Shocks, Frictions, and Inequality in US Business Cycles*

Christian Bayer, Benjamin Born, and Ralph Luetticke

January 30, 2020

Abstract

How much 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 taxes. We find that adding data on inequality does not materially change the estimated shocks and frictions driving the US business cycle. The estimated shocks, however, have significantly contributed to the evolution of US wealth and income inequality. The systematic components of monetary and fiscal policy are important for inequality as well.

JEL codes: C11, D31, E32, E63

Keywords: Bayesian Estimation, Business Cycles, Income Inequality, Incomplete Markets, Monetary and Fiscal Policy, Wealth Inequality

*Bayer: University of Bonn, CEPR, CESifo, and IZA, christian.bayer@uni-bonn.de, Born: Frankfurt School of Finance & Management, CEPR, and CESifo, b.born@fs.de, Luetticke: University College London, CEPR, and CFM, r.luetticke@ucl.ac.uk. We would like to thank Gregor Boehl, Edouard Challe, Miguel Leon-Ledesma, Johannes Pfeifer, Mu-Chun Wang, and participants at various conferences and seminars 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. Codes are available at https://github.com/BenjaminBorn/HANK_BusinessCycleAndInequality.

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] Much of this literature so far has focused on specific channels of transmission, shocks, or puzzles. In contrast, the present paper takes a more encompassing approach and asks how our view of the business cycle and of inequality dynamics changes when we bring this model to the data. In particular, we aim to answer two questions: First, do data on inequality change the estimated shocks and frictions driving the US business cycle? Second, how important are business cycle shocks for the evolution of US inequality?

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 progressive taxation, 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 progressivity of taxes 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, which gives rise to an endogenous liquidity premium, and the presence of occasional hand-to-mouth consumers leads the HANK model to have rich distributional dynamics in response to aggregate shocks.

[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) study the effectiveness of forward guidance.

[2] Auclert et al. (2018) and Hagedorn et al. (2019) 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.

To infer the importance of inequality for the business cycle, we estimate the HANK model with and without data on inequality. We first estimate the model on the same observables as in Smets and Wouters (2007) (plus proxies for income risk and taxes) covering the time period of 1954 to 2015.[4] We then 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 Inequality Database. We focus on the top 10% shares because this measure is most consistent across alternative data sources such as the Survey of Consumer Finances, where available.[5] With respect to the first question, we find that the addition of distributional data does not change what we infer about the aggregate shocks and frictions driving the US business cycle. The answer to the second question rationalizes this result. We find that business cycle shocks generate very persistent movements in wealth and income inequality that are consistent with the U-shaped evolution of US inequality 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 and thus business cycle shocks persistently redistribute across households with different portfolios. To our knowledge, this paper is the first to quantify the distributional consequences of all standard business cycle shocks and estimate their importance in explaining US inequality. The impulse response functions show that wealth inequality responds in particular to technology, fiscal, and markup shocks. Wage and price markups directly affect the distribution of income shares, while technology and fiscal shocks affect the return spread between illiquid capital and liquid bonds. For income inequality, income risk shocks are important as well. The drivers of consumption inequality are a mixture of the drivers of current income and wealth.

The historical decomposition of US inequality reveals that changing markups and technology are the main contributors to the rise of wealth and income inequality from the 1980s to today. Yet, fiscal policies also play their role both by providing liquid assets for self-insurance through government deficits and by changing the incentives to self-insure through progressive taxation. Both move liquidity premia and thus affect the savings incentives of the rich and the poor differentially. Quantitatively, we find deficits to be less important than changes in progressive taxation. For consumption and income inequality, fluctuations in income risk play a significant role and this role goes beyond increasing the dispersion of income once the higher risk is realized. Wealth poor, and thus badly insured, households react to an increase in uncertainty by cutting consumption particularly strongly, while for well-insured households, which are already consumption rich, behavior changes little. Consequently, these shocks account for 20% of the cyclical variations in consumption inequality. They also account for 20% of aggregate consumption fluctuations in US recessions.

[4] We use the estimates of income risk for the US provided by Bayer et al. (2019) and top marginal tax rates as a proxy for progressivity.

[5] See Kopczuk (2015) and Bricker et al. (2016) for detailed comparisons of all available data sources.

Given the estimated shocks, we assess the importance of policy rules in shaping inequality over the business cycle. We find that policy rules are more important in shaping inequality than policy mistakes. Broadly speaking, output stabilization is key to reducing fluctuations in inequality. A more hawkish monetary policy, i.e., a stronger reaction to inflation, would have increased inequality in the 1970s and today. Both periods – through the lens of our model – are characterized by high markups such that hawkish policy leads to output losses and increases inequality. Countercyclical fiscal policy affects inequality not only by stabilizing output, but also through its effect on returns. Larger and more persistent deficits after the Great Recession would have depressed the liquidity premium and reduced wealth inequality.

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.[6] Auclert et al. (2020) and Hagedorn et al. (2018) 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. Using a state-space approach is key for us, as we need to deal with mixed-frequency data.

Our findings provide new insights into the literature on the drivers of inequality that focuses on long-run trends such as the rising skill premium or changes in taxes.[7] Kaymak and Poschke (2016) and Hubmer et al. (2019), which are most closely related to our approach, use quantitative models to study permanent changes in the US tax and transfer system and the distribution of income. They find that these changes can explain a significant part of the recent increase in wealth inequality. We complement their findings by showing that business cycles have very persistent effects on inequality and can account for 50% of the rise in US wealth inequality from 1980 to 2015. In our estimation, tax progressivity is of secondary importance for wealth inequality relative to changes in markups and technology. In terms of methods, these papers solve for steady-state transitions of calibrated models, while we estimate our model on US macro and micro time series data.

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) is also related. 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]

[6] See, for example, Auclert et al. (2018); Ahn et al. (2018); Bayer et al. (2019); Broer et al. (2019); 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).

[7] There is a growing literature on inequality dynamics. On the theory side, see, e.g., Gabaix et al. (2016). On the empirical side, see, e.g., Heathcote et al. (2010), Piketty and Saez (2003) or Saez and Zucman (2016).

[8] Our approach is different and simpler than the method suggested by Liu and Plagborg-Møller (2019),

Focusing on the methodological contribution, Auclert et al. (2019) provide a fast estimation method for heterogeneous-agent models that requires a sequence space representation of the model and thus does not allow us 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 the fact 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.

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 prior distributions for the remaining parameters that we estimate. It also gives an overview over the data we employ in our estimation. Section 5 discusses the estimated 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 Calvo (1983).

Households earn income from supplying (raw) labor and capital and from owning the firm sector, absorbing all its rents that stem from the 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 taxes on labor income and distributed profits, issues government bonds, and adjusts expenditures to stabilize debt in the long run and aggregate demand in the short run. The monetary authority sets the nominal interest rate on government bonds according to a Taylor rule.

which includes full cross-sectional information in the estimation of a heterogeneous-agent DSGE model. We, in contrast, only use the model to fit certain generalized cross-sectional moments.

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 asset (capital). Trading illiquid assets 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 earning interest on bonds, b_{it}, and potentially from profits or union transfers. Households pay taxes on labor and profit income.

2.1.1 Productivity, labor supply and labor income

A household's gross labor income $w_t n_{it} h_{it}$ 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 endogenous feedback to aggregate effective hours N_{t+1} (hats denote log-deviations from the 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 entrepreneur obtains a fixed share of the pure rents (aside from 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. For tractability, we assume union profits to be taxed at the average income tax rate of the economy.

This modeling strategy serves two purposes. First and foremost, it generally solves the problem of the allocation of pure rents without distorting factor returns and without introducing another tradable asset.[9] Second, we use the entrepreneur state in particular – a transitory state in which incomes are very high – to match the income and wealth distribution following the idea by Castaneda et al. (1998). The entrepreneur state does not change the asset returns or investment opportunities available to households.

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

$$\mathbb{E}_0 \max_{\{c_{it}, n_{it}\}} \sum_{t=0}^{\infty} \beta^t u \left[c_{it} - G(h_{it}, n_{it}) \right]. \tag{4}$$

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

[9] There are basically three possibilities for dealing with the pure rents. One attributes them to capital and labor, but this affects their factor prices; or one introduces a third asset that pays out rents as dividends and is priced competitively; one distributes the rents in the economy to an exogenously determined group of households. The latter has the advantage that factor supply decisions remain the same as in any standard New-Keynesian framework and still avoids the numerical complexity of dealing with three assets.

[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 the labor supply; see, e.g., Born and Pfeifer (2014). It also simplifies the numerical analysis somewhat. Unfortunately, it is not feasible to estimate the flexible form of preference of Jaimovich and Rebelo (2009), which also encompasses 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.

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

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

Assuming a (progressive) income-tax schedule (which we borrow from Benabou, 2002; Heathcote et al., 2017), a household's net labor income, y_{it}, is given by

$$y_{it} = (1 - \tau_t^L)(w_t h_{it} n_{it})^{1-\tau_t^P}, \tag{5}$$

where w_t is the aggregate wage rate and τ_t^L and τ_t^P determine the level and the progressivity of the tax code. Given net labor income, the first-order condition for labor supply is

$$\frac{\partial G(h_{it}, n_{it})}{\partial n_{it}} = (1 - \tau_t^P)(1 - \tau_t^L)(w_t h_{it})^{1-\tau_t^P} n_{it}^{-\tau_t^P} = (1 - \tau_t^P) \frac{y_{it}}{n_{it}}. \tag{6}$$

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 this first-order condition (6) and substitute $G(h_{it}, n_{it})$ out of the individual planning problem:

$$x_{it} = c_{it} - G(h_{it}, n_{it}) = c_{it} - \frac{1 - \tau_t^P}{1 + \gamma} y_{it}. \tag{7}$$

When the Frisch elasticity of labor supply is constant and the tax schedule has the form (5), the disutility of labor is always a fraction of labor income, constant across households. 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.

What remains to be determined is individual and aggregate effective labor supply. Without further loss of generality, we assume $G(h_{it}, n_{it}) = h_{it}^{1-\bar{\tau}^P} \frac{n_{it}^{1+\gamma}}{1+\gamma}$, where $\bar{\tau}^P$ is the stationary equilibrium level of progressivity of the tax code. This functional form simplifies the household problem in the stationary equilibrium 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 we normalized $\int h_{it} di = 1$.

Importantly, this means that we can read off average productivity risk directly from the

estimated income risk series of Bayer et al. (2019). Without scaling the labor disutility by productivity we would need to translate productivity risk to income risk through the endogenous hour response. When tax progressivity does not coincide with its stationary equilibrium value, individual hours worked differ across agents and are given by

$$n_{it} = \left[(1-\tau_t^P)(1-\tau_t^L)\right]^{\frac{1}{\gamma+\tau_t^P}} h_{it}^{\frac{\bar{\tau}^P-\tau_t^P}{\gamma+\tau_t^P}} w_t^{\frac{1-\tau_t^P}{\gamma+\tau_t^P}} , \tag{8}$$

such that aggregate effective hours are given by

$$N_t = \int n_{it} h_{it} = \left[(1-\tau_t^P)(1-\tau_t^L)\right]^{\frac{1}{\gamma+\tau_t^P}} w_t^{\frac{1-\tau_t^P}{\gamma+\tau_t^P}} \underbrace{\int h_{it}^{\frac{\gamma+\bar{\tau}^P}{\gamma+\tau_t^P}}}_{:=H_t} . \tag{9}$$

Here H_t measures how the tax progressivity influences the (hours-weighted) average labor productivity. Scaling of the disutility of labor by $h_{it}^{1-\bar{\tau}}$ is thus a normalization of H_t to one in the stationary equilibrium. It implies that, despite the progressive income tax, changes in the distribution of productivity h have no first-order effect on effective labor supply and thus, as in Bayer et al. (2019), shocks to income risk do not directly move effective labor.

Household after-tax labor income, plugging in the optimal supply of hours, is then:

$$y_{it} = (1-\tau_t^L)(w_t h_{it} n_{it})^{1-\tau_t^P} = (1-\tau_t^L)^{\frac{1+\gamma}{\gamma+\tau_t^P}} (1-\tau_t^P)^{\frac{1-\tau_t^P}{\gamma+\tau_t^P}} w_t^{\frac{1+\gamma}{\gamma+\tau_t^P}(1-\tau_t^P)} h_{it}^{\frac{\gamma+\bar{\tau}^P}{\gamma+\tau_t^P}(1-\tau_t^P)} . \tag{10}$$

2.1.2 Consumption, savings, and portfolio choice

Given this labor income, households optimize intertemporally subject to their budget constraint:

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

where Π_t^U is union profits, Π_t^F is firm profits (both net of taxes), 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 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. Holdings of bonds

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

Substituting the expression $c_{it} = x_{it} + \frac{1-\tau_t^P}{1+\gamma} y_{it}$ for consumption, we obtain the budget constraint for the composite leisure-consumption good:

$$x_{it} + b_{it+1} + q_t k_{it+1} = b_{it} \frac{R(b_{it}, R_t^b, A_t)}{\pi_t} + (q_t + r_t) k_{it} + \frac{\tau_t^P + \gamma}{1+\gamma} y_{it} + \mathbb{I}_{h_{it} \neq 0} \Pi_t^U + \mathbb{I}_{h_{it}=0} \Pi_t^F, \qquad (11)$$
$$k_{it+1} \geq 0, \quad b_{it+1} \geq \underline{B}.$$

Households make their savings choices and their portfolio choice between liquid bonds and illiquid capital in light of a capital market friction that renders capital illiquid because participation in the capital market is random and i.i.d. in the sense that only a fraction, λ, of households is selected to be able to adjust their capital holdings in a given period.

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

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

The extra wedge for unsecured borrowing creates a mass of households with zero unsecured credit but with the possibility to borrow, though at a penalty rate. The intermediation efficiency wedge A_t can be thought of as a cost of a banking sector turning government bonds into deposits. This cost follows an AR(1) process in logs and fluctuates in response to shocks, ϵ_t^A. If A_t goes down, households will implicitly 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).

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: value function V^a for the case where the household adjusts its capital holdings, the function V^n

for the case in which it does not adjust, and the expected envelope value, $\mathbb{E}V$, over both:

$$\begin{aligned} V_t^a(b,k,h) =& \max_{k',b'_a} u[x(b,b'_a,k,k',h)] + \beta \mathbb{E}_t V_{t+1}(b'_a,k',h) \\ V_t^n(b,k,h) =& \max_{b'_n} u[x(b,b'_n,k,k,h)] + \beta \mathbb{E}_t V_{t+1}(b'_n,k,h) \\ \mathbb{E}_t V_{t+1}(b',k',h) =& \mathbb{E}_t \left[\lambda V_{t+1}^a(b',k',h)\right] + \mathbb{E}_t \left[(1-\lambda) V_{t+1}^n(b',k,h)\right] \end{aligned} \tag{12}$$

Expectations about the continuation value are taken with respect to all stochastic processes conditional on the current states, including time-varying income risk. Maximization is subject to the corresponding budget constraint.

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 and then sell them to goods bundlers, who finally sell them as consumption goods to households, and to (d) capital goods 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.[11] 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 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, each of whom offers a different variety of labor to labor packers who then provide labor services to intermediate goods producers. Labor packers produce final labor services according

[11] 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.

to the production function

$$N_t = \left(\int \hat{n}_{jt}^{\frac{\zeta_t - 1}{\zeta_t}} dj \right)^{\frac{\zeta_t}{\zeta_t - 1}}, \tag{13}$$

out of labor varieties $\hat{n}_{jt}$. Cost minimization by labor packers implies that each variety of labor, each union j, faces a downward-sloping demand curve

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

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. However, they face a Calvo-type (1983) of adjustment friction with indexation with the probability λ_w to keep wages constant. They therefore maximize

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \lambda_w^t \frac{W_t^F}{P_t} N_t \left\{ \left(\frac{W_{jt} \bar{\pi}_W^t}{W_t^F} - \frac{W_t}{W_t^F} \right) \left(\frac{W_{jt} \bar{\pi}_W^t}{W_t^F} \right)^{-\zeta_t} \right\}, \tag{14}$$

by setting W_{jt} in period t and keeping it constant except for indexation to $\bar{\pi}_W$, the steady-state wage inflation rate.

Since all unions are symmetric, we focus on a symmetric equilibrium and obtain the linearized 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:

$$\log \left(\frac{\pi_t^W}{\bar{\pi}_W} \right) = \beta \mathbb{E}_t \log \left(\frac{\pi_{t+1}^W}{\bar{\pi}_W} \right) + \kappa_w \left(\frac{w_t}{w_t^F} - \frac{1}{\mu_t^W} \right), \tag{15}$$

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 and $\kappa_w = \frac{(1-\lambda_w)(1-\lambda_w \beta)}{\lambda_w}$. This target fluctuates in response to markup shocks, $\epsilon_t^{\mu W}$, and follows a log AR(1) process.[12]

[12] Including the first-order irrelevant terms, the Phillips curve reads

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

where τ_t is the average income tax.

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

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

for each good j and buy the intermediate good at the nominal price MC_t. As we do for unions, we assume price adjustment frictions à la Calvo (1983) with indexation.

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

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \lambda_Y^t (1-\tau_t^L) Y_t^{1-\tau_t^P} \left\{ \left(\frac{p_{jt}\bar{\pi}_Y^t}{P_t} - \frac{MC_t}{P_t} \right) \left(\frac{p_{jt}\bar{\pi}^t}{P_t} \right)^{-\eta_t} \right\}^{1-\tau_t^P}, \tag{16}$$

with a time constant discount factor.

