Quantitative Economics 14 (2023), 717–751 1759-7331/20230717

Unemployment risk, MPC heterogeneity, and business cycles

Daeha Cho
College of Economics and Finance, Hanyang University

This paper uses an estimated Heterogeneous Agent New Keynesian (HANK) model to evaluate the quantitative importance of two channels in driving aggregate consumption fluctuations in the US: (i) precautionary savings against unemployment risk and (ii) MPC heterogeneity. I find that MPC heterogeneity is the dominant channel because a large fraction of households are close to the borrowing limit. The empirical average MPC target in HANK generates counterfactually volatile aggregate consumption, and thus makes it more difficult for the estimated model to match the persistence of the aggregate data, indicating an MPC puzzle. This is because the likelihood-based estimation favors a low degree of nominal rigidity and responsive monetary policy in the HANK model to reduce the discrepancy between consumption volatility in the model and in the data. The low degree of nominal rigidity and responsive monetary policy reduce the persistence of endogenous variables in the model.

Keywords. Heterogeneous Agent New Keynesian model, Bayesian estimation, precautionary savings, marginal propensity to consume.

JEL classification. E20, E32.

1. Introduction

Do household heterogeneity and incomplete markets matter for aggregate fluctuations? This has been one of the most fundamental questions in the quantitative macroeconomics literature at least since Krusell and Smith (1998). Emerging literature revisits this question using Heterogeneous Agent New Keynesian (HANK hereafter) models that combine incomplete markets with nominal rigidities. The HANK literature stresses two key channels that make aggregate consumption dynamics in HANK models different from those in Representative Agent New Keynesian (RANK hereafter) models: (i) precautionary savings motives against countercyclical uninsurable risk and (ii) marginal propensity to consume (MPC hereafter) heterogeneity.[1]

Daeha Cho: daehac@hanyang.ac.kr

An earlier version of the paper was circulated under the title “Unemployment Risk and Business Cycles.” I thank three anonymous referees for their helpful comments and suggestions that have improved the paper substantially. I also benefited from Alisdair McKay, Robert King, Stephen Terry, Simon Gilchrist, Jianjun Miao, Susanto Basu, Ben Moll, Bruce Preston, Chris Edmond, and participants at numerous seminars and conferences. This work was supported by the research fund of Hanyang University (HY-202300000001166).

[1] Two examples of countercyclical uninsurable risks considered in the literature are unemployment risk and risk to long-term income growth. Bayer, Luetticke, Pham-Dao, and Tjaden (2019) document the importance of the asset portfolio adjustment channel that arises from changes in precautionary motives.

Available at http://qeconomics.org. https://doi.org/10.3982/QE1550

However, little is known about the extent to which these two channels are supported by the aggregate data because most HANK literature uses calibrated models to quantify each channel. Moreover, these two channels are typically studied in isolation, making it hard to evaluate the relative strength of each channel. For example, many HANK models that document the importance of countercyclical unemployment risk assume that all employed households save and all unemployed households live hand-to-mouth, implying a counterfactually low average MPC.[2] In contrast, most studies that highlight the role of MPC heterogeneity either do not isolate cyclical uninsurable risk or completely exclude such a risk.[3]

The goal of this paper is to assess the relative importance of each channel in a HANK model that embeds both the empirical average MPC and unemployment risk and evaluate the fit of the model to the aggregate data. To this end, I build on the model of Challe et al. (2017), matching two key moments: the empirical average MPC and the average consumption differential between employed and unemployed households. The model features nominal and real frictions and aggregate shock processes that are understood to improve the fitness of DSGE models to aggregate data. Exploiting a recent method for solving and estimating heterogeneous agent models introduced by Winberry (2018), I estimate the structural parameters using the US aggregate time series with a Bayesian technique in the style of Justiniano, Primiceri, and Tambalotti (2010).

In the model, two channels contribute to the aggregate consumption responses. First, a change in the current macroeconomic conditions endogenously alters individuals' perceived probability of becoming unemployed, thus affecting current aggregate consumption via precautionary motives. This is the time-varying unemployment risk channel. Second, a change in the current macroeconomic conditions induces a change in real wages and the number of unemployed households. The resulting change in the current aggregate household income affects the current aggregate consumption via high average MPCs. This is the MPC heterogeneity channel.

