Which idiosyncratic risk matters for business cycles?

Antoine Frech*

November 22, 2025

Abstract

I estimate a medium-scale Heterogeneous-Agent New Keynesian model where the effects of household heterogeneity are encapsulated in a wedge within the Euler equation. Estimating the model enables the computation of this wedge over the past decades, facilitating an investigation into how individual risk and precautionary savings evolve throughout the business cycle. By comparing the estimated wedge with empirical measures of individual risk, I demonstrate that the model captures realistic dynamics of individual risk and precautionary savings. Notably, the estimated wedge shows a strong correlation with unemployment during normal times. Furthermore, incorporating this wedge into an otherwise standard macroeconomic model significantly enhances the fit of macroeconomic time series, underscoring the importance of idiosyncratic risk in shaping business cycles.

Keywords: New Keynesian, Business Cycles, Bayesian Estimation, Incomplete Markets, Idiosyncratic Risk, Precautionary Savings.
JEL codes: D52, E20, E32, E52.

*Université Paris-Saclay, Centre for Economics at Paris-Saclay. Email: antoine.frech@universite-paris-saclay.fr
I am grateful to my advisors, Edouard Challe and Michel Guillard, for their guidance and support. I also thank Sushant Acharya, Adrien Auclert, Hubert Kempf, and Xavier Ragot for their helpful comments and discussions.

1 Introduction

In this paper, I estimate the contribution of cyclical idiosyncratic risk and associated precautionary savings to aggregate fluctuations. My framework of analysis is a medium-scale Heterogeneous-Agent New Keynesian (HANK) model with rich cross-sectional heterogeneity, based on Acharya and Dogra (2020) and Acharya et al. (2023), in which the contribution of precautionary savings to aggregate demand is summarized by a wedge in the aggregate Euler equation. This feature allows me to isolate the role of time-varying precautionary savings for aggregate dynamics and to compare how alternative popular measures of idiosyncratic risk fare in terms of model fit.

My estimation strategy differs from the recent contributions of Auclert et al. (2020) and Bayer et al. (2024), who also estimate medium-scale HANK models. In conventional HANK frameworks such as theirs, uninsured idiosyncratic risk generates not only a precautionary saving motive but also high-MPC agents. As a result, redistribution, profit allocation, and income cyclicality jointly interact with precautionary savings in shaping consumption dynamics. By contrast, the distinctive feature of my approach is that it isolates and quantifies the cyclical strength of the precautionary savings channel. To this end, I use a Representative-Agent New Keynesian (RANK) model—where the precautionary motive is shut down—as a benchmark. Introducing cyclical precautionary savings into this framework substantially improves the model's empirical fit relative to the benchmark. Furthermore, the estimated cyclical variation in precautionary savings closely tracks standard measures of idiosyncratic risk, such as unemployment risk.

The primary contribution of this paper is to quantify the effects of idiosyncratic risk on precautionary savings over recent decades using the CARA-Normal framework and to compare the results with empirical measures of individual risk. Developed by Acharya and Dogra (2020) and Acharya et al. (2023), this framework rests on two key assumptions: idiosyncratic income shocks are normally distributed, and households exhibit Constant Absolute Risk Aversion (CARA) preferences. These assumptions allow for a closed-form solution of the aggregate Euler equation—a feature not attainable in conventional HANK models such as those in Auclert et al. (2020) and Bayer et al. (2024). The CARA-Normal framework is therefore particularly well suited to isolating and clarifying the link between idiosyncratic risk and precautionary savings over the business cycle. This tractability enables a clearer and more transparent account of the dynamics of precautionary savings and their aggregate implications.

My approach also differs from that of Bilbiie et al. (2023), who aim to quantify the relative importance of the different transmission channels in HANK models. In their framework, precautionary savings are closely tied to the presence of high-MPC agents, since idiosyncratic risk is modeled as the risk of becoming financially constrained. In this setting, precautionary savings are governed by the marginal utility gap between savers and high-MPC agents and therefore depend critically on redistribution policies and on the cyclicality of high-MPC agents' income. As a result, precautionary savings cannot be separated from the behavior of high-MPC agents.

By contrast, the challenge I address is to measure the effects of idiosyncratic risk on precautionary savings while remaining agnostic about the underlying source of risk. My framework isolates the impact of countercyclical risk without committing to a specific mechanism generating it. The aggregate effects are summarized by a single statistic that does not depend on whether the underlying risk stems from unemployment or other forces. Embedding this statistic in a macroeconomic model estimated on aggregate time series then allows me to trace its evolution over past decades and to directly compare it with empirical measures of individual risk.

This approach delivers three main insights. First, it allows for an assessment of whether HANK models can reproduce empirically observed dynamics of individual risk without directly targeting these measures in estimation. Second, by comparing the estimated dynamics of idiosyncratic risk with empirical evidence, it provides a clear answer to the central question of the paper: the dimension of idiosyncratic risk that matters for business cycles is one that is strongly countercyclical, with unemployment serving as a good proxy. Third, I show that introducing a single wedge into an otherwise standard macroeconomic model markedly improves its fit to macroeconomic time series relative to a benchmark without idiosyncratic risk.

The estimated model also shows that the aggregate consequences of precautionary savings depend critically on the type of shock affecting the economy. In particular, the interaction between investment dynamics and precautionary savings plays a central role in shaping how idiosyncratic risk influences output responses, especially following supply shocks. While Acharya and Dogra (2020) show that countercyclical risk amplifies aggregate fluctuations and procyclical risk dampens them, I find the opposite relationship for supply shocks once capital and investment are incorporated: countercyclical risk dampens output fluctuations, whereas procyclical risk amplifies them. This reversal arises because, under countercyclical risk, the dampening of investment responses induced by precautionary savings outweighs the amplification of consumption responses, leading to overall muted output fluctuations. The opposite mechanism operates under procyclical risk, where investment amplification dominates.

Alongside capital and investment, the model incorporates other features of medium-scale DSGE models (Christiano et al., 2005; Smets and Wouters, 2007), such as capital utilization, investment adjustment costs, and nominal rigidities in both prices and wages. The model is estimated over seven key U.S. time series using Bayesian techniques. By employing the Kalman smoother, I extract the dynamics of the wedge over past decades. To ensure comparability, I use the same dataset as Bayer et al. (2024), allowing a direct performance comparison of my model against theirs in fitting macroeconomic time series. Despite the simplified HANK framework with CARA utility, the model performs comparably to the Constant Relative Risk Aversion (CRRA) benchmark up to a first-order approximation, as discussed extensively by Acharya and Dogra (2020) and Acharya et al. (2023). Notably, my model achieves a higher marginal data density than the one developed by Bayer et al. (2024), indicating that the CARA-Normal framework is not less effective in reproducing macroeconomic data.

Related literature

This paper contributes to the recent literature examining the ability of HANK models to replicate macroeconomic time series while simultaneously aligning with cross-sectional data on inequality and time series capturing the dynamics of MPCs and individual risk.

Using a medium-scale HANK model, Auclert et al. (2020) assess its ability to fit business cycle data. They first match empirical impulse response functions following Christiano et al. (2005) and then estimate their model on key U.S. time series using Bayesian techniques, as in Smets and Wouters (2007). The authors find that while their HANK model performs well, additional assumptions—such as consumption habits or sticky expectations—are required to replicate the hump-shaped response of consumption following a monetary policy shock.

Bayer et al. (2024) also estimate a HANK model using Bayesian techniques, demonstrating that their model can simultaneously fit macroeconomic time series and cross-sectional data on economic inequalities and wealth distributions. In an extension, they incorporate an empirical measure of income idiosyncratic risk as an observable to discipline individual risk within their model. Since I use their dataset, I compare the results of my model with their measure of income idiosyncratic risk.

Challe et al. (2017) highlight precautionary savings as a key driver of business cycles within a HANK&SaM framework,[1] which attributes individual risk to the likelihood of becoming or remaining unemployed. In these models, the precautionary saving motive arises from the probabilities of both job loss and prolonged unemployment. Their framework is sufficiently tractable to allow for a finite state-space representation, facilitating straightforward estimation via a maximum likelihood estimator. They find that precautionary savings play a pivotal role during recessions.

In related work, Cho (2023) estimate a HANK&SaM model in a quantitative setting, emphasizing the dual importance of precautionary savings against unemployment risk and MPC heterogeneity in explaining aggregate fluctuations. However, the author identifies a "MPC puzzle," which emerges from the tension between the volatility of consumption responses driven by high-MPC agents and the inherent persistence in macroeconomic data.

Bilbiie et al. (2023) estimate a medium-scale version of the tractable HANK model introduced in Bilbiie (2025), which allows them to disentangle the effects of various HANK features and assess their relative impact on business cycles. In this framework, there are two types of agents: savers and hand-to-mouth consumers, with individual risk modeled as the probability of switching between these types. As a result, precautionary savings are inherently connected to the presence of high-MPC agents. By aggregating the consumption of both groups, the authors derive an aggregate Euler equation that captures the effects of MPC heterogeneity and precautionary savings in response to the risk of becoming hand-to-mouth. Through counterfactual analysis, the authors

[1]Such models incorporate elements from the search-and-matching (SaM) literature into a HANK framework, resulting in HANK&SaM models where individual risk is directly tied to employment status (Ravn and Sterk, 2017, 2021; Den Haan et al., 2018).

demonstrate that consumption inequalities amplify aggregate fluctuations. They identify long-run inequalities and cyclical precautionary saving behavior as the primary drivers of this amplification.

In contrast to the tractable approaches of Challe et al. (2017) and Bilbiie et al. (2023), the CARA-Normal framework minimizes assumptions about the sources of individual risk. It assumes only that households experience normally distributed idiosyncratic income shocks, without specifying the origin of this risk, thereby allowing for a precautionary saving motive independent of high-MPC agents' consumption. This agnostic approach enables a fair comparison between the estimated wedge and empirical measures of individual risk, providing a more general answer to the central question posed in the present paper's title.

This paper is also closely related to Berger et al. (2023), who conduct a "wedge accounting" exercise in a more general HANK model. The authors exploit the fact that the effects of heterogeneity in their model can be summarized by two wedges: one in the Euler equation and another in the Phillips curve. These wedges are functions of households' consumption shares and relative wages. Using microeconomic data, they estimate these wedges, allowing them to evaluate the effects of heterogeneity on consumption and output. My approach complements that of Berger et al. (2023): while they use microeconomic data to discipline the wedges and infer the effects of heterogeneity on macroeconomic outcomes, I estimate the model on macroeconomic data to compute the wedge and compare the results with empirical measures of individual risk.

Layout

The remainder of the paper is organized as follows: Section 2 details the construction of the model. Section 3 describes the data and estimation procedure. Section 4 presents and discusses the estimation results. Section 5 evaluates the most suitable empirical measure of idiosyncratic risk according to the model. Section 6 explores the effects of idiosyncratic risk on aggregate fluctuations. Section 7 compares the baseline results with alternative specifications. Finally, Section 8 concludes.

2 Model

This section introduces a medium-scale DSGE model featuring incomplete markets with both idiosyncratic and aggregate risk. Unlike quantitative HANK models, the goal here is to maintain simplicity to isolate the interplay between individual risk and precautionary savings. Following Acharya and Dogra (2020), two key assumptions underpin this HANK framework: idiosyncratic shocks are normally distributed, and households' utility function is a CARA function. To ensure comparability with Bayer et al. (2024), this model incorporates the main features they considered while omitting those they excluded.

The model includes government spendings and bonds, capital accumulation, investment adjustment costs, capital utilization, sticky prices, and sticky nominal wages. However, as in Bayer et al. (2024), it excludes consumption habits[2] and price or wage indexation. Hand-to-mouth (HtM) agents are also excluded, as the focus is on the relationship between idiosyncratic risk and precautionary savings, and these agents do not save by assumption.[3] Unlike Bayer et al. (2024), this model is log-linearized, aligning it more closely with the approaches of Christiano et al. (2005) and Smets and Wouters (2007).

2.1 Households

2.1.1 Preferences, idiosyncratic risk and budget constraint

To introduce idiosyncratic risk and precautionary savings, the model adopts a Bewley-Huggett framework where households face uninsurable idiosyncratic risk due to incomplete financial markets. Unlike the standard DSGE literature, this model assumes a CARA utility function following Acharya and Dogra (2020). This assumption enables aggregation of individual consumption and derivation of a closed-form solution for the aggregate MPC. The utility function for household i is:

$$\mathbb{U}_t(i) = -\frac{1}{\gamma}e^{-\gamma C_t(i)} - \rho e^{\frac{1}{\rho}\left(L_t(i)-\bar{\xi}\right)}$$

where γ is the coefficient of absolute risk aversion, ρ governs labor disutility and $\bar{\xi}$ is a scaling parameter that will be used to pin down the labor supply at steady state.

Household i chooses consumption $C_t(i)$ and investment $A_t(i)$ in actuarial bonds. Households delegate wage[4] negotiations to unions, which optimize the aggregate labor supply L_t (more details in Appendix A.2.3). Households are the sole owners of the firms. Thus, a household i receives dividends $D_t(i)$ from the intermediate firms and $D_t^K(i)$ from the capital firms. As in Acharya and Dogra (2020), individual income is denoted $\mathcal{Y}_t(i)$ and is made of labor income, dividends and taxes. This part of household i's income is subject to idiosyncratic risk. The second key assumption is that $\mathcal{Y}_t(i)$ is normally distributed, so that:

$$\mathcal{Y}_t(i) \sim \mathcal{N}(\mathcal{Y}_t, \sigma_t^2)$$

where $\mathcal{Y}_t$ is made of aggregate labor income, dividends and taxes (i.e., $\mathcal{Y}_t = \mathcal{W}_t L_t + D_t + D_t^K - T_t$). σ_t^2 is the variance of idiosyncratic income risk, which may vary with economic activity. The

[2]While Christiano et al. (2005) and Auclert et al. (2020) show that sluggish consumption adjustments are necessary to match impulse response functions (IRFs) following a monetary policy shock, Bayer et al. (2024) demonstrate that their model performs well without consumption habits when estimated using Bayesian techniques. Excluding consumption habits avoids irrelevant distortions in intertemporal MPCs.

[3]I discuss how high-MPC agents would affect the results in Subsection 7.3.

[4]I denote W_t the nominal wage and P_t the price level. The real wage is defined as: $\mathcal{W}_t \equiv \frac{W_t}{P_t}$.

household's utility maximization problem is:

$$\max_{C_t(i),A_t(i)} \mathbb{E}_t \sum_{k=0}^{\infty} \beta^k \mathbb{U}_{t+k}(i)$$
$$s.t. \quad C_t(i) + A_t(i) = \mathcal{Q}_{t-1} A_{t-1}(i) + \mathcal{Y}_t(i) \tag{1}$$

where β is the discount factor and $\mathcal{Q}_t$ is the real return on actuarial bonds.

2.1.2 Consumption

Household i maximizes utility by choosing $C_t(i)$, considering her budget constraint and the ability to smooth consumption through borrowing or lending. Let $x_t(i)$ denote the demeaned cash-on-hand of household i. Differentiating the maximization program (1) with respect to $A_t(i)$ and $C_t(i)$ and combining the First Order Conditions (FOCs), one obtains the following Euler equation:

$$e^{-\gamma C_t(i)} = \beta \mathcal{Q}_t \mathbb{E}_t \int \left(\sqrt{2\pi}\sigma_{t+1}\right)^{-1} e^{-\gamma C_{t+1}(i) - \frac{1}{2}\left(\frac{\mathcal{Y}_{t+1}(i) - \mathcal{Y}_t}{\sigma_{t+1}}\right)^2} dx_{t+1}(i) \tag{2}$$

This formulation accounts for aggregate risk, unlike Acharya and Dogra (2020), who assume no aggregate shocks.[5] In Appendix A.2.5, I demonstrate that the aggregate Euler equation remains unchanged under first-order approximation around the steady state.[6]

As it is common in the HANK literature, idiosyncratic risk is assumed to vary with economic activity. Following Acharya et al. (2023), σ_t^2 is correlated with the deviation of output from its steady state, such that:

$$\sigma_t^2 = \tilde{\sigma}^2 e^{2\varphi(Y_t - Y)}$$

where $\tilde{\sigma}^2$ is the steady state variance of $\mathcal{Y}_t(i)$ and φ determines the cyclicality of idiosyncratic risk ($\varphi < 0$ implies countercyclical risk, $\varphi > 0$ implies procyclical risk, and $\varphi = 0$ implies acyclicality). We recover the RA paradigm when there is no idiosyncratic risk at all (i.e., $\tilde{\sigma} = 0$). In order to ease the interpretation of the parameter $\tilde{\sigma}$, I assume that it is proportional to steady state output, such that: $\tilde{\sigma} = \sigma Y$. This way, the standard deviation of individual incomes at steady state can be interpreted in units of aggregate income.

2.1.3 Financial intermediaries

The financial sector is similar to the one in Acharya and Dogra (2020) and Acharya et al. (2023). Financial intermediaries operate in perfectly competitive markets, earning no profits. They pur-

[5] Acharya and Dogra (2020) assume that there is no aggregate risk, leading to a simpler formulation of the Euler equation. As my economy features small aggregate shocks, I have to take into account the effects of aggregate risk in the Euler equation.

[6] This approach aligns with other HANK model estimation studies, such as Challe et al. (2017), Auclert et al. (2020), Bilbiie et al. (2023), Cho (2023), and Bayer et al. (2024).

chase government debt B_t and receive interest payments. Then, they sell actuarial bonds A_t to households with a real return $\mathcal{Q}_t$. The decision of financial intermediaries can be represented by the following maximization program, featuring their actualized constraint:

$$\max_{A_t, B_t} B_t - A_t \quad s.t. \quad \frac{B_t}{\tilde{R}_t} - \frac{A_t}{\mathcal{Q}_t} \leq 0$$

where $\tilde{R}_t$ is the real return on government debt which is equal to the gross real interest rate R_t, given by the Fisher equation, multiplied by a stochastic process generating a wedge between the interest rate controlled by the central bank and the return on government debt.[7] Hence, the real return on government debt is such that:

$$\tilde{R}_t = V_t R_t$$

where V_t is the stochastic process and is assumed to follow the following AR(1) process:

$$\log V_t \equiv v_t = \varrho_v v_{t-1} + \sigma_v \epsilon_t^v$$

where ϱ_v is the persistence parameter and σ_v is the standard deviation of the shock.

The zero-profit condition implies that $A_t = B_t$, yielding the following solution for the maximization program of financial intermediaries:

$$\mathcal{Q}_t = \tilde{R}_t$$

2.2 Firm sector

2.2.1 Final goods and intermediate goods producers

This component of the model adheres to standard practices in the DSGE literature. The final goods sector operates under perfect competition, with households' preferences for varieties represented by a CES aggregator (details are provided in Appendix A.2.1). The elasticity of substitution between varieties is denoted by θ. The intermediate goods sector features monopolistic competition à la Dixit and Stiglitz (1977). Firms in this sector utilize a Cobb-Douglas production function with constant returns to scale, where α represents the capital share. This production function is subject to total factor productivity (TFP) shocks, which follow an AR(1) process. The persistence of these shocks is governed by the parameter ϱ_z, and their standard deviation is denoted by σ_z. Prices are assumed to be sticky in the manner of Calvo (1983), where a fraction ζ_p of firms cannot adjust their prices optimally at a given time t (additional details are in Appendix A.2.2). This price rigidity

[7] This allows to introduce a risk premium shock as in Smets and Wouters (2007) and Bayer et al. (2024). This shock produces aggregate effects similar to those of a "preference shock" (i.e., a shock to households' discount factor β).

results in the emergence of a standard New Keynesian Phillips Curve (NKPC). Furthermore, the firms' mark-up is assumed to follow an ARMA(1,1) process, with ϱ_p as the persistence parameter, σ_p as the standard deviation of the shocks, and ϑ_p as the weight of the moving average (MA) component.

2.2.2 Labor market

This aspect of the model is also standard. The labor market operates under monopolistic competition. Households' labor supplies are assumed to be differentiated and imperfectly substitutable, granting them market power to negotiate higher nominal wages. It is further assumed that households delegate wage negotiations to unions. These unions take into account the preferences of a hypothetical representative agent who consumes aggregate consumption, as in Auclert et al. (2021). Nominal wages are assumed to be sticky à la Erceg et al. (2000), meaning that a fraction ζ_w of unions cannot adjust wages optimally at a given time t (additional details are provided in Appendix A.2.3). This wage rigidity results in the derivation of a Wage NKPC. Additionally, the labor market mark-up is assumed to follow an ARMA(1,1) process, with ϱ_w as the persistence parameter, σ_w as the standard deviation of the shocks, and ϑ_w as the weight of the MA component.

2.2.3 Capital firms and investment

These elements also align with other medium-scale DSGE models (further details are provided in Appendix A.2.4). Following Auclert et al. (2020), capital firms provide capital K_t to intermediate firms in exchange for $R_t^K K_t$, where R_t^K is the rental rate of the capital. These firms aim to maximize their profits, D_t^K, by deciding how much installed capital $\bar{K}_t$ to own, how much to invest in new capital I_t, and the utilization rate u_t of the capital in the production function.

Capital firms convert final goods into new capital using an investment-specific technology. Consistent with the DSGE literature, this technology is subject to adjustment costs, modeled as in Christiano et al. (2005). The adjustment cost function S satisfies $S(1) = S'(1) = 0$ and $\chi \equiv S''(1) > 0$. The law of motion of capital is influenced by investment-specific technology shocks, which follow an AR(1) process. The persistence of these shocks is governed by ϱ_I, and their standard deviation is denoted by σ_I.

Capital utilization is modeled following Smets and Wouters (2007). Capital firms choose the degree of utilization u_t while accounting for the associated cost $a(u_t)$. The capital used in the production function K_t is defined as:

$$K_t = u_t \bar{K}_{t-1}$$

Higher utilization rates increase costs for capital firms. The cost function $a(u_t)$ satisfies $a(1) = 0$, and the parameter λ is defined as $\frac{a''(1)}{a'(1)} \equiv \frac{\lambda}{1-\lambda}$. In steady state, all available capital is utilized, so that $u = 1$ and the cost associated to capital utilization is null.

2.3 Government

2.3.1 Monetary policy

The nominal interest rate is determined by a Taylor rule, expressed as:

$$1 + i_t = (1 + r)^{1-\varrho_r}(1 + i_{t-1})^{\varrho_r}\Pi_t^{(1-\varrho_r)\phi_\pi}\left(\frac{Y_t}{Y_{t-1}}\right)^{(1-\varrho_r)\phi_y} F_t$$

where i_t is the nominal interest rate and $r = R - 1$ represents the steady state interest rate. The parameter ϱ_r captures the degree of inertia in the nominal interest rate. The central bank seeks to stabilize inflation, with ϕ_π representing the sensitivity of the nominal interest rate to inflation. Additionally, the central bank responds to output growth, with ϕ_y measuring the sensitivity of the nominal interest rate to changes in output growth.

Monetary policy innovations are represented by F_t, which is assumed to follow an AR(1) process:

$$\log F_t \equiv f_t = \varrho_f f_{t-1} + \sigma_f \epsilon_t^f$$

where ϱ_f is the persistence parameter, σ_f is the standard deviation of the monetary policy shock.