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

$$\log\left(\frac{\pi_t}{\bar{\pi}}\right) = \beta \mathbb{E}_t \log\left(\frac{\pi_{t+1}}{\bar{\pi}}\right) + \kappa_Y \left(mc_t - \tfrac{1}{\mu_t^Y} \right), \tag{17}$$

where we again dropped all terms irrelevant for a first-order approximation and have $\kappa_Y = \frac{(1-\lambda_Y)(1-\lambda_Y\beta)}{\lambda_Y}$. Here, π_t is the gross inflation rate of final goods, $\pi_t := \frac{P_t}{P_{t-1}}$, $mc_t := \frac{MC_t}{P_t}$ is the real marginal costs, $\bar{\pi}$ is 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.

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 the 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 goods producers. The intermediate goods producer maximizes profits,

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

where r_t^F and q_t are the rental rate of firms and the (producer) price of capital goods respectively. The intermediate goods producer 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:

$$w_t^F = \alpha mc_t Z_t \left(\frac{u_t K_t}{N_t}\right)^{1-\alpha}, \tag{18}$$

$$r_t + q_t \delta(u_t) = u_t (1-\alpha) mc_t Z_t \left(\frac{N_t}{u_t K_t}\right)^{\alpha}. \tag{19}$$

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

$$q_t \left[\delta_1 + \delta_2 (u_t - 1)\right] = (1-\alpha) mc_t Z_t \left(\frac{N_t}{u_t K_t}\right)^{\alpha}, \tag{20}$$

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

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t I_t \left\{ \Psi_t q_t \left[1 - \frac{\phi}{2} \left(\log \frac{I_t}{I_{t-1}} \right)^2 \right] - 1 \right\}, \tag{21}$$

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}.[13]

Optimality of the capital goods production requires (again dropping all terms irrelevant

[13] 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.

up to first order)

$$\Psi_t q_t \left[1 - \phi \log \frac{I_t}{I_{t-1}}\right] = 1 - \beta \mathbb{E}_t \left[\Psi_{t+1} q_{t+1} \phi \log \left(\frac{I_{t+1}}{I_t}\right)\right], \tag{22}$$

and each capital goods producer will adjust its production until (22) is fulfilled.

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

$$K_t - (1 - \delta(u_t)) K_{t-1} = \Psi_t \left[1 - \frac{\phi}{2} \left(\log \frac{I_t}{I_{t-1}}\right)^2\right] I_t \ . \tag{23}$$

The functional form assumption implies that investment adjustment costs are minimized and equal to 0 in the 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, chooses both the average tax rate in the economy as well as tax progressivity, 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-type (1993) rule with interest rate smoothing:

$$\frac{R^b_{t+1}}{\bar{R}^b} = \left(\frac{R^b_t}{\bar{R}^b}\right)^{\rho_R} \left(\frac{\pi_t}{\bar{\pi}}\right)^{(1-\rho_R)\theta_\pi} \left(\frac{Y_t}{Y^*_t}\right)^{(1-\rho_R)\theta_Y} \epsilon^R_t. \tag{24}$$

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 the output gap, where the gap, $\frac{Y_t}{Y^*_t}$, is defined relative to what output would be at stationary equilibrium markups, Y^*_t. $\rho_R \geq 0$ captures interest rate smoothing.

We assume that the government runs a budget deficit and hence accumulates debt governed by a rule (c.f. Woodford, 1995):

$$\frac{B_{t+1}}{B_t} = \left(\frac{B_t}{\bar{B}}\right)^{-\gamma_B} \left(\frac{\pi_t}{\bar{\pi}}\right)^{\gamma_\pi} \left(\frac{Y_t}{Y^*_t}\right)^{\gamma_Y} D_t \ , \qquad D_t = D^{\rho_G}_{t-1} \epsilon^G_t, \tag{25}$$

where D_t is a persistent shock to the government's structural deficit. Besides issuing bonds, the government uses tax revenues T_t, defined below, to finance government consumption, G_t,

and interest on debt. The parameters γ_B, γ_Y, and γ_π measure, respectively, how the deficit reacts to outstanding debt, the ouput gap, and inflation.

The government sets the average tax rate in the economy according to a similar rule

$$\frac{\tau_t}{\bar{\tau}} = \left(\frac{\tau_{t-1}}{\bar{\tau}}\right)^{\rho_\tau} \left(\frac{B_t}{\bar{B}}\right)^{(1-\rho_\tau)\gamma_B^\tau} \left(\frac{Y_t}{Y_t^*}\right)^{(1-\rho_\tau)\gamma_Y^\tau} \epsilon_t^\tau \ . \tag{26}$$

The parameter τ_t^P that governs the progressivity of the tax schedule evolves according to

$$\frac{\tau_t^P}{\bar{\tau}^P} = \left(\frac{\tau_{t-1}^P}{\bar{\tau}^P}\right)^{\rho_P} \epsilon_t^P \ . \tag{27}$$

The level parameter of the tax code τ_t^L adjusts such that the average tax rate on income equals this target level:

$$\tau_t = \frac{\mathbb{E}_t\left(w_t n_{it} h_{it} + \mathbb{I}_{h_{it}=0}\Pi_t^F\right) - \tau_t^L \mathbb{E}_t\left(w_t n_{it} h_{it} + \mathbb{I}_{h_{it}=0}\Pi_t^F\right)^{\tau_t^P}}{\mathbb{E}_t w_t n_{it} h_{it} + \mathbb{I}_{h_{it}=0}\Pi_t^F} \ , \tag{28}$$

where $\mathbb{E}_t$ is the expectation operator, which here gives the cross-sectional average. Total taxes T_t are then $T_t = \tau_t\left(w_t n_{it} h_{it} + \mathbb{I}_{h_{it}\neq 0}\Pi_t^U + \mathbb{I}_{h_{it}=0}\Pi_t^F\right)$ and the government budget constraint determines government spending residually: $G_t = B_{t+1} + T_t - R_t^b/\pi_t B_t$.

There are thus four shocks to government rules: monetary policy shocks, ϵ_t^R, tax progressivity shocks ϵ_t^P, tax level shocks ϵ_t^τ, and structural deficit, i.e., government spending, shocks, ϵ_t^G. We assume these shocks to be log normally distributed with mean zero.

2.4 Goods, Bonds, Capital, and Labor Market Clearing

The labor market clears at the competitive wage given in (18). The bond market clears whenever the following equation holds:

$$B_{t+1} = B^d(R_t^b, A_t, r_t, q_t, \Pi_t^F, \Pi_t^U, w_t, \pi_t, \tau_t, \tau_t^P, \Theta_t, V_{t+1}) := \mathbb{E}_t\left[\lambda b_{a,t}^* + (1-\lambda) b_{n,t}^*\right], \tag{29}$$

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 (and tax rates) $(R_t^b, A_t, r_t, q_t, \Pi_t^F, \Pi_t^U, w_t, \pi_t, \tau_t, \tau_t^P)$. 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:

$$K_{t+1} \quad = \quad K^d(R_t^b, A_t, r_t, q_t, \Pi_t^F, \Pi_t^U, w_t, \pi_t, \tau_t, \tau_t^P, \Theta_t, V_{t+1}) := \mathbb{E}_t[\lambda {k_t}^* + (1-\lambda)k] \ , \quad (30)$$

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, λk_t^*, and those that do not, $(1-\lambda)k$. 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 $\{w_t, w_t^F, \Pi_t^F, \Pi_t^U, q_t, r_t, R_t^b, \pi_t, \pi_t^W, \tau_t, \tau_t^P\}$, a sequence of stochastic states A_t, Ψ_t, Z_t and shocks $\epsilon_t^R, \epsilon_t^G, \epsilon_t^P, \epsilon_t^\tau, \epsilon_t^A, \epsilon_t^Z, \epsilon_t^\Psi, \epsilon_t^{\mu W}, \epsilon_t^{\mu Y}, \epsilon_t^\sigma$, 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 equation (12).

2. Distributions of wealth and income evolve according to households' policy functions.

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

4. Expectations are model consistent.

3 Numerical Solution and Estimation Technique

We solve the model by perturbation methods. We choose a first-order Taylor expansion around the stationary equilibrium following the method of Bayer and Luetticke (2018). This method replaces the value functions with linear interpolants and the distribution functions with 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.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.

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

$$\mathbb{E}_t F(x_t, X_t, x_{t+1}, X_{t+1}, \sigma\Sigma\epsilon_{t+1}), \tag{31}$$

where x_t is "idiosyncratic" states and controls: the value and distribution functions, and X_t is 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-algorithm (2000) (see also Schmitt-Grohé and Uribe, 2004) to obtain a linear state-space representation, 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 $\rho_h, \bar{\sigma}_h, \lambda, \beta, \gamma$, and ξ do, but these parameters we do not estimate but calibrate from the stationary equilibrium.[14] 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

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

functions and the histograms) only affect the aggregate equations through their parameter-free effect on summary variables, such that this part of the Jacobian also does not need to be 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 are 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 is still much larger because of the many additional "idiosyncratic" states (histograms) and controls (marginal 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 filter to obtain the likelihood from the state-space representation of the model solution[15] 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, Data, and Priors

We follow a two-step procedure to estimate the model. First, we calibrate or fix all parameters that affect the steady state of the model. Second, we estimate by full-information methods all parameters that only matter for the dynamics of the model, i.e., the aggregate shocks and frictions. Table 1 summarizes the calibrated and externally chosen parameters and Table 3 lists the estimated parameters. One period in the model refers to a quarter of a year.

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). We set the Frisch elasticity to 0.5; see Chetty et al. (2011). We take estimates

[15]The Kalman filter allows us to deal with missing values and mixed frequency data quite naturally. For a one-frequency data set 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 two-asset HANK 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	0.981	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. (2014)
$\bar{R}$	1.65%	Borrowing penalty	see Table 2
Firms			
α	0.68	Share of labor	62% labor income
δ_0	1.75%	Depreciation rate	7.0% p.a.
$\bar{\eta}$	11	Elasticity of substitution	Price markup 10%
$\bar{\zeta}$	11	Elasticity of substitution	Wage markup 10%
Government			
$\bar{\tau}^L$	0.12	Tax rate level	$G/Y = 20\%$
$\bar{\tau}^P$	0.18	Tax progressivity	Heathcote et al. (2017)
$\bar{R^b}$	1.004	Nominal rate	1.6% p.a.
$\bar{\pi}$	1.00	Inflation	0% p.a.

for idiosyncratic income risk from Storesletten et al. (2004), $\rho_h = 0.98$ and $\bar{\sigma}_h = 0.12$. Guvenen et al. (2014) provide the probability that a household will fall out of the top 1% of the income distribution in a given year, which we take as the transition probability from entrepreneur to worker, $\iota = 1/16$.

Table 2 summarizes the calibration of the remaining household parameters. We match 4 targets: 1) average illiquid assets (K/Y=11.44), 2) average liquidity (B/Y=1.58), 3) the fraction of borrowers, 16%, and 4) the average top 10% share of wealth, which is 67%. This yields a discount factor of 0.981, a portfolio adjustment probability of 6.5%, borrowing penalty of 1.65% quarterly (given a borrowing limit of two times average quarterly income), and a transition probability from worker to entrepreneur of 1/5000.[16]

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 1.75% per quarter. 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%. All are standard values in

[16]Detailed data sources can be found in Appendix A.

Table 2: Calibrated parameters

Targets	Model	Data	Source	Parameter
Mean illiquid assets (K/Y)	11.44	11.44	NIPA	Discount factor
Mean liquidity (B/Y)	1.58	1.58	FRED	Port. adj. probability
Top10 wealth share	0.67	0.67	WID	Fraction of entrepreneurs
Fraction borrowers	0.16	0.16	SCF	Borrowing penalty

the literature.

The government taxes labor and profit income using a non-linear tax schedule that approximates the progressivity of the US tax system; see Heathcote et al. (2017). The progressivity parameter, $\tau^P = 0.18$, is taken from Heathcote et al. (2017). The level of taxes, τ^L, is set to clear the government budget constraint that corresponds to a government share of $G/Y = 20\%$. The policy rate is set to an annualized rate of 1.6%. 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 eight observable time series: the growth rates of per capita GDP, private consumption, investment, federal tax receipts, and wages, all in real terms, the logarithm of the level of per capita hours worked, the log difference of the GDP deflator, and the (shadow) federal funds rate. We add more data with shorter and/or non-quarterly availability. Idiosyncratic income uncertainty (estimated from panel data in the Survey of Income and Program Participation (SIPP) in Bayer et al., 2019) is available as quarterly series from 1983Q1 to 2013Q1 and included in log-levels. We proxy the progressivity of the US tax and transfer system by the highest bracket of the US individual income tax. This series is available at annual frequency from 1954 to 2015. Wealth and income shares of the top 10% are included at annual frequency and available from 1954 to 2014 from the World Inequality Database (drawing on work by Piketty, Saez, and Zucman; see, e.g., Piketty and Saez (2003) or Saez and Zucman (2016)).

4.3 Prior Distributions

Columns 1-4 of Table 3 presents the parameters we estimate and their assumed prior distributions. The posterior distribution is discussed in the next section. Where available, we use prior values that are standard in the literature and independent of the underlying data. 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_Y 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. 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% and standard deviation 2%. The only exception is the uncertainty shock, where, given the evidence in Bayer et al. (2019), we use a higher prior mean of 1.0. The employment feedback parameter in the uncertainty process is assumed to follow a normal prior with large variance.

Regarding policy, 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. In the bond rule, the debt-feedback parameter γ_B is assumed to follow a gamma distribution with mean 0.10 and standard deviation 0.08, such that the prior for the autocorrelation of debt is centered around 0.9, implying a half-life of a deviation in debt of between one and eight years. The parameters governing feedback to inflation and output, γ_π and γ_Y, follow standard normal distributions. Similarly, the autoregressive parameters, in the tax rules, ρ_i where $i \in \{P, \tau\}$, are assumed to follow beta distributions (with mean 0.5 and standard deviation 0.2), while the feedback parameters, γ_Y^τ and γ_B^τ, follow standard normal distributions.

The standard deviations of the measurement errors are assumed to have inverse-gamma prior distributions. For the uncertainty series, we set a relatively high prior mean of 5.00%, while for all other measurement errors we set lower prior means of 0.05%. As the top marginal tax rate does not directly map into how we model progressivity, we allow for a linear scaling parameter $\propto^\tau$ (Boivin and Giannoni, 2006), following a standard normal prior, in addition to the measurement error.

Table 3: Prior and posterior distributions of estimated parameters

Parameter		Prior		Posterior HANK				Posterior HANK*			
	Distribution	Mean	Std. Dev.	Mean	Std. Dev.	5 %	95 %	Mean	Std. Dev.	5%	95%
					Frictions						
δ_s	Gamma	5.00	2.00	1.483	0.137	1.265	1.717	1.420	0.031	1.371	1.472
ϕ	Gamma	4.00	2.00	0.233	0.032	0.182	0.288	0.218	0.036	0.162	0.278
κ	Gamma	0.10	0.02	0.101	0.014	0.080	0.126	0.105	0.014	0.083	0.129
κ_w	Gamma	0.10	0.02	0.128	0.017	0.100	0.157	0.133	0.019	0.103	0.164
					Debt and monetary policy rules						
ρ_R	Beta	0.50	0.20	0.800	0.016	0.774	0.825	0.803	0.014	0.779	0.826
σ_R	Inv.-Gamma	0.10	2.00	0.265	0.015	0.241	0.291	0.266	0.016	0.242	0.293
θ_π	Normal	1.70	0.30	2.603	0.125	2.404	2.817	2.614	0.053	2.520	2.695
θ_y	Normal	0.13	0.05	0.086	0.020	0.054	0.119	0.078	0.018	0.048	0.107
γ_B	Gamma	0.10	0.08	0.144	0.018	0.116	0.174	0.157	0.021	0.119	0.190
γ_π	Normal	0.00	1.00	-1.134	0.031	-1.185	-1.083	-1.175	0.025	-1.215	-1.135
γ_Y	Normal	0.00	1.00	-0.716	0.030	-0.765	-0.668	-0.697	0.026	-0.741	-0.655
ρ_G	Beta	0.50	0.20	0.988	0.008	0.972	0.997	0.992	0.005	0.981	0.998
σ_G	Inv.-Gamma	0.10	2.00	0.260	0.015	0.237	0.286	0.263	0.016	0.238	0.289
					Tax rules						
ρ_τ	Beta	0.50	0.20	0.550	0.029	0.502	0.597	0.552	0.058	0.457	0.648
σ_τ	Inv.-Gamma	1.00	2.00	0.127	0.156	0.022	0.537	0.122	0.171	0.021	0.544
γ_B^τ	Normal	0.00	1.00	0.777	0.062	0.674	0.883	0.833	0.075	0.697	0.953
γ_Y^τ	Normal	0.00	1.00	2.646	0.176	2.360	2.943	2.496	0.134	2.303	2.755
ρ_P	Beta	0.50	0.20	0.986	0.006	0.974	0.995	0.988	0.006	0.978	0.996
σ_P	Inv.-Gamma	1.00	2.00	2.220	0.629	1.354	3.422	2.014	0.532	1.289	3.027
$\propto^\tau$	Normal	0.00	1.00	2.226	0.513	1.503	3.177	2.408	0.510	1.648	3.309