Using the estimated model, I compare the volatilities of consumption in HANK and RANK. I decompose the difference in consumption volatility into the portion that arises from unemployment risk and the portion that arises from MPC heterogeneity. I find that the contribution of unemployment risk to the difference between consumption volatilities in HANK and RANK is 12%. Most of this difference arises from MPC heterogeneity. The significant contribution of MPC heterogeneity arises from the large fraction of households that are close to the borrowing limit. Intuitively, households that are constrained or near the borrowing limits barely precautionary-save as they live close to hand-to-mouth. As these wealth-poor households have high MPCs, their spending is largely determined by their income realization today. The dominant role of MPC heterogeneity in driving aggregate consumption volatility is robust to less responsive monetary

[2] Prominent examples for studies that emphasize unemployment risk are Kreamer (2016), Ravn and Sterk (2017, 2018), Challe, Matheron, Ragot, and Rubio-Ramírez (2017), Den Haan, Rendahl, and Riegler (2018), Heathcote and Perri (2018), McKay and Reis (2019), Challe (2020), and Oh and Rogantini Picco (2019). Although Kreamer (2016) and Den Haan, Rendahl, and Riegler (2018) allow employed households to hit the borrowing constraint, they do not discuss the role of MPCs.

[3] See Galí, López-Salido, and Vallés (2007), Oh and Reis (2012), Kaplan, Moll, and Violante (2018), Auclert (2019), Kekre (2019), and Bilbiie (2020) for studies that focus on MPC heterogeneity.

policy and alternative investment adjustment costs, average employed–unemployed consumption differences, and liquid wealth to GDP ratios.

I then evaluate the extent to which the estimated HANK model explains the aggregate data by comparing its model fit, measured by log marginal likelihood, with that of the estimated RANK model. I find that the estimated HANK underperforms the estimated RANK in fitting the aggregate data. The reason for this underperformance is the empirical average MPC target in HANK, which generates counterfactually volatile aggregate consumption. To reduce consumption volatility, the likelihood-based estimation delivers a low degree of nominal rigidity and responsive monetary policy in HANK relative to those in RANK. The more flexible prices are, the more quickly the prices return to the steady state, causing quantities to converge faster. Therefore, it is more difficult for the estimated HANK to match the persistent aggregate data than the estimated RANK, revealing an *MPC puzzle*.

The present paper contributes to the active literature on the Bayesian estimation of HANK models. Bayer, Born, and Luetticke (2020) use an estimated two-asset HANK model with portfolio choice to study the sources of the evolution of US wealth and income inequality. The earnings risk in their paper is exogenous, whereas it evolves endogenously in my model due to search and matching frictions. Auclert, Rognlie, and Straub (2020) estimate only the aggregate shock processes using a likelihood-based method. In contrast, I also estimate the parameters of price and wage stickiness and the coefficients in the monetary policy rule. The work closest to mine is Challe et al. (2017). They study the importance of precautionary savings against unemployment risk in aggregate dynamics using an estimated model. However, they make assumptions on risk-sharing and market structure in order to construct an analytically tractable equilibrium with the wealth distribution of finite support, which restricts them from matching MPCs. Drawing on Winberry's (2018) method, I match the empirical average MPC, distinguishing the unemployment risk channel from the MPC heterogeneity channel.

Moreover, the present paper builds on the work by Eusepi and Preston (2015), who emphasize the effect of compositional changes between employed and unemployed households on aggregate consumption. Because the average consumption of employed households is higher than that of unemployed households, the compositional changes directly affect aggregate consumption. They show that this effect is powerful enough to solve the comovement problem that arises after nonproductivity shocks, namely the Barro and King (1984) problem. My model includes the composition effect, which is part of the MPC heterogeneity channel. However, they assume asset markets are complete, so households do not have a precautionary savings motive.

The paper proceeds as follows. Section 2 presents the baseline HANK model. Section 3 discusses how the model is solved. Section 4 describes the fixed and estimated parameters, properties of the stationary equilibrium, and estimation results. In Section 5, I compare aggregate dynamics in HANK and RANK, focusing on consumption. Section 6 compares HANK and RANK focusing on parameter estimates, shock decompositions, and model fits. Section 7 discusses the property of two-agent New Keynesian models. Section 8 concludes. The Appendices to the paper can be found in the Online Supplementary Material (Cho (2023)).

2. Model

The economy is populated by households that self-insure against the idiosyncratic incidence of unemployment. Similar to Krusell and Smith (1998), McKay and Reis (2016), and Hagedorn, Manovskii, and Mitman (2019), I allow for two groups of households, which permanently differ in the discount factor. I label the households with a low discount factor as impatient households and those with a high discount factor as patient households. As will be discussed in Section 4.1, heterogeneous discount factors are necessary to match the realistic MPC and the supply of liquid assets available in the economy. The remaining parts consist of standard ingredients present in medium-scale DSGE models with a frictional labor market. Such ingredients are a representative final goods firm, a continuum of wholesale firms producing differentiated goods subject to price rigidity, a representative intermediate goods firm that hires and invests, a monetary authority, and a fiscal authority. The model resembles that of Challe et al. (2017) but has an important difference in terms of wealth distribution. The assumption on wealth distribution in their model is restrictive in the sense that all employed households have the same amount of wealth, regardless of employment histories. However, in my model, different individual employment histories generate heterogeneity in wealth across households.

2.1 *Households*

A fraction $1-\Omega$ of households are impatient and indexed by $i \in [0, 1-\Omega]$. Impatient households transition between two states: employment and unemployment. They receive real wages when employed and real unemployment benefits when unemployed. As will be discussed in Section 4.1, I choose the level of unemployment benefits to match the difference in average consumption between employed and unemployed households. Therefore, in my model, the level of unemployment benefits should be interpreted as the degree of partial insurance that includes other insurance devices such as home production.

Impatient households can self-insure only through trading riskless liquid assets, but they cannot take short positions.[4] The budget constraint of impatient household i at period t is given by

$$C_{i,t} + a_{i,t+1} = (1-\tau_t)\frac{W_t}{P_t}e_{i,t} + (1-\tau_t)b^u\frac{W_t}{P_t}(1-e_{i,t}) + \frac{R_{t-1}}{\pi_t}a_{i,t}, \tag{2.1}$$

together with borrowing constraint, $a_{i,t+1} \geq 0$, where $C_{i,t}$ denotes the consumption of impatient household i. π_t denotes the gross inflation rate, R_{t-1} is the gross nominal interest rate paid on liquid assets purchased in period $t-1$. W_t is the nominal wage, and P_t is the aggregate price index. b^u is the replacement rate, and τ_t denotes the tax rate

[4]An alternative way to generate a realistic MPC is to work with an incomplete markets model with two assets, liquid and illiquid, as in Kaplan and Violante (2014). The assumption on asset holdings in my model can be interpreted as households' strong preference for using liquid savings over costly illiquid savings to insure against unemployment risk, which is relatively short-lived.

on labor and transfer income. $e_{i,t}$ refers to an indicator for employment status, where $e_{i,t} = 1$ if the household is employed and $e_{i,t} = 0$ if it is unemployed.

Impatient households choose a stream of consumption and savings that maximizes

$$\mathbb{E}_0 \sum_{t=0}^{\infty} (\beta_L^t \zeta_t) \left[\frac{C_{i,t}^{1-\sigma}}{1-\sigma} \right],$$

subject to the budget constraint (2.1) and a borrowing constraint. β_L is the discount factor for impatient households, and ζ_t is a common preference shock and evolves according to

$$\log(\zeta_t) = \rho_\zeta \log(\zeta_{t-1}) + \epsilon_t^\zeta, \quad \epsilon_t^\zeta \overset{\text{iid}}{\sim} \mathcal{N}(0, {\sigma_\zeta}^2).$$

A fraction Ω of households are patient and are assumed to be unconstrained, and thus their liquid asset Euler equation holