coefficients, we obtain the state-transition equations

$$\begin{bmatrix} Y_t - Y_* \\ a_t \end{bmatrix} = \begin{bmatrix} \Phi_{yy} & \Phi_{ya} \\ \Phi_{ay} & \Phi_{aa} \end{bmatrix} \begin{bmatrix} Y_{t-1} - Y_* \\ a_{t-1} \end{bmatrix} + \begin{bmatrix} u_{y,t} \\ u_{a,t} \end{bmatrix}. \quad (9)$$

Under the assumption that the innovations are normally distributed (9) can be written more compactly as

$$W_t = \Phi_1 W_{t-1} + u_t, \quad u_t \sim N(0, \Sigma), \quad (10)$$

where $W_t = [(Y_t - Y_*)', a_t']'$.

Third, we use a quadratic approximation to $\mathcal{L}^{(K)}(\alpha_t|X_t)$ in (3) which leads to a measurement equation of the form

$$\hat{a}_t = a_t + N^{-1/2}\eta_t, \quad \eta_t \sim N(0, (\hat{\Lambda}\hat{V}_t^{-1}\hat{\Lambda}')^{-1}). \quad (11)$$

Because N is large in our application, the measurement error $N^{-1/2}\eta_t$ is close to zero and for the empirical analysis we simply replace a_t in (9) and the subsequent definition of W_t by \hat{a}_t .

3 Identification and Bayesian Estimation

3.1 Structural Shocks

We include external instruments for the structural shocks of interest in the definition of W_t , assuming that the instruments are ordered first. Formally, we partition $W_t = [W'_{1,t}, W'_{2,t}]'$, where the $n_1 \times 1$ vector $W_{1,t}$ contains the instruments. In our application n_1 equals either one or two. Let n_2 be the number of elements of $W_{2,t}$. We write the structural form of the VAR in (10) as – omitting the 1 subscript from Φ and using subscripts to indicate the obvious partitions

$$W_{1,t} = \Phi_1.W_{t-1} + \Phi_{11}^\epsilon \epsilon_{1,t} \quad (12)$$

$$W_{2,t} = \Phi_2.W_{t-1} + \Phi_{22}^{\tilde{\epsilon}} \tilde{\epsilon}_{2,t}. \quad (13)$$

Here, the instrument innovations $\epsilon_{1,t}$ are also of dimension n_1 and they are standardized such that $\mathbb{E}[\epsilon_{1,t}\epsilon'_{1,t}] = I$. We impose two exclusion restrictions: the vector of structural shocks $\tilde{\epsilon}_{2,t}$ does not affect the instruments $W_{1,t}$ contemporaneously, and $W_{2,t}$ does not directly depend on the instrument innovations $\epsilon_{1,t}$. Importantly, however, we allow for correlation between $\epsilon_{1,t}$ and $\tilde{\epsilon}_{2,t}$. We further partition the structural shock vector $\tilde{\epsilon}_{2,t} = [\tilde{\epsilon}'_{2.1,t}, \tilde{\epsilon}'_{2.2,t}]'$, where $\tilde{\epsilon}_{2.1,t}$ is the $n_1 \times 1$ subvector of for which we are computing impulse response functions. Conformingly, we also partition $\Phi_{22}^{\tilde{\epsilon}} = [\Phi_{22.1}^{\tilde{\epsilon}}, \Phi_{22.2}^{\tilde{\epsilon}}]$. Our goal is to identify $\Phi_{22.1}^{\tilde{\epsilon}}$ which determines the impact of $\tilde{\epsilon}_{2.1,t}$ the structural shocks of interest.

We now make the following additional assumptions: first, the structural shocks of interest $\tilde{\epsilon}_{2.1,t}$ are correlated with the instrument innovations $\epsilon_{1,t}$:

$$\tilde{\epsilon}_{2.1,t} = \Gamma_1 \epsilon_{1,t} + \Gamma_{2.1} \epsilon_{2.1,t}, \quad (14)$$

where Γ_1 is diagonal and the vector $[\epsilon'_{1,t}, \epsilon'_{2.1,t}]'$ has an identity covariance matrix. Second, $\Gamma_1 \neq 0$ (instrument relevance) and $\epsilon_{1,t}$ is uncorrelated with $\tilde{\epsilon}_{2.2,t}$ (instrument validity). Third, $\epsilon_{2.1,t}$ is also uncorrelated with $\tilde{\epsilon}_{2.2,t}$ which ensures that the structural shocks $\tilde{\epsilon}_{2,t}$ have a diagonal covariance matrix. Using (14), the previously defined vector and matrix partitions we can rewrite (13) as follows:

$$W_{2,t} = \Phi_2.W_{t-1} + (\Phi_{22.1}^{\tilde{\epsilon}} \Gamma_1) \epsilon_{1,t} + \Phi_{22.1}^\epsilon \epsilon_{2.1,t} + \Phi_{22.2}^\epsilon \epsilon_{2.2,t}, \quad (15)$$

where $\Phi_{22.1}^\epsilon = \Phi_{22.1}^{\tilde{\epsilon}} \Gamma_{2.1}$, $\Phi_{22.2}^\epsilon = \Phi_{22.2}^{\tilde{\epsilon}} (\mathbb{E}[\tilde{\epsilon}_{2.2,t} \tilde{\epsilon}'_{2.2,t}])^{1/2}$ and $\epsilon_{2.2,t} = (\mathbb{E}[\tilde{\epsilon}_{2.2,t} \tilde{\epsilon}'_{2.2,t}])^{-1/2} \tilde{\epsilon}_{2.2,t}$. Thus, $\Phi_{22.1}^{\tilde{\epsilon}}$ can be obtained up to scale Γ_1 from the contemporaneous effect of $\epsilon_{1,t}$ on $W_{2,t}$.

By combining (12) with (15) we obtain a block triangular system. Let Σ_{tr} be the lower-triangular Cholesky factor of Σ with partitions conforming to the partitions of W_t . Write $u_t = \Phi^\epsilon \epsilon_t$ where $\Phi^\epsilon = [(\Phi_{22,1}^\epsilon \Gamma_1), \Phi_{22,1}^\epsilon, \Phi_{22,2}^\epsilon]$ and $\epsilon_t = [\epsilon'_{1,t}, \epsilon'_{2,1,t}, \epsilon'_{2,2,t}]'$. Now factorize $\Phi^\epsilon = \Sigma_{tr} \Omega$, where Ω is an orthonormal matrix. The block diagonal structure of the system implies that $\Omega_{12} = 0$ and $\Omega_{21} = 0$ and we can deduce that $\Phi_{22,1}^\epsilon \Gamma_1 = \Sigma_{21,tr} \Omega_{11}$. Thus, conditional on choosing an Ω_{11} we can identify the contemporaneous effects of the structural shocks on $W_{2,t}$ up to a scale factor from the responses to the instrument innovations. This implementation of VAR shock identification through instrumental variables has been used, for instance, in Anderson, Inoue, and Rossi (2016) for fiscal policy shocks and Jarocinski and Karadi (2020) for monetary policy shocks.⁴

3.2 Bayesian Estimation

We estimate the model by setting the measurement error η_t in (11) equal to zero. Thus, W_t is treated as fully observed. We write the VAR in matrix form as

$$W = Z\Phi + U, \quad (16)$$

where the matrix W has rows W'_t , U has rows u'_t , Z has rows Z'_t , and $Z'_t = [Z'_{t-1}, \dots, Z'_{t-p}, 1]$. While we previously considered a specification with only one lag, in general one can include p lags and an intercept. The coefficient matrix $\Phi = [\Phi'_1, \dots, \Phi'_p, \Phi_c]$. We use the same prior as in Jarocinski and Karadi (2020), which is a version of the Minnesota prior. The prior imposes independence of Σ and $\phi = \text{vec}(\Phi)$ and takes the form

$$\Sigma \sim IW(\underline{\nu}, \underline{S}), \quad \phi | \lambda \sim N(\underline{\mu}_\phi, \underline{P}_\phi^{-1}(\lambda)), \quad (17)$$

where λ is a vector of hyperparameters. To tune the prior distribution, we first estimate univariate AR(1) models for the $W_{j,t}$ series and let \hat{s}_j denote the estimate of the innovation standard deviation. We set $\nu = 2 + n_w$ and \bar{S} is a diagonal matrix with entries \hat{s}_j^2 . The elements of the prior mean vector are mostly zero, except that for some variables the prior mean for the coefficient on the first own lag is one, implying a univariate random walk representation. The precision matrix $\underline{P}_\phi(\lambda)$ is diagonal. It implies that the precision for the coefficient on lag s of variable j in equation i is $(1/\lambda_1)(\sigma_j/\sigma_i)^2 s^{\lambda_2}$. We set $\lambda_1 = 5$ and $\lambda_2 = 1$. A Gibbs sampler is used to generate draws from the posterior distributions of $\Sigma | (\phi, \lambda)$ and

⁴Alternative implementations of Bayesian estimation of structural VARs identified by external instruments are discussed in Caldara and Herbst (20219) and Arias, Rubio-Ramirez, and Waggoner (2022).

