EUROPEAN CENTRAL BANK

EUROSYSTEM

Working Paper Series

Hanno Kase, Rodolfo Dinis Rigato

Beyond averages: heterogeneous effects of monetary policy in a HANK model for the euro area

No 3086

Abstract

We introduce an estimated medium scale Heterogeneous-Agent New Keynesian model for forecasting and policy analysis in the Euro Area and discuss the applications of this type of models in central banks, focusing on two main exercises. First, we examine an alternative scenario for monetary policy during the early 2020s inflationary episode, showing that earlier hikes in interest rates would have affected more strongly households at the lower end of the wealth distribution, whose consumption our model suggests was already depressed relative to the rest of the population. To provide intuition for this result, we introduce a new decomposition of the effects of monetary policy on consumption across the wealth distribution. Second, we show that introducing heterogeneous households does not come at the cost of forecasting accuracy by comparing the performance of our model to its exact representative-agent counterpart and demonstrating nearly identical results in predicting key aggregate variables.

JEL classification codes: D31, E12, E21, E52

Keywords: Heterogeneous-Agent New Keynesian Models, Monetary Policy, Inequality, Forecasting

Non-Technical Summary

In this paper, we introduce a Heterogeneous-Agent New Keynesian (HANK) model for policy analysis in the Euro Area and explore the usefulness and applications of these models in central banks. Traditional economic models often assume that all households are alike, potentially overlooking the varied impacts of policy decisions on different segments of the population. HANK models, which have become a cornerstone of contemporary macroeconomic analysis, allow us to account for the diverse financial situations of households, such as differences in wealth and income, thereby providing a more nuanced understanding of policy impacts. Our primary motivation is to equip central banks, such as the European Central Bank (ECB), with clearer insights into how these models can be utilized, both by changing the transmission channels of monetary policy, as well as by providing a framework to understanding its effects across the wealth and income distributions.

The main findings of our research are as follows: First, we explore a hypothetical scenario during the early 2020s inflationary period, finding that earlier interest rate hikes would have disproportionately impacted lower-income households, whose consumption had not recovered as quickly after the pandemic compared to other wealth groups. To better explain the differences in how monetary policy affects households across the wealth spectrum, we propose a new decompositoin of the effects of monetary policy on consumption across different wealth levels. Using this decomposition, we show that wealthier households respond differently to interest rate changes compared to less wealthy households. For less wealthy households, the negative effects are primarily driven by labor market conditions, whereas wealthier households are predominantly influenced by asset prices. Finally, we demonstrate that introducing heterogeneous households into the model does not come at the cost of forecasting accuracy. By comparing the performance of our HANK model to a corresponding representative-agent model that does not consider household differences. We find that both models exhibit similar forecasting accuracy in key aggregate variables, while the HANK model provides richer insights into distributional effects.

Methodologically, our model incorporates several realistic features, such as households saving in both liquid and illiquid assets and encountering information frictions that affect how swiftly they perceive economic changes. The model is estimated using Euro Area data from 2000 to 2019 to ensure it accurately reflects real-world conditions.

The policy implications of our work are significant for central banks. By understanding the varied effects of monetary policies on different household groups, the ECB and other central banks can design policies that stabilize the economy while minimizing adverse impacts on wealth inequality. Our model serves as a tool to guide more equitable monetary policy decisions, particularly relevant for the diverse economic landscape of the Euro Area

1 Introduction

Heterogeneous-Agent New Keynesian models (HANK) have emerged as a cornerstone of contemporary macroeconomic analysis. These models incorporate agents with high marginal propensities to consume, in line with empirical observations, thereby significantly altering the transmission channels of monetary and fiscal policies.[1] In addition to providing researchers and policy makers with new insights into these long-standing questions, HANK models also open the door to addressing new questions that were inaccessible with the previous generation of representative agent (RANK) models. These include the distributional impacts of monetary policy and the efficacy of widely used policies such as unemployment insurance extensions.[2]

In this paper, we introduce an estimated medium-scale HANK model designed for policy analysis in the Euro Area, and use it to study the applications of this type of models in central banks. Our model incorporates households facing idiosyncratic income risk and allows them to save in both liquid and illiquid assets. The two-asset structure enables the model to accurately reflect the observed levels of asset holdings and wealth inequality, while retaining high marginal propensities to consume, as the majority of assets are illiquid.[3] Therefore, the model is well-suited for understanding the effects of monetary policy and other macroeconomic shocks on wealth inequality. Additionally, households face information frictions when forming expectations about the economic environment, updating their information sets only infrequently.[4] These frictions result in hump-shaped responses of aggregate consumption to monetary policy and other aggregate shocks, in line with time series estimates.[5]

The model is estimated using Bayesian methods and quarterly time series data from 2000Q1 through 2019Q4. It is well-equipped to perform counterfactual scenario analyses of various policies as well as other typical applications of DSGE models using aggregate data, such as forecasting. We apply our model to two distinct areas, each relevant to the use of DSGE models in central banks.

First, we use our model to generate a counterfactual scenario for monetary policy at the

[1]See Kaplan et al. (2018) and Auclert et al. (2024).

[2]See Lee (2024), Kekre (2023), and Fernandes and Rigato (2024) for papers that address these questions.

[3]See Kaplan and Violante (2014) and Kaplan et al. (2018) for a discussion on two-asset models and marginal propensities to consume.

[4]See Auclert et al. (2020) for a detailed examination of these frictions in heterogenenous-agent models, which builds on the framework established by Mankiw and Reis (2002).

[5]See Christiano et al. (2005) and Auclert et al. (2020) for examples of estimated hump-shaped responses using time series data, or Ramey (2016) for a survey of the related literature.

onset of the high inflation episode in the early 2020s. Specifically, we examine the implications of a hypothetical more hawkish response by the European Central Bank, where it starts raising the short-term interest rate one quarter earlier (2022Q1) and maintains it at a higher than observed level until 2023Q2, at which point it returns to the actual policy path. Our model allows us to understand the consequences of this alternative policy stance across the wealth distribution. In the absence of aggregate data on consumption by different wealth brackets, we use our model to generate filtered values of these time series, which form the basis for our counterfactual scenario. Our findings suggest that a more aggressive monetary stance would have had a more pronounced impact on the consumption of households in the lower quartile. Moreover, these households' filtered consumption levels were already depressed relative to the rest of the population at the time, likely due to a slower recovery from the COVID-19 pandemic.

To better understand the distributional effects of monetary policy, we introduce a novel decomposition of the effects of monetary policy on consumption along the wealth distribution. In our model, the dynamic effect of an unexpected increase in interest rates can be broken down into three components: (i) a disposable income term, primarily related to lower real wages and reduced labor demand; (ii) an unexpected capital gains term, capturing the decline in the prices of long-duration assets following an unexpected increase in interest rates; and (iii) an expected asset returns channel, which relates to higher anticipated future returns driven by persistently higher interest rates. The lower end of the wealth distribution is predominantly influenced by the labor income channel (i), whereas capital income channels (ii) and (iii) have very persistent effects on the consumption of wealthier households. Furthermore, we find that the hump-shaped responses of aggregate consumption are largely driven by the disposable labor income channel at the lower end of the wealth distribution.

Second, we evaluate the out-of-sample forecasting accuracy of our model against that of a comparable representative agent (RANK) model, finding nearly identical performance. To isolate the impact of heterogeneity, we ensure both models are as similar as possible, with the only difference being the replacement of the heterogeneous households block with a representative consumer in the RANK model. Importantly, the RANK model incorporates habit persistence in consumption, which is the conventional method for producing hump-shaped responses in this setting.[6] We estimate the RANK model and re-estimate the HANK model using only the first half of our sample (up to 2009Q4), and then generate unconditional, out-of-sample forecasts

[6]See Smets and Wouters (2007).

for the growth rates of real output, consumption, investment, and the consumer price index up to eight quarters ahead, starting from 2010Q1. The forecasting accuracy is similar across all four variables, notably for consumption, where the two models differ the most.

We begin by introducing our model. In addition to featuring heterogeneous households, the model incorporates the standard components of medium-scale DSGE models. Both prices and wages are sticky. Monopolistically competitive firms set prices subject to nominal rigidities and partial indexation to past inflation, while unions face a similar problem for setting wages. A representative firm utilizes capital and labor to produce final goods. Capital goods producers convert final goods into capital, subject to investment adjustment costs. The monetary authority sets interest rates according to a standard Taylor (1993) rule, while a fiscal authority collects taxes to service public debt and finance government consumption.

Two features of our model are distinctive to the heterogeneous-agent framework. First, there is a fiscal rule linking the labor income tax rate to public debt. Since households in our model are not Ricardian, the timing and manner in which the government levies taxes influence consumer behavior and, consequently, the entire model dynamics. Second, a financial intermediary holds government debt and firm equity while supplying both liquid and illiquid assets to households. The nominal return on liquid assets is assumed to be the nominal interest rate set by the monetary authority, akin to a bank deposit. In contrast, illiquid assets represent claims on the financial intermediary's equity, reflecting the returns on equity and government bonds (the intermediary's assets) and liquid assets (its liabilities). The financial intermediary incurs a unit cost for supplying liquid assets to households, which generates a spread between its returns and illiquid assets.

We solve the model using the Sequence Space Jacobian method, developed by Auclert et al. (2021). This approach involves a first-order perturbation around the deterministic steady state, obtained by setting all aggregate shocks to zero. In a heterogenenous-agent framework, computing the deterministic steady state is computationally costly. Therefore, we follow a two-step approach for parameterizing the model. We calibrate parameters influencing the steady state, such as the degree of risk aversion of households, and estimate only those parameters that exclusively impact the model's dynamics, such as the coefficients in the Taylor rule.[7]

Our findings indicate that the primary contribution of HANK models to policy-making lies in leveraging the novel dimensions they introduce, rather than using them for the same

[7] Auclert et al. (2020) and Bayer et al. (2024) follow the same approach.

applications as their representative agent predecessors. When estimated with identical data, both HANK and RANK models yield remarkably similar dynamics for aggregate variables, as demonstrated by our forecast comparison. This suggests that the added value of HANK models is not in replicating the results of traditional models but in offering unique insights that stem from their inherent heterogeneity.

Related literature. There is a large and growing literature on the implications of household heterogeneity for monetary policy. To better link to it, we organize the discussion on related work around our main findings. The discussion on the transmission channels of monetary shocks in HANK models was pioneered by Kaplan et al. (2018). They found that monetary policy affects consumption in these models mostly through indirect channels, like general equilibrium effects via labor supply, due to households having high MPCs. This contrasts with RANK models, where changes in consumption are mainly driven by intertemporal substitution. Auclert et al. (2020) expand on this by highlighting the significant role of investment in HANK models, which, by influencing labor markets, affects the response of consumption to monetary policy .

The closest paper to out setting is Lee (2024), who studies the distributional effects of unconventional monetary policy.[8] However, our approach diverges in two significant ways. First, Lee (2024) primarily focuses on unconventional monetary policy. Second, and more critically, Lee provides a static decomposition of welfare gains across different household groups. In contrast, we focus on the dynamic implications for consumption. We argue that our model is more adept at addressing positive, rather than normative, question about consumption dynamics, as it incorporates frictions that produce a hump-shaped consumption response to a monetary shocks – an outcome consistent with empirical evidence in the time series literature. This distinction allows us to better capture the evolution of consumption in response to policy changes.

In terms of forecasting comparison between HANK and RANK, our study is most closely related to the work of Acharya et al. (2023). Their analysis similarly compares these models and concludes that that HANK models perform poorly, particularly concerning aggregate consumption. However, our approach differs in a crucial aspect: we incorporate information stickiness into the household block. This feature helps to disciplines the estimated impulse re-

[8]Del Negro et al. (2025) use the same model and report similar results.

sponse functions of consumption relative to various shocks, playing a role analogous to habit formation in RANK models. Our results suggest that the poorer performance of the HANK model observed by Acharya et al. (2023) may be attributed to the absence of this valuable degree of freedom relative to its RANK competitor.

Lastly, several studies have explored counterfactual scenarios for monetary and other types of policies in HANK models. For instance, Kekre (2023) examines how different unemployment insurance policies could have affected output and employment in the US following the Great Recession. Similarly, Lee (2024) investigates the effects of different combinations of conventional and unconventional monetary policies on welfare and inequality. To the best of our knowlege, our study is the first to explore alternative scenarios for consumption along the wealth distribution, as well as the specific episode of the interest rate hike in the Euro Area at the onset of the 2020s inflationary episode. This novel focus allows us to provide unique insights into how monetary policy changes can differentially impact consumption patterns across wealth spectrum, contributing to a deeper understanding of policy effects in heterogeneous agent settings.

2 A HANK Model for the Euro Area

In this section, we introduce the main features of our model. Time, indexed by t, is discrete and each period corresponds to a quarter. There are several types of agents: households, firms, unions, a financial intermediary, and monetary and fiscal authorities. We describe each type in the following subsections.

Households. There is a continuum of households indexed by i. Households are heterogeneous in three dimensions: labor productivity z_{it}, illiquid asset holdings a_{it}, and liquid asset holdings b_{it}. They derive utility from consumption and disutility from hours worked according to the preferences

$$\mathbb{E} \sum_{t=0}^{\infty} \beta^t \exp(-\varepsilon_t^C) \left[u(c_{i,t}) - v(n_{i,t}) \right]. \tag{1}$$

Households discount future utility with a factor β and are subject to a preference shock ε_t^C, which affects all households identically and follows an AR(1) process. In addition, they face information frictions as in Auclert, Rognlie and Straub (2020): at every point in time they know the current values of their idiosyncratic states, but only learn about the values of aggregate variables with i.i.d. probability $1-\theta^C$.[9] Therefore, the expectation operator in (1) depends on the information set of each agent.

Household after-tax labor income is given by

$$e_{it} = (1-\tau_t^L) w_t n_{it} z_{it}. \tag{2}$$

The income from labor is the product of the real wage w_t, the number of hours worked n_{it}, and labor productivity z_{it}, and is subject to proportional taxation at a rate τ_t^L. We assume that labor productivity z_{it} follows an AR(1) process in logs:

$$\log z_{it} = \mu^z + \rho^z \log z_{i,t-1} + \sigma^z \varepsilon_{it}.$$

The constant μ^z is normalized so that average labor productivity $\int z_{it} di$ equals 1.

Households do not directly choose the number of hours they work. Instead, there are unions that set nominal wages, and households commit to supplying as many hours as demanded by firms at the prevailing wage. For simplicity, we follow Auclert, Rognlie and Straub (2020) in assuming that unions split the number of hours equally across agents, i.e., $n_{it} = n_t$.

Households can save in liquid or illiquid asset. We model portfolio adjustment frictions as in Auclert, Rognlie and Straub (2020) and Bayer, Born and Luetticke (2024): households can only move funds across liquid and illiquid accounts with exogenous probability χ. This lottery is i.i.d. across agents and happens every period. Real ex-post returns on illiquid and liquid assets are denoted, respectively, by r_t^a and r_t^b. In steady state, households will only hold illiquid assets if they yield higher returns, which must be the case in equilibrium. However, following an aggregate shock, r_t^a may fall temporarily below r_t^b. Households face the following budget constraint:

$$a_{it} + b_{it} + c_{it} = e_{it} + (1+r_t^a) a_{i,t-1} + (1+r_t^b) b_{i,t-1}.$$

A household who is able to rebalance his or her portfolio can freely choose illiquid (a_{it}) and liquid (b_{it}) asset holdings. Otherwise, they must choose $a_{it} = (1+r_t^a) a_{i,t-1}$. Additionally, they are also subject to the constraints

$$a_{it} \geq 0 \quad b_{it} \geq 0.$$

[9]Formally, this introduces another dimension of heterogeneity, namely the number of periods an agent has remain uninformed. Auclert, Rognlie and Straub (2020) show, however, that it is not necessary to keep track of this additional state when solving for the first-order dynamics of the model.

Firms. There are three types of firms in our model: intermediate goods producers, final goods producers, and capital goods producers. Intermediate goods producers are identical and perfectly competitive, and operate according to the following Cobb-Douglas production function:

$$Y_t = \exp(\varepsilon_t^Z) Z K_t^{\alpha} N_t^{1-\alpha} - \Phi.$$

The total amounts of capital and labor input are denoted by K_t and N_t, respectively, while total factor productivity is determined by its steady state level Z and transitory aggregate shocks ε_t^Z. Firms also face a fixed cost Φ, which will be used in the calibration to match total household wealth.

Intermediate goods producers rent capital from capital producers at rate r_t^K and hire labor from unions at real wage w_t. Their real marginal cost, which due to perfect competition equals the real price they charge, is given by

$$mc_t = \frac{1}{\alpha^{\alpha}(1-\alpha)^{1-\alpha}} \frac{(r_t^K)^{\alpha} w_t^{1-\alpha}}{\exp(\varepsilon_t^Z) Z}.$$

Final goods producers use intermediate goods to produce different varieties of final goods. Their output, given by a CES aggregator as in Dixit and Stiglitz (1977) with elasticity of substitution ζ^p, is used for both final consumption and investment. Final goods producers face nominal price stickiness as in Calvo (1983). Every period, a firm can adjust its price with i.i.d. probability $1-\theta^p$. The ones who cannot adjust their prices partially index them to past inflation, with the degree of indexation denoted by ι^p. After log linearizing the model around the steady state, we arrive at the same Phillips Curve formulation as in Smets and Wouters (2007):

$$\pi_t - \iota^p \pi_{t-1} = \frac{1}{1+r} \mathbb{E}_t(\pi_{t+1} - \iota^p \pi_t) + \kappa^p \widehat{mc}_t + \varepsilon_t^p,$$

where π_t is the inflation rate and ε_t^p is a price markup shock. We use hats to denote log deviations from steady state and variables without subscripts to refer to steady state values, e.g. r is the steady state real interest rate, according to which risk neutral firms discount future profits. The Phillips curve slope is $\kappa^p = (1-\theta^p)(1-\theta^p/(1+r))/\theta^p$.

Capital goods producers transform final goods into investment goods subject to investment adjustment costs. They own the capital stock and rent it to firms at the rate r_t^K. They maximize the net present value of future profits

$$\mathbb{E}_t \sum_{s=0}^{\infty} \frac{1}{R_{t,t+s}} \left(r^K_{t+s} K_{t+s} - I_{t+s} \right),$$

where $R_{t,t} = 1$ and, for $s \geq 1$,

$$R_{t,t+s} = \prod_{k=1}^{s} (1 + r_{t+k}).$$

The capital stock evolves according to

$$K_{t+1} = (1-\delta)K_t + I_t \left[1 - \frac{1}{2}\gamma^I \left(\frac{I_t}{I_{t-1}} - 1 \right)^2 \right] \exp(\varepsilon^I_t).$$

Above, I_t denotes investment, δ is the depreciation rate, γ^I is the adjustment cost parameter, and ε^I_t is an investment-specific technology shock.

Unions. We model unions as in Hagedorn, Manovskii and Mitman (2019) and Auclert, Rognlie and Straub (2024). There is a continuum of unions, each one providing differentiated labor services under monopolistic competition. Each household supplies labor to all unions. Labor services are bundled according to a CES aggregator with elasticity of substitution ζ^w. Similarly to firms, unions set nominal wages subject to Calvo frictions with partial indexation to past inflation. We assume that unions maximize the welfare of a fictitious average agent with consumption $C_t = \int c_{it} di$, which results in the following linearized Phillips Curve:

$$\pi^w_t - \iota^w \pi_{t-1} = \beta \mathbb{E}_t (\pi^w_{t+1} - \iota^w \pi_t) + \kappa^w \widehat{\mu}^w_t,$$

where

$$\mu^w_t = \frac{v'(N_t)}{(1-\tau_t) w_t u'(C_t)} \tag{3}$$

and π^w_t denotes wage inflation. As is standard in sticky-wage models, the intercept μ^w balances the disutility of working an extra hour and the utility of consuming the resulting earnings. The Phillips curve slope is $\kappa^w = {}^{(1-\theta^w)(1-\beta\theta^w)}/_{\theta^w}$, where $1 - \theta^w$ is the probability of wage adjustment.

Financial intermediary. We model household balance sheets in the same way as Auclert, Rognlie and Straub (2024). There is a risk-neutral, perfectly competitive financial intermediary that holds government bonds and firm equity in illiquid form, and provides liquid assets to households at a unit cost ξ.

The ex-dividend real value of firm equity is given by

$$p_t^e = \mathbb{E}_t \frac{d_{t+1} + p_{t+1}^e}{1 + r_{t+1}}, \tag{4}$$

where d_t is the sum of the real dividends of all firms and is given by

$$d_t = Y_t - w_t N_t - I_t,$$

and r_t is the real interest rate.

Government bonds pay coupons that decreases exponentially by a factor ρ^{bonds}. This is quantitatively important since long-term nominal bonds may display substantial variation in real prices following a monetary shock, which may generate sizable wealth effects on households. Nominal bond prices are given by

$$P_t^g = \frac{1 + \rho^{bonds} \mathbb{E}_t P_{t+1}^g}{1 + i_t}, \tag{5}$$

where i_t is the nominal interest rate set by the monetary authority. In turn, nominal and real interest rates are related by the usual Fisher equation, expressed below in terms of ex-post real returns:

$$1 + r_t = \frac{1 + i_{t-1}}{1 + \pi_t}. \tag{6}$$

We assume that liquid assets provided by the financial intermediary have zero maturity, and thereforetheir real returns are given by $r_t^b = r_t - \xi$. The real returns on illiquid assets are then simply equal to the rate of return on the financial intermediary's portfolio, which is a weighted average of the returns on government bonds, equity, and liquid assets. The latter enters with a negative weight, as liquid assets corresponds to the financial intermediary's liabilities.

Monetary and fiscal policies. The fiscal authority levies proportional taxes on households' labor income to finance government consumption G_t and to service its debt. The government budget constraint is given by

$$B_t^g = (1 + r_t^g) B_{t-1}^g + G_t - \tau_t w_t N_t,$$

where B_t^g is the real market value of outstanding debt and r_t^g its real ex-post return, and $G_t$is real government consumption that is subject to shocks:

$$G_t = G + Y\varepsilon_t^G.$$

The government spending shock is scaled by steady state output.

To ensure determinacy, we assume the fiscal authority reacts to the debt level according to the same rule as in Auclert, Rognlie and Straub (2020):

$$\tau_t = \tau + \phi^B \frac{B_{t-1}^g - B^g}{Y}.$$

The monetary authority, in turn, follows a standard Taylor rule:

$$i_t = \phi^i i_{t-1} + (1 - \phi_i)\left(i + \phi^\pi \pi_t\right) + \phi^{\Delta y} \frac{Y_t - Y_{t-1}}{Y_{t-1}} + \varepsilon_t^i,$$

where ε_t^i is a monetary shock.

Market clearing. Our model has two market clearing conditions. Asset markets clear when the market value of government debt and firm equity equals the aggregate holdings of illiquid and liquid assets, denoted A_t and B_t, respectively:

$$B_t^g + p_t^e = A_t + B_t.$$

The goods market clearing condition states that output equals the sum of aggregate consumption C_t, investment I_t, government spending G_t, and the physical costs incurred by the financial intermediary to provide liquid assets ξB_{t-1}:

$$Y_t = C_t + I_t + G_t + \xi B_{t-1}.$$

3 Calibration and Estimation

In this section, we discuss the calibration and estimation of the model parameters. We follow a two step approach, as is typical in models with heterogeneous agents, where computing the deterministic steady state of the model is computationally costly[10]. First, we calibrate the model parameters that affect the steady state. Second, we estimate the remaining parameters using Bayesian methods and aggregate time series data. We solve for the first-order dynamics around the deterministic steady state using the Sequence Space Jacobian toolbox introduced by Auclert, Bardóczy, Rognlie and Straub (2021).

Calibration. Consumer preferences are given by

$$u(c) - v(n) = \log c - \psi \frac{n^{1+\varphi}}{1+\varphi}.$$

In the absence of wage rigidity, the parameter φ corresponds to the inverse Frisch elasticity of labor supply. We set $\varphi = 2$ and normalize ψ so that in steady state $n = 1$. We calibrate the persistence of household earnings to $\rho^z = 0.978$, based on estimates from Floden and Lindé (2001)[11]. Following Auclert, Rognlie and Straub (2024), we calibrate the standard deviation of the distribution of log income to be $\sigma^z / \sqrt{1-(\rho^z)^2}$=0.92.

The quarterly real interest rate on illiquid assets is $r^a = 1\%$, which in steady state is also the returns on equity and government bonds. Given this, we then jointly calibrate the household discount factor β, the return on liquid assets r^b (or, equivalently, the cost of providing liquidity ξ), and the probability of accessing the illiquid assets χ in order to match three moments: (i) total household wealth of 285% of annual GDP, (ii) liquid asset holdings of 45% of annual GDP, and (iii) an average quarterly marginal propensity to consume (MPC) of 20%. Moments (i) and (ii) correspond to estimates from the Euro Area using data from the Household Finance and Consumption Survey (HFCS)[12][13]. We calibrate the average MPC to be close to the the ones typically found in the literature (Johnson et al., 2006; Kekre, 2023). We obtain parameter values $\beta = 0.985$, $\chi = 11.9\%$, and $r^b = -0.09\%$. The real return on liquid assets is slightly negative, as in Auclert et al. (2024).

Government debt B^g is set at 50% of annual GDP, corresponding to the total amount of held by household according to the HFCS. The factor of decay of government bonds ρ^{bonds} is chosen to generate an average bond duration of 10 years, as in Coenen et al. (2018), which requires

[10] See for instance Auclert, Rognlie and Straub (2020) andAcharya et al. (2023).

[11] Floden and Lindé (2001) estimate $\rho^{z,annual} = 0.9136$ at annual frequency, which we convert to quarterly frequency by setting $\rho^z = (\rho^{z,annual})^{1/4}$.

[12] European Central Bank (2023) Household finance and consumption survey - Results from the 2021 wave.

[13] We thank Michal Brzoza-Brzezina for proving us with these numbers.

$\rho^{bonds} = 0.985$. Government purchases in our model amount to 20.9% of annual GDP, the average observed value between years 2000 and 2019. These fiscal parameters imply a steady state labor income tax rate $\tau = 33.1\%$.

Turning to the supply side of the model, we calibrate the elasticity of substitution across varieties $\zeta^p = 11$, consistent with a 10% price markup over marginal costs, and use the same value for the elasticity of substitution between labor varieties provided by different unions ζ^w. Aggregate investment I is set at 21.4% of output and the depreciation rate is $\delta = 2.5\%$. Finally, the fixed cost parameter Φ is chosen to target a market value of firm equity equal to 235% of annual GDP, consistent with the calibrated values of total household wealth and government debt. This entails a value of 9.1% of GDP. Table 1 summarizes the calibration of the model.

Parameter	Description	Value
Households: Preferences and earnings		
φ	Labor disutility curvature	2
β	Discount factor	0.985
ρ^z	Idiosyncratic labor productivity persistence	0.978
$\sigma^z/\sqrt{1-(\rho^z)^2}$	Standard deviation of labor productivity	0.92
Households: Balance sheets		
r^a	Returns on illiquid assets	1%
r^b	Returns on liquid assets	-0.09%
χ	Probability of accessing illiquid assets	11.9%
Government		
B^g/Y	Government debt	50%
G/Y	Government spending	20.9%
τ	Labor income taxes	33.1%
ρ^{bonds}	Bonds coupon decay factor	0.985
Firms and unions		
$\varepsilon^p, \varepsilon^w$	Elasticity of substitution	11
δ	Depreciation rate	2.5%
I/Y	Aggregate investment	21.4%
Φ/Y	Fixed costs in production	9.1%

Table 1: Calibrated parameters

As some of the exercises conducted in the next section are related to the wealth distribution, it is important to understand whether our model generates realistic levels of wealth inequality.

Table 2 shows wealth inequality statistics from Distributional Wealth Accounts (DWA) produced by the European System of Central Banks and compares it to our model. Household wealth in the model is defined as the sum of liquid and illiquid wealth:

$$\text{wealth} = (1 + r_t^a)a_{i,t-1} + (1 + r_t^b)b_{i,t-1}.$$

Except for underestimating the wealth share of the top 5%, the model captures remarkably well the wealth distribution observed in the data.It is particularly noteworthy because we do not target any moments related to it in our calibration.

Wealth shares (%)	Model	Data (2024 Q1)
Top 5%	31.3	44.3
Next 5%	16.1	13.0
Bottom 50%	6.3	5.1
Gini coefficient	0.66	0.72

Table 2: Wealth distribution

Estimation. We estimate our model using Bayesian methods and time series data. First, it is necessary to specify stochastic processes for the exogenous shocks that determine the dynamics of the model. Consumer demand (ε_t^C), investment (ε_t^I), government spending (ε_t^G), total factor productivity (ε_t^Z), price markup (ε_t^p), and wage markup shock (ε_t^w) are all assumed to follow AR(1) processes: $\varepsilon_t^x = \rho_t^x \varepsilon_{t-1}^x + \sigma^x \eta_t^x$, where $x \in \{C, I, G, Z, p, w\}$ and the disturbances $\eta^x \sim NID(0, 1)$.The monetary shock (ε_t^i) is assumed to be i.i.d. as in Coenen et al. (2018).

Our estimation sample ranges from 2000Q1 to 2019Q4, stopping before the COVID pandemic, and includes eight Euro Area time series: the quarterly growth rates of (i) real GDP, (ii) consumption, (iii) aggregate investment, (iv) government consumption, (v) total hours worked, and (vi) nominal compensation per employee, as well as (vii) quarterly inflation and (viii) nominal interest rates, measured by the 3-month Euribor rate. The mappings from these time series to model objects are straightforward and follow Smets and Wouters (2007). As we have one more observable than shocks, we introduce measurement error in the GDP growth rate. This also helps accommodate other types of model misspecification, e.g. due to the absence of net exports or inventories. The standard deviation of this measurement error, denoted $\sigma^{\Delta y}$ is fixed at 0.25% p.a.