2.3.2 Fiscal policy

Government spending is assumed to follow a simple rule:

$$G_t = \Gamma_t G$$

where G represents the level of government spendings at steady state. Following Bayer et al. (2024), G is set to one-fifth of the steady-state output Y, implying $G = \frac{Y}{5}$. The stochastic process Γ_t introduces fiscal stimuli and follows an AR(1) process:

$$\log \Gamma_t \equiv g_t = \varrho_g g_{t-1} + \sigma_g \epsilon_t^g + \varrho_{gz}\sigma_z \epsilon_t^z$$

where ϱ_g is the persistence parameter, σ_g is the standard deviation of the shock, and ϱ_{gz} captures the correlation between government spending and the TFP shock, as in Smets and Wouters (2007).

To ensure a balanced government budget, lump-sum taxes paid by households adjust accordingly. The government budget constraint is given by:

$$G_t + \tilde{R}_{t-1}B_{t-1} = B_t + T_t$$

2.4 Equilibrium and aggregate consumption

In equilibrium, all markets clear. By combining the government's budget constraint with the aggregate budget constraint of households, the goods market clearing condition is derived as:

$$C_t + I_t + G_t + a(u_t)\bar{K}_{t-1} = Y_t$$

By solving the household's optimization problem in the presence of aggregate risk, the log-linear expression for aggregate consumption is obtained (details in Appendix A.2.5):

$$Cc_t = C\mathbb{E}_t c_{t+1} - \frac{1}{\gamma}\tilde{r}_t - \left(\gamma\mu^2\tilde{\sigma}^2\mathbb{E}_t\hat{\mu}_{t+1} + \varphi\gamma\mu^2\tilde{\sigma}^2 Y\mathbb{E}_t y_{t+1}\right) \tag{3}$$

where c_t, y_t $\tilde{r}_t$ and $\hat{\mu}_t$ are the log-linear expressions[8] for consumption, output, the real interest rate, and the aggregate MPC respectively. The terms in parentheses on the right-hand side of (3) represent the "wedge in the (aggregate) Euler equation." When idiosyncratic risk is absent, these terms vanish, and the model reduces to the RA case. When idiosyncratic risk is present, the aggregate MPC influences aggregate consumption, with the log-linear form given by:

$$\hat{\mu}_t = \frac{1}{R}\left(\tilde{r}_t + \mathbb{E}_t\hat{\mu}_{t+1}\right) \tag{4}$$

2.5 Model setup

The main equations of the model are summarized in Appendix A.3. The steady-state conditions are outlined in Appendix A.4, and the log-linearized model is presented in Appendix A.5.

3 Data and estimation procedure

The model is estimated using seven U.S. time series commonly referenced in the DSGE literature. These include the (shadow) federal funds rate,[9] hours worked per capita, and the growth rates of the GDP deflator, real GDP per capita, real consumption per capita, real investment per capita, and real wages. The series used for estimation are sourced from FRED and span the period from 1954Q3 to 2019Q4, capturing a wide range of economic conditions, from the Great Inflation to the Great Recession. The dataset utilized is identical to that of Bayer et al. (2024), facilitating a straightforward comparison of results (see Appendix A.1 for further details). To reproduce these seven time series, I incorporate seven aggregate shocks: a TFP shock, an investment-specific

[8] Let N_t denote an endogenous variable of the model, its log-linear expression is expressed as $n_t \equiv \frac{N_t - N}{N} \approx \log\left(\frac{N_t}{N}\right)$.

[9] As in Bayer et al. (2024), I use the Wu and Xia (2016) shadow federal funds rate for the period 2009Q1 to 2015Q4.

technology (IST) shock, a price mark-up shock, a wage mark-up shock, a financial intermediation shock, a government spending shock, and a monetary policy shock.

The model is estimated using full-information Bayesian techniques in a state-space framework, leveraging a log-linear representation. Bayesian estimation, standard in the DSGE literature, aligns with methodologies employed by Smets and Wouters (2007) and is also prevalent in the HANK literature (Challe et al., 2017; Auclert et al., 2020; Bilbiie et al., 2023; Cho, 2023; Bayer et al., 2024). The procedure begins with setting prior distributions for the parameters by selecting a prior mean and standard deviation for each parameter of interest. The mode of the posterior distribution is then estimated by maximizing the log posterior function, which combines the prior distributions with the likelihood of the data. To obtain a comprehensive view of the posterior distribution and evaluate the model's Log Data Density (LDD), I employ the Metropolis-Hastings algorithm to generate a single chain of draws. After a long burn-in, 1,000,000 draws from the posterior are used to compute the posterior statistics. Appendix A.6 reports convergence diagnostics following the method of Geweke (1992), which indicate convergence for all parameters.

Initially, I estimate nearly all model parameters. Subsequently, I calibrate certain parameters to assess the sensitivity of the results to key parameters (see Section 7).

4 Estimation results

In this section, I present and analyze the results of my model's estimation. Almost all parameters are estimated, with the exception of $\bar{\xi}$,[10] the elasticity of substitution between goods θ, and the capital depreciation rate δ.[11]

4.1 Prior distribution

Table 1 provides the prior and posterior distributions of the parameters associated with the shock processes. The prior distributions are consistent with those commonly used in the DSGE literature (Smets and Wouters, 2007; Justiniano et al., 2010, 2011). Similarly, Table 2 outlines the prior and posterior distributions for the structural parameters, which are also largely in line with established literature.

For parameters related to idiosyncratic risk, I set priors based on the calibration choices in Acharya and Dogra (2020) and Acharya et al. (2023). Drawing from Guvenen et al. (2014), they argue that the appropriate value for the standard deviation of individual incomes slightly exceeds 0.5.[12] Accordingly, I set the prior mean of σ at 0.6 with a standard deviation of 0.15. Additionally,

[10] $\bar{\xi}$ is calibrated to ensure that labor supply equals unity in the steady state.

[11] In line with standard practices in the DSGE literature, these parameters are calibrated. Following Bayer et al. (2024), I set $\theta = 11$ and $\delta = 0.0175$.

[12] In Acharya and Dogra (2020) and Acharya et al. (2023), steady-state output is normalized to 1, allowing σ to also be interpreted in units of aggregate income.

Table 1: Prior and posterior distribution of shock processes

	Prior distribution			Posterior distribution			
	Distr.	Mean	St. Dev.	Mode	Mean	95 percent	5 percent
ϱ_v	Beta	0.5	0.2	0.9305	0.9269	0.9028	0.9508
ϱ_z	Beta	0.5	0.2	0.9779	0.9757	0.9667	0.9849
ϱ_p	Beta	0.5	0.2	0.9888	0.9858	0.9759	0.9964
ϑ_p	Beta	0.5	0.2	0.6460	0.6206	0.4973	0.7470
ϱ_w	Beta	0.5	0.2	0.9904	0.9901	0.9858	0.9946
ϑ_w	Beta	0.5	0.2	0.9289	0.9145	0.8827	0.9481
ϱ_I	Beta	0.5	0.2	0.0878	0.1120	0.0312	0.1871
ϱ_f	Beta	0.5	0.2	0.2436	0.2459	0.1573	0.3347
ϱ_g	Beta	0.5	0.2	0.9768	0.9748	0.9613	0.9889
ϱ_{gz}	Beta	0.5	0.2	0.9025	0.8778	0.7876	0.9760
σ_v	Invgamma	0.001	2	0.0021	0.0022	0.0017	0.0028
σ_z	Invgamma	0.001	2	0.0064	0.0065	0.0060	0.0069
σ_p	Invgamma	0.001	2	0.0015	0.0016	0.0013	0.0018
σ_w	Invgamma	0.001	2	0.0083	0.0085	0.0075	0.0094
σ_I	Invgamma	0.001	2	0.0196	0.0195	0.0175	0.0215
σ_f	Invgamma	0.001	2	0.0020	0.0021	0.0019	0.0022
σ_g	Invgamma	0.001	2	0.0098	0.0099	0.0091	0.0106

Note: The posterior distribution is obtained using the Random Walk Metropolis-Hastings algorithm.

they propose that the calibration of risk cyclicality φ should yield strongly countercyclical risk, corresponding to values below -5. I set the prior mean at -5 and the standard deviation at 2. In Section 7, I examine the implications of alternative calibrations of the parameters specific to the CARA-Normal framework.

4.2 Posterior distribution

For parameters common to other medium-scale DSGE models, I find that the posterior distribution closely aligns with results from previous studies (Smets and Wouters, 2007; Justiniano et al., 2010, 2011). It is also consistent with posterior distributions obtained in estimated HANK models (Auclert et al., 2020; Bilbiie et al., 2023; Bayer et al., 2024). Both price and nominal wage rigidities are estimated to be pretty high. This aligns with the typical finding in estimated DSGE models that both nominal rigidities are important. The capital share, α, is estimated to be relatively low, around 0.23, which is consistent with the literature, where estimates for α often fall below 0.25. The posterior mean of the steady-state interest rate is approximately 0.01, corresponding to a 4% annual rate. This aligns with the average federal funds rate over the sample period.[13]

[13] One potential concern is the consistency of the estimated steady-state parameters, as the model is estimated using demeaned data. However, given that the parameter estimates are reasonable and consistent with the litera-

Table 2: Prior and posterior distribution of structural parameters

	Prior distribution			Posterior distribution			
	Distr.	Mean	St. Dev.	Mode	Mean	95 percent	5 percent
r	Gamma	0.012	0.001	0.0102	0.0103	0.0090	0.0115
α	Beta	0.3	0.05	0.2322	0.2325	0.2196	0.2455
ϱ_r	Beta	0.5	0.2	0.8059	0.8055	0.7747	0.8356
ϕ_π	Normal	1.5	0.25	1.7223	1.7765	1.5367	2.0153
ϕ_y	Gamma	0.12	0.05	0.1983	0.2070	0.1103	0.3026
ζ_p	Beta	0.7	0.2	0.8192	0.8144	0.7828	0.8477
ζ_w	Beta	0.7	0.2	0.8662	0.8527	0.8138	0.8922
λ	Beta	0.5	0.15	0.6777	0.6476	0.5677	0.7824
χ	Gamma	4	1	3.7288	3.8377	2.8559	4.7872
γ	Gamma	3	1.5	3.6546	3.6498	2.8603	4.3953
ρ	Gamma	1	0.5	1.7310	1.9840	1.0712	2.8422
φ	Normal	-5	2	-7.5312	-7.4956	-10.0259	-4.9102
σ	Gamma	0.6	0.15	0.8657	0.8909	0.6803	1.0947

Note: The posterior distribution is obtained using the Random Walk Metropolis-Hastings algorithm.

For the parameter governing the degree of risk aversion, the posterior mean of γ is slightly higher than its prior mean (the posterior mean is 3.6498 compared to a prior mean of 3). High risk aversion is common in the HANK literature, where the relative risk aversion parameter typically ranges from 3 to 4 (e.g., Bayer et al. (2024) set this parameter at 4). Higher risk aversion strengthens the precautionary savings motive. In Section 7, I calibrate γ to 2 to evaluate the sensitivity of the results to a lower degree of risk aversion.

Having established the conformity of the posterior distribution for the standard parameters, I now turn to the results for the parameters related to idiosyncratic risk. The posterior mean of the standard deviation of individual risk in units of aggregate income, σ, is estimated at 0.8909, significantly above zero. This implies that the standard deviation of individual incomes is approximately 90% of aggregate income, a notably high value.[14] The confidence interval for σ decisively excludes values near zero. As σ approaches zero, individual risk diminishes, and the model converges towards the RANK framework. The fact that the data strongly supports a high σ underscores the importance of accounting for individual risk in a DSGE model.

The cyclicality of idiosyncratic risk, φ, is estimated to be significantly more countercyclical than suggested by the prior. The posterior mean is -7.4956, with a confidence interval well below

ture, I argue that estimating these parameters does not harm the quality of the estimation. The posterior mean of α is close to the calibration used in Auclert et al. (2020), and the posterior distribution of r implies that $R^{-1} \simeq 0.99$, which is a standard quarterly calibration in the literature.

[14] The model assumes that idiosyncratic shocks are normally distributed, which allows for the aggregation of individual consumptions. However, this assumption prevents the model from capturing the skewness in the actual distribution of individual incomes. Consequently, the high value of σ partly reflects the fact that individual incomes are not normally distributed but are highly dispersed in reality.

zero. This result unambiguously indicates that the data favors strongly countercyclical risk. This finding contributes to the debate in the HANK literature regarding the Forward Guidance Puzzle (FGP) and the cyclicality of individual risk. McKay et al. (2016) argued that HANK models solve the FGP but implicitly assumed procyclical individual risk. As demonstrated by Acharya and Dogra (2020), HANK models can only solve the FGP under the assumption of procyclical risk. Similarly, Werning (2015) contended that HANK models fail to solve the FGP when risk is countercyclical or when liquidity is procyclical. My results align with the conclusion that HANK models, when calibrated in an empirically relevant manner, do not solve the FGP.

In summary, the findings indicate that the data supports high risk aversion, a significant role for individual risk, and strongly countercyclical idiosyncratic risk. These factors collectively result in a pronounced precautionary saving motive over the business cycle.

4.3 Estimated dynamics of individual risk over the business cycle

Estimating the model using macroeconomic time series data enables the recovery of unobserved variables. In particular, the wedge in the Euler equation—the last two terms on the right-hand side of (3)—is extracted using the Kalman smoother. This smoothed variable captures the influence of idiosyncratic risk and precautionary savings on business cycle dynamics over the past decades.

The wedge consists of two components: the "MPC component" ($\gamma\mu^2\tilde{\sigma}^2\mathbb{E}_t\hat{\mu}_{t+1}$) and the "Cyclical component" ($\varphi\gamma\mu^2\tilde{\sigma}^2Y\mathbb{E}_t y_{t+1}$). The correlation between the wedge and its cyclical component is 0.999, while the correlation with the MPC component is only 0.22. This indicates that the wedge is driven almost entirely by its cyclical component.

Table 3: Unconditional variance decompositions of the wedge and its components

In %	Risk prem.	TFP	Mon. pol.	Price MU	IST	Gov. spend.	Wage MU
Wedge	8.49	13.75	2.95	37.63	4.19	1.32	31.67
Cyc. comp.	7.32	12.87	2.48	40.08	4.20	1.87	31.17
MPC comp.	24.62	12.03	10.64	23.88	9.47	12.06	7.30

To further explore the determinants of the wedge and its components, it is useful to examine their variance decompositions. Table 3 reports the contribution of each structural shock to the variance of the wedge and its components. As shown, the variance decomposition of the wedge closely mirrors that of the cyclical component.

Similarly, Figure A.1 in the appendix presents the historical decompositions of the wedge and its components. The historical decompositions of the wedge and its cyclical component are nearly identical, whereas the one of the MPC component exhibits a distinct pattern.

These findings confirm that the shocks influencing expected output are also those driving the wedge. Since the objective of this paper is to quantify the effects of cyclical idiosyncratic risk, the fact that the cyclical component is the dominant force behind the wedge is a desirable property of the model.

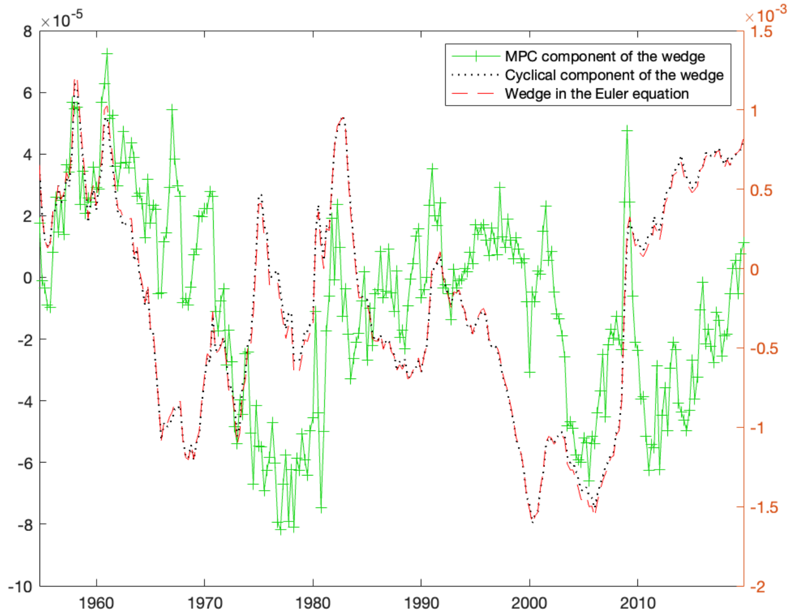

Figure 1: The estimated wedge in the Euler equation compared to its cyclical (right axis) and its MPC components (left axis)

To further illuminate the dynamics of the wedge and its components, Figure 1 plots the wedge (red dashed line) and its cyclical component (black dotted line) against the right axis, while the MPC component (solid green line) is plotted against the left axis. It is evident from the figure that the cyclical component accounts for the majority of fluctuations in the wedge. The next section delves deeper into the estimated dynamics of the wedge throughout the business cycle.

5 Does the model deliver realistic dynamics of individual risk and precautionary savings?

As highlighted in the previous section, the data strongly supports the presence of significant individual risk that is highly countercyclical, leading to substantial precautionary savings over the business cycle. This raises an important question: Is this a realistic measure of individual risk? Since microeconomic data is not used to directly inform the estimation of individual risk, concerns may arise regarding the consistency of the estimated parameters that govern the wedge in the Euler equation. This section addresses these concerns and seeks to answer the central question posed by the title of this paper.

To evaluate the realism of the estimated individual risk, I compare the estimated wedge in the

Euler equation with empirical measures of individual risk. This analysis aims to determine which empirical measure of individual risk aligns most closely with the estimated wedge. It also allows to assess whether the estimated wedge and its associated parameters are consistent with these empirical measures. The results of this exercise will provide crucial insights into the validity and reliability of the model's representation of individual risk.

5.1 Labor market conditions

I begin by comparing the estimated wedge with various measures related to labor market conditions and the associated individual risk. Since the relationship between employment conditions and individual risk is central to HANK&SaM models, this analysis also provides valuable insights for this literature.

The first comparison is between the estimated wedge and the unemployment rate. Figure 2 illustrates the U.S. unemployment rate, represented by the solid blue line and plotted against the left axis. The estimated wedge is depicted by the red dashed line and plotted against the right axis. Visually, the unemployment rate and the wedge appear strongly correlated. The estimated wedge demonstrates a good fit with unemployment, as both series share similar peaks and troughs. The correlation between the unemployment rate and the estimated wedge is 0.4.

The model, which imposes minimal assumptions about the sources of individual risk, captures unemployment fluctuations over the past decades with impressive accuracy. However, the fit begins to deteriorate after 2013, when unemployment continued to decline sharply while output growth and inflation remained subdued. Focusing on the period from 1954Q4 to 2013Q1, the correlation between the unemployment rate and the estimated wedge increases to 0.57. This correlation is even stronger over the 1968Q1–2013Q1 subperiod, reaching a peak of 0.86. These findings suggest that the period following the Great Recession is an outlier. Possible explanations for this phenomenon are discussed in Subsection 5.4.

These findings suggest that countercyclical idiosyncratic risk, measured by the wedge, has a dynamics similar to the one of unemployment in normal times. The exceptionally strong fit from the early 1950s to the early 2010s indicates that unemployment risk is a likely key driver of precautionary savings. This result aligns with the predictions of HANK&SaM models, which derive a precautionary saving motive from the risk of unemployment. Empirically, this connection is well-supported by the evidence presented here.

In HANK&SaM models, it is not the unemployment rate itself that drives precautionary savings. Instead, countercyclical risk is governed by the job-finding and job-loss rates, which determine the conditional probabilities of becoming unemployed, finding a new job, remaining unemployed, or staying employed. In estimated HANK&SaM models, these rates are often used as observables. Challe et al. (2017) and Cho (2023) compute these rates following the method of Shimer (2005) and Shimer (2012). Following this precedent, I also compare these measures with the estimated

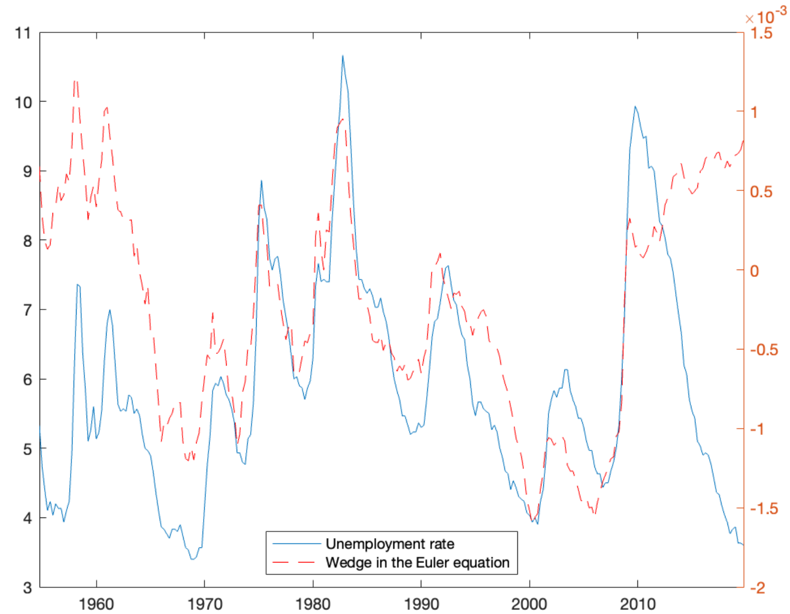

Figure 2: The estimated wedge in the Euler equation (right axis) compared to the unemployment rate (left axis)

wedge. Appendix A.1 details precisely how these rates are computed.

In Figure 3, the job-finding rate is represented by a solid blue line, while the estimated wedge is plotted as in Figure 2. A clear negative correlation emerges between the job-finding rate and the estimated wedge. As the probability of finding a new job increases, the risk of becoming unemployed decreases, which explains the negative relationship. Specifically, the correlation between the job-finding rate and the wedge is -0.39. This indicates a fairly strong negative correlation. From 1968Q1 to 2013Q1, the correlation increases at -0.54.

In Figure 4, the job-loss rate is shown as a solid blue line, plotted alongside the estimated wedge as in previous figures. The two series appear positively correlated: when the job-loss rate increases, the likelihood of becoming unemployed rises, heightening individual risk and thereby increasing the wedge. Quantitatively, the correlation between the job-loss rate and the estimated wedge is 0.24 over the full sample, rising to 0.51 during the 1968Q1–2013Q1 subperiod.

These findings indicate that the estimated wedge is strongly positively correlated with the job-loss rate and strongly negatively correlated with the job-finding rate. This suggests that both statistics are meaningful indicators of individual risk. However, the unemployment rate proves to be a better predictor of the wedge than either the job-loss or job-finding rates. As with unemployment, the correlations between these rates and the wedge decline after 2013. In particular, the job-finding rate rises sharply and the job-loss rate falls to historically low levels, yet precautionary savings remain elevated.

I argue that the primary reason the unemployment rate outperforms the job-finding and job-

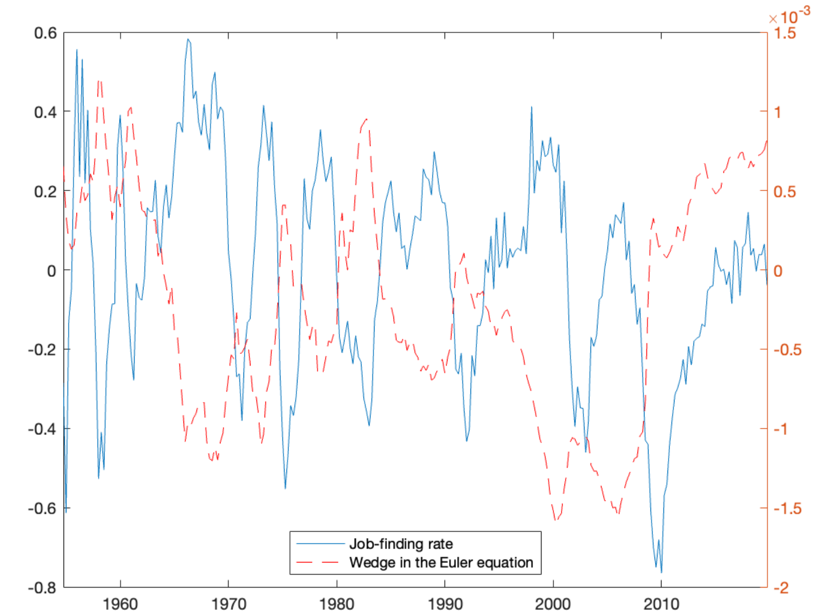

Figure 3: The estimated wedge in the Euler equation (right axis) compared with the job-finding rate (left axis)

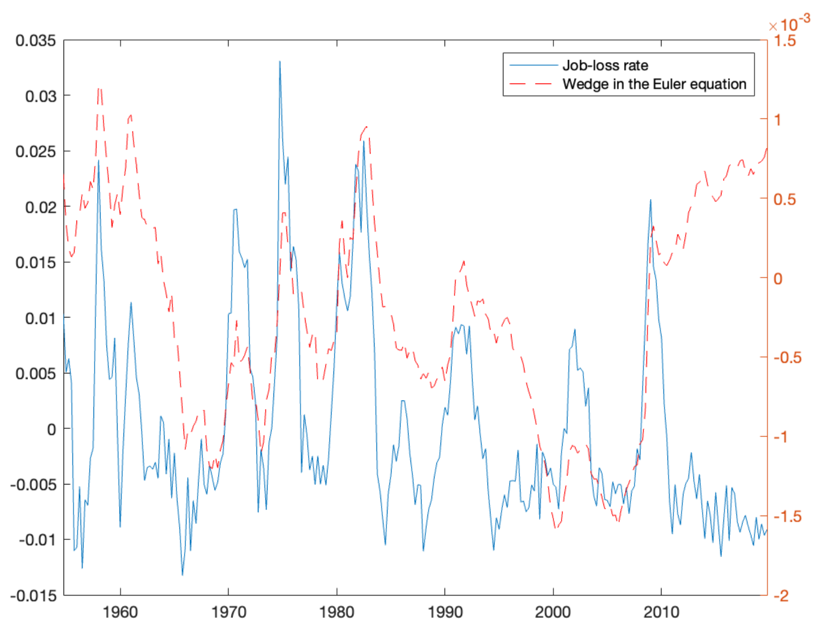

Figure 4: The estimated wedge in the Euler equation (right axis) compared with the job-loss rate (left axis)

loss rates in tracking the wedge lies in its cyclicality. As discussed in Subsection 4.3, the wedge is primarily driven by its cyclical component. Therefore, the relative cyclicality of the unemployment rate compared to the job-finding and job-loss rates is central to explaining its superior performance in capturing the dynamics of the wedge.

Table 4: Correlations of the wedge and the labor market statistics with output

Period considered	Wedge	Unemployment	Job-finding rate	Job-loss rate
Full sample	−0.9985	−0.4134	0.3956	−0.2290
From 1968Q1 to 2013Q1	−0.9977	−0.8778	0.5479	−0.4906

Note: Correlations are computed using the output y_t series, extracted via the Kalman smoother.

To investigate this further, I compute the cross-correlations between these observables and smoothed output y_t. Table 4 reports these statistics for both the full sample and the subperiod during which unemployment best tracks the wedge. Over the full sample, the unemployment rate shows a slightly stronger correlation with output than the job-finding rate, while the job-loss rate appears to be the least cyclical of the three. In the restricted sample from 1968Q1 to 2013Q1, the absolute correlation between output and the unemployment rate increases to 0.88, closely approaching the correlation between the wedge and output. During this same period, the correlation between output and the job-finding rate is 0.55, while the job-loss rate becomes more countercyclical, with its correlation with output rising to 0.49 in absolute value.

These findings suggest that the stronger cyclicality of the unemployment rate explains its superior performance in capturing the dynamics of the wedge, which is itself primarily driven by its cyclical component. Although the job-finding and job-loss rates offer valuable information about individual risk, the unemployment rate proves to be a more effective and practical proxy. This underscores the usefulness of the unemployment rate as a simplified yet robust measure of countercyclical individual risk.

5.2 Income idiosyncratic risk measured in Bayer et al. (2019)

Next, I turn to an empirical measure of income idiosyncratic risk estimated in Bayer et al. (2019), which aligns more closely with the concept of individual risk commonly used in HANK models. In Bayer et al. (2024), the authors use this time series as an observable, demonstrating that their HANK model can simultaneously fit both macroeconomic time series and this measure of individual risk (though their model explicitly targets this measure during estimation, which mine does not). The dataset spans from 1983Q1 to 2013Q1 and is based on the Survey of Income and Program Participation.

For comparison with my results, I plot the measure targeted in Bayer et al. (2024) in Figure 5, alongside the estimated wedge as in previous graphs. The measure of income idiosyncratic risk is represented by the solid blue line and plotted against the left axis.

Figure 5: The estimated wedge in the Euler equation (right axis) compared to an empirical measure of idiosyncratic income risk (Bayer et al., 2019) (left axis)

Examining the fit between the wedge and the empirical measure from Bayer et al. (2019), the two series appear to follow a similar pattern, sharing broadly comparable peaks and troughs. However, there is a slight delay in the response of precautionary savings during the early 1990s and after the Great Recession. The correlation between the measure from Bayer et al. (2019) and the estimated wedge is 0.28.

Overall, the fit between this empirical measure and the estimated wedge is reasonably strong, although it is less pronounced than the fit observed with labor market-related measures. This suggests that while the measure from Bayer et al. (2019) captures aspects of individual risk relevant to precautionary savings, labor market conditions may provide a more direct and robust proxy for the dynamics of countercyclical idiosyncratic risk.

5.3 Consumer sentiment

Finally, I compare the estimated wedge with a measure of consumer sentiment. The dataset, compiled by researchers from the University of Michigan, is based on the Surveys of Consumers and spans from 1959Q4 to 2019Q4. While consumer sentiment is not a direct source of individual risk, it is reasonable to expect a negative correlation with precautionary savings. Intuitively, when individuals become pessimistic about the present and future states of the economy, they are likely to save more as a precautionary measure.

In Figure 6, consumer sentiment is represented by the solid blue line and plotted against the left axis, while the estimated wedge is displayed as in previous graphs. Examining the fit between

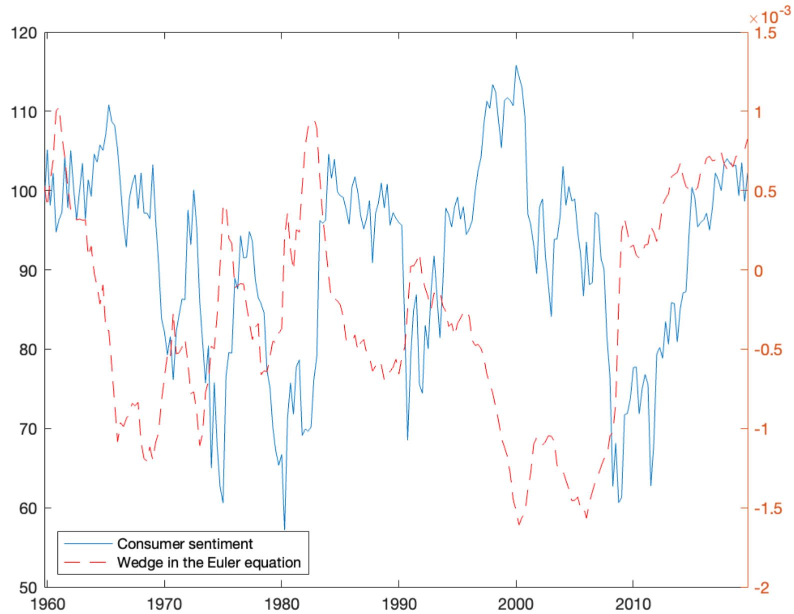

Figure 6: The estimated wedge in the Euler equation (right axis) compared to consumer sentiment (left axis)

the wedge and consumer sentiment, one observes that when consumer sentiment moves in one direction, the estimated wedge generally moves in the opposite direction. For instance, during recessions, when consumer sentiment declines, precautionary savings tend to increase. Conversely, when consumer sentiment improves, precautionary savings decrease. As expected, the correlation between consumer sentiment and the smoothed version is negative, at -0.23.

These results confirm the anticipated negative relationship between consumer sentiment and precautionary savings. While consumer sentiment is not a direct measure of individual risk, its correlation with the estimated wedge suggests it captures broader economic conditions that influence precautionary behavior.

5.4 Individual risk in normal times and after the Great Recession

The effects of idiosyncratic risk, as summarized by the wedge, appear to have realistic amplitudes and dynamics. The fit between unemployment and the estimated wedge is particularly striking over the period from the early 1950s to the early 2010s. This strong correlation underscores the fact that my model favors a form of idiosyncratic risk that is strongly countercyclical, like unemployment risk in normal times. These findings demonstrate that HANK&SaM models are empirically relevant in capturing the effects of idiosyncratic risk, as the dynamics of unemployment fluctuations closely mirror those of the wedge in normal periods.

However, the relationship between labor market conditions and precautionary savings seems to have shifted during the 2010s. The Great Recession and its aftermath appear to have had a lasting impact on precautionary savings, disrupting the typical relationship with unemployment. My model estimates indicate that precautionary savings remained elevated long after the initial downturn, even as unemployment and the job-loss rate declined sharply and the job-finding rate increased dramatically.

Figure A.1 and Figure A.2 in the appendix display the historical decompositions of the wedge, its components, and output. Since the wedge is primarily driven by expected output, the prolonged period of below-trend output from the Great Recession to the end of the sample explains the persistently high level of precautionary savings. Therefore, the disconnect observed in the 2010s between labor market statistics and the wedge reflects a structural shift in the relationship between labor market outcomes and output.

The literature on liquidity traps in HANK&SaM models provides valuable insights into this phenomenon. A seminal contribution by Ravn and Sterk (2021) explores the interplay between unemployment, precautionary savings, and liquidity traps. Their framework demonstrates that the interaction between nominal rigidities and uninsurable unemployment risk can lead to three distinct steady states: one for normal times, another characterized by an unemployment trap with persistently high unemployment and depressed demand, and a third representing a liquidity trap caused by the Zero Lower Bound. Importantly, their framework suggests that a liquidity trap does not necessarily result in deflation, which aligns with the experiences of many developed economies during the 2010s.

Nevertheless, explaining why precautionary savings remained persistently high and aggregate demand stayed durably depressed over the past decade—despite a sharp decline in unemployment risk—remains an open question. Aggregate demand may be subdued for reasons beyond the scope of my model, such as population aging, rising wealth inequality, and secular stagnation. The decline in labor force participation may also have contributed to this phenomenon, as it mechanically reduced the unemployment rate without necessarily diminishing idiosyncratic risk. Investigating how such structural shifts influence individual risk and precautionary savings represents an important direction for future research.

6 The effects of idiosyncratic risk on business cycles

In this section, I show that the effects of precautionary savings on consumption and output vary considerably depending on the type of shock. My model incorporates seven aggregate shocks: four demand shocks and three supply shocks.[15] Unlike most HANK models, my model does not include agents with high MPCs, allowing for an isolated examination of the effects of precautionary

[15] Demand shocks include monetary policy, government spending, financial intermediation, and IST shocks. Supply shocks include TFP, wage mark-up, and price mark-up shocks. IST shocks behave like demand shocks, but they are sometimes classified as supply shocks. For the analysis of my results, it is more convenient to categorize them as demand shocks, as their effects are similar to those of the monetary policy and the risk premium shocks.

savings on business cycles. In conventional HANK models, the volatility generated by high-MPC agents can overshadow the effects of precautionary savings.

Acharya and Dogra (2020) and Acharya et al. (2023) show that countercyclical risk amplifies aggregate fluctuations compared to the RA case, while procyclical risk attenuates output and consumption responses. This applies to aggregate shocks, forward guidance, and fiscal multipliers. However, their model does not include investment. This means that, for most shocks, amplification of consumption responses automatically translates into amplified output responses. In this section, I demonstrate that this result does not always hold in a similar model that includes capital accumulation and investment. To do so, I compute the IRFs of the model following each shock, calibrating parameters to the posterior mean (as shown in Table 1 and Table 2). The IRFs for the seven aggregate shocks are plotted in Appendix A.8, alongside a counterfactual RA economy calibrated with $\sigma = 0$, while keeping all other parameters unchanged.

For comparison, I also include the IRFs of a reestimated RANK model, with the posterior means of the parameters reported in Table 5. Given the substantial differences in parameter values between the baseline and the reestimated RANK models, I focus primarily on the comparison between the baseline and the RA counterfactual with $\sigma = 0$. This approach allows for a clearer assessment of the aggregate effects of idiosyncratic risk.[16]

6.1 Demand shocks

For demand shocks, countercyclical risk generates modest amplification of both consumption and output responses for monetary policy innovations and risk premium shocks (in comparison to the RA counterfactual with $\sigma = 0$). These results align with Acharya and Dogra (2020) and Acharya et al. (2023): countercyclical risk amplifies consumption and output responses compared to the RANK counterfactual. A closer look reveals that capital and investment responses are slightly attenuated, while consumption responses are more strongly amplified, leading to overall output amplification. As discussed in Challe et al. (2017), precautionary savings can attenuate investment responses, as households' additional savings generate downward pressures on the real interest rate.

Examining (3), one can see that the wedge imposes additional pressures on the real interest rate, $\tilde{r}_t$. Since the shadow value of capital, q_t, depends negatively on the real rate, a lower $\tilde{r}_t$ exerts upward pressure on the price of capital, and vice versa. Because investment responds positively to q_t, these pressures on the price of capital are transmitted to investment dynamics. In the case of negative demand shocks, the rise in precautionary savings mitigates the decline in investment—without this mechanism, as shown in the RA counterfactual, the drop in investment would have been more pronounced.[17] Combining the log-linearized Euler equations for actuarial

[16]Therefore, when I state that idiosyncratic risk amplifies or dampens aggregate fluctuations, the comparison is made relative to the RA counterfactual, not the reestimated RANK model.

[17]It is important to emphasize that in this framework, households do not use capital as a vehicle for self-insurance against idiosyncratic risk. Instead, they rely exclusively on actuarial bonds for this purpose. Consequently, pre-

bonds and capital with the investment optimality condition yields the following expression for log-linear investment ι_t (see Appendix A.2.6 for derivations):

$$
\begin{aligned}
\iota_t = &\frac{1}{1+\tilde{\beta}}\iota_{t-1} + \frac{\tilde{\beta}}{1+\tilde{\beta}}\mathbb{E}_t\iota_{t+1} + \hat{\psi}_t \\
&+ \frac{\gamma}{\chi(1+\tilde{\beta})}\mathbb{E}_t\sum_{s=0}^{\infty}\tilde{\beta}^s(1-\delta)^s\left[\frac{\tilde{\beta}}{\gamma}R^K r^K_{t+1+s} - C(c_{t+1+s}-c_{t+s}) + \underbrace{\gamma\mu^2\tilde{\sigma}^2\hat{\mu}_{t+1+s} + \varphi\gamma\mu^2\tilde{\sigma}^2 Y y_{t+1+s}}_{\text{Wedge in } t+s}\right]
\end{aligned}
$$

where $\hat{\psi}_t$ is the renormalized IST process and r_t^K is the (log-linearized) rental rate of capital. This equation illustrates well how the contemporaneous wedge and the expected path of future wedges influence investment. Since the cyclical component is the dominant driver of the wedge (and φ is negative), precautionary savings exert upward pressure on investment during recessions and downward pressure during expansions. This mechanism reflects how expectations of future output propagate through the wedge to influence capital accumulation dynamics.

For investment-specific technology shocks, the effects on investment are similar in both specifications. Consumption responses are initially amplified but do not persist, resulting in minimal differences between the baseline model and the RA counterfactual for output responses. The low estimated persistence of this shock explains the limited divergence between the two models.

The case of increased government spending is more nuanced, as it is the only shock where output and consumption are negatively correlated. In both the baseline and the RANK models, higher government spending reduces consumption. However, precautionary savings decrease as output rises, given the countercyclical nature of individual risk. This decline in precautionary savings amplifies the crowding out of investment compared to the RA counterfactual but attenuates the initial drop in consumption. Ultimately, the response of output is slightly attenuated in comparison to the RA counterfactual.

Comparing the baseline with the reestimated RANK model, the risk premium and monetary policy shocks exhibit greater and more persistent effects in the baseline. In contrast, the effects of the IST shock are significantly more pronounced in the reestimated RANK model, owing to the much higher estimated persistence of this shock (see Table 5). Regarding the government spending shock, the decline in consumption is considerably stronger in the reestimated RANK model, whereas the reduction in investment is more pronounced in the baseline. Ultimately, the overall output responses are quite similar across both specifications.

cautionary savings affect the capital stock only indirectly, through their influence on the equilibrium real interest rate.

6.2 Supply shocks

Supply shocks reveal a more complex interaction between precautionary savings and investment. For a TFP shock, consumption responses are initially amplified compared to the RA counterfactual, but overall output responses are attenuated. This occurs because the attenuation of investment responses is stronger and more persistent than the amplification of consumption responses. As output rises, precautionary savings decline, which dampens the rise in investment. The amplification of consumption is short-lived, as the smaller output increase generates a shorter drop in precautionary savings and reduces the rise in permanent income. Eventually, consumption responses weaken compared to the RA counterfactual due to the diminished rise in permanent income. Thus, precautionary savings create a persistent attenuation of investment responses but only a brief amplification of consumption responses, resulting in an overall attenuation of output responses relative to the RA counterfactual.

A similar dynamic occurs for price and wage mark-up shocks, albeit in the opposite direction. These shocks reduce output, increasing precautionary savings. The rise in precautionary savings attenuates the decline in investment and output compared to the RA counterfactual. Consumption responses are initially amplified but eventually attenuate as precautionary savings stabilize. As with the TFP shock, there is a turning point where consumption responses become weaker than in the RA counterpart, reflecting a smaller decline in permanent income.

Fluctuations in precautionary savings are markedly more persistent following supply shocks than in response to demand shocks. This reflects the fact that supply shocks have more enduring effects on output dynamics. Since the wedge is primarily driven by its cyclical component, supply shocks also lead to more persistent variations in precautionary savings.

As shown in Appendix A.2.6, the price of capital, q_t, depends on the expected trajectory of future wedges. Because supply shocks have long-lasting effects on the wedge, it exerts a stronger and more sustained influence on the current price of capital and, ultimately, on investment. In contrast, demand shocks elicit less persistent changes in precautionary savings (with the exception of the government spending shock), which diminishes the influence of the expected path of future wedges on q_t and results in a more limited effect on investment.

Regarding the IRFs of the reestimated RANK model for supply shocks, the responses to the TFP and wage mark-up shocks are quite similar to those observed in the RA counterfactual. In the case of the price mark-up shock, however, the effects are generally smaller in the reestimated RANK model.

6.3 Comparison with counterfactual risk cyclicalities

To further illustrate the ambiguous effects of risk cyclicality on output, I introduce an additional counterfactual scenario where individual risk is procyclical. In this scenario, I calibrate φ to the opposite value of its posterior mean, resulting in $\varphi = 7.4956$ (see Table 2). I then compute and

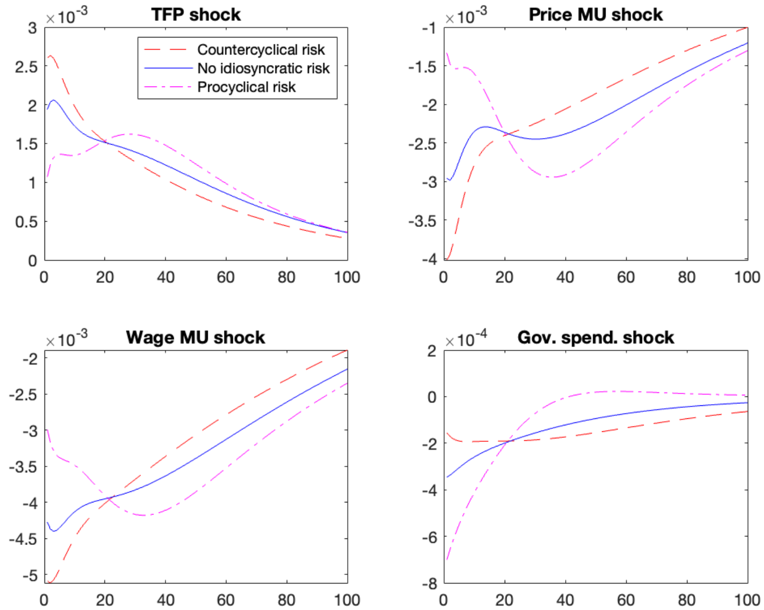

Figure 7: Consumption responses depending on the cyclicality of the idiosyncratic risk

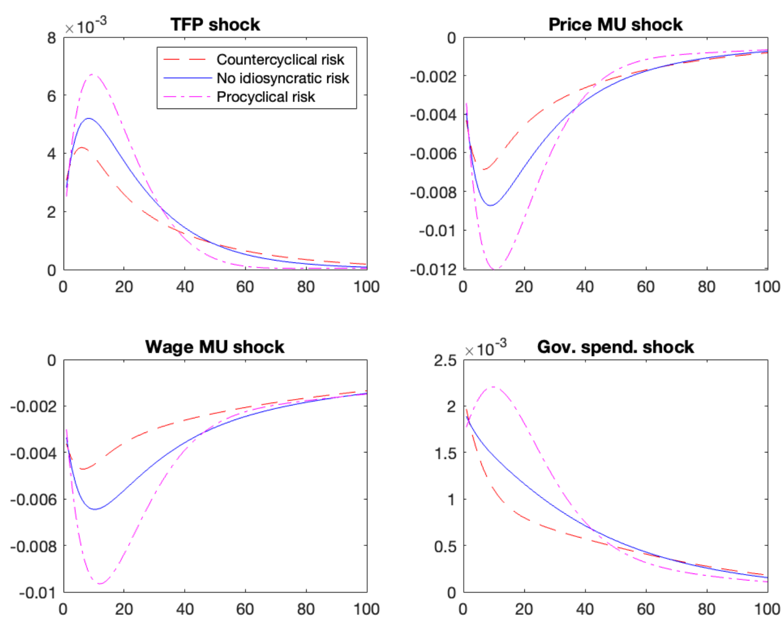

Figure 8: Output responses depending on the cyclicality of the idiosyncratic risk

plot the IRFs for this counterfactual case, focusing on the four shocks that generated significant differences between the estimated model and the RA counterfactual: government spending, TFP, wage mark-up, and price mark-up shocks.

In Figure 7, I plot the consumption responses for this new counterfactual alongside the baseline model with countercyclical risk and the RA counterfactual. As shown in the graphs, procyclical risk leads to an initial attenuation of consumption responses for the TFP and the two mark-up shocks. These findings are consistent with the results of Acharya and Dogra (2020) and Acharya et al. (2023), which demonstrate that procyclical risk dampens consumption responses. For government spending shocks, however, procyclical risk increases precautionary savings, resulting in a larger drop in consumption compared to the countercyclical and RA cases.