$\phi(\Sigma, \lambda)$. As in Jarocinski and Karadi (2020), we modify the model to impose that $\Phi_1 = 0$ in (12). In specifications with $n_1 > 1$ we use sign restrictions described in Section 4 to set-identify the responses of $W_{1,t}$ to $\epsilon_{1,t}$. We use a truncated prior for Ω_{11} that is uniform on the space of orthogonal matrices, denoted by $p(\Omega_{11}|\Phi, \Sigma)$. Because Ω_{11} does not enter the likelihood function, this prior does not get updated.

4 Data

We are using three types of data: high-frequency instruments for monetary policy shocks, macroeconomic time series, and cross-sectional data on labor earnings and consumption expenditures. The functional VAR with earnings data is estimated based on monthly data, whereas the model with consumption data is estimated based on quarterly data.

High-frequency instruments. The instruments are taken from two previous studies. The first set of instrumental variables ($n_1 = 2$) is taken from Jarocinski and Karadi (2020), who consider two surprise variables that allow them to separate unanticipated changes in monetary policy (monetary policy shocks) from the central bank’s revelation of information about the state of the economy that is conveyed through interest rates (information shocks). The variables are surprises in the three-month fed funds futures (*ff4.hf*), surprises in the S&P 500 stock market index (*sp500.hf*). Sign restrictions are used to separate the two shocks of interest. It is assumed that a contractionary monetary policy shock generates an interest rate increase and a drop in stock prices, whereas a positive information shock is associated with an increase in both interest rates and stock prices. The second instrumental variable ($n_1 = 1$) is obtained from Nakamura and Steinsson (2018). They consider unexpected changes in interest rates over a 30-minute window surrounding scheduled FOMC announcements. Tick-by-tick data on Fed funds futures and eurodollar futures are used to construct the instrument.

Aggregate Variables. Following Jarocinski and Karadi (2020) we use six monthly macroeconomic variables in the empirical model: (i) the monthly average of the one-year constant-maturity Treasury yield serves as the monetary policy indicator. The advantage of using a one-year rate is that it remains a valid measure of monetary policy stance also when the federal funds rate is constrained by the zero lower bound (ZLB). (ii) The monthly average of the S&P 500 stock price index in log levels. (iii,iv) Real GDP and GDP deflator in log levels interpolated to monthly frequency based on Stock and Watson (2010). (v) The excess

bond premium (EBP) as indicator of financial conditions. (vi) An aggregate employment rate constructed from the micro data (see below). (vii) The functional VAR with micro-level consumption data also includes personal consumption expenditures per capita from the National Income and Product Accounts (NIPA).

Micro Data on Earnings. Weekly earnings (PRERNWA) are obtained from the monthly Current Population Survey (CPS) through the website of the National Bureau of Economic Research (NBER). Weekly earnings are scaled to annual earnings by multiplying with 52. Based on the CPS variable PREXPLF “Experienced Labor Force Employment” we construct an employment indicator which is one if the individual is employed and zero otherwise. This indicator is used to compute the aggregate employment rate. We standardize individual-level earnings by $(2/3)$ of nominal per-capita GDP. Rather than taking a logarithmic transformation of the earnings data, we apply the inverse hyperbolic sine transformation, which is given by

$$x = g(z|\theta) = \frac{\ln(\theta z + (\theta^2 z^2 + 1)^{1/2})}{\theta} = \frac{\sinh^{-1}(\theta z)}{\theta}, \quad z = \frac{\text{Earnings}}{(2/3) \cdot \text{per-capita GDP}} \quad (18)$$

with $\theta = 1$. For small values of z the function is approximately equal to z and for large values of z it is equal to $\ln(z) + \ln(2)$. This transformation avoids the thorny issue of applying a log transformation to earnings that are close to zero. Below we will refer to x as transformed data and to z as original data.

Micro Data on Consumption. Our data set is based on the public use micro data of the Consumer Expenditure Survey conducted by the Bureau of Labor Statistics. Individual-level quarterly per capita consumption expenditures c_{it} are computed by summing within-quarter monthly expenditures by respondent and dividing by the number of members with age 16 and over belonging to the consumption unit (e.g., household or family). Let C_t be aggregate per capita consumption from NIPA. We calculate $\frac{1}{T} \sum_{t=1}^T \text{median}(c_{1t}, \dots, c_{Nt})/C_t \approx 0.45$. We then define $z_{it} = c_{it}/(0.45 \cdot C_t)$. Thus, if $z_{it} = 1$ the individual approximately consumes at the level of aggregate consumption per capita. Finally, as for the earnings data, we apply the inverse hyperbolic sine transformation to obtain x_{it} .

Sample Periods. For the Jarocinski and Karadi (2020) instruments, we use the sample periods 1990:M2 to 2016:M12 (earnings) and 1990:Q2 to 2016:Q4 (consumption). This sample has only one missing value (the financial market disruption after the 9/11 terrorist attack in 2001:M9), which is replaced with zero. The sizes are 323 and 108, respectively. For the Nakamura and Steinsson (2018) instrument, the sample periods are 1995:M2 to 2014:M3

(earnings) and 1995:Q1 to 2014:Q1 (consumption). The number of observations are 230 and 76, respectively.

5 Empirical Analysis with Earnings Data

The estimation follows the empirical analysis in Chang, Chen, and Schorfheide (2021). We assume that the transformed earnings are located on the interval $[0, \bar{x}]$ and use a cubic spline as basis functions. We construct the spline from $x = \bar{x}$ to $x = 0$, using a linear element for the right tail:

$$\begin{aligned}\zeta_K(x) &= \max\{\bar{x} - x, 0\} \\ \zeta_k(x) &= [\max\{x_{k-1} - x, 0\}]^3, \quad k = K - 1, \dots, 1.\end{aligned}\tag{19}$$

In a first step, we estimate for each month t a cross-sectional density for the transformed earnings-to-GDP ratio. Based on results in our earlier work we set $K = 10$. Moreover, we set the lag length to $p = 1$. The cross-sectional density estimation delivers the sequence of transformed spline coefficient estimates \hat{a}_t , which is then used in the Bayesian VAR estimation described in Section 3.2. The posterior sampler generates a set of parameter draws $\{(\Phi^i, \Sigma^i, \Omega^i)\}_{i=1}^N$. Based on these parameter draws, we use the VAR law of motion in (9) to generate impulse response functions (IRFs) for (Y_t, a_t) . The a_t IRFs are converted into α_t responses by undoing the compression and standardization. Subsequently, the α_t IRFs are converted into density IRFs using

$$p^{(K)}(x|\alpha_t) = \frac{\exp\{\zeta'(x)\alpha_t\}}{\int \exp\{\zeta'(x)\alpha_t\} dx}.\tag{20}$$

In a final step, we apply a change-of-variable to convert the density for the inverse-hyperbolic-sine transformed earnings, see (18), into actual earnings. In addition to estimating a functional VAR that includes the earnings data, we also estimate a VAR for Y_t , excluding the cross-sectional data.