We estimate the remaining parameters of the model, which include the degree of information stickiness of households θ^c, the parameters that govern price and wage rigidity, the investment adjustment cost γ^I, and the autocorrelation and standard deviation of all exogenous shocks. The only exception is the fiscal rule parameter, which we calibrate as $\phi^B = 0.1$. This falls within the range of empirical estimates surveyed by Auclert et al. (2020).[14] Importantly, we estimate directly the slopes of the price and wage Phillips curves (κ^p and κ^w, respectively), instead of the frequencies of price and wage changes. Table 3 shows moments of the prior and posterior distribution. The latter is obtained using a Multi-proposal Parallel Metropolis-Hastings algorithm from (Calderhead, 2014)[15].

To illustrate the aggregate dynamics implied by our estimation, Figure 1 shows selected impulse response functions to a one standard deviation monetary shock. The solid line corresponds to the average across 1000 draws from the posterior distribution, while the shaded areas range from the 10th and 90th percentiles. Results are broadly similar to the ones from Coenen et al. (2018), with the main difference being the smaller response of output. This results from our closed economy setting and, therefore, from the lack of an exchange rate channel of monetary policy. Impulse responses to other shocks are shown in Appendix A, together with shock decompositions and other related results.

4 Distributional Effects of Tighter Monetary Policy

We now turn to the applications of our model for policy analysis, focusing on the recent inflationary episode in the Euro Area. Specifically, we explore a counterfactual monetary policy scenario, to examine how an alternative, more hawkish monetary policy stance would have impacted households across the wealth distribution. We create counterfactual monetary policy scenarios following the inflationary surge of the early 2020s. In particular, we examine a scenario where the European Central Bank begins raising interest rates in the second quarter of 2022 – one quarter earlier than it actually did – and maintained higher rates than the observed policy until the first quarter of 2023. We implement this scenario by backing out a sequence of

[14] According to Auclert et al. (2020), empirical estimates of the response of tax receipts to public debt over (annual) GDP typically fall in the range 0.015 – 0.3. In our model tax revenues are given by $\tau w N$, so the implied sensitivity is $4\phi^B w N \approx 0.28$, where the factor 4 comes from the conversion from quarterly to annual GDP.

[15] The resulting series of posterior samples of estimated parameters are shown in the appendix.

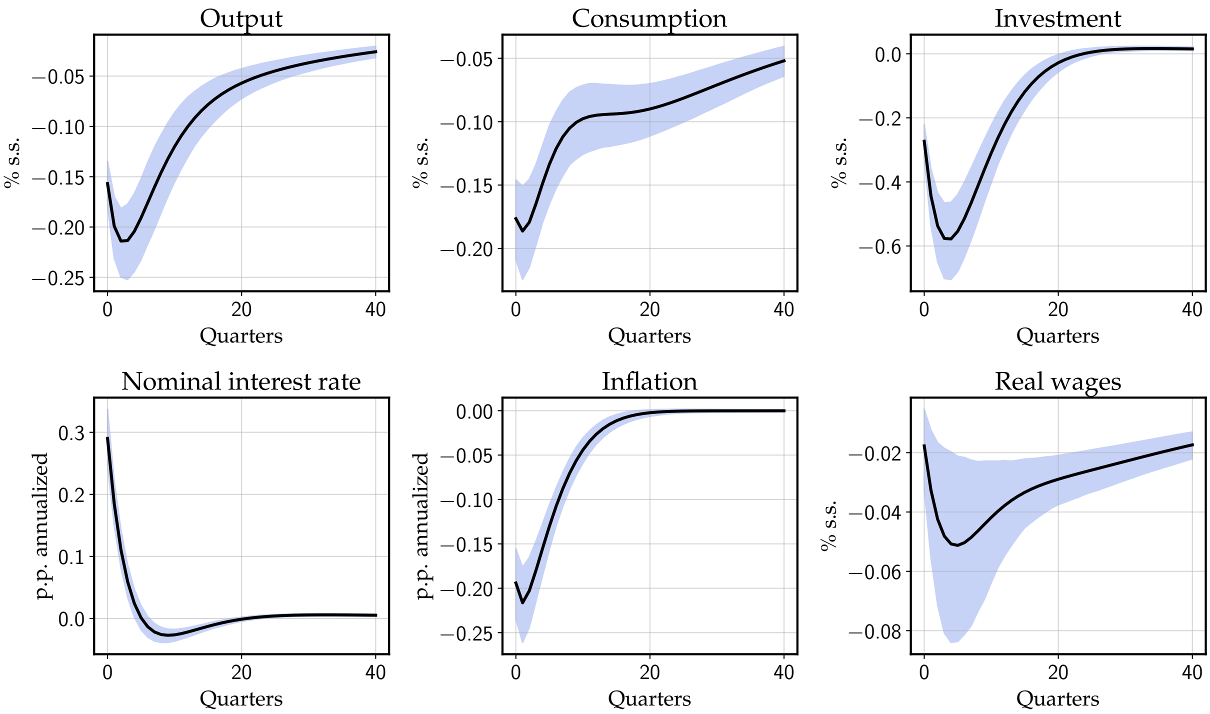

Figure 1: Impulse response functions to a monetary shock

unanticipated monetary shocks that would implement the desired interest rate path. Table 4 shows the differences between actual and counterfactual average quarterly interest rates. Although arbitrary, this interest rate path allows us to illustrate well the properties of our model.

Figure 2 shows results for selected aggregate variables, normalizing output and consumption to be 1 in 2022Q1, right before our counterfactual scenario begins. In the top panels, we observe that output and consumption would have been lower if the European Central Bank had started raising interest rates earlier, as expected. Inflation, shown in the bottom left panel, would have been slightly lower than observed, consistent with the impulse response functions from Figure 1 and in line with a relatively flat Phillips curve.

With an estimated HANK model, we can simulate the effects of alternative monetary policy scenarios not only on aggregate variables but also across the wealth distribution. To do so, we define a set of thresholds $0 = \overline{w}_0 < \overline{w}_1 < \cdots < \overline{w}_N < \overline{w}_{N+1} = \infty$ and divide households into groups according to their wealth. Let $\overline{C}_{kt}$ be the per capita consumption of households entering period t with wealth in the interval $[\overline{w}_{k-1}, \overline{w}_k)$. Since we lack data on the $\overline{C}_{kt}$ variables, we use our model to filter their values based on the observed time series used in the estimation and

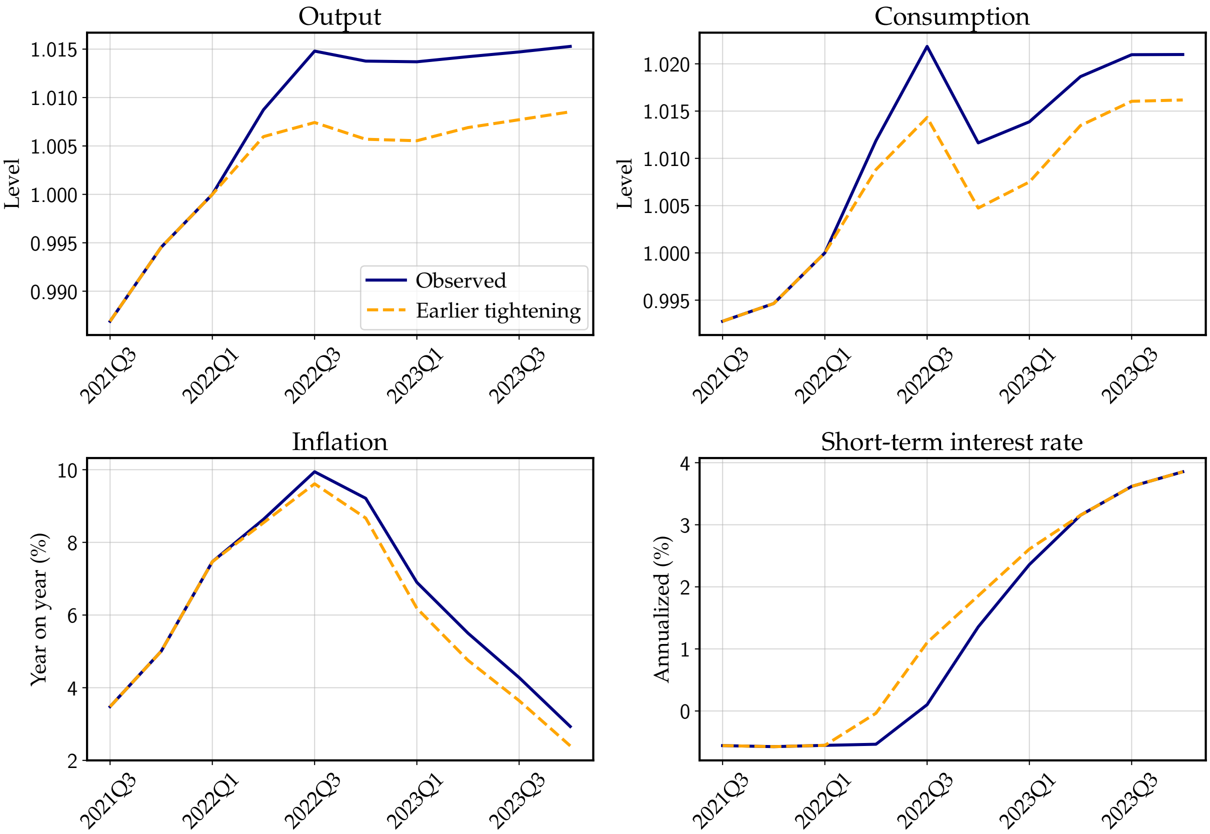

Figure 2: Counterfactual paths for aggregate variables

Note: Output and consumption are presented as normalized levels, with 2022Q1, right before the start of the counterfactual scenario, set as the baseline value of one.

then construct counterfactual paths around the filtered series.

Figure 3 presents the results using the steady state quartiles of the wealth distribution as thresholds. All consumption series are shown in terms of deviations from the steady state consumption of each quartile. Two key findings emerge. First, the filtered average consumption of the first quartile of the wealth distribution, shown in the upper left panel, starts from a level lower than those of the other groups.This suggests that, according to our model, consumption at the lower end of the wealth distribution took longer to recover from the COVID recession. Second, and most importantly, the gap between the filtered and simulated values is larger for poorer households, consistent with the results from Figure 5. Our model indicates that an earlier tightening of interest rates would have had more severe effects on the consumption at the bottom of the wealth distribution.

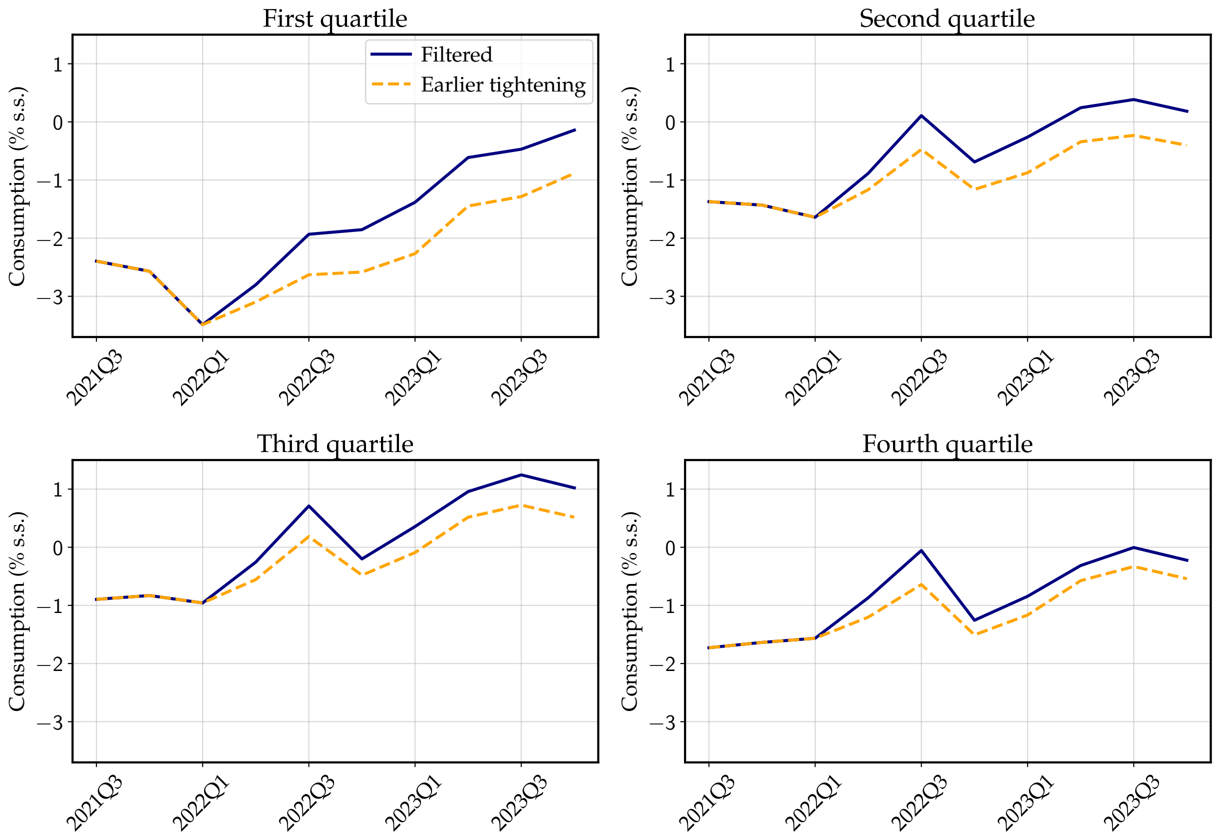

Figure 3: Counterfactual paths of per capita consumption across the wealth distribution
Note: Each line shows the deviation of consumption to the corresponding quartile's steady state under a different scenario. Steady state consumption varies across wealth quartiles.

4.1 The transmission channels of monetary policy across the wealth distribution

To shed light on why these differences arise, we introduce a decomposition of the effects of monetary policy on consumption along the wealth distribution. As already mentioned, we solve our model using the Sequence Space Jacobian method, which hinges on the first-order equivalence between aggregate shocks and perfect foresight transitions. We focus, therefore, on a perfect foresight monetary shock happening at $t = 0$. Following such a shock, household consumption depends on three sequences of aggregate variables: the aggregate component of disposable labor income $\{y_t\}_{t=0}^{\infty}$, defined as $y_t = (1 - \tau_t)w_t n_t$, returns on illiquid assets $\{r_t^a\}_{t=0}^{\infty}$, and returns on liquid assets $\{r_t^b\}_{t=0}^{\infty}$.[16]

Figure 4 shows the responses of disposable labor income and asset returns to a monetary shock computed at the posterior mode. Disposable income displays a negative, hump-shaped

[16]This holds for all shocks but the preference one, for which the sequence $\{\varepsilon_t^C\}_{t=0}^{\infty}$ must also be included in the decomposition.

path, as is common in DSGE models. In order to understand the apparent "discontinuity" in the paths of assets returns at $t = 0$, recall that these are real *ex-post* returns. On impact, they are, therefore, affected by unexpected changes in asset prices or inflation that arises as a consequence of the shock. This effect is particularly strong for the return on illiquid assets, since it depends largely on the returns on equity and government bonds, which can be obtained from equations (4) and (5), respectively. Since both assets display long duration in our model – government bonds by assumption and equity because it represents a stream of discounted future dividends – r_t^a displays a sharp initial drop. However, for $t \geq 1$, the perfect foresight nature of the shock ensures that both r_t^a and r_t^b equal the real interest rate defined in (6).