Figure 8 shows the output responses under different risk cyclicality scenarios, plotted in the same manner as for consumption. Here, procyclical risk amplifies output responses for the TFP and the two mark-up shocks, while countercyclical risk attenuates their effects. This supports the earlier discussion: when investment is included in the model, procyclical risk does not necessarily lead to an attenuation of output responses. In fact, investment responses are more amplified under procyclical risk. For government spending shocks, procyclical risk increases precautionary savings, which mitigates the crowding out of investment, ultimately amplifying the output response.

These figures highlight the intricate interplay between precautionary savings and investment. This mechanism is often obscured in quantitative HANK models because the responses of high-MPC consumers tend to dominate. However, the strong attenuation of investment responses following supply shocks, as observed here, is also a key feature of HANK&SaM models, as emphasized by Challe et al. (2017) and Den Haan et al. (2018). Using counterfactual scenarios, Challe et al. (2017) demonstrate the significance of this channel during recessions.

7 Alternative specifications and the quantitative relevance of idiosyncratic risk

In this section, I calibrate the parameters specific to the CARA-Normal framework to assess their relative importance for the precautionary saving motive and the model's ability to fit the data. Specifically, I examine how these calibrations affect the LDD, which evaluates the model's fit to the seven macroeconomic time series. As previously noted, the case of lower risk aversion is particularly significant, as it substantially weakens the precautionary saving motive. Additionally, I analyze the effects of acyclical individual risk and a reduced standard deviation of individual incomes. Furthermore, I compare the baseline model to its RA counterpart to highlight the role of individual risk and precautionary savings.

The methodology is straightforward: I calibrate each parameter of interest individually and reestimate the model for each scenario. The four cases I focus on are:

- $\sigma = 0$: Eliminates the effect of φ on consumption, representing the RA case.
- $\gamma = 2$: Represents a lower degree of risk aversion.
- $\sigma = 0.5$: Implies a standard deviation of individual incomes equal to 50% of aggregate income.
- $\varphi = 0$: Assumes acyclical individual risk.

The results of this analysis are summarized in Table 5, which reports the LDD and the mean of the posterior distribution for each scenario. For comparison, the table also includes the corresponding values from the baseline model (referred to as "Baseline," as shown in Table 1 and Table 2) and the ones from an estimated Two-agent New Keynesian (TANK) model (see Subsection 7.3).

7.1 Does HANK better fit the data?

One significant advantage of the CARA-Normal framework is that it nests the RANK model, which is recovered when $\sigma = 0$. Reestimating the model without individual risk allows us to assess whether the differences between the two specifications are substantial. To quantify this, I compare the LDD obtained without idiosyncratic risk to that of the baseline model.

As shown in Table 5, the LDD drops to 6766.7 in the RANK model, compared to 6774.3 in the baseline model. Incorporating countercyclical risk improves the LDD by 7.6—a modest yet meaningful improvement. This result demonstrates that, despite its simplicity, the baseline model fits the data better than its RA counterpart by a notable margin. For comparison, Bayer et al. (2024) obtained a LDD of 6588 with their HANK model and 6590 with a RANK version on the same dataset. Hence, surpassing the RA framework in terms of LDD is no trivial feat. The CARA-Normal model achieves this by estimating parameters that influence idiosyncratic risk and risk aversion (more details in the next subsection).

As shown in the variance and historical decompositions in Appendix A.7, the baseline model attributes little explanatory power to IST shocks in accounting for business cycle fluctuations, whereas these shocks emerge as the primary driver in the RANK model. This contrast is mainly due to the estimated persistence of the IST shock, which is very low in the baseline but significantly higher in the RANK model (see Table 5). Because the IST shock is largely muted in the baseline, the resulting variance decompositions differ substantially from those of the RANK model. In the baseline, risk premium and supply shocks play a central role, while in the RANK model, the IST shock dominates the dynamics of most variables.

Since I use the same dataset as Bayer et al. (2024), it is also possible to compare the fit of my baseline model to their HANK model. The baseline achieves a LDD of 6774.3, significantly outperforming the 6588 obtained by their HANK model. This substantial difference stems primarily from the fact that the log-linear structure of my model makes parameter estimation straightforward,

Table 5: Comparison across specifications

	Baseline	RANK	$\gamma = 2$	$\sigma = 0.5$	$\varphi = 0$	TANK
Log Data Density	6774.3	6766.7	6767.7	6766.1	6767.1	6773.4
Mean of the parameters						
σ_v S.d. Risk prem.	0.0022	0.0018	0.0018	0.0018	0.0018	0.0018
σ_z S.d. TFP	0.0065	0.0065	0.0065	0.0065	0.0065	0.0065
σ_p S.d. Price MU	0.0016	0.0017	0.0017	0.0017	0.0017	0.0017
σ_w S.d. Wage MU	0.0085	0.0090	0.0093	0.0092	0.0088	0.0078
σ_I S.d. IST	0.0195	0.0129	0.0128	0.0130	0.0128	0.0201
σ_f S.d. Mon. Pol.	0.0021	0.0023	0.0023	0.0023	0.0023	0.0023
σ_g S.d. Gov. Spend.	0.0099	0.0097	0.0097	0.0097	0.0097	0.0097
ϱ_v Autocorr. Risk Prem.	0.9269	0.9407	0.9389	0.9399	0.9407	0.9395
ϱ_z Autocorr. TFP	0.9757	0.9724	0.9708	0.9707	0.9725	0.9742
ϱ_p Autocorr. Price MU	0.9858	0.9752	0.9768	0.9760	0.9750	0.9665
ϑ_p MA comp. Price MU	0.6206	0.6297	0.6108	0.6140	0.6350	0.6012
ϱ_w Autocorr. Wage MU	0.9901	0.9752	0.9757	0.9749	0.9745	0.9599
ϑ_w MA comp. Wage MU	0.9145	0.7965	0.7760	0.7812	0.8118	0.8935
ϱ_I Autocorr. IST	0.1120	0.9035	0.9006	0.9028	0.9003	0.9442
ϱ_f Autocorr. Mon. Pol.	0.2459	0.1879	0.2005	0.1945	0.1859	0.1769
ϱ_g Autocorr. Gov. Spend.	0.9748	0.9752	0.9749	0.9750	0.9756	0.9801
ϱ_{gz} Corr. TFP Gov. Spend.	0.8778	0.8737	0.8734	0.8729	0.8733	0.8710
r Steady state interest rate	0.0103	0.0102	0.0100	0.0101	0.0102	0.0100
α Capital share	0.2325	0.2392	0.2372	0.2388	0.2390	0.2368
ϱ_r Taylor rule inertia	0.8055	0.7834	0.7788	0.7780	0.7850	0.7929
ϕ_π Taylor rule inflation	1.7765	2.1610	2.1742	2.1645	2.1623	2.2446
ϕ_y Taylor rule output growth	0.2070	0.4419	0.4304	0.4380	0.4475	0.4871
ζ_p Price stickiness	0.8144	0.7800	0.7712	0.7721	0.7820	0.7722
ζ_w Nom. wage stickiness	0.8527	0.7600	0.7355	0.7448	0.7696	0.8569
λ Capital utilization	0.6476	0.7603	0.7665	0.7681	0.7581	0.7647
χ Investment adjustment	3.8377	1.7166	1.7268	1.6326	1.8129	1.0256
γ Risk aversion	3.6498	1.9934	$\emptyset$	1.9541	1.9934	2.0640
ρ Disutility of labor	1.9840	1.7330	1.8031	1.7799	1.7497	1.7501
φ Risk cyclicality	-7.4956	$\emptyset$	-4.1108	-4.3133	$\emptyset$	$\emptyset$
σ Weight of individual risk	0.8909	$\emptyset$	0.5426	$\emptyset$	0.6006	$\emptyset$
η Share of HtM agents	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	0.2617
o Cyclicality of HtM income	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	$\emptyset$	1.1116

Note: The RANK specification is obtained by calibrating $\sigma = \varphi = 0$. The TANK specification is the RANK model augmented by a share η of HtM agents. The cyclicality of HtM agents' income is measured by o. The prior distribution for η is a beta distribution with a mean of 0.25 and a standard deviation of 0.1. The prior for o follows a gamma distribution with a mean of 1.5 and a standard deviation of 0.75.

allowing to estimate easily the parameters related to utility and idiosyncratic risk, as well as those affecting the steady state. In contrast, estimating these parameters in the quantitative model of Bayer et al. (2024) is not as easy, as it requires recomputing the stationary equilibrium of the HANK economy for each possible value of the steady-state parameters during estimation—a computationally intensive process. Due to this technical difficulty, they do not estimate the parameters influencing the steady state and utility.[18]

7.2 Sensitivity of the results to key parameters

The precautionary saving motive is likely sensitive to the calibration of the parameter governing risk aversion. Higher risk aversion naturally strengthens the precautionary saving motive, while a lower degree of risk aversion should intuitively weaken it. To test this intuition, I reestimate the model with γ calibrated at 2, compared to its posterior mean of approximately 3.6. As shown in Table 5, the LDD drops to 6767.7. The mean of σ is much lower in this case and the mean of φ indicates less countercyclical risk. These results show that lower risk aversion has a significant effect on the estimates of the parameters governing idiosyncratic risk: the mean of σ and φ are nearly halved. Lower risk aversion dampens the precautionary saving motive, yielding a model that performs only marginally better than its RA counterpart in terms of LDD. This confirms that risk aversion is critical for the interaction between individual risk and precautionary savings. Nevertheless, σ remains well above zero, and φ stays significantly negative, indicating that the precautionary saving motive is not entirely muted.

Taking advantage of the tractability of my model, I also estimate the parameters governing idiosyncratic risk, which are notably high in the baseline, enhancing the precautionary saving motive. Calibrating these parameters to lower values is expected to weaken the motive and worsen the model's fit to the data. First, I calibrate σ at 0.5, compared to its baseline mean of approximately 0.9. In this case, the LDD drops to 6766.1, a significant decrease compared to the baseline. The mean of φ is also lower, but still significantly negative. This confirms that a lower σ dampens the precautionary saving motive as expected. Estimating this parameter is then crucial to enhance the effects of precautionary savings.

Imposing acyclical idiosyncratic risk also diminishes the effects of precautionary savings on consumption. When φ is set to zero, the LDD drops significantly at 6767.1. This result highlights the importance of countercyclical risk, as acyclical risk fails to improve meaningfully the model's fit to the data. The countercyclicality of individual risk is thus a central feature driving the effectiveness of precautionary savings. Through counterfactual analysis, Bilbiie et al. (2023) also find that the countercyclicality of individual risk is crucial to the impact of agent heterogeneity on business cycles.

[18]Similarly, Acharya et al. (2024) find that the model of Smets and Wouters (2007) outperforms that of Bayer et al. (2024) in terms of data fit, further highlighting the advantage of the computational simplicity of log-linear models.

These findings emphasize the importance of estimating the parameters governing idiosyncratic risk to enhance the effects of precautionary savings. In most HANK models, these parameters are typically calibrated, which may explain why relatively few studies have found that HANK models consistently outperform RANK models in terms of LDD.

7.3 MPC heterogeneity

MPC heterogeneity is the other central feature of quantitative HANK models. It arises in incomplete markets frameworks when households face borrowing constraints. As discussed in the introduction, the presence of high-MPC agents brings inequality and redistribution to the forefront, making them key determinants of consumption dynamics alongside precautionary savings.

To assess whether the precautionary savings motive is as important as MPC heterogeneity in shaping business cycles, I estimate a variant of the model that features tractable MPC heterogeneity but excludes idiosyncratic risk. TANK models offer a tractable alternative for capturing the effects of MPC heterogeneity. In this setup, a share η of consumers are HtM agents who consume their entire income. Following Bilbiie (2020), the income of these agents is assumed to be proportional to aggregate income y_t, with a proportionality factor captured by the parameter o.[19]

Table 5 presents the estimation results of the TANK model. The resulting LDD is slightly lower than in the baseline, at 6773.4. This suggests that while MPC heterogeneity plays an important role, it may be somewhat less influential than the precautionary savings motive—though the two are close. Importantly, the TANK model still performs significantly better than the RANK model, indicating that both features are essential components of HANK models.

Appendix A.7 also provides the variance and historical decompositions for the TANK model. These decompositions closely resemble those of the RANK model, but with an even greater role assigned to the IST shock. This result aligns with the findings of Auclert et al. (2020) and Bayer et al. (2024), who show that IST shocks are even more influential in HANK models than in their RANK counterparts. Thus, the shock decomposition in the TANK model is closer to that of a quantitative HANK model, as high-MPC agents tend to overshadow the impact of precautionary savings in these models.

Because the baseline and the TANK model exhibit fundamentally different sources of business cycle fluctuations, a version of the model in which both features—countercyclical risk and MPC heterogeneity—jointly improve the fit to the data does not seem to exist. Estimating a model that includes both features results in a bimodal likelihood, with each mode corresponding to one feature dominating the other. When precautionary savings dominate, the posterior mode of η is low, and IST shocks play a minimal role, as in the baseline model. Conversely, when MPC

[19] The (log-linearized) consumption of HtM agents c_t^h is given by: $c_t^h = oy_t$. Transfers and taxes ensure this proportionality holds, and that there is no inequality in the steady state.

heterogeneity dominates, the posterior mode of σ is low, and IST shocks become the primary driver of fluctuations, as in the RANK and TANK models. Whether this tension reflects a broader structural limitation of HANK models remains an open question, but it represents a promising direction for future research.

8 Conclusion

This paper makes three main contributions. First, it demonstrates that a HANK framework can generate empirically realistic dynamics of individual risk and precautionary savings, and identifies unemployment as the most suitable proxy for capturing the effects of idiosyncratic risk on savings behavior. In doing so, it shows that the dimension of idiosyncratic risk most relevant for business cycles is one that is strongly countercyclical, such as unemployment risk. Second, it establishes that incorporating the interaction between individual risk and precautionary savings substantially improves the model's ability to match aggregate time series. Third, it shows that once capital accumulation and countercyclical risk are jointly considered, precautionary savings can have ambiguous effects on macroeconomic dynamics.

A central takeaway is that a simple heterogeneous-agent model can capture the realistic dynamics of precautionary savings over the business cycle without directly targeting any measure of individual risk. This result underscores that the foundational theory of HANK models aligns with empirical evidence while highlighting the crucial role of the countercyclicality of individual risk. Furthermore, the strong link between individual risk and unemployment risk, as emphasized in this paper, supports the relevance of HANK&SaM models. These models effectively capture the most significant aspects of individual risk according to my results.

Regarding the model's ability to replicate macroeconomic time series, the findings indicate that incorporating countercyclical idiosyncratic risk into a DSGE framework significantly enhances its performance. Notably, the heterogeneous-agent model outperforms its representative-agent counterpart by a meaningful margin. While the effects of precautionary savings may not be large, they are nonetheless impactful, particularly in the context of supply shocks. For these shocks, countercyclical risk generates a short-lived amplification of consumption responses but attenuates output responses due to an even greater attenuation of investment responses.

A Appendix

A.1 Data description

The dataset used for estimation is the same as in Bayer et al. (2024).[20] All elements about the observables are given in the appendix of their paper. The original series span from 1954Q3 to 2019Q4 and are sourced from the St.Louis FED – FRED. Output, consumption and investment are expressed per capita and in real terms. Below, I detail how each observable is tailored for estimation[21] and how the empirical measures of idiosyncratic risk are computed:

- Output growth: $\Delta \log(Y_t) - \overline{\Delta \log(Y_t)}$
- Consumption growth: $\Delta \log(C_t) - \overline{\Delta \log(C_t)}$
- Investment growth: $\Delta \log(I_t) - \overline{\Delta \log(I_t)}$
- Real wage growth: $\Delta \log(\mathcal{W}_t) - \overline{\Delta \log(\mathcal{W}_t)}$
- Nominal interest rate: $\log(1 + i_t) - \overline{\log(1 + i_t)}$
- Hours worked: $\log(L_t) - \overline{\log(L_t)}$
- Inflation: $\log(\Pi_t) - \overline{\log(\Pi_t)}$
- Income idiosyncratic risk in Figure 5: estimated time series for the variance of idiosyncratic income from Bayer et al. (2019), based on the Survey of Income and Program Participation. Available from 1983Q1 to 2013Q1. Used in Bayer et al. (2024) as an observable.
- Unemployment rate: quarterly average of the series sourced from FRED, from 1954Q4 to 2019Q4 (FRED mnemonics: "UNRATE").
- Job-finding and job-loss rates: "As in Challe et al. (2017) and Cho (2023), I follow the method of Shimer (2005) and Shimer (2012) to compute the job-finding and job-loss rates. The monthly transition matrices across employment statuses are constructed using seasonally adjusted data on employment (BLS series LNS12000000), unemployment (BLS series LNS13000000), and short-run unemployment (BLS series LNS13008396). Then, these matrices are multiplied over the three consecutive months of each quarter to obtain quarterly transition probabilities. The final series are detrended using the method of Hamilton (2018)." (the dataset spans from 1954Q4 to 2019Q4).

[20]The dataset can be found at: https://github.com/BASEforHANK/HANK_BusinessCycleAndInequality/blob/master/src/bbl_data_inequality.csv.

[21]Δ denotes the temporal difference operator and $\overline{N}$ denotes the average of variable N.

- Consumer sentiment (FRED mnemonics: "UMCSENT"): "Surveys of Consumers, University of Michigan, University of Michigan: Consumer Sentiment, retrieved from FRED". From 1959Q4 to 2019Q4. Index 2017Q3 = 100. Not seasonally adjusted.

A.2 Modeling details

A.2.1 Final goods producers

The final goods sector is assumed to be perfectly competitive. A firm combines intermediate goods to produce final goods using the following CES function:

$$Y_t = \left(\int_0^1 Y_t(j)^{\frac{\theta_t - 1}{\theta_t}} dj \right)^{\frac{\theta_t}{\theta_t - 1}}$$

where θ_t is the elasticity of substitution across intermediate goods. This elasticity is not constant because the price mark-up is subject to stochastic shocks. I denote $\mathcal{M}_t^p$, the price mark-up that is such that: $\mathcal{M}_t^p = \frac{\theta_t}{\theta_t - 1}$.

The final goods firm pays a price $P_t(j)$ for every $Y_t(j)$ that it buys to produce the final goods. The final goods firm takes the price of the final goods P_t as given since it evolves in a perfectly competitive market. Cost-minimization together with the zero-profit condition imply that:

$$Y_t(j) = \left(\frac{P_t(j)}{P_t} \right)^{-\theta_t} Y_t$$

where $P_t = \left(\int_0^1 P_t(j)^{1-\theta_t} dj \right)^{\frac{1}{1-\theta_t}}$.

A.2.2 Intermediate goods producers

There is a unit continuum of producers of intermediate goods. These firms compete in monopolistic competition à la Dixit and Stiglitz (1977). Each firm j is assumed to have the same production technology. Then, firm j's production function is assumed to take the following form:

$$Y_t(j) = Z_t L_t(j)^{1-\alpha} K_t(j)^{\alpha}$$

where $K_t(j)$ is capital used, $L_t(j)$ is effective labor and Z_t is the Total Factor Productivity (TFP) subject to aggregate shocks. The TFP is assumed to follow the following AR(1) process:

$$\log Z_t \equiv z_t = \varrho_z z_{t-1} + \sigma_z \epsilon_t^z$$

where ϱ_z is the persistence parameter and σ_z is the standard deviation of the shock.

Intermediate firms have to pay households for their labor supplies and they rent capital from capital owners. Under monopolistic competition, the real profit of a firm j $D_t(j)$ can be written such that:

$$D_t(j) = \frac{P_t(j)}{P_t} Y_t(j) - \mathcal{W}_t L_t(j) - R_t^K K_t(j)$$

It is assumed that the intermediate goods producers hire capital and labor from a common market. As a result, all firms share the same capital-to-labor ratio (i.e., $\frac{K_t(j)}{L_t(j)} = \frac{K_t}{L_t}$). This implies that the optimal demand for production factors is identical across all intermediate goods producers. Cost minimization results in:

$$\mathcal{W}_t = S_t (1-\alpha) Z_t \left(\frac{K_t}{L_t} \right)^{\alpha} \tag{A.1}$$

$$R_t^K = S_t \alpha Z_t \left(\frac{L_t}{K_t} \right)^{1-\alpha}$$

where S_t is the real marginal cost. Let's denote MPL_t, the marginal productivity of labor which is defined as:

$$MPL_t \equiv \frac{\partial Y_t}{\partial L_t} = \frac{(1-\alpha) Y_t}{L_t}$$

From (A.1), one obtains:

$$S_t = \frac{\mathcal{W}_t}{MPL_t}$$

Intermediate firms possess a degree of market power due to consumers' preference for variety. Under monopolistic competition à la Dixit and Stiglitz (1977), a firm j can set a price $P_t(j)$ above the market average price P_t while still facing positive demand. Consequently, firm j seeks to determine the optimal price $P_t(j)$ that maximizes its profit $D_t(j)$.

Since all firms share the same production technology and marginal cost, the profit-maximizing price is identical across firms and is denoted as P_t^*. However, I assume the presence of price stickiness à la Calvo (1983). Specifically, a fraction ζ_p of firms cannot adjust their prices in period t, while the remaining fraction $(1-\zeta_p)$ can reset their prices and set them at the optimal level P_t^*. As a result, a fraction ζ_p of firms must retain the prices they last set when they had the opportunity to adjust.

Under this pricing mechanism, the probability that a given price remains unchanged for one period is $(1-\zeta_p)\zeta_p$, for two periods is $(1-\zeta_p)\zeta_p^2$, and for k periods is $(1-\zeta_p)\zeta_p^k$. The resulting equation governing the dynamics of the aggregate price level under sticky prices is then given by:

$$P_t = \left[\zeta_p P_{t-1}^{1-\theta_t} + (1-\zeta_p) {P_t^*}^{1-\theta_t} \right]^{\frac{1}{1-\theta_t}}$$

or equivalently:

$$\Pi_t^{1-\theta_t} = \zeta_p + (1-\zeta_p) \left(\frac{P_t^*}{P_{t-1}} \right)^{1-\theta_t} \tag{A.2}$$

where Π_t is inflation in level, that is defined such that: $\Pi_t = \frac{P_t}{P_{t-1}}$.

The profit maximization program of a firm j is:

$$\begin{aligned} P_t^* &= \arg\max_{P_t(j)} \mathbb{E}_t \sum_{k=0}^{\infty} \zeta_p^k \mathbb{Q}_{t,t+k} Y_{t+k}(j) \left[\frac{P_t(j)}{P_{t+k}} - S_{t+k} \right] \\ s.t. \quad Y_{t+k}(j) &= \left(\frac{P_t(j)}{P_{t+k}} \right)^{-\theta_{t+k}} Y_{t+k} \end{aligned}$$

where $\mathbb{Q}_{t,t+k} = \prod_{s=t}^{t+k-1} \frac{1}{\tilde{R}_s}$ is the stochastic discount factor as in Acharya and Dogra (2020).