5.1 Response of Aggregate Variables to a Monetary Policy Shock

We compute IRFs from two functional VARs that differ only with respect to the sample period and the instrumental variables for the identification of the monetary policy shock, which are taken from Jarocinski and Karadi (2020) (JK) and Nakamura and Steinsson (2018)

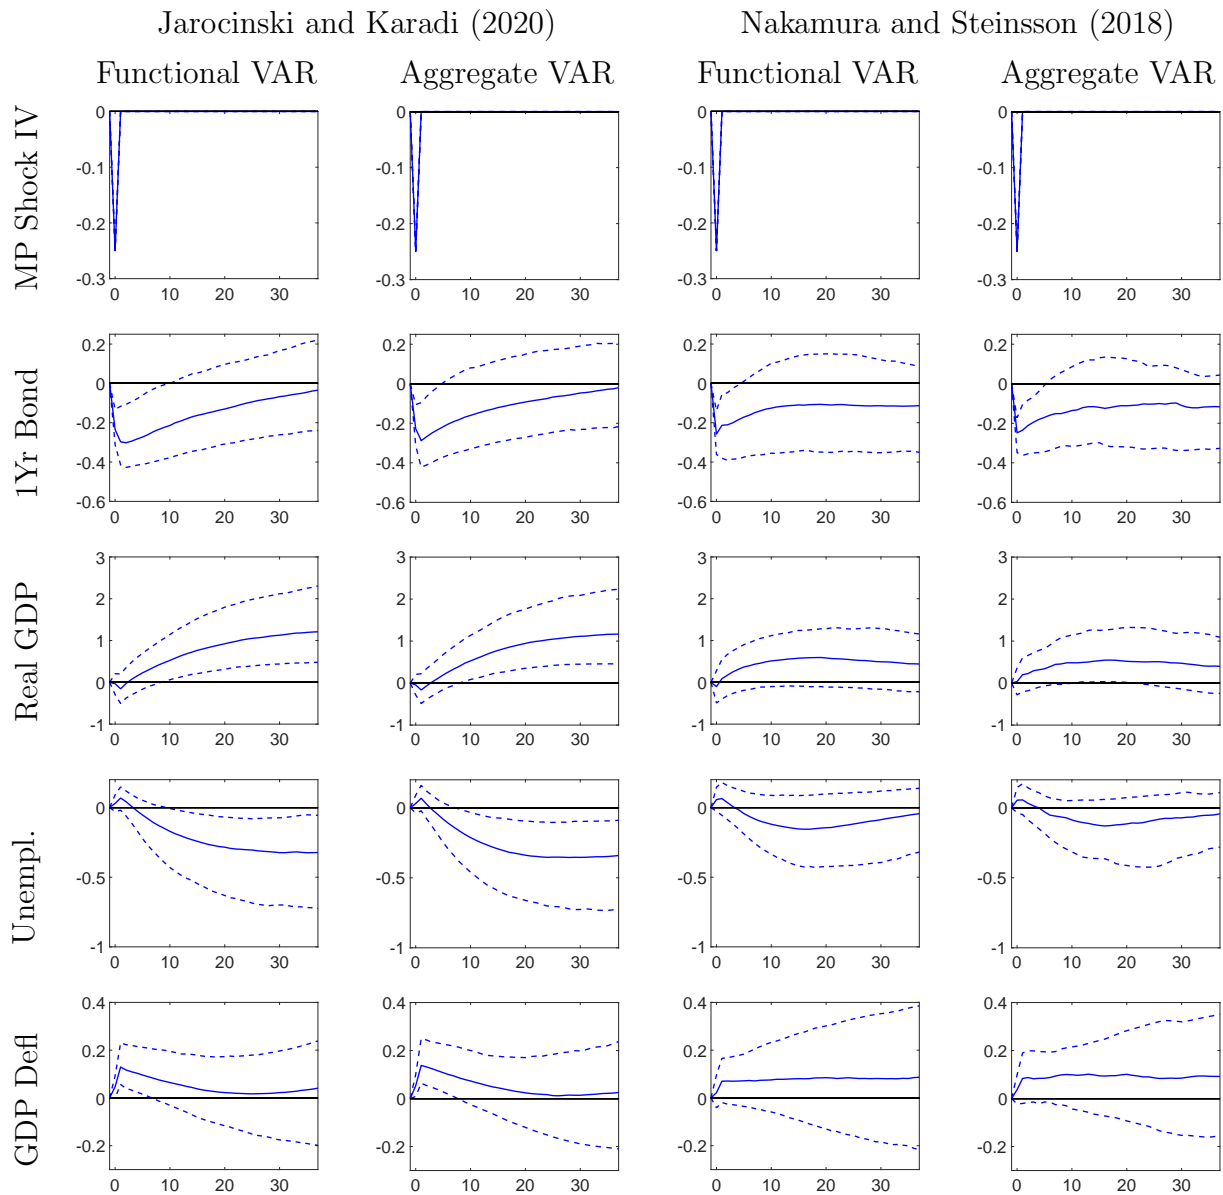
(NS), respectively. In addition to the functional VARs, we also report results from aggregate VARs that exclude the cross-sectional information. We begin with IRFs for the aggregate variables, which are plotted in Figure 1. The responses are normalized such that the surprise reduction in the three-month federal funds rate is 25 basis points (bp). The first two columns are based on the JK instruments, whereas the last two columns are obtained from the NS instrument. The time period for the IRFs is a month.

The responses in the four columns are very similar. Neither the choice of instruments and sample (JK versus NS) nor the presence of the cross-sectional variables has a substantial effect on the impulse response inference. The 80% bands constructed from the functional VAR are slightly wider, reflecting uncertainty generated by the inclusion of additional explanatory variables. At the posterior median, the one-year bond rate moves approximately one-for-one with the federal funds rate surprise and then reverts back to steady state after three years. Under the JK identification, the expansionary monetary policy leads to an increase of real GDP by about 1.2% and a reduction of the unemployment rate by 0.3 percentages after three years at the posterior median. Under the NS identification the real effects of the expansion are smaller in absolute value (approximately half the size) and more mean-reverting. However, there is considerable overlap of the credible bands. Prices increase by about 0.1% upon impact and revert back to the initial level under the JK specification. Under the NS specification the price increase is persistent. Overall, the long-run behavior under the NS identification is more in line with the prediction of economic theory: a monetary policy surprise leads to a temporary change in real activity and a permanent change in the price level.

5.2 Response of Earnings Distribution to a Monetary Policy Shock

Panel (i) of Figure 2 depicts the response of the continuous part of the earnings distribution to a monetary policy shock. The panels show the difference between the steady state earnings density and the shocked density for $h = 0$ (impact of the shock), $h = 4$, $h = 8$, and $h = 12$. The x -axis in these plots correspond to the level of earnings. Recall that a value of one means that the earnings of the individual are equal to $2/3$ (approximately the labor share) of GDP per capita. The earnings densities are normalized to integrate to one minus the unemployment rate. Because the unemployment rate drops in response to an expansionary monetary policy shock, the probability mass increases along the density response, relative to the steady state density.

Figure 1: Responses of Aggregate Variables to Monetary Policy Shock



Notes: Columns 1 and 2: responses to a one-standard-deviation monetary policy shock based on Jarocinski and Karadi (2020). Columns 3 and 4: responses to a one-standard-deviation monetary policy shock based on Nakamura and Steinsson (2018). The Aggregate VAR uses aggregate variables only. The Functional VAR uses the cross-sectional data on earnings in addition. The system is in steady state at $h = -1$ and the shock occurs at $h = 0$. The plots depict 10th (dashed), 50th (solid), and 90th (dashed) percentiles of the posterior distribution. GDP deflator and real GDP responses are percentage deviations from the steady state, whereas the other responses are absolute percentages.

Earnings above 2 are essentially not affected by the monetary policy intervention. Upon impact, the IRFs provide some evidence that the probability mass of individuals earning


between 0.5 and 1 times GDP per capita drops and the mass of individuals earning between 1 and 2 times GDP increases upon impact. However, the 80% bands are wide and the sign of the responses is mostly ambiguous. The density differential for earnings between 1 and 2 reverts quickly to zero, whereas the differential for earnings between 0.5 and 1 becomes positive in the medium run before the earnings density reverts back to its steady state.

An important advantage is that the density response can be easily converted into IRFs for statistics derived from the earnings distribution. Panel (ii) of Figure 2 shows the responses of the percentiles of the earnings distribution as a function of the horizon h . The percentile responses account for a point mass of zero labor earnings that corresponds to the number of individuals that are unemployed. The percentile responses are reported as percentage changes relative to the base level. For instance, suppose in steady state the earnings level is 0.2 times GDP per capita at the 10th percentile and after the shock earnings rise to 0.21. This corresponds to a 5% increase.

According to the plots in Panel (ii), in percentage terms, the monetary policy shock has the largest impact on the earnings distribution at the 10th percentile, capturing in part the individuals moving from unemployment into employment. The posterior median response of the 10th percentile under the JK specification implies a 2.5% increase in earnings at $h = 1$ with an 80% band ranging from about 0 to 5%. For the 20th percentile the response ranges from 0 to 1 percent and for the 80th and 90th percentiles the responses are essentially zero. Using the NS identification scheme and sample leads to wider credible bands and leaves the direction of the response ambiguous. As we have seen from Figure 1, using the NS instrument and sample, the real effects of the monetary policy shock are not quite as strong and are transitory. Accordingly, the NS responses of the 10th and 20th percentiles are smaller than the JK responses and revert back to zero after 3 years.

We proceed by computing four measures of earnings inequality from the cross-sectional densities (accounting for the pointmass at zero): the fraction of individuals earning less than the labor share of GDP, the Gini coefficient, the ratio of the 90th and the 10th percentile of the income distribution (90-10 ratio), and the cross-sectional standard deviation. Impulse responses for the inequality measures are depicted in Figure 3. At the posterior median the fraction of individuals earning less than GDP per capita slightly falls, from 44% to 43.5% under the NS specification, and the cross-sectional standard deviation rises, from 1.07 to 1.08, upon impact of the monetary policy shock. However, the 80% credible bands are wide, leaving the signs of the responses ambiguous, and the effect is in three out of four

Figure 2: Response of Earnings Monetary Policy Shock




Notes: Responses to a 25bp monetary policy shock based on Jarocinski and Karadi (2020) (JK shock) and Nakamura and Steinsson (2018) (NS shock). The system is in steady state at $h = -1$ and the shock occurs at $h = 0$. The plots depict 10th (dashed), 50th (solid), and 90th (dash-dot) percentiles of the posterior distribution. Panel (i): As distributional responses we depict differences between the shocked and the steady state cross-sectional density (continuous part, normalized to $1 - UR_t$) of earnings / GDP per capita at various horizons. Panel (ii): The percentile responses are computed from distribution of actual earnings, accounting for the pointmass at zero.

cases short-lived. Only under the JK instrument, there seems to be a long-run effect on the cross-sectional standard deviation, lowering it to 1.065.

The IRFs from the JK specification imply that the Gini coefficient and the 90-10 ratio fall in response to the expansionary monetary policy shock. At the posterior median, the 90-10

Figure 3: Responses of Inequality Measures to Monetary Policy Shock



Notes: Responses to a 1-standard-deviation monetary policy shock based on Jarocinski and Karadi (2020) (JK shock) and Nakamura and Steinsson (2018) (NS shock). The system is in steady state at $h = -1$ and the shock occurs at $h = 0$. The plots depict 10th (dashed), 50th (solid), and 90th (dashed) percentiles of the posterior distribution. The inequality measures are computed from the distribution of actual earnings and account for the probability mass at zero.

ratio drops from 12.25 to 11.75 after 36 months. The path of the Gini coefficient resembles the path of the 90-10 ratio, falling from 0.431 to 0.428. The comparable IRFs under the NS shock are weaker, have an ambiguous sign according to the 80% credible bands and are less persistent.


5.3 Responses to an Information Shock

The JK instruments also allow us to compute IRFs for an information shock. The results are summarized in Figure 4. Instead of being associated with an unanticipated expansionary monetary policy shock, an unexpected drop in the federal funds rate could simply be a signal from the central bank that it has private information indicating that aggregate output and prices will be lower than expected by the public. The shock is set identified through the assumption that surprises in interest rates have the same sign as surprises to the stock market index.


Panel (i) summarizes the responses of the aggregate variables: a reduction of bond yields indicates that real activity and prices will be below expectation. At the posterior mean real

Figure 4: Responses to Information Shock


Panel (i): Response of Aggregate Variables




Panel (ii): Response of Earnings Density



Panel (iii): Response of Percentiles



Panel (iv): Response of Inequality Measures



Notes: Responses to an information shock based on Jarocinski and Karadi (2020). The system is in steady state at $h = -1$ and the shock occurs at $h = 0$. The plots depict 10th (dashed), 50th (solid), and 90th (dashed) percentiles of the posterior distribution. Panel (i): GDP deflator and real GDP responses are percentage deviations from the steady state, whereas the other responses are absolute percentages. Panel (ii): we depict differences between the shocked and the steady state cross-sectional density (continuous part, normalized to $1 - UR_t$) of earnings / GDP per capita at various horizons. Panels (iii,iv): The responses are computed from distribution of actual earnings, accounting for the pointmass at zero.

GDP drops 20 bp, prices fall by 10 bp and the unemployment rate increases up to 0.15 percentages after one year. The posterior mean responses revert back to zero after three years. Panel (ii) contains the responses of the earnings density. Most notably, the density associated with individuals earning between 0.5 and 1 times the labor share of GDP per capita falls. This effect is most pronounced upon impact of the information shock and then slowly dies out.


Panel (iii) illustrates that the density responses translate into a persistent drop of the earnings at the 10th percentile of about 2% at the posterior median, whereas there is a small and short-lived increase for the other percentiles. Finally, Panel (iv) contains responses of inequality measures. The fraction of individuals earning less than GDP per capita drops in the short run, which is consistent with the density differentials depicted in Panel (i). The cross-sectional standard deviation increases and eventually reverts back to its initial level. According to the response of the 90-10 ratio and the Gini coefficient inequality rises in the long run.

5.4 Discussion

Our empirical analysis focused, in the terminology of Kaplan, Moll, and Violante (2018), on indirect effects of household heterogeneity on the propagation of monetary policy interventions. An expansionary monetary policy lowers the real interest rate temporarily, which creates a disincentive to save and stimulates economic activity in the current period. Labor demand rises and earnings increase. As we have shown in Section 5.2, this increase is most pronounced at the 10th percentile which is consistent with the notion that low-productivity workers move out of unemployment. In turn inequality as measured through the 90-10 ratio and the Gini coefficient falls.

This result mirrors the finding in regard to the response of the earnings distribution to a technology shock reported in Chang, Chen, and Schorfheide (2021). It is broadly consistent with a heterogeneous agent model with indivisible labor supply as in Chang and Kim (2006). This class of models generates a negative correlation between idiosyncratic productivity and reservation wage. In turn, it is low-skill workers who enter the labor market during booms, when the demand for labor is sufficiently high such that the wage per efficiency unit exceeds their reservation wage. At this point the labor earnings switches from zero to a positive value, which reduces labor earnings inequality. Ma (2021) incorporates this mechanism into a HANK model and shows that in his model an expansionary monetary policy shock raises

Figure 5: Responses of Inequality Measures: With and Without Pointmass



Notes: Responses to a 1-standard-deviation monetary policy shock based on Jarocinski and Karadi (2020) (JK shock) and Nakamura and Steinsson (2018) (NS shock). The system is in steady state at $h = -1$ and the shock occurs at $h = 0$. The plot depicts 10th (dashed), 50th (solid), and 90th (dashed) percentiles of the posterior distribution. The inequality measures in the top row are computed from the continuous part of the distribution of actual earnings, not assigning labor earnings of zero for the unemployed. The bottom row reproduces four of the panels in Figure 3.

wages and more low productivity individuals are starting to work, which raises earnings in the left tail of the distribution. Under his calibration the Gini coefficient for labor earnings (on a scale from 0 to 1) drops by approximately 0.001 upon impact.⁵ In our estimated VAR the drop is between 0.001 to 0.002, which is very similar.

As in Ma (2021)'s HANK model, the reduction in the earnings inequality is mainly driven by the fall in unemployment. We recomputed the response of the earnings distribution and the derived inequality measures by excluding the pointmass at zero and normalizing the continuous part of the earnings density to one. The results are plotted in Figure 5. The comparison of the IRFs in the top row (no pointmass at zero for the unemployed) to the IRFs in the bottom row (which are identical to the ones previously shown in Figure 5), shows that the effect of monetary policy shocks on earnings inequality is mostly driven by individuals switching between unemployment and employment. The effect of the monetary policy shocks is small and short-lived.

⁵See Figure 3 in Ma (2021). He considers a 100 bp monetary policy shock and measures the Gini coefficient on a scale from 0 to 100. Thus, $-0.4/(4 \cdot 100) = 0.001$.