Table 3: Prior and posterior distributions of estimated parameters - continued

Parameter		Prior		Posterior HANK				Posterior HANK*			
	Distribution	Mean	Std. Dev.	Mean	Std. Dev.	5 %	95 %	Mean	Std. Dev.	5%	95%
					Structural Shocks						
ρ_A	Beta	0.50	0.20	0.977	0.011	0.957	0.992	0.983	0.009	0.965	0.995
σ_A	Inv.-Gamma	0.10	2.00	0.160	0.014	0.139	0.184	0.159	0.013	0.140	0.181
ρ_Z	Beta	0.50	0.20	0.995	0.002	0.991	0.999	0.992	0.003	0.986	0.996
σ_Z	Inv.-Gamma	0.10	2.00	0.601	0.028	0.558	0.649	0.608	0.029	0.564	0.658
ρ_Ψ	Beta	0.50	0.20	0.976	0.008	0.963	0.988	0.965	0.009	0.950	0.978
σ_Ψ	Inv.-Gamma	0.10	2.00	2.723	0.229	2.362	3.107	2.531	0.214	2.197	2.895
ρ_μ	Beta	0.50	0.20	0.889	0.022	0.852	0.923	0.900	0.021	0.863	0.933
σ_μ	Inv.-Gamma	0.10	2.00	1.695	0.149	1.471	1.958	1.645	0.139	1.435	1.889
$\rho_{\mu w}$	Beta	0.50	0.20	0.909	0.020	0.873	0.938	0.909	0.020	0.875	0.938
$\sigma_{\mu w}$	Inv.-Gamma	0.10	2.00	5.355	0.513	4.605	6.272	5.216	0.541	4.441	6.158
					Income Risk Process						
ρ_s	Beta	0.50	0.20	0.657	0.032	0.602	0.708	0.639	0.038	0.577	0.699
σ_s	Inv.-Gamma	1.00	2.00	63.184	4.257	56.561	70.559	61.443	4.177	55.144	68.804
Σ_N	Normal	0.00	100.00	0.834	0.127	0.615	1.038	0.634	0.086	0.469	0.756
					Measurement Errors						
σ_s^{me}	Inv.-Gamma	5.00	10.00	4.267	3.208	1.217	11.126	4.076	3.164	1.195	10.412
σ_T^{me}	Inv.-Gamma	0.05	0.01	3.695	0.170	3.427	3.985	3.707	0.174	3.433	3.999
$\sigma_{\tau p}^{me}$	Inv.-Gamma	0.05	0.01	0.056	0.065	0.012	0.193	0.038	0.027	0.012	0.091
σ_{WI}^{me}	Inv.-Gamma	0.05	0.01	–	–	–	–	3.772	0.402	3.168	4.475
σ_{II}^{me}	Inv.-Gamma	0.05	0.01	–	–	–	–	8.074	0.819	6.845	9.514

Notes: The standard deviations of the shocks and measurement errors have been transformed into percentages by multiplying by 100. HANK and HANK* denote posterior estimates for the model without and with observable inequality series, respectively.

5 US Business Cycles

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, does the inclusion of measures of inequality change what the model infers about shocks and frictions in business cycles? Second, to what extent do business cycle shocks explain the movements in inequality measures? Third, how would inequality have developed if government business cycle policies had been different?

To answer these questions, we estimate the HANK model with and without 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 Inequality Database. The reason we focus on the top 10% wealth and income share is that this measure is most consistent across alternative, but less frequently available, data sources such as the Survey of Consumer Finances (SCF); see Kopczuk (2015).[17]

In this section, we answer the first question by comparing parameter estimates, variance decompositions, and historical decompositions of US business cycles for the estimated HANK model with and without data on inequality. We postpone the model implications for US inequality to Section 6.

5.1 Parameter Estimates w/ and w/o Inequality

Table 3 reports the posterior distributions across the two main estimation variants: HANK without data on inequality and HANK* with data on inequality.[18] Strikingly, the parameter estimates with and without data on inequality are basically the same; none of the estimated parameters is substantially different across the two estimations. Only some parameters of the fiscal policy rule $(\gamma_B^\tau, \gamma_Y^\tau)$ have significantly different posterior distributions. Using inequality data leads the model to view tax policy as slightly more aggressive in stabilizing government debt.

This implies that both inequality measures provide little additional identification of busi-

[17] We abstain from including other cross-sectional data in the estimation, such as the Panel Study of Income Dynamics (PSID) or the Survey of Consumer Finances (SCF) to avoid dealing with two measurements of the same model variable. Since the income risk data we use for both sets of estimation exercises do contain some cross-sectional information on income inequality (it is constructed from SIPP data), we also consider a third estimation exercise where we also drop the income risk data. Results are available in Appendix E and are in general highly similar to the other two sets of estimation results.

[18] The estimation is conducted with 5 parallel RWMH chains started from an over-dispersed target distribution after an extensive mode search. After burn in, 300,000 draws from the posterior are used to compute the posterior statistics. The acceptance rates across chains are between 26% and 29%. Appendix F provides Gelman and Rubin (1992) convergence statistics. Traceplots of individual parameters are available on request.

ness cycle shocks and frictions. In the next section on US inequality, we show that already the model estimated only on aggregate data implies a U-shaped evolution of inequality from 1950 to 2015 in line with the data. This explains why adding data on inequality has little effect on the estimated parameters. The estimated shocks and frictions do a good job in matching the evolution of wealth and income inequality over the last 60 years. We will explore this result in detail in the next section.

The parameter estimates are broadly in line with the representative-agent literature (which corresponds to our priors that are taken from this literature). Real frictions are an exception. They are up to one order of magnitude smaller in our estimation. In particular, investment adjustment costs are substantially smaller. This reflects the portfolio adjustment costs at the household level that generate inertia in aggregate investment. Our estimates for nominal frictions are standard and close to the priors, with price and wage stickiness being less than 4 quarters on average. In terms of shocks, the estimated persistence and variance for the seven "standard" shocks are comparable to the results of Smets and Wouters (2007). The persistence ranges from 0.995 for TFP to 0.889 for wage markups. The variance ranges from 0.2% for risk premium shocks to 5% for wage-markup shocks.

Idiosyncratic income risk as a driver of portfolio allocations and consumption demand has been highlighted in Bayer et al. (2019). Our estimate for the shock series mostly coincides with the observed series for income risk, i.e., we find income risk shocks to be slightly less persistent than the risk series itself $\rho_\sigma = 0.66$ and the variance slightly larger at 63%. While our estimates imply that there is endogenous amplification of uncertainty, the feedback is negligible in economic terms; see Appendix C for the historical time series of income risk implied by the model. It is important to note here that the income process we use focuses on the uncertainty of persistent shocks to income. As we do not model job-search and unemployment, we abstract from the short-run and transitory income risk fluctuations that others have discussed (e.g., Ravn and Sterk, 2017; Den Haan et al., 2017).

In terms of the estimated policy coefficients, the estimated Taylor rule is in line with the literature. The coefficients on inflation and output deviations are 2.6 and 0.1, and there is substantial inertia 0.8. The fiscal rule that governs deficits and hence government spending exhibits a countercyclical response to inflation and output deviations, -1.1 and -0.7, and features persistence as well. The tax rule that governs average taxation has similar properties. Average tax rates rise when output or debt is high, but not very persistently. Changes to tax progressivity, by contrast, are very persistent, 0.99.

Figure 1: Variance decompositions: Output growth and its components

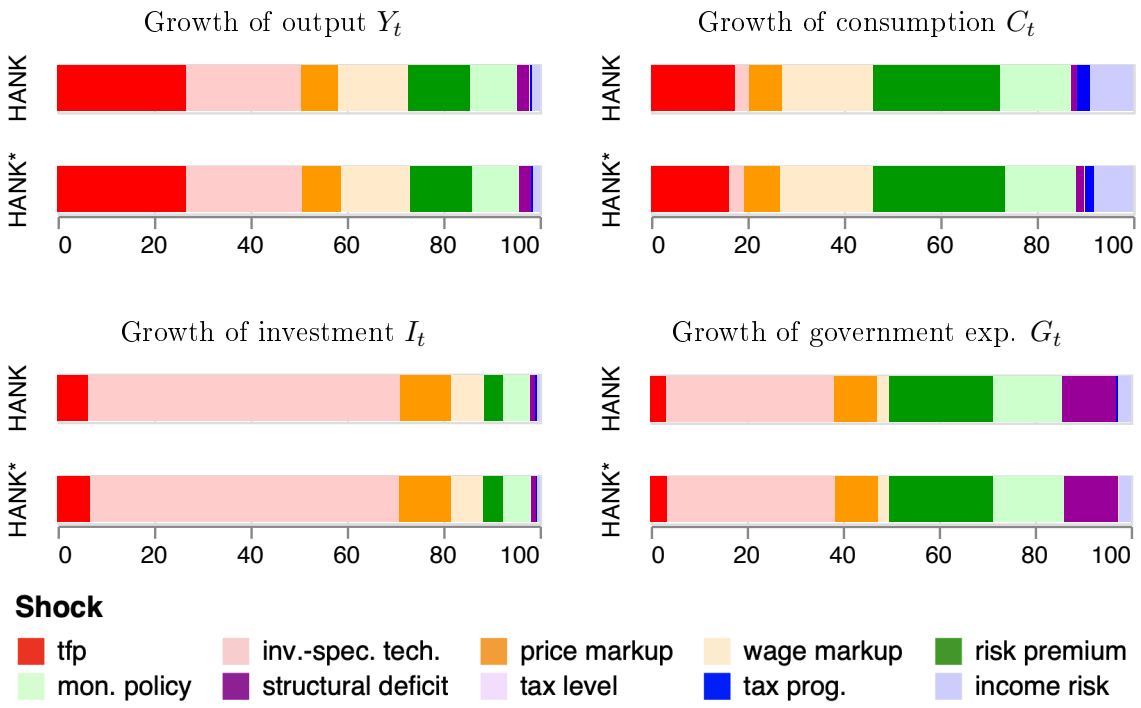

Notes: Conditional variance decompositions at a 4-quarter forecast horizon. HANK* [HANK] corresponds to the estimated HANK model with inequality data [w/o inequality data].

5.2 Variance Decompositions of Business Cycles

Next, we show what the estimated parameters imply for our view of US business cycles by looking at variance decompositions at business cycle frequency.[19] Figure 1 shows these decompositions for the growth rates of output, consumption, investment, and government spending. Unsurprisingly, we find very similar decompositions for the estimations with and without using inequality data. As in the representative-agent literature, TFP and investment-specific technology shocks are the most important drivers of output growth each explaining roughly 25% of its conditional variance at business cycle frequency. All together, supply side shocks (the two markup and the two productivity shocks)[20] account for almost 75% of output volatility.

Demand side shocks (i.e., shocks to uncertainty, fiscal policy, monetary policy, and the

[19] We report the conditional variances at 4 quarters.

[20] The grouping of the shocks is based on the question of 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.

Table 4: Contribution of shocks to US recessions

	Output growth		Consumption growth	
Shocks	HANK	HANK*	HANK	HANK*
TFP, ϵ^Z	-0.23	-0.27	-0.16	-0.18
Inv.-spec. tech., ϵ^Ψ	-0.25	-0.22	0.14	0.15
Price markup, $\epsilon^{\mu Y}$	-0.03	-0.03	-0.10	-0.11
Wage markup, $\epsilon^{\mu W}$	-0.31	-0.31	-0.52	-0.50
Risk premium, ϵ^A	-0.27	-0.27	-0.39	-0.38
Monetary policy, ϵ^R	0.12	0.13	0.23	0.23
Structural deficit, ϵ^G	0.02	0.02	-0.03	-0.07
Tax level, ϵ^τ	0.00	0.00	0.00	0.00
Tax progressivity, ϵ^{τ^P}	0.01	0.01	0.00	0.00
Income risk, ϵ^σ	-0.06	-0.05	-0.18	-0.14

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% for output and -0.5% for consumption.

risk premium), by contrast, explain around 55% of consumption volatility. A new source of fluctuations in consumption is income risk shocks. They explain 10% of the volatility of consumption growth. This is so because income risk is mostly exogenous.

5.3 Historical Decompositions of US Business Cycles

While the variance decompositions help us understand the average cycle implied by the model, a historical decomposition tells us how the model views the actual cycles that the US economy has gone through. We summarize the historical decomposition of NBER-dated recessions in Table 4.[21] During the average NBER-dated recession quarter, 0.82% of a 1.0% decline in output results from technology and markup shocks. In particular, investment-specific technology shocks are important with a contribution of 0.25%. Risk premium and income risk shocks explain another 0.33% of a 1.0% decline in output. Monetary policy

[21] We report historical decompositions for output, consumption, investment, and government expenditures from 1954 to 2015 in Appendix C.

Figure 2: US inequality – data vs. model

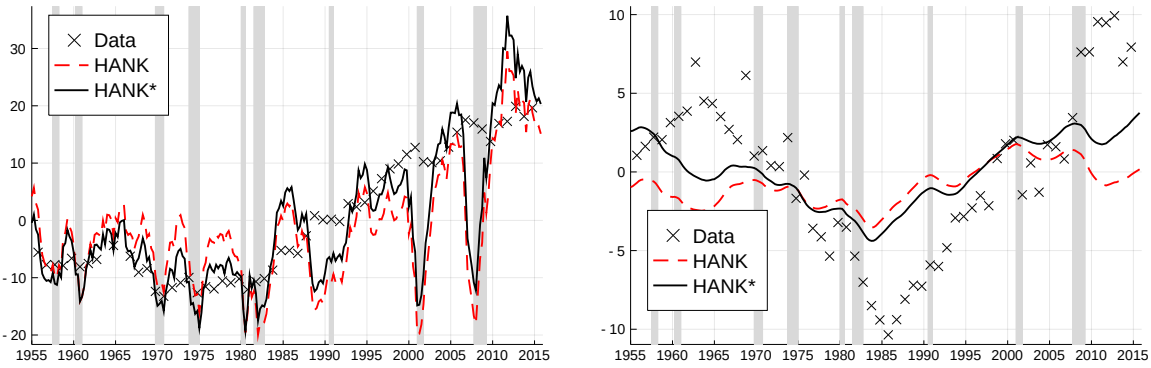

(a) Top 10% income share (b) Top 10% wealth share

Notes: Data (crosses) correspond to log-deviations of the annual observations of the share of pre-tax income and wealth held by the top 10% in each distribution in the US taken from the World Inequality Database. HANK* (solid) [HANK (dashed)] corresponds to the smoothed states of both implied by the estimated HANK model with inequality data [w/o inequality data]. Shaded areas correspond to NBER-dated recessions.

shocks are estimated to be, on average, slightly more expansionary during recessions than the Taylor rule implied rate would suggest – policy shocks contribute positively to output growth in recessions.

This fact is even stronger for consumption. Similarly, investment-specific technology shocks stabilize consumption in recessions as households invest less upon an adverse investment-specific technology shock. In consequence, the decline in consumption during recessions is to a larger extent driven by risk premia and income risk – 0.6% out of a 1% decline.

6 US Inequality

Now, we show that the estimated business cycle shocks can explain the movements of wealth and income inequality in the US. Figure 2 plots the inequality data and the model implied smoothed states for the estimation with and without inequality data. Both data series are available on an annual basis almost throughout our whole sample period (1954-2014). 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 12 percentage points from 1980 to 2015, and the model gets 50% of this increase. The top 10% (pre-tax) income share increases by 32 percentage points over the same time period in the data, and the model predicts an increase by almost

40 percentage points.[22] Business cycle shocks can move inequality along the lines of what we observe in the data. This matching of the distributional data, on top of the "standard" macroeconomic time series, does not change significantly what we infer about shocks and frictions; see the previous section. Section 6.2 will provide a more detailed account of the driving forces behind this. The business cycle analysis requires that both wage markups and price markups be increasing in lockstep up until the mid 1970s. The decade after, both markups fall, leading to the protracted boom of the 1980s. As a result of this comovement, income inequality remains roughly constant and below its long-term average. The result is a fall in wealth inequality. The increase in inequality over the most recent three decades is estimated to be a result of rising price markups that are not accompanied by higher wage markups this time. In addition, income risks, i.e., idiosyncratic productivity risks, have been increasing over the last three decades.

6.1 Propagation of Inequality

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 3 shows, as an example, for markup and income risk shocks.[23] 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 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 dissave. 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.

Income risk shocks move the consumption Gini immediately and strongly. Poor households cut back their consumption to accumulate liquidity in response to the shock. As a result, wealth becomes initially more equally distributed. Yet, as demand falls and markups increase, income inequality goes up immediately, too. Later, when the increase in income risk leads to more dispersed realizations of income, income inequality remains high and wealth inequality overshoots its long-run level.

[22] The fact that the model fits the actual movement of income inequality well reaffirms our modeling choices regarding the distribution of income from various sources – in particular to attribute rents to the income

Figure 3: Impulse responses of inequality

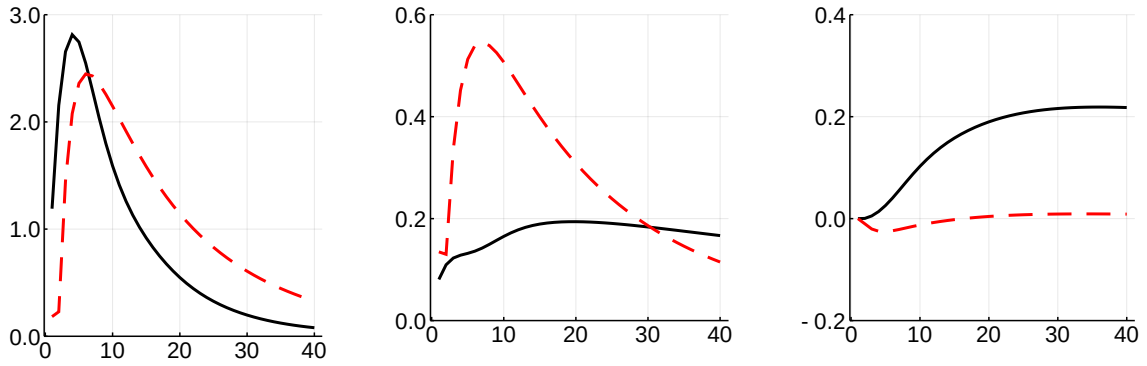

(a) Top 10% income share (b) Consumption Gini (c) Top 10% wealth share

Notes: Response (in percent) of the top 10% pre-tax income share, the consumption Gini, and the top 10% wealth share to a one standard deviation price-markup shock (solid, black) and income-risk shock (dashed, red). The response of the Gini coefficient/top 10% shares is calculated by including them as a generalized moment in the linearized model.

6.2 Historical Decompositions of US Inequality

To dig into the details of the evolution of inequality, Figure 4(a) plots the historical decomposition of the top 10% income share. 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. Looking at 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.[24]

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 in 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 onward. 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

rich.

[23] In Appendix D, we report the IRFs of wealth, income, and consumption inequality for all 10 shocks.

[24] This finding is consistent with other papers that study the Great Recession in particular; see, e.g., Perri and Steinberg (2012) or McKay (2017).

Figure 4: Historical decompositions of US inequality

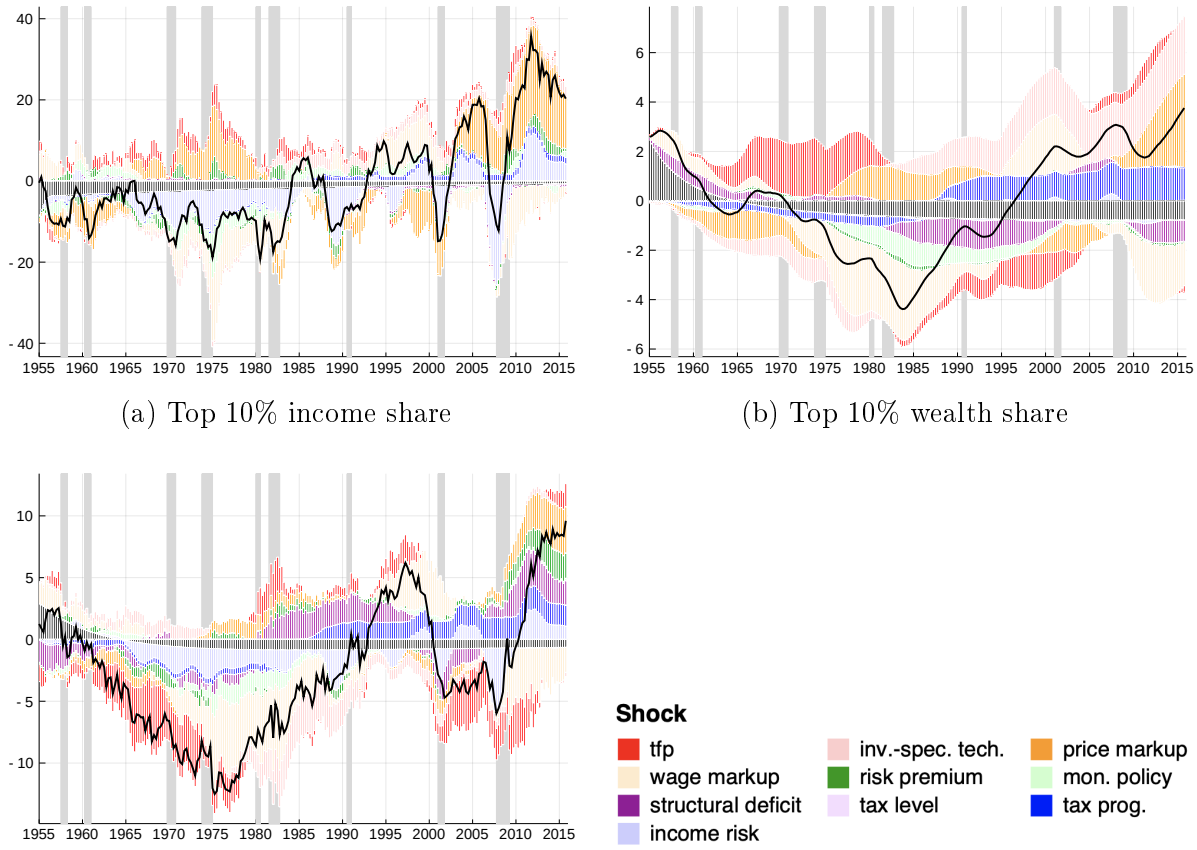

(a) Top 10% income share

(b) Top 10% wealth share

(c) Gini coefficient of consumption

Notes: Historical decomposition of the log-deviations of the top 10% share in pre-tax income (top left), top 10% wealth share (top right), and the Gini coefficient of consumption (bottom). These inequality measures are treated as generalized moments that are included as controls into the state-space representation of the model. Shaded areas correspond to NBER dated recessions.

use of completely different data sources, the historical decomposition thus is in line with the evidence by De Loecker and Eeckhout (2020) on the evolution of markups in the US.[25]

For the evolution of wealth inequality other shocks are important as well. Figure 4(b) shows the historical decomposition of the top 10% wealth share. 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 that up until the end of the 1970s, wage-markup, investment-specific technology, and monetary shocks are the strongest downward drivers of wealth inequality. From the 1980s on, it is then mainly shocks to investment-

[25]There is a growing literature on the rise of markups; see, e.g., Karabarbounis and Neiman (2019), Barkai (2019), Hall (2018), or Kehrig and Vincent (2018).

Table 5: Contribution of shocks to US inequality 1980-2015

Shock	Top 10% Income	Top 10% Wealth	Gini Consumption
TFP, ϵ^Z	-0.30	-1.14	-0.33
Inv.-spec. tech., ϵ^Ψ	1.32	2.63	1.62
Price markup, $\epsilon^{\mu Y}$	11.99	2.44	2.50
Wage markup, $\epsilon^{\mu W}$	8.83	0.90	3.46
Risk premium, ϵ^A	-0.51	0.06	1.72
Monetary policy, ϵ^R	6.81	0.72	2.10
Structural deficit, ϵ^G	-0.84	-0.94	1.55
Tax level, τ	0.00	0.00	0.00
Tax progressivity, τ^P	2.06	1.70	2.23
Income risk, ϵ^σ	8.99	0.07	2.35
Sum of shocks	39.56	6.25	17.40

Notes: The table displays the contribution (in p.p.) of the various shocks to the increase in the top 10% share of pre-tax income, top 10% share of wealth, and the Gini coefficient of consumption from 1980 to 2015 based on our historical shock decompositions.

specific technology and fiscal policy (deficits and tax progressivity) that drive up wealth inequality. Investment-specific technology and fiscal deficits matter for wealth inequality because they affect the return spread between the liquid and illiquid asset. Only since the 2000s have rising price markups become a strong positive contributor to wealth inequality.

Finally, Figure 4(c) plots the historical decomposition of the Gini coefficient of consumption. 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 markups and fiscal policy. 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 5 summarizes the driving forces behind the increase in all three inequality measures from 1980 to 2015. The business cycle shocks in our model capture virtually all the observed

Figure 5: Variance decompositions: Inequality

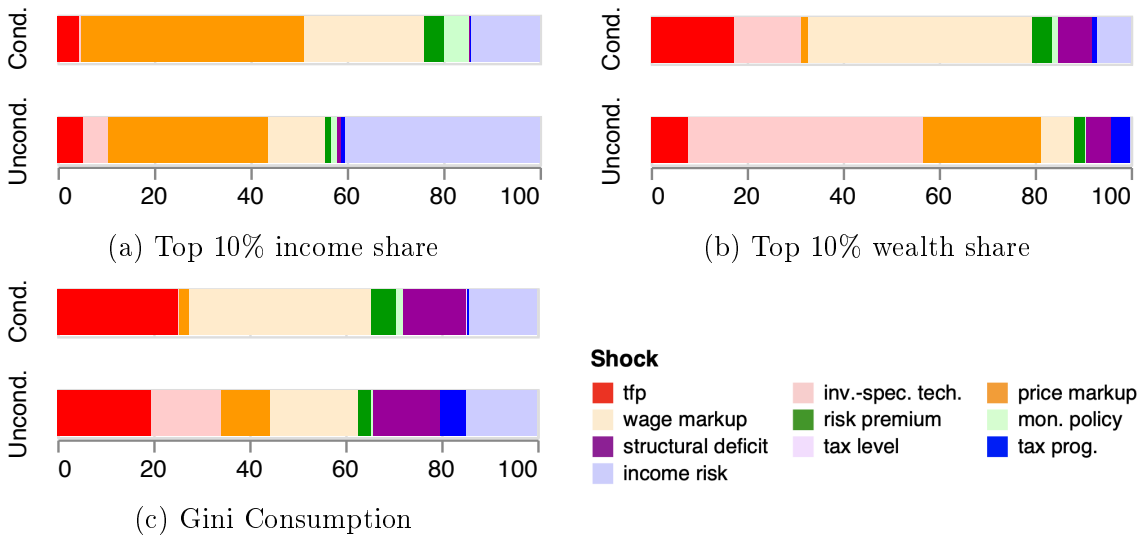

(a) Top 10% income share

(b) Top 10% wealth share

(c) Gini Consumption

Notes: Unconditional and conditional (4-quarter horizon) variance decompositions of the top 10% share of pre-tax income, top 10% share of wealth, and the Gini coefficient of consumption.

increase in income inequality and roughly half of the increase in wealth inequality. The estimated model somewhat misses the trough in the 1980s and does not fully capture the rising wealth inequality after the Great Recession. The likely main reason seems that our model misses the full complexity of household portfolios and the late divergent returns on houses and other forms of capital; see Kuhn et al. (2020). Figure 5 shows that, in general, our findings from the historical decompositions also hold true for the average business cycle in terms of variance decompositions.

These findings challenge the prevailing literature on the drivers of inequality that focuses on long-run trends such as the rising skill premium or changes in the tax and transfer system (see, e.g., Kaymak and Poschke (2016) or Hubmer et al. (2019)). We add to this literature by showing that the business cycle has very persistent effects on inequality and can account for up to 50% of the rise in US wealth inequality from 1980 to 2015. In the estimation of our model, we allow for very persistent shocks to tax progressivity that capture the decline in progressivity from the 1970s onward.[26] We find that the decline of tax progressivity explains 1.7 p.p. of the increase in wealth inequality since the 1980s; see Table 5. This makes the design of the tax system more important for the historical decomposition of wealth inequality than government spending or monetary shocks. Changes in tax progressivity are the

[26]See Appendix C for the implied time series.

third most important driver of wealth inequality after price markups and investment-specific technology, which account for 2.5 p.p. and 2.1 p.p respectively. One reason for the muted effect on inequality is that lower progressivity leads to a higher capital stock through its interaction with portfolio heterogeneity, because wealthy households have a higher propensity to invest in capital; see Luetticke (2018). This reduces dividends and increases wages, which mitigates the effect on inequality.

6.3 Inequality Counterfactuals

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 monetary and fiscal policy experiments based on the estimated model. The results of these experiments are displayed in Figures 6 and 7. In detail, the figures display 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; see Figure 6. 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 the fact 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, importantly, a substantial fraction of shocks during this time are demand shocks. For the Great Recession, which we estimate to be followed by a persistent increase in the demand for government bonds (risk premium shock) alongside an increase in markups, the more hawkish monetary policy would have helped the recovery and lowered inequality. Interestingly, the commitment to a more aggressive policy reaction requires a smaller interest rate movement in response to the increase in the demand for bonds. Second, we consider a dovish policy where we double the monetary policy response to output fluctuations. This leads in general to more stable 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 to fluctuations in output, not inflation. In fact, it generates output fluctuations in the 1990s, while it does little to change all series in the 1970s. For the Great Recession, a more dovish policy stance would have led to an earlier recovery and in particular lower income inequality. The effects on wealth inequality are milder.

Figure 6: Counterfactual evolution of output and income, wealth, and consumption inequality: Monetary policy

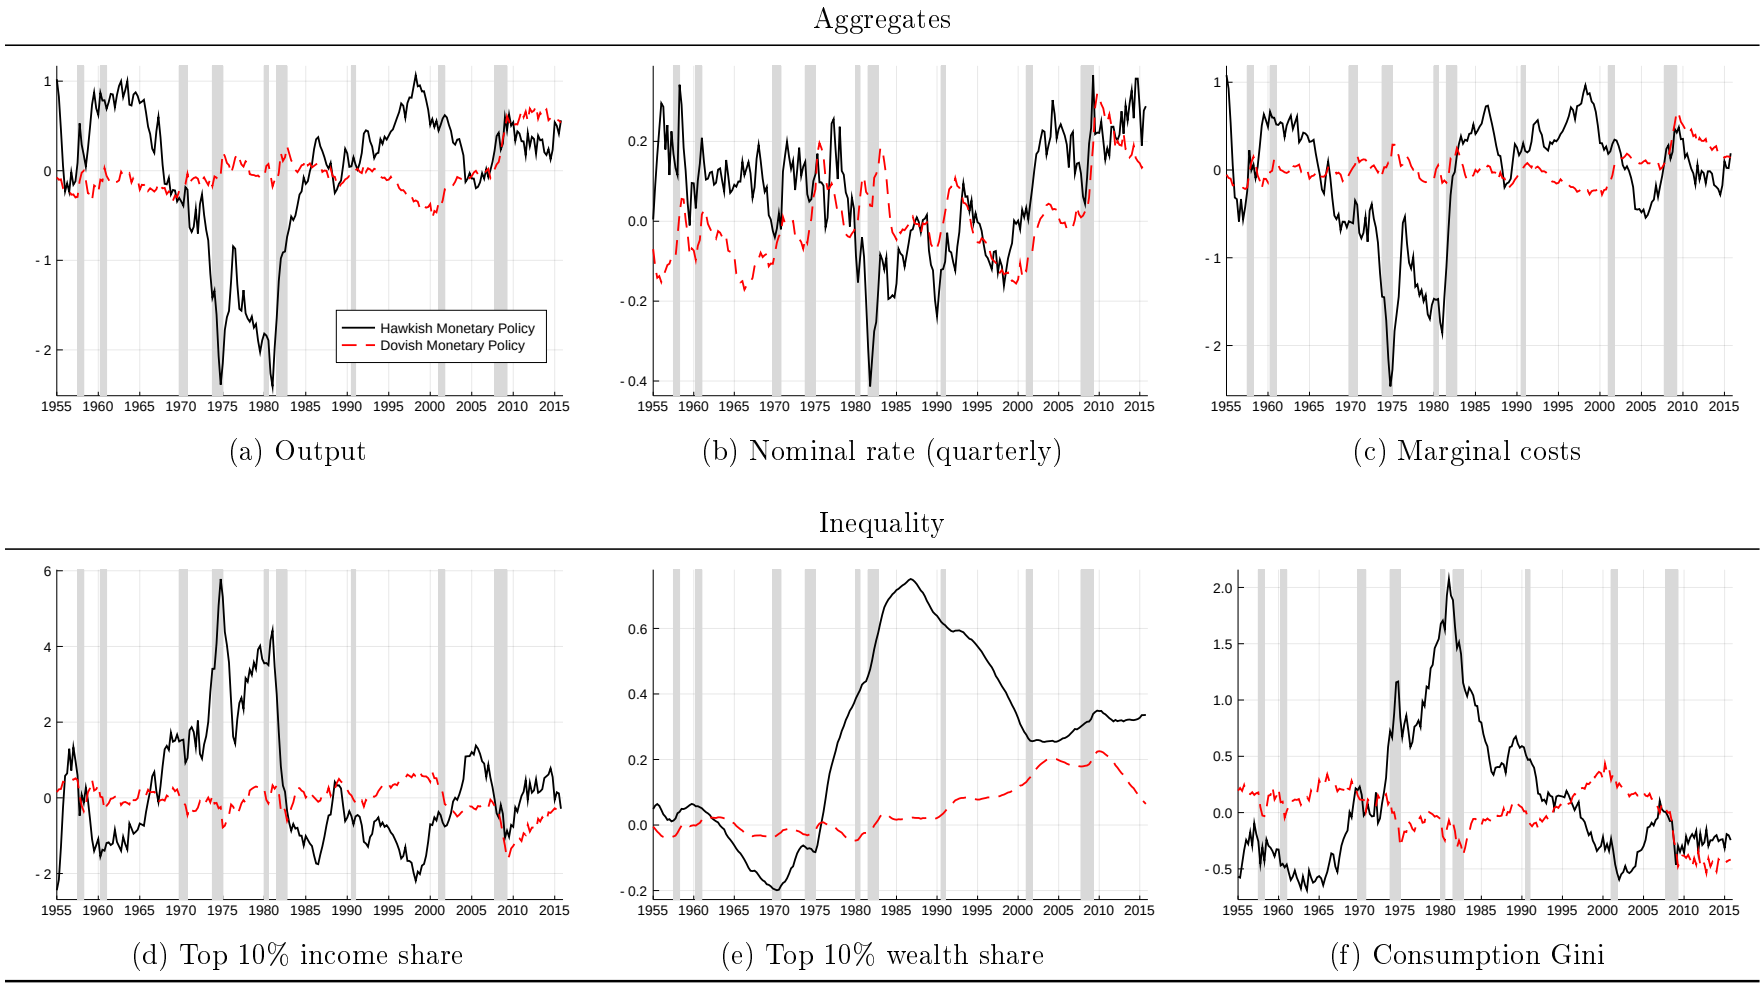

(a) Output (b) Nominal rate (quarterly) (c) Marginal costs

(d) Top 10% income share (e) Top 10% wealth share (f) Consumption Gini

Notes: The panels display the evolution of output, the nominal interest rate, and marginal costs as well as wealth, income, and consumption inequality that the model would counterfactually predict had the government policies been different, feeding the smoothed sequence of shocks (as in Figure 4) through the model. The lines represent the difference (in p.p.) in the evolution compared to feeding the same shocks through the baseline model. The solid line corresponds to a setup where we double the inflation response θ_π. The dashed line reflects the counterfactual where we double the estimated response to output, θ_y. Shaded areas correspond to NBER-dated recessions.

