Shocks, Frictions, and Inequality in US Business Cycles*
Preliminary and Incomplete

Christian Bayer, Benjamin Born, and Ralph Luetticke

September 12, 2019

Abstract

In how far does inequality matter for the business cycle and vice versa? Using a Bayesian likelihood approach, we estimate a heterogeneous-agent New-Keynesian (HANK) model with incomplete markets and portfolio choice between liquid and illiquid assets. The model enlarges the set of shocks and frictions in Smets and Wouters (2007) by allowing for shocks to income risk and portfolio liquidity. We find income risk to be an important driver of output and consumption. This makes US recessions more demand driven relative to the otherwise identical complete markets benchmark (RANK). The HANK model further implies that business cycle shocks and policy responses have significantly contributed to the evolution of US wealth and income inequality.

JEL codes: E32, E52

Keywords: Incomplete Markets, Business Cycles

*Bayer: University of Bonn, CEPR and IZA, christian.bayer@uni-bonn.de, Born: Frankfurt School of Finance & Management, CEPR and CESifo, b.born@fs.de, Luetticke: University College London and CEPR, r.luetticke@ucl.ac.uk. We would like to thank Gregor Boehl, Miguel Leon-Ledesma, Johannes Pfeifer, and seminar and conference participants at AMCM 2019 in Lillehammer, the 3rd MMCN Conference in Frankfurt, the 2019 meeting of the VfS Committee on Monetary Economics, the 2019 meetings of the EEA in Manchester, the 2019 EABCN meetings in Mannheim, Junior Finance and Macro Conference at U Chigaco, SED 2019, Bocconi, U Kent, the ECB for helpful comments and suggestions. Christian Bayer gratefully acknowledges support through the ERC-CoG project Liquid-House-Cycle funded by the European Union's Horizon 2020 Program under grant agreement No. 724204. We would like to thank Sharanya Pillai for excellent research assistance.

1 Introduction

A new generation of monetary business cycle models has become popular featuring heterogeneous agents and incomplete markets (known as HANK models). This new class of models implies new transmission channels of monetary[1] and fiscal[2] policy, as well as new sources of business cycle fluctuations working through household portfolio decisions.[3] While much of this literature so far has focused on specific channels of transmission, shocks, or puzzles, the present paper asks how our overall view of the business cycle and inequality, of the underlying aggregate shocks and frictions, changes when we bring such a model to the data.

For this purpose, we study the business cycle using a technique that has become standard at least since Smets' and Wouters' (2007) seminal paper, extending this technique to the analysis of HANK models: We estimate an incomplete markets model by a full information Bayesian likelihood approach using the state-space representation of the model. Specifically, we estimate an extension of the New-Keynesian incomplete markets model of Bayer et al. (2019). We add features such as capacity utilization, a frictional labor market with sticky wages, and time variation in the liquidity of assets, as well as the usual plethora of shocks that drive business cycle fluctuations in estimated New-Keynesian models: aggregate and investment-specific productivity shocks, wage- and price-markup shocks, monetary- and fiscal-policy shocks, risk premium shocks, and, as two additional incomplete-market-specific ones, shocks to the liquidity of assets and shocks to idiosyncratic productivity risk.

In this model, precautionary motives play an important role for consumption-savings decisions. Since individual income is subject to idiosyncratic risk that cannot be directly insured and borrowing is constrained, households structure their savings decisions and portfolio allocations to optimally self-insure and achieve consumption smoothing. In particular, we assume that households can either hold liquid nominal bonds or invest in illiquid physical capital. Capital is illiquid because its market is segmented and households participate only from time to time. This portfolio-choice component and the presence of occasional hand-to-mouth consumers leads the HANK model to interpret data in a different manner than its complete-markets, representative-agent twin (in short: RANK), which is otherwise identical except for market completeness and all assets being perfectly liquid.

To infer the importance of household heterogeneity for the business cycle, we first estimate

[1] Auclert (2019) analyzes the redistributive effects of monetary policy, Kaplan et al. (2018) show the importance of indirect income effects, and Luetticke (2018) analyzes the portfolio rebalancing channel of monetary policy. McKay et al. (2016) studies the effectiveness of forward guidance.

[2] Auclert et al. (2018) and Hagedorn et al. (2018a) discuss the fiscal multiplier, McKay and Reis (2016) discuss the role of automatic stabilizers.

[3] Bayer et al. (2019) quantify the importance of shocks to idiosyncratic income risk, and Guerrieri and Lorenzoni (2017) look at the effects of shocks to the borrowing limit

both models on the same observables as in Smets and Wouters (2007) (plus proxies for income risk and liquidity) covering the time period of 1954 to 2015.[4] We find that with incomplete markets—compared to the complete markets benchmark—demand shocks are more important for business cycle fluctuations.[5] This is true for both output growth as well as for its components. Relative to RANK, demand shocks explain roughly 30% more of output volatility. The increased importance of demand is driven by shocks to income uncertainty, which explain almost 20% of consumption volatility. This reflects the fact that portfolio choices in our HANK model—even up to a first order approximation in aggregates—react to income and risk positions of households.

The difference between HANK and RANK is even more pronounced in the historical decomposition of US recessions. Through the lens of the HANK model, 42% of output losses in US recessions come from demand shocks. This number drops to 7% when the same data is viewed through the RANK model.

To analyze US inequality, we re-estimate the model with two additional observables for the shares of wealth and income held by the top-10% of households in each dimension, which are taken from the World-Income-Database.[6] The addition of distributional data does not significantly change what we infer about shocks and frictions. However, we find that business cycle shocks can explain the very persistent movements in wealth and income inequality in the US over 1954-2015. In the HANK model, even transitory shocks have very persistent effects on inequality, because wealth is a slowly moving variable that accumulates past shocks.

The historical decomposition of US inequality reveals that TFP, markups and fiscal policy are the main contributors to the rise of wealth and income inequality from the 1990s to today. We find that a more expansionary fiscal policy that would have driven up the rates on government bonds and driven down the liquidity premium could have decreased wealth and income inequality substantially. Income risk shocks play a significant role for consumption inequality, because wealth poor, and thus badly insured, households react to an uncertainty increase by cutting consumption particularly strongly, while for well insured households, that are already consumption rich, behavior changes little. Consequently, these shocks account for 20% of the cyclical variations in consumption inequality.

[4] We use the estimates of income risk for the US provided by Bayer et al. (2019) and months of housing supply as proxy for liquidity.

[5] Demand side shocks are shocks to liquidity, uncertainty, government spending, monetary policy and the risk premium, and supply side shocks are the two markup and the two productivity shocks. The grouping of the shocks is based on the question whether they directly affect the Phillips curve as the relevant aggregate supply function or primarily work through affecting the bond-market clearing condition, as the aggregate demand function.

[6] Since these data come at mixed frequencies and with observational gaps, it is key that we obtain a state-space representation of our model, which allows us to use a standard Kalman filter to obtain the likelihood of the model.

Overall, this shows that fluctuations in idiosyncratic income risk and asset-market liquidity are important elements to understand the cyclical behavior both of aggregates and of inequality. This is line with ample evidence that income risk and liquidity are both negatively correlated with the cycle.[7] We strengthen this evidence and show that fluctuations in both are to a large extent the result of exogenous shocks. We do so by estimating alongside shocks to the two respective variables also feedback parameters for liquidity and uncertainty on other aggregate state variables. The estimated feedback implies counter-cyclical fluctuations in both, but is quantitatively unimportant.

To our knowledge, our paper is the first to provide an encompassing estimation of shocks and frictions using a HANK model with portfolio choice. Most of the literature on monetary heterogeneous agent models has used a calibration approach (see for example Auclert et al., 2018; Ahn et al., 2018; Bayer et al., 2019; Broer et al., 2016; Challe and Ragot, 2015; Den Haan et al., 2017; Gornemann et al., 2012; Guerrieri and Lorenzoni, 2017; McKay et al., 2016; McKay and Reis, 2016; Ravn and Sterk, 2017; Sterk and Tenreyro, 2018; Wong, 2019). Auclert et al. (2019b) and Hagedorn et al. (2018b) both go beyond calibration but use one-asset HANK models. The latter provide parameter estimates based on impulse-response function matching, while the former estimate the model using the MA-∞ representation in the sequence space.

Focusing on the methodological contribution, Auclert et al. (2019a) provide a fast estimation method for heterogeneous agent models that, however, requires a sequence space representation of the model and thus does not allow to deal with missing or mixed frequency data as we need to do here, when combining cross-sectional and aggregate data. Since this is the setup we are facing, we build on the solution method of Reiter (2009) using the dimensionality reduction approach of Bayer and Luetticke (2018) to make this feasible for estimation. We further exploit that only a small fraction of the Jacobian of the non-linear difference equation that represents the model needs to be re-calculated during the estimation.

Related in the sense that it estimates a state-space model of both distributional (cross sectional) data and aggregates is also the paper by Chang et al. (2018). They find that, in an SVAR sense, shocks to the cross sectional distribution of income have only a mild impact on aggregate time-series. Our finding of structural estimates being relatively robust to the inclusion or exclusion of cross sectional information resembles their results.[8]

[7] Storesletten et al. (2001) estimate that for the US the variance of persistent income shocks to disposable household income almost doubles in recessions. Similarly, Guvenen et al. (2014b) find a sizable increase in the left skewness of the income distribution in recessions. Various measures of liquidity are counter-cyclical as well. Hedlund (2016) documents a sharp increase in the time to sell a house in the US during the Great Recession. Similarly, also credit spreads rise in recessions, too; see Gilchrist and Zakrajšek (2012).

[8] Our approach is different and simpler than the method suggested by Liu and Plagborg-Møller (2019) which includes full cross-sectional information into the estimation of a heterogeneous agent DSGE model.

We also contribute to the study of inequality. Previous studies that use quantitative models to understand the evolution of inequality consider permanent changes in the US tax and transfer system and solve for steady state transitions; see, e.g., Kaymak and Poschke (2016) or Hubmer et al. (2016). They find that these changes can explain a significant part of the recent increase in wealth inequality. Complementary to this literature, we estimate a HANK model to study in how far the conduct of fiscal and monetary policy over the business cycle contributes to inequality. Compared to this literature, we analyze not only policy rules but allow for a wide range of other business cycle shocks.

The remainder of this paper is organized as follows: Section 2 describes our model economy, its sources of fluctuations and its frictions. Section 3 provides details on the numerical solution method and estimation technique. Section 4 presents the parameters that we calibrate to match steady-state targets and our main estimation results for all other parameters, and it gives an overview over the data we employ in our estimation. Section 5 discusses the shocks and frictions driving the US business cycle. Section 6 does so for US inequality. Section 7 concludes. An appendix follows.

2 Model

We model an economy composed of a firm sector, a household sector and a government sector. The firm sector comprises (a) perfectly competitive intermediate goods producers who rent out labor services and capital; (b) final goods producers that face monopolistic competition, producing differentiated final goods out of homogeneous intermediate inputs; (c) producers of capital goods that turn consumption goods into capital subject to adjustment costs; (d) labor packers that produce labor services combining differentiated labor from (e) unions that differentiate raw labor rented out from households. Price setting for the final goods as well as wage setting by unions is subject to a pricing friction à la Rotemberg (1982).[9]

Households earn income from supplying (raw) labor and capital and from owning the firm sector, absorbing all its rents that stem from market power of unions and final-goods producers, and decreasing returns to scale in capital goods production.

The government sector runs both a fiscal authority and a monetary authority. The fiscal authority levies a time-constant tax on labor income and distributed profits, issues government bonds, and adjusts expenditures to stabilize debt in the long run and aggregate demand

We in contrast only use the model to fit certain generalized cross-sectional moments.

[9] We choose Rotemberg (1982) over Calvo (1983) price adjustment costs as this implies all firms to have the same profits and avoids introducing cross-sectional profit risk. In terms of the implied Phillips curve, both assumptions are identical for our estimation because we solve the model using a first-order perturbation in aggregates.

in the short run. The monetary authority sets the nominal interest rate on government bonds according to a Taylor rule.

2.1 Households

The household sector is subdivided into two types of agents: workers and entrepreneurs. The transition between both types is stochastic. Both rent out physical capital, but only workers supply labor. The efficiency of a worker's labor evolves randomly exposing worker-households to labor-income risk. Entrepreneurs do not work, but earn all pure rents in our economy except for the rents of unions which are equally distributed across workers. All households self-insure against the income risks they face by saving in a liquid nominal asset (bonds) and a less liquid physical asset (capital). Trading illiquid capital is subject to random participation in the capital market.

To be specific, there is a continuum of ex-ante identical households of measure one, indexed by i. Households are infinitely lived, have time-separable preferences with time-discount factor β, and derive felicity from consumption c_{it} and leisure. They obtain income from supplying labor, n_{it}, from renting out capital, k_{it}, and from interest on bonds, b_{it}, and potentially profit income or union transfers.

A household's labor income $w_t h_{it} n_{it}$ is composed of the aggregate wage rate on raw labor, w_t, the household's hours worked, n_{it}, and its idiosyncratic labor productivity, h_{it}. We assume that productivity evolves according to a log-AR(1) process with time-varying volatility and a fixed probability of transition between the worker and the entrepreneur state:

$$\tilde{h}_{it} = \begin{cases} \exp\left(\rho_h \log \tilde{h}_{it-1} + \epsilon_{it}^h\right) & \text{with probability } 1-\zeta \text{ if } h_{it-1} \neq 0, \\ 1 & \text{with probability } \iota \text{ if } h_{it-1} = 0, \\ 0 & \text{else,} \end{cases} \tag{1}$$

with individual productivity $h_{it} = \frac{\tilde{h}_{it}}{\int \tilde{h}_{it} di}$ such that $\tilde{h}_{it}$ is scaled by its cross-sectional average, $\int \tilde{h}_{it} di$, to make sure that average worker productivity is constant. The shocks ϵ_{it}^h to productivity are normally distributed with time-varying variance that follows a log AR(1) process with endogeneous feedback to aggregate hours N_{t+1} (hats denote log-deviations from steady state):

$$\sigma_{h,t}^2 = \bar{\sigma}_h^2 \exp \widehat{s}_t, \tag{2}$$

$$\widehat{s}_{t+1} = \rho_s \widehat{s}_t + \Sigma_Y \widehat{N}_{t+1} + \epsilon_t^\sigma \ , \tag{3}$$

i.e., at time t households observe a change in the variance of shocks that drive the next period's productivity. With probability ζ households become entrepreneurs ($h = 0$). With probability ι an entrepreneur returns to the labor force with median productivity. An entrepreneurial household obtains a fixed share of the pure rents (aside union rents), Π_t^F, in the economy (from monopolistic competition in the goods sector and the creation of capital). We assume that the claim to the pure rent cannot be traded as an asset. Union rents, Π_t^U are distributed lump-sum across workers, leading to labor-income compression.

With respect to leisure and consumption, households have Greenwood et al. (1988) (GHH) preferences and maximize the discounted sum of felicity:[10]

$$E_0 \max_{\{c_{it}, n_{it}\}} \sum_{t=0}^{\infty} \beta^t u \left[c_{it} - G(h_{it}, n_{it}) \right]. \tag{4}$$

The maximization is subject to the budget constraints described further below. The felicity function u exhibits a constant relative risk aversion (CRRA) with risk aversion parameter $\xi > 0$,

$$u(x_{it}) = \frac{1}{1-\xi} x_{it}^{1-\xi},$$

where $x_{it} = c_{it} - G(h_{it}, n_{it})$ is household i's composite demand for goods consumption c_{it} and leisure and G measures the disutility from work. Goods consumption bundles varieties j of differentiated goods according to a Dixit-Stiglitz aggregator:

$$c_{it} = \left(\int c_{ijt}^{\frac{\eta_t - 1}{\eta_t}} dj \right)^{\frac{\eta_t}{\eta_t - 1}}.$$

Each of these differentiated goods is offered at price p_{jt}, so that for the aggregate price level, $P_t = \left(\int p_{jt}^{1-\eta_t} dj \right)^{\frac{1}{1-\eta_t}}$, the demand for each of the varieties is given by

$$c_{ijt} = \left(\frac{p_{jt}}{P_t} \right)^{-\eta_t} c_{it}.$$

The disutility of work, $G(h_{it}, n_{it})$, determines a household's labor supply given the ag-

[10] The assumption of GHH preferences is mainly motivated by the fact that many estimated DSGE models of business cycles find small aggregate wealth effects in labor supply, see e.g. Born and Pfeifer (2014). It also simplifies the numerical analysis somewhat. Unfortunately, it is not feasible to estimate a flexible Jaimovich and Rebelo (2009) preference form, which encompasses also King et al. (1988) preferences. This would require solving the stationary equilibrium in every likelihood evaluation, which is substantially more time consuming than solving for the dynamics around this equilibrium.

gregate wage rate, w_t, and a labor income tax, τ, through the first-order condition:

$$\frac{\partial G(h_{it}, n_{it})}{\partial n_{it}} = (1-\tau) w_t h_{it}. \tag{5}$$

Assuming that G has a constant elasticity w.r.t. n, $\frac{\partial G(h_{it}, n_{it})}{\partial n_{it}} = (1+\gamma)\frac{G(h_{it}, n_{it})}{n_{it}}$ with $\gamma > 0$, we can simplify the expression for the composite consumption good x_{it} making use of the first-order condition (5) and substitute $G(h, n)$ out of the individual planning problem:

$$x_{it} = c_{it} - G(h_{it}, n_{it}) = c_{it} - \frac{(1-\tau) w_t h_{it} n_{it}}{1+\gamma}. \tag{6}$$

When the Frisch elasticity of labor supply is constant, the disutility of labor is always a constant fraction of labor income. Therefore, in both the budget constraint of the household and its felicity function only after-tax income enters and neither hours worked nor productivity appears separately.[11]

The households optimize subject to their budget constraint:

$$\begin{aligned} c_{it} + b_{it+1} + q_t k_{it+1} = (1-\tau)(h_{it} w_t N_t + \mathbb{I}_{h_{it} \neq 0} \Pi_t^U + \mathbb{I}_{h_{it}=0} \Pi_t^F) \\ + b_{it} \tfrac{R(b_{it}, R_t^b, A_t)}{\pi_t} + (q_t + r_t) k_{it}, \quad k_{it+1} \geq 0, b_{it+1} \geq \underline{B}, \end{aligned}$$

where b_{it} is real bond holdings, k_{it} is the amount of illiquid assets, q_t is the price of these assets, r_t is their dividend, $\pi_t = \frac{P_t - P_{t-1}}{P_{t-1}}$ is realized inflation, and R is the nominal interest rate on bonds, which depends on the portfolio position of the household and the central bank's interest rate R_t^b, which is set one period before. All households that do not to participate in the capital market ($k_{it+1} = k_{it}$) still obtain dividends and can adjust their bond holdings. Depreciated capital has to be replaced for maintenance, such that the dividend, r_t, is the net return on capital.

Market participation is random and i.i.d. in the sense that a fraction, λ_t, of households is selected to adjust their capital holdings in a given period. This fraction, λ_t, itself follows an autoregressive process with endogenous feedback to the bond rate R_{t+1}^b:

$$\widehat{\lambda}_{t+1} = \rho_\lambda \widehat{\lambda}_t + \Lambda_R \widehat{R}_{t+1}^B + \epsilon_t^\lambda \ . \tag{7}$$

[11] This implies that we can assume $G(h_{it}, n_{it}) = h_{it} \frac{n_{it}^{1+\gamma}}{1+\gamma}$ without further loss of generality as long as we treat the empirical distribution of labor income as a calibration target. This functional form simplifies the household problem as h_{it} drops out from the first-order condition and all households supply the same number of hours $n_{it} = N(w_t)$. Total effective labor input, $\int n_{it} h_{it} di$, is hence also equal to $N(w_t)$ because $\int h_{it} di = 1$. This means that we can read off productivity risk directly from the estimated income risk series of Bayer et al. (2019).

Holdings of bonds have to be above an exogenous debt limit $\underline{B}$, and holdings of capital have to be non-negative.

We assume that there is a wasted intermediation cost that drives a wedge between the government bond yield R^b_t an the interest paid by/to households R_t. This wedge is given by a time varying wedge, A_t, plus a constant, $\overline{R}$, when households resort to unsecured borrowing. Therefore, we specify:

$$R(b_{it}, R^b_t, A_t) = \begin{cases} R^b_t A_t & \text{if } b_{it} \geq 0 \\ R^b_t A_t + \overline{R} & \text{if } b_{it} < 0. \end{cases}$$

The assumption of a borrowing wedge creates a mass of households with zero unsecured credit but with the possibility to borrow, though at a penalty rate. The efficiency wedge A_t can be thought of as a cost of intermediating government debt to households. It follows an AR(1) process in logs and fluctuates in response to shocks, ϵ^A_t. If A_t goes down, household will demand less government bonds and find it more attractive to save in (illiquid) real capital, akin to the "risk-premium shock" in Smets and Wouters (2007).

Substituting the expression $c_{it} = x_{it} + \frac{(1-\tau)w_t h_{it} N_t}{1+\gamma}$ for consumption, we obtain:

$$\begin{aligned} x_{it} + b_{it+1} + q_t k_{it+1} = b_{it}\tfrac{R(b_{it}, R^b_t, A_t)}{\pi_t} + (q_t + r_t)k_{it} + (1-\tau)\tfrac{\gamma}{1+\gamma} w_t h_{it} N_t \\ + (1-\tau)\left(\mathbb{I}_{h_{it}\neq 0}\Pi^U_t + \mathbb{I}_{h_{it}=0}\Pi^F_t\right), \quad k_{it+1} \geq 0, \quad b_{it+1} \geq \underline{B}. \end{aligned} \tag{8}$$

Since a household's saving decision will be some non-linear function of that household's wealth and productivity, inflation and all other prices will be functions of the joint distribution, Θ_t, of (b, k, h) in t. This makes Θ a state variable of the household's planning problem and this distribution evolves as a result of the economy's reaction to aggregate shocks. For simplicity, we summarize all effects of aggregate state variables, including the distribution of wealth and income, by writing the dynamic planning problem with time-dependent continuation values.

This leaves us with three functions that characterize the household's problem: The value function V^a for the case where the household adjusts its capital holdings, the value function V^n for the case in which it does not adjust, and the expected envelope value, $\mathbb{E}V$, over both:

$$\begin{aligned} V^a_t(b, k, h) =& \max_{k', b'_a} u[x(b, b'_a, k, k', h)] + \beta \mathbb{E}_t V_{t+1}(b'_a, k', h) \\ V^n_t(b, k, h) =& \max_{b'_n} u[x(b, b'_n, k, k, h)] + \beta \mathbb{E}_t V_{t+1}(b'_n, k, h) \\ \mathbb{E}_t V_{t+1}(b', k', h) =& \mathbb{E}_t\left[\lambda_{t+1} V^a_{t+1}(b', k', h)\right] + \mathbb{E}_t\left[(1-\lambda_{t+1}) V^n_{t+1}(b', k, h)\right] \end{aligned} \tag{9}$$

Expectations about the continuation value are taken with respect to all stochastic processes conditional on the current states, including time-varying income risk and liquidity.

2.2 Firm Sector

The firm sector consists of four sub-sectors: (a) a labor sector composed of "unions" that differentiate raw labor and labor packers who buy differentiated labor and then sell labor services to intermediate goods producers, (b) intermediate goods producers who hire labor services and rent out capital to produce goods, (c) final goods producers who differentiate intermediate goods, selling these then to goods bundlers, who finally sell them as consumption goods to households and to (d) capital good producers, who turn bundled final goods into capital goods.

When profit maximization decisions in the firm sector require intertemporal decisions (i.e. in price and wage setting and in producing capital goods), we assume for tractability that they are delegated to a mass-zero group of households (managers) that are risk neutral and compensated by a share in profits.[12] They do not participate in any asset market and have the same discount factor as all other households. Since managers are a mass-zero group in the economy, their consumption does not show up in any resource constraint and all, but the unions', profits – net of price adjustment costs – go to the entrepreneur households (whose $h = 0$). Union profits go lump sum to worker households.

2.2.1 Labor Packers and Unions

Worker households sell their labor services to a mass-one continuum of unions indexed by j, who each offer a different variety of labor to labor packers who then provide labor services to intermediate goods producers. Labor packers produce final labor services according to the production function

$$N_t = \left(\int n_{jt}^{\frac{\zeta_t - 1}{\zeta_t}} dj \right)^{\frac{\zeta_t}{\zeta_t - 1}}, \tag{10}$$

out of labor varieties n_{jt}. Cost minimization by labor packers implies that each variety of labor, each union j, faces a downward sloping demand curve

$$n_{jt} = \left(\frac{W_{jt}}{W_t^F} \right)^{-\zeta_t} N_t,$$

[12]Since we solve the model by a first order perturbation in aggregate shocks, the assumption of risk-neutrality only serves as a simplification in terms of writing down the model. With a first-order perturbation we have certainty equivalence and fluctuations in stochastic discount factors become irrelevant.

where W_{jt} is the *nominal* wage set by union j and W_t^F is the nominal wage at which labor packers sell labor services to final goods producers.