The FOC of this program is:

$$\mathbb{E}_t \sum_{k=0}^{\infty} \zeta_p^k \mathbb{Q}_{t,t+k} Y_{t+k}(j) \left(\frac{P_t^*}{P_{t+k}} - \mathcal{M}_{t+k}^p S_{t+k} \right) = 0 \tag{A.3}$$

Log-linearizing (A.2) and (A.3) while assuming that inflation is null at steady state and combining these two expressions yields the New Keynesian Phillips Curve:

$$\pi_t = \tilde{\beta} \mathbb{E}_t \pi_{t+1} + \kappa_p s_t + m_t^p$$

where $\tilde{\beta} = \frac{1}{R}$, $\kappa_p = \frac{(1-\zeta_p)(1-\tilde{\beta}\zeta_p)}{\zeta_p}$ and m_t^p is the renormalized version of $\log\left(\frac{\mathcal{M}_t^p}{\mathcal{M}^p}\right)$. As in Smets and Wouters (2007), I assume that m_t^p follows an ARMA(1,1) process such that:

$$m_t^p = \varrho_p m_{t-1}^p + \sigma_p \epsilon_t^p - \vartheta_p \epsilon_{t-1}^p$$

where ϱ_p is the persistence parameter, σ_p the standard deviation of the shock and ϑ_p measures the weight of the MA component. At steady state, the price mark-up returns to its initial value: $\mathcal{M}^p = \frac{\theta}{\theta-1}$.

A.2.3 Labor packers and unions

The labor market is assumed to be affected by real and nominal rigidities as the goods market. For expositional purposes, I decompose the utility function of an agent i such that:

$$\mathbb{U}_t(i) = u(C_t(i)) + h(L_t(i))$$

where $u(C_t(i)) = -\frac{1}{\gamma} e^{-\gamma C_t(i)}$ and $h(L_t(i)) = -\rho e^{\frac{1}{\rho}\left(L_t(i) - \bar{\xi}\right)}$. The marginal rate of substitution between labor and consumption of an agent i is:

$$MRS_t(i) = \left| \frac{h'(L_t(i))}{u'(C_t(i))} \right|$$

Following Erceg et al. (2000), each household supplies a differentiated labor service, which grants them market power to negotiate higher wages. The labor market operates under monopolistic competition à la Dixit and Stiglitz (1977). Additionally, labor packers aggregate these differentiated labor supplies using the following CES function:

$$L_t = \left(\int_0^1 L_t(i)^{\frac{\nu_t - 1}{\nu_t}} di \right)^{\frac{\nu_t}{\nu_t - 1}}$$

where ν_t is the elasticity of substitution across differentiated labor services. Similarly to the price mark-up, this elasticity is not constant because it is assumed that the wage mark-up is subject to stochastic shocks. The wage mark-up is denoted such that: $\mathcal{M}_t^w = \frac{\nu_t}{\nu_t - 1}$.

I assume that labor packers evolve in a perfectly competitive market so that they make no profits. Then, the mission of unions is to sell households' labor supplies to the labor packers. Similarly to the final goods producers, labor packers buy a differentiated labor service $L_t(i)$ for a wage $W_t(i)$ and they take the average nominal wage W_t as given. Cost minimization coupled with the zero-profit condition yields the following demand function:

$$L_t(i) = \left(\frac{W_t(i)}{W_t} \right)^{-\nu_t} L_t$$

where $W_t = \left(\int_0^1 W_t(i)^{1-\nu_t} di \right)^{\frac{1}{1-\nu_t}}$.

As in Erceg et al. (2000), unions negotiate wages, but each union cannot renegotiate in every period, following the Calvo (1983) framework. Specifically, each union has a probability ζ_w of not being able to renegotiate in period t and a probability $(1 - \zeta_w)$ of being able to do so. Due to the presence of wage stickiness, nominal wages are not always set at their optimal level W_t^* in every period. This leads to nominal wage inflation in the model, which is defined as follows:

$$\Pi_t^w = \frac{W_t}{W_{t-1}}$$

or equivalently:

$$\frac{\mathcal{W}_t}{\mathcal{W}_{t-1}} = \frac{\Pi_t^w}{\Pi_t}$$

where $\Pi^w = 1$ since Π is assumed to be equal to unity. Nominal wages evolve in a similar fashion as prices do, that is:

$$W_t = \left[\zeta_w W_{t-1}^{1-\nu_t} + (1 - \zeta_w) W_t^{*1-\nu_t} \right]^{\frac{1}{1-\nu_t}} \tag{A.4}$$

Since unions cannot renegotiate wages in every period, they must account for the possibility that they may not be able to do so in the future when formulating their maximization program, following a similar logic to that of intermediate firms. Given agent heterogeneity, I follow Auclert et al. (2021) in assuming that unions base their decisions on the preferences of a representative middleman who consumes aggregate consumption C_t while negotiating the optimal wage W_t^*. The unions' maximization problem is then given by:

$$
\begin{aligned}
W_t^* &= \arg\max_{W_t(i)} \mathbb{E}_t \sum_{k=0}^{\infty} (\tilde{\beta}\zeta_w)^k \left[u'(C_{t+k}) \frac{W_t(i)}{P_{t+k}} L_{t+k}(i) - h(L_{t+k}(i)) \right] \\
s.t. \quad L_{t+k}(i) &= \left(\frac{W_t(i)}{W_{t+k}} \right)^{-\nu_t} L_{t+k}
\end{aligned}
$$

The FOC of the previous program is:

$$
\mathbb{E}_t \sum_{k=0}^{\infty} (\tilde{\beta}\zeta_w)^k u'(C_{t+k}) L_{t+k}(i) \left[\frac{W_t^*}{P_{t+k}} - \mathcal{M}_{t+k}^w \frac{h'(L_{t+k}(i))}{u'(C_{t+k})} \right] = 0 \tag{A.5}
$$

Log-linearizing[22] (A.4) and (A.5) and combining these two expressions, one obtains the New Keynesian Wage Phillips Curve:

$$
\pi_t^w = \tilde{\beta} \mathbb{E}_t \pi_{t+1}^w - \kappa_w (\omega_t - mrs_t) + m_t^w
$$

where $\kappa_w = \frac{(1-\zeta_w)(1-\tilde{\beta}\zeta_w)}{\zeta_w}$ and m_t^w is the renormalized version of $\log\left(\frac{\mathcal{M}_t^w}{\mathcal{M}^w}\right)$. As in Smets and Wouters (2007), I assume that m_t^w follows an ARMA(1,1) process such that:

$$
m_t^w = \varrho_w m_{t-1}^w + \sigma_w \epsilon_t^w - \vartheta_w \epsilon_{t-1}^w
$$

where ϱ_w is the persistence parameter, σ_w the standard deviation of the shock and ϑ_w measures the weight of the MA component. At steady state, the wage mark-up returns to its initial value: $\mathcal{M}^w = \frac{\nu}{\nu - 1}$.

A.2.4 Capital firms and investment

Following Auclert et al. (2020), capital firms own the capital stock and they rent it to intermediate firms. In exchange for this service, they receive $R_t^K K_t$. Capital firms have then to decide how much capital to own, how much to invest in new capital and what utilization rate is optimal. The profit of capital firms D_t^K is:

$$
D_t^K = R_t^K K_t - a(u_t) \bar{K}_{t-1} - I_t
$$

Capital firms are in charge of turning final goods into new capital using an investment-specific technology. As it is standard in the DSGE literature, I assume that this technology is affected by adjustment costs. I model these adjustment costs as in Christiano et al. (2005), relying on a function S satisfying $S(1) = S'(1) = 0$ and $\chi \equiv S''(1) > 0$. Hence, capital firms create new capital

[22] The log-linearized real wage is denoted such that: $\omega_t \equiv \frac{\mathcal{W}_t - \mathcal{W}}{\mathcal{W}} \approx \log\left(\frac{\mathcal{W}_t}{\mathcal{W}}\right)$.

according to the following law of motion:

$$\bar{K}_t = (1-\delta)\bar{K}_{t-1} + \Psi_t I_t \left(1 - S\left(\frac{I_t}{I_{t-1}}\right)\right) \tag{A.6}$$

where Ψ_t is an investment-specific technology shock. As in Smets and Wouters (2007), this shock is renormalized in the log-linear model.[23] I assume that it follows the following AR(1) process:

$$\hat{\psi}_t = \varrho_I \hat{\psi}_{t-1} + \sigma_I \epsilon_t^I$$

where $\hat{\psi}_t$ is the renormalized shock in the log-linear model, ϱ_I is the persistence parameter and σ_I the standard deviation of the shock.

All the capital stock might not be used in the production function. This depends on the degree of capital utilization, as stated by the following equation:

$$K_t = u_t \bar{K}_{t-1} \tag{A.7}$$

Equation (A.7) describes how the degree of capital utilization u_t affects the capital ultimately used in the production function. In steady state, all the capital stock is used, so that $u = 1$. $a(u_t)$ is the cost associated to the degree of capital utilization. Following Smets and Wouters (2007), I assume that $a(1) = 0$ and denote $\frac{a''(1)}{a'(1)} \equiv \frac{\lambda}{1-\lambda}$.

Similarly to intermediate firms, I assume that capital firms use $\mathbb{Q}_{t,t+k}$ to discount the future. The maximization program of capital firms is then:

$$\begin{aligned} &\max_{\bar{K}_t, I_t, u_t} \mathbb{E}_t \sum_{k=0}^{\infty} \mathbb{Q}_{t,t+k} D_{t+k}^K \\ &s.t. \quad \bar{K}_t = (1-\delta)\bar{K}_{t-1} + \Psi_t I_t \left(1 - S\left(\frac{I_t}{I_{t-1}}\right)\right) \\ &\qquad\; K_t = u_t \bar{K}_{t-1} \end{aligned} \tag{A.8}$$

Differentiating (A.8) with respect to $\bar{K}_t$, it follows that:

$$\tilde{R}_t Q_t = \mathbb{E}_t R_{t+1}^K u_{t+1} - \mathbb{E}_t a(u_{t+1}) + (1-\delta)\mathbb{E}_t Q_{t+1} \tag{A.9}$$

where Q_t is the Lagrange multiplier (or the capital price) associated to (A.6). Differentiating (A.8) with respect to I_t, one finds:

$$Q_t \Psi_t \left[1 - S\left(\frac{I_t}{I_{t-1}}\right) - \frac{I_t}{I_{t-1}} S'\left(\frac{I_t}{I_{t-1}}\right)\right] = 1 - \mathbb{E}_t \frac{Q_{t+1}}{\tilde{R}_t} \Psi_{t+1} \left(\frac{I_{t+1}}{I_t}\right)^2 S'\left(\frac{I_{t+1}}{I_t}\right) \tag{A.10}$$

[23] In the log-linear model, this implies: $\hat{\psi}_t \equiv \frac{1}{\chi(1+\tilde{\beta})}\psi_t$, where $\log \Psi_t \equiv \psi_t$.

Differentiating (A.8) with respect to u_t, one obtains:

$$R_t^K = a'(u_t)$$

A.2.5 Derivation of aggregate consumption

Demeaned cash-on-hand is defined as:

$$x_t(i) = \tilde{R}_{t-1}(A_{t-1}(i) - A_{t-1}) + \mathcal{Y}_t(i) - \mathcal{Y}_t$$

Let $C_t(x_t(i))$ denote the consumption of household i as a function of her demeaned cash-on-hand $x_t(i)$. By iterating forward the previous equation and incorporating household i's budget constraint, the demeaned cash-on-hand in period $t+1$ can be expressed as:

$$x_{t+1}(i) = \tilde{R}_t \left(x_t(i) - C_t(x_t(i)) + C_t\right) + \mathcal{Y}_{t+1}(i) - \mathcal{Y}_{t+1}$$

As $\mathcal{Y}_{t+1}(i)$ is normally distributed, we have $x_{t+1}(i) \sim \mathcal{N}_t(\tilde{R}_t \left(x_t(i) - C_t(x_t(i)) + C_t\right), \sigma_{t+1}^2)$. Combining the expression of the demeaned cash-on-hand in $t+1$ with (2), one obtains:

$$e^{-\gamma C_t(x_t(i))} = \beta \tilde{R}_t \mathbb{E}_t \int e^{-\gamma C_{t+1}(x_{t+1}(i))} \left\{ \frac{1}{\sqrt{2\pi}\sigma_{t+1}} e^{-\frac{1}{2}\left(\frac{x_{t+1}(i) - \tilde{R}_t[x_t(i) - C_t(x_t(i)) + C_t]}{\sigma_{t+1}}\right)^2} \right\} dx_{t+1}(i)$$

As shown in Acharya and Dogra (2020), individual consumption can be expressed in the absence of aggregate shocks as:

$$C(x) = C + \mu x \quad \forall x$$

where $\mu = 1 - \frac{1}{R}$. Let's define $\tilde{r}_t \equiv \log \frac{\tilde{R}_t}{R}$, $\hat{C}_t(x_t(i)) \equiv C_t(x_t(i)) - C(x)$ and $\hat{Y}_t \equiv Y_t - Y$. As discussed in the main text, σ_t is assumed to covary with output, such that: $\sigma_t = \tilde{\sigma} e^{\varphi \hat{Y}_t}$. Rewriting the Euler equation as a function of $\hat{C}_t(x_t(i))$, one gets:

$$\begin{aligned} e^{-\gamma\left[C + \mu x + \hat{C}_t(x_t(i))\right]} = {} & \beta R e^{\tilde{r}_t} \mathbb{E}_t \frac{1}{\tilde{\sigma} e^{\varphi \mathbb{E}_t \hat{Y}_{t+1}}} \\ & \times \int e^{-\gamma\left[C + \mu x_{t+1}(i) + \hat{C}_{t+1}(x_{t+1}(i))\right]} \left\{ \frac{1}{\sqrt{2\pi}} e^{-\frac{1}{2}\left(\frac{x_{t+1}(i)}{\tilde{\sigma} e^{\varphi \mathbb{E}_t \hat{Y}_{t+1}}} - \frac{R e^{\tilde{r}_t}\left[(1-\mu)x - \hat{C}_t(x_t(i)) + \hat{C}_t\right]}{\tilde{\sigma} e^{\varphi \mathbb{E}_t \hat{Y}_{t+1}}}\right)^2} \right\} dx_{t+1}(i) \end{aligned}$$

When there is only idiosyncratic risk, the previous equation reduces to:

$$1 = \frac{\beta R}{\tilde{\sigma}} \int e^{-\gamma\mu(x_{t+1}(i) - x)} f\left(\frac{x_{t+1}(i) - x}{\tilde{\sigma}}\right) dx_{t+1}(i)$$

where $f\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right) = \frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)}$. Solving the integral yields the following expression for the steady state interest rate:

$$R = \frac{1}{\beta}e^{-\frac{\gamma^2\mu^2\tilde{\sigma}^2}{2}}$$

A first-order Taylor expansion of the individual Euler equation around the steady state yields:

$$\begin{aligned}
0 = &\underbrace{\left\{\frac{\beta R}{\tilde{\sigma}}\int e^{-\gamma\mu(x_{t+1}(i)-x)}f\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)dx_{t+1}(i)\right\}}_{=1}\tilde{r}_t \\
&+\underbrace{\left\{\frac{\beta R}{\tilde{\sigma}}\int e^{-\gamma\mu(x_{t+1}(i)-x)}f\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)dx_{t+1}(i)\right\}}_{=1}(-\varphi)\,\mathbb{E}_t\hat{Y}_{t+1} \\
&+\underbrace{\left\{\frac{\beta R}{\tilde{\sigma}}\int e^{-\gamma\mu(x_{t+1}(i)-x)}f\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)dx_{t+1}(i)\right\}}_{=1}\gamma\hat{C}_t\left(x_t(i)\right) \\
&+\underbrace{\frac{\beta R}{\tilde{\sigma}}\left\{(-\gamma)\int e^{-\gamma\mu(x_{t+1}(i)-x)}f\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)\mathbb{E}_t\left[\hat{C}_{t+1}\left(x_{t+1}(i)\right)\right]dx_{t+1}(i)\right\}}_{\equiv\varsigma} \\
&+\frac{\beta R}{\tilde{\sigma}}\underbrace{\int e^{-\gamma\mu(x_{t+1}(i)-x)}f'\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)dx_{t+1}(i)}_{=-\tilde{\sigma}e^{\frac{\gamma^2\mu^2\tilde{\sigma}^2}{2}}(1+\gamma^2\mu^2\tilde{\sigma}^2)}(-\varphi)\,\mathbb{E}_t\hat{Y}_{t+1} \\
&-\beta\left(\frac{R}{\tilde{\sigma}}\right)^2\underbrace{\int e^{-\gamma\mu(x_{t+1}(i)-x)}f'\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)dx_{t+1}(i)}_{=\gamma\mu\tilde{\sigma}^2e^{\frac{\gamma^2\mu^2\tilde{\sigma}^2}{2}}}\left(\hat{C}_t-\hat{C}_t\left(x_t(i)\right)\right) \\
&+\frac{\beta R}{\tilde{\sigma}}\underbrace{\int e^{-\gamma\mu(x_{t+1}(i)-x)}f'\left(\frac{x_{t+1}(i)-x}{\tilde{\sigma}}\right)dx_{t+1}(i)}_{=\gamma\mu\tilde{\sigma}^2e^{\frac{\gamma^2\mu^2\tilde{\sigma}^2}{2}}}\left(-\frac{x}{\tilde{\sigma}}\right)\tilde{r}_t
\end{aligned}$$

To simplify the Euler equation further, one needs to specify the form of households' policy functions (otherwise there is no simpler formulation for ς). I guess and verify that individual consumptions take the form:

$$\hat{C}_t(x) = \hat{C}_t + \mu\hat{\mu}_t x \quad \forall x$$

where $\hat{\mu}_t = \frac{\mu_t-\mu}{\mu}$. Substituting this guess in the previous equation, one obtains an expression for the aggregate Euler equation:

$$\begin{aligned}
0 = &\tilde{r}_t - \varphi\mathbb{E}_t\hat{Y}_{t+1} + \gamma\hat{C}_t + \gamma\mu\hat{\mu}_t x - \gamma\mathbb{E}_t\hat{C}_{t+1} - \gamma\mu\left(\mathbb{E}_t\hat{\mu}_{t+1}\right)x + \gamma^2\mu^2\tilde{\sigma}^2\mathbb{E}_t\hat{\mu}_{t+1} \\
&-\left(1+\gamma^2\mu^2\tilde{\sigma}^2\right)(-\varphi)\,\mathbb{E}_t\hat{Y}_{t+1} + R\gamma\mu^2\hat{\mu}_t x - \gamma\mu x\tilde{r}_t
\end{aligned}$$

Matching coefficients in x, one finds:

$$\hat{\mu}_t = \frac{1}{R}\left(\tilde{r}_t + \mathbb{E}_t\hat{\mu}_{t+1}\right)$$

that is (4) in the main text. Combining the two previous equations, the aggregate Euler equation boils down to:

$$0 = \tilde{r}_t - \varphi\mathbb{E}_t\hat{Y}_{t+1} + \gamma\hat{C}_t - \gamma\mathbb{E}_t\hat{C}_{t+1} + \gamma^2\mu^2\tilde{\sigma}^2\mathbb{E}_t\hat{\mu}_{t+1} - \left(1 + \gamma^2\mu^2\tilde{\sigma}^2\right)(-\varphi)\,\mathbb{E}_t\hat{Y}_{t+1}$$

Let's define the log-linearized variables $c_t \equiv \frac{\hat{C}_t}{C}$ and $y_t \equiv \frac{\hat{Y}_t}{Y}$. Rewriting the aggregate Euler equation as a function of these log-linearized variables, one finally gets:

$$Cc_t = C\mathbb{E}_tc_{t+1} - \frac{1}{\gamma}\tilde{r}_t - \gamma\mu^2\tilde{\sigma}^2\mathbb{E}_t\hat{\mu}_{t+1} - \varphi\gamma\mu^2\tilde{\sigma}^2Y\mathbb{E}_ty_{t+1}$$

that is (3) in the main text.

A.2.6 Derivation of investment as a function of the wedge

Log-linearizing the Euler equation for capital (A.9) yields the following expression:

$$q_t = \tilde{\beta}R^K\mathbb{E}_tr^K_{t+1} + \tilde{\beta}(1-\delta)\mathbb{E}_tq_{t+1} - \tilde{r}_t$$

Isolating $\tilde{r}_t$ in (3), the aggregate Euler equation for actuarial bonds becomes:

$$\tilde{r}_t = \gamma C(\mathbb{E}_tc_{t+1} - c_t) - \gamma\left(\gamma\mu^2\tilde{\sigma}^2\mathbb{E}_t\hat{\mu}_{t+1} + \varphi\gamma\mu^2\tilde{\sigma}^2Y\mathbb{E}_ty_{t+1}\right)$$

Substituting this into the equation for q_t gives:

$$q_t = \tilde{\beta}R^K\mathbb{E}_tr^K_{t+1} + \tilde{\beta}(1-\delta)\mathbb{E}_tq_{t+1} - \gamma C(\mathbb{E}_tc_{t+1} - c_t) + \gamma\left(\gamma\mu^2\tilde{\sigma}^2\mathbb{E}_t\hat{\mu}_{t+1} + \varphi\gamma\mu^2\tilde{\sigma}^2Y\mathbb{E}_ty_{t+1}\right)$$

Iterating forward the previous expression, one obtains:

$$q_t = \mathbb{E}_t\sum_{s=0}^{\infty}\tilde{\beta}^s(1-\delta)^s\left[\tilde{\beta}R^Kr^K_{t+1+s} - \gamma C(c_{t+1+s} - c_{t+s}) + \gamma\left(\gamma\mu^2\tilde{\sigma}^2\hat{\mu}_{t+1+s} + \varphi\gamma\mu^2\tilde{\sigma}^2Y\mathbb{E}_ty_{t+1+s}\right)\right]$$

The previous expression illustrates well how the current price of capital is influenced by the expected path of future wedges. This is crucial to understand how precautionary savings shape investment. Log-linearizing (A.10) provides the investment equation:

$$\iota_t = \frac{1}{1+\tilde{\beta}}\iota_{t-1} + \frac{\tilde{\beta}}{1+\tilde{\beta}}\mathbb{E}_t\iota_{t+1} + \frac{1}{\chi(1+\tilde{\beta})}q_t + \hat{\psi}_t$$

where $\iota_t \equiv \frac{I_t - I}{I} \approx \log\left(\frac{I_t}{I}\right)$ and $\hat{\psi}_t \equiv \frac{1}{\chi(1+\tilde{\beta})}\psi_t$. Substituting the forward-looking expression for q_t into the investment equation yields:

$$\iota_t = \frac{1}{1+\tilde{\beta}}\iota_{t-1} + \frac{\tilde{\beta}}{1+\tilde{\beta}}\mathbb{E}_t\iota_{t+1} + \hat{\psi}_t$$
$$+ \frac{1}{\chi(1+\tilde{\beta})}\mathbb{E}_t\sum_{s=0}^{\infty}\tilde{\beta}^s(1-\delta)^s\left[\tilde{\beta}R^K r^K_{t+1+s} - \gamma C(c_{t+1+s} - c_{t+s}) + \gamma\left(\gamma\mu^2\tilde{\sigma}^2\hat{\mu}_{t+1+s} + \varphi\gamma\mu^2\tilde{\sigma}^2 Y y_{t+1+s}\right)\right]$$

Putting γ in factor of the whole sum, one finally obtains:

$$\iota_t = \frac{1}{1+\tilde{\beta}}\iota_{t-1} + \frac{\tilde{\beta}}{1+\tilde{\beta}}\mathbb{E}_t\iota_{t+1} + \hat{\psi}_t$$
$$+ \frac{\gamma}{\chi(1+\tilde{\beta})}\mathbb{E}_t\sum_{s=0}^{\infty}\tilde{\beta}^s(1-\delta)^s\left[\frac{\tilde{\beta}}{\gamma}R^K r^K_{t+1+s} - C(c_{t+1+s} - c_{t+s}) + \gamma\mu^2\tilde{\sigma}^2\hat{\mu}_{t+1+s} + \varphi\gamma\mu^2\tilde{\sigma}^2 Y y_{t+1+s}\right]$$