Figure 7: Counterfactual evolution of output and income, wealth, and consumption inequality: Fiscal policy

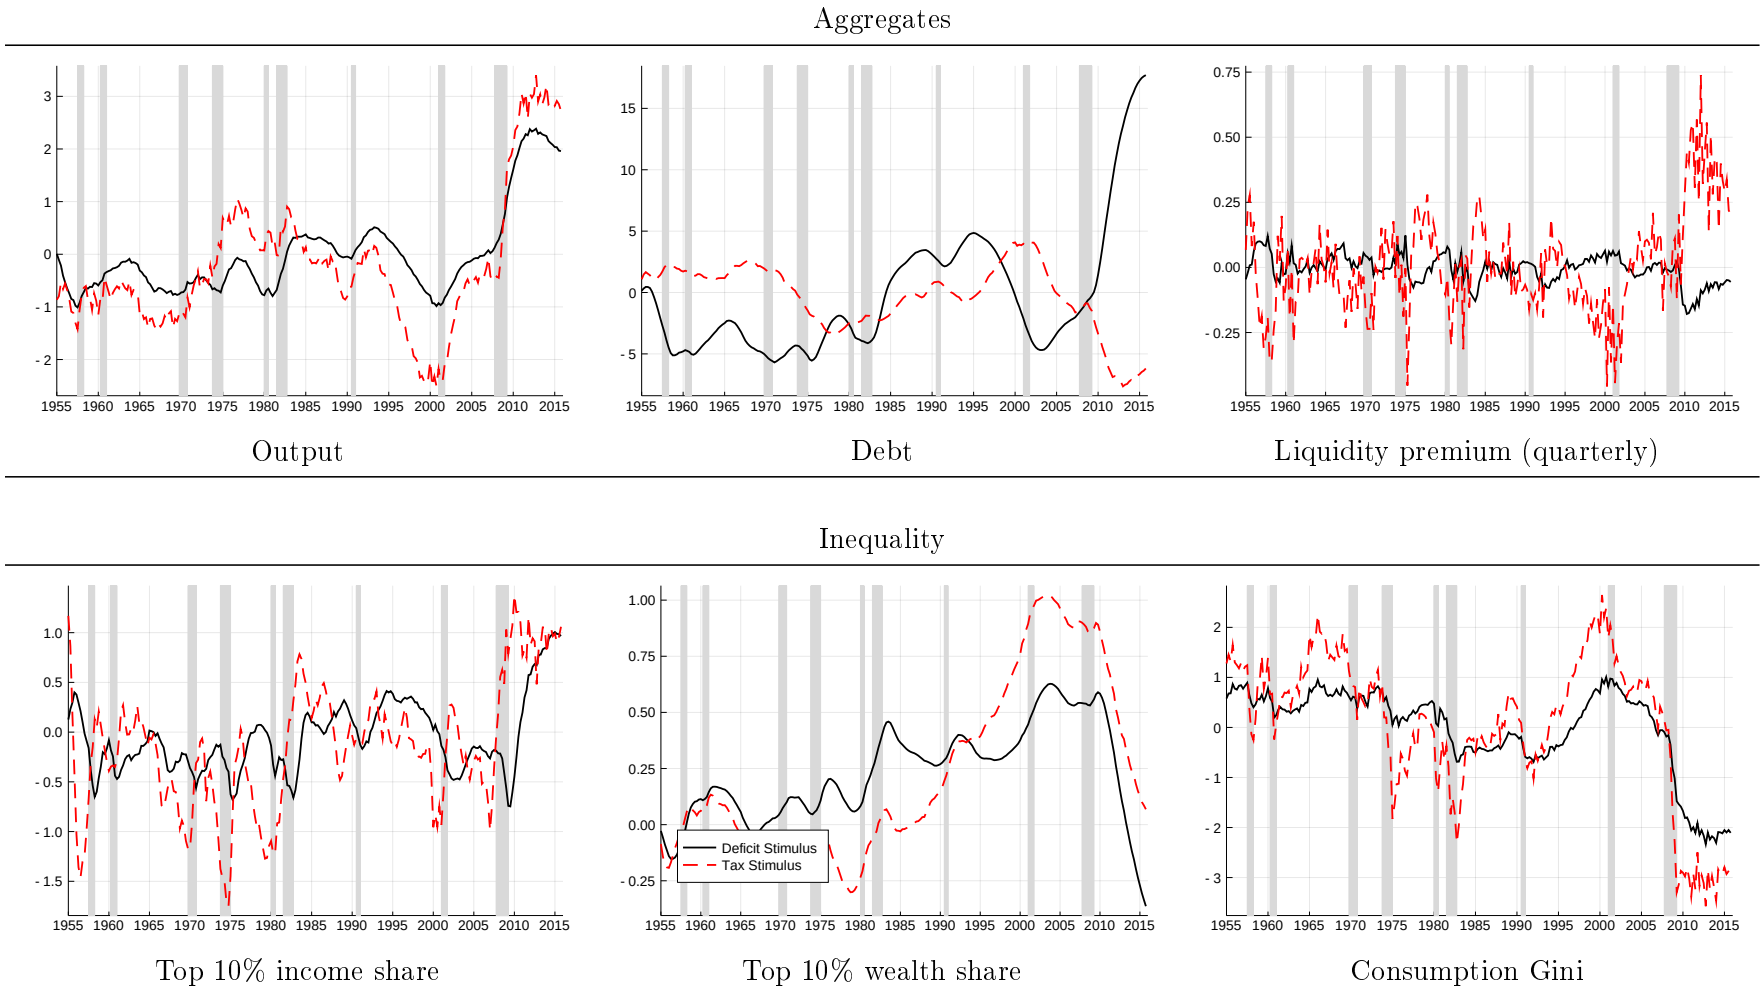

Notes: The panels display the evolution of output, government bonds, and the liquidity premium as well as, wealth, income, and consumption inequality that the model would counterfactually predict had the government policies been different, feeding the smoothed sequence of shocks (as in Figure 4) through the model. The lines represent the difference (in p.p.) in the evolution compared to feeding the same shocks through the baseline model. The solid line corresponds to a setup where we allow for persistence deviations of government debt by setting $\gamma_B = 0.1$ and $\gamma_B^\tau = 0.4$. The dotted line reflects the counterfactual where we double the estimated tax response, γ_Y^τ. Shaded areas correspond to NBER-dated recessions.

Finally, we consider alternative fiscal policy scenarios; see Figure 7. First, we assume more aggressive deficit (spending) policies. In particular, we allow debt to increase more persistently, by lowering γ_B and γ_B^τ such that the average tax rate path is roughly kept as in the baseline. Second, we consider a policy that adjusts taxes more heavily, leaving the overall deficit and hence debt as in the baseline. That is, we consider a policy that lowers taxes rather than raises government consumption when fighting a recession with a government deficit.

The two alternative policies fare particularly differently in their response to the Great Recession. In the first scenario, the more active deficit scenario, government debt rises almost by another 20% and output is initially more stable after 2008 until the government starts to raise taxes to bring back the government debt to its steady-state level. In fact, the aggressive fiscal policy lifts the nominal rate (not displayed) by almost 1 percentage points (annualized), and inflation up by 0.5 percentage points, because households are willing to hold the extra liquidity only at higher returns. The equilibrium result is a substantially lower liquidity premium. This, in turn, reduces wealth inequality substantially, but drives up income and consumption inequality. The depressed liquidity premium means that there is redistribution to wealth-poor households, which predominantly save in liquid assets at the expense of wealth-rich households, which mostly hold illiquid assets. Cutting taxes more aggressively during the Great Recession while holding the deficit constant would have stabilized output and hence consumption inequality more strongly. However, this would have increased the liquidity premium, putting further downward pressure on nominal rates, such that our abstraction from the effective lower bound becomes even more binding.

7 Conclusion

How much 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. Otherwise, we find that household heterogeneity and the inclusion of micro data in the estimation do 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 50% of the increase in US wealth inequality and virtually all of the increase in 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 expansionary fiscal policy that would have allowed government debt to increase substantially more after the Great Recession would have had a positive impact on interest rates and thus helped the economy to escape the effective lower bound earlier and boosted the recovery. At the same time, this evolution of government debt would have eroded the return difference between illiquid and liquid assets, helping in particular poor households to accumulate wealth, driving down wealth inequality.

These findings suggest that future research on inequality should take business cycles into account. A synthesis of the previous literature that focuses on permanent changes, e.g. in the tax and transfer system or the skill premium, with the forces that we highlight will be an important area of research. Our findings further suggest exploring 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.

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

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. (2019). Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. NBER Working Papers 26123.

Auclert, A., Rognlie, M., and Straub, L. (2018). The intertemporal Keynesian Cross. NBER Working Papers 25020.

Auclert, A., Rognlie, M., and Straub, L. (2020). Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model. NBER Working Papers 26647.

Barkai, S. (2019). Declining labor and capital shares. *Journal of Finance (forthcoming)*.

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.

Benabou, R. (2002). Tax and education policy in a heterogeneous-agent economy: What levels of redistribution maximize growth and efficiency? *Econometrica*, 70(2):481–517.

Boivin, J. and Giannoni, M. (2006). DSGE Models in a Data-Rich Environment. NBER Working Papers 12772.

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

Bricker, J., Henriques, A., Krimmel, J., and Sabelhaus, J. (2016). Measuring income and wealth at the top using administrative and survey data. *Brookings Papers on Economic Activity*, 2016(1):261–331.

Broer, T., Hansen, N.-J. H., Krusell, P., and Öberg, E. (2019). The New Keynesian transmission mechanism: A heterogenous-agent perspective. *The Review of Economic Studies*, 87(1):77–101.

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

Castaneda, A., Diaz-Gimenez, J., and Rios-Rull, J.-V. (1998). Exploring the income distribution business cycle dynamics. *Journal of Monetary Economics*, 42(1):93 – 130.

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. Mimeo, 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. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics (forthcoming)*.

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.

Gabaix, X., Lasry, J.-M., Lions, P.-L., and Moll, B. (2016). The dynamics of inequality. *Econometrica*, 84(6):2071–2111.

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

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. (2014). How risky are recessions for top earners? *The American Economic Review*, 104(5):148–153.

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

Hagedorn, M., Manovskii, I., and Mitman, K. (2019). The fiscal multiplier. NBER Working Papers 25571.

Hall, R. E. (2018). New evidence on the markup of prices over marginal costs and the role of mega-firms in the US economy. NBER Working Papers 24574.

Heathcote, J., Perri, F., and Violante, G. L. (2010). Unequal we stand: An empirical analysis of economic inequality in the United States, 1967–2006. *Review of Economic Dynamics*, 13(1):15–51.

Heathcote, J., Storesletten, K., and Violante, G. L. (2017). Optimal tax progressivity: An analytical framework. *The Quarterly Journal of Economics*, 132(4):1693–1754.

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. (2019). Sources of U.S. wealth inequality: Past, present, and future. Mimeo, University of Pennsylvania.

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.

Karabarbounis, L. and Neiman, B. (2019). Accounting for factorless income. *NBER Macroeconomics Annual*, 33(1):167–228.

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.

Kehrig, M. and Vincent, N. (2018). The micro-level anatomy of the labor share decline. NBER Working Papers 25275.

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.

Kopczuk, W. (2015). What do we know about the evolution of top wealth shares in the united states? *Journal of Economic Perspectives*, 29(1):47–66.

Kuhn, M., Schularick, M., and Steins, U. I. (2020). Income and wealth inequality in America, 1949-2013. *Journal of Political Economy (forthcoming)*.

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

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

McKay, A. (2017). Time-varying idiosyncratic risk and aggregate consumption dynamics. *Journal of Monetary Economics*, 88:1–14.

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.

Perri, F. and Steinberg, J. (2012). Inequality and redistribution during the Great Recession. Economic Policy Paper 12-1, Federal Reserve Bank of Minneapolis.

Piketty, T. and Saez, E. (2003). Income inequality in the United States, 1913–1998. *The Quarterly Journal of Economics*, 118(1):1–41.

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.

Saez, E. and Zucman, G. (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, 131(2):519–578.

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. (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. Mimeo, Princeton University.

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. *Journal of Money, Credit and Banking*, 48(2-3):253–291.

A Data

A.1 Data for Calibration

Mean illiquid assets. Fixed assets (NIPA table 1.1) over quarterly GDP (excluding net exports; see below), averaged over 1954-2015.

Mean liquidity. Gross federal debt held by the public as percent of GDP (FYPUGDA188S). Available from 1954-2015.

Fraction of borrowers. Taken from the Survey of Consumer Finances (1983-2013); see Bayer et al. (2019) for more details.

Average top 10% share of wealth. Source is the World Inequality Database (1954-2015).

A.2 Data for Estimation

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

$$OBS_t = \begin{bmatrix} \Delta \log\left(Y_t\right) \\ \Delta \log\left(C_t\right) \\ \Delta \log\left(I_t\right) \\ \Delta \log\left(w_t^F\right) \\ \Delta \log\left(T_t\right) \\ \log\left(\hat{N}_t\right) \\ \log\left(\hat{R}_{t+1}^b\right) \\ \log\left(\hat{\pi}_t\right) \\ \log\left(\hat{s}_t\right) \\ \propto^\tau \times \log\left(\hat{\tau}_t^P\right) \\ \log\left(\hat{p90p100}_t^{wealth}\right) \\ \log\left(\hat{p90p100}_t^{income}\right) \end{bmatrix}$$

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

Federal tax receipts. Federal government current tax receipts (FEDT) 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.

Tax progressivity. US Individual Income Tax: Tax Rates for Regular Tax: Highest Bracket in Percent (IITTRHB). Available annually 1954 to 2015.

Wealth inequality. p90p100 of US net personal wealth from the World Inequality Database. Available annually 1954 to 2014.

Income inequality. p90p100 of US pre-tax national income from the World Inequality Database. Available annually 1954 to 2014.

B Inspecting Model Mechanisms

B.1 Impulse Response Functions of the Estimated HANK* Model

Figure 8: IRFs to structural deficit and monetary policy shocks

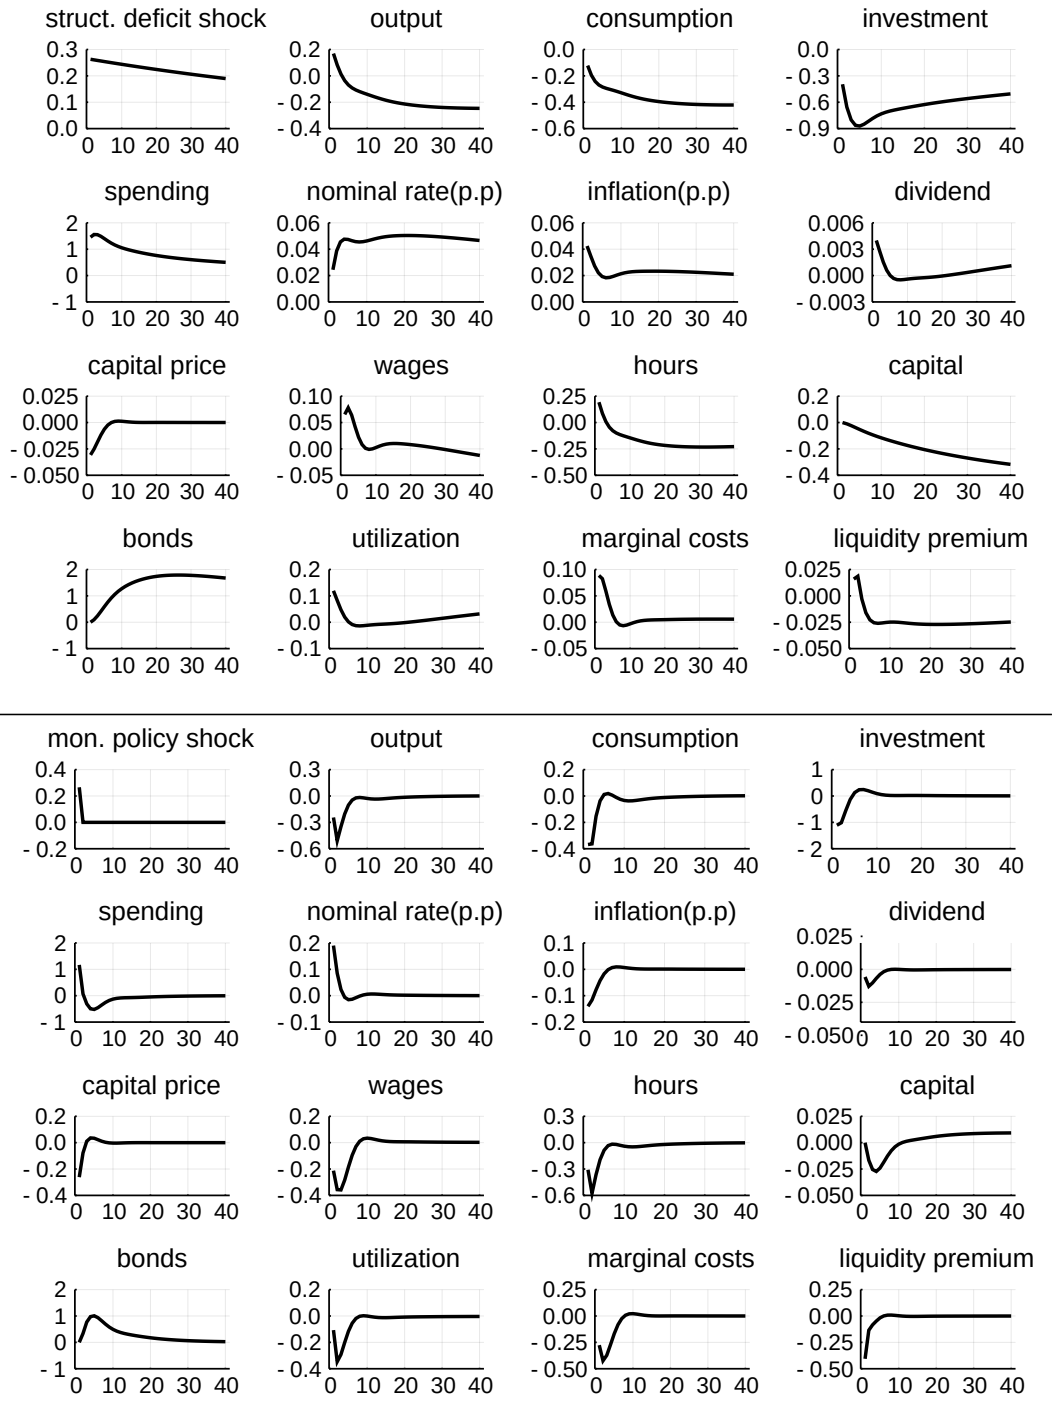

Notes: Top: IRF to a structural deficit shock. Bottom: IRF to a monetary policy shock.

Figure 9: IRFs to markup shocks

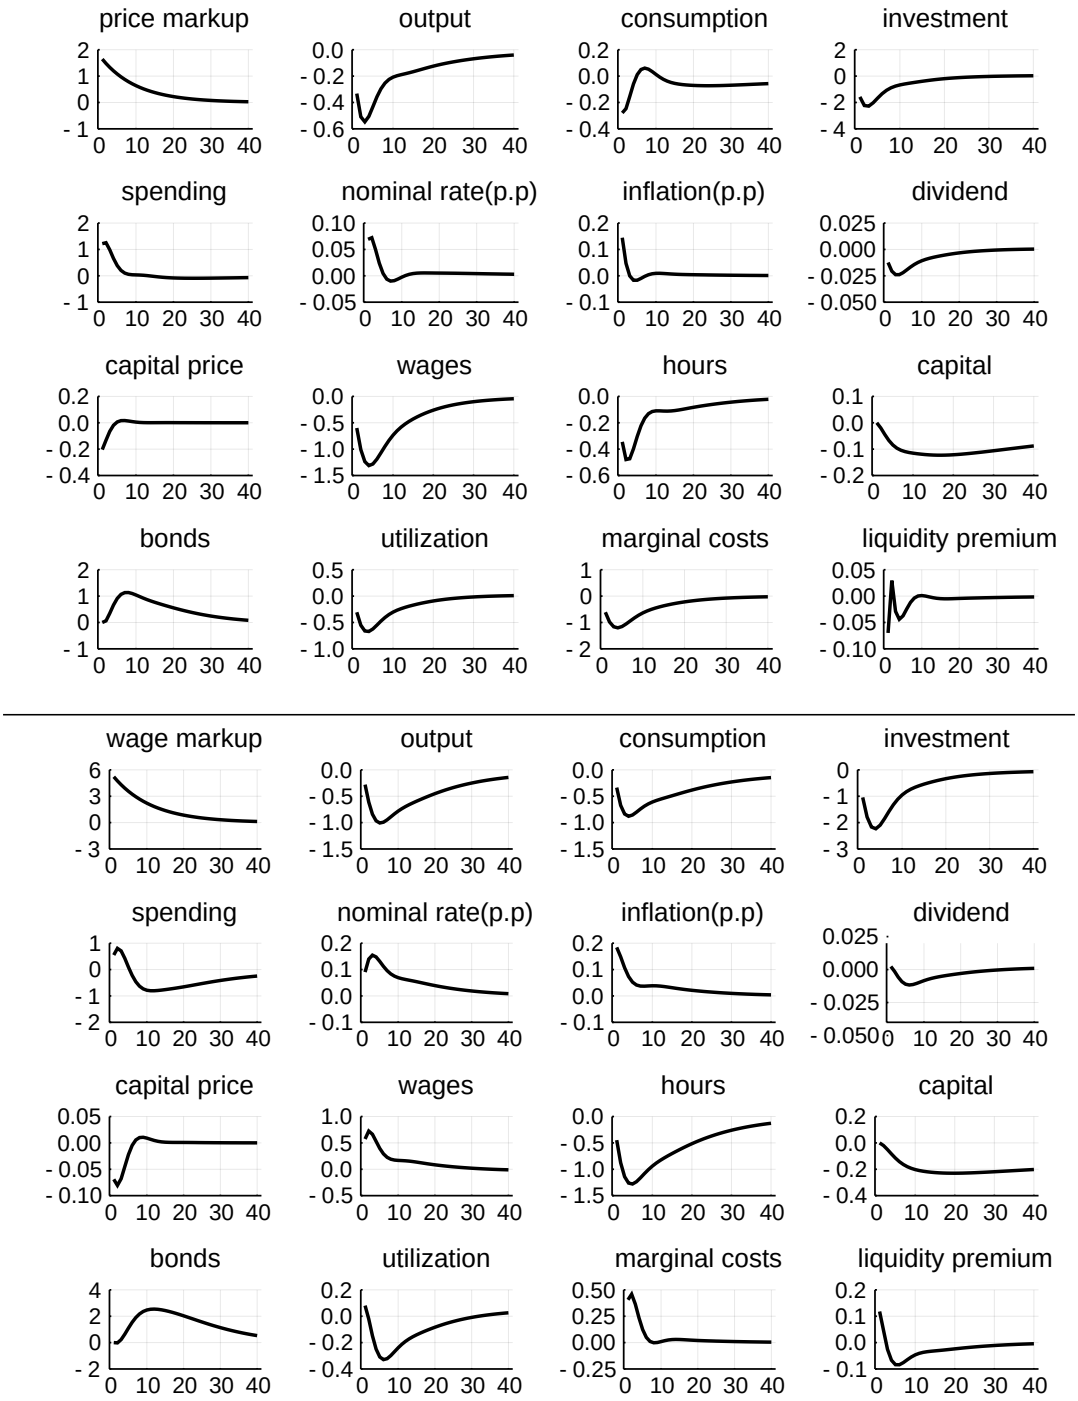

Notes: Top: IRF to a price-markup shock. Bottom: IRF to a wage-markup shock.

Figure 10: IRFs to technology shocks

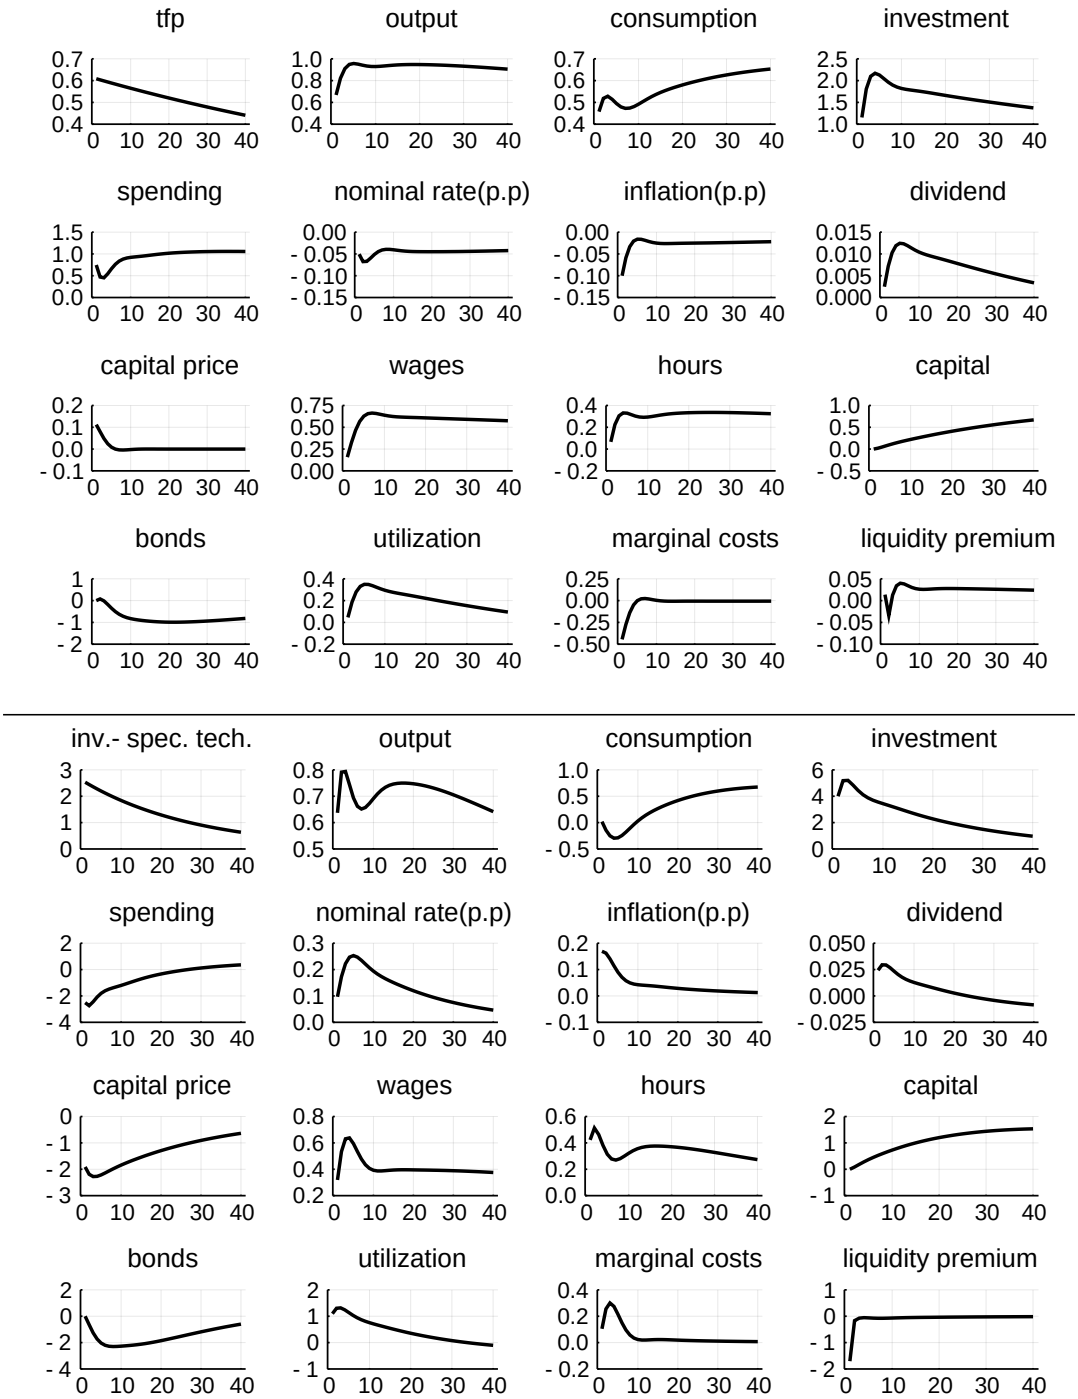

Notes: Top: IRF to a TFP shock. Bottom: IRF to an MEI shock.

Figure 11: IRFs to risk premium and income risk shocks

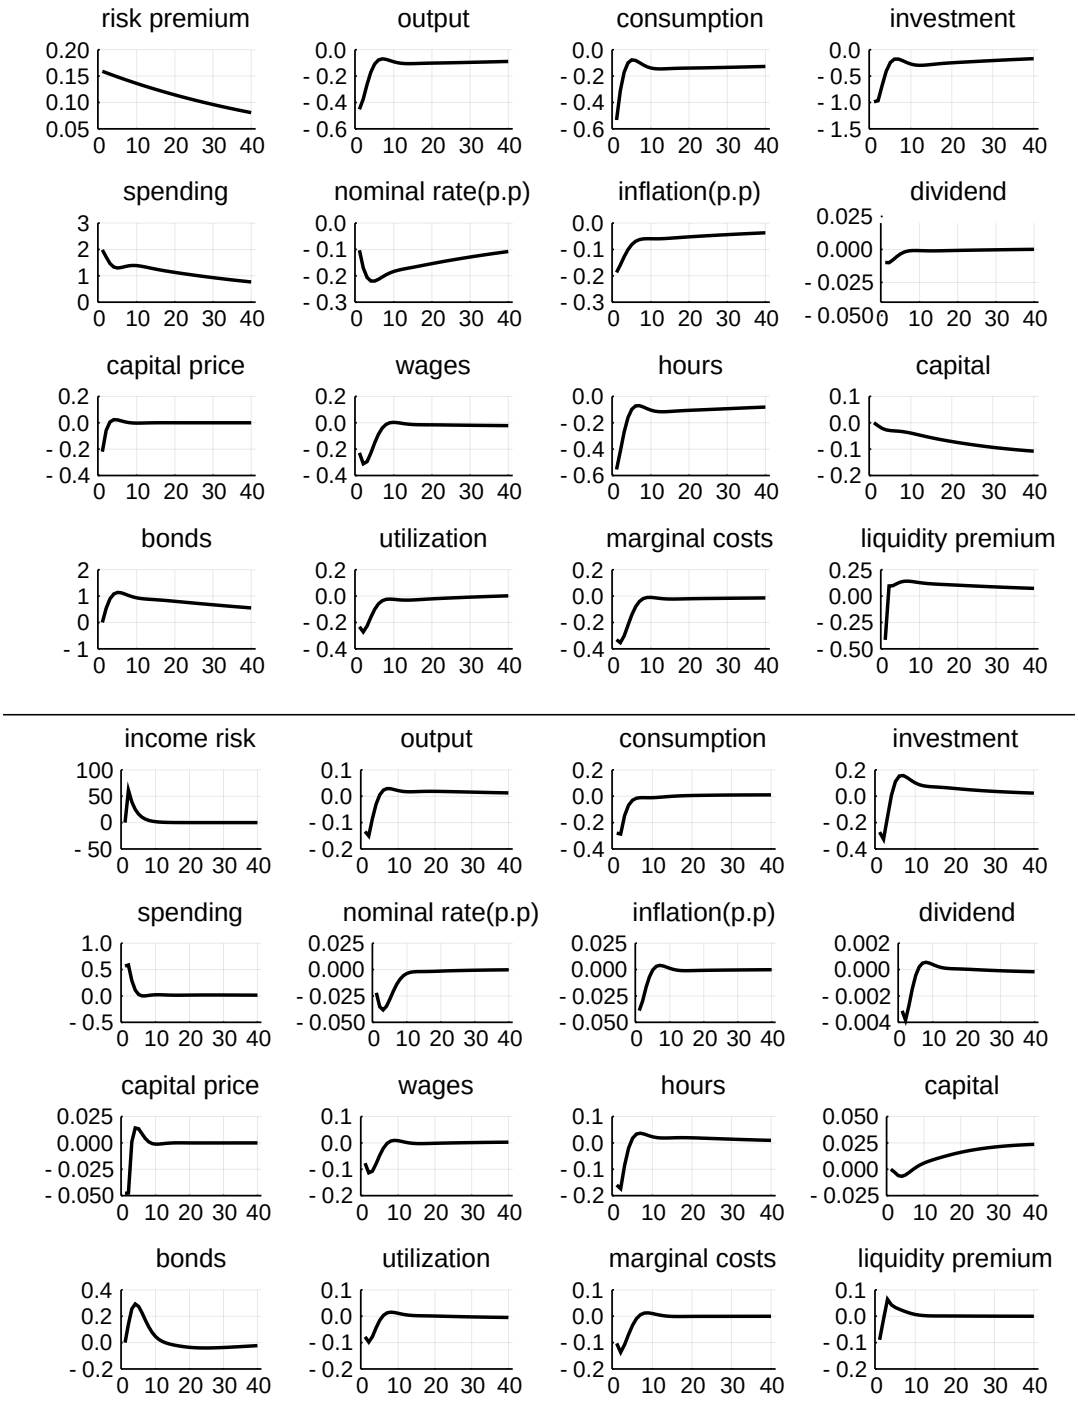

Notes: Top: IRF to a risk premium shock. Bottom: IRF to an income risk shock.