Since unions have market power, they pay the households a wage lower than the price at which they sell labor to labor packers. Given the nominal wage W_t at which they buy labor from households and given the *nominal* wage index W_t^F, unions seek to maximize their discounted stream of profits. In doing so, they face costs of adjusting wages charged from the labor packers, W_t^F, which are quadratic in the log rate of wage change and proportional to the wage sum in the economy, $N_t \frac{W_t^F}{P_t} \frac{\zeta_t}{2\kappa_w} \left(\log \frac{W_{jt}}{W_{jt-1}}\right)^2$. They therefore maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t \frac{W_t^F}{P_t} N_t \left\{ \left(\frac{W_{jt}}{W_t^F} - \frac{W_t}{W_t^F} \right) \left(\frac{W_{jt}}{W_t^F} \right)^{-\zeta_t} - \frac{\zeta_t}{2\kappa_w} \left(\log \frac{W_{jt}}{W_{jt-1}\bar{\pi}^W} \right)^2 \right\}, \tag{11}$$

by adjusting W_{jt} every period; $\bar{\pi}^W$ is steady state wage inflation and the fact that it shows up in wage adjustment costs reflects indexation.

Since all unions are symmetric, we focus on a symmetric equilibrium and obtain the wage Phillips curve from the corresponding first order condition as follows, leaving out all terms irrelevant at a first order approximation around the stationary equilibrium:

$$\log \left(\frac{\pi_t^W}{\bar{\pi}^W} \right) = \beta E_t \log \left(\frac{\pi_{t+1}^W}{\bar{\pi}^W} \right) + \kappa_w \left(\frac{w_t}{w_t^F} - \frac{1}{\mu_t^W} \right), \tag{12}$$

with $\pi_t^W := \frac{W_t^F}{W_{t-1}^F} = \frac{w_t^F}{w_{t-1}^F} \pi_t^Y$ being wage inflation, w_t and w_t^F being the respective *real* wages for households and firms, and $\frac{1}{\mu_t^W} = \frac{\zeta_t - 1}{\zeta_t}$ being the target mark-down of wages the unions pay to households, W_t, relative to the wages charged to firms, W_t^F. This target fluctuates in response to markup shocks, $\epsilon_t^{\mu W}$, and follows a log AR(1) process.[13]

2.2.2 Final Goods Producers

Similar to unions, final-goods producers differentiate a homogeneous intermediate good and set prices. They face a downward sloping demand curve

$$y_{jt} = (p_{jt}/P_t)^{-\eta_t} Y_t$$

[13]Up to the first order approximation around the steady state, the Phillips curve is identical the Phillips curve of a model with Calvo adjustment costs. Including the first-order irrelevant terms, the Phillips curve reads

$$\log \left(\frac{\pi_t^W}{\bar{\pi}^W} \right) = \beta E_t \left[\log \left(\frac{\pi_{t+1}^W}{\bar{\pi}^W} \right) \frac{\zeta_{t+1}}{\zeta_t} \frac{W_{t+1}^F P_t}{W_t^F P_{t+1}} \frac{N_{t+1}}{N_t} \right] + \kappa_w \left(\frac{w_t}{w_t^F} - \frac{1}{\mu_t^W} \right).$$

for each good j and buy the intermediate good at the nominal price MC_t. As we do for unions, we assume price adjustment costs à la Rotemberg (1982).

Under this assumption, the firms' managers maximize the present value of real profits given this costs of price adjustment, i.e., they maximize:

$$E_0 \sum_{t=0}^{\infty} \beta^t Y_t \left\{ \left(\frac{p_{jt}}{P_t} - \frac{MC_t}{P_t} \right) \left(\frac{p_{jt}}{P_t} \right)^{-\eta_t} - \frac{\eta_t}{2\kappa} \left(\log \frac{p_{jt}}{p_{jt-1}\bar{\pi}} \right)^2 \right\}, \tag{13}$$

with a time constant discount factor.

The corresponding first-order condition for price setting implies again a Phillips curve

$$\log \left(\frac{\pi_t^Y}{\bar{\pi}} \right) = \beta E_t \log \left(\frac{\pi_{t+1}^Y}{\bar{\pi}} \right) + \kappa \left(mc_t - \tfrac{1}{\mu_t^Y} \right), \tag{14}$$

where we dropped again all terms irrelevant for a first order approximation. Here, π_t^Y is the gross inflation rate of final goods, $\pi_t^Y := \frac{P_t}{P_{t-1}}$, $mc_t := \frac{MC_t}{P_t}$ is the real marginal costs, $\bar{\pi}$ steady state inflation and $\mu_t^Y = \frac{\eta_t}{\eta_t - 1}$ is the target markup. As for the unions, this target fluctuates in response to markup shocks, $\epsilon^{\mu Y}$, and follows a log AR(1) process. We choose the cost to vary with the target markup to create a Phillips curve with a constant steepness as under Calvo adjustment. The price adjustment then creates real costs $\frac{\eta_t}{2\kappa} Y_t \log(\pi_t/\bar{\pi})^2$.

2.2.3 Intermediate Goods Producers

Intermediate goods are produced with a constant returns to scale production function:

$$Y_t = Z_t N_t^{\alpha} (u_t K_t)^{(1-\alpha)},$$

where Z_t is total factor productivity and follows an autoregressive process in logs, and $u_t K_t$ is the effective capital stock taking into account utilization u_t, i.e., the intensity with which the existing capital stock is used. Using capital with an intensity higher than normal results in increased depreciation of capital according to $\delta(u_t) = \delta_0 + \delta_1 (u_t - 1) + \delta_2/2 (u_t - 1)^2$, which, assuming $\delta_1, \delta_2 > 0$, is an increasing and convex function of utilization. Without loss of generality, capital utilization in steady state is normalized to 1, so that δ_0 denotes the steady-state depreciation rate of capital goods.

Let mc_t be the relative price at which the intermediate good is sold to final-good producers. The intermediate-good producer maximizes profits,

$$mc_t Z_t Y_t - w_t^F N_t - [r_t + q_t \delta(u_t)] K_t,$$

but it operates in perfectly competitive markets, such that the real wage and the user costs of capital are given by the marginal products of labor and effective capital:

$$w_t^F = \alpha mc_t Z_t \left(\frac{u_t K_t}{N_t}\right)^{1-\alpha}, \tag{15}$$

$$r_t + q_t \delta(u_t) = u_t(1-\alpha) mc_t Z_t \left(\frac{N_t}{u_t K_t}\right)^{\alpha}. \tag{16}$$

We assume that utilization is decided by the owners of the capital goods, taking the aggregate supply of capital services as given. The optimality condition for utilization is given by

$$q_t \left[\delta_1 + \delta_2(u_t - 1)\right] = (1-\alpha) mc_t Z_t \left(\frac{N_t}{u_t K_t}\right)^{\alpha}, \tag{17}$$

i.e., capital owners increase utilization until the marginal maintenance costs equal the marginal product of capital services.

2.2.4 Capital Goods Producers

Capital goods producers take the relative price of capital goods, q_t, as given in deciding about their output, i.e. they maximize

$$E_0 \sum_{t=0}^{\infty} \beta^t I_t \left\{ \Psi_t q_t \left[1 - \frac{\phi}{2}\left(\log \frac{I_t}{I_{t-1}}\right)^2\right] - 1 \right\}, \tag{18}$$

where Ψ_t governs the marginal efficiency of investment à la Justiniano et al. (2010, 2011), which follows an AR(1) process in logs and is subject to shocks ϵ_t^{Ψ}.[14]

Optimality of the capital goods production requires (again dropping all terms irrelevant up to first order)

$$\Psi_t q_t \left[1 - \phi \log \frac{I_t}{I_{t-1}}\right] = 1 - \beta E_t \left[\Psi_{t+1} q_{t+1} \phi \log\left(\frac{I_{t+1}}{I_t}\right)\right], \tag{19}$$

and each capital goods producer will adjust its production until (19) is fulfilled.

Since all capital goods producers are symmetric, we obtain as the law of motion for

[14] This shock has to be distinguished from a shock to the relative price of investment, which has been shown in the literature (Justiniano et al., 2011; Schmitt-Grohé and Uribe, 2012) to not be an important driver of business cycles as soon as one includes the relative price of investment as an observable. We therefore focus on the MEI shock.

aggregate capital

$$K_t - (1-\delta(u_t))\, K_{t-1} = \Psi_t \left[1 - \frac{\phi}{2}\left(\log \frac{I_t}{I_{t-1}}\right)^2\right] I_t \ . \tag{20}$$

The functional form assumption implies that investment adjustment costs are minimized and equal to 0 in steady state.

2.3 Government

The government operates a monetary and a fiscal authority. The monetary authority controls the nominal interest rate on liquid assets, while the fiscal authority issues government bonds to finance deficits and adjusts expenditures to stabilize debt in the long run and output in the short run.

We assume that monetary policy sets the nominal interest rate following a Taylor (1993)-type rule with interest rate smoothing:

$$\frac{R^b_{t+1}}{\bar{R}^b} = \left(\frac{R^b_t}{\bar{R}^b}\right)^{\rho_R} \left(\frac{\pi_t}{\bar{\pi}}\right)^{(1-\rho_R)\theta_\pi} \left(\frac{Y_t}{Y_{t-1}}\right)^{(1-\rho_R)\theta_Y} \epsilon^R_t. \tag{21}$$

The coefficient $\bar{R}^b \geq 0$ determines the nominal interest rate in the steady state. The coefficients $\theta_\pi, \theta_Y \geq 0$ govern the extent to which the central bank attempts to stabilize inflation and output growth around their steady-state values. $\rho_R \geq 0$ captures interest rate smoothing.

We assume that the government issues bonds according to the rule (c.f. Woodford, 1995):

$$\frac{B_{t+1}}{\bar{B}} = \left(\frac{B_t R^b_t/\pi_t}{\bar{B}\bar{R}^b/\bar{\pi}}\right)^{\rho_B} \left(\frac{Y_t}{\bar{Y}}\right)^{\gamma_Y} \varepsilon^G_t, \tag{22}$$

using tax revenues $T_t = \tau(N_t w_t + \Pi^U_t + \Pi^F_t)$ to finance government consumption, G_t, and interest on debt. We treat the tax rate, τ, as fixed over the cycle.

There are thus two shocks to government rules: monetary policy shocks, ϵ^R_t and government spending shocks, ϵ^G_t.[15] The government budget constraint then determines government spending $G_t = B_{t+1} + T_t - R^b_t/\pi_t B_t$.

[15] Note that we allow for first-order autocorrelation in the government spending shock, such that $\log \varepsilon^G_t = \rho_G \log \varepsilon^G_{t-1} + \epsilon^G_t$.

2.4 Goods, Bonds, Capital, and Labor Market Clearing

The labor market clears at the competitive wage given in (15). The bond market clears whenever the following equation holds:

$$B_{t+1} = B^d(R_t^b, A_t, r_t, q_t, \Pi_t^F, \Pi_t^U, w_t, \lambda_t, \Theta_t, V_{t+1}) := \mathbb{E}_t\left[\lambda_t b_{a,t}^* + (1-\lambda_t) b_{n,t}^*\right], \tag{23}$$

where $b_{a,t}^*, b_{n,t}^*$ are functions of the states (b, k, h), and depend on how households value asset holdings in the future, $V_{t+1}(b, k, h)$, and the current set of prices $(R_t^b, A_t, r_t, q_t, \Pi_t^F, \Pi_t^U, w_t)$. Future prices do not show up because we can express the value functions such that they summarize all relevant information on the expected future price paths. Expectations in the right-hand-side expression are taken w.r.t. the distribution $\Theta_t(b, k, h)$. Equilibrium requires the total *net* amount of bonds the household sector demands, B^d, to equal the supply of government bonds. In gross terms there are more liquid assets in circulation as some households borrow up to $\underline{B}$.

Last, the market for capital has to clear:

$$K_{t+1} \quad = \quad K^d(R_t^b, A_t, r_t, q_t, \Pi_t^F, \Pi_t^U, w_t, \lambda_t, \Theta_t, V_{t+1}) := \mathbb{E}_t[\lambda_t k_t{}^* + (1-\lambda_t)k], \tag{24}$$

where the first equation stems from competition in the production of capital goods, and the second equation defines the aggregate supply of funds from households – both those that trade capital, $\lambda_t k_t^*$, and those that do not, $(1-\lambda_t)k$. Again k_t^* is a function of the current prices and continuation values. The goods market then clears due to Walras' law, whenever labor, bonds, and capital markets clear.

2.5 Equilibrium

A *sequential equilibrium with recursive planning* in our model is a sequence of policy functions $\{x_{a,t}^*, x_{n,t}^*, b_{a,t}^*, b_{n,t}^*, k_t^*\}$, a sequence of value functions $\{V_t^a, V_t^n\}$, a sequence of prices $\{w_t, w_t^F,\ \Pi_t^F, \Pi_t^U, q_t, r_t, R_t^b, \pi_t^Y, \pi_t^W\}$, a sequence of stochastic states A_t, Ψ_t, Z_t and shocks $\epsilon_t^R, \epsilon_t^G, \epsilon_t^A, \epsilon_t^Z, \epsilon_t^\Psi, \epsilon_t^{\mu W}, \epsilon_t^{\mu Y}, \epsilon_t^\lambda, \epsilon_t^\sigma$, aggregate capital and labor supplies $\{K_t, N_t\}$, distributions Θ_t over individual asset holdings and productivity, and expectations Γ for the distribution of future prices, such that

1. Given the functional $\mathbb{E}_t V_{t+1}$ for the continuation value and period-t prices, policy functions $\{x_{a,t}^*, x_{n,t}^*, b_{a,t}^*, b_{n,t}^*, k_t^*\}$ solve the households' planning problem, and given the policy functions $\{x_{a,t}^*, x_{n,t}^*, b_{a,t}^*, b_{n,t}^*, k_t^*\}$, prices, and the value functions $\{V_t^a, V_t^n\}$ are a solution to the Bellman equations (9).

2. The distribution of wealth and income evolves according to the households policy functions.

3. The labor, the final goods, the bond, the capital and the intermediate good markets clear in every period, interest rates on bonds are set according to the central bank's Taylor rule, fiscal policy is set according to the fiscal rule, and stochastic processes evolve according to their law of motion.

4. Expectations are model consistent.

2.6 Representative-agent version and other simplified variants

Since one goal of this paper is to compare the incomplete markets, heterogeneous agent model with two assets to other simpler variants, including a representative agent model version that features complete markets, we describe next how these model variants look like. The first variant we consider is a model which only differs in that all assets are perfectly liquid. This implies that the bonds market and the capital market clearing conditions are replaced with the following two Euler equation for bonds and capital respectively:

$$
\begin{aligned}
u'(x_{it}) &\geq \mathbb{E}_t \left[\beta R(R^b_{t+1}, A_{t+1}, b_{it+1})/\pi_{t+1} u'(x_{it+1})\right] && (25)\\
q_t u'(x_{it}) &\geq \mathbb{E}_t \left[\beta (r_{t+1} + q_{t+1}) u'(x_{it+1})\right]. && (26)
\end{aligned}
$$

These hold with equality if the borrowing constraint is not binding. Since we solve the aggregate dynamics of the model using first-order perturbation and at least one household needs to hold both types of assets, we can simplify the condition to the no-arbitrage condition and a single set of individual consumption Euler equations

$$
\begin{aligned}
u'(x_{it}) &\geq \mathbb{E}_t \left[\beta R(R^b_{t+1}, A_{t+1}, b_{it+1})/\pi_{t+1} u'(x_{it+1})\right] && (27)\\
\mathbb{E}_t \frac{R(R^b_{t+1}, A_{t+1})}{\pi_{t+1}} &= \mathbb{E}_t \frac{r_{t+1} + q_{t+1}}{q_t}. && (28)
\end{aligned}
$$

Since households are indifferent between the two assets in equilibrium, we assume that all households hold the same fraction of bonds in their portfolio. The steady state properties of this model have been discussed for example in Aiyagari and McGrattan (1998). In practice, this is a model of only one asset, which we therefore abbreviate in the following as HANK-1. Otherwise the model is identical to our baseline HANK-2 model. Comparing these two model highlights thus the role of portfolio choices.

Secondly, we consider a variant, where there are only two fixed, exogenously given types of households instead of endogenous heterogeneity: a saver and a spender type as in Campbell and Mankiw (1989). Following the nomenclature introduced by Kaplan et al. (2018), we abbreviate this model as “TANK”. The model abstracts from idiosyncratic risk and occasionally binding borrowing constraints. Savers make their intertemporal decisions according to a consumption Euler equation as above, while spenders use all their income (net of taxes) for consumption. We assume all profit income goes to savers. Spenders receive only wage income and union profits. The idea is that this model picks up the fact that some households have higher marginal propensities to consume and aggregate shocks redistribute across households with different marginal consumption propensities without the need to model this endogenously. As a consequence, the model not only abstracts from the portfolio-choice channel but removes also the precautionary savings channel that is still operative in HANK-1. Finally, we use a representative agent model, which features only a single type of agent and complete markets (“RANK”). Also here, a single aggregate consumption Euler equation determines the expected rate of return on bonds and capital.

3 Numerical Solution and Estimation Technique

We solve all model variants by perturbation methods, and choose a first-order Taylor expansion around the stationary equilibrium/steady state. For the model with household heterogeneity, we follow the method of Bayer and Luetticke (2018). This method replaces the value functions by linear interpolants and the distribution functions by histograms to calculate a stationary equilibrium. Then it performs dimensionality reduction before linearization but after calculation of the stationary equilibrium. The dimensionality reduction is achieved by using discrete cosine transformations (DCT) for the value functions and perturbing only the largest coefficients of this transformation[16] and by approximating the joint distributions through distributions with a fixed Copula and flexible marginals. We solve the model originally on a grid of 80x80x22 points for liquid assets, illiquid assets, and income, respectively. The dimensionality-reduced number of states and controls in our system is roughly 900.

[16]Specifically, we proceed in two steps. First, we solve a version of the model using the estimated parameters from the RANK model where we keep the coefficients from the DCT that represent 99.9999% of the entropy of the value functions. Second, we calculate the forecasting variances of all these coefficients at all horizons between 1 and 1000 under this baseline calibration. We then keep the union of all those coefficients that represent 99% of the sum of all variances of the coefficients at every forecast horizon. All coefficients we do not keep are set to their stationary equilibrium values. This means that, in an R^2 sense, we make sure that the value functions as used reflect 99% of the variation of the value functions around the stationary equilibrium at every horizon. This leaves us with roughly 350 coefficients for each marginal-value function which are perturbed out of around 280,000 for the marginal values of liquid and illiquid assets, which are the controls we use instead of working with the value function itself.

Approximating the sequential equilibrium in a linear state-space representation then boils down to the linearized solution of a non-linear difference equation

$$\mathbb{E}_t F(x_t, X_t, x_{t+1}, X_{t+1}, \sigma\Sigma\epsilon_{t+1}), \tag{29}$$

where x_t are "idiosyncratic" states and controls: the value and distribution functions, and X_t are aggregate states and controls: prices, quantities, productivities, etc. The error term ϵ_t represents fundamental shocks. Importantly, we can also order the equations in a similar way. The law of motion for the distribution and the Bellman equations describe a non-linear difference equation for the idiosyncratic variables, and all other optimality and market clearing conditions describe a non-linear difference equation for the aggregate variables. By introducing auxiliary variables that capture the mean of b, k, and h, we make sure that the distribution itself does not show up in any aggregate equation other than in the one for the summary variables. Yet, these equations are free of all model parameters.

This helps substantially in estimating the model. For each parameter draw, we need to calculate the Jacobian of F and then use the Klein (2000)-algorithm (see also Schmitt-Grohé and Uribe, 2004) to obtain a linear state-space representation,[17] which we then feed into a Kalman filter to obtain the likelihood of the data given our model. However, most model parameters do not show up in the Bellman equation. Only $\rho_h, \bar{\sigma}_h, \bar{\lambda}, \beta, \gamma$, and ξ do, but these parameters we calibrate from the stationary equilibrium already.[18] Therefore, the Jacobian of the "idiosyncratic equations" is unaltered by all parameters that we estimate and we only need to calculate it once. Similarly, "idiosyncratic variables" (i.e. the value functions and the histograms) only affect the aggregate equations through their parameter free effect on summary variables, such that also this part of the Jacobian does not need to get updated during the estimation. This leaves us with the same number of derivatives to be calculated for every parameter draw during the estimation as in a representative agent model. Still, solving for the state-space representation and evaluating the likelihood is substantially more time consuming and computing the likelihood of a given parameter draw takes roughly 4 to 5 seconds on a workstation computer, 90% of the computing time goes into the Schur decomposition, which still is much larger because of the many additional "idiosyncratic" states (histograms) and controls (value functions) the system contains.

We use a Bayesian likelihood approach as described in An and Schorfheide (2007) and Fernández-Villaverde (2010) for parameter estimation. In particular, we use the Kalman

[17] We also experimented with the Anderson and Moore (1985) algorithm. While it is more than twice as fast as Klein's method for the HANK model with two assets in many cases, it appears to produce less numerically stable results in a setting such as ours, where the Jacobians are not very sparse.

[18] Note that the scaling of idiosyncratic risk, s_t, like the liquidity of assets, λ_t, shows up in the Bellman equation, but similar to a price and not as a parameter.

filter to obtain the likelihood from the state-space representation of the model solution[19] and employ a standard *Random Walk Metropolis-Hastings* algorithm to generate draws from the posterior likelihood. Smoothed estimates of the states at the posterior mean of the parameters are obtained via a Kalman smoother of the type described in Koopman and Durbin (2000) and Durbin and Koopman (2012).

4 Calibration, Priors, and Estimated Parameters

One period in the model refers to a quarter of a year. Tables 1 summarizes the calibrated and externally chosen parameters and columns 2-4 of Table 3 list the prior distributions of the estimated parameters.

4.1 Calibrated Parameters

We fix a number of parameters either following the literature or targeting steady-state ratios; see Table 1 (all at quarterly frequency of the model). For the household side, we set the relative risk aversion to 4, which is common in the incomplete markets literature; see Kaplan et al. (2018), and the Frisch elasticity to 0.5; see Chetty et al. (2011). We take estimates for idiosyncratic income risk from Storesletten et al. (2004), $\rho_h = 0.98$ and $\bar{\sigma}_h = 0.12$. Guvenen et al. (2014a) provide the probability to fall out of the top 1% of the income distribution in a given year, which we take as transition probability from entrepreneur to worker, $\iota = 1/16$.

Table 2 summarizes the calibration of the remaining household parameters. We match 4 targets: 1) Mean illiquid assets (K/Y=12.68) equals Fixed Assets and Durables over quarterly GDP (excluding net exports) averaged over 1954-2015 (NIPA tables 1.1 and 1.1.5). 2) Mean liquidity (B/Y=2.76) equals average Federal Debt over quarterly GDP from 1966-2015 (FRED: GFDEBTN). 3) The fraction of borrowers, 16%, is taken from the Survey of Consumer Finances (1983-2013); see Bayer et al. (2019) for more details. Finally, the average top-10% share of wealth from 1954-2015, which is 69%, comes from the World Inequality Database. This yields a discount factor of 0.98, a portfolio adjustment probability of 6.5%, borrowing penalty of 0.74% (given a borrowing limit of one time average income), and a transition probability from worker to entrepreneur of 1/5000.

[19] The Kalman filter allows us to deal with missing values and mixed frequency data quite naturally. For a one-frequency dataset without missing values, one can speed up the estimation by employing so-called "Chandrasekhar Recursions" for evaluating the likelihood. These recursions replace the costly updating of the state variance matrix by multiplications involving matrices of much lower dimension (see Herbst, 2014, for details). This is especially relevant for the HANK-2 model as the speed of evaluating the likelihood is dominated by the updating of the state variance matrix, which involves the multiplication of matrices that are quadratic in the number of states.

Table 1: External/Calibrated parameters (quarterly frequency)

Parameter	Value	Description	Target
Households			
β	0.98	Discount factor	see Table 2
ξ	4	Relative risk aversion	Kaplan et al. (2018)
γ	2	Inverse of Frisch elasticity	Chetty et al. (2011)
λ	0.065	Portfolio adj. prob.	see Table 2
ρ_h	0.98	Persistence labor income	Storesletten et al. (2004)
σ_h	0.12	STD labor income	Storesletten et al. (2004)
ζ	1/5000	Trans.prob. from W. to E.	see Table 2
ι	1/16	Trans.prob. from E. to W.	Guvenen et al. (2014a)
$\bar{R}$	0.74%	Borrowing penalty	see Table 2
Firms			
α	0.68	Share of labor	62% labor income
δ_0	1.6%	Depreciation rate	NIPA
$\bar{\eta}$	11	Elasticity of substitution	Price markup 10%
$\bar{\zeta}$	11	Elasticity of substitution	Wage markup 10%
Government			
τ	0.3	Tax rate	$G/Y = 20\%$
$\bar{R}^b$	1.004	Nominal rate	1.6% p.a.
$\bar{\pi}$	1.00	Inflation	0% p.a.