A.3 Summary of the main equations

A.3.1 Consumption

$$e^{-\gamma C_t(x_t(i))} = \beta\mathcal{Q}_t\mathbb{E}_t\int\left(\sqrt{2\pi}\sigma_{t+1}\right)^{-1} e^{-\gamma C_{t+1}(x_{t+1}(i)) - \frac{1}{2}\left(\frac{\mathcal{Y}_{t+1}(i)-\mathcal{Y}_t}{\sigma_{t+1}}\right)^2} dx_{t+1}(i)$$

$$\tilde{R}_t = \frac{V_t(1+i_t)}{\mathbb{E}_t\Pi_{t+1}}$$

$$\sigma_t^2 = (\sigma Y)^2 e^{2\varphi(Y_t - Y)}$$

A.3.2 Labor market

$$W_t = \left[\zeta_w W_{t-1}^{1-\nu_t} + (1-\zeta_w)W_t^{*1-\nu_t}\right]^{\frac{1}{1-\nu_t}}$$

$$\mathbb{E}_t\sum_{k=0}^{\infty}(\tilde{\beta}\zeta_w)^k u'(C_{t+k})L_{t+k}(i)\left[\frac{W_t^*}{P_{t+k}} - \mathcal{M}^w_{t+k}\frac{h'(L_{t+k}(i))}{u'(C_{t+k})}\right] = 0$$

$$MRS_t = e^{\gamma C_t + \frac{L_t - \bar{\xi}}{\rho}}$$

$$\frac{\mathcal{W}_t}{\mathcal{W}_{t-1}} = \frac{\Pi^w_t}{\Pi_t}$$

A.3.3 Capital utilization

$$K_t = u_t\bar{K}_{t-1}$$

$$R^K_t = a'(u_t)$$

A.3.4 Capital and investment

$$\tilde{R}_t Q_t = \mathbb{E}_t R^K_{t+1} u_{t+1} - \mathbb{E}_t a(u_{t+1}) + (1-\delta)\mathbb{E}_t Q_{t+1}$$

$$Q_t \Psi_t \left[1 - S\left(\frac{I_t}{I_{t-1}}\right) - \frac{I_t}{I_{t-1}} S'\left(\frac{I_t}{I_{t-1}}\right)\right] = 1 - \mathbb{E}_t \frac{Q_{t+1}}{\tilde{R}_t} \Psi_{t+1} \left(\frac{I_{t+1}}{I_t}\right)^2 S'\left(\frac{I_{t+1}}{I_t}\right)$$

$$\bar{K}_t = (1-\delta)\bar{K}_{t-1} + \Psi_t I_t \left(1 - S\left(\frac{I_t}{I_{t-1}}\right)\right)$$

A.3.5 Production and market clearing

$$Y_t = Z_t L_t^{1-\alpha} K_t^{\alpha}$$

$$C_t + I_t + G_t + a(u_t)\bar{K}_{t-1} = Y_t$$

A.3.6 Production factors

$$\mathcal{W}_t = S_t MPL_t$$

$$MPL_t = \frac{(1-\alpha)Y_t}{L_t}$$

$$R^K_t = S_t \alpha Z_t \left(\frac{L_t}{K_t}\right)^{1-\alpha}$$

A.3.7 Price dynamics and inflation

$$\Pi_t^{1-\theta_t} = \zeta_p + (1-\zeta_p)\left(\frac{P^*_t}{P_{t-1}}\right)^{1-\theta_t}$$

$$\mathbb{E}_t \sum_{k=0}^{\infty} \zeta_p^k \mathbb{Q}_{t,t+k} Y_{t+k}(j) \left(\frac{P^*_t}{P_{t+k}} - \mathcal{M}^p_{t+k} S_{t+k}\right) = 0$$

A.3.8 Monetary policy

$$1 + i_t = R^{1-\varrho_r}(1+i_{t-1})^{\varrho_r} \Pi_t^{(1-\varrho_r)\phi_\pi} \left(\frac{Y_t}{Y_{t-1}}\right)^{(1-\varrho_r)\phi_y} F_t$$

A.3.9 Fiscal policy

$$G_t = \Gamma_t G$$

A.4 Steady state of the model

A.4.1 Consumption

$$R = \frac{1}{\tilde{\beta}}$$

$$\mu = \frac{R-1}{R}$$

$$C = Y - I - G$$

A.4.2 Labor market

$$MRS = \frac{\nu - 1}{\nu}\mathcal{W}$$

$$L = 1$$

$$\Pi^w = 1$$

A.4.3 Capital utilization

$$u = 1$$

$$a(u) = 0$$

A.4.4 Capital and investment

$$R^K = \frac{1}{\tilde{\beta}} - (1 - \delta)$$

$$Q = 1$$

$$I = \delta K$$

A.4.5 Production and market clearing

$$Z = 1$$

$$Y = K^\alpha$$

A.4.6 Production factors

$$S = \frac{\theta - 1}{\theta}$$

$$MPL = (1 - \alpha)Y$$

$$\mathcal{W} = SMPL$$

$$K = \left(\frac{\alpha S}{R^K}\right)^{\frac{1}{1-\alpha}}$$

A.4.7 Price dynamics and inflation

$$\Pi = 1$$

$$\mathcal{M}^p = \frac{\theta}{\theta - 1}$$

A.4.8 Monetary policy

$$1 + i = R$$

A.4.9 Fiscal policy

$$G = \frac{Y}{5}$$

A.5 Log-linear model

The aggregate marginal propensity to consume μ_t is denoted in log-linear terms such that:

$$\hat{\mu}_t \equiv \frac{\mu_t - \mu}{\mu} \approx \log\left(\frac{\mu_t}{\mu}\right)$$

Likewise, capital utilization u_t is denoted in log-linear terms such that:

$$\hat{u}_t \equiv \frac{u_t - u}{u} \approx \log\left(\frac{u_t}{u}\right)$$

Log-linearized investment ι_t is such that:

$$\iota_t \equiv \frac{I_t - I}{I} \approx \log\left(\frac{I_t}{I}\right)$$

The log-linearized real wage ω_t is such that:

$$\omega_t \equiv \frac{\mathcal{W}_t - \mathcal{W}}{\mathcal{W}} \approx \log\left(\frac{\mathcal{W}_t}{\mathcal{W}}\right)$$

The other log-linearized variables are denoted such that: $n_t \equiv \frac{N_t - N}{N} \approx \log\left(\frac{N_t}{N}\right)$.

A.5.1 Consumption

$$\tilde{r}_t = i_t + v_t - \mathbb{E}_t\pi_{t+1} - r$$

$$\hat{\mu}_t = \tilde{\beta}\left(\mathbb{E}_t\hat{\mu}_{t+1} + i_t + v_t - \mathbb{E}_t\pi_{t+1} - r\right)$$

$$Cc_t = -\frac{1}{\gamma}\left(i_t + v_t - \mathbb{E}_t\pi_{t+1} - r\right) + C\mathbb{E}_t c_{t+1} - \gamma\mu^2(\sigma Y)^2\varphi Y\mathbb{E}_t y_{t+1} - \gamma\mu^2(\sigma Y)^2\mathbb{E}_t\hat{\mu}_{t+1}$$

where the wedge is equal to: $\gamma\mu^2(\sigma Y)^2\varphi Y\mathbb{E}_t y_{t+1} + \gamma\mu^2(\sigma Y)^2\mathbb{E}_t\hat{\mu}_{t+1}$.

$$v_t = \varrho_v v_{t-1} + \sigma_v\epsilon_t^v$$

A.5.2 Labor market

$$\pi_t^w = \tilde{\beta}\mathbb{E}_t\pi_{t+1}^w - \kappa_w(\omega_t - mrs_t) + m_t^w$$

$$mrs_t = \frac{l_t}{\rho} + \gamma C c_t$$

$$m_t^w = \varrho_w m_{t-1}^w + \sigma_w \epsilon_t^w - \vartheta_w \epsilon_{t-1}^w$$

$$\omega_t = \omega_{t-1} + \pi_t^w - \pi_t$$

A.5.3 Capital utilization

$$k_t = \hat{u}_t + \bar{k}_{t-1}$$

$$r_t^K = \frac{\lambda}{1-\lambda}\hat{u}_t$$

A.5.4 Capital and investment

$$q_t = \tilde{\beta} R^K \mathbb{E}_t r_{t+1}^K + \tilde{\beta}(1-\delta)\mathbb{E}_t q_{t+1} - (i_t + v_t - \mathbb{E}_t \pi_{t+1} - r)$$

$$\bar{k}_t = (1-\delta)\bar{k}_{t-1} + \delta \iota_t + \delta\chi(1+\tilde{\beta})\hat{\psi}_t$$

$$\iota_t = \frac{1}{1+\tilde{\beta}}\iota_{t-1} + \frac{\tilde{\beta}}{1+\tilde{\beta}}\mathbb{E}_t \iota_{t+1} + \frac{1}{\chi(1+\tilde{\beta})} q_t + \hat{\psi}_t$$

$$\hat{\psi}_t \equiv \frac{1}{\chi(1+\tilde{\beta})}\psi_t$$

$$\hat{\psi}_t = \varrho_I \hat{\psi}_{t-1} + \sigma_I \epsilon_t^I$$

A.5.5 Production and market clearing

$$y_t = z_t + (1-\alpha) l_t + \alpha k_t$$

$$z_t = \varrho_z z_{t-1} + \sigma_z \epsilon_t^z$$

$$y_t = \frac{C}{Y} c_t + \frac{I}{Y}\iota_t + \frac{G}{Y} g_t + \frac{K}{Y} R^K \hat{u}_t$$

A.5.6 Production factors

$$\omega_t = s_t + mpl_t$$

$$mpl_t = y_t - l_t$$

$$r_t^K = s_t + z_t + (1-\alpha)(l_t - k_t)$$

A.5.7 Price dynamics and inflation

$$\pi_t = \tilde{\beta}\mathbb{E}_t \pi_{t+1} - \kappa_p (mpl_t - \omega_t) + m_t^p$$

$$m_t^p = \varrho_p m_{t-1}^p + \sigma_p \epsilon_t^p - \vartheta_p \epsilon_{t-1}^p$$

A.5.8 Monetary policy

$$i_t = \varrho_r i_{t-1} + (1 - \varrho_r)(r + \phi_\pi \pi_t + \phi_y(y_t - y_{t-1})) + f_t$$
$$f_t = \varrho_f f_{t-1} + \sigma_f \epsilon_t^f$$

A.5.9 Fiscal policy

$$g_t = \varrho_g g_{t-1} + \sigma_g \epsilon_t^g + \varrho_{gz} \sigma_z \epsilon_t^z$$

A.6 MCMC diagnostics

After a long burn-in, 1,000,000 posterior draws are used to compute the posterior statistics. The jumping scale parameter is tuned to ensure that the acceptance rate remains between 25% and 33% across all specifications. To provide evidence of chain convergence, I follow Geweke (1992) by testing the equality of means between the first 20% and the last 50% of the draws for each parameter. The table below reports the p-values for each parameter across the baseline, RANK, and TANK models. At the 5% significance level, the null hypothesis of equal means cannot be rejected for any parameter.

Table A.1: Geweke (1992) convergence diagnostics

Parameter	Description	P-value		
		Baseline	RANK	TANK
σ_v	S.d. Risk prem.	0.343	0.208	0.447
σ_z	S.d. TFP	0.058	0.467	0.292
σ_p	S.d. Price MU	0.332	0.376	0.325
σ_w	S.d. Wage MU	0.698	0.369	0.329
σ_I	S.d. IST	0.595	0.572	0.313
σ_f	S.d. Mon. Pol.	0.474	0.174	0.543
σ_g	S.d. Gov. Spend.	0.377	0.975	0.216
ϱ_v	Autocorr. Risk Prem.	0.484	0.639	0.321
ϱ_z	Autocorr. TFP	0.244	0.811	0.419
ϱ_p	Autocorr. Price MU	0.925	0.612	0.599
ϑ_p	MA comp. Price MU	0.147	0.226	0.810
ϱ_w	Autocorr. Wage MU	0.776	0.756	0.265
ϑ_w	MA comp. Wage MU	0.908	0.494	0.478
ϱ_I	Autocorr. IST	0.644	0.854	0.227
ϱ_f	Autocorr. Mon. Pol.	0.664	0.487	0.900
ϱ_g	Autocorr. Gov. Spend.	0.969	0.831	0.252
ϱ_{gz}	Corr. TFP Gov. Spend.	0.646	0.304	0.907
r	Steady state interest rate	0.253	0.571	0.211
α	Capital share	0.896	0.942	0.243
ϱ_r	Taylor rule inertia	0.818	0.378	0.975
ϕ_π	Taylor rule inflation	0.577	0.136	0.059
ϕ_y	Taylor rule output growth	0.467	0.617	0.392
ζ_p	Price stickiness	0.328	0.300	0.816
ζ_w	Nom. wage stickiness	0.925	0.617	0.398
λ	Capital utilization	0.615	0.500	0.199
χ	Investment adjustment	0.602	0.758	0.173
γ	Risk aversion	0.653	0.274	0.516
ρ	Disutility of labor	0.938	0.755	0.443
φ	Risk cyclicality	0.796	$\emptyset$	$\emptyset$
σ	Weight of individual risk	0.561	$\emptyset$	$\emptyset$
η	Share of HtM agents	$\emptyset$	$\emptyset$	0.443
o	Cyclicality of HtM income	$\emptyset$	$\emptyset$	0.638

Note: P-values are computed based on the 15% Taper window of Newey–West standard errors.

A.7 Variance and historical decompositions of the baseline, the RANK, and the TANK models

Table A.2: Unconditional variance decompositions in the baseline

In %	Risk prem.	TFP	Mon. pol.	Price MU	IST	Gov. spend.	Wage MU
Output	7.70	12.71	2.73	39.19	5.28	1.94	30.44
Inflation	68.61	3.42	6.53	16.69	1.57	0.76	2.42
Investment	15.21	6.60	5.44	44.46	21.04	0.91	6.35
Consumption	3.08	7.55	1.10	24.62	0.43	0.17	63.06
Hours worked	9.58	15.79	3.63	23.59	7.97	4.00	35.44
Real wage	1.26	13.32	0.25	83.49	0.55	0.10	1.03
Nominal rate	73.56	1.17	15.58	5.89	2.14	1.05	0.63

Table A.3: Unconditional variance decompositions in the RANK model

In %	Risk prem.	TFP	Mon. pol.	Price MU	IST	Gov. spend.	Wage MU
Output	1.94	17.12	1.59	18.71	35.44	1.23	23.97
Inflation	39.37	3.70	7.09	9.59	36.07	0.08	4.11
Investment	1.34	5.89	1.10	11.78	70.15	0.02	9.73
Consumption	1.43	15.81	1.09	17.31	24.56	0.50	39.30
Hours worked	2.52	4.68	2.28	14.27	49.12	2.02	25.10
Real wage	0.77	24.74	0.40	49.18	21.38	0.06	3.46
Nominal rate	42.45	0.96	4.14	1.44	49.00	0.14	1.87

Table A.4: Unconditional variance decompositions in the TANK model

In %	Risk prem.	TFP	Mon. pol.	Price MU	IST	Gov. spend.	Wage MU
Output	3.77	19.79	3.42	17.08	34.69	1.42	19.83
Inflation	30.62	3.46	6.87	9.29	45.78	0.13	3.85
Investment	2.64	6.20	2.53	11.17	66.56	0.11	10.79
Consumption	3.18	19.76	2.47	15.74	30.28	0.72	27.84
Hours worked	4.82	4.74	4.77	15.30	45.30	2.63	22.43
Real wage	0.69	27.86	0.36	37.67	29.70	0.15	3.57
Nominal rate	33.48	0.54	2.80	1.11	60.72	0.21	1.14

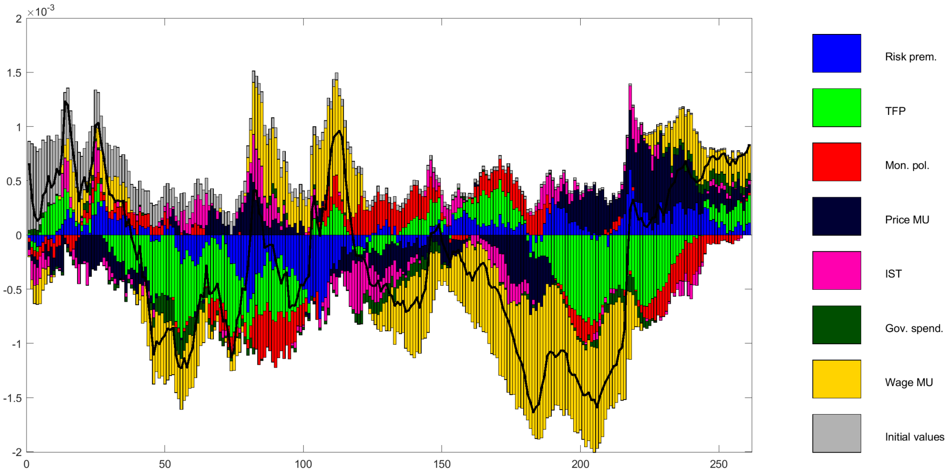

(a) Wedge in the Euler equation

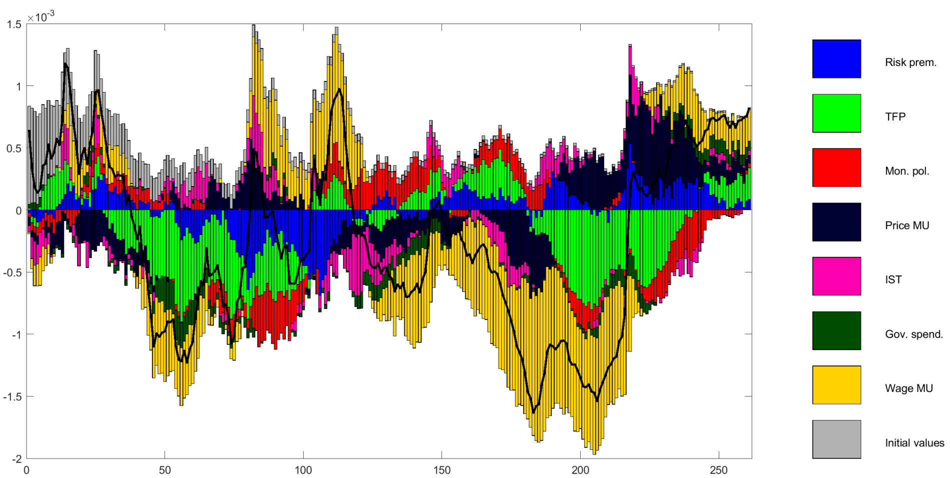

(b) Cyclical component

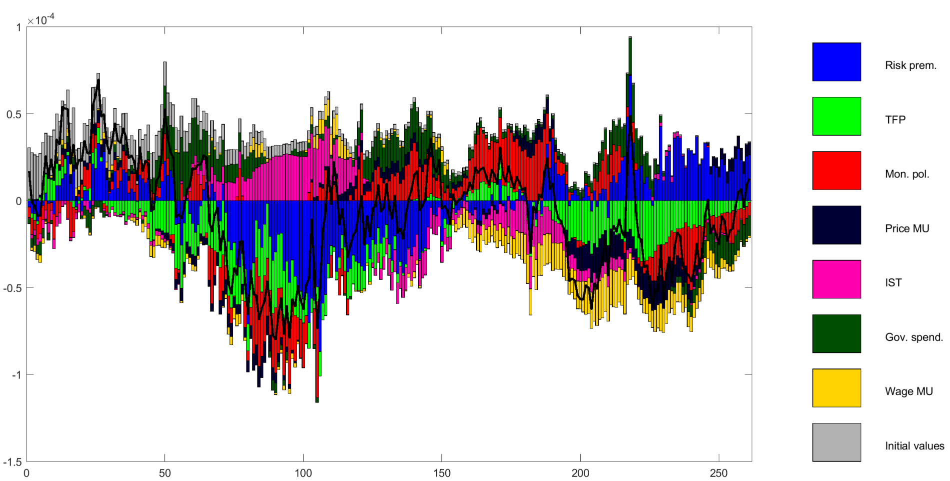

(c) MPC component

Figure A.1: Historical decompositions of the wedge in the Euler equation and its components

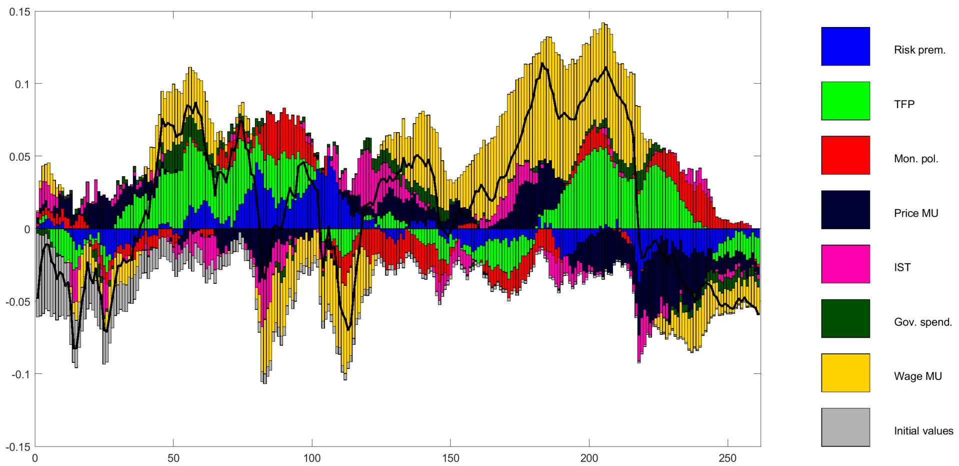

(a) Baseline

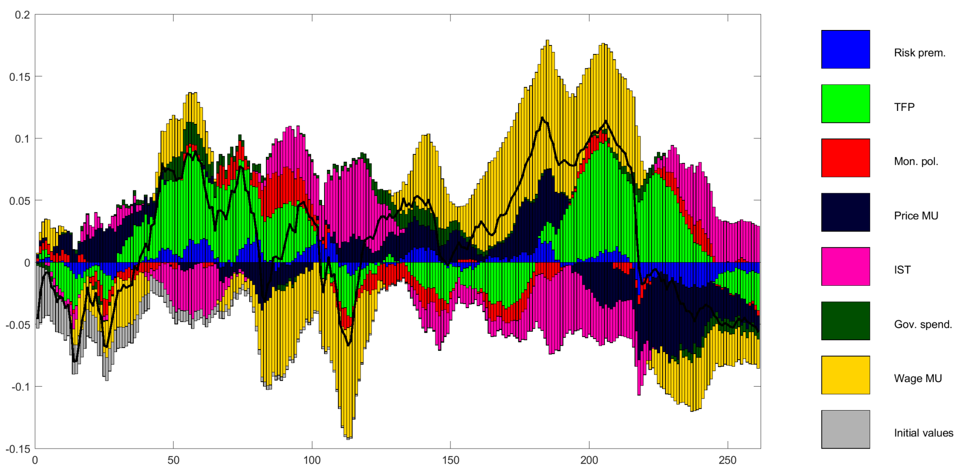

(b) RANK

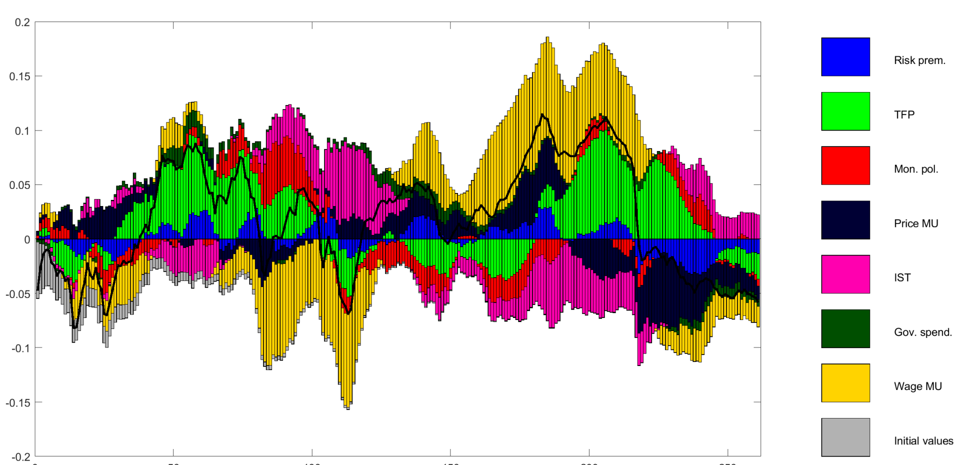

(c) TANK

Figure A.2: Historical decomposition of output in different specifications

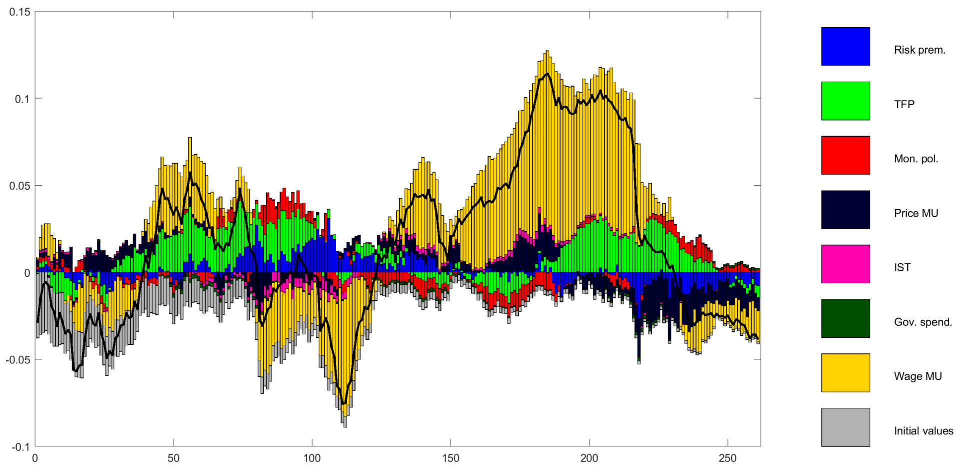

(a) Baseline

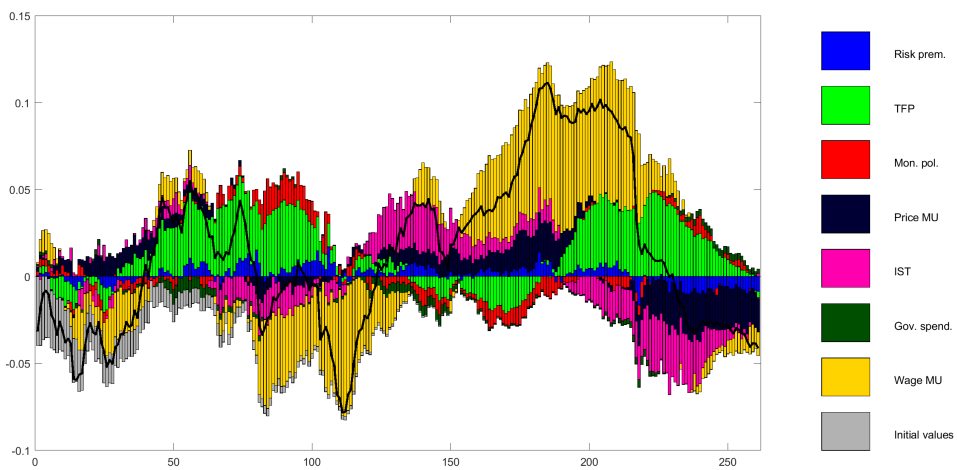

(b) RANK

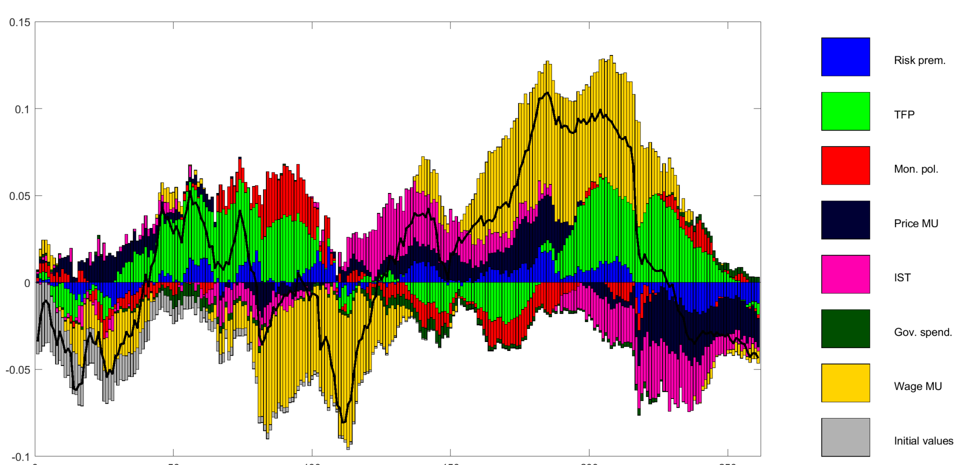

(c) TANK

Figure A.3: Historical decomposition of consumption in different specifications

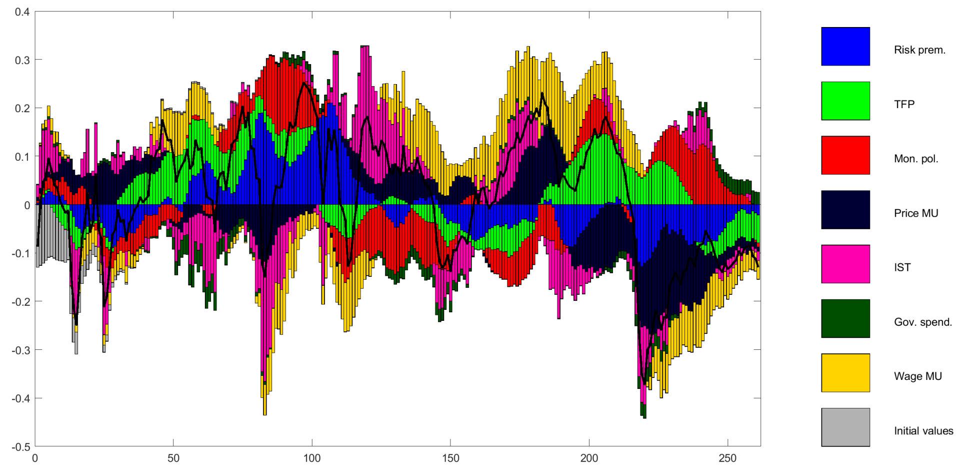

(a) Baseline

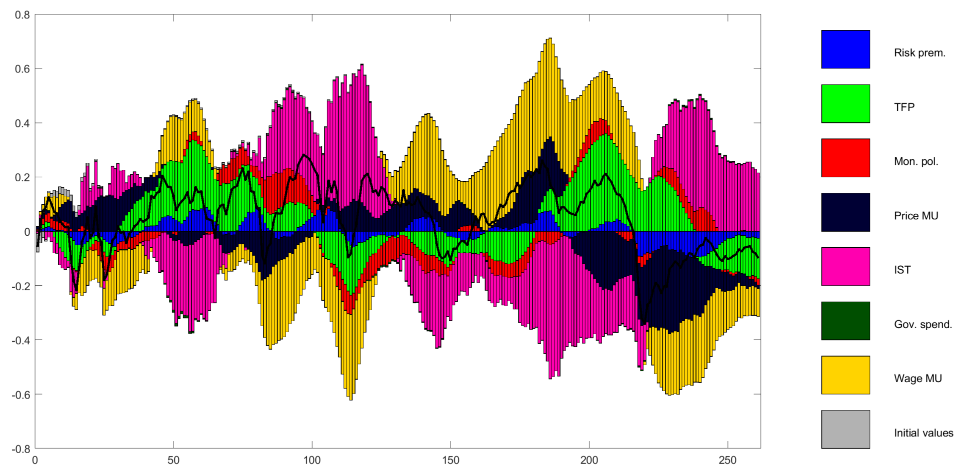

(b) RANK

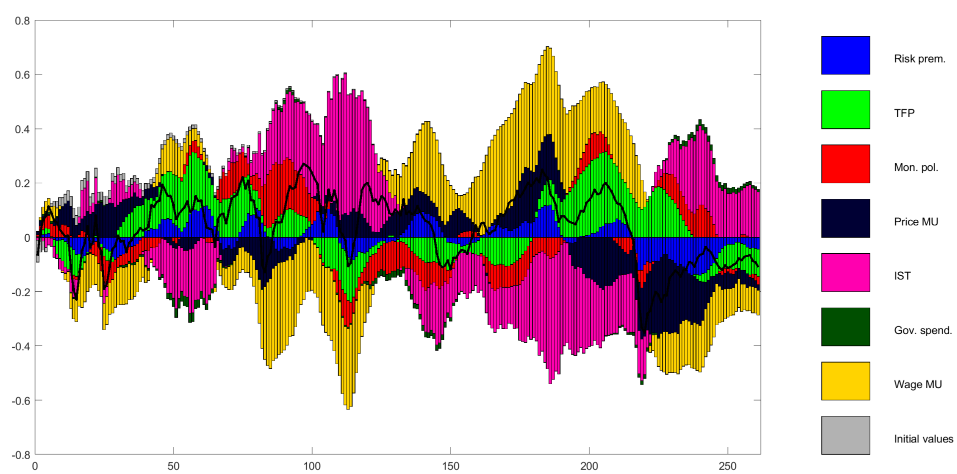

(c) TANK

Figure A.4: Historical decomposition of investment in different specifications

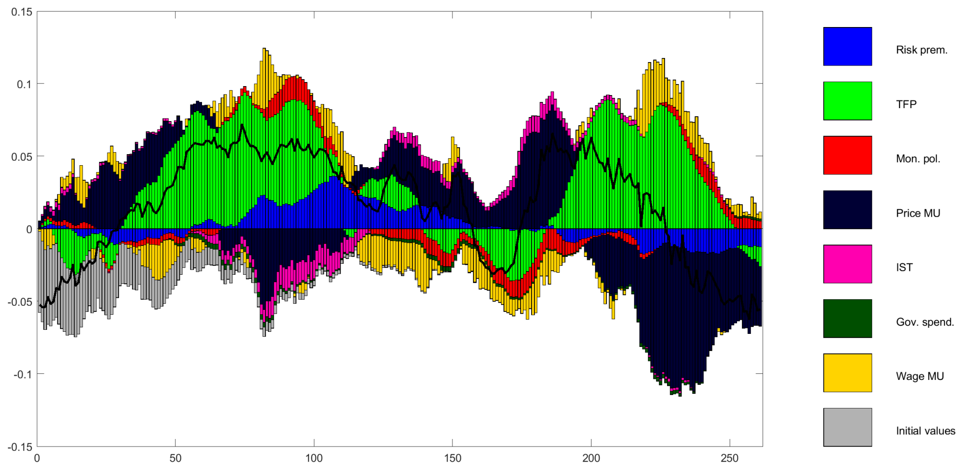

(a) Baseline

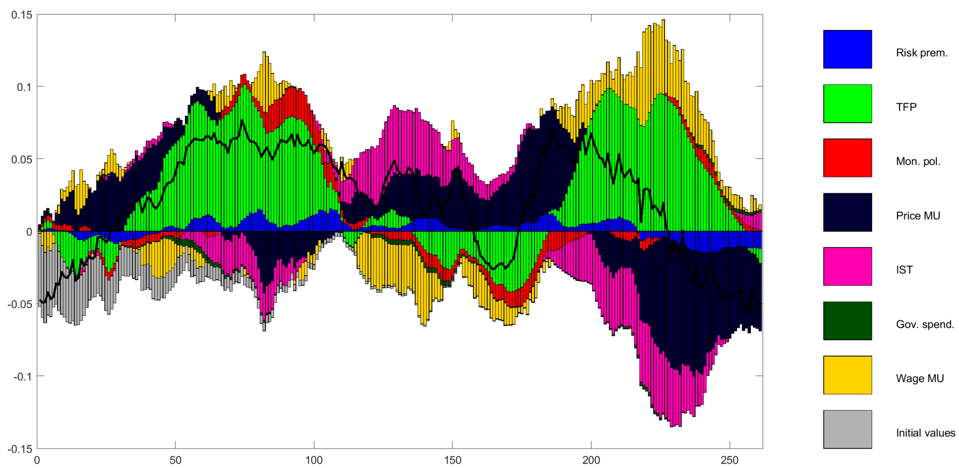

(b) RANK

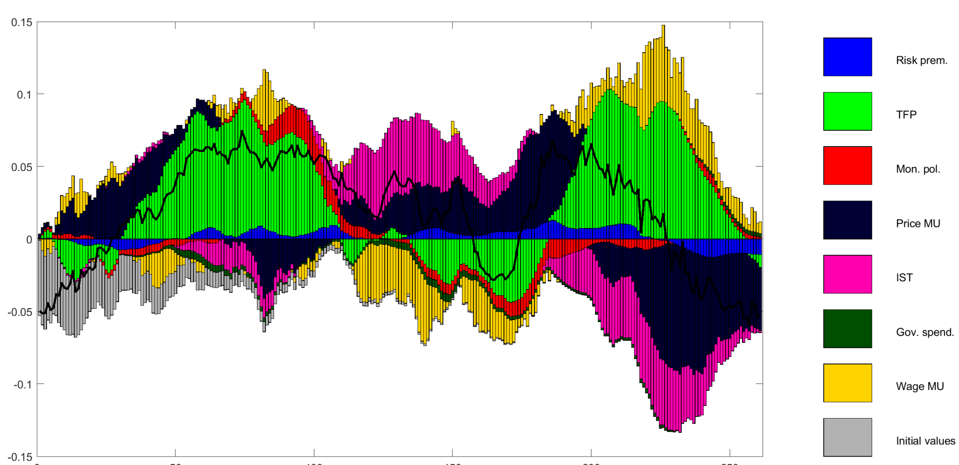

(c) TANK

Figure A.5: Historical decomposition of the real wage in different specifications

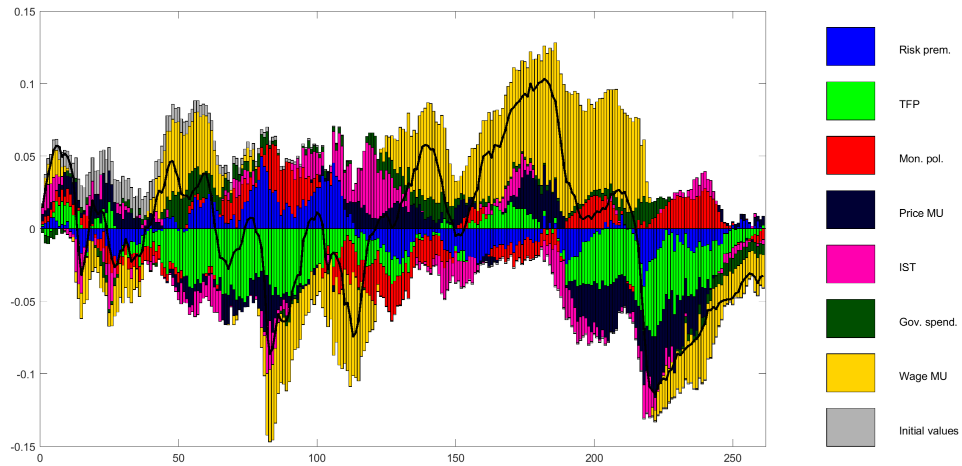

(a) Baseline

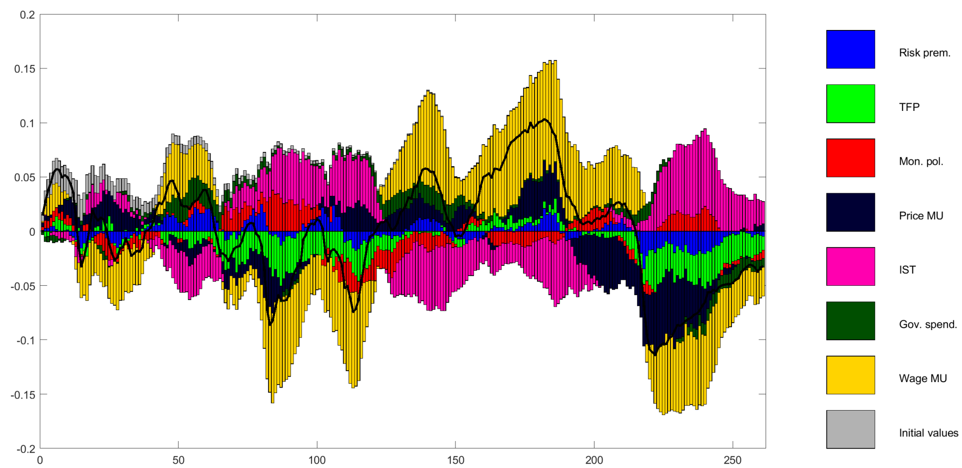

(b) RANK

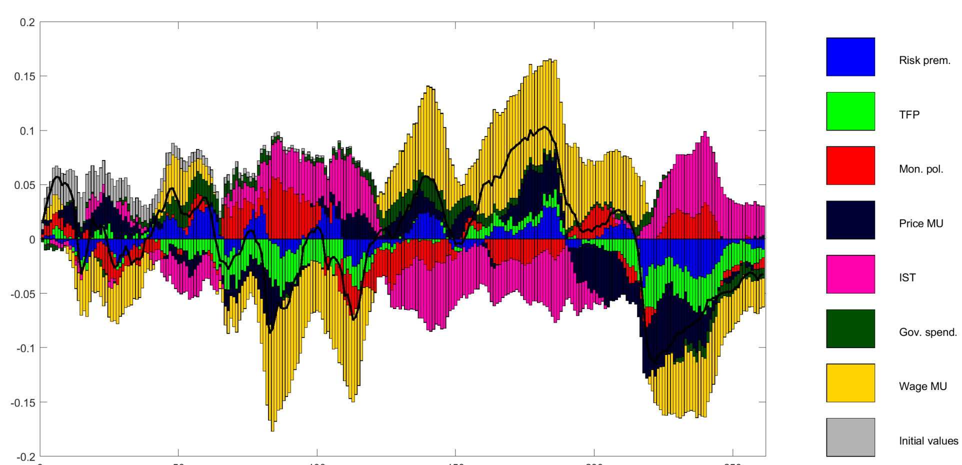

(c) TANK

Figure A.6: Historical decomposition of hours worked in different specifications

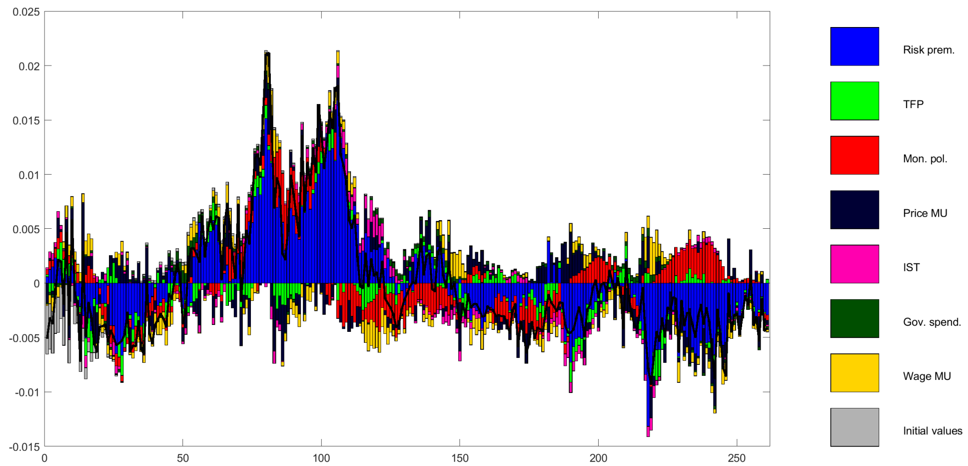

(a) Baseline

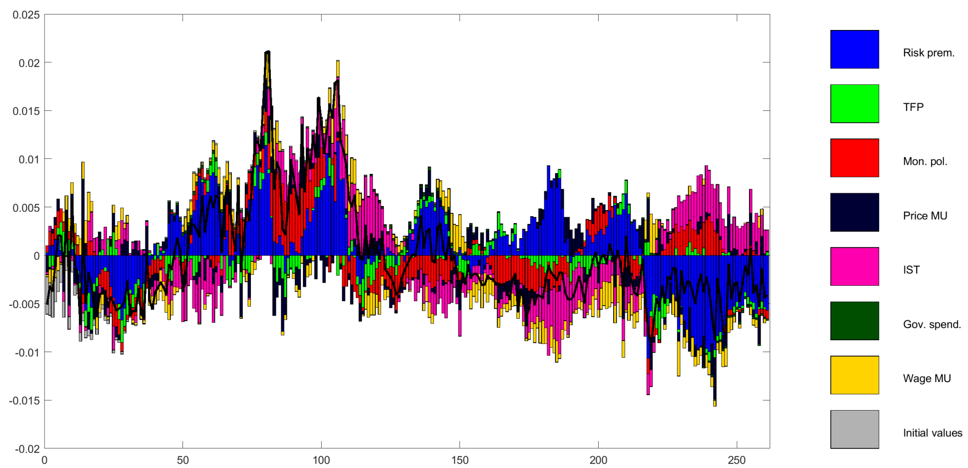

(b) RANK

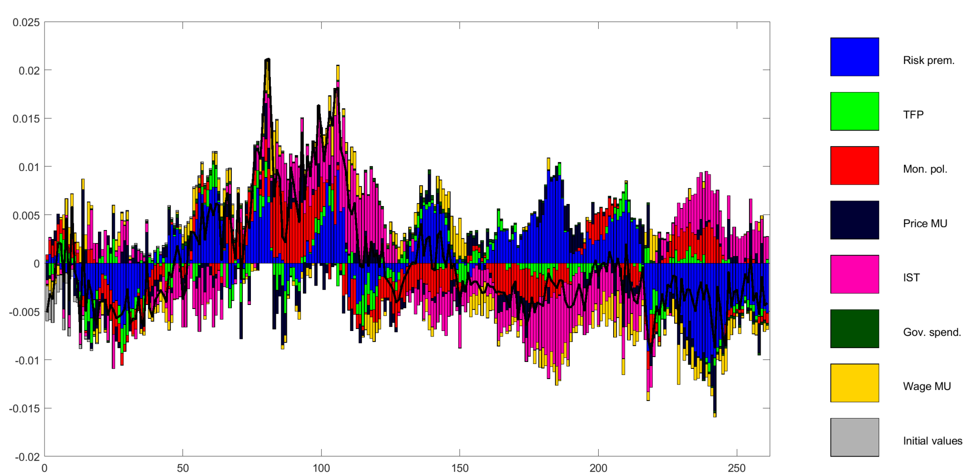

(c) TANK

Figure A.7: Historical decomposition of inflation in different specifications

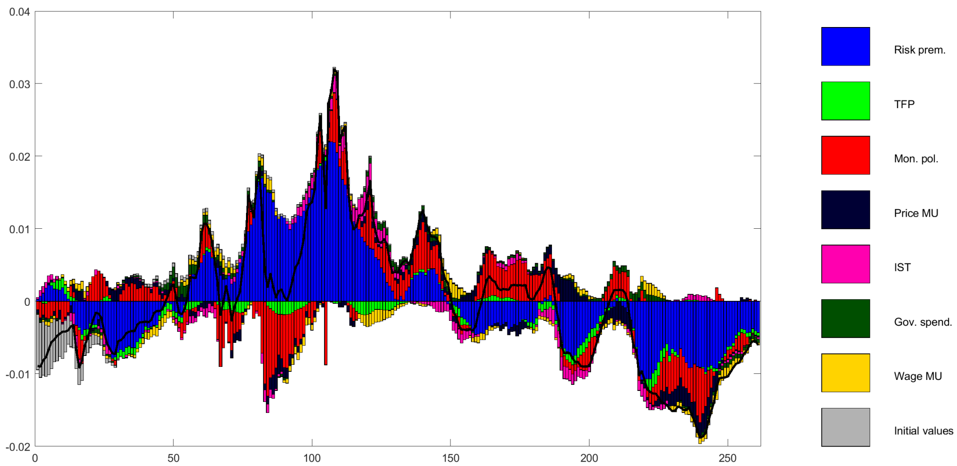

(a) Baseline

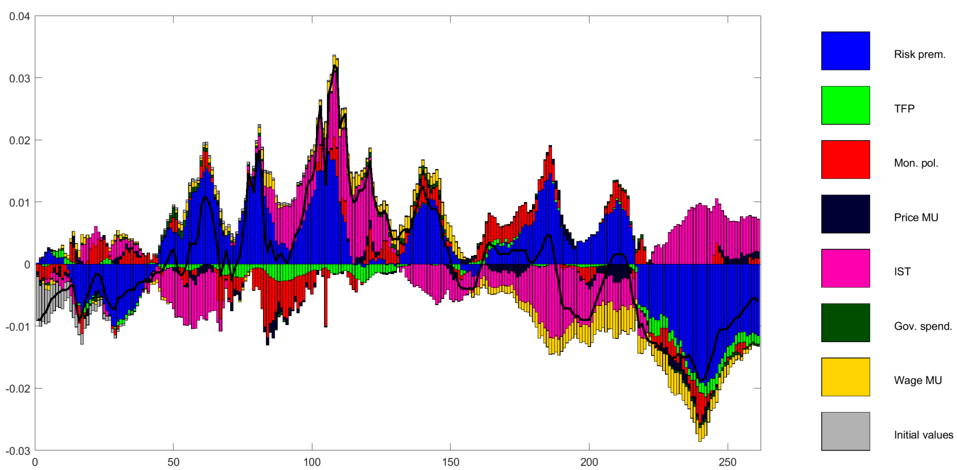

(b) RANK

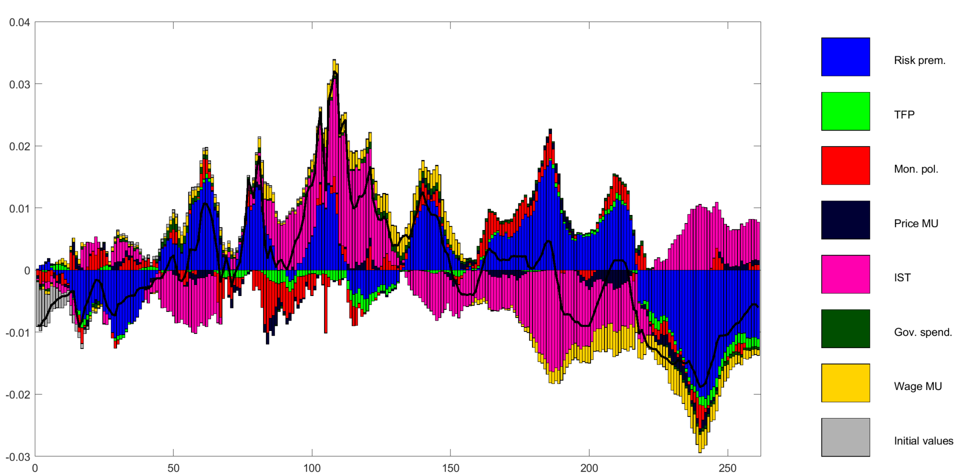

(c) TANK

Figure A.8: Historical decomposition of the nominal rate in different specifications

A.8 Estimated response functions

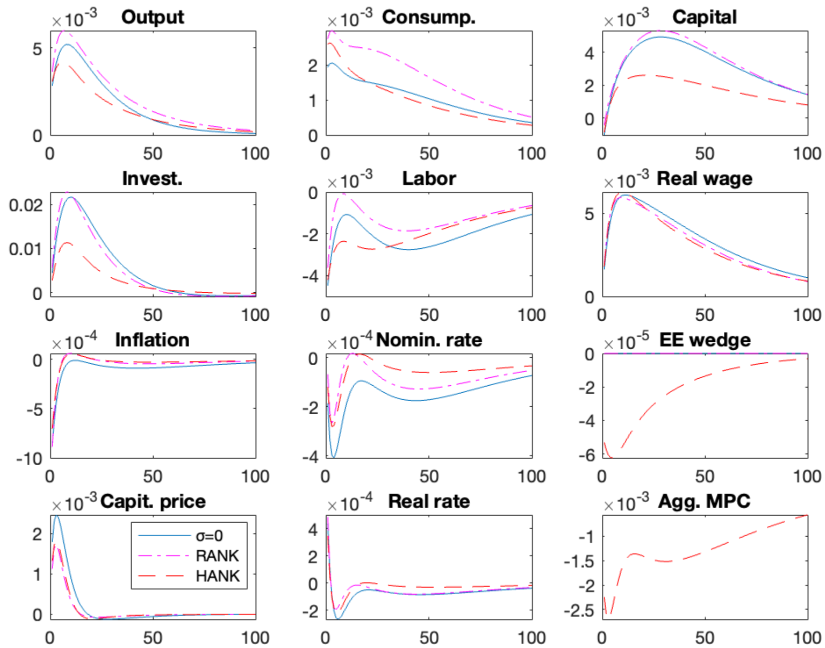

Figure A.9: Estimated response to a TFP shock: comparison with the RA counterfactual ($\sigma = 0$) and the reestimated RANK model.

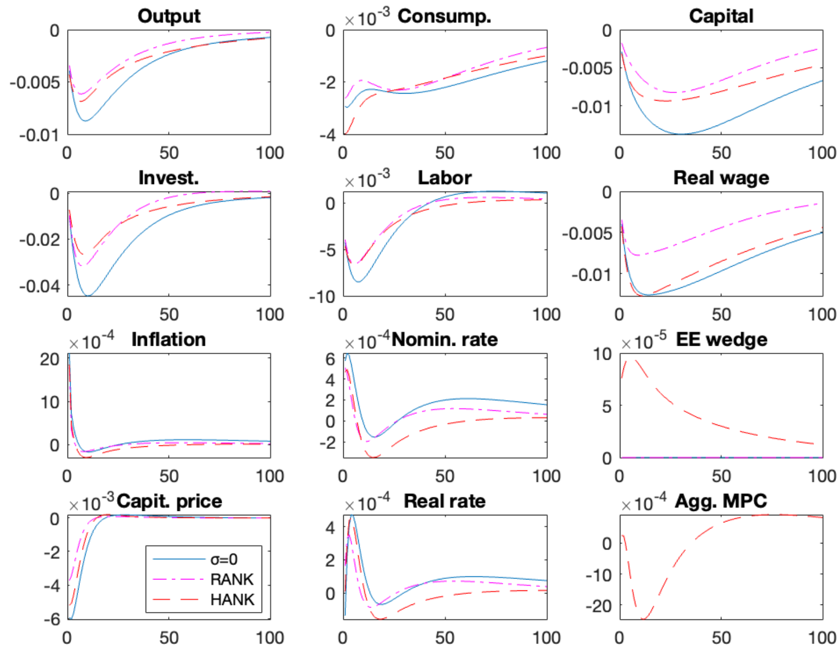

Figure A.10: Estimated response to a price mark-up shock: comparison with the RA counterfactual ($\sigma = 0$) and the reestimated RANK model.

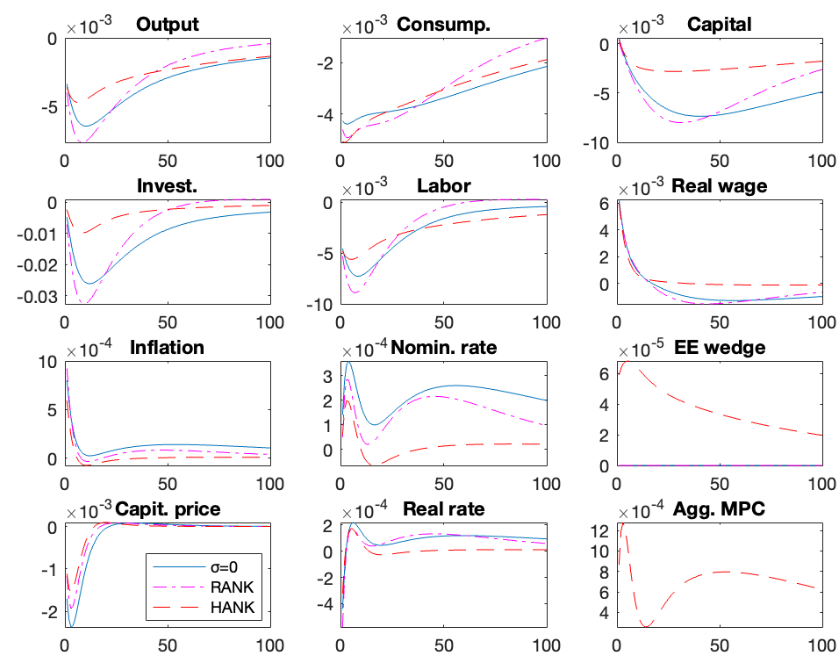

Figure A.11: Estimated response to a wage mark-up shock: comparison with the RA counterfactual ($\sigma = 0$) and the reestimated RANK model.

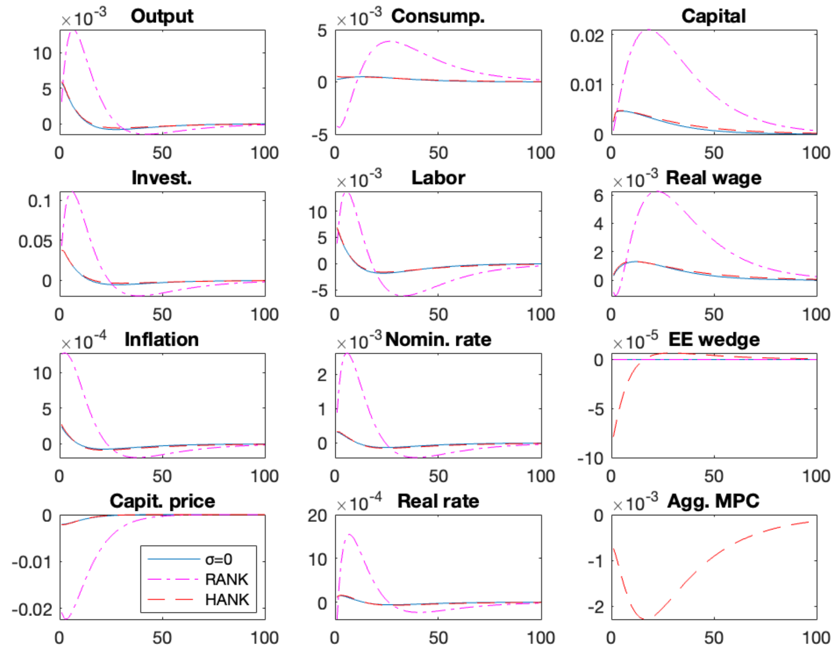

Figure A.12: Estimated response to an investment-specific technology shock: comparison with the RA counterfactual ($\sigma = 0$) and the reestimated RANK model.

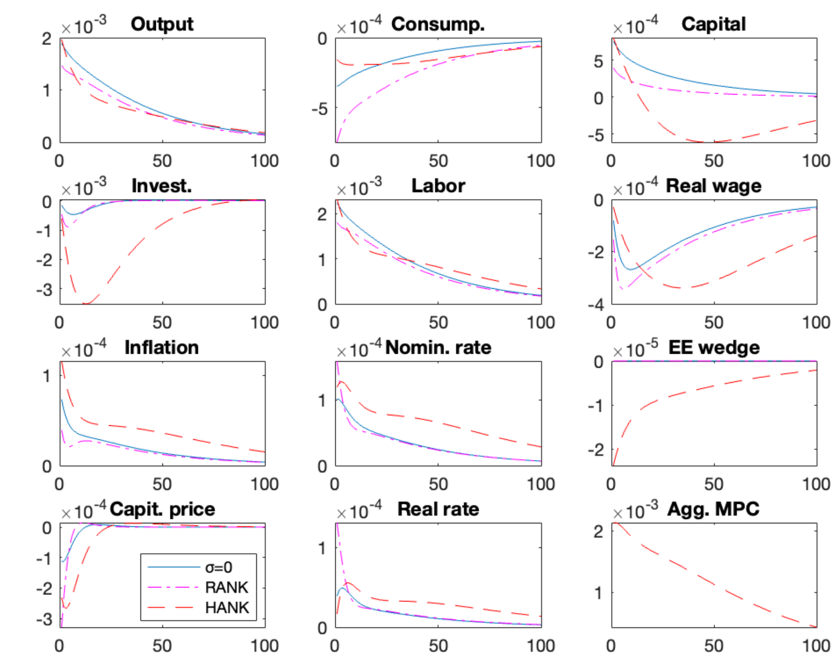

Figure A.13: Estimated response to an increase in government spendings: comparison with the RA counterfactual ($\sigma = 0$) and the reestimated RANK model.

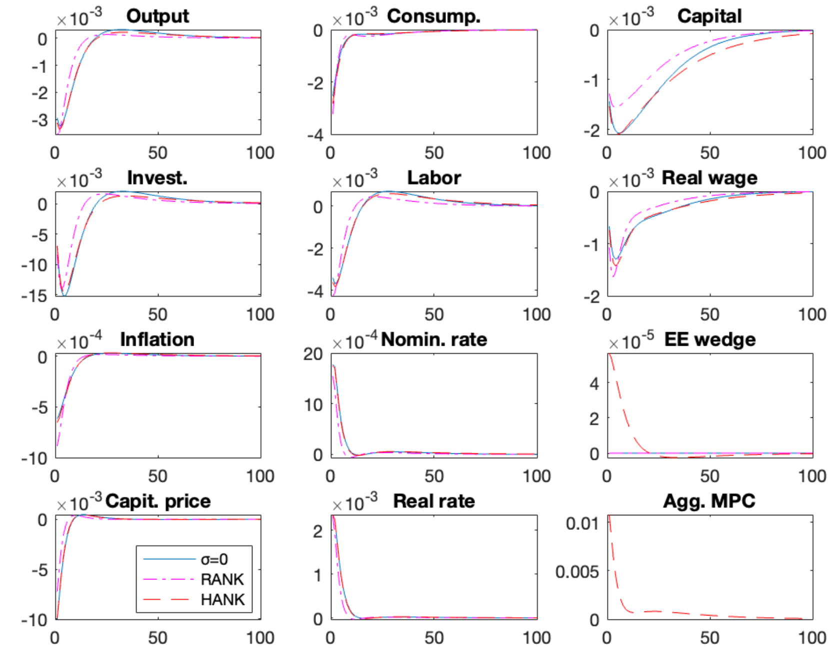

Figure A.14: Estimated response to a monetary policy innovation: comparison with the RA counterfactual ($\sigma = 0$) and the reestimated RANK model.

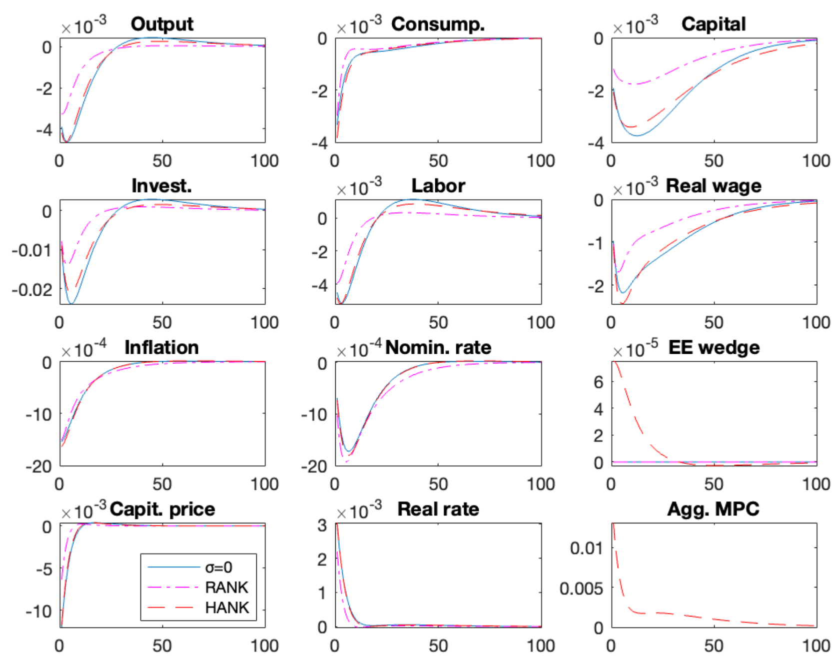

Figure A.15: Estimated response to a risk premium shock: comparison with the RA counterfactual ($\sigma = 0$) and the reestimated RANK model.

References

Acharya, Sushant, Edouard Challe, and Keshav Dogra (2023) "Optimal Monetary Policy According to HANK," *American Economic Review*, 113 (7), 1741–1782.

Acharya, Sushant, William Chen, Marco Del Negro, Keshav Dogra, Aidan Gleich, Shlok Goyal, Donggyu Lee, Ethan Matlin, Reca Sarfati, and Sikata Sengupta (2024) "Estimating HANK for Central Banks," in Sofía Bauducco, Andrés Fernández and Giovanni L. Violante eds. *Heterogeneity in Macroeconomics: Implications for Monetary Policy*, 181–225: Central Bank of Chile.

Acharya, Sushant and Keshav Dogra (2020) "Understanding HANK: Insights From a PRANK," *Econometrica*, 88 (3), 1113–1158.

Auclert, Adrien, Matthew Rognlie, Martin Souchier, and Ludwig Straub (2021) "Exchange Rates and Monetary Policy with Heterogeneous Agents: Sizing up the Real Income Channel," NBER Working Papers 28872, National Bureau of Economic Research.

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2020) "Micro Jumps, Macro Humps: Monetary Policy and Business Cycles in an Estimated HANK Model," NBER Working Papers 26647, National National Bureau of Economic Research.

Bayer, Christian, Benjamin Born, and Ralph Luetticke (2024) "Shocks, Frictions, and Inequality in US Business Cycles," *American Economic Review*, 114 (5), 1211–47.

Bayer, Christian, Ralph Luetticke, Lien Pham-Dao, and Volker Tjaden (2019) "Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk," *Econometrica*, 87 (1), 255–290.

Berger, David, Luigi Bocola, and Alessandro Dovis (2023) "Imperfect Risk Sharing and the Business Cycle," *The Quarterly Journal of Economics*, 138 (3), 1765–1815.

Bilbiie, Florin O. (2020) "The New Keynesian cross," *Journal of Monetary Economics*, 114 (C), 90–108.

——— (2025) "Monetary Policy and Heterogeneity: An Analytical Framework," *Review of Economic Studies*, 92 (4), 2398–2436.

Bilbiie, Florin O., Giorgio Primiceri, and Andrea Tambalotti (2023) "Inequality and Business Cycles," NBER Working Papers 31729, National Bureau of Economic Research.

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

Challe, Edouard, Julien Matheron, Xavier Ragot, and Juan F. Rubio-Ramirez (2017) "Precautionary Saving and Aggregate Demand," *Quantitative Economics*, 8 (2), 435–478.

Cho, Daeha (2023) "Unemployment risk, MPC heterogeneity, and business cycles," *Quantitative Economics*, 14 (2), 717–751.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005) "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, 113 (1), 1–45.

Den Haan, Wouter J., Pontus Rendahl, and Markus Riegler (2018) "Unemployment (Fears) and Deflationary Spirals," *Journal of the European Economic Association*, 16 (5), 1281–1349.

Dixit, Avinash K. and Joseph E. Stiglitz (1977) "Monopolistic Competition and Optimum Product Diversity," *American Economic Review*, 67 (3), 297–308.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin (2000) "Optimal monetary policy with staggered wage and price contracts," *Journal of Monetary Economics*, 46 (2), 281–313.

Geweke, John (1992) "Evaluating the Accuracy of Sampling-Based Approaches to the Calculation of Posterior Moments," in *Bayesian Statistics*, 4: Oxford:Clarendon Press.

Guvenen, Fatih, Serdar Ozkan, and Jae Song (2014) "The Nature of Countercyclical Income Risk," *Journal of Political Economy*, 122 (3), 621–660.

Hamilton, James D. (2018) "Why You Should Never Use the Hodrick-Prescott Filter," *The Review of Economics and Statistics*, 100 (5), 831–843.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2010) "Investment shocks and business cycles," *Journal of Monetary Economics*, 57 (2), 132–145.

——— (2011) "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, 14 (1), 101–121.

U.S. Bureau of Labor Statistics, Labor Force Statistics from the Current Population Survey (accessed May 22, 2025) https://www.bls.gov/cps/.

U.S. Bureau of Labor Statistics, Unemployment Rate [UNRATE] retrieved from FRED, Federal Reserve Bank of St. Louis (accessed Nov. 2, 2024) https://fred.stlouisfed.org/series/UNRATE.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson (2016) "The Power of Forward Guidance Revisited," *American Economic Review*, 106 (10), 3133–3158.

University of Michigan, University of Michigan: Consumer Sentiment [UMCSENT] retrieved from FRED, Federal Reserve Bank of St. Louis (accessed Nov. 2, 2024) https://fred.stlouisfed.org/series/UMCSENT.

Ravn, Morten O. and Vincent Sterk (2017) "Job uncertainty and deep recessions," *Journal of Monetary Economics*, 90 (C), 125–141.

——— (2021) "Macroeconomic Fluctuations with HANK & SAM: an Analytical Approach," *Journal of the European Economic Association*, 19 (2), 1162–1202.

Shimer, Robert (2005) "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95 (1), 25–49.

——— (2012) "Reassessing the Ins and Outs of Unemployment," *Review of Economic Dynamics*, 15 (2), 127–148.

Smets, Frank and Rafael Wouters (2007) "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 97 (3), 586–606.

U.S. Census Bureau, Survey of Income and Program Participation (SIPP) (accessed Nov. 2, 2024) https://www.census.gov/programs-surveys/sipp.html.

Werning, Iván (2015) "Incomplete Markets and Aggregate Demand," NBER Working Papers 21448, National Bureau of Economic Research.

Wu, Jing Cynthia and Fan Dora Xia (2016) "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," *Journal of Money, Credit and Banking*, 48 (2-3), 253–291.