Figure 12: IRFs to tax shocks

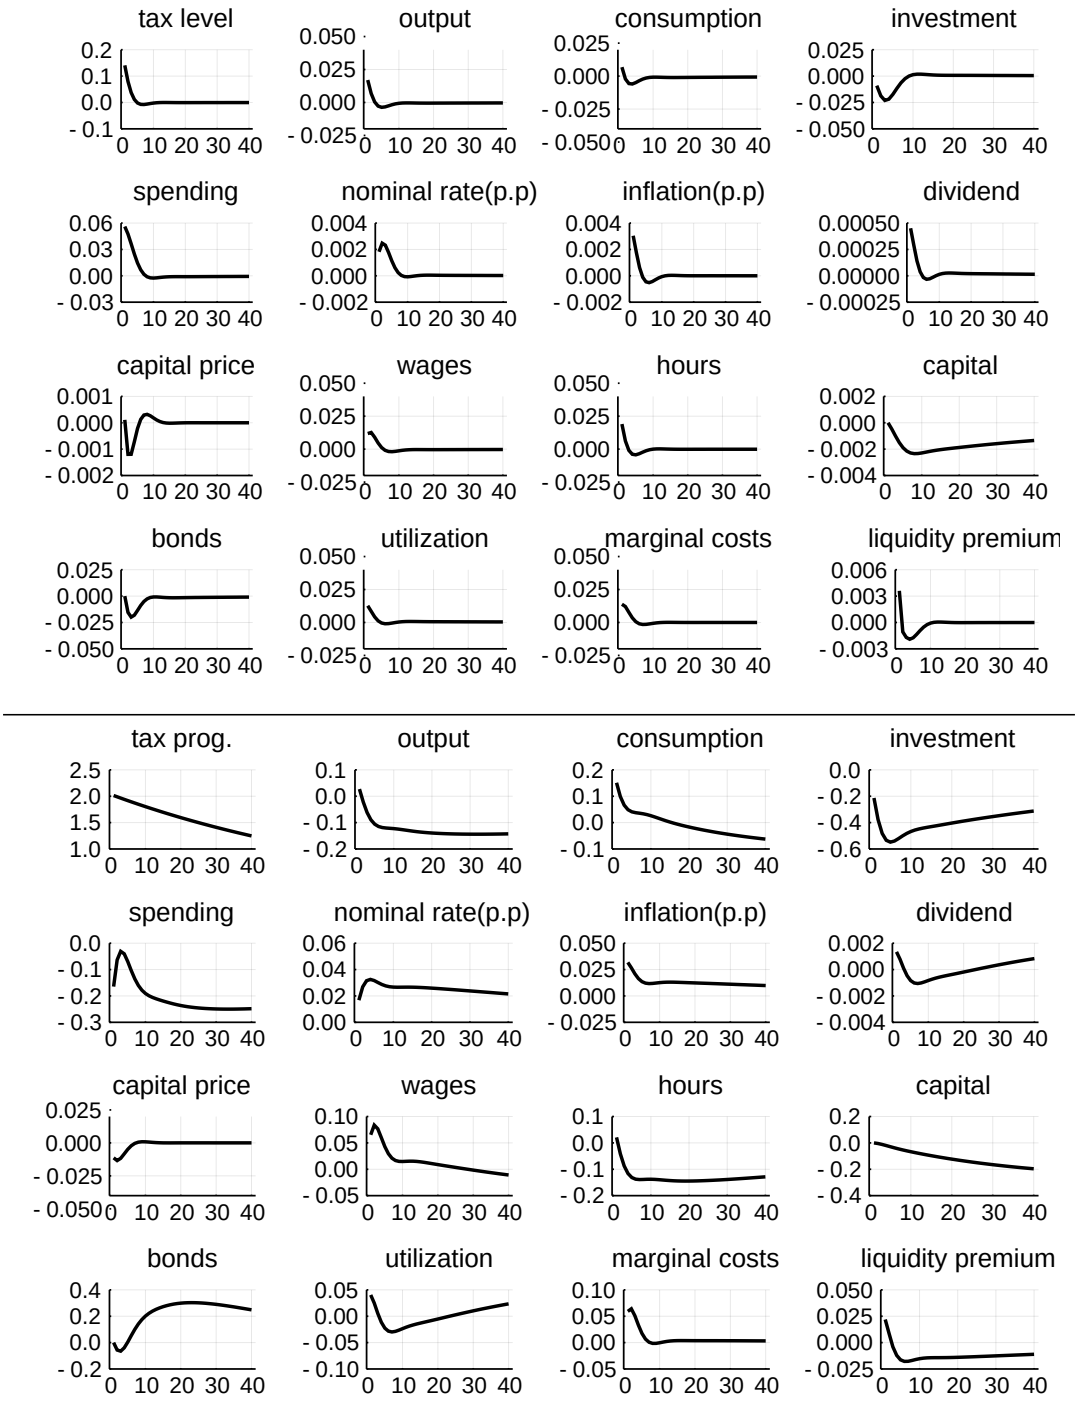

Notes: Top: IRF to tax level shock. Bottom: IRF to a tax progessivity shock.

Figure 13: Historical decompositions: Taxes and income risk

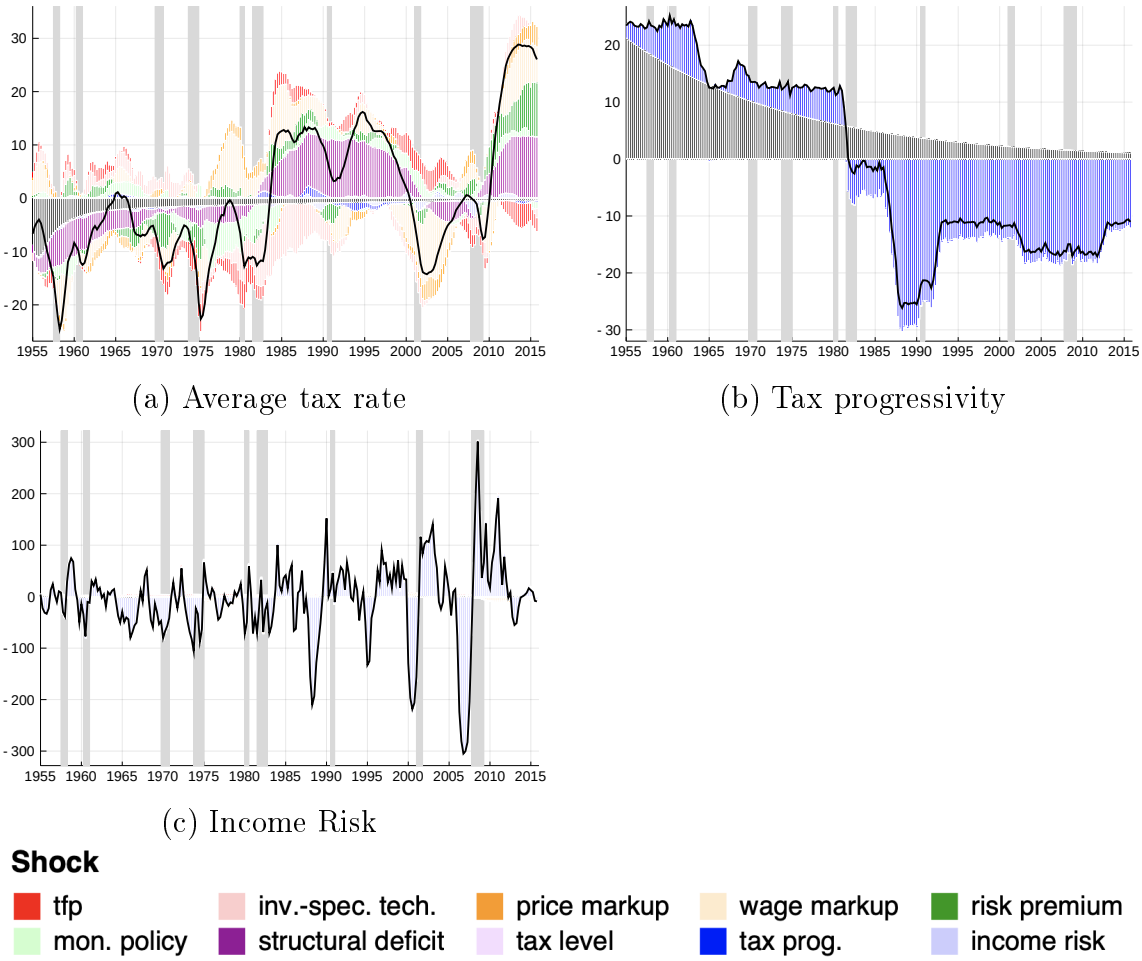

(a) Average tax rate

(b) Tax progressivity

(c) Income Risk

Notes: Historical decompositions of the log-deviations of the average tax rate (top left), tax progressivity (top right), and income risk (bottom left). Shaded areas correspond to NBER-dated recessions.

C Further Historical Decompositions

Figure 13 shows the historical decomposition of taxes and income risk. The historical decomposition shows that income risk is mostly driven by exogenous shocks and not endogenous feedback. Figure 14 shows the historical decomposition of the growth rates of output, consumption, investment, and government spending for the HANK* model (with inequality data). Figure 15 shows the historical decomposition of the log-level of output, consumption, investment, and government spending for the HANK* model. Figure 16 shows the historical decomposition of the nominal rate, inflation, marginal costs, and government debt.

Figure 14: Historical decompositions: Output, consumption, investment and government spending (growth rates)

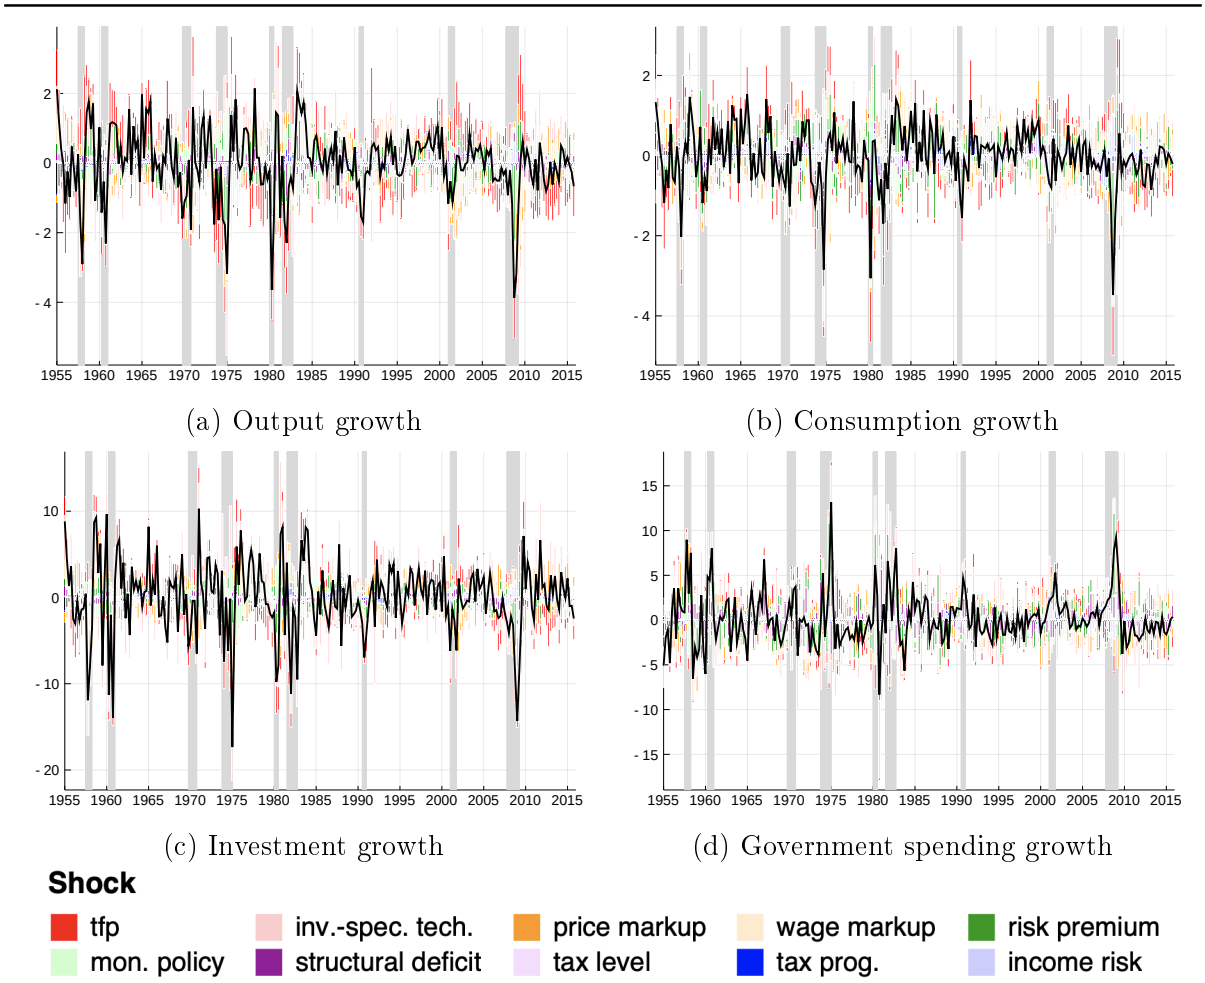

Notes: Historical decompositions of the log-deviations of output growth (top left), consumption growth (top right), investment growth (bottom left), and government spending growth (bottom right). Shaded areas correspond to NBER-dated recessions.

Figure 15: Historical decompositions: Output, consumption, investment and government spending (log levels)

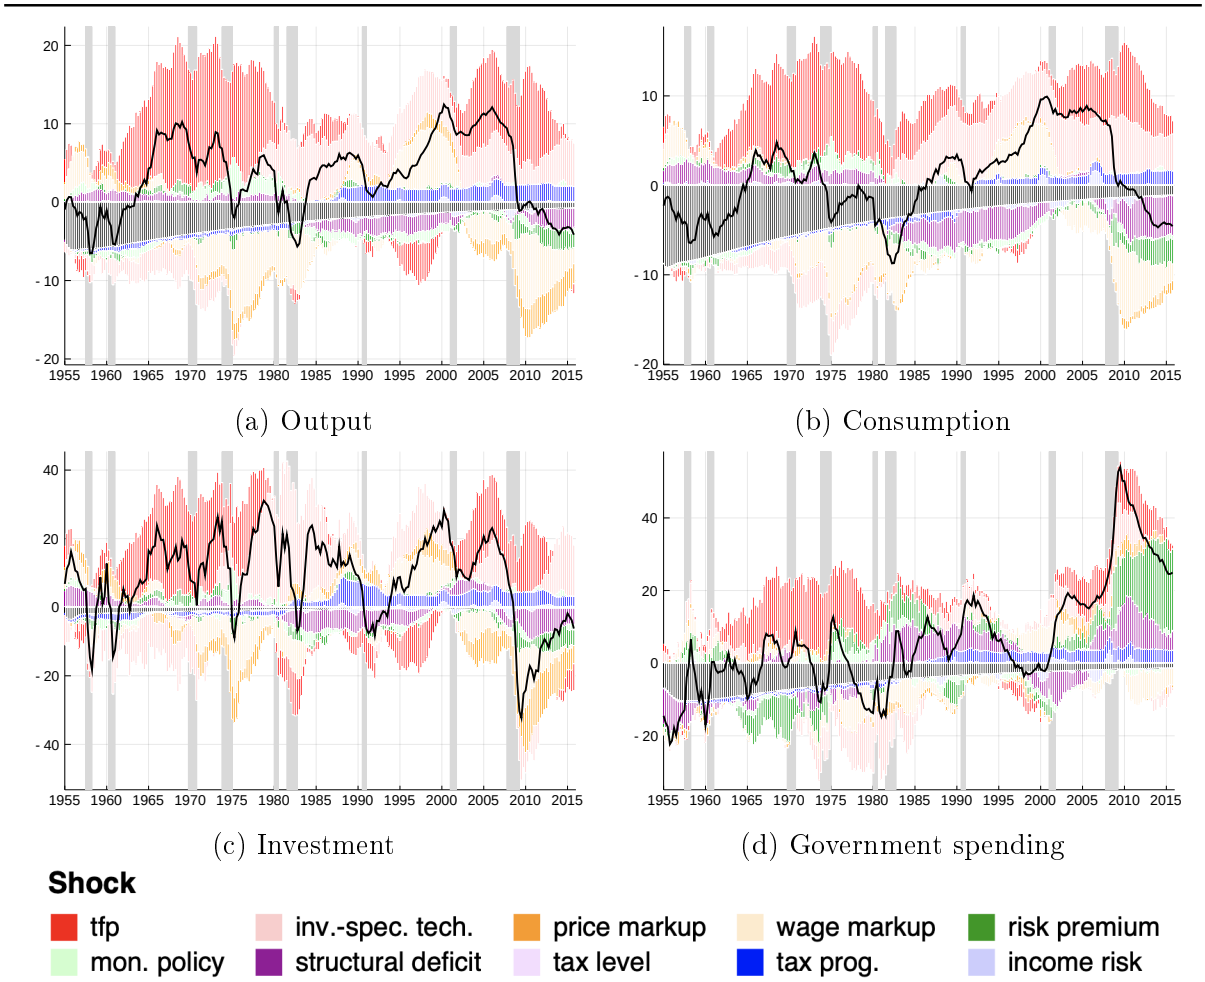

Notes: Historical decompositions of the log-deviations of output (top left), consumption (top right), investment (bottom left), and government spending (bottom right). Shaded areas correspond to NBER-dated recessions.