For the firm side, we set the labor share in production, α, to 68% to match a labor income share of 62%, which corresponds to the average BLS labor share measure over 1954-2015. The depreciation rate is 1.6% per quarter (NIPA tables 1.1 and 1.1.3). An elasticity of substitution between differentiated goods of 11 yields a markup of 10%. The elasticity of substitution between labor varieties is also set to 11, yielding a wage markup of 10%. Both are standard values in the literature.

The government levies a proportional tax rate of 30% on labor and profit income. This corresponds to a government share of $G/Y = 20\%$. The policy rate is set to 1.6% annualized rate. This corresponds to the average Federal Funds Rate in real terms over 1954-2015. We set steady state inflation to zero as we have assumed indexation to the steady state inflation rate in the Phillips curves.

4.2 Estimation Data

We use quarterly US data from 1954Q3 to 2015Q4 and include the following seven observable time series: the growth rates of per capita GDP, private consumption, investment, and wages, all in real terms, the logarithm of the level of per capita hours worked, the log difference

Table 2: Calibrated parameters

Targets	Model	Data	Source	Parameter
Mean illiquid assets (K/Y)	12.68	12.68	NIPA	Discount factor
Mean liquidity (B/Y)	2.76	2.76	FRED	Port. adj. probability
Top10 wealth share	0.69	0.69	WID	Fraction of entrepreneurs
Fraction borrowers	0.16	0.16	SCF	Borrowing penalty

of the GDP deflator, and the (shadow) federal funds rate. We augment the dataset with further data with shorter and/or non-quarterly availability. Idiosyncratic income uncertainty (estimated in Bayer et al., 2019) (1983Q1-2013Q1) and the month supply of homes as proxy of liquidity (1963Q1-2015Q4) are available as quarterly series and included in log-levels. Wealth and income shares of the top 10 are included at annual frequency and available from 1954 to 2015 from the World Inequality Database (drawing on work from Piketty, Saez, and Zucman; see, e.g., Alvaredo et al. (2017)).[20]

4.3 Prior and Posterior Distributions

Table 3: Prior and Posterior Distributions of Estimated Parameters

Parameter		Prior		Posterior			
	Distribution	Mean	Std. Dev.	Mean	Std. Dev.	5 Percent	95 Percent
				Frictions			
ϕ	Gamma	4.00	2.00	0.334	0.026	0.289	0.377
δ_2/δ_1	Gamma	5.00	2.00	0.165	0.029	0.118	0.215
κ	Gamma	0.10	0.02	0.067	0.009	0.053	0.083
κ_w	Gamma	0.10	0.02	0.170	0.024	0.133	0.210
				Fiscal and monetary policy rules			
ρ_R	Beta	0.50	0.20	0.749	0.017	0.720	0.776
σ_R	Inv.-Gamma	0.10	2.00	0.253	0.015	0.231	0.278
θ_π	Normal	1.70	0.30	1.934	0.053	1.850	2.024

[20]Detailed data sources and the observation equation that describes how the empirical time series are matched to the corresponding model variables can be found in Appendix A.

Table 3: Prior and Posterior Distributions of Estimated Parameters - Continued

Parameter		Prior		Posterior			
	Distribution	Mean	Std. Dev.	Mean	Std. Dev.	5 Percent	95 Percent
θ_y	Normal	0.13	0.05	0.422	0.025	0.381	0.464
ρ_B	Beta	0.50	0.20	0.983	0.003	0.977	0.987
ρ_G	Beta	0.50	0.20	0.990	0.006	0.978	0.997
σ_G	Inv.-Gamma	0.10	2.00	0.169	0.010	0.154	0.186
γ_Y	Normal	0.00	1.00	-0.168	0.012	-0.189	-0.149

Table 3: Prior and Posterior Distributions of Estimated Parameters - Continued

Parameter		Prior		Posterior			
	Distribution	Mean	Std. Dev.	Mean	Std. Dev.	5 Percent	95 Percent
				Structural Shocks			
ρ_A	Beta	0.50	0.20	0.996	0.002	0.992	0.999
σ_A	Inv.-Gamma	0.10	2.00	0.188	0.010	0.172	0.205
ρ_Z	Beta	0.50	0.20	0.977	0.006	0.967	0.987
σ_Z	Inv.-Gamma	0.10	2.00	0.645	0.030	0.599	0.697
ρ_Ψ	Beta	0.50	0.20	0.997	0.002	0.994	0.999
σ_Ψ	Inv.-Gamma	0.10	2.00	1.428	0.079	1.302	1.560
ρ_μ	Beta	0.50	0.20	0.990	0.005	0.980	0.997
σ_μ	Inv.-Gamma	0.10	2.00	1.328	0.084	1.197	1.473
$\rho_{\mu w}$	Beta	0.50	0.20	0.872	0.020	0.837	0.902
$\sigma_{\mu w}$	Inv.-Gamma	0.10	2.00	4.666	0.457	3.993	5.476
				Risk and Liquidity Process			
ρ_s	Beta	0.50	0.20	0.643	0.029	0.593	0.687
σ_s	Inv.-Gamma	1.00	2.00	85.23	5.951	75.708	95.004
Σ_N	Normal	0.00	1.00	-0.521	0.030	-0.569	-0.471
ρ_λ	Beta	0.50	0.20	0.901	0.018	0.870	0.930
σ_λ	Inv.-Gamma	1.00	2.00	8.847	0.429	8.171	9.572
Λ_R	Normal	0.00	1.00	-0.626	0.030	-0.674	-0.574
				Measurement Errors			
σ_λ^{me}	Inv.-Gamma	0.05	0.10	0.040	0.027	0.012	0.097

Notes: The standard deviations of the shocks and measurement errors have been transformed into percentages by multiplying with 100. Posterior estimates are for the HANK-2 model without observable inequality series.

Columns 2-4 of Table 3 present the initial prior distributions. The posterior distribution is discussed in the next section, Section 5.1. Where available, we use prior values that are standard in the literature and independent of the underlying data. Following Smets and Wouters (2007), the autoregressive parameters of the shock processes are assumed to follow a beta distribution with mean 0.5 and standard deviation 0.2. The standard deviations of the shocks follow inverse-gamma distributions with prior mean 0.1 percent and standard deviation 2 percent. The only exceptions are the uncertainty and liquidity shocks, where we use a prior mean of 1.0 percent, and the measurement error, for which we assume an inverse-gamma prior with a lower prior mean of 0.05 percent. The employment and interest feedback parameters in the uncertainty and liquidity processes are assumed to follow Standard Normal priors. The autoregressive and feedback parameters in the bond rule, ρ_B and γ_Y, are assumed to follow Beta (with mean 0.5 and standard deviation 0.2) and Standard Normal distributions, respectively. For the inflation and output feedback parameters in the Taylor-rule, θ_π and θ_Y, we impose normal distributions with prior means of 1.7 and 0.13, respectively, while the interest rate smoothing parameter ρ_R has the same prior distribution as the persistence parameters of the shock processes. Following Justiniano et al. (2011), we impose a Gamma distribution with prior mean of 5.0 and standard deviation of 2.0 for δ_2/δ_1, the elasticity of marginal depreciation with respect to capacity utilization, and a Gamma prior with mean 4.0 and standard deviation of 2.0 for the parameter controlling investment adjustment costs, ϕ. For the slopes of price and wage Phillips curves, κ and κ_w, we assume Gamma priors with mean 0.1 and standard deviation 0.02, which corresponds to price and wage contracts having an average length of one year if adjustment costs were Calvo.

5 US Business Cycles

In this section, we compare parameter estimates, variance decompositions, and historical decompositions of US business cycles across the estimated RANK and HANK model. We postpone the model implications for US inequality to the next section, Section 6.

5.1 Parameter Estimates Across Models

Table 4 reports the posterior distributions across the two main model variants (RANK, HANK-2).[21] By and large, the parameter estimates are very similar; only few estimated parameters are significantly different across the two models.[22] First and foremost, we estimate real rigidities (investment adjustment costs and utilization) to be smaller and nominal rigidities to be somewhat larger using the HANK model rather than RANK – the estimated frequency of price adjustment changes from roughly 4 to roughly 5 quarters if κ is inter-

[21] The estimation is conducted with 5 parallel RWMH chains started from an over-dispersed target distribution after an extensive mode search. After burn in, 150000 draws from the posterior are used to compute the posterior statistics. The average acceptance across chains is 21.46%. Appendix E provides visual and statistical convergence checks.

[22] Appendix B shows that it is the response to idiosyncratic uncertainty, the medium term response of the real rate, and the imperfect crowding out of government debt and capital that discriminates the models. We do so by adding also the saver-spender model (TANK) and a one asset HANK model to the picture.

preted as if from a Calvo adjustment cost. Second, the HANK model views the Fed's policy as more reactive to output than what the RANK model infers and finds fiscal policy to be less shock driven. Third, the HANK model estimates shocks to price-markups and investment efficiency to be more persistent than the RANK model. Finally, the HANK model estimate the feedback coefficients of idiosyncratic uncertainty and liquidity are such that they amplify fluctuations. Uncertainty goes up when employment falls and liquidity goes down, when interest rates rise. However, the feedback coefficients are small compared to the variance of uncertainty and liquidity. As a result, the feedback is negligible in economic terms; see Appendix C for the historical uncertainty and liquidity time series implied by the model.

Table 4: Comparison of model estimates

Parameter	RANK	HANK-2		RANK	HANK-2
	Real Frictions			Nominal Frictions	
ϕ	0.517	0.334	κ	0.110	0.067
	(0.484, 0.549)	(0.289, 0.377)		(0.088, 0.133)	(0.053, 0.083)
δ_2/δ_1	0.759	0.165	κ_w	0.166	0.170
	(0.625, 0.896)	(0.118, 0.215)		(0.126, 0.211)	(0.133, 0.210)
	Monetary policy rules			Fiscal policy rules	
ρ_R	0.751	0.749	ρ_B	0.996	0.983
	(0.723, 0.776)	(0.720, 0.776)		(0.992, 0.999)	(0.977, 0.987)
σ_R	0.248	0.253	ρ_G	0.982	0.990
	(0.228, 0.270)	(0.231, 0.278)		(0.968, 0.994)	(0.978, 0.997)
θ_π	1.828	1.934	σ_G	0.246	0.169
	(1.682, 1.972)	(1.850, 2.024)		(0.221, 0.273)	(0.154, 0.186)
θ_y	0.267	0.422	γ_Y	-0.288	-0.168
	(0.206, 0.328)	(0.381, 0.464)		(-0.322, -0.256)	(-0.189, -0.149)
	Risk process			Liquidity process	
ρ_s	0.799	0.643	ρ_λ	0.924	0.901
	(0.720, 0.874)	(0.593, 0.687)		(0.889, 0.958)	(0.870, 0.930)
σ_s	58.905	85.23	σ_λ	8.823	8.847
	(53.103, 65.231)	(75.708, 95.004)		(8.144, 9.562)	(8.171, 9.572)
Σ_N	-0.011	-0.521	Λ_R	-0.221	-0.626
	(-0.065, 0.046)	(-0.569, -0.471)		(-0.373, -0.069)	(-0.674, -0.574)

Table 4: Comparison of model estimates - Continued

Parameter	RANK	HANK-2		RANK	HANK-2
			Structural Shocks		
ρ_A	0.992	0.996	ρ_μ	0.880	0.990
	(0.983, 0.997)	(0.992, 0.999)		(0.850, 0.907)	(0.980, 0.997)
σ_A	0.116	0.188	σ_μ	1.649	1.328
	(0.099, 0.132)	(0.172, 0.205)		(1.438, 1.892)	(1.197, 1.473)
ρ_Z	0.999	0.977	$\rho_{\mu w}$	0.863	0.872
	(0.997, 1.000)	(0.967, 0.987)		(0.816, 0.905)	(0.837, 0.902)
σ_Z	0.512	0.645	$\sigma_{\mu w}$	5.161	4.666
	(0.474, 0.554)	(0.599, 0.697)		(4.347, 6.193)	(3.993, 5.476)
ρ_Ψ	0.969	0.997			
	(0.955, 0.982)	(0.994, 0.999)			
σ_Ψ	2.418	1.428			
	(2.217, 2.631)	(1.302, 1.560)			
	Measurement Errors				
σ_λ^{me}	0.044	0.040			
	(0.012, 0.115)	(0.012, 0.097)			

Notes: Parentheses contain the 90% highest posterior density interval. The standard deviations of the shocks and measurement errors have been transformed into percentages by multiplying with 100.

5.2 Variance Decompositions

Next, we ask, if and how the differences in internal propagation and in estimated parameters across models change our view of US business cycles by looking at variance decompositions at business cycle frequency. Figure 1 shows these decompositions for the growth rates of output, consumption, investment, and government spending. Again we find, by and large, similar decompositions across the models with important differences in the details though.

We summarize these differences here and in the following in terms of lumping together demand side shocks (i.e. shocks to liquidity, uncertainty, government spending, monetary policy and the risk premium) and supply side shocks (the two markup and the two productivity shocks). The grouping of the shocks is based on the question whether they directly affect the Phillips curve as the relevant aggregate supply function or primarily work through affecting the bond-market clearing condition, as the aggregate demand function. In terms

Figure 1: Variance Decompositions: Output growth and its components

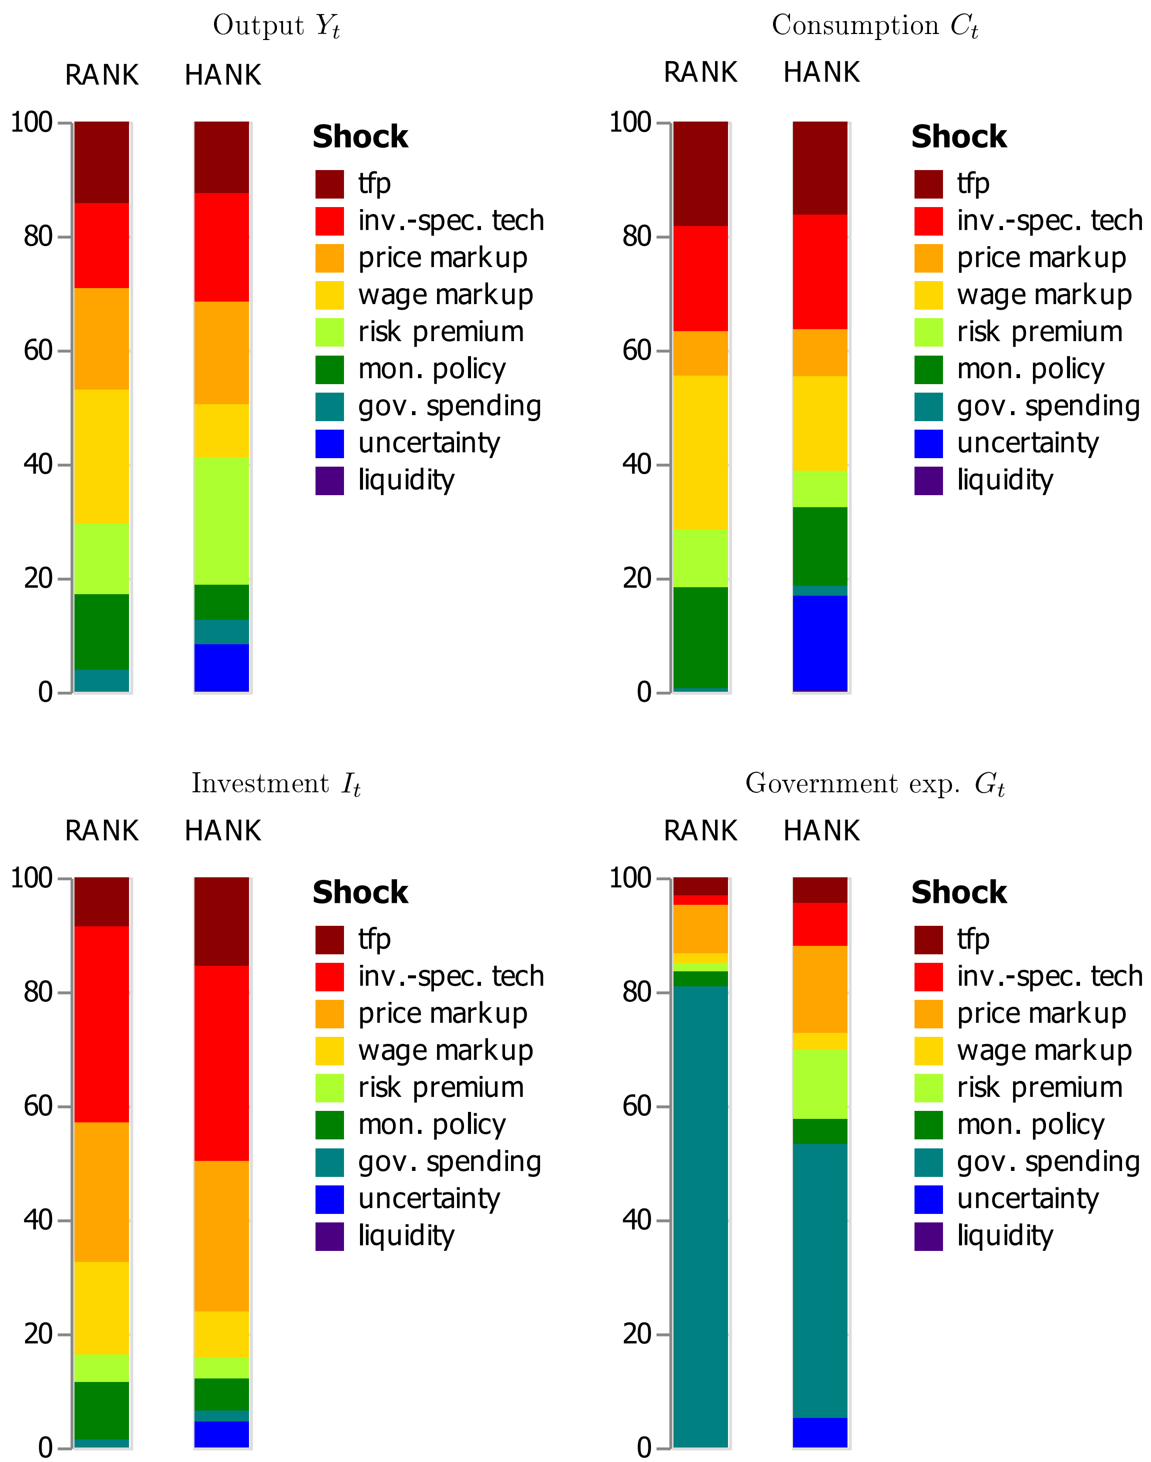

Notes: Conditional variance decompositions at a 4-quarter forecast horizon.

Figure 2: Variance Decompositions: Inflation and Policy Rate

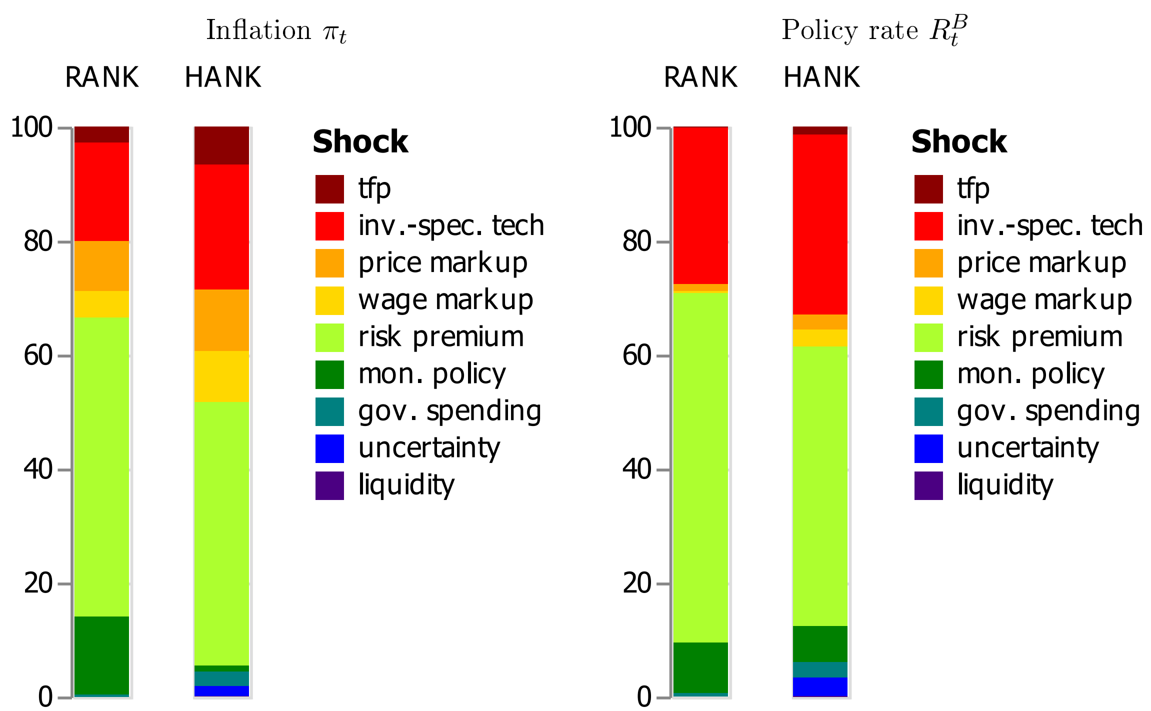

Notes: Conditional variance decompositions at a 4-quarter forecast horizon.

of this grouping, the HANK model views demand side shocks as more important for output, consumption and investment than does the RANK model. One key reason for this is that uncertainty shocks enter as a new additional driver of the business cycle. The differences are the strongest for consumption where shocks to income risk alone explain almost 20% of consumption volatility in the HANK model. This increases the importance of demand shocks relative to RANK by 30% for consumption and similarly for output.

Income risk shocks are an important driver of consumption, because income risk is mostly exogenous. The estimated endogenous feedback parameter shows that income risk goes up in recessions, but the endogenous feedback effect is small. Fluctuations in the liquidity of assets plays only a minor role even though we find that a decrease in liquidity can lead to a contraction in the HANK model; see Appendix B.2. Yet, the empirical fluctuations of liquidity are too small for it to be a substantive contributer to the cycle. As the HANK model estimates a much smaller variance of government expenditure shocks while the response rate to other variables is estimated to be of roughly the same size as in the RANK model, we find that government expenditures appear to be much more systematic and driven by other

Figure 3: Historical Decompositions: Output Growth - Demand vs Supply

RANK	HANK

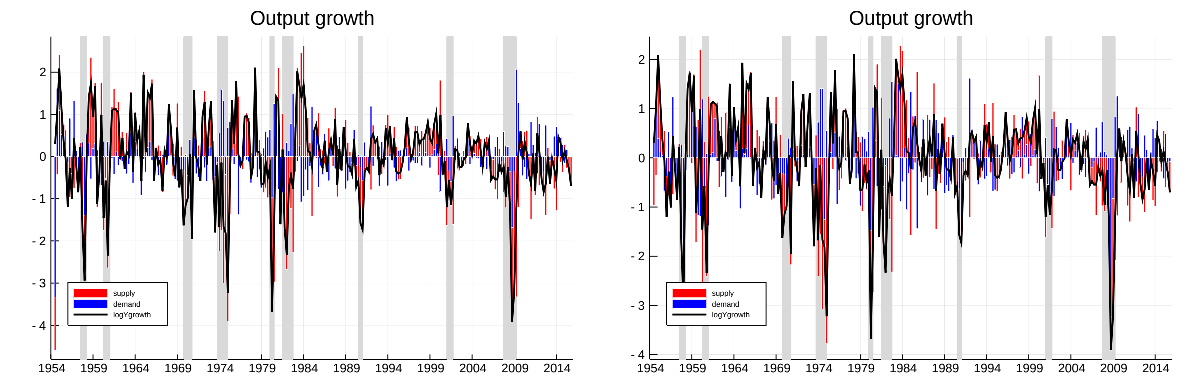

Notes: Demand shocks is the sum of risk premium, monetary, fiscal, income risk, and liquidity shocks. Supply shocks is the sum of TFP, price markup, wage markup, and investment-specific technology shocks.

shocks to the economy.

In terms of nominal variables, we find the opposite result as demand shocks become less important. Figure 2 shows the variance decomposition of inflation and the policy rate across both models. Here we find that supply side shocks are more important drivers of the interest rate and inflation in HANK than in RANK. In particular, the risk premium shock becomes less important for the nominal side. This reflects that the HANK model estimates monetary policy to be more reactive to output fluctuations.

In summary, the estimated HANK model changes our view on the average business cycle relative to the RANK model in that income risk fluctuations increase the importance of demand shocks for aggregate quantities and in particular for aggregate consumption. At the same time, inflation and the policy rate are driven to a larger extend by supply-side shocks.

5.3 Historical Decompositions

While the variance decompositions help us understand the average cycle implied by the model, a historical decomposition tells us how the two models view the actual cycles that the US economy has gone through differently.

Figure 3 starts by summarizing the decomposition of output growth into demand and supply shocks. Figure 4 plots the contribution of the various shocks both for growth rates and levels separately. We report historical decompositions for consumption, investment, and

Table 5: Contribution of shocks to US recessions

	Supply Shocks			Demand Shocks	
Shock	RANK	HANK		RANK	HANK
TFP, ϵ^Z	-0.12	-0.28	Risk premium, ϵ^A	-0.24	-0.42
Inv.-spec. tech, ϵ^Ψ	-0.32	-0.01	Mon. policy, ϵ^R	0.15	0.06
Price markup, $\epsilon^{\mu Y}$	-0.12	-0.05	Fisc. policy, ϵ^G	0.02	0.13
Wage markup, $\epsilon^{\mu W}$	-0.37	-0.24	Uncertainty, ϵ^σ	0.00	-0.21
			Liquidity, ϵ^λ	0.00	0.02
Sum of shocks	-0.94	-0.58	Sum of shocks	-0.07	-0.42

Notes: The table displays the average contribution of the various shocks during an NBER-dated recession that result from our historical shock decomposition. Values are calculated by averaging the value of each shock component over all NBER recession quarters. To improve readability, we normalized the size of the overall contraction to -1%. In the data, the average is -1.24%.

government expenditures in Appendix C.

In the historical decompositions again there is an apparently larger role of demand side shocks, in particular of income risk. Also in terms of levels, i.e. in terms of accumulated shocks, the RANK and HANK models paint similar picture with important differences in details. For example, the HANK model views the long expansion of the Great Moderation period even more strongly characterized by liberalization, i.e. falling markups, making up for slower productivity growth than the RANK model does. In general, the HANK model finds slightly larger shocks that just happen to offset each other in comparison to the RANK model.

As the graphs are potentially hard to read, given the many quarters of data, we summarize the historical decomposition of NBER dated recessions in Table 5. We find a substantially larger role for demand shocks in US recessions through the lens of the HANK model compared to the RANK model. In RANK, 0.93% of a 1.0% decline in output results from technology and mark-up shocks, while the HANK model suggests that 0.58% out of a 1.0% decline result from risk-premium and uncertainty shocks during the average NBER-dated recession quarter.

Figure 4: Historical Decompositions: Output Growth and Levels

RANK	HANK

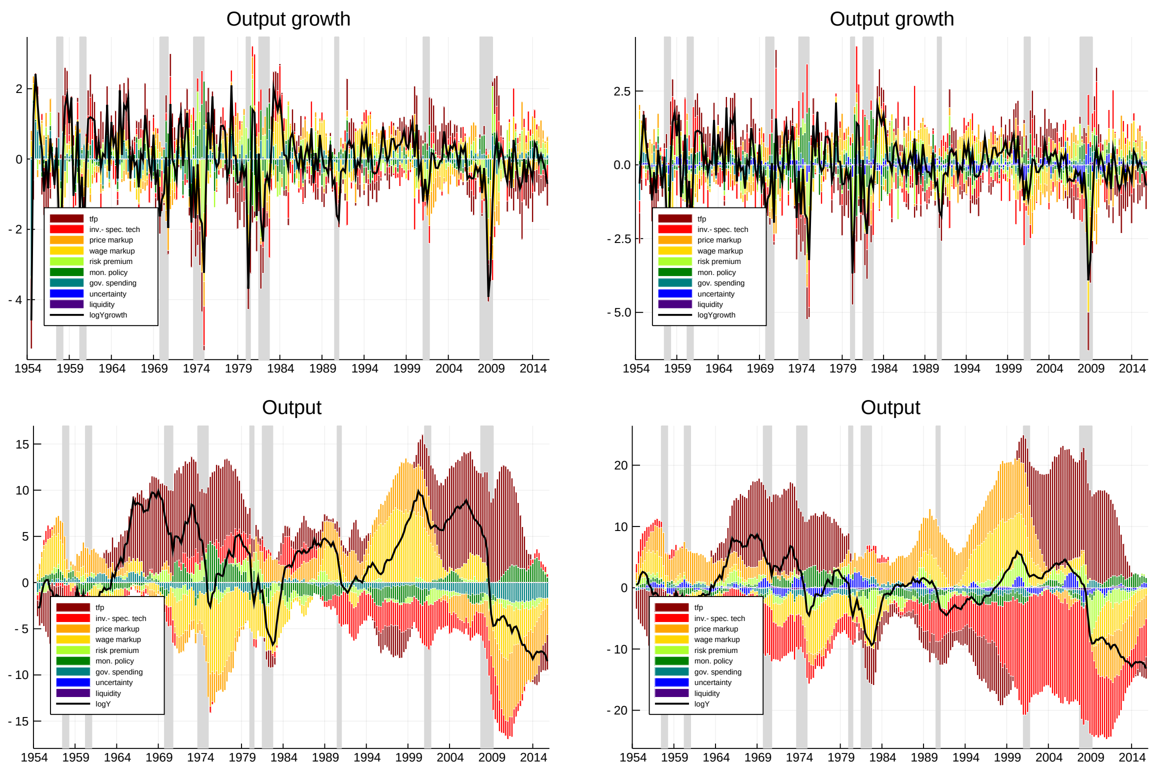

Notes: The top panel shows the historical decomposition of output growth into the contribution of various shocks. The bottom panel shows the same for output levels. The left column is for the RANK estimates the right column for the HANK estimates. The contribution of the smoothed initial state has been omitted.

Figure 5: US Inequality – Data vs Model

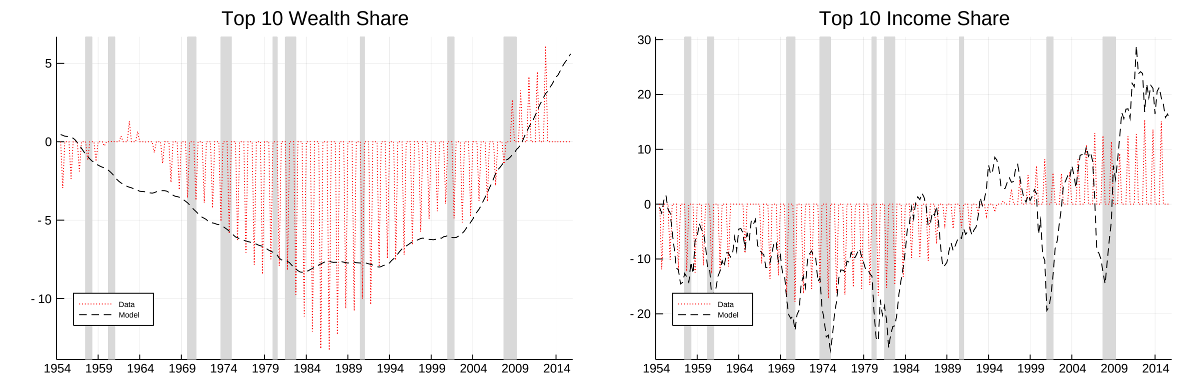

Notes: Data (red) corresponds to log-deviations of the annual observations of the share of income and wealth held by the top 10% in each distribution in the US taken from the World Inequality Database. Model (black) corresponds to the smoothed states of both implied by the estimated HANK model.

6 US Inequality

One key advantage of HANK models is that we can use them to understand the distributional consequences of business cycle shocks and policies. This raises three questions. First, to what extent do business cycle shocks explain the movements in inequality measures? Second, does the inclusion of measures of inequality change what the model infers about shocks and frictions in business cycles? Third, how would inequality have developed if government business cycle policies had been different.

To answer these questions, we re-estimate the HANK model with additional observables (plus measurement error) for the shares of wealth and income held by the top-10% of households in each dimension, which are taken from the World-Income-Database. The reason we focus on the top 10% wealth and income share is that this measure is most consistent across alternative data sources such as the SCF, where available; see Hubmer et al. (2016).

Figure 5 plots the new data and the model implied smoothed states. Both data series are available on an annual basis throughout our whole sample period 1954-2015. The top-10% wealth and income shares are both U-shaped and trough around 1980 in the data. The model implied top-10% wealth and income shares match the data well. In the data, the top-10% wealth share increases by 20 percentage points from 1980 to 2015, and the model gets 3/4 of this increase. The top-10% income share increases by almost 30 percentage points over the same time period in the data, and the model predicts in increase by 40 percentage points.

This means that with regard to the first question, our answer is affirmative. Business cycle

Figure 6: Variance Decompositions: Output growth and its components

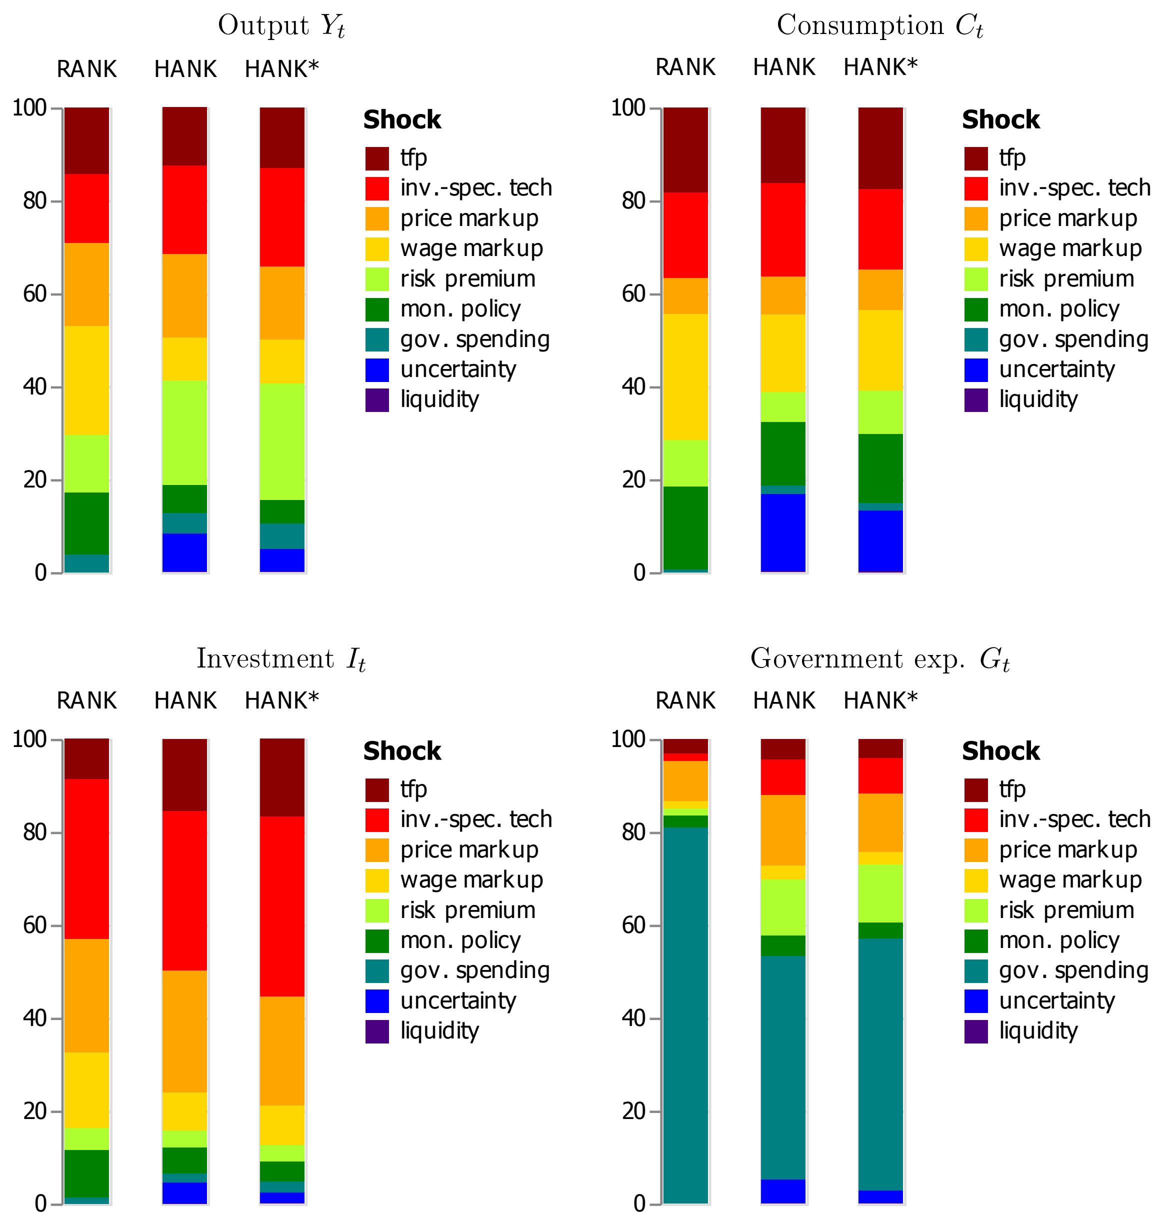

Notes: Conditional variance decompositions at 4 quarters. HANK* refers to the estimation with two additional observables (plus measurement errors): Top-10% wealth and income shares from WID.

Figure 7: Impulse responses of inequality

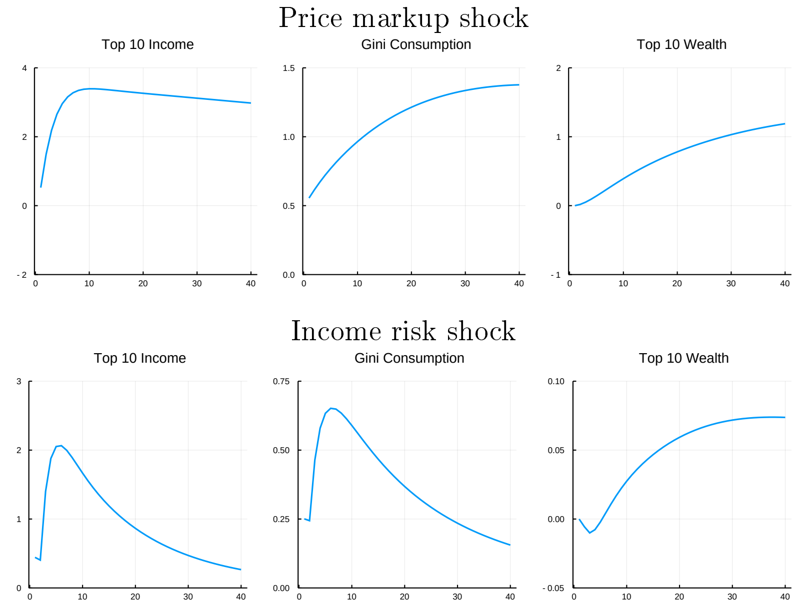

Notes: Response of the top 10% income share, the consumption Gini, and the top-10% wealth share to a one standard deviation shock in the price markups and income risk. The response of the Gini coefficients/Top 10% shares are calculated by including them as a generalized moment in the linearized model.

shocks can move inequality along the lines of what we observe in the data. This matching of the distributional data, on top, does not change significantly what we infer about shocks and frictions. Figure 6 shows that these additional observables, by and large, do not change the estimation results and the corresponding variance decompositions. In the Appendix, Table 7 shows that this holds also true for all individual parameter estimates. This implies that both inequality measures provide little additional identification of business cycle shocks and frictions. The estimated shocks and frictions, on the other hand, do a good job in matching the evolution of wealth and income inequality over the last 60 years.

Why is the model able to explain the slow moving inequality dynamics? Our model implies that business cycle shocks have very persistent effects on the wealth distribution, as Figure 7 shows exemplary for markup and risk shocks. The response to either shock is the least persistent for income inequality, is more persistent for consumption inequality and most persistent for wealth inequality. Consider, for example a price-markup shock. This drives up the income of entrepreneurs, the income richest households in our model. However, because of

Figure 8: Historical Decompositions: Income Inequality

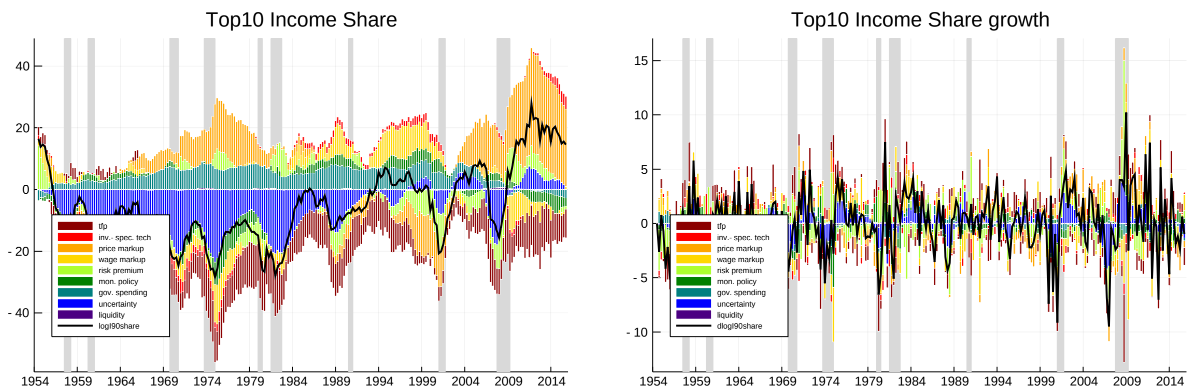

Notes: Historical decomposition of the level and growth rate of the top-10% share in income. The top-10% share is treated as a generalized moment that is included as a control into the state-space representation of the model. The contribution of the initial state has been discarded for the graph.

sticky prices, the increase in inequality is staggered. Therefore, we see initially a greater rise in consumption than in income inequality because entrepreneurs foresee their future incomes increasing and dis-save. Once markups reach their now increased target, entrepreneurs save part of their higher income to smooth consumption. Consequently, consumption inequality peaks later than income inequality, and the rise in markups slowly translates into wealth inequality, which then peaks last. This makes it possible for transitory business cycle shocks to explain persistent deviations in inequality.

To dig into the details of the evolution of inequality, Figure 8 plots the historical decomposition of the top-10% income share in terms of its level and its growth rate. The decomposition of the level of income inequality shows that medium term trends of income inequality primarily result from markup shocks and fluctuations in income risk. The decomposition of growth rates reveals that income risk is also an important driver of income inequality at business cycle frequency, and in the Great Recession in particular.

With respect to particular historical episodes, our decomposition suggests the following. Rising wage markups and low idiosyncratic productivity risks are mainly responsible for the decrease of income inequality throughout the 1960s until the 1970s. The 1980s are seen as a period of liberalization through the lens of our model (both in terms of output cycles and inequality). Wage markups fell, which increased income inequality, but this was partly offset by falling price markups. This picture changes throughout the 1990s but most clearly from the early 2000s onwards. Through the lens of our model, it is larger income risks and sharply increasing price markups that best explain aggregate fluctuations and the sharp rise in income inequality these years have witnessed. Interestingly and despite the

Figure 9: Historical Decompositions: Wealth Inequality

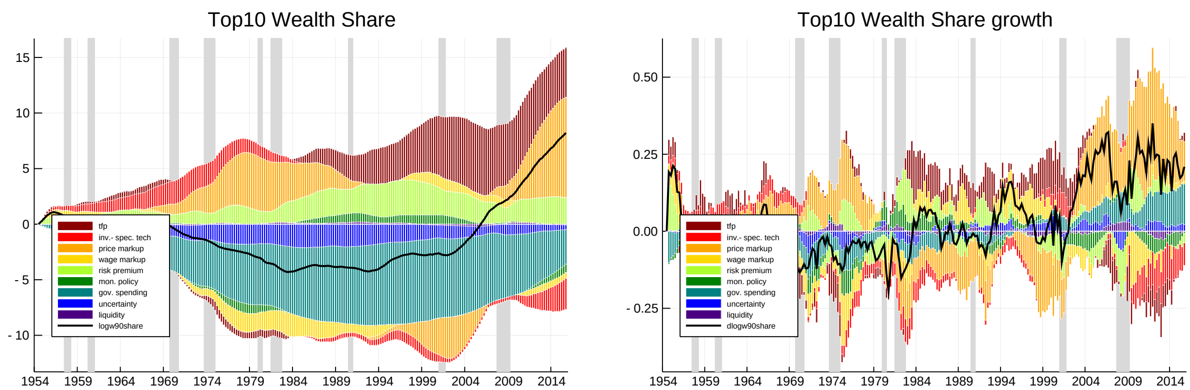

Notes: See Figure 8. Here, the top-10% wealth share is displayed.

use of completely different data sources, the historical decomposition thus is in line with the evidence by De Loecker and Eeckhout (2017) on the evolution of markups in the US.

For the evolution of wealth inequality other shocks are important as well. Figure 9 shows the historical decomposition of the top-10% wealth share in level and growth rate. Wealth inequality fell in the first half of the sample and then increased. The pattern is similar in shape to income inequality, but smoother. Yet, the drivers of wealth inequality are not the same as the drivers of income inequality. The decomposition shows up until the end of the 1970s government expenditure shocks are the strongest downward driver of wealth inequality. From the 1980s on, it is then mainly technological shocks that drive up wealth inequality. A series of positive technological shocks give rise to particular high ex-post returns on illiquid wealth. Only since the 2000s rising price markups have become a strong positive contributor to wealth inequality.

Figure 10 plots the historical decomposition of the Gini coefficient of consumption in level and growth rate. Income risk is the most important driver of short run fluctuations in consumption inequality. The long run trend in consumption inequality is primarily due to TFP and fiscal policy, as well as markups. The reason why income risk is an important driver of consumption inequality lies in the portfolio choice problem of the households. In general, poor households react more strongly to changes in income risk when rebalancing their portfolios (both in the data and in the model, see Bayer et al., 2019). This means that when income risk goes up, the poor more severely cut back consumption to acquire more liquid funds. Therefore, an increase in income risk decreases the consumption of the poor more strongly than the consumption of the wealthy. Table 6 summarizes the driving forces behind the increase in all three inequality measures from 1980 to 2015. Figure 11 shows

Figure 10: Historical Decompositions: Consumption Inequality

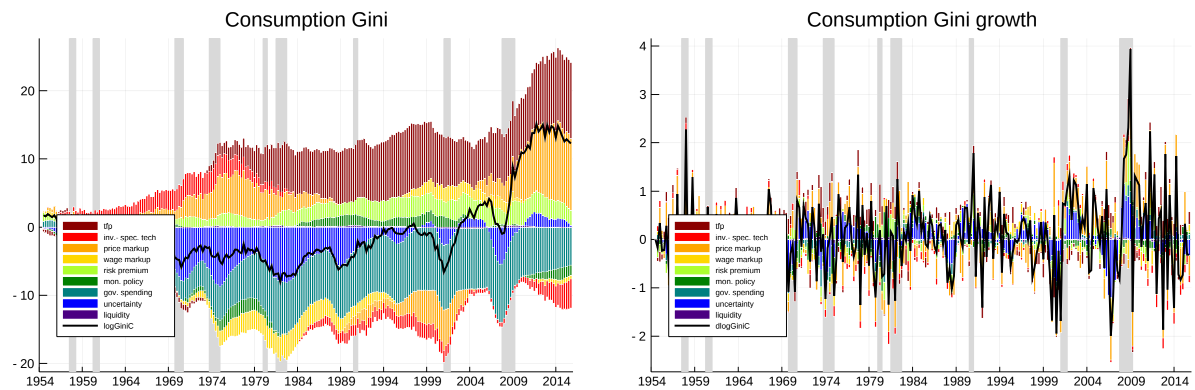

Notes: See Figure 8. Here, the consumption Gini is displayed.

that in general our findings from the historical decompositions also hold true for the average business cycle in terms of variance decompositions.

How important are the estimated policy coefficients for the evolution of inequality? To understand the role of systematic business cycle policies in shaping inequality, we run a set of counterfactual policy experiments based on the estimated model. The results of these experiments are displayed in Figure 12. In detail, the figure displays the difference in the evolution of output, income inequality, wealth inequality and consumption inequality between running the estimated shock sequence through our baseline estimate and through the solution with the counterfactually set policy parameters.

First, we consider an experiment, where the Fed reacts very aggressively to inflation. This creates large output losses after markup shocks, but stabilizes very effectively after demand shocks. Given the series of shocks, output would have been lower in the 70s, and income, wealth, and consumption inequality would have been substantially higher. This reflects that our model attributes a substantial fraction of the fluctuations of the 70s to markup (cost push) shocks. In the 1980s and especially the 1990s, the same policy would have led to higher output, and lower inequality, however, because markups were falling and on top a substantial fraction of shocks during this time are demand shocks. For the period after the Great Recession, which we estimate again to be characterized by a sharp markup increases, the hawkish monetary policy would have delayed the recovery and increased inequality.

Second, we consider a dovish policy, where we triple the monetary policy response to output fluctuations. This leads in general to stabler markups and output at the expense of higher inflation volatility, see also Gornemann et al. (2012). It is not fully the mirror image of the hawkish policy we looked at before because this experiment changes the response

Table 6: Contribution of shocks to US inequality 1980-2015

Shock	Top 10 Wealth	Top 10 Income	Gini Consumption
TFP, ϵ^Z	5.1	6.04	6.07
Inv.-spec. tech, ϵ^Ψ	-3.86	3.52	-5.41
Price markup, $\epsilon^{\mu Y}$	4.2	19.37	6.12
Wage markup, $\epsilon^{\mu W}$	1.03	5.08	2.69
Risk premium, ϵ^A	1.27	3.46	1.19
Mon. policy, ϵ^R	-0.07	-2.12	-0.39
Fisc. policy, ϵ^G	2.38	-10.27	1.89
Uncertainty, ϵ^σ	1.33	10.97	5.17
Liquidity, ϵ^λ	-0.12	-0.07	-0.16
Sum of shocks	12.69	35.14	18.45

Notes: The table displays the contribution (in p.p.) of the various shocks to the increase in inequality from 1980 to 2015 based on our historical shock decompositions.

to output, not inflation, fluctuations. In fact, it exaggerates the output fluctuations of the 1970s, while it does little change to all series in the 1990s. For the Great Recession, a more dovish policy stance would have lead to an earlier recovery and in particular lower income inequality. The effects on wealth inequality are mild.

Finally, we consider a fiscal policy that is more concerned with balancing the budget and keeping the debt level at bay. Here we set the autoregressive coefficient ρ_B to 0.5, which implies a half-live of one quarter instead of the original 35 quarter half-live of a deviation of government debt from its steady state level. This effectively produces a balanced budget at the annual level and eliminates most fluctuations in government debt. The result is that debt would have been lower in the 1960s (it takes out the Vietnam war debt) and this would have led to an increase in wealth inequality because it drives up the liquidity premium (from which wealthy households profit the most). In the 1970s the reverse holds true, where inflation surprises eroded the government debt position. Under our counterfactual, the government uses this for a more expansionary fiscal stance, driving down the liquidity premium. For the expansion of debt under the Reagan administration, we again find that this helped to limit the increase in wealth inequality, while the fiscally “sound” policies of the Clinton administration tended to increase the liquidity premium and thereby wealth inequality.

Figure 11: Variance Decompositions: Inequality

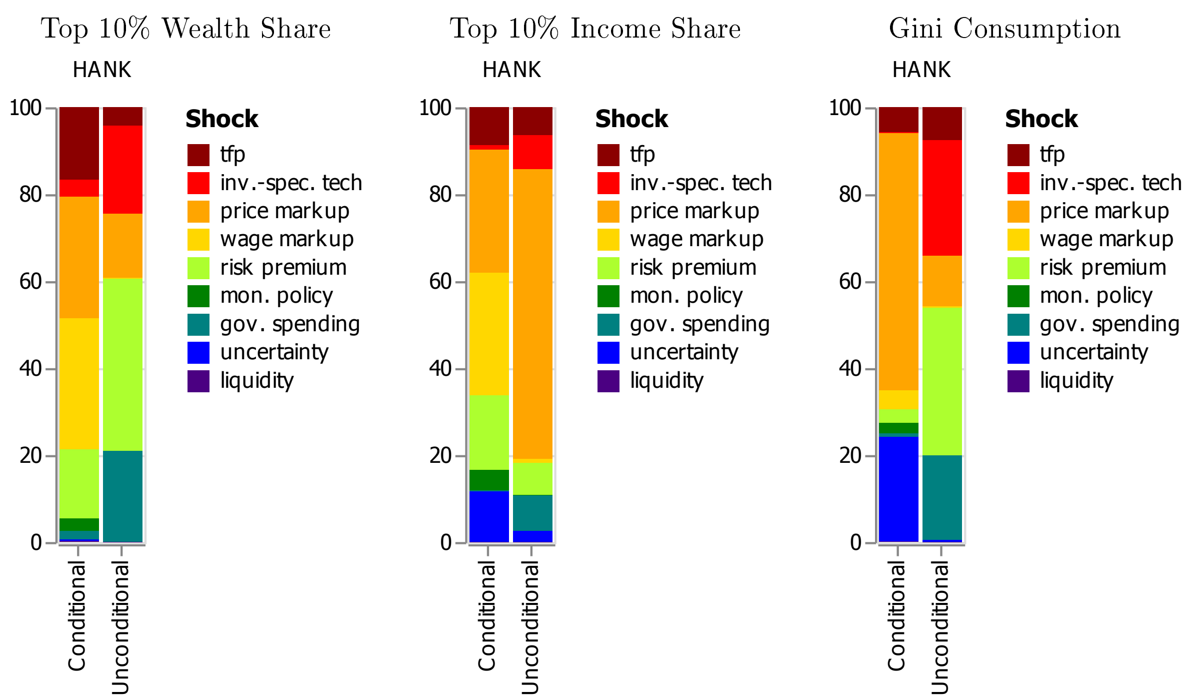

Notes: Conditional refers to a 4 quarter horizon, unconditional refers to a 1000 quarter horizon.

Figure 12: Counterfactual evolution of output and income, wealth, and consumption inequality

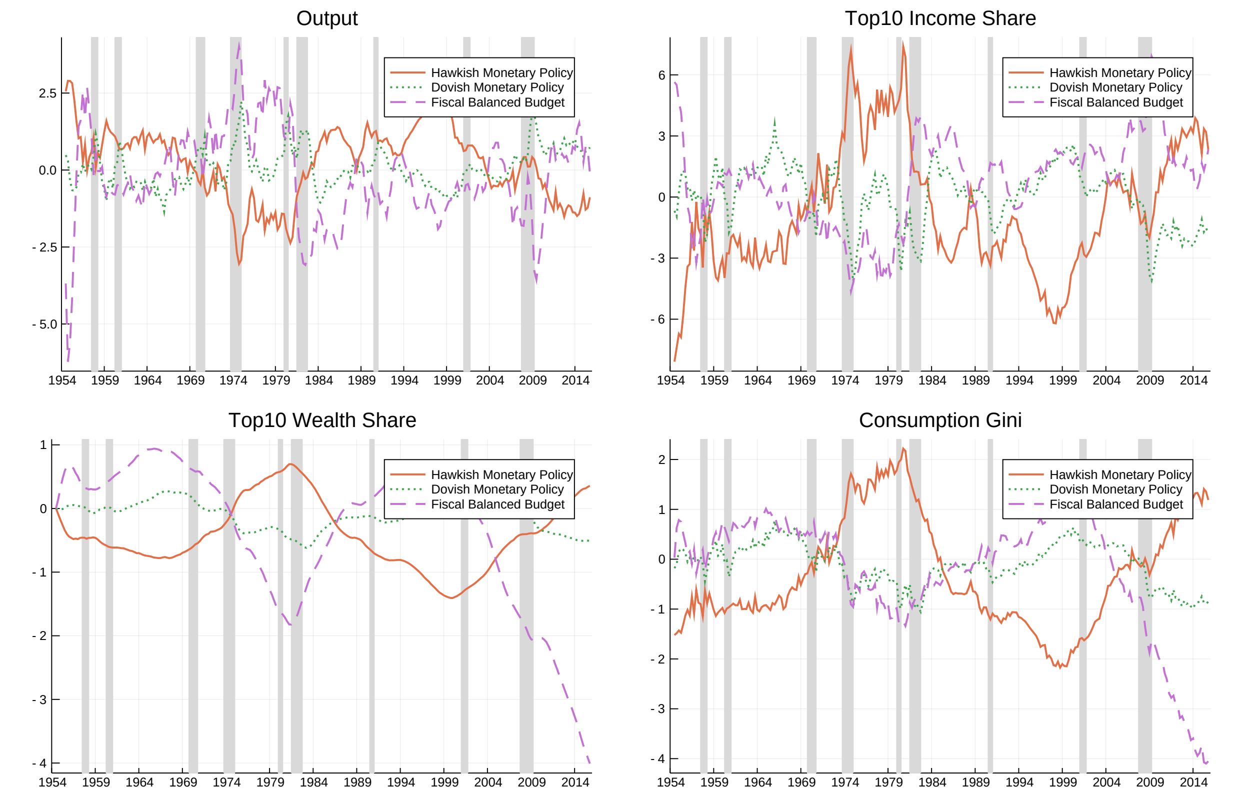

Notes: The panels display the evolution of wealth and income inequality the model would counterfactually predict had the government policies been different, feeding the smoothed sequence of shocks (as in Figures 8 and 9) through the model. The top-left panel displays the evolution of output, the top right panel the evolution of the top 10% income share. The bottom left panel displays the evolution of the top 10 % wealth share, the bottom right panel the evolution of consumption inequality in terms of Gini coefficients. The lines represent the difference in the evolution compared to feeding the same shocks through the baseline model[23] The solid line corresponds to a setup, where we set the inflation response $\theta_\pi = 10$. The dotted line reflects the counterfactual, where we triple the estimated response to output, θ_y. Finally, the dashed dotted line corresponds to a counterfactual, where the fiscal authority seeks more actively to keep debt at its steady state level, setting $\rho_B = 0.5$, effectively balancing the budget at the annual level.

7 Conclusion

In how far does inequality matter for the business cycle and vice versa? To shed light on this two-way relationship, this paper estimates a state-of-the-art New Keynesian business cycle model with household heterogeneity and portfolio choice on macro and micro data. We find household income risk to be an important driver of output and consumption; in particular in US recessions. This strengthens the role of aggregate demand for recessions. Otherwise, we find that household heterogeneity and the inclusion of micro data in the estimation does not materially alter the shocks and frictions in US business cycles.

However, we find that business cycles are important to understand the evolution of US inequality. We show that business cycle shocks and policy responses can account for the increase in US wealth and income inequality since the 1980s. The reason behind this is that wealth (inequality) is a slowly moving variable that accumulates past shocks. Our analysis suggests that price markups have substantially increased over the last two decades. This has driven down output and has increased income, consumption and wealth inequality. A more hawkish monetary policy stance would have exaggerated the increase in markups and hence inequality during and after the Great Recession. The evolution of government debt is of substantial importance as well. An increase of government debt erodes the return difference between illiquid and liquid assets, which helps poor households to accumulate wealth, driving down wealth inequality.

These findings suggest that future research of inequality should take business cycles into account. A synthesis of the previous literature that focuses on permanent changes in the tax and transfer system with the forces that we highlight will be an important area of research. Our findings further suggest to explore the role of shocks that affect household insurance for the business cycle. Including a micro-foundation for income risk, as e.g. via search and matching, is of first order importance to understand how the business cycle and policies work differently by affecting income risk itself.

References

Ahn, S., Kaplan, G., Moll, B., Winberry, T., and Wolf, C. (2018). When inequality matters for macro and macro matters for inequality. *NBER Macroeconomics Annual*, 32(1):1–75.

Aiyagari, S. R. and McGrattan, E. R. (1998). The optimum quantity of debt. *Journal of Monetary Economics*, 42(3):447–469.

Alvaredo, F., Chancel, L., Piketty, T., Saez, E., and Zucman, G. (2017). Global inequality dynamics: New findings from wid. *American Economic Review*, 107(5):404–09.

An, S. and Schorfheide, F. (2007). Bayesian analysis of DSGE models. *Econometric Reviews*, 26(2-4):113–172.

Anderson, G. and Moore, G. (1985). A linear algebraic procedure for solving linear perfect foresight models. *Economics Letters*, 17(3):247–252.

Auclert, A. (2019). Monetary policy and the redistribution channel. *American Economic Review*, 109(6):2333–67.

Auclert, A., Bardóczy, B., Rognlie, M., and Straub, L. (2019a). Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *National Bureau of Economic Research, No. w26123.*

Auclert, A., Rognlie, M., and Straub, L. (2018). The intertemporal keynesian cross. *NBER Working Paper No. 25020.*

Auclert, A., Rognlie, M., and Straub, L. (2019b). Micro jumps, macro humps: monetary policy and business cycles in an estimated hank model. In *2019 Meeting Papers*, number 1449. Society for Economic Dynamics.

Bayer, C. and Luetticke, R. (2018). Solving heterogeneous agent models in discrete time with many idiosyncratic states by perturbation methods. CEPR Discussion Papers 13071.

Bayer, C., Luetticke, R., Pham-Dao, L., and Tjaden, V. (2019). Precautionary savings, illiquid assets, and the aggregate consequences of shocks to household income risk. *Econometrica*, 87(1):255–290.

Born, B. and Pfeifer, J. (2014). Policy risk and the business cycle. *Journal of Monetary Economics*, 68:68–85.

Broer, T., Hansen, N.-J. H., Krusell, P., and Öberg, E. (2016). The New Keynesian Transmission Mechanism: A Heterogenous-Agent Perspective. *NBER Working Paper No. 22418.*

Calvo, G. A. (1983). Staggered prices in a utility-maximizing framework. *Journal of Monetary Economics*, 12(3):383–398.

Campbell, J. Y. and Mankiw, N. G. (1989). Consumption, income and interest rates: Reinterpreting the time series evidence. In *NBER Macroeconomics Annual 1989, Volume 4*, pages 185–246. MIT Press.

Challe, E. and Ragot, X. (2015). Precautionary saving over the business cycle. *The Economic Journal*, 126(590):135–164.

Chang, M., Chen, X., and Schorfheide, F. (2018). Heterogeneity and aggregate fluctuations. *Manuscript, University of Pennsylvania.*

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *The American Economic Review*, 101(3):471–475.

De Loecker, J. and Eeckhout, J. (2017). The rise of market power and the macroeconomic implications. Working Paper 23687, National Bureau of Economic Research.

Den Haan, W. J., Rendahl, P., and Riegler, M. (2017). Unemployment (fears) and deflationary spirals. *Journal of the European Economic Association*, 16(5):1281–1349.

Durbin, J. and Koopman, S. J. (2012). *Time Series Analysis by State Space Methods*. Oxford University Press.

Fernández-Villaverde, J. (2010). The econometrics of DSGE models. *SERIEs*, 1(1):3–49.

Gelman, A. and Rubin, D. B. (1992). Inference from iterative simulation using multiple sequences. *Statistical Science*, 7(4):457–472.

Geweke, J. (1992). Evaluating the accuracy of sampling-based approaches to the calculation of posterior moments. In Bernardo, J. M., Berger, J. O., Dawid, A. P., and Smith, A. F. M., editors, *Bayesian Statistics*, volume 4, pages 641–649. Clarendon Press, Oxford.

Gilchrist, S. and Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4):1692–1720.

Gornemann, N., Kuester, K., and Nakajima, M. (2012). Monetary policy with heterogeneous agents. *Philadelphia Fed Working Paper 12-21.*

Greenwood, J., Hercowitz, Z., and Huffman, G. W. (1988). Investment, capacity utilization, and the real business cycle. *American Economic Review*, 78(3):402–417.

Guerrieri, V. and Lorenzoni, G. (2017). Credit crises, precautionary savings, and the liquidity trap. *The Quarterly Journal of Economics*, 132(3):1427–1467.

Guvenen, F., Kaplan, G., and Song, J. (2014a). How risky are recessions for top earners? *The American Economic Review*, 104(5):148–153.

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

Hagedorn, M., Manovskii, I., and Mitman, K. (2018a). The fiscal multiplier.

Hagedorn, M., Manovskii, I., and Mitman, K. (2018b). Monetary policy in incomplete market models: Theory and evidence. In *University of Pennsylvania Working Paper*.

Hedlund, A. (2016). Illiquidity and its discontents: Trading delays and foreclosures in the housing market. *Journal of Monetary Economics*, 83:1–13.

Herbst, E. (2014). Using the "Chandrasekhar Recursions" for likelihood evaluation of DSGE models. *Computational Economics*, 45(4):693–705.

Hubmer, J., Krusell, P., and Smith Jr, A. A. (2016). The historical evolution of the wealth distribution: A quantitative-theoretic investigation. Technical report, National Bureau of Economic Research.

Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.

Justiniano, A., Primiceri, G. E., and Tambalotti, A. (2010). Investment shocks and business cycles. *Journal of Monetary Economics*, 57(2):132 – 145.

Justiniano, A., Primiceri, G. E., and Tambalotti, A. (2011). Investment shocks and the relative price of investment. *Review of Economic Dynamics*, 14(1):101–121.

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

Kaymak, B. and Poschke, M. (2016). The evolution of wealth inequality over half a century: The role of taxes, transfers and technology. *Journal of Monetary Economics*, 77:1–25.

King, R. G., Plosser, C. I., and Rebelo, S. T. (1988). Production, growth and business cycles: I. The basic neoclassical model. *Journal of Monetary Economics*, 21(2):195 – 232.

Klein, P. (2000). Using the generalized schur form to solve a multivariate linear rational expectations model. *Journal of economic dynamics and control*, 24(10):1405–1423.

Koopman, S. J. and Durbin, J. (2000). Fast filtering and smoothing for multivariate state space models. *Journal of Time Series Analysis*, 21(3):281–296.

Liu, L. and Plagborg-Møller, M. (2019). Full-information estimation of heterogeneous agent models using macro and micro data.

Luetticke, R. (2018). Transmission of monetary policy with heterogeneity in household portfolios. *CFM Discussion Paper Series*, (CFM-DP2018-19).

McKay, A., Nakamura, E., and Steinsson, J. (2016). The power of forward guidance revisited. *The American Economic Review*, 106(10):3133–3158.

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

Ravn, M. O. and Sterk, V. (2017). Job uncertainty and deep recessions. *Journal of Monetary Economics*, 90:125–141.

Reiter, M. (2009). Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665.

Rotemberg, J. J. (1982). Sticky prices in the United States. *Journal of Political Economy*, pages 1187–1211.

Schmitt-Grohé, S. and Uribe, M. (2004). Solving dynamic general equilibrium models using a second-order approximation to the policy function. *Journal of Economic Dynamics and Control*, 28(4):755–775.

Schmitt-Grohé, S. and Uribe, M. (2012). What's news in business cycles. *Econometrica*, 80(6):2733–2764.

Smets, F. and Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606.

Sterk, V. and Tenreyro, S. (2018). The transmission of monetary policy through redistributions and durable purchases. *Journal of Monetary Economics*, 99:124–137.

Storesletten, K., Telmer, C. I., and Yaron, A. (2001). The welfare cost of business cycles revisited: Finite lives and cyclical variation in idiosyncratic risk. *European Economic Review*, 45(7):1311–1339.

Storesletten, K., Telmer, C. I., and Yaron, A. (2004). Cyclical dynamics in idiosyncratic labor market risk. *Journal of Political Economy*, 112(3):695–717.

Taylor, J. B. (1993). Discretion versus policy rules in practice. In *Carnegie-Rochester Conference Series on Public Policy*, volume 39, pages 195–214. Elsevier.

Wong, A. (2019). Refinancing and the transmission of monetary policy to consumption. Technical report, Mimeo.

Woodford, M. (1995). Price-level determinacy without control of a monetary aggregate. *Carnegie-Rochester Conference Series on Public Policy*, 43(Supplement C):1 – 46.

Wu, J. C. and Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound. 48(2-3):253–291.

A Data

The observation equation describes how the empirical times series are matched to the corresponding model variables:

$$OBS_t = \begin{bmatrix} \Delta \log(Y_t) \\ \Delta \log(C_t) \\ \Delta \log(I_t) \\ \Delta \log\left(w_t^F\right) \\ \log\left(\hat{N}_t\right) \\ \log\left(\hat{R}_{t+1}^b\right) \\ \log(\hat{\pi}_t) \\ \log(\hat{s}_t) \\ \log\left(\hat{\lambda}_t\right) \end{bmatrix}$$

where Δ denotes the temporal difference operator and the hats above the variables denote relative deviations from steady state.

Unless otherwise noted, all series available at quarterly frequency from 1954Q3 to 2015Q4 from the St.Louis FED - FRED Database (mnemonics in parentheses).

Output. Sum of gross private domestic investment (GPDI), personal consumption expenditures for nondurable goods (PCND), durable goods (PCDG), and services (PCESV), and government consumption expenditures and gross investment (GCE) divided by the GDP deflator (GDPDEF) and the civilian noninstitutional population (CNP16OV).

Investment. Sum of gross private domestic investment (GPDI) and personal consumption expenditures for durable goods (PCDG) divided by the GDP deflator (GDPDEF) and the civilian noninstitutional population (CNP16OV).

Consumption. Sum of personal consumption expenditures for nondurable goods (PCND) and services (PCESV) divided by the GDP deflator (GDPDEF) and the civilian noninstitutional population (CNP16OV).

Real wage. Hourly compensation in the nonfarm business sector (COMPNFB) divided by the GDP deflator (GDPDEF).

Inflation. Computed as the log-difference of the GDP deflator (GDPDEF).

Nominal interest rate. Quarterly average of the effective Federal Funds Rate (FEDFUNDS). From 2009Q1 till 2015Q4 we use the Wu and Xia (2016) shadow federal funds rate.

Hours worked. Nonfarm business hours worked (COMPNFB) divided by the civilian noninstitutional population (CNP16OV).

Idiosyncratic income risk. Based on Bayer et al. (2019) and available from 1983Q1 till 2013Q1.

Liquidity. Inverse of quarterly average of monthly supply of houses in the United States (MSACSR) divided by 3. Available since 1963Q1.

B Inspecting Model Mechanisms

We find that the HANK model with portfolio choice interprets the data in a different way than the RANK model. Real rigidities are smaller, nominal rigidities larger, and demand shocks are more important.

B.1 Model Comparison

To understand the importance of high marginal propensities to consume and the existence of precautionary savings for this results versus the role of adding the portfolio choice problem between liquid and illiquid assets, we estimate two additional model variants that shut down the last (two) channels. We go into more details by looking into differences in shock propagation (IRF) in the next subsection.

Table 7 displays the parameter estimates from the model variants. Figure 13 compares the variance decompositions. The TANK model, with a saver and a spender type, obtains very similar parameter estimates compared to the representative agent model. If anything, real rigidities appear to be larger. In terms of variance decomposition, we find that demand shocks play a slightly larger role in the TANK model compared to RANK, but the differences are mild and nowhere near the HANK model with two assets.

The HANK-1 model where all assets are liquid shows mostly a picture that is similar to the TANK model. The investment adjustment costs fall, while the utilization cost increase, nominal rigidities appear smaller. Notwithstanding, in the variance decomposition the importance of demand shocks increases compared to the TANK model, but still is lower than in the HANK-2 model. The main reason for this is that shocks to income uncertainty play little role in the fluctuation of output and consumption. This reflects the fact that the main mechanism through which uncertainty shocks operate in the two-asset HANK model is portfolio re-balancing, see Bayer et al. (2019).

Table 7: Comparison of model estimates, further model variants

Parameter	TANK	HANK-1	HANK-2	HANK-2*
		Frictions		
ϕ	0.482	0.231	0.334	0.445
	(0.361, 0.632)	(0.207, 0.254)	(0.289, 0.377)	(0.390, 0.504)
δ_2/δ_1	1.023	1.701	0.165	0.269
	(0.823, 1.231)	(1.535, 1.877)	(0.118, 0.215)	(0.220, 0.325)
κ	0.129	0.138	0.067	0.069
	(0.100, 0.159)	(0.132, 0.144)	(0.053, 0.083)	(0.053, 0.089)

Table 7: Comparison of model estimates - Continued

Parameter	TANK	HANK-1	HANK-2	HANK-2*
κ_w	0.140	0.172	0.170	0.158
	(0.101, 0.180)	(0.141, 0.208)	(0.133, 0.210)	(0.122, 0.200)
		Monetary policy rules		
ρ_R	0.757	0.680	0.749	0.756
	(0.724, 0.789)	(0.656, 0.702)	(0.720, 0.776)	(0.730, 0.781)
σ_R	0.250	0.280	0.253	0.247
	(0.228, 0.274)	(0.257, 0.306)	(0.231, 0.278)	(0.226, 0.271)
θ_π	1.936	1.723	1.934	1.885
	(1.771, 2.117)	(1.689, 1.762)	(1.850, 2.024)	(1.831, 1.941)
θ_y	0.269	0.306	0.422	0.349
	(0.206, 0.334)	(0.272, 0.338)	(0.381, 0.464)	(0.320, 0.378)
		Fiscal policy rules		
ρ_B	0.997	0.992	0.983	0.984
	(0.994, 0.999)	(0.986, 0.995)	(0.977, 0.987)	(0.978, 0.987)
ρ_G	0.985	0.989	0.990	0.995
	(0.970, 0.996)	(0.978, 0.997)	(0.978, 0.997)	(0.989, 0.998)
σ_G	0.261	0.211	0.169	0.184
	(0.232, 0.292)	(0.194, 0.229)	(0.154, 0.186)	(0.168, 0.203)
γ_Y	-0.295	-0.236	-0.168	-0.166
	(-0.334, -0.258)	(-0.256, -0.216)	(-0.189, -0.149)	(-0.187, -0.145)
		Structural Shocks		
ρ_A	0.991	0.985	0.996	0.996
	(0.981, 0.997)	(0.971, 0.995)	(0.992, 0.999)	(0.992, 0.999)
σ_A	0.125	0.102	0.188	0.198
	(0.108, 0.143)	(0.091, 0.114)	(0.172, 0.205)	(0.181, 0.216)
ρ_Z	0.999	0.981	0.977	0.986
	(0.998, 1.000)	(0.972, 0.990)	(0.967, 0.987)	(0.978, 0.993)
σ_Z	0.545	0.576	0.645	0.640
	(0.504, 0.588)	(0.532, 0.625)	(0.599, 0.697)	(0.594, 0.692)
ρ_Ψ	0.975	0.972	0.997	0.997
	(0.961, 0.987)	(0.964, 0.980)	(0.994, 0.999)	(0.994, 0.999)

Table 7: Comparison of model estimates - Continued

Parameter	TANK	HANK-1	HANK-2	HANK-2*
σ_Ψ	2.526	2.190	1.428	1.626
	(2.263, 2.822)	(1.961, 2.432)	(1.302, 1.560)	(1.479, 1.788)
ρ_μ	0.877	0.914	0.990	0.988
	(0.846, 0.904)	(0.891, 0.936)	(0.980, 0.997)	(0.976, 0.996)
σ_μ	1.621	1.404	1.328	1.338
	(1.408, 1.879)	(1.295, 1.523)	(1.197, 1.473)	(1.206, 1.489)
$\rho_{\mu w}$	0.852	0.870	0.872	0.870
	(0.797, 0.897)	(0.837, 0.902)	(0.837, 0.902)	(0.833, 0.906)
$\sigma_{\mu w}$	5.606	4.996	4.666	4.876
	(4.582, 7.045)	(4.342, 5.759)	(3.993, 5.476)	(4.108, 5.781)
Risk and Liquidity Process				
ρ_s	0.804	0.749	0.643	0.611
	(0.724, 0.882)	(0.701, 0.784)	(0.593, 0.687)	(0.565, 0.658)
σ_s	58.904	60.079	85.23	72.998
	(52.852, 65.503)	(57.566, 62.533)	(75.708, 95.004)	(66.153, 79.848)
Σ_N	0.124	-0.053	-0.521	-0.652
	(-0.606, 0.823)	(-0.083, -0.027)	(-0.569, -0.471)	(-0.70, -0.598)
ρ_λ	0.924	0.932	0.901	0.915
	(0.887, 0.957)	(0.902, 0.960)	(0.870, 0.930)	(0.883, 0.943)
σ_λ	8.804	8.804	8.847	8.833
	(8.116, 9.557)	(8.135, 9.532)	(8.171, 9.572)	(8.142, 9.593)
Λ_R	-0.314	-0.137	-0.626	-0.661
	(-1.009, 0.353)	(-0.327, 0.045)	(-0.674, -0.574)	(-0.728, -0.591)
Measurement Errors				
σ_λ^{me}	0.041	0.038	0.040	0.038
	(0.013, 0.099)	(0.012, 0.093)	(0.012, 0.097)	(0.012, 0.093)
σ_{income}^{me}	-	-	-	7.925
	(-, -)	(-, -)	(-, -)	(6.702, 9.360)
σ_{wealth}^{me}	-	-	-	2.302
	(-, -)	(-, -)	(-, -)	(1.905, 2.779)

Notes: Parentheses contain the 90% highest posterior density interval. The standard deviations of the shocks and measurement errors have been transformed into percentages by multiplying with 100.

Figure 13: Variance Decompositions: Model Comparison

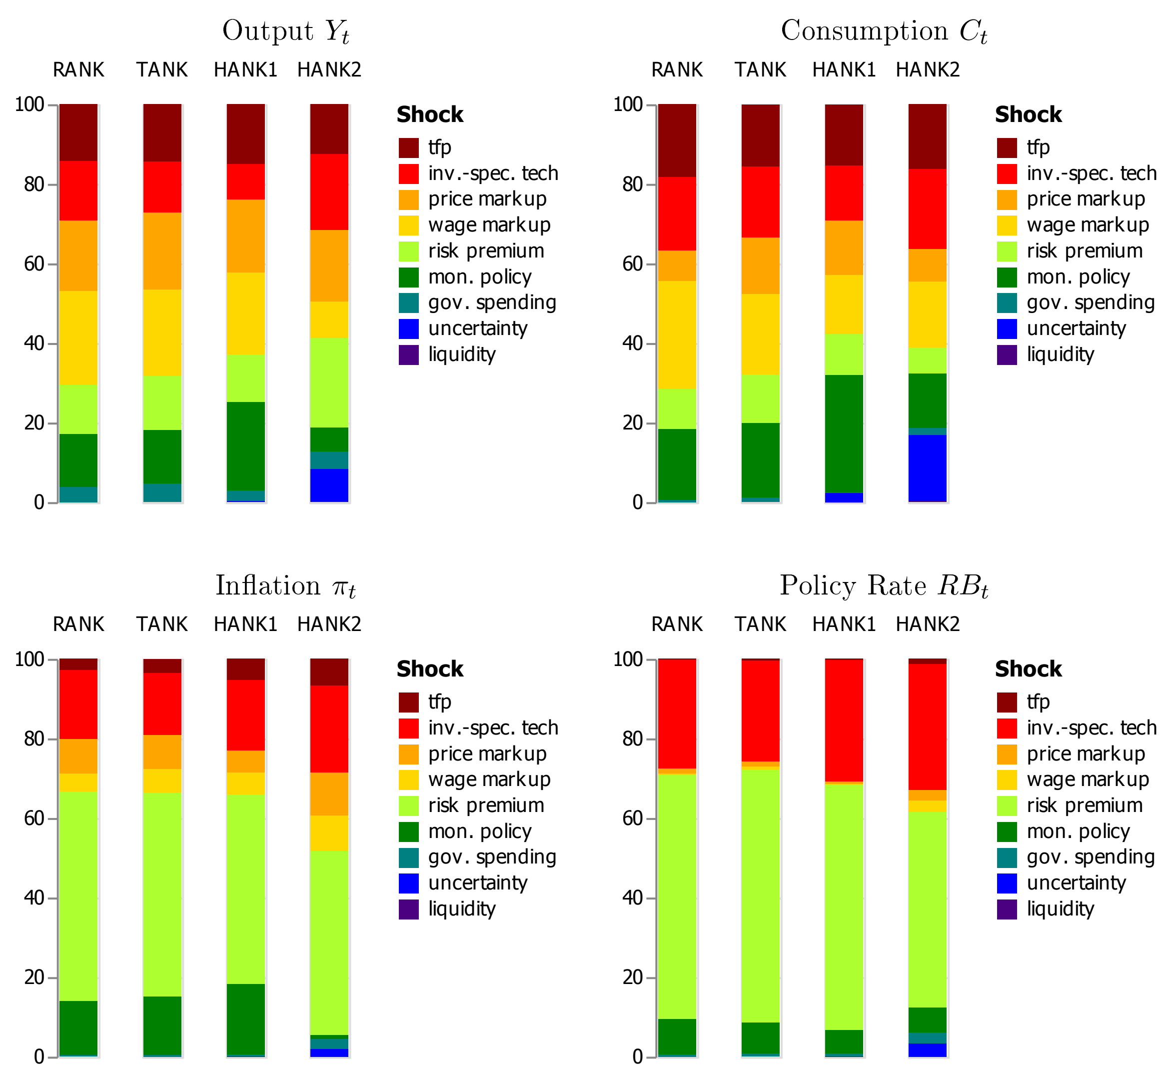

Notes: Variance decomposition at a 4-quarter forecast horizon.

B.2 Comparing Impulse Response Functions across Models

To understand, where the differences in the estimation come from, it is useful to first compare the impulse response functions across models, for the same set of parameters. As we will see, the key difference is the response of the real interest rate, and in particular the liquidity premium, that is the expected difference in return between holding an illiquid asset and a liquid asset, to an increase in government debt. Higher government debt leads to an increased availability of liquid assets in the HANK model. Since their usefulness for consumption smoothing is decreasing in their quantity, the liquidity premium falls. At the same time, along the lines of Aiyagari and McGrattan (1998), the increase in government debt crowds out physical capital. The result is a striking difference in the medium term behavior of the HANK model on the one hand and the RANK and TANK model on the other hand, see the top panel in Figure 14 which displays the impulse response to a government spending shock using the same parameter values across models (HANK estimates).

The bottom panel for the same figure shows the response to a monetary policy shock. While there is a difference in the impact behavior of consumption – the elasticity of intertemporal substition is muted in incomplete markets, see Kaplan et al. (2018) – the big difference is again the medium term behavior which diverges as soon as government debt deviates sufficiently strongly from the stationary equilibrium.

Figure 15 runs the same experiment for the markup shocks. Again, we see differences in the consumption behavior on impact – here because markup shocks have direct effects on income risk and income inequality, but there is also the medium term difference that emerges with government debt building up, pushing up real rates and crowding out investment.

Figure 16 shows that the HANK model also differs in another important dimension from the RANK model. When productivity increases we see that investment reacts more strongly. The reason is that households hold capital not only for intertemporal but also for precautionary motives. As output goes up, households would like to spend more resources on self-insurance. At the same time, government debt falls as output goes up. In consequence households invest more into capital because of higher income. What is more, also the relative returns between liquid and illiquid assets change which leads to a rebalancing of portfolios towards the illiquid asset. As a result, the capital accumulation response to productivity shocks is substantially different and thus the business cycle dynamics differs.

Figure 14: Impulse responses to a government spending and monetary policy – keeping parameters fixed –

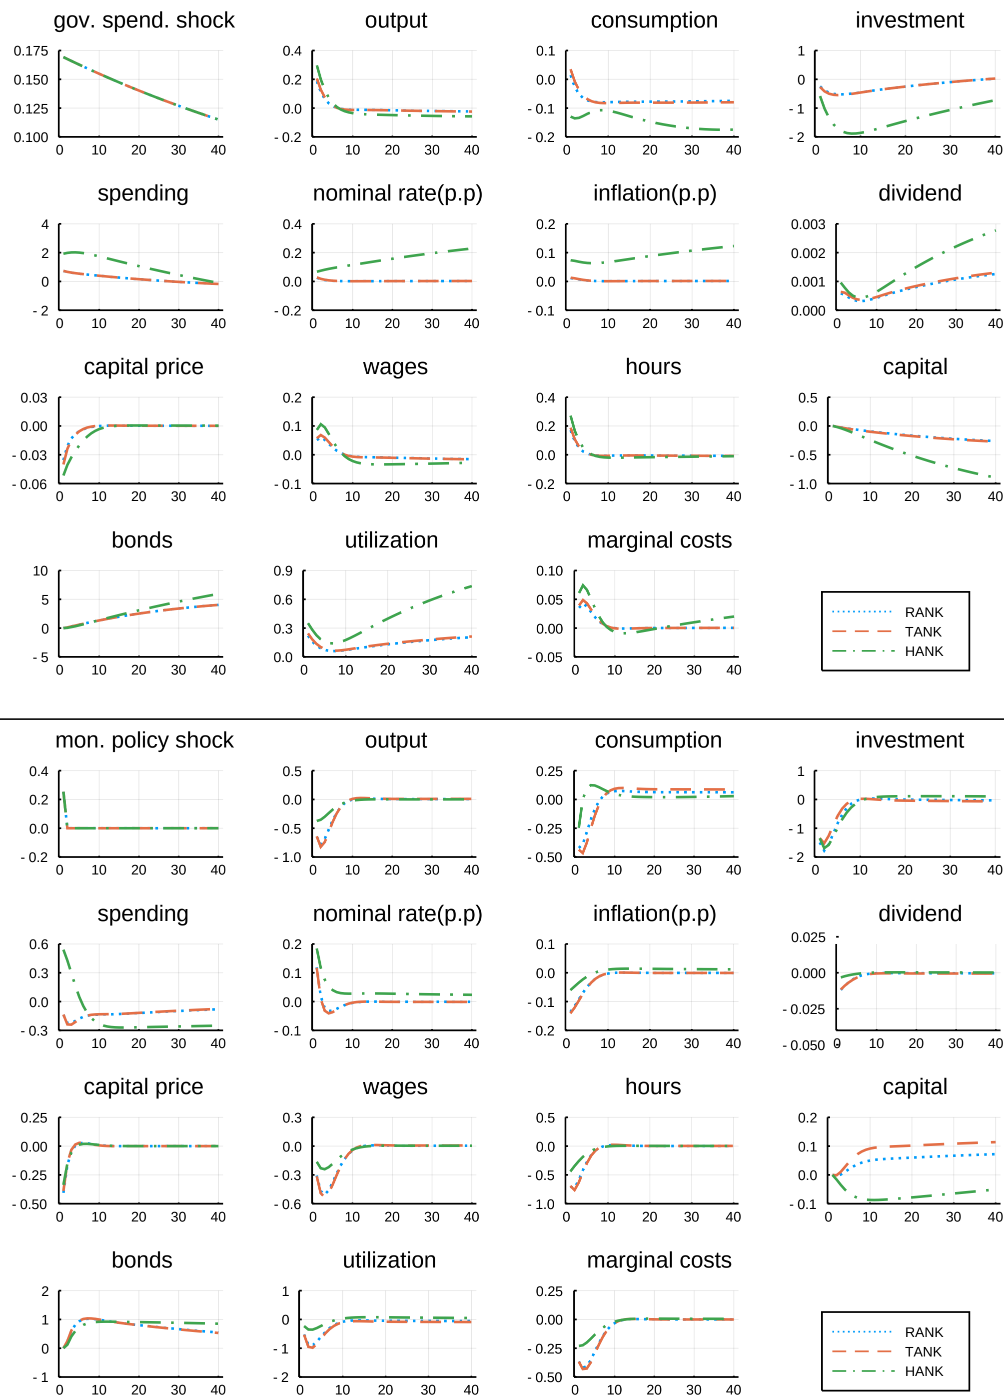

Notes: Top: Impulse responses a government spending shock. Bottom: Impulse response to a monetary policy shock. The parameters are the same across models and equal our baseline HANK model estimates.

Figure 15: Impulse responses to markup shocks
– keeping parameters fixed –

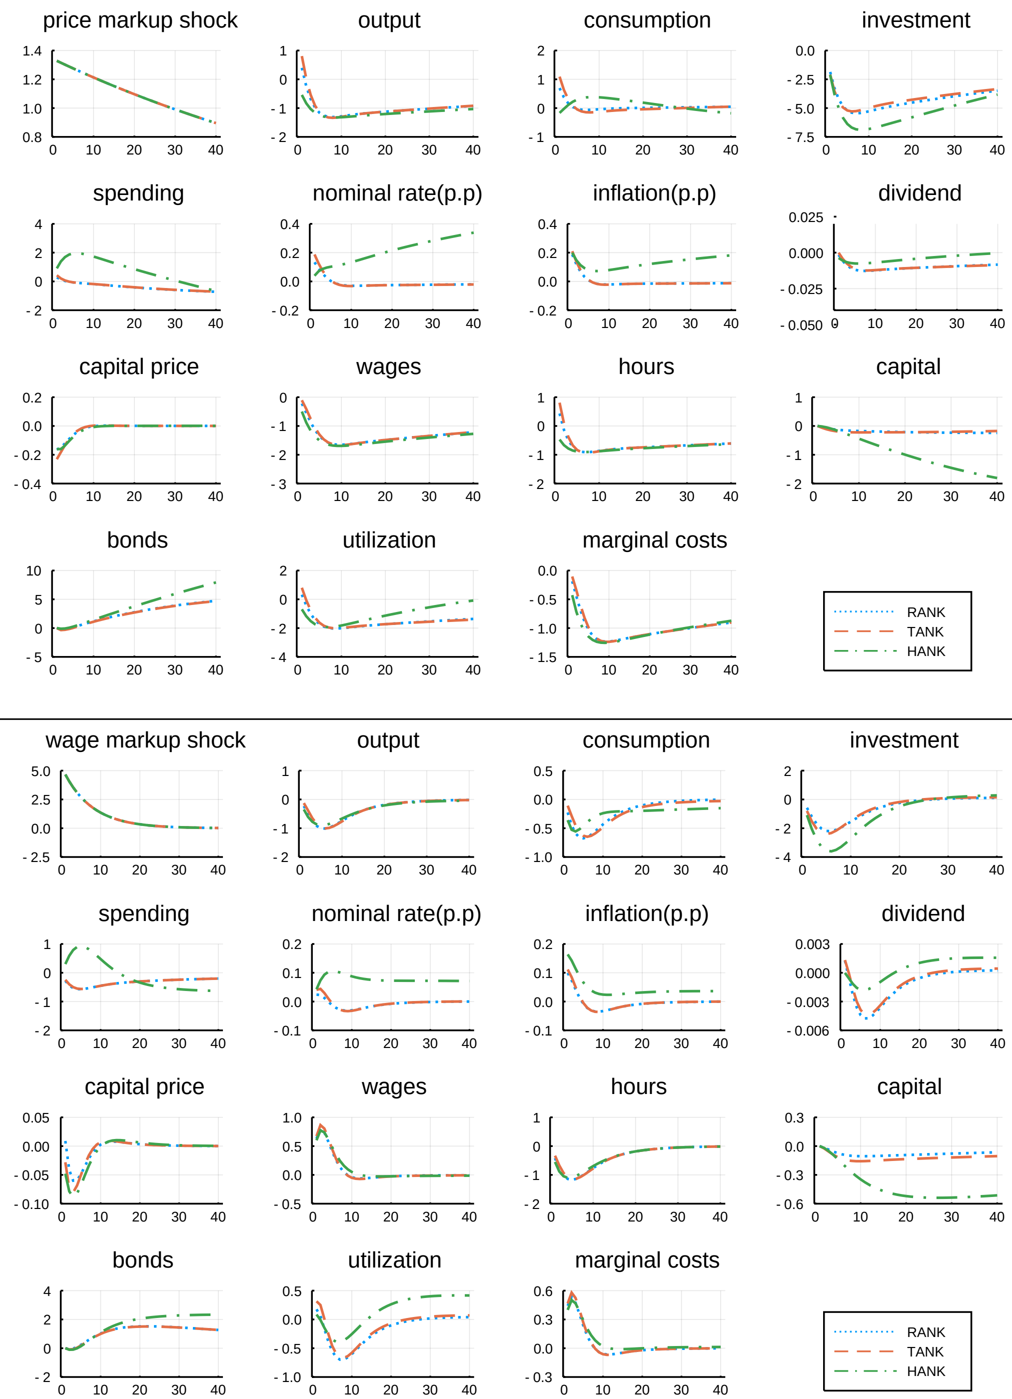

Notes: Top: Impulse responses a price-markup shock. Bottom: Impulse response to a wage-markup shock. The parameters are the same across models and equal our baseline HANK model estimates.

Figure 16: Impulse responses to technology shocks
– keeping parameters fixed –

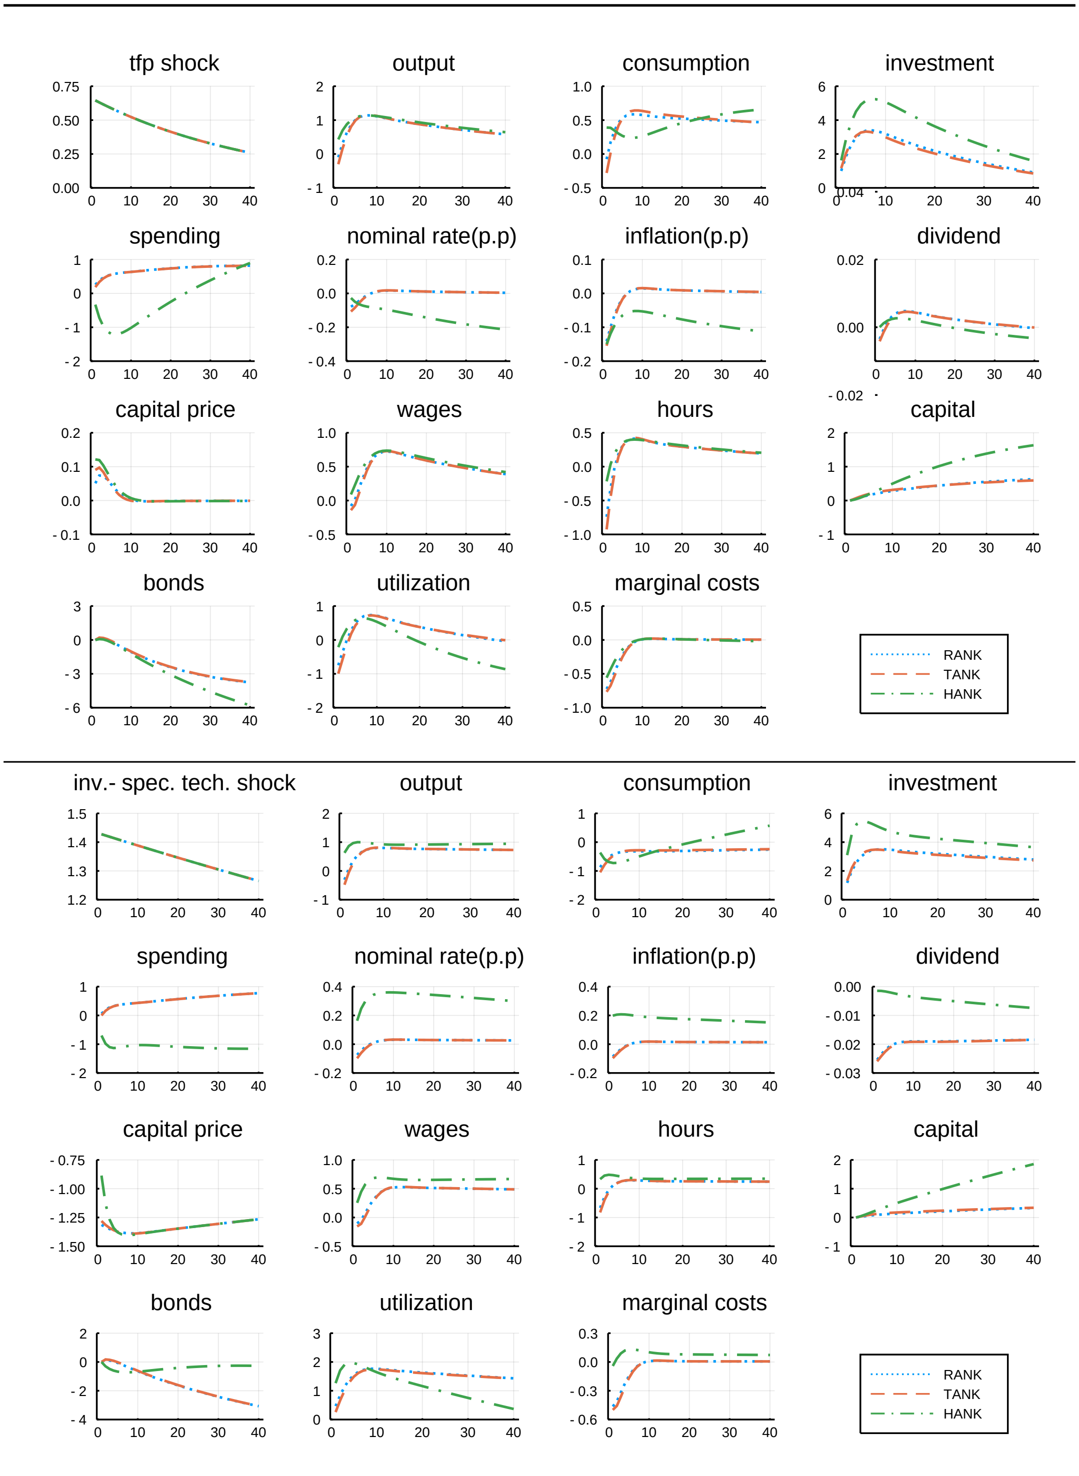

Notes: Top: Impulse responses a TFP shock. Bottom: Impulse response to a MEI shock. The parameters are the same across models and equal our baseline HANK model estimates.

Of course, the data constrains impulse responses and, therefore, once we estimate the models, impulse responses look more comparable for those variable/shock combinations where actual data closely constrains the estimation, like monetary policy shocks and consumption responses, see Figures 17 to 19. Figure 20 does the comparison for the risk premium shock, both for keeping the parameters constant as well as for re-estimating the parameters.

Finally, Figure 21 shows the responses of the HANK model to uncertainty and liquidity. An increase in uncertainty as well as an increase in the time-to-sell leads to a portfolio rebalancing towards liquid assets and therefore to a simultaneous decline in consumption and investment (more savings are required for the same level of consumption smoothing).

Figure 17: Impulse responses to a government spending and monetary policy – re-estimated –

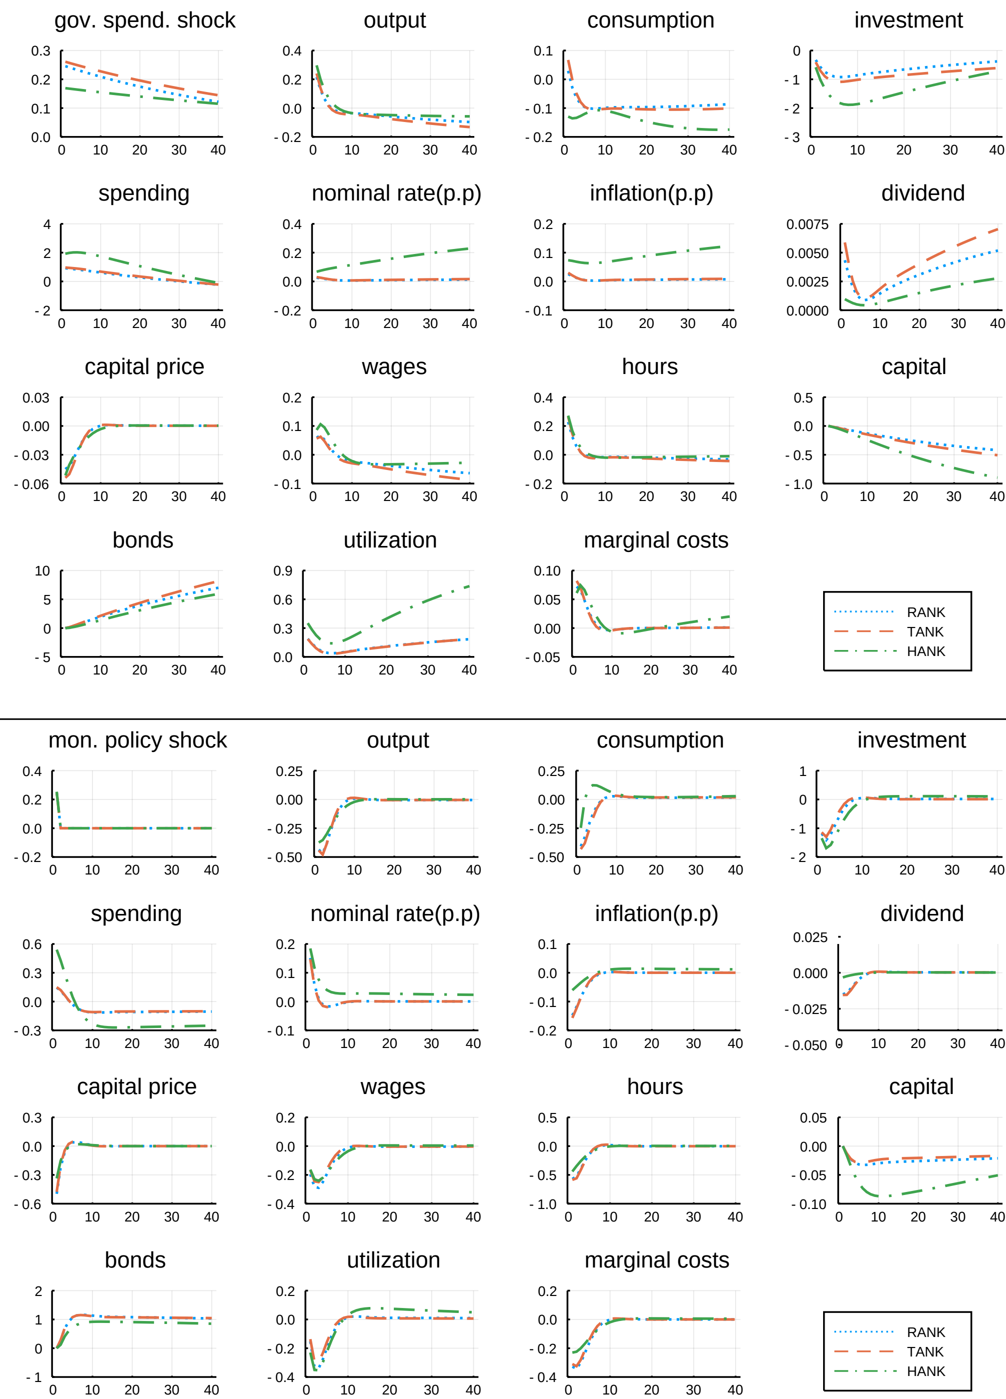

Notes: Top: Impulse responses a government spending shock. Bottom: Impulse response to a monetary policy shock. The parameters are each estimated using Bayesian maximum likelihood as described in the main text.

Figure 18: Impulse responses to markup shocks
– re-estimated –

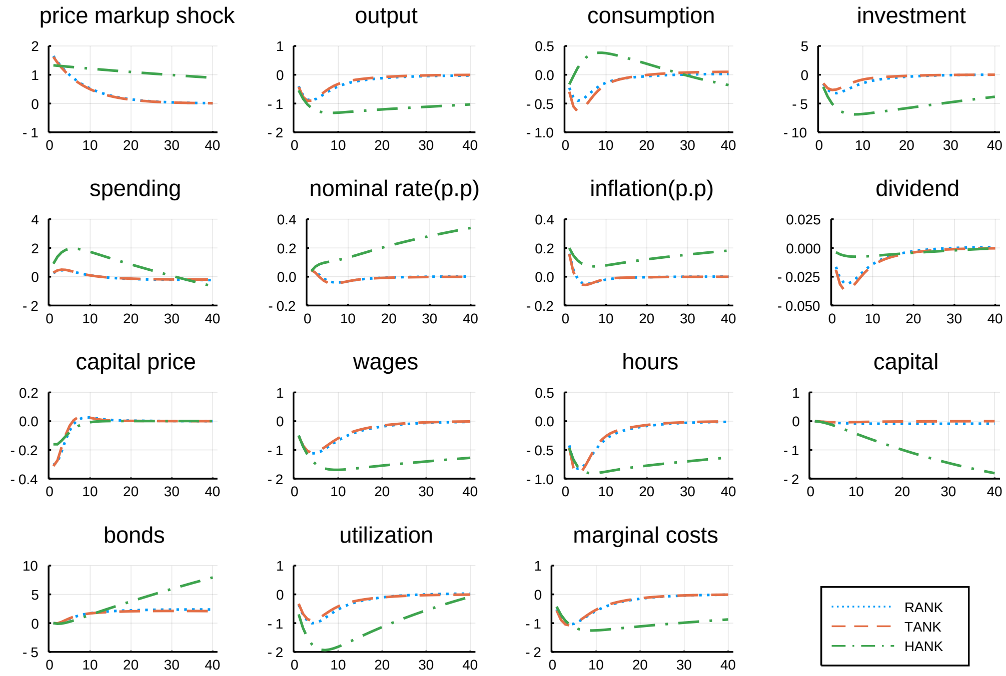

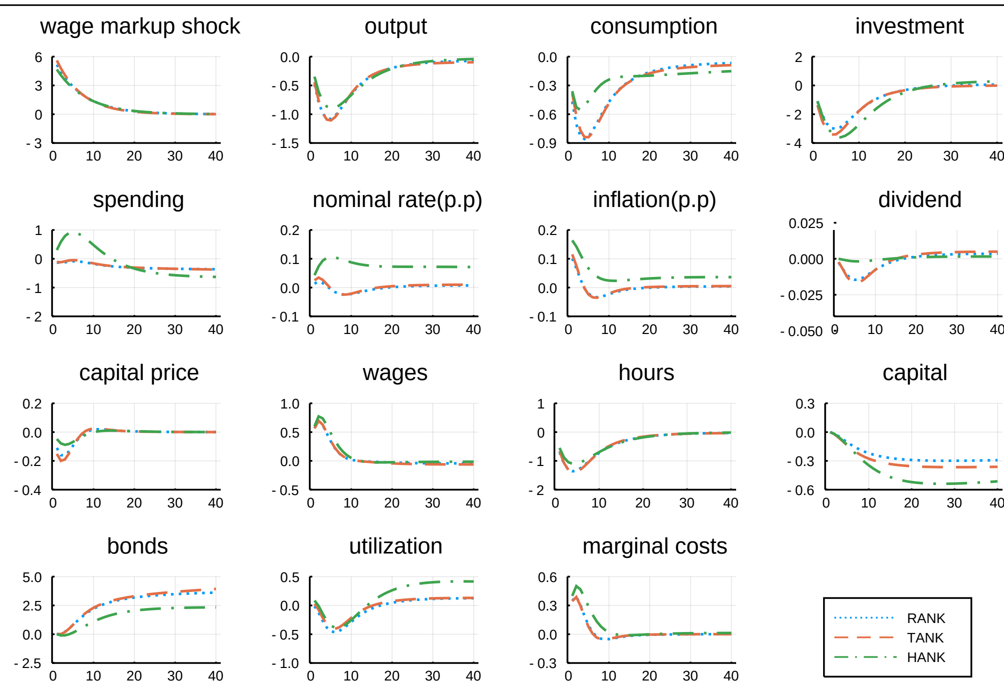

Notes: Top: Impulse responses a price-markup shock. Bottom: Impulse response to a wage-markup shock. The parameters are each estimated using Bayesian maximum likelihood as described in the main text.

Figure 19: Impulse responses to technology shocks
– re-estimated –

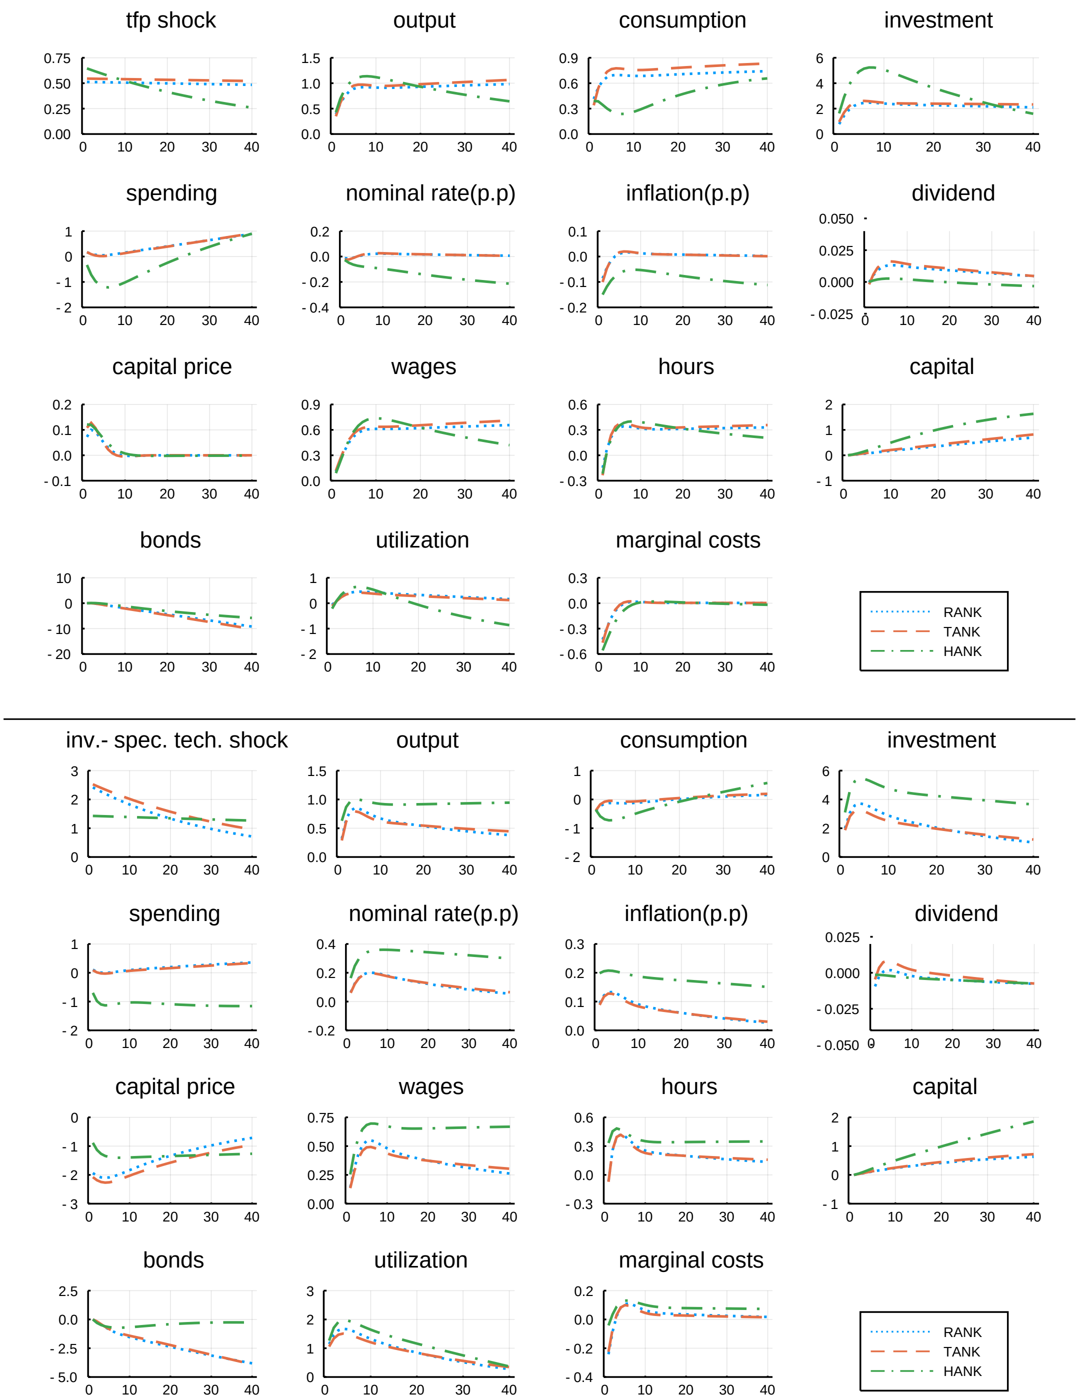

Notes: Top: Impulse responses a TFP shock. Bottom: Impulse response to a MEI shock. The parameters are each estimated using Bayesian maximum likelihood as described in the main text.

Figure 20: Impulse responses to a risk premium shock

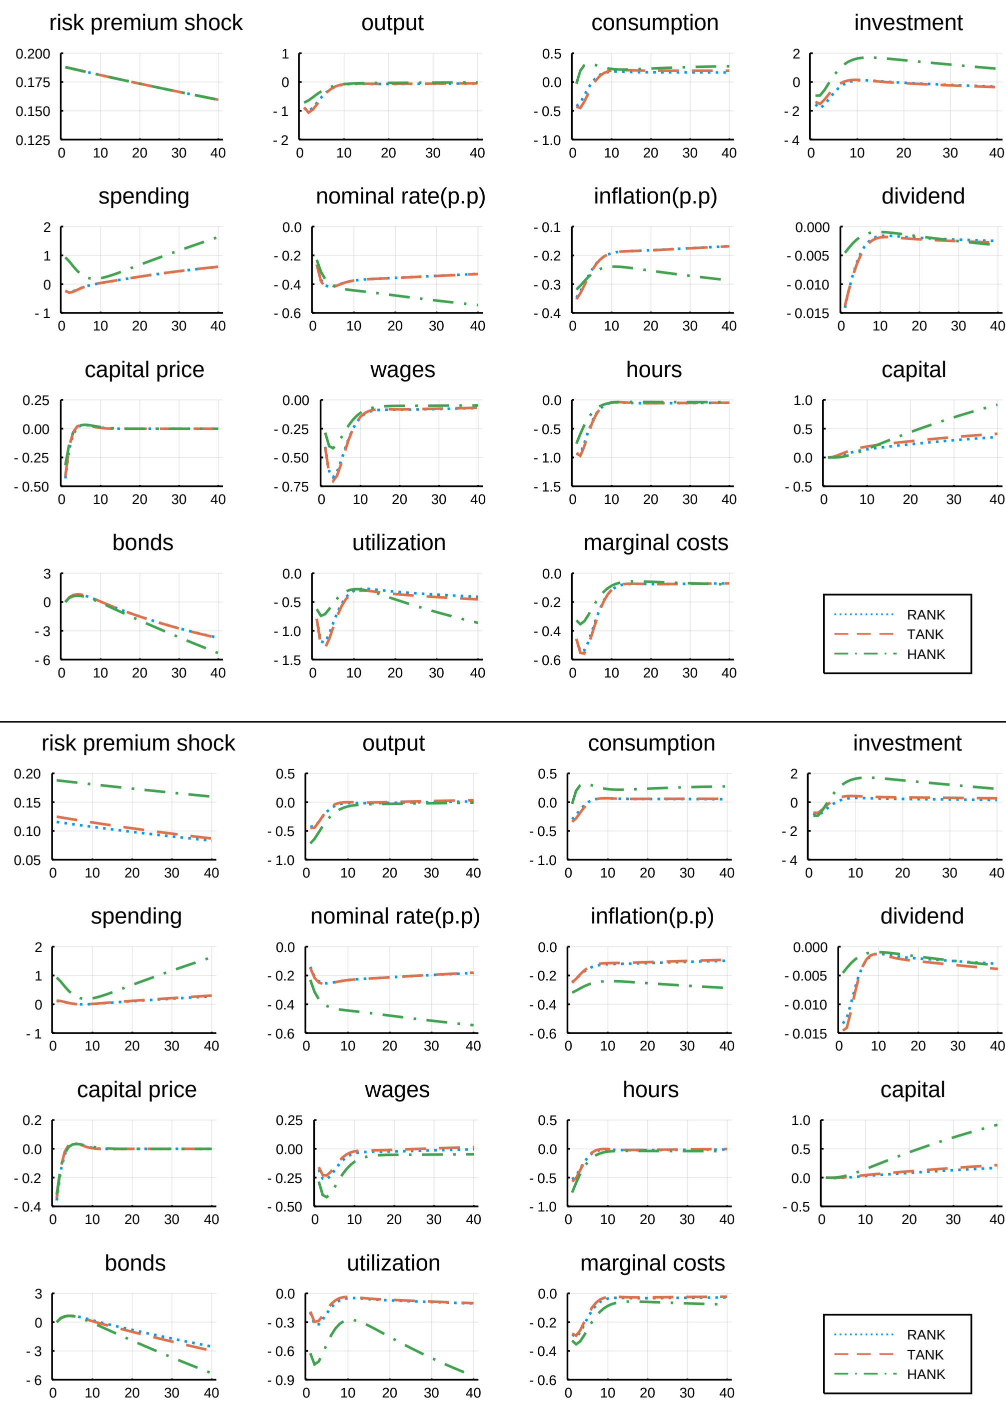

Notes: Top: Impulse responses when all models have the same parameters (as in HANK). Bottom: Impulse responses with estimated parameters.

Figure 21: Impulse responses to an income uncertainty and liquidity
– estimated –

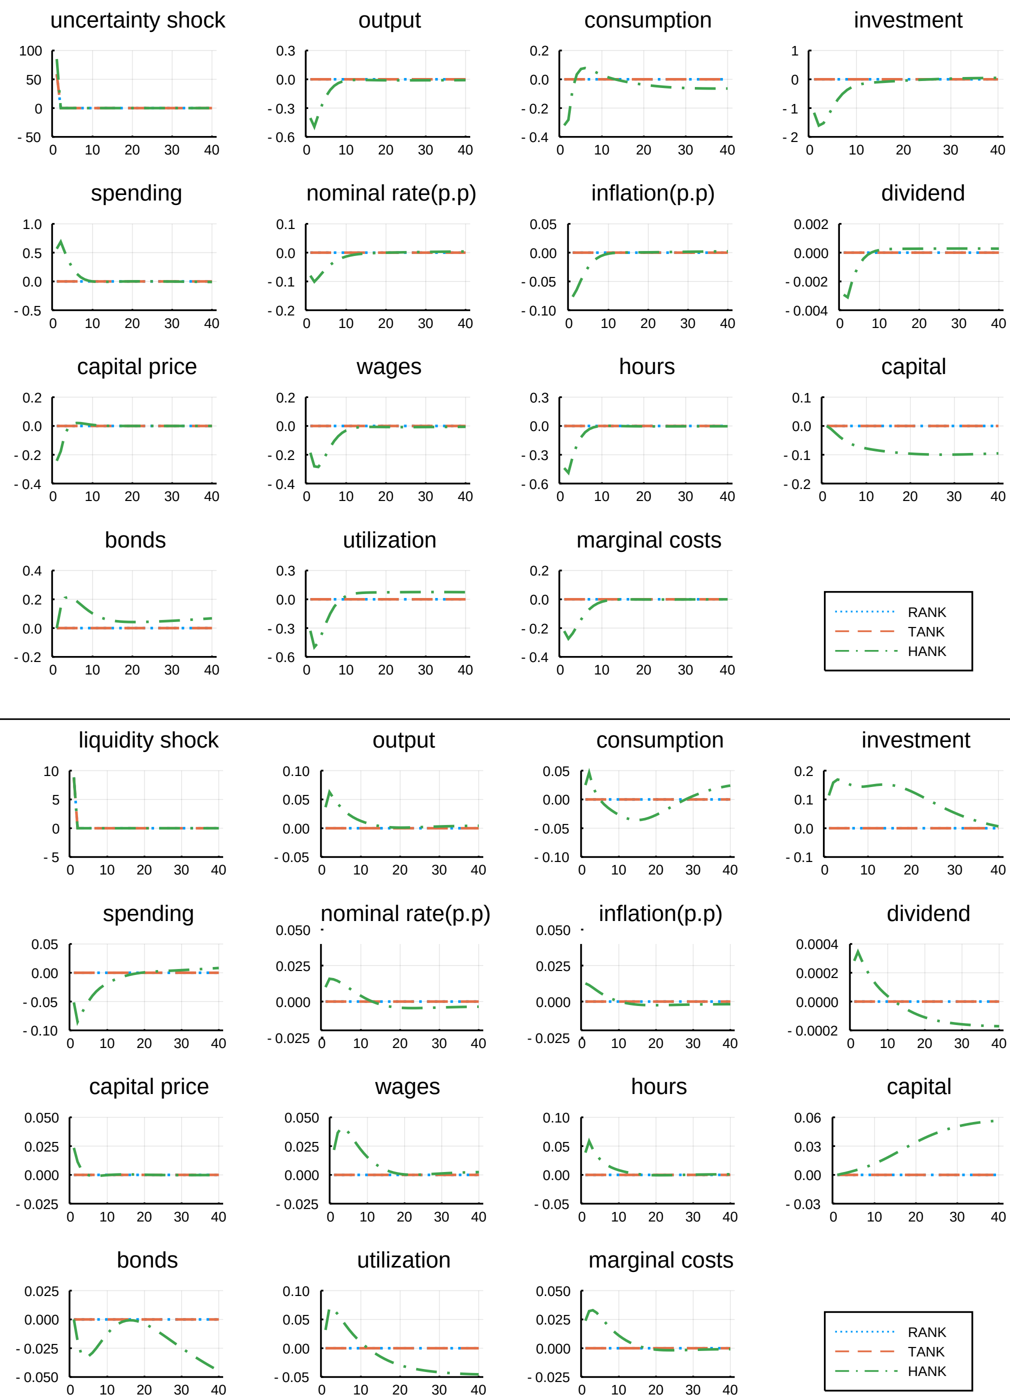

Notes: Top: Impulse responses a shock to uncertainty. Bottom: Impulse response to a shock to the time-to-sell (liquidity). The parameters are each estimated using Bayesian maximum likelihood as described in the main text.

C Further historical decompositions

Figure 22 shows the historical decomposition of income risk and liquidity. The historical decompositions show that both are mostly driven by exogenous shocks and not endogenous feedback.

Figure 22: Historical Decompositions: Income Risk and Liquidity

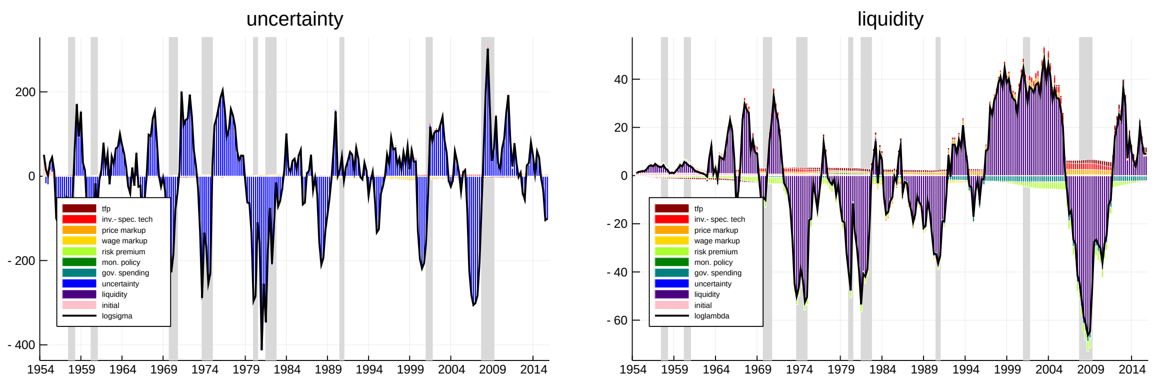

Notes: Historical decomposition of income risk and liquidity.

Figure 23 shows the historical decomposition of the growth rate of consumption, investment, and government spending for the HANK and RANK model.

Figure 23: Historical Decompositions: Consumption, Investment and Government Spending Growth

RANK	HANK

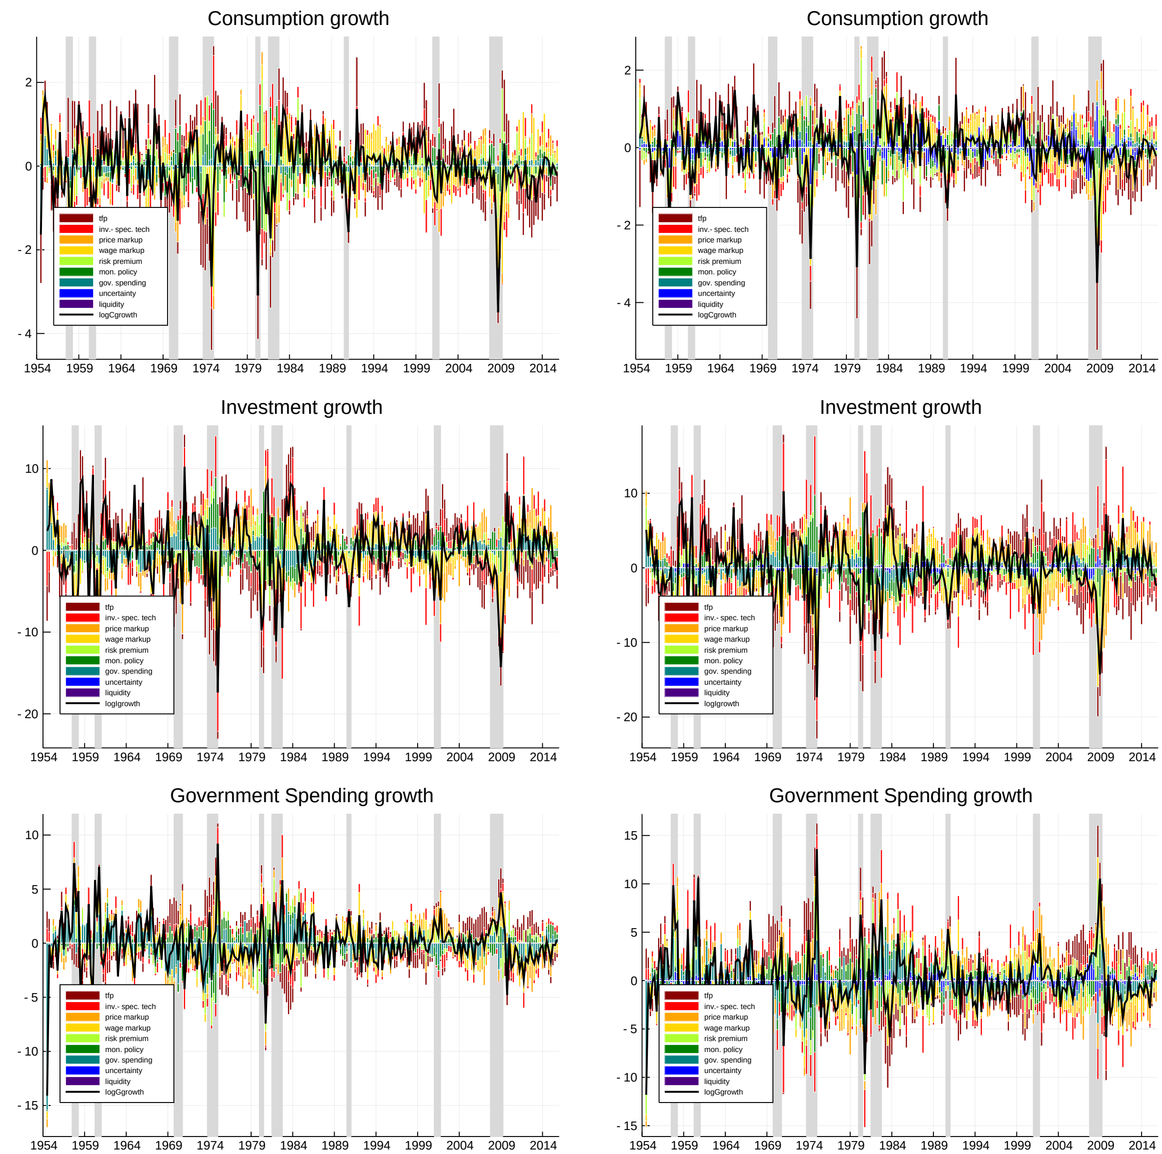

Notes: The top panel shows the historical decomposition of consumption growth into the contribution of various shocks. The middle panel shows the same for investment growth. The bottom panel shows the same for government spending growth. The left column is for the RANK estimates the right column for the HANK estimates. The contribution of the smoothed initial state has been omitted.

D Further results on inequality dynamics

Figure 24 presents the estimated impulse responses of inequality on the other shocks not reported in the main text.

Figure 24: Impulse responses of inequality

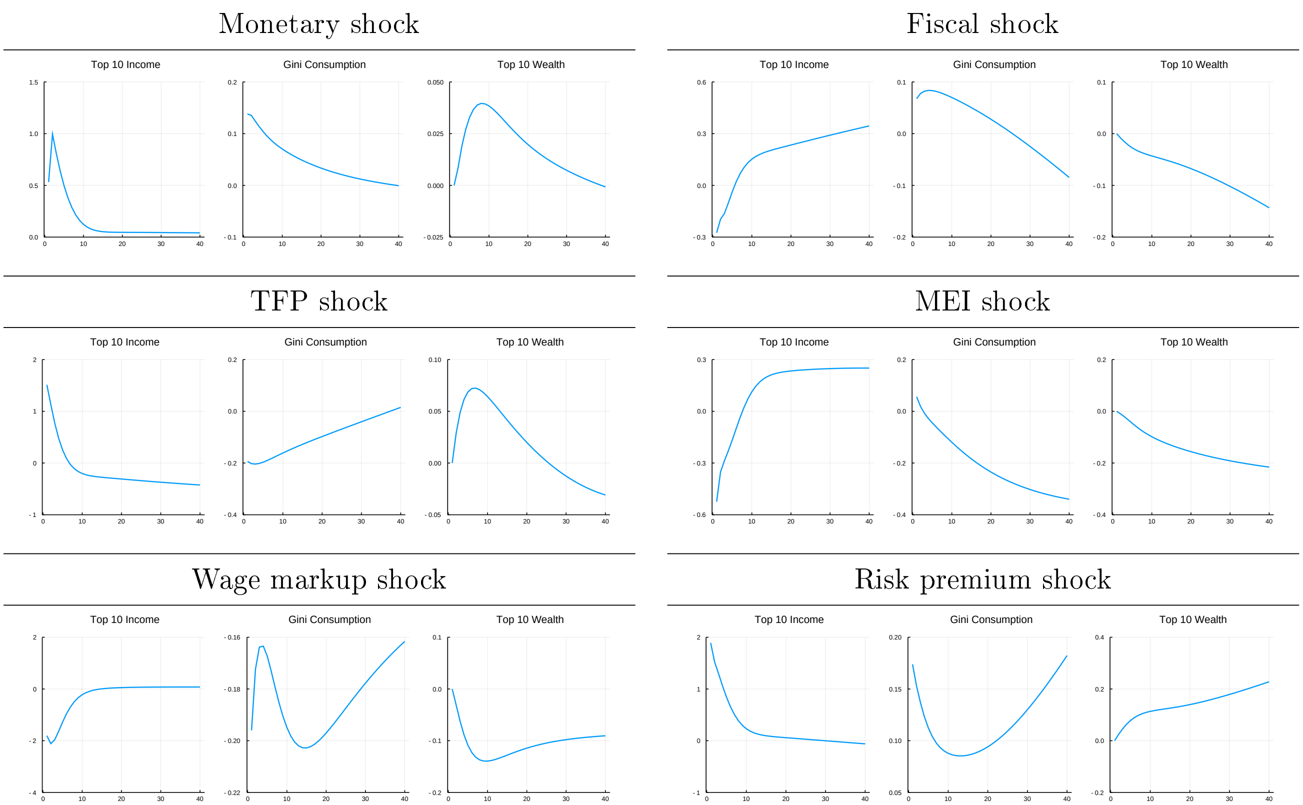

Notes: The figures display the impulse response of income, consumption, and wealth inequality in response to the shocks labeled above. Parameter estimates from HANK*. See main text for further details.

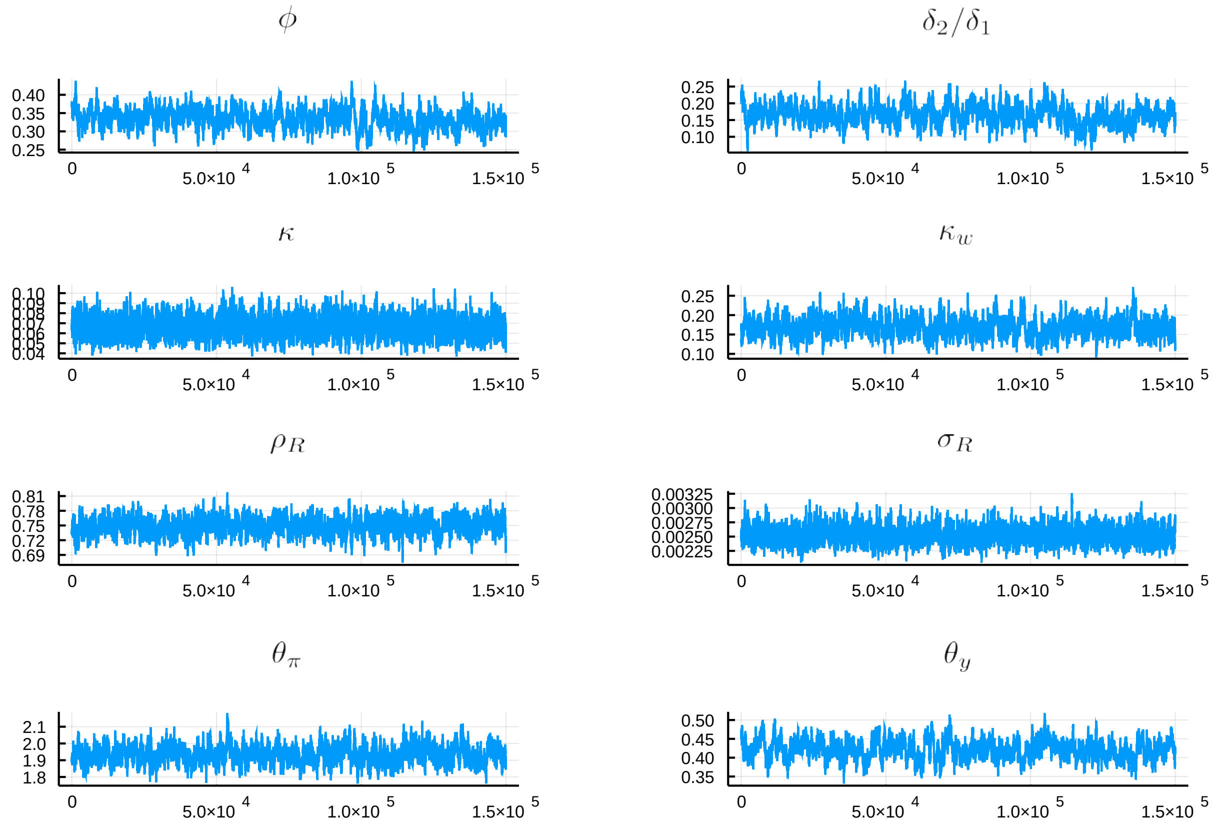

Figure 25: MCMC draws of HANK-2 model

E MCMC Diagnostics

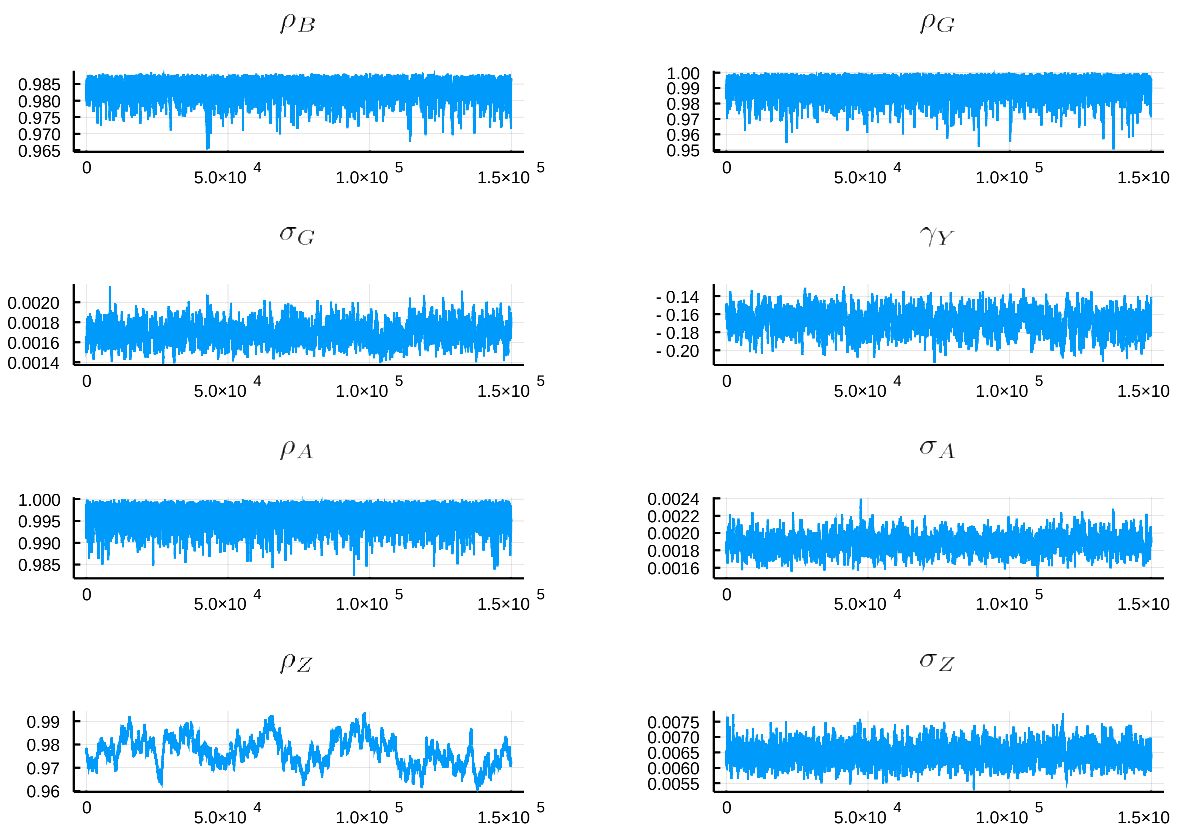

Figure 26: MCMC draws of HANK-2 model

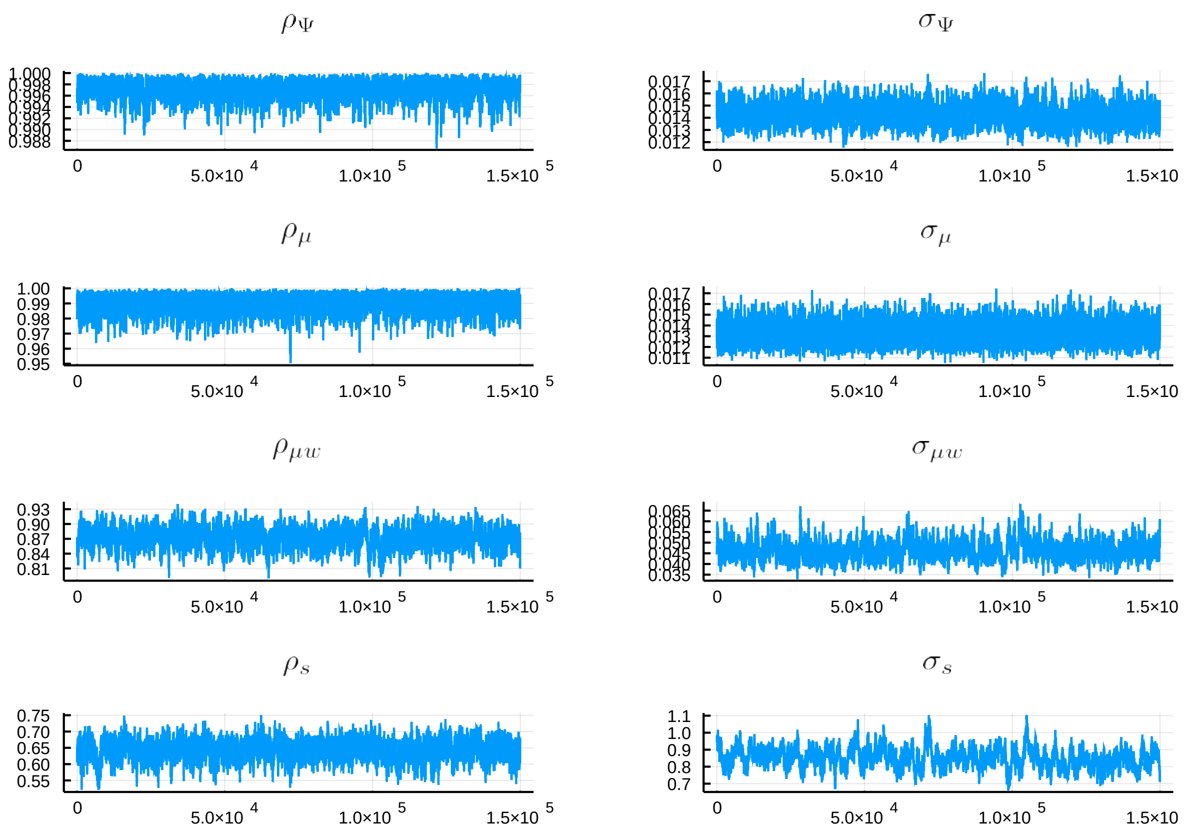

Figure 27: MCMC draws of HANK-2 model

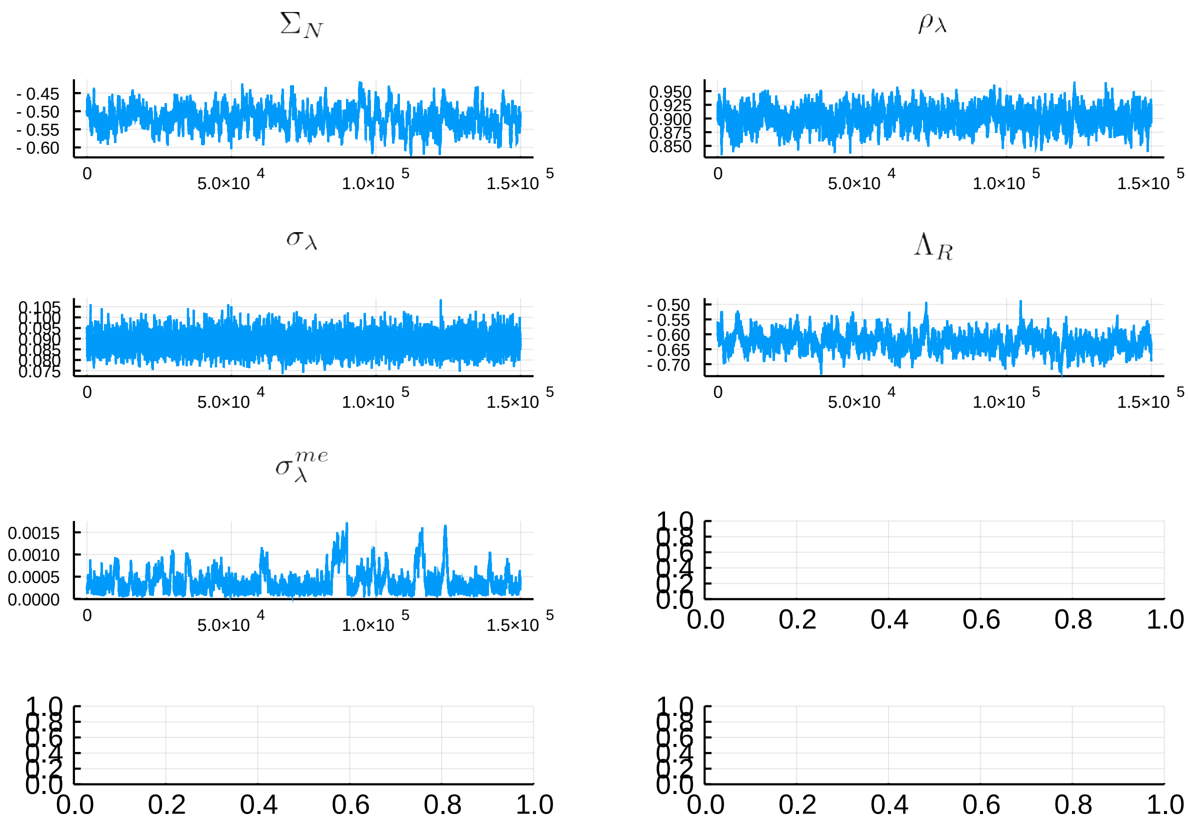

Figure 28: MCMC draws of HANK-2 model

Table 8: Convergence diagnostics

	Geweke		Gelman and Rubin	
Parameter	test statistic	p-value	PSRF	97.5%
ϕ	2.257	0.024	1.045	1.114
δ_2/δ_1	1.488	0.137	1.026	1.068
κ	-0.003	0.997	1.004	1.011
κ_w	0.169	0.866	1.009	1.022
ρ_R	-0.816	0.415	1.003	1.007
σ_R	0.818	0.413	1.002	1.005
θ_π	-0.31	0.756	1.005	1.011
θ_y	1.117	0.264	1.016	1.04
ρ_B	0.927	0.354	1.003	1.007
ρ_G	0.242	0.809	1.005	1.012
σ_G	-0.539	0.59	1.005	1.012
γ_Y	1.359	0.174	1.009	1.024
ρ_A	-0.07	0.944	1.001	1.002
σ_A	-0.644	0.52	1.007	1.017
ρ_Z	0.721	0.471	1.09	1.236
σ_Z	-0.824	0.41	1.003	1.006
ρ_Ψ	2.64	0.008	1.007	1.015
σ_Ψ	1.328	0.184	1.011	1.028
ρ_μ	-1.94	0.052	1.004	1.01
σ_μ	0.182	0.856	1.001	1.001
$\rho_{\mu w}$	1.023	0.306	1.01	1.023
$\sigma_{\mu w}$	-0.512	0.609	1.015	1.033
ρ_s	-0.906	0.365	1.005	1.011
σ_s	1.878	0.06	1.041	1.107
Σ_N	0.706	0.48	1.014	1.029
ρ_λ	-0.623	0.534	1.005	1.013
σ_λ	0.048	0.962	1.002	1.004
Λ_R	1.999	0.046	1.035	1.093
σ_λ^{me}	-1.376	0.169	1.054	1.114

Note: Columns 1-2: Geweke (1992) test of equality of means of the first 10% of draws last 50% of draws (after burn-in); columns 3-4: Gelman and Rubin (1992) potential scale reduction factor and its 97.5% quantile based on 5 chains.