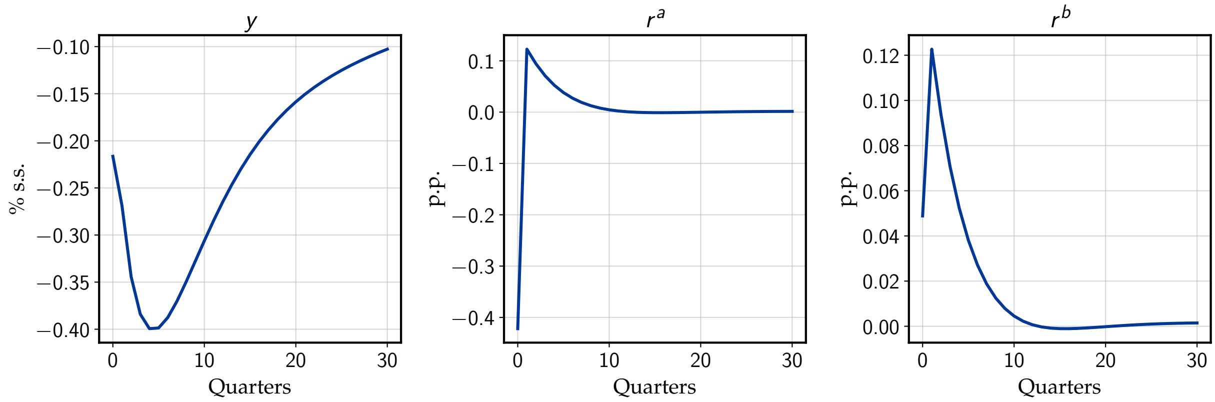

Figure 4: Responses of disposable labor income and asset returns following a monetary shock

This discussion motives a decomposition of the consumption response to a monetary shock in three components: (i) a *disposable labor income* term, (ii) a *capital gains* term, which reflects the unanticipated asset returns at $t = 0$, and (iii) an *expected asset returns* term, capturing the higher returns starting at $t = 1$. More specifically, we have the decomposition

$$dC_t = \underbrace{\sum_{s=0}^{\infty} \frac{\partial C_t}{\partial y_s} dy_s}_{\text{disposable income}} + \underbrace{\frac{\partial C_t}{\partial r_0^a} dr_0^a + \frac{\partial C_t}{\partial r_0^b} dr_0^b}_{\text{capital gains}} + \underbrace{\sum_{s=1}^{\infty} \left(\frac{\partial C_t}{\partial r_s^a} dr_s^a + \frac{\partial C_t}{\partial r_s^b} dr_s^b \right)}_{\text{expected asset returns}},$$

where d denotes deviations from steady state. Note that while the capital gains term captures only wealth effects, the expected returns component also reflects intertemporal substitution. Figure 5 shows the result of this decomposition. To better understand the effects of monetary policy across the wealth distribution, we divide the model population into six groups by wealth: four quartiles, plus the bottom and top 10%. Each line represents the change in the consumption of a given group in terms of its steady state value relative to the path that would have prevailed without the monetary shock.[17]

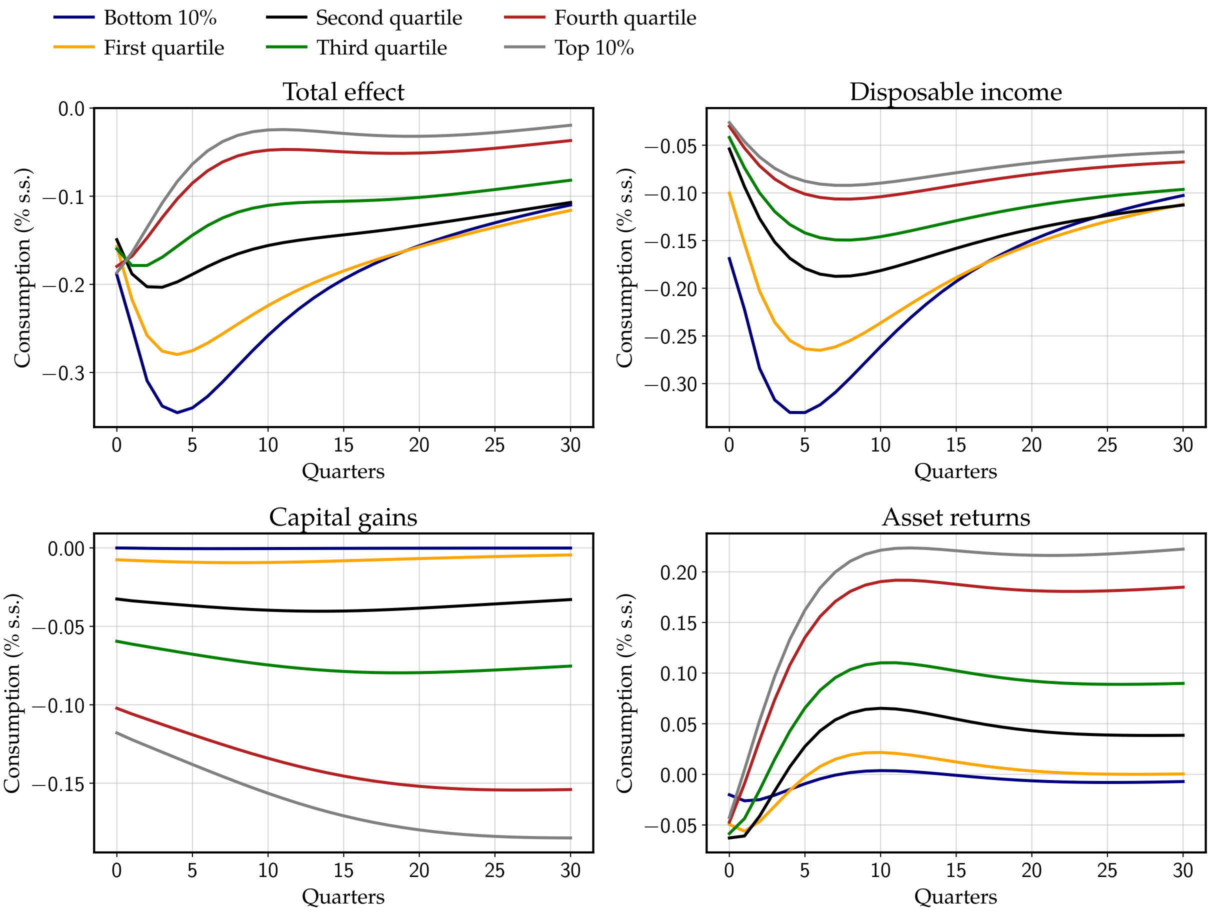

Figure 5: Decomposition of the consumption response to a monetary shock

Our results offer a clear picture to understand the transmission channels of monetary policy across the wealth distribution. For instance, the transmission through disposable income, displayed in the upper right panel of Figure 5, is stronger at the bottom of the wealth distribution. This is intuitive since these households have few assets that can be used to smooth consumption following an unexpected drop in income. In fact, for the bottom 10% the consumption drop is very close to the drop in disposable income displayed in Figure 4, consistent with MPCs close to 1. In contrast, the consumption of the top 10% is much less affected through this channel.

[17] In our framework, as in most heterogenous-agent models, the distribution of agent displays mixing, i.e., the average consumption of any positive mass of agents converges over time to the average steady state consumption in the absence of aggregate shocks, hence the need to consider differences from the paths obtained in the absence of the monetary shock.

On the other hand, the bottom of the wealth distribution is barely affected by the capital gains and asset returns channels, shown in the bottom two panels of Figure 5. For all other groups, especially at the top, those two channels have very persistent effects of broadly opposite signs. The unexpected drop in assets prices depresses the consumption of households, with stronger effects on those with greater exposure to it. Conversely, higher asset returns have the opposite effect, except for a small negative effect in the initial periods following the shock. This is due to the stronger incentives to save in response to persistently higher returns – the intertemporal substitution effect.

These results are broadly in line with the findings from Lee (2024) and Del Negro et al. (2025), which also highlight the varied effects of monetary policy on different households. First, although tighter monetary policy can prevent inflation from eroding real wages, a more aggressive response to inflation disproportionately harms the bottom of the wealth distribution by suppressing output and labor demand, thereby reducing real wages. Additionally, in their model, poorer households suffer from higher interest rates on debt. In contrast, our model does not include borrowing, so the main impact is through the labor market. Second, a contractionary monetary policy shock benefits wealthier households through higher interest rates on assets. However, in our model, this effect is not offset by strongly procyclical and declining profits. As a result, the total negative effect on consumption is mostly decreasing in wealth, contrasting with the inverted U-shaped pattern documented by Lee (2024) and Del Negro et al. (2025).

5 Forecasting accuracy of HANK relative to RANK

To demonstrate that the inclusion of heterogeneous households in macroeconomic models does not inherently compromise their ability to deliver accurate forecasts of aggregate variables, we compare the forecasting performance of our HANK model with that of a comparable representative-agent (RANK) model. To ensure a clear focus on the role of heterogeneity, we construct the closest possible RANK model to our setting by replacing the heterogeneous household block, in the terminology of Auclert et al. (2021), by a representative agent with preferences given by

$$\mathbb{E}\sum_{t=0}^{\infty}\beta^t \exp(-\varepsilon_t^C)\left[u(C_t - hC_{t-1}) - v(N_t)\right].$$

Importantly, here we allow for habit formation in preferences with a parameter $0 \leq h < 1$, rather than information frictions, as this is the standard approach to generating hump-shaped impulse responses in this setting. All other model components remain unchanged, except for a minor adjustment to the unions' problem (3) to account for habit formation in preferences. We estimate this model using the same data as the HANK specification.

In this exercise we conduct out-of-sample forecasts for the growth rates of GDP, consumption, and investment, and consumer price inflation, and calculate Root Mean Squared Errors (RMSE) for various forecast horizons. To perform the out-of-sample forecasting, we compute posterior modes of parameters from both models using only the first half of our data (ranging from 2000Q1 to 2009Q4) and the same set of prior distributions and generate forecasts using these values. Details are shown in Appendix B.

We proceed iteratively by computing forecasts 8 periods ahead for each period starting from 2010Q1 up to 2019Q4. At every step, we only use data available up to that point in time. However, because we demean variables for estimation, at each step we subtract the means of the variables up to that point. We then use these demeaned variables in our models to obtain forecasts. Therefore, parameter values in this exercise are derived using data up to 2009Q4, except for time series means, which are interpreted as being estimated in real time.

Figure 6 shows the RMSE for the four variables using both models. The HANK model deliverers slightly more accurate forecasts for inflation and short-term GDP growth. However, the difference to the RANK model is very small, and may not be robust to parameter uncertainty or different sample periods. More surprising is the fact that both models generate equally accurate forecasts for consumption growth, despite the substantial differences in the way households are modeled. This highlights the importance of a mechanism that generates hump-shaped impulse response functions of aggregate consumption in a HANK model. The lack of such frictions could be one of the reasons why Acharya et al. (2023) find contrasting results, in which their RANK model displays superior forecasting performace to the HANK model that is missing this crucial feature. Regarding investment, both models perform similarly.

To further clarify the differences and similarities between the models, Figure 7 displays the one-quarter-ahead forecasts generated by the HANK model compared to those from the RANK model. We have maintained the models as comparable as possible, differing only in the modeling of the household sector. As a result, the forecasts show a very high correlation,

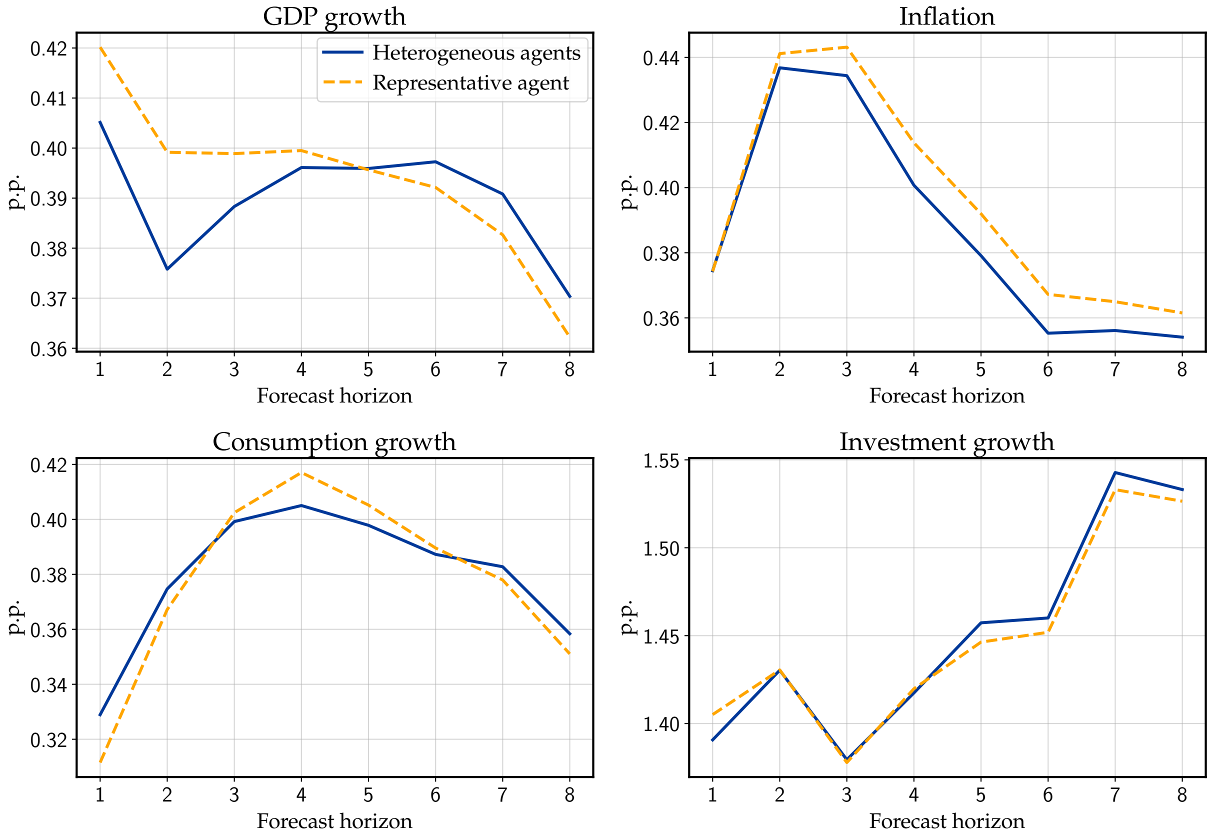

Figure 6: RMSE of forecasts for HANK and RANK models

particularly in predictions for investment, inflation, and GDP growth. The most significant structural differences between the models arise from the household modeling approach, which is evident in the figure's bottom left panel as a notably lower correlation between the HANK and RANK predictions. Despite these differences, the forecasting performance for the one-period-ahead horizon is similar for both models in terms of root mean square error (RMSE), as shown in Figure 6. This, combined with the lower correlation in predictions of consumption growth, highlights the complementary roles of the models in economic forecasting.

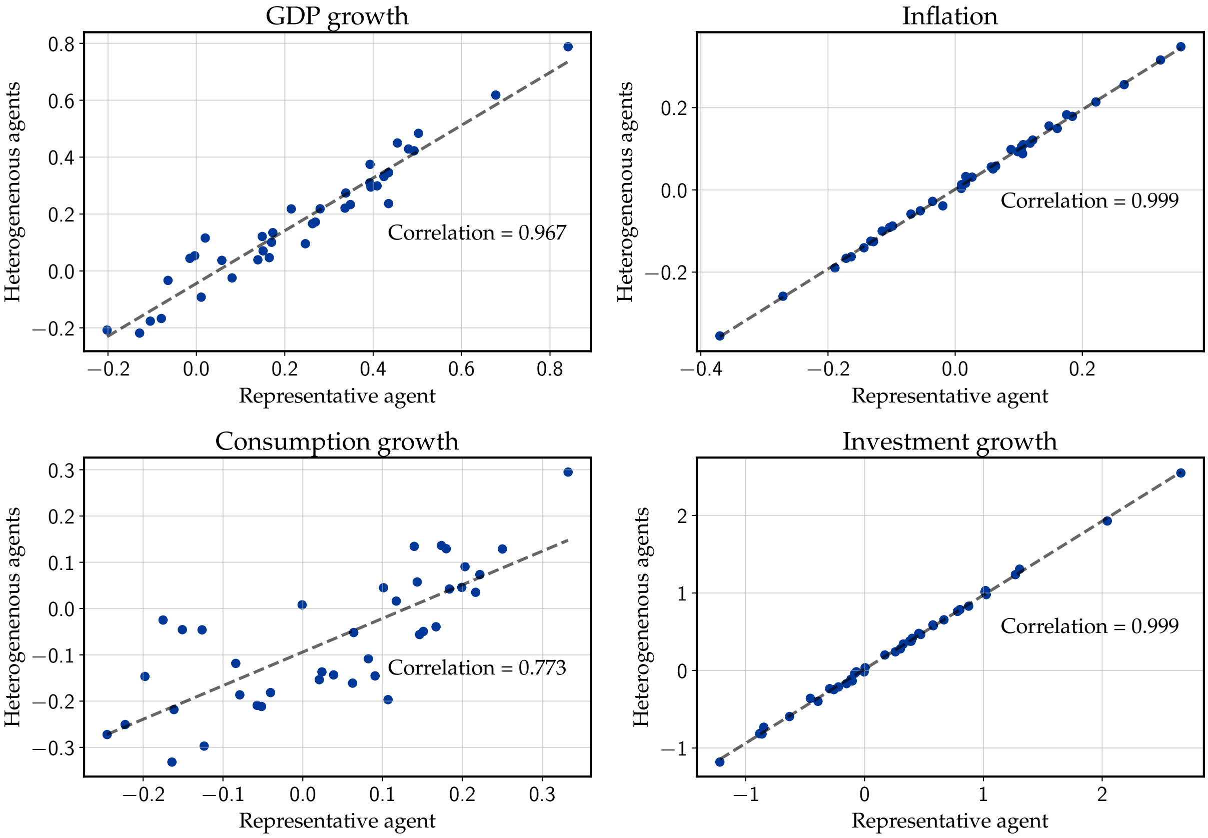

Figure 7: Correlations between one-period-ahead forecasts in the HANK and RANK models

Parameter	Description	Prior			Posterior		
		Distribution	Mean	S.d.	Mode	10%	90%
Households							
θ^c	Information stickiness	Beta	0.5	0.10	0.79	0.71	0.85
Firms and unions							
κ^p	Price rigidity	Gamma	0.05	0.015	0.045	0.034	0.068
κ^w	Wage rigidity	Beta	0.05	0.015	0.014	0.009	0.024
ι^p	Price indexation	Beta	0.5	0.15	0.16	0.10	0.29
ι^w	Wage indexation	Beta	0.5	0.15	0.19	0.12	0.33
γ^I	Investment adj. cost	Normal	4	1	4.4	3.6	5.7
Monetary policy							
ϕ^i	Interest rate smoothing	Beta	0.7	0.05	0.89	0.86	0.90
ϕ^π	Taylor rule π coefficient	Normal	2	0.1	1.85	1.72	1.97
$\phi^{\Delta y}$	Taylor rule Δy coefficient	Normal	0.05	0.05	0.14	0.12	0.17
Shock autocorrelation							
ρ^C	Consumption shock	Beta	0.5	0.1	0.89	0.86	0.91
ρ^I	Investment shock	Beta	0.5	0.1	0.59	0.49	0.65
ρ^G	Gov. spending shock	Beta	0.5	0.1	0.93	0.90	0.94
ρ^Z	TFP shock	Beta	0.5	0.1	0.85	0.80	0.89
ρ^p	Price markup shock	Beta	0.5	0.1	0.46	0.35	0.55
ρ^w	Wage markup shock	Beta	0.5	0.1	0.74	0.64	0.81
Shock standard deviation							
σ^C	Consumption shock	Inv. Gamma	0.1	2	0.15	0.12	0.19
σ^I	Investment shock	Inv. Gamma	0.1	2	3.7	3.0	5.1
σ^G	Gov. spending shock	Inv. Gamma	0.1	2	0.079	0.072	0.090
σ^Z	TPF shock	Inv. Gamma	0.1	2	0.30	0.28	0.34
σ^p	Price markup shock	Inv. Gamma	0.1	2	0.30	0.27	0.37
σ^w	Wage markup shock	Inv. Gamma	0.1	2	0.20	0.17	0.26
σ^i	Monetary shock	Inv. Gamma	0.1	2	0.100	0.092	0.120

Table 3: Estimated parameters

Period	Difference (annualized p.p.)
2022Q2	0.5
2022Q3	1.0
2022Q4	0.5
2023Q1	0.25

Table 4: Difference between observed and counterfactual interest rates

6 Conclusion

In this paper, we explore the use and advantages of Heterogeneous-Agent New Keynesian (HANK) models for central banks, with a particular emphasis on policy analysis and forecasting in the Euro Area. We demonstrate how these models provide deeper insights into the transmission of monetary policy and its distributional effects, offering a richer perspective than their representative-agent counterparts.

First, we focus on the recent monetary policy tightening cycle and analyze a counterfactual scenario in which the European Central Bank had initiated the tightening cycle one quarter earlier. We emphasize the additional dimension that the introduction of a heterogeneous agent model adds to the suite of models: the differentiated effects along the income and wealth distribution. Our findings reveals that post-COVID pandemic, the consumption of households at the lower end of the asset distribution was already depressed relative to the rest of the population. An earlier and more aggressive tightening would have disproportionately affected these households, further compressing their consumption compared to the rest of the distribution, for whom the negative impact from the labor income channel would be at least partially offset by higher asset returns. These applications suggest that heterogeneous models can provide central banks with deeper insights, contributing to a more nuanced understanding of monetary policy transmission.

To better understand how monetary policy affects households differently, we introduce a novel decomposition of the effects of monetary policy on consumption across the wealth distribution—an analysis that is inherently impossible within the framework of representative agent models. Our decomposition breaks down the dynamic effect of monetary policy shocks into three key channels: (i) the disposable labor income channel, which captures the contractionary effects on wages and labor demand; (ii) the unexpected capital gains channel, which reflects declines in asset prices; and (iii) the expected asset returns channel, driven by persistently higher interest rates. We show that the total effect on consumption is most pronounced for households at the lower end of the wealth distribution. These households, possessing minimal assets, do not gain from increased asset returns and suffer from the contractionary effects of tighter monetary policy on output, labor demand, and real wages. Conversely, for wealthier households, the labor income channel is less significant while they benefit from higher asset returns that mitigate the initial adverse impact of capital gains. This decomposition sheds light on the mechanisms driving the heterogeneous effects of monetary policy and provides policymakers with a clearer understanding of how monetary shocks propagate through the economy.

Finally, we assess the forecasting accuracy of our HANK model relative to a comparable representative-agent (RANK) model. To isolate the role of heterogeneity, we carefully construct the RANK model by replacing the heterogeneous household block with a representative agent, while keeping all other features identical. Our results show that both models exhibit nearly identical forecasting performance for key macroeconomic variables over various horizons, from one quarter to two years. This finding suggests that introducing heterogeneity does not compromise the ability of the model to forecast aggregate variables, while adding valuable insights into the distributional effects of monetary policy.

These applications demonstrate the value of HANK models for central banks. By explicitly incorporating household heterogeneity, these models provide a more nuanced understanding of monetary policy transmission, capturing both aggregate dynamics and the distributional effects across wealth and income groups. As central banks increasingly consider inequality and distributional concerns in their policy frameworks, HANK models offer a powerful tool to complement traditional representative-agent approaches and enhance the policy-making process.

References

Acharya, Sushant, William Chen, Marco Del Negro, Keshav Dogra, Aidan Gleich, Shlok Goyal, Ethan Matlin, Donggyu Lee, Reca Sarfati, and Sikata Sengupta, "Estimating HANK for central banks," *FRB of New York Staff Report*, 2023, (1071).

Auclert, Adrien, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub, "Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models," *Econometrica*, 2021, *89* (5), 2375–2408.

_ , Matthew Rognlie, and Ludwig Straub, "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model," Technical Report, National Bureau of Economic Research 2020.

_ , _ , and _ , "The intertemporal keynesian cross," *Journal of Political Economy*, 2024, *132* (12), 4068–4121.

Bayer, Christian, Benjamin Born, and Ralph Luetticke, "Shocks, Frictions, and Inequality in US Business Cycles," *American Economic Review*, 2024, *114* (5), 1211–1247.

Calderhead, Ben, "A general construction for parallelizing Metropolis- Hastings algorithms," *Proceedings of the National Academy of Sciences*, 2014, *111* (49), 17408–17413.

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

Christiano, Lawrence J, Martin Eichenbaum, and Charles L Evans, "Nominal rigidities and the dynamic effects of a shock to monetary policy," *Journal of political Economy*, 2005, *113* (1), 1–45.

Coenen, Günter, Peter Karadi, Sebastian Schmidt, and Anders Warne, "The New Area-Wide Model II: an extended version of the ECB's micro-founded model for forecasting and policy analysis with a financial sector," Technical Report, ECB working paper 2018.

Dixit, Avinash K and Joseph E Stiglitz, "Monopolistic competition and optimum product diversity," *The American economic review*, 1977, *67* (3), 297–308.

European Central Bank, *Household finance and consumption survey - Results from the 2021 wave*, European Central Bank, 2023.

Fernandes, Adriano and Rodolfo Rigato, "Precautionary Savings and Stabilization Policy in a Present-Biased Economy," *Manuscript*, 2024.

Floden, Martin and Jesper Lindé, "Idiosyncratic risk in the United States and Sweden: Is there a role for government insurance?," *Review of Economic dynamics*, 2001, *4* (2), 406–437.

Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman, "The fiscal multiplier," Technical Report, National Bureau of Economic Research 2019.

Johnson, David S, Jonathan A Parker, and Nicholas S Souleles, "Household expenditure and the income tax rebates of 2001," *American Economic Review*, 2006, *96* (5), 1589–1610.

Kaplan, Greg and Giovanni L Violante, "A model of the consumption response to fiscal stimulus payments," *Econometrica*, 2014, *82* (4), 1199–1239.

_ , Benjamin Moll, and Giovanni L Violante, "Monetary policy according to HANK," *American Economic Review*, 2018, *108* (3), 697–743.

Kekre, Rohan, "Unemployment insurance in macroeconomic stabilization," *Review of Economic Studies*, 2023, *90* (5), 2439–2480.

Lee, Donggyu, "Unconventional Monetary Policies and Inequality," *Federal Reserve Bank of New York Staff Reports*, 2024, (1108).

Mankiw, N Gregory and Ricardo Reis, "Sticky information versus sticky prices: a proposal to replace the New Keynesian Phillips curve," *The Quarterly Journal of Economics*, 2002, *117* (4), 1295–1328.

Negro, Marco Del, Ibrahima Diagne, Keshav Dogra, Pranay Gundam, Donggyu Lee, and Brian Pacula, "Tradeoffs for the Poor, Divine Coincidence for the Rich," *FRB of New York Staff Report*, 2025, (1147).

Ramey, Valerie A, "Macroeconomic shocks and their propagation," *Handbook of macroeconomics*, 2016, *2*, 71–162.

Smets, Frank and Rafael Wouters, “Shocks and frictions in US business cycles: A Bayesian DSGE approach,” *American economic review*, 2007, *97* (3), 586–606.

Taylor, John B, “Discretion versus policy rules in practice,” in “Carnegie-Rochester conference series on public policy,” Vol. 39 Elsevier 1993, pp. 195–214.

A Additional Estimation Results

In this appendix A, we present a comprehensive set of additional estimation results, impulse response functions, and shock decompositions for the HANK model. Figure 8 illustrates the sequences of posterior draws for the estimated parameters, obtained using a Metropolis-Hastings sampler. Figure 9 displays the resulting posterior densities of these parameters.

In addition to the monetary shock already detailed in the main text, we provide impulse response functions for all other shocks. Specifically, the price markup shock is depicted in Figure 10, the wage markup shock in Figure 11, and the preference (or discount factor) shock in Figure 12. The investment shock is shown in Figure 13, the government spending shock in Figure 14, and, finally, the total factor productivity (TFP) shock is illustrated in Figure 15.

To further illustrate the HANK model, we also provide shock decompositions for the period from 2000Q1 to 2019Q4, the main estimation period of the model, for GDP growth in Figure 16, consumer price inflation in Figure 17, household consumption growth in Figure 18, and investment growth in Figure 19.

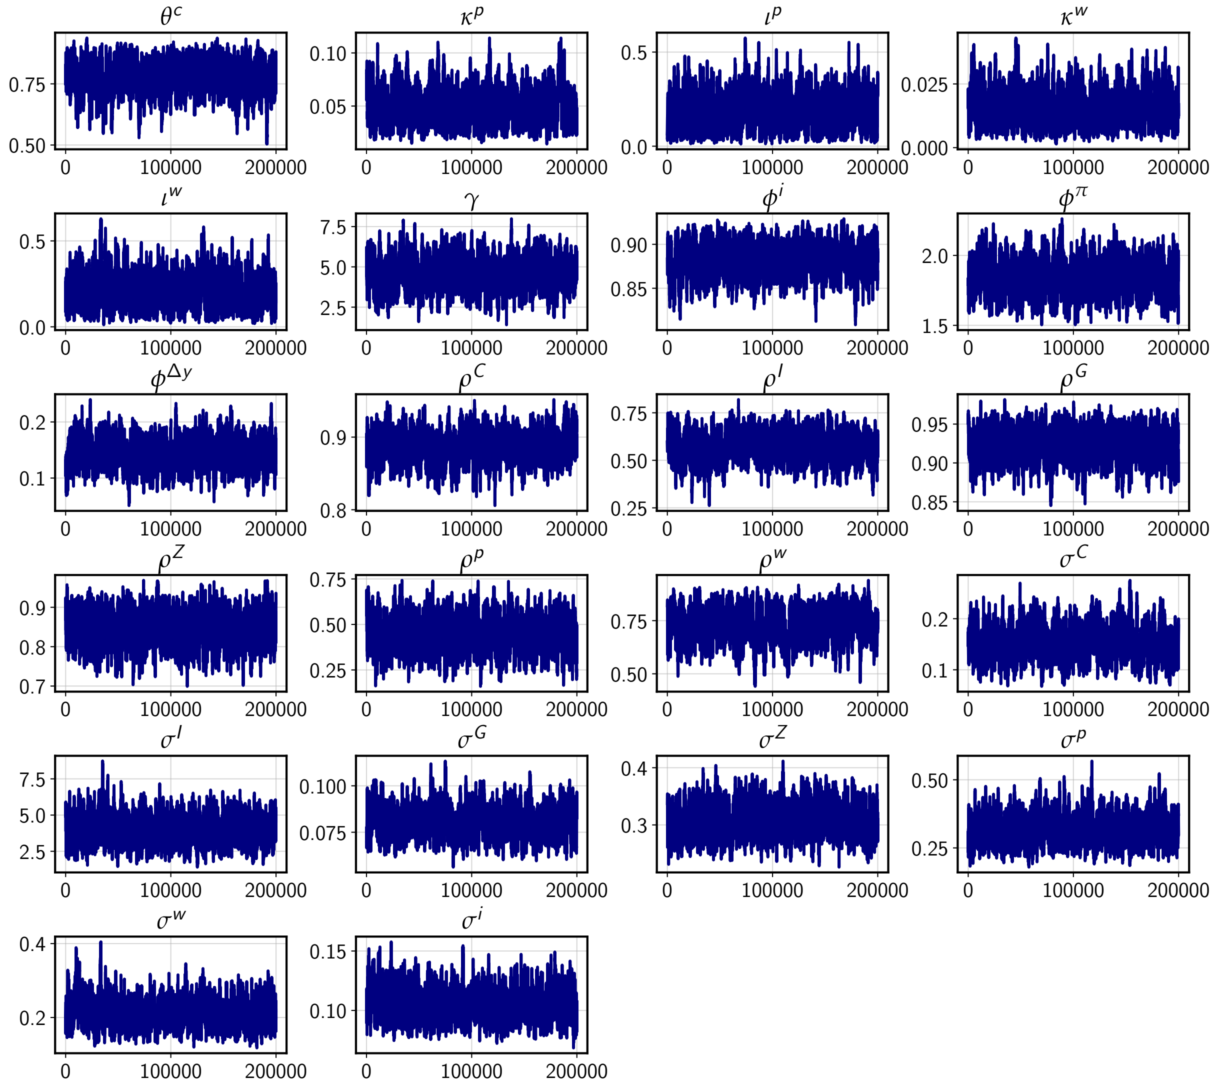

Figure 8: Simulated samples of estimated parameters

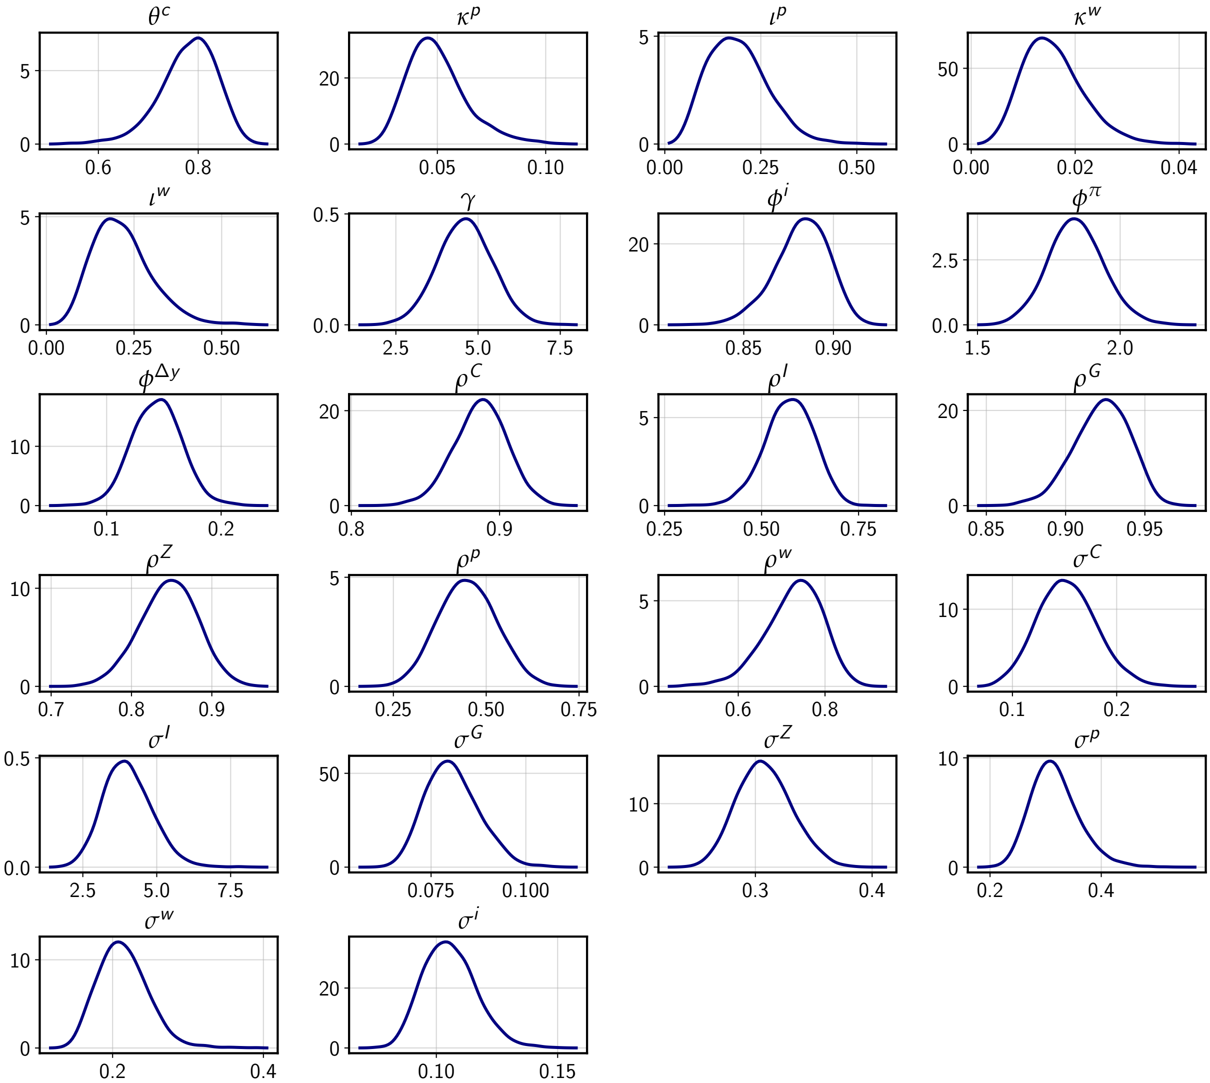

Figure 9: Posterior densities of estimated parameters

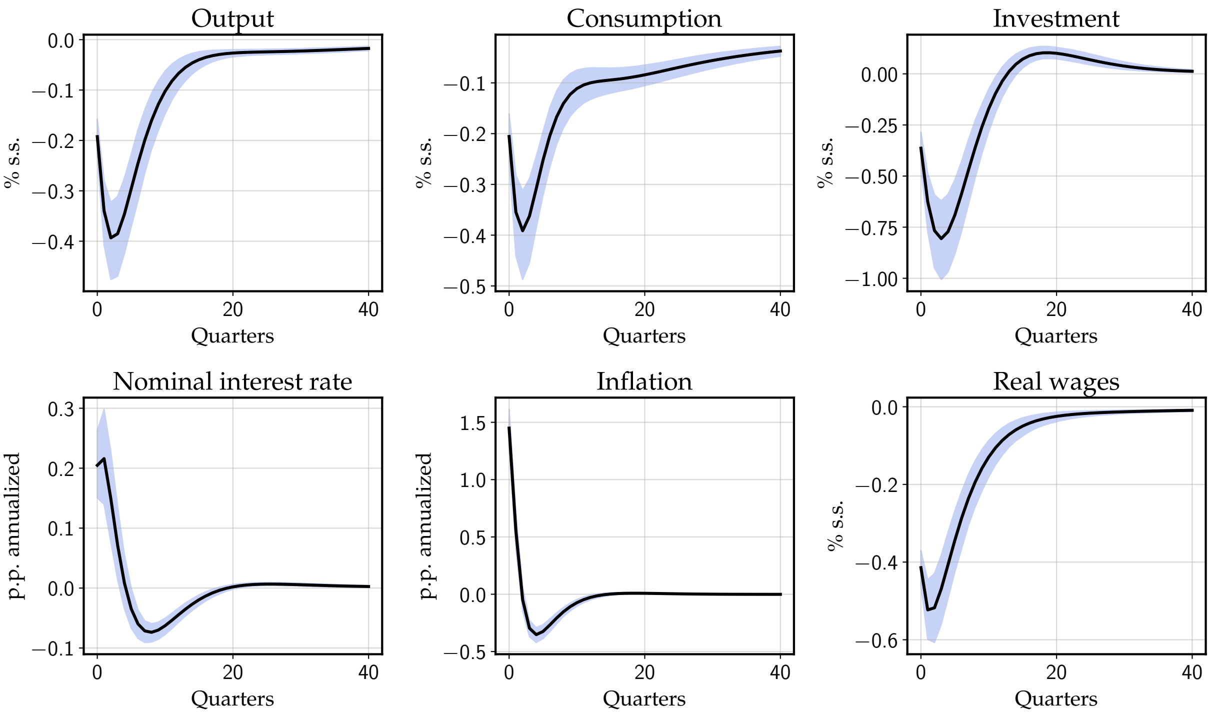

Figure 10: Impulse response functions to a price markup shock

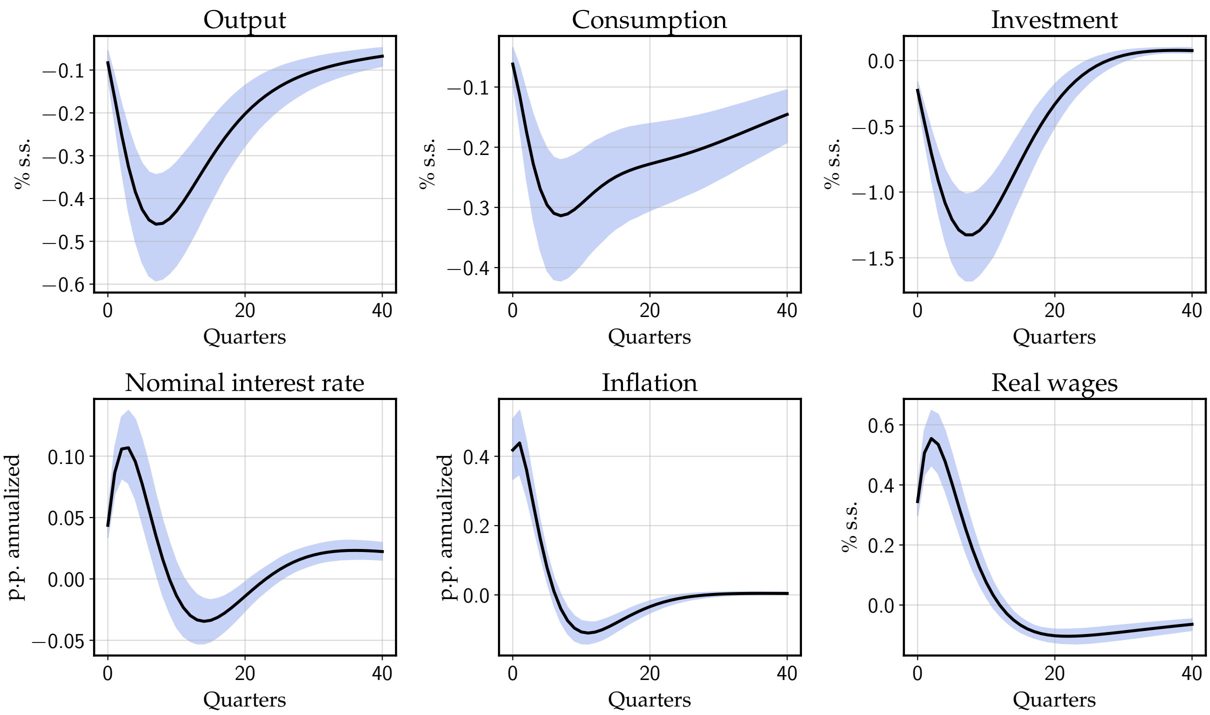

Figure 11: Impulse response functions to a wage markup shock

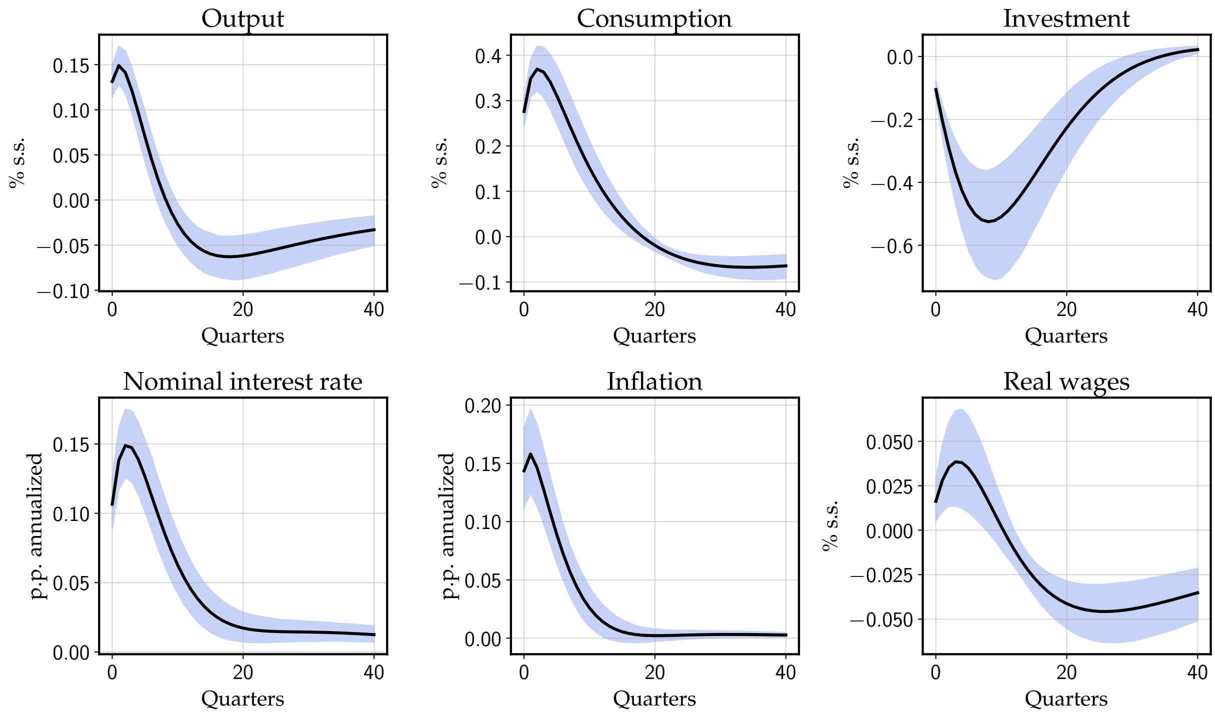

Figure 12: Impulse response functions to a consumer discount factor shock

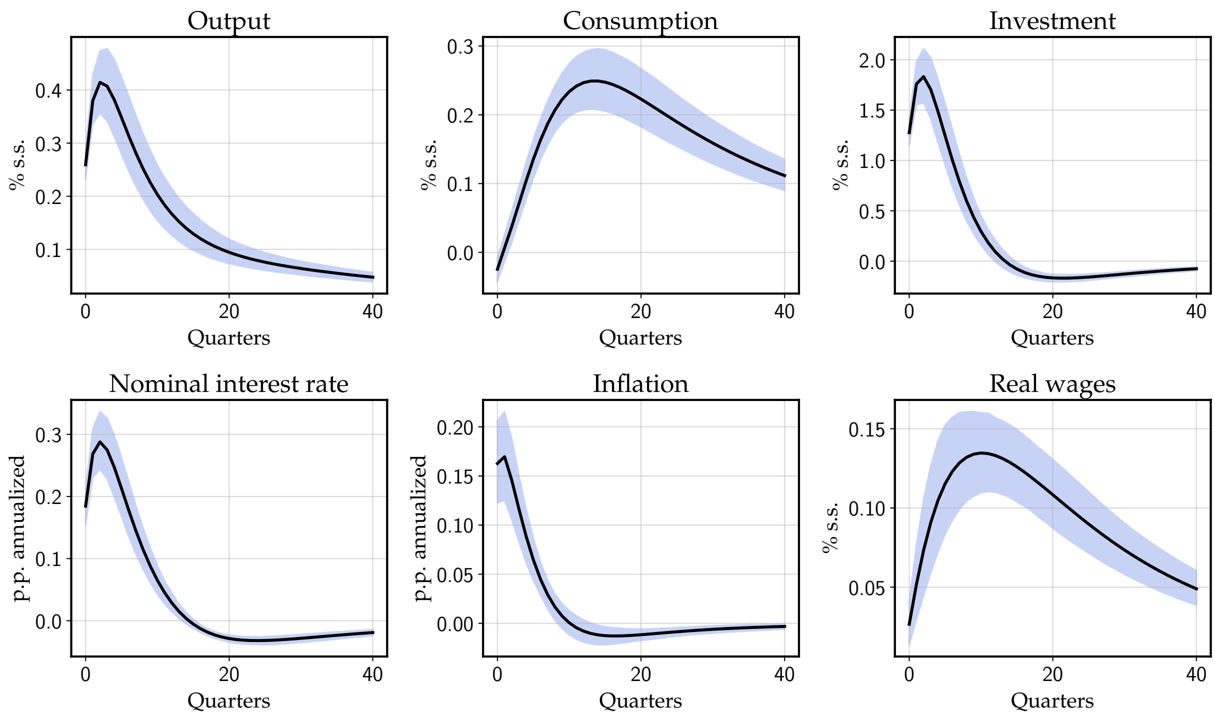

Figure 13: Impulse response functions to a investment shock

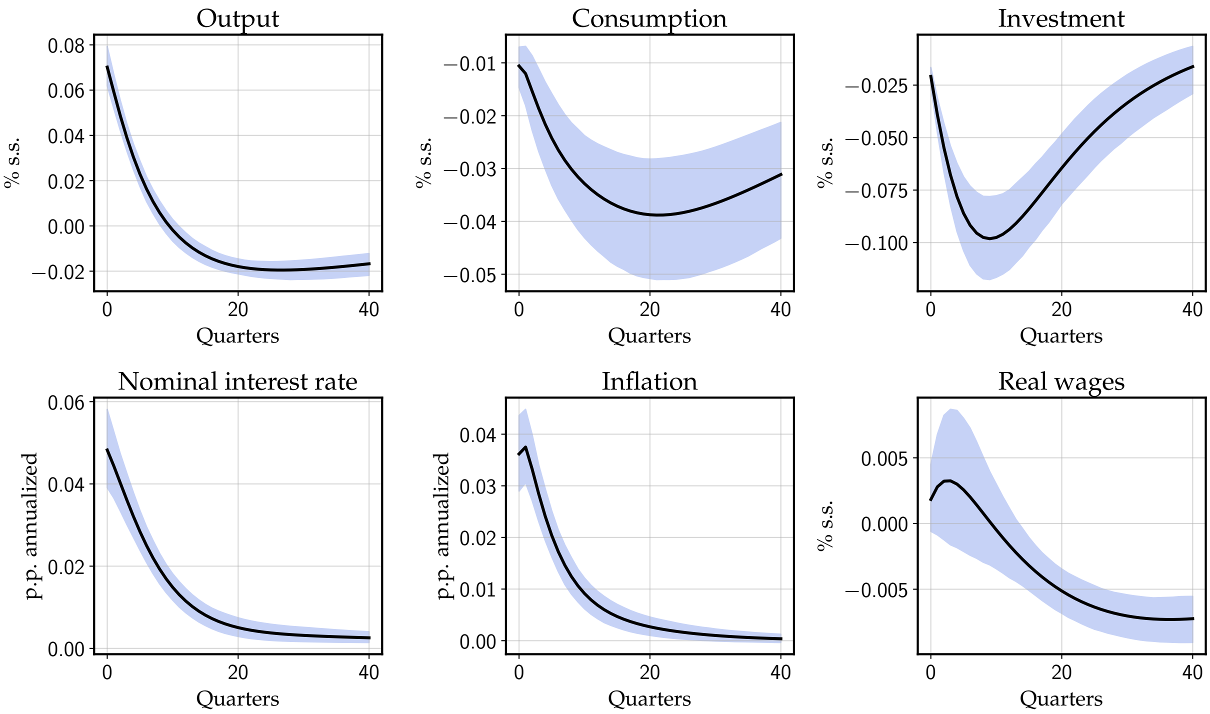

Figure 14: Impulse response functions to a government spending shock

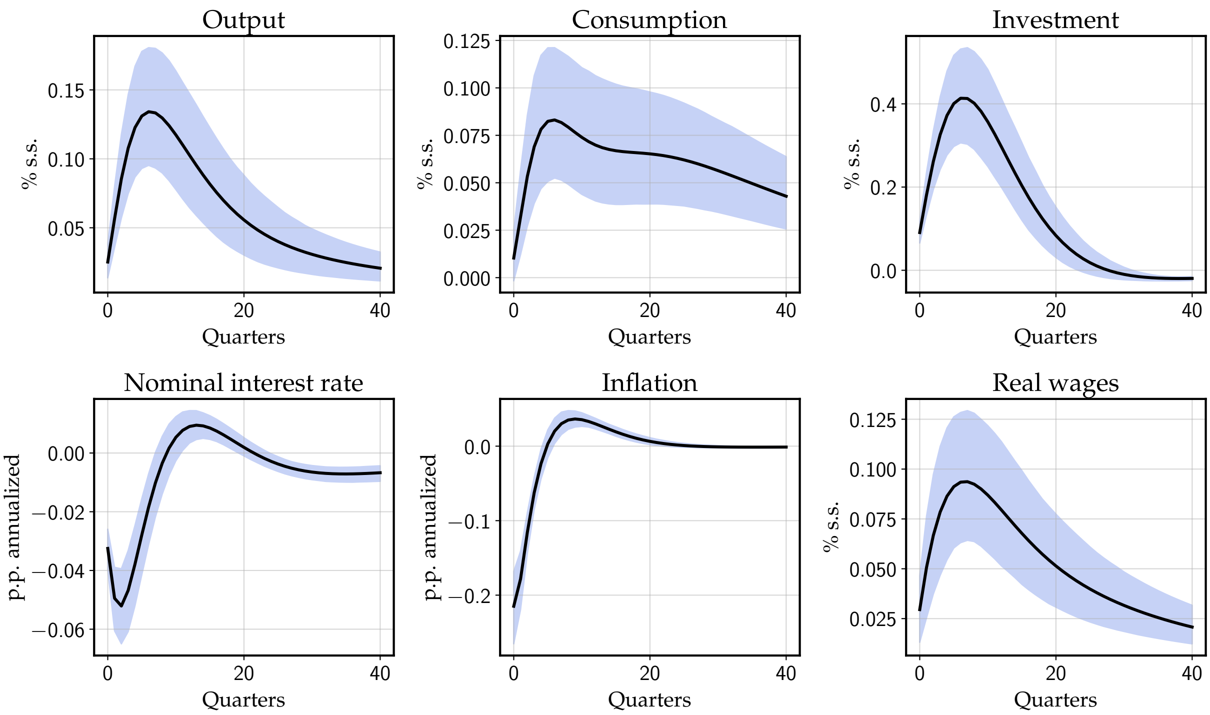

Figure 15: Impulse response functions to a TFP shock

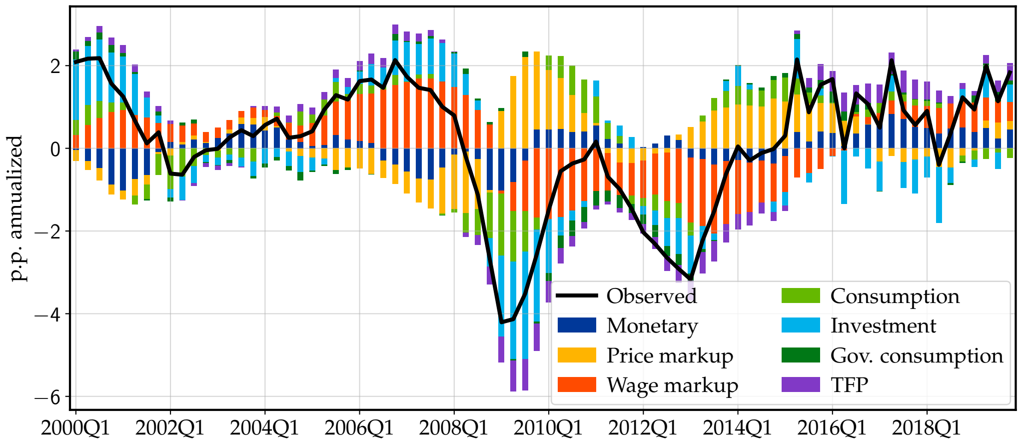

Figure 16: Shock decomposition of y-o-y GDP growth rate

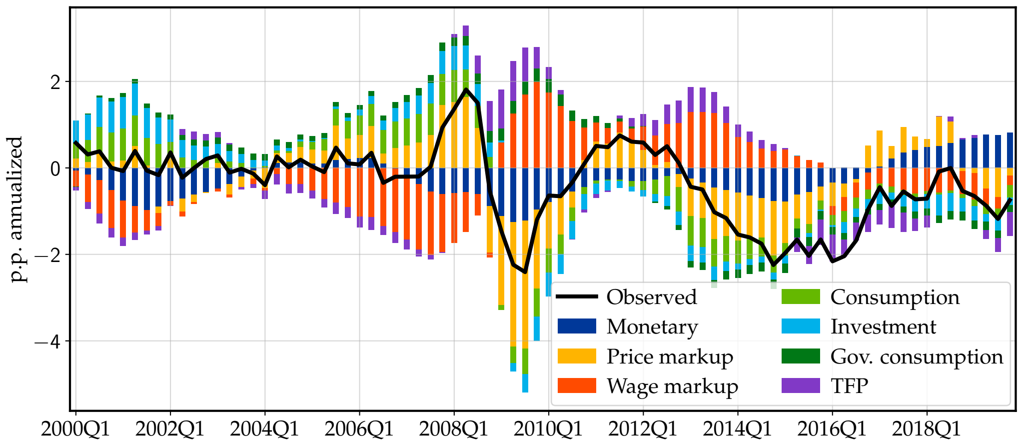

Figure 17: Shock decomposition of y-o-y consumer price inflation

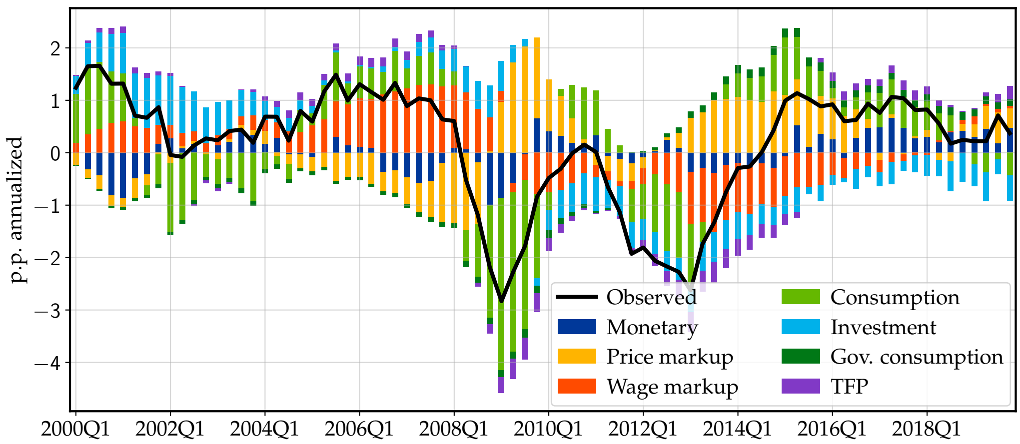

Figure 18: Shock decomposition of y-o-y consumption growth

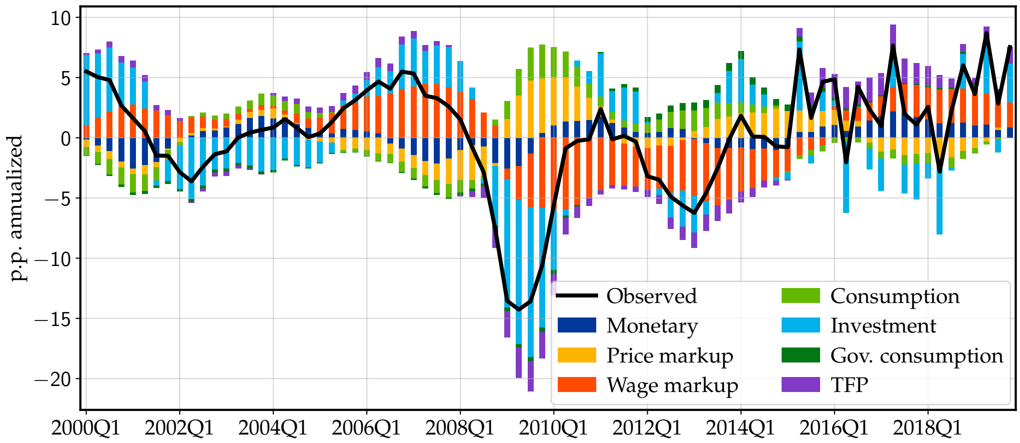

Figure 19: Shock decomposition of y-o-y investment growth

B Estimation of the Representative Agent Model

In this appendix B, we present supplementary findings related to the estimated RANK model. The estimation procedure and dataset employed closely mirror those used for the HANK model to ensure a fair comparison between the models. Figure20 illustrates the sequence of posterior draws obtained using the same Metropolis-Hastings sampler as applied to the HANK model. Figure 21 depicts the resulting posterior densities of the model parameters. Table 5 provides a detailed overview of the estimation outcomes for the RANK model, including the priors, posterior modes of the parameters, and the 10% and 90% credible intervals. In Subsection 5, we assess forecasting accuracy by estimating the models on a truncated sample from 2000Q1 to 2009Q4, using the remaining data until 2019Q4 for out-of-sample forecasting accuracy evaluation. The estimation results for both models based on the shorter sample are presented in Table 6.

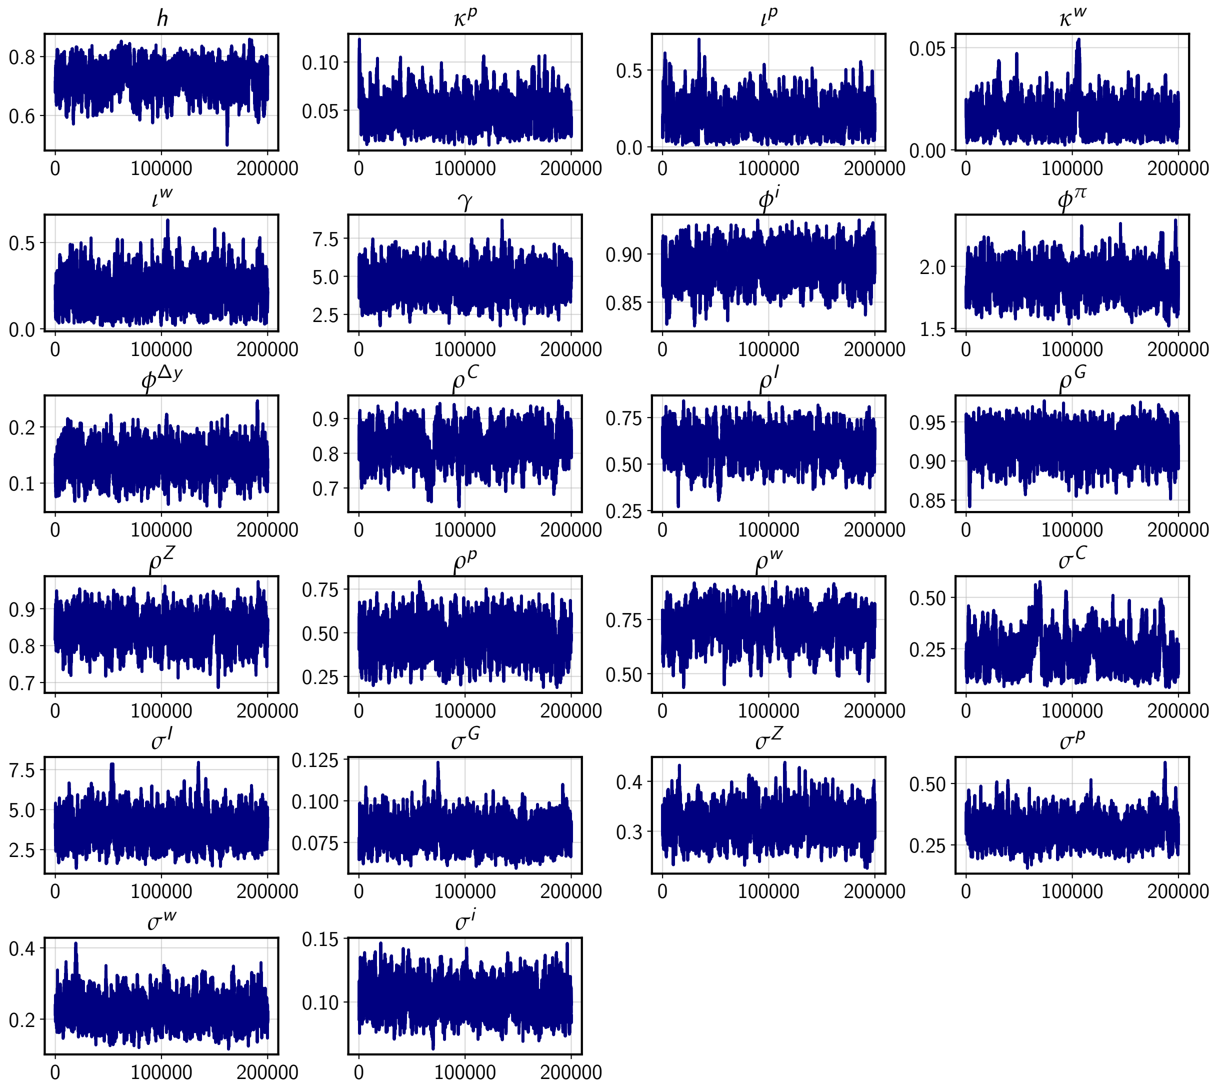

Figure 20: Simulated samples of estimated parameters of the RANK model

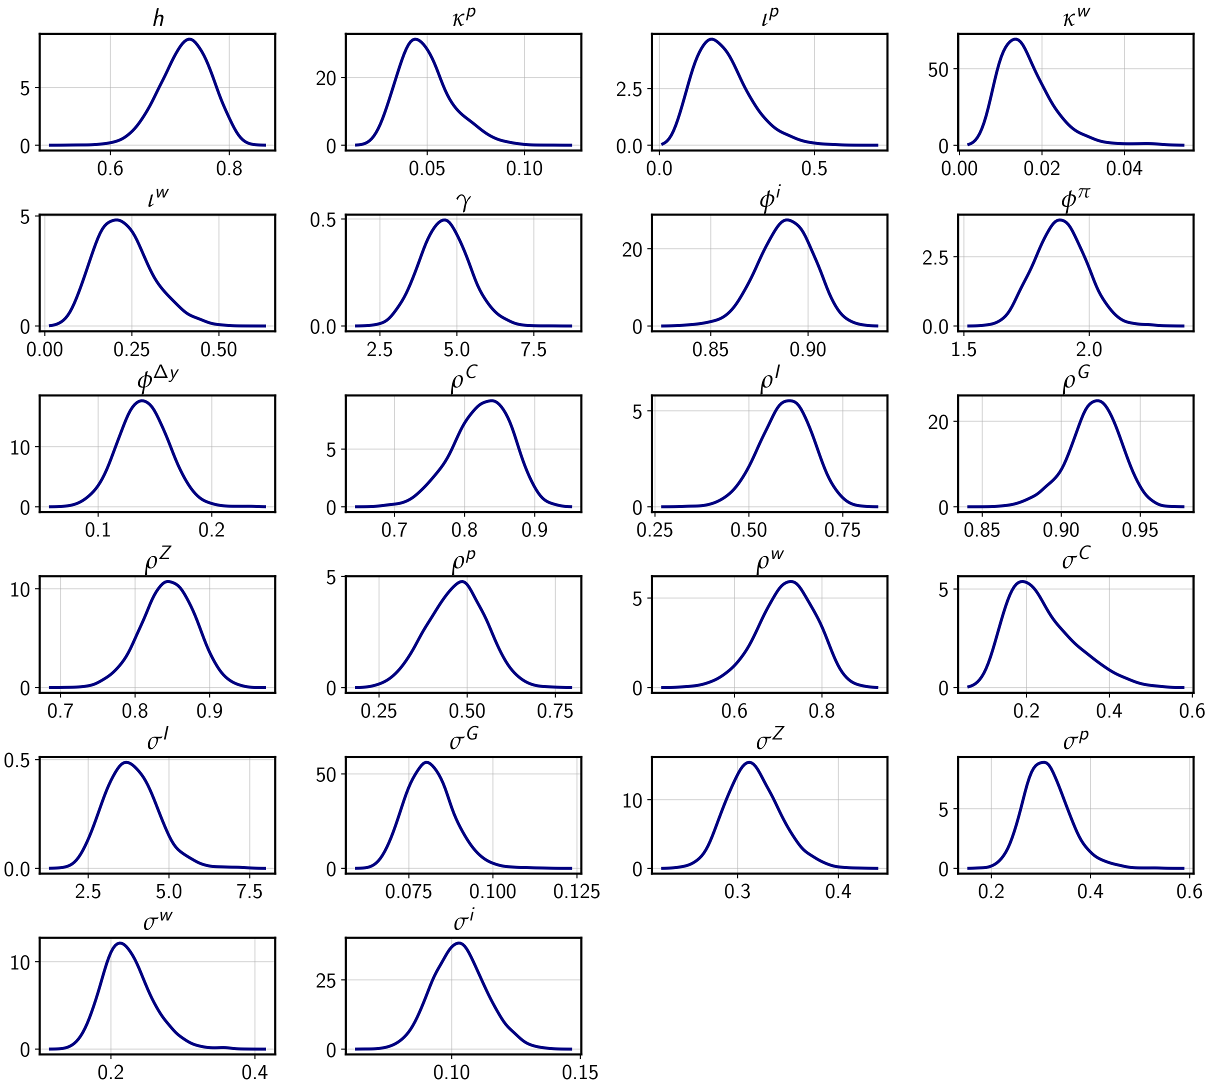

Figure 21: Posterior densities of estimated parameters of the RANK model

Parameter	Description	Prior			Posterior		
		Distribution	Mean	S.d.	Mode	10%	90%
Households							
h	Habit formation	Beta	0.7	0.05	0.71	0.67	0.78
Firms and unions							
κ^p	Price rigidity	Gamma	0.05	0.015	0.043	0.033	0.068
κ^w	Wage rigidity	Beta	0.05	0.015	0.014	0.009	0.025
ι^p	Price indexation	Beta	0.5	0.15	0.16	0.10	0.32
ι^w	Wage indexation	Beta	0.5	0.15	0.21	0.13	0.34
γ	Investment adj. cost	Normal	4	1	4.4	3.6	5.6
Monetary policy							
ϕ^i	Interest rate smoothing	Beta	0.7	0.05	0.89	0.87	0.91
ϕ^π	Taylor rule π coefficient	Normal	2	0.1	1.89	1.75	2.02
$\phi^{\Delta y}$	Taylor rule Δy coefficient	Normal	0.05	0.05	0.14	0.11	0.17
Shock autocorrelation							
ρ^C	Consumption shock	Beta	0.5	0.1	0.85	0.77	0.88
ρ^I	Investment shock	Beta	0.5	0.1	0.53	0.51	0.68
ρ^G	Gov. spending shock	Beta	0.5	0.1	0.92	0.90	0.94
ρ^Z	TFP shock	Beta	0.5	0.1	0.85	0.80	0.89
ρ^p	Price markup shock	Beta	0.5	0.1	0.48	0.36	0.57
ρ^w	Wage markup shock	Beta	0.5	0.1	0.73	0.63	0.80
Shock standard deviation							
σ^C	Consumption shock	Inv. Gamma	0.1	2	0.19	0.14	0.35
σ^I	Investment shock	Inv. Gamma	0.1	2	3.4	2.8	4.9
σ^G	Gov. spending shock	Inv. Gamma	0.1	2	0.080	0.072	0.091
σ^Z	TPF shock	Inv. Gamma	0.1	2	0.31	0.28	0.35
σ^p	Price markup shock	Inv. Gamma	0.1	2	0.29	0.26	0.37
σ^w	Wage markup shock	Inv. Gamma	0.1	2	0.20	0.18	0.27
σ^i	Monetary shock	Inv. Gamma	0.1	2	0.099	0.090	0.117

Table 5: Estimated parameters of the RANK model

Parameter	Description	Het. agents	Rep. agent
Households			
θ^c	Information stickiness	0.66	–
h	Habit formation	–	0.70
Firms and unions			
κ^p	Price rigidity	0.056	0.051
κ^w	Wage rigidity	0.021	0.023
ι^p	Price indexation	0.21	0.21
ι^w	Wage indexation	0.20	0.22
γ	Investment adj. cost	3.9	4.0
Monetary policy			
ϕ^i	Interest rate smoothing	0.84	0.85
ϕ^π	Taylor rule π coefficient	1.89	1.90
$\phi^{\Delta y}$	Taylor rule Δy coefficient	0.12	0.11
Shock autocorrelation			
ρ^C	Consumption shock	0.81	0.69
ρ^I	Investment shock	0.58	0.61
ρ^G	Gov. spending shock	0.66	0.66
ρ^Z	TFP shock	0.82	0.81
ρ^p	Price markup shock	0.38	0.40
ρ^w	Wage markup shock	0.76	0.75
Shock standard deviation			
σ^C	Consumption shock	0.18	0.36
σ^I	Investment shock	2.6	2.5
σ^G	Gov. spending shock	0.076	0.076
σ^Z	TPF shock	0.30	0.31
σ^p	Price markup shock	0.36	0.34
σ^w	Wage markup shock	0.22	0.23
σ^i	Monetary shock	0.14	0.13

Table 6: Posterior modes for the 2000Q1–2009Q4 sample

Acknowledgements

We would like to thank Marta Banbura, Nikola Bokan, Matteo Ciccarelli, Günter Coenen, and Antoine Kornprobst, as well as the anonymous referee, for their thoughtful comments, corrections, and suggestions. We also thank other colleagues from the Forecasting and Policy Modelling Division, along with various seminar participants for their valuable feedback.
The views expressed here are solely those of the authors and do not necessarily reflect the views of the European Central Bank.

Hanno Kase
European Central Bank, Frankfurt am Main, Germany; email: hanno.kase@ecb.europa.eu

Rodolfo Dinis Rigato
European Central Bank, Frankfurt am Main, Germany; email: rodolfo.dinis_rigato@ecb.europa.eu

Postal address 60640 Frankfurt am Main, Germany
Telephone +49 69 1344 0
Website www.ecb.europa.eu

PDF ISBN 978-92-899- 7404-2 ISSN 1725-2806 doi:10.2866/4558236 QB-01-25-179-EN-N