Figure 16: Historical decompositions: Nominal rate, inflation, marginal costs and government debt (log levels)

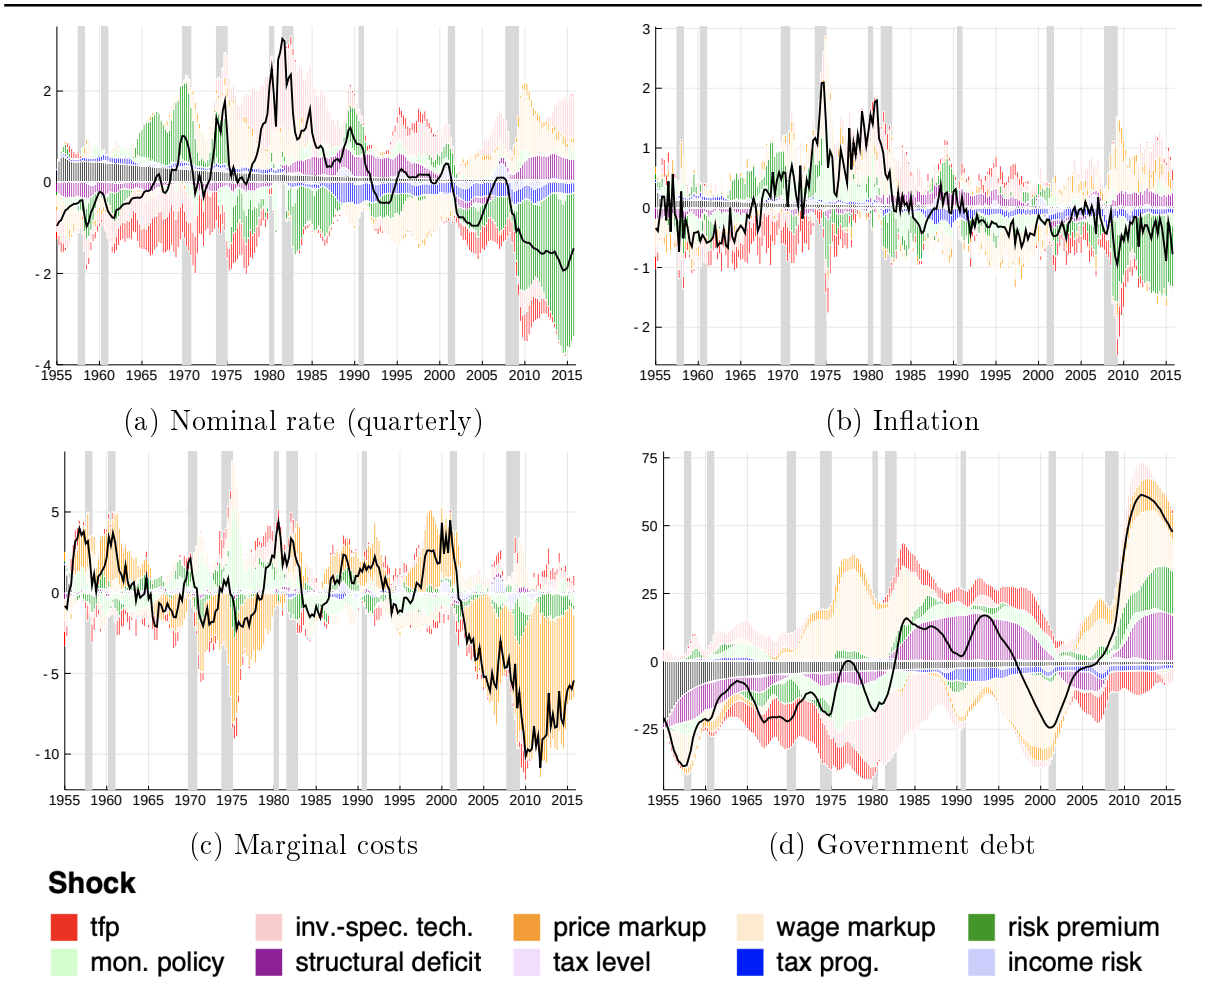

Notes: Historical decompositions of the log-deviations of the nominal bond rate (top left), inflation (top right), marginal costs (bottom left), and government debt (bottom right). Shaded areas correspond to NBER-dated recessions.

D Further Results on Inequality Dynamics

Figure 17 presents the estimated impulse responses of inequality on all 10 shocks. The general picture is that income inequality shows the most transitory movements. Wealth inequality moves most persistently after all shocks because it is driven by accumulation decisions. Consumption, driven by both income and wealth, shows both short- and long-run dynamics.

Monetary policy and risk premium shocks both increase inequalities through a rise in price markups that benefits mostly the (wealthy) entrepreneurs. The effects on income and consumption inequality are rather transitory, while wealth inequality increases persistently (more so for the risk premium shock).

Turning to fiscal policy, increasing the structural deficit lowers income and wealth inequality but increases consumption inequality. On the tax side, a shock to the average tax rate has only very small overall effects on inequality, but tends to lower income inequality (at least in the very short run) but increases consumption and wealth inequality marginally. A shock to the progressiveness of the tax schedule, on the other hand, causes persistently lower income, consumption, and wealth inequality.

Price and wage markup shocks both have persistent effects on income, consumption, and wealth inequality; however, they differ markedly in the sign of their effects. While price markups increase inequality because they raise profits that go to entrepreneurs, a rise in the target wage markup lowers inequalities by increasing wage compression.

Higher income risk leads to a quick increase in income and consumption inequality as poorer households over-proportionally increase their liquid asset holdings for insurance purposes. Wealth inequality initially falls because poor households react strongly by accumulating extra, mostly liquid, assets. After about 2.5 years, wealth inequality increases persistently when higher income risks have been realized and have led to more dispersed incomes.

The response of wealth inequality to TFP and investment-specific technology shocks looks very similar. In both cases, wealth inequality initially increases but eventually declines, and persistently so, after about 5 years. While both shocks lead to a persistent decline in consumption inequality, the investment-specific technology shock initially raises it for the first year or so. Both shocks also raise income inequality, the TFP shock by raising productivity and therefore the returns to capital held by the wealthy households, and the investment-specific technology shock by making investment in capital more efficient, which benefits the wealthy as well.

Figure 17: Impulse responses of inequality

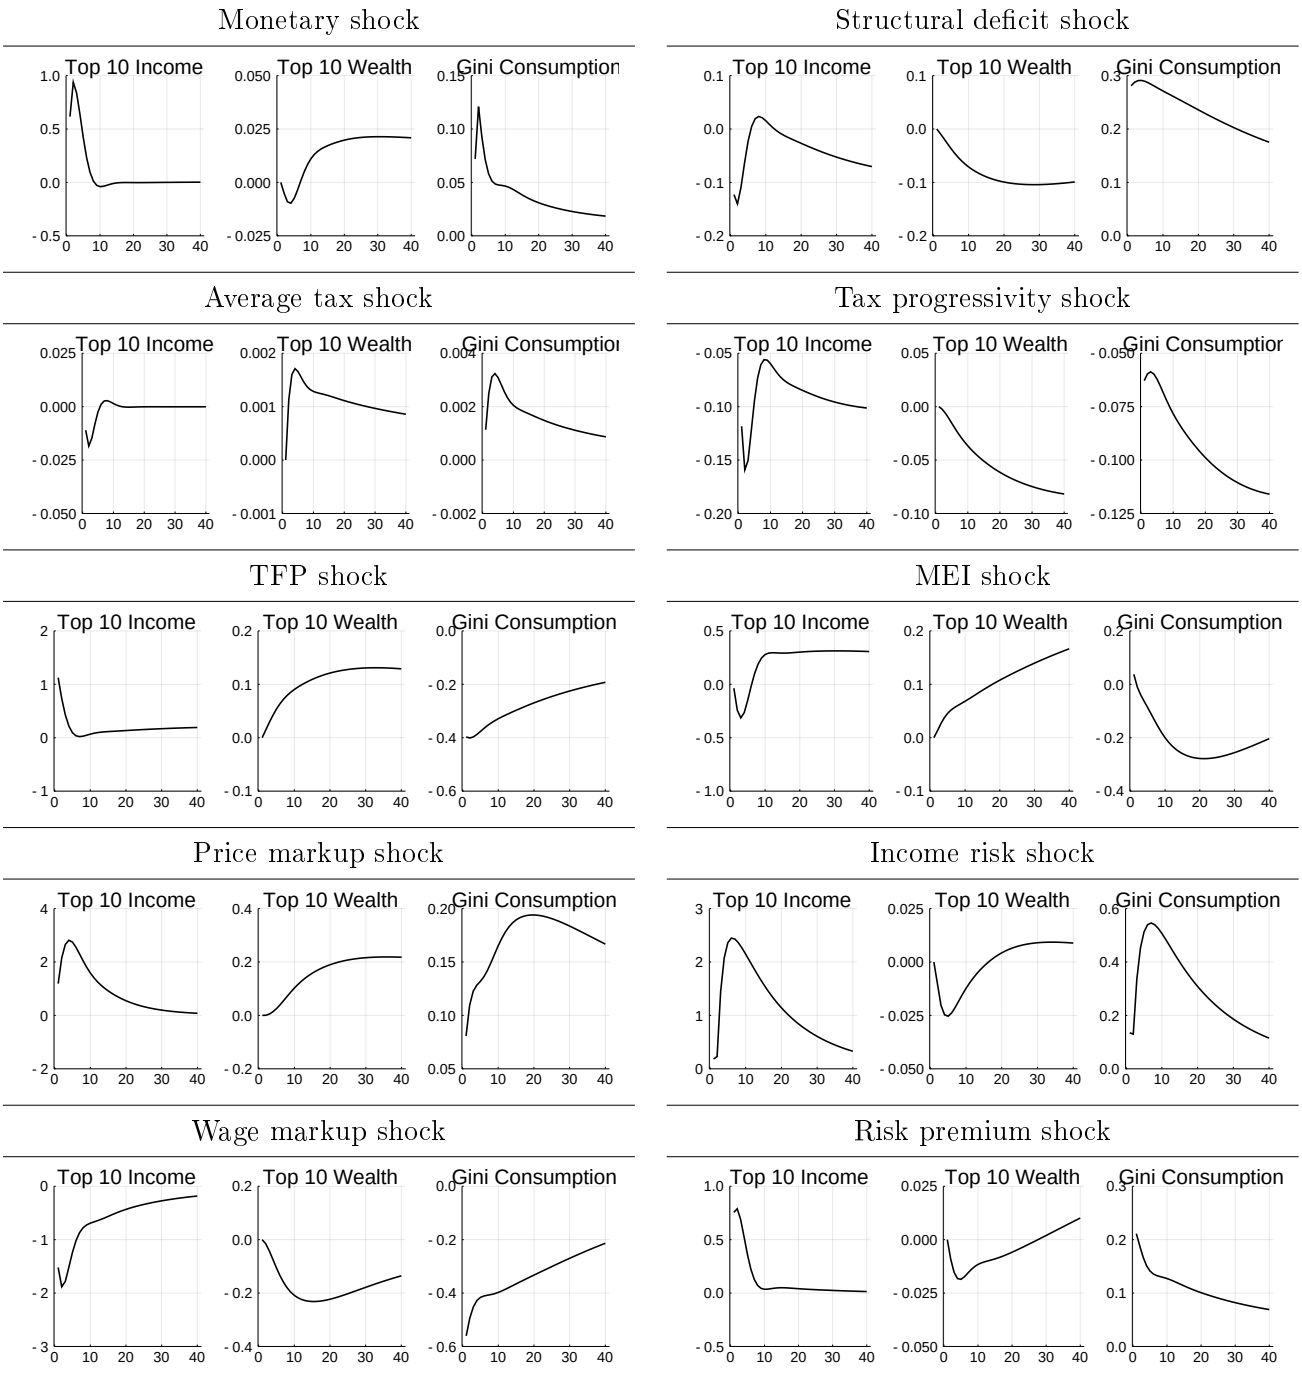

Notes: The figures display the impulse responses of income, consumption, and wealth inequality in response to the shocks labeled above. Parameter estimates from HANK*. See main text for further details.

Figure 18: US inequality – data vs. model

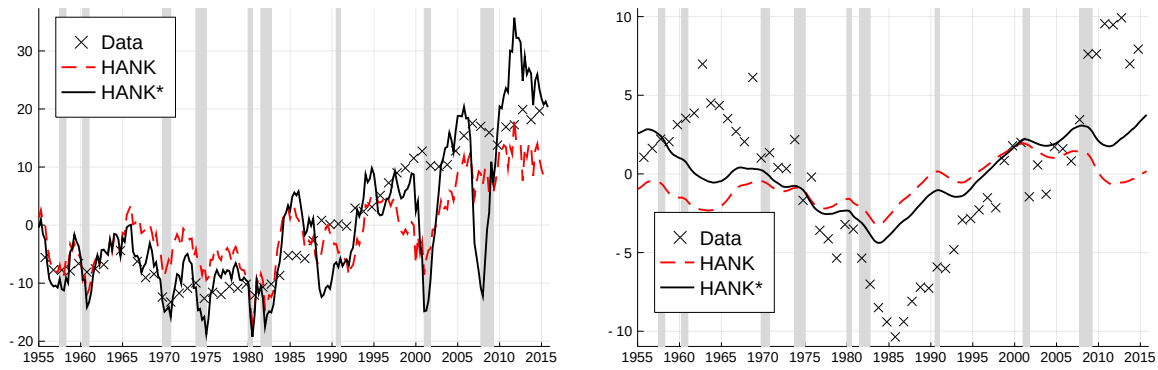

(a) Top-10% income share (b) Top-10% wealth share

Notes: Data (crosses) corresponds to log-deviations of the annual observations of the share of pre-tax income and wealth held by the top 10% in each distribution in the US taken from the World Inequality Database. HANK* (solid) [HANK (dashed)] corresponds to the smoothed states of both implied by the estimated HANK model with inequality data [w/o inequality data and income risk shocks]. Shaded areas correspond to NBER dated recessions.

E Estimation Without Income Risk Data

We re-estimate the model without movements in income risk as this measure already contains some cross-sectional information about income inequality. For this purpose, we do not allow for shocks to income risk, i.e., we keep the variance of idiosyncratic income risk at its steady-state value, and drop the time series for income risk from Bayer et al. (2019) from the estimation. Figure 18 shows that the model without movements in income risk (denoted by HANK) still implies a U-shaped evolution of income and wealth inequality. The fit is less good then for the model with all inequality data (HANK*). More persistent fluctuations in markups partly compensate for the missing movements in income risk.

F MCMC Diagnostics

Table 6: Gelman and Rubin (1992) convergence diagnostics

	HANK		HANK*	
Parameter	PSRF	97.5%	PSRF	97.5%
δ_s	1.042	1.107	1.011	1.030
ϕ	1.022	1.055	1.011	1.028
κ	1.004	1.011	1.007	1.018
κ_w	1.005	1.010	1.007	1.018
ρ_A	1.002	1.004	1.003	1.007
σ_A	1.003	1.006	1.001	1.003
ρ_Z	1.003	1.006	1.002	1.005
σ_Z	1.001	1.002	1.000	1.000
ρ_Ψ	1.001	1.003	1.004	1.010
σ_Ψ	1.013	1.033	1.013	1.035
ρ_μ	1.002	1.005	1.001	1.002
σ_μ	1.004	1.010	1.003	1.008
$\rho_{\mu w}$	1.004	1.009	1.005	1.013
$\sigma_{\mu w}$	1.002	1.005	1.011	1.020
ρ_s	1.023	1.062	1.007	1.008
σ_s	1.006	1.017	1.002	1.004
ρ_R	1.011	1.029	1.002	1.004
σ_R	1.002	1.005	1.001	1.003
θ_π	1.026	1.070	1.019	1.049
θ_y	1.003	1.008	1.002	1.003
γ_B	1.011	1.029	1.025	1.055
σ_G	1.021	1.057	1.016	1.039
γ_π	1.001	1.002	1.009	1.019
γ_Y	1.034	1.087	1.005	1.012
ρ_τ	1.032	1.081	1.013	1.032
σ_τ	1.276	1.895	1.428	2.980
γ_B^τ	1.013	1.032	1.039	1.090
γ_Y^τ	1.021	1.058	1.035	1.080
ρ_P	1.000	1.001	1.001	1.003
σ_P	1.010	1.024	1.019	1.050
Σ_N	1.041	1.103	1.032	1.074
$\propto^\tau$	1.010	1.025	1.015	1.040
ρ_G	1.002	1.004	1.012	1.024
σ_s^{me}	1.013	1.025	1.026	1.046
σ_T^{me}	1.003	1.007	1.007	1.020
$\sigma_{\tau p}^{me}$	1.317	2.149	1.121	1.255
σ_{WI}^{me}	–	–	1.001	1.003
σ_{II}^{me}	–	–	1.001	1.003

Note: Gelman and Rubin (1992) potential scale reduction factor (PSRF) and 97.5% quantile.