There are a number of earlier studies that examined the effect of monetary policy shocks on earnings inequality. The authors of these studies typically included the inequality measures directly into a VAR or linear projections. For instance, Coibion, Gorodnichenko, Kueng, and Silvia (2017) report IRFs to a 100 bp increase in the monetary policy rate, estimated on U.S. data. They find that in the medium run the Gini coefficient on earnings rises by about 0.0025. Adjusting for the different shock size, their estimate is a slightly smaller than ours. Furceri, Loungani, and Zdzienicka (2018) consider a panel of 32 advanced and emerging market economies. They report an estimate (converted into our scale) of 0.005, which is larger, but in the same order of magnitude as our estimates.

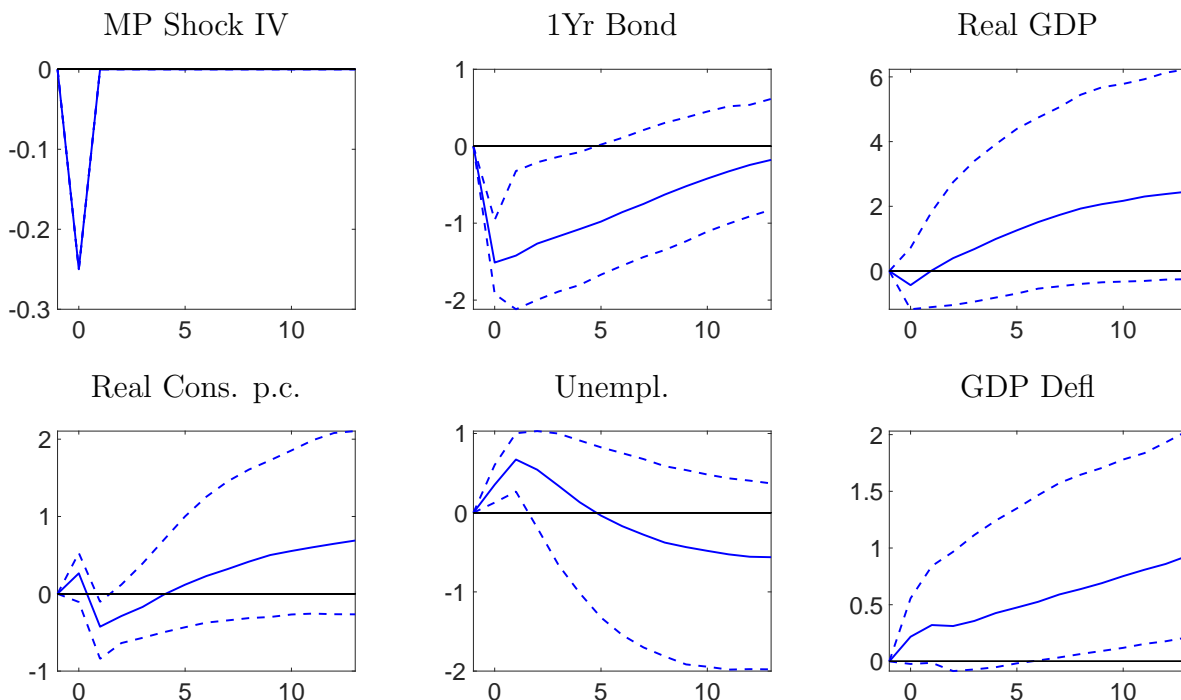
6 Empirical Analysis with Consumption Data

Responses of aggregate variables. Because our cross-sectional consumption data are at quarterly frequency, we time-aggregate the monthly aggregate data and estimate the functional VAR on quarterly data from 1990:Q2 to 2016:Q4. In addition to the aggregate series that were used in Section 5 we now also include real per capita consumption from NIPA. Throughout this section we focus on responses to monetary policy shocks identified using the JK instruments.

Figure 6 shows the responses of the aggregate variables to a monetary policy shock. Because the model is estimated on quarterly data, the x -axes in the figure now refer to quarters. Qualitatively, the aggregate responses from the functional consumption VAR are very similar to the responses functional earnings VAR in Figure 1. Quantitatively, there is a difference in magnitude. The reason is that a 25 basis point surprise at monthly frequency is only roughly a third of a 25 basis point surprise over an entire quarter. In the Online Appendix we time aggregate the responses from the monthly earnings VAR to quarterly frequency and rescale the IRFs such that the instrument moves by 25 basis points over the quarter. After this re-scaling, the magnitude of the aggregate responses from the VARs estimated with monthly and quarterly data, respectively, are indeed very similar.

Response of Cross-sectional Distribution and Inequality Measures. Figure 7 shows the response of the cross-sectional consumption distribution to a monetary policy shock. Panel (i) shows density differentials, Panel (ii) shows the responses of percentiles of the consumption distribution, and Panel (iii) depicts the response of inequality measures. Recall

Figure 6: Responses of Aggregate Variables to JK Monetary Policy Shock

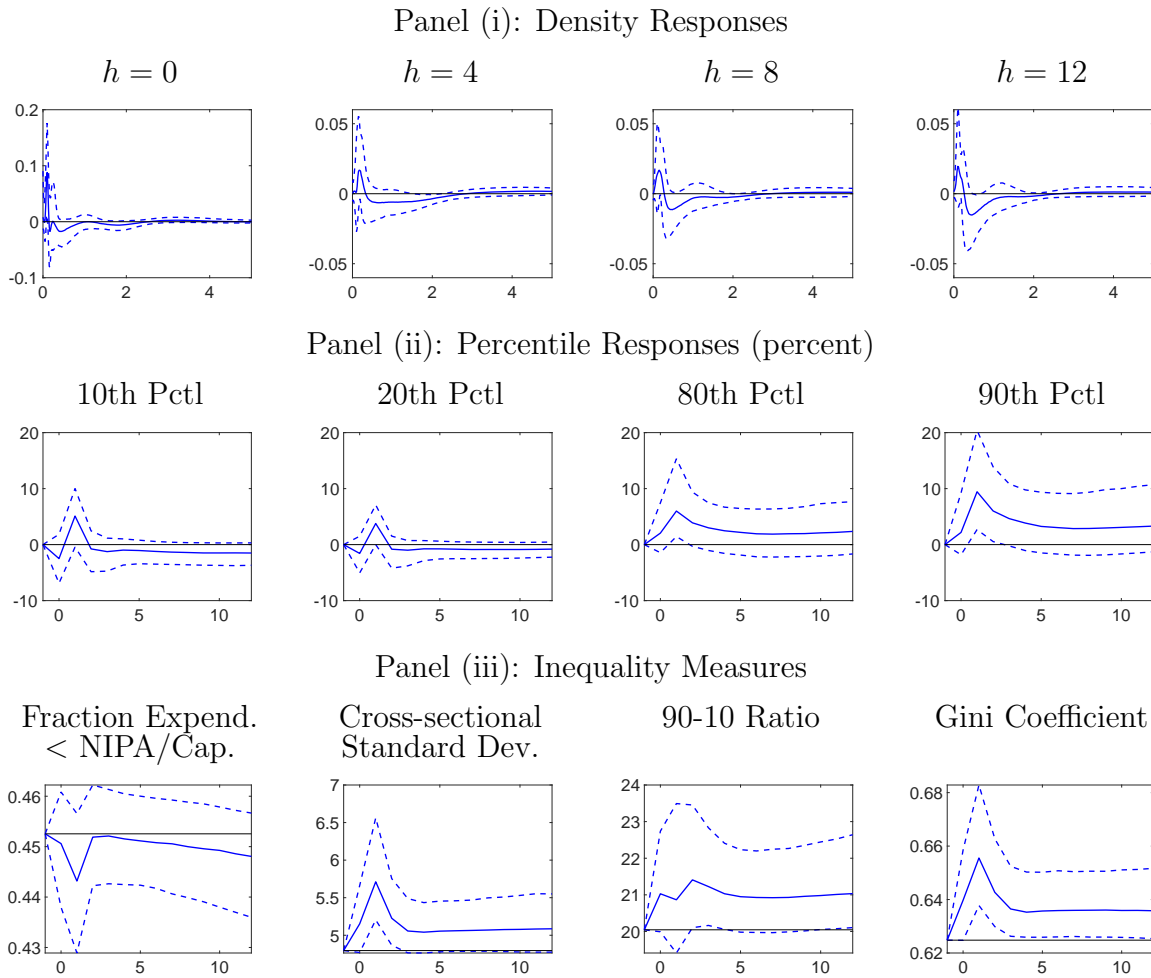


Notes: Responses to a monetary policy shock based on Jarocinski and Karadi (2020). The x -axes refer to quarters. The system is in steady state at $h = -1$ and the shock occurs at $h = 0$. The plots depict 10th (dashed), 50th (solid), and 90th (dashed) percentiles of the posterior distribution.

that the individual-level consumption expenditures are normalized by NIPA per capita consumption, meaning that the consumption expenditures of an individual with a value of one approximately equals aggregate per capita consumption. The differentials in Panel (i) are with respect to a steady state distribution with a median of 1.0. The percentile responses reported in Panel (ii) are in percentage deviations from percentiles of the steady state distribution. For instance, a 5% increase at the median would imply that consumption rises from 1 to 1.05.

We begin with the results shown in Panel (ii). Based on the posterior mean the 10th and 20th percentiles slightly drop on impact ($h = 0$), then rise above their steady state values in the subsequent period, before they fall below their initial values in the second quarter after impact. Over the subsequent quarters the responses stay 0.8% (20th percentile) to 1.3% (10th percentile) below their baseline values. The responses of the 80th and 90th percentiles are markedly different: after an initial overshooting the percentiles stay persistently around above their steady state values: 2.4% for the 80th and 3.5% for 90th percentile. While the

Figure 7: Response of Consumption to JK Monetary Policy Shock



Notes: Responses to a monetary policy shock based on Jarocinski and Karadi (2020). The system is in steady state at $h = -1$ and the shock occurs at $h = 0$. The plots depict 10th (dashed), 50th (solid), and 90th (dashed) percentiles of the posterior distribution. Panel (i): we depict differences between the shocked and the steady state cross-sectional density of consumption expenditures/aggregate consumption per capita at various horizons. Panels (ii, iii): The x -axes refer to quarters. The responses are computed from distribution of consumption expenditures/aggregate consumption per capita. Percentile responses are in percent from steady state value.

80% credible intervals do not exclude zero, the zero line lies close to the boundary of the credible bands.

The percentile responses are computed from the density responses, depicted as differentials with respect to the steady state density in Panel (i). To reconcile the information in the two panels, consider horizons $h = 4$ and $h = 8$. The 10th and 20th percentiles of the steady state distribution are 0.18 and 0.3, respectively. Because density is increasing in response

to the monetary policy shock for consumption levels between 0 and 0.3, the percentiles shift to the left, which leads to a negative response in terms of percentage deviations from the steady state value, as can be seen in Panel (ii).

The 80th and 90th percentiles of the steady state distribution are 2.4 and 3.6, respectively. In the Online Appendix we provide a figure that zooms into the density response for values of $x > 2$. These figures show that the density response is positive for values of $x > 2.5$. In combination with a long right tail of the steady state distribution, this causes a shift of the percentiles to the right, resulting in a consumption increase at the 80th and 90th percentiles.

Panel (iii) of Figure 7 depicts the responses of our four inequality measures, which are also derived from the response of the cross-sectional density of consumption. The fraction of households consuming less than the aggregate per capita amount falls, while the other three inequality measures – cross-sectional standard deviation, the 90-10 ratio, and the Gini coefficient rise. Thus, overall, consumption inequality unambiguously increases in response to interest rate cut.

Discussion. Our main finding is that in response to an expansionary monetary policy shock earnings inequality decreases while consumption inequality increases. In theoretical models, the consumption response is determined by the income response in combination with the marginal propensity to consume (MPC). In this regard, the increase in consumption inequality is a “puzzle” because typically low-income households are associated with a high MPC.

There are three possible channels that could explain the result. First, we only consider labor income which is indirectly affected by monetary policy. Rich households typically also have a significant amount of unearned income, which may be directly affected the monetary policy intervention, e.g., rising stock and bond prices in response to an interest rate cut. Our current analysis did not use direct observations on the response of the unearned income distribution. Second, even wealthy households who have a relatively high level of consumption may be close to their borrowing constraint (wealthy hand-to-mouth consumers) and have a relatively high marginal propensity to consume which could explain the larger response in the right tail of the distribution. Third, the direct effect of monetary policy on consumption may be heterogeneous, whereby high consumption/income households’ consumption decisions are more interest rate sensitive.

7 Conclusion

We estimated a functional VAR that stacks macroeconomic aggregates and the cross-sectional distributions of earnings to provide semi-structural evidence about the effect of monetary policy shocks on earnings inequality. We found that an expansionary monetary policy shock reduces inequality. The reduction is generated by what we call the employment channel. At the left end of the earnings distribution, the expansion lifts individuals out of unemployment and thereby reduces the earnings dispersion. For consumption we obtain the opposite result: the expansionary policy shock raises the mass in the right tail of the consumption distribution, thereby increasing consumption inequality. We offered some potential explanations for these findings, but leave a more granular decomposition of the effects into the relative importance of competing propagation channels for future research.

References

- ACHARYA, S., W. CHEN, M. DEL NEGRO, K. DOGRA, E. MATLIN, AND R. SAFARTI (2021): “Estimating HANK: Macro Time Series and Micro Moments,” *Working Paper, FRB New York*.
- ANDERSON, E., A. INOUE, AND B. ROSSI (2016): “Heterogeneous Consumers and Fiscal Policy Shocks,” *Journal of Money, Credit, and Banking*, 48(8), 1877–1888.
- ARIAS, J. E., J. RUBIO-RAMIREZ, AND D. WAGGONER (2022): “Inference in Bayesian Proxy-SVARs,” *Journal of Econometrics*, forthcoming.
- AUCLERT, A. (2019): “Monetary Policy and the Redistribution Channel,” *American Economic Review*, 109(6), 2333–2367.
- BAYER, C., B. BORN, AND R. LUETTICKE (2020): “Shocks, Frictions, and Inequality in US Business Cycles,” *Working paper, University of Bonn*.
- BOSQ, D. (2000): *Linear Processes in Function Spaces*. Springer Verlag, New York.
- CALDARA, D., AND E. HERBST (2021): “Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs,” *American Economic Journal: Macroeconomics*, 11(1), 157–192.

- CHANG, M., X. CHEN, AND F. SCHORFHEIDE (2021): “Heterogeneity and Aggregate Fluctuations,” *NBER Working Paper*, 28853.
- CHANG, Y., AND S.-B. KIM (2006): “From Individual to Aggregate Labor Supply: A Quantitative Analysis Based on a Heterogeneous Agent Macroeconomy,” *International Economic Review*, 47, 1–27.
- COIBION, O., Y. GORODNICHENKO, L. KUENG, AND J. SILVIA (2017): “Innocent Bystanders? Monetary Policy and Inequality,” *Journal of Monetary Economics*, 88, 70–89.
- DOEPKE, M., AND M. SCHNEIDER (2006): “Inflation and the Redistribution of Nominal Wealth,” *Journal of Political Economy*, 114(6), 1069–1097.
- FISHER, I. (1933): “The Debt-Deflation Theory of Great Depressions,” *Econometrica*, 1(4), 337–357.
- FURCERI, D., P. LOUNGANI, AND A. ZDZIENICKA (2018): “The Effects of Monetary Policy Shocks on Inequality,” *Journal of International Money and Finance*, 85, 168–186.
- GUERELLO, C. (2018): “Conventional and Unconventional Monetary Policy vs. Household Income Distribution: An Empirical Analysis for the Euro Area,” *Journal of International Money and Finance*, 85, 187–214.
- HORVATH, L., AND P. KOKOSZKA (2012): *Inference for Functional Data with Applications*. Springer Verlag, New York.
- JAROCINSKI, M., AND P. KARADI (2020): “Deconstructing Monetary Policy Surprises - The Role of Information Shocks,” *American Economic Journal: Macroeconomics*, 12, 1–43.
- KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): “Monetary Policy According to HANK,” *American Economic Review*, 108(3), 697–743.
- KAPLAN, G., AND L. VIOLANTE, GIOVANNI (2018a): “Microeconomic Heterogeneity and Macroeconomic Shocks,” *NBER Working Paper*, 24734.
- (2018b): “Microeconomic Heterogeneity and Macroeconomic Shocks,” *Journal of Economic Perspective*, 32, 167–194.
- MA, E. (2021): “Monetary Policy and Inequality: How Does One Affect the Other,” *Working Paper*, Louisiana State University.

NAKAMURA, E., AND J. STEINSSON (2018): “High-Frequency Identification of Monetary Non-Neutrality: The Information Effect,” *Quarterly Journal of Economics*, 133, 1283–1330.

RAMSEY, J. O., AND B. W. SILVERMAN (2005): *Functional Data Analysis*. Springer Verlag, New York, 2nd edn.

STOCK, J. H., AND M. W. WATSON (2010): “Monthly GDP and GNI – Research Memorandum,” *Manuscript, Princeton University*.

VARIN, C., N. REID, AND D. FIRTH (2011): “An Overview of Composite Likelihood Methods,” *Statistica Sinica*, 21, 5–42.

Online Appendix: On the Effects of Monetary Policy Shocks on Earnings and Consumption Heterogeneity

Minsu Chang, and Frank Schorfheide

This Appendix consists of the following sections:

- A. Data Used in the Empirical Analysis
- B. Additional Empirical Results

A Data Used in the Empirical Analysis

Aggregate Data. Following Jarocinski and Karadi (2020) we use six monthly macroeconomic variables in the empirical model: (i) the monthly average of the one-year constant-maturity Treasury yield serves as the monetary policy indicator. (ii) The monthly average of the S&P 500 stock price index in log levels. (iii,iv) Real GDP and GDP deflator in log levels interpolated to monthly frequency based on Stock and Watson (2010). (v) The excess bond premium (EBP) as indicator of financial conditions. (vi) An aggregate employment rate constructed from the micro data (see below). The functional VAR with micro-level consumption data includes an additional aggregate variable: real personal consumption expenditures per capita in log levels. We use Personal Consumption Expenditures (*PCE* from FRED) and divide it by population level (*CNP16OV* from FRED) to get per capita values. Then we use GDP deflator to get the real values.

Micro-Level Earnings Data. The CPS raw data are downloaded from <http://www.nber.org/data/cps-basic.html>.

The raw data files are converted into STATA using the do-files available at: http://www.nber.org/data/cps-basic_progs.html.

We use the series PREXPLF (“Experienced Labor Force Employment”), which is the same as in the raw data, and the series PRERNWA (“Weekly Earnings”), which is constructed

as PEHRUSL1 (“Hours Per Week at One’s Main Job”) times PRHERNAL (“Hourly Earnings”) for hourly workers, and given by PRWERNAL for weekly workers. STATA dictionary files are available at:

<http://www.nber.org/data/progs/cps-basic/>

We pre-process the cross-sectional data as follows. We drop individuals if (i) the employment indicator is not available; and (ii) if they are coded as “employed” but the weekly earnings are missing. In addition, we re-code individuals with non-zero earnings as employed and set earnings to zero for individuals that are coded as not employed. A CPS-based unemployment rate is computed as the fraction of individuals that are coded as not employed. By construction this is one minus the fraction of individuals with non-zero weekly earnings, which is used to normalize the cross-sectional density of earnings. It turns out that the CPS-based unemployment rate tracks the aggregate unemployment rate (*UNRATE* from FRED) very closely. The levels of the two series are very similar, but the CPS unemployment rate exhibits additional high-frequency fluctuations, possibly due to seasonals that have been removed from the aggregate unemployment rate.

Micro-Level Consumption Data. We use public use microdata (PUMD) from the Consumer Expenditure Survey (CEX) conducted by the Bureau of Labor Statistics (BLS). Quarterly expenditure (by respondent ID and expenditure category) are computed as the sum over three months in a quarter. We divide the expenditures by the number of individuals with age 16 and over belonging to the consumption unit (i.e. household or family) to obtain per capita expenditures. Because CEX per capita expenditures capture less than 50% of NIPA per capita expenditures we rescale them as follows. Let C_t be aggregate per capita consumption from NIPA. We calculate $\frac{1}{T} \sum_{t=1}^T \text{median}(c_{1t}, \dots, c_{Nt})/C_t \approx 0.45$. We then define $z_{it} = c_{it}/(0.45 \cdot C_t)$. Thus, if $z_{it} = 1$ the individual approximately consumes at the level of aggregate consumption per capita. Finally, as for the earnings data, we apply the inverse hyperbolic sine transformation to obtain x_{it} .

Transformation of Micro Data. We transform the micro data (earnings-GDP ratio, or consumption relative to aggregate per capita consumption), denoted by z below, using the

inverse hyperbolic sine transformation, which is given by

$$x = g(z|\theta) = \frac{\ln(\theta z + (\theta^2 z^2 + 1)^{1/2})}{\theta} = \frac{\sinh^{-1}(\theta z)}{\theta} \quad (\text{A.1})$$

with $\theta = 1$. Note that $g(0|\theta) = 0$ and $g^{(1)}(0|\theta) = 1$, that is, for small values of z the transformation is approximately linear. For large values of z the transformation is logarithmic:

$$g(z|\theta) \approx \frac{1}{\theta} \ln(2\theta z) = \frac{1}{\theta} \ln(2\theta) + \frac{1}{\theta} \ln(z).$$

The inverse of the transformation takes the form

$$z = g^{-1}(x|\theta) = \frac{1}{\theta} \sinh(\theta x) = \frac{1}{2\theta} (e^{\theta x} - e^{-\theta x}).$$

Most of the calculations in the paper are based on $p_x(x)$. But in some instances, it is desirable to report for $p_z(z)$. From a change of variables (omitting the θ), we get

$$p_z(z) = p_x(g(z)) |g'(z)|,$$

where

$$g'(z) = \frac{1 + \frac{\theta z}{(\theta^2 z^2 + 1)^{1/2}}}{\theta z + (\theta^2 z^2 + 1)^{1/2}} = \frac{1}{(\theta^2 z^2 + 1)^{1/2}}.$$

Whenever we do convert the estimated densities back from z to x , we recycle the density evaluations at x_j . Thus, we evaluate $p_z(z)$ for grid points $z_j = g^{-1}(x_j)$, which leads to

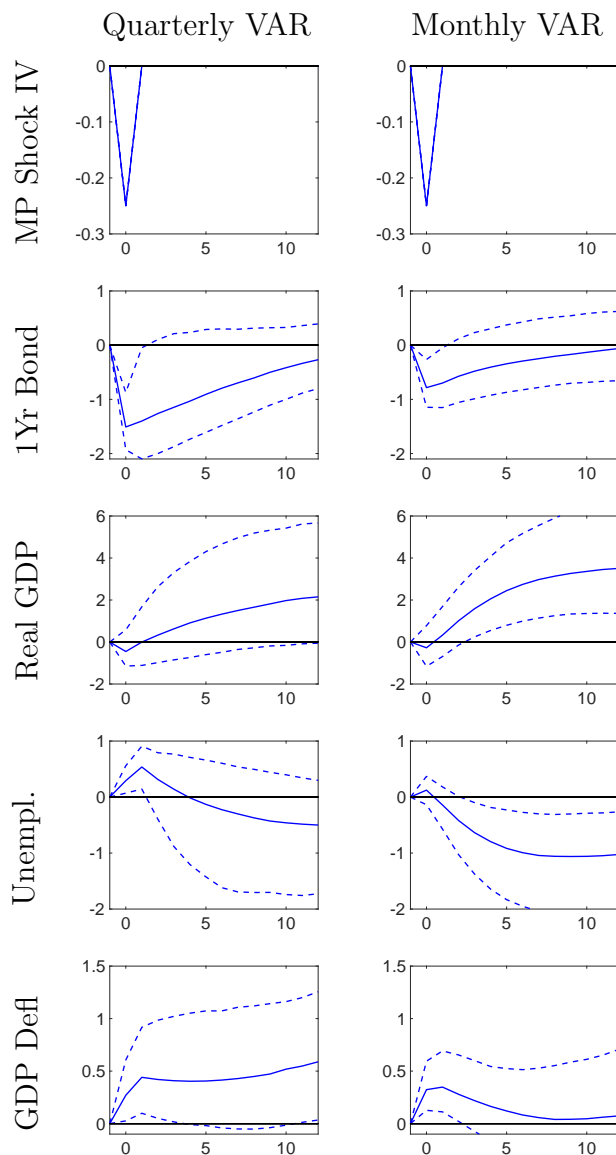
$$p_z(z_j) = p_x(x_j) |g'(g^{-1}(x_j))|,$$

where

$$|g'(g^{-1}(x_j))| = \frac{1}{\left(\frac{1}{4}(e^{\theta x_j} - e^{-\theta x_j})^2 + 1\right)^{1/2}} = \frac{2}{(e^{2\theta x_j} + e^{-2\theta x_j} + 2e^{2\theta x_j} e^{-2\theta x_j})^{1/2}} = \frac{2}{e^{\theta x_j} + e^{-\theta x_j}}.$$

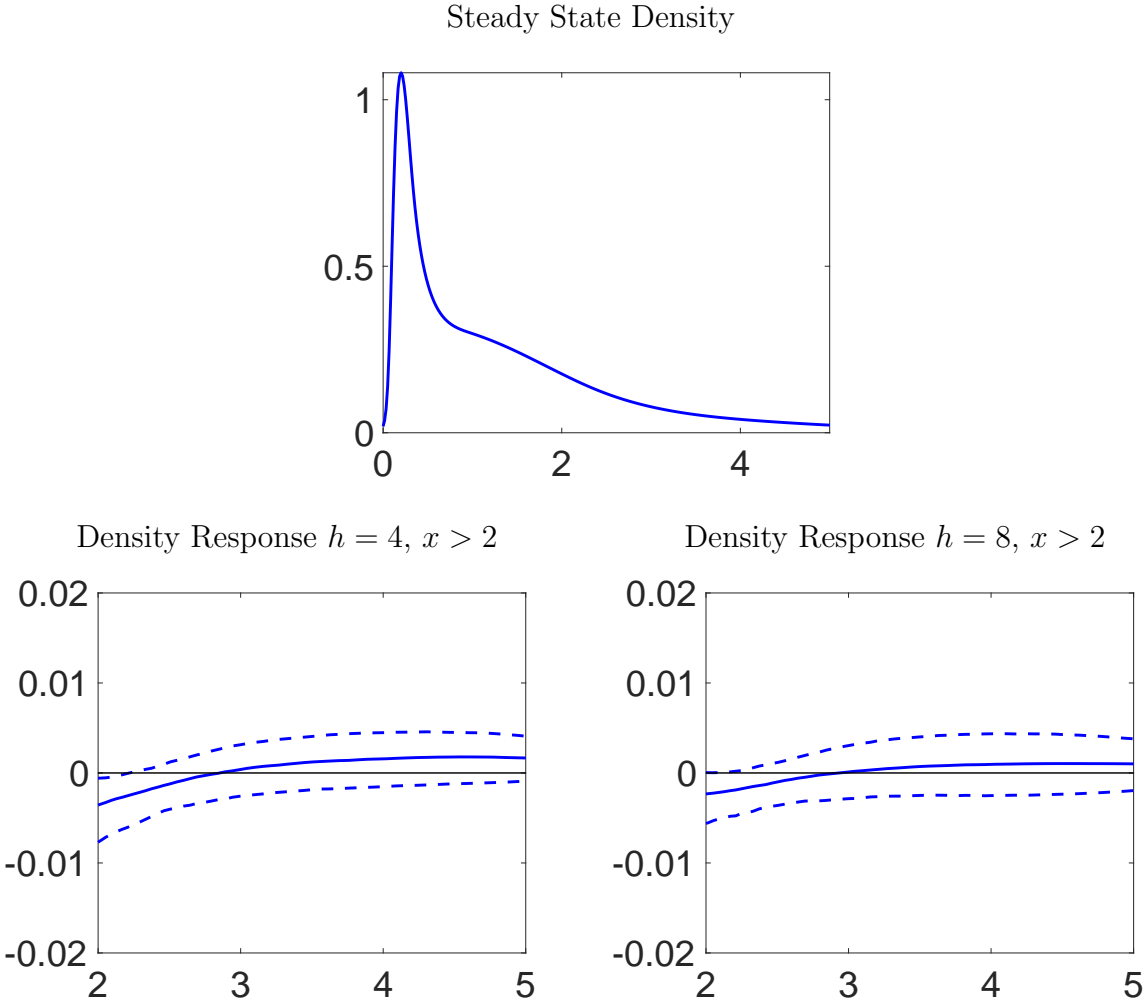
B Additional Empirical Results

Figure A-1: Responses of Aggregate Variables to JK Monetary Policy Shock



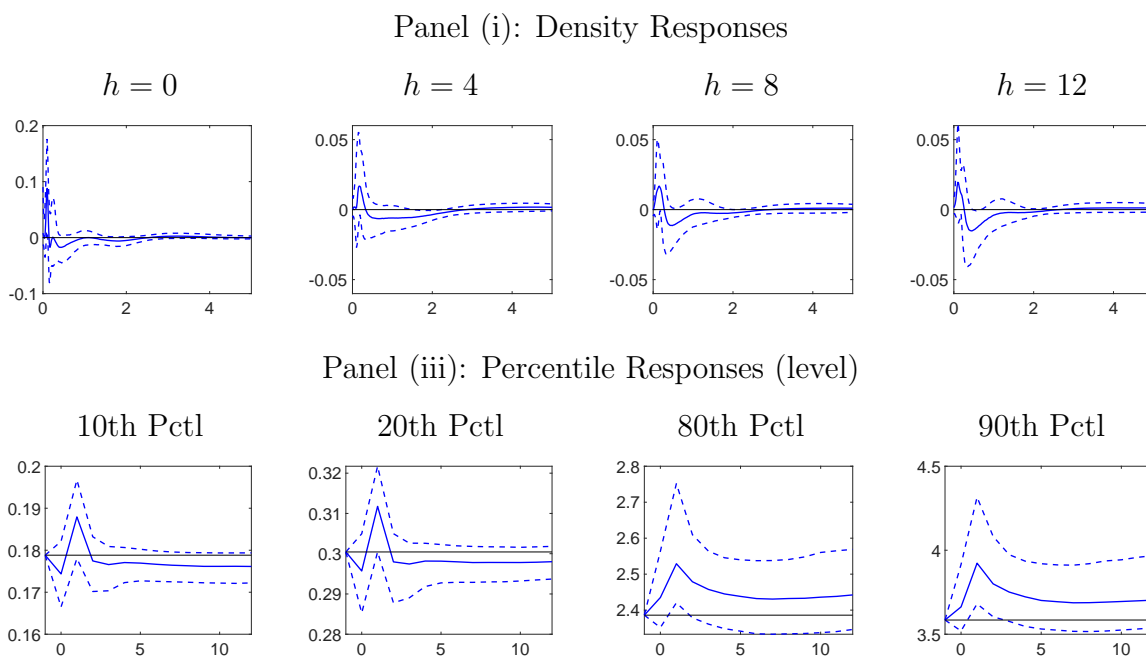
Notes: This figure compares IRFs from aggregate VARs (without aggregate consumption) estimated on monthly and quarterly data. We use a quarterly scale for the x-axis of the monthly VAR responses and plot the average response over 3 months in each quarter. The responses in the *Monthly VAR* column are scaled so that the MP Shock IV impact is the same as in the *Quarterly VAR* column.

Figure A-2: Steady State Density of Consumption and Response



Notes: Left panel: estimated steady state density of consumption. The quantiles are: 10% = 0.2, 20%=0.3, 50%=1.00, 80%=2.4, and 90%=3.6. Right panel: density responses to JK monetary policy shock at horizon $h = 4$, reported as differential with respect to steady state.

Figure A-3: Response of Consumption to JK Monetary Policy Shock



Notes: Responses to a monetary policy shock based on Jarocinski and Karadi (2020). The system is in steady state at $h = -1$ and the shock occurs at $h = 0$. The plots depict 10th (dashed), 50th (solid), and 90th (dashed) percentiles of the posterior distribution. Panel (i) depicts differences between the shocked and the steady state cross-sectional density of consumption expenditures/aggregate consumption per capita at various horizons. It is identical to Panel (i) in Figure 7. Panel (ii): The x -axes refer to quarters. The quantile responses are computed from distribution of consumption expenditures/aggregate consumption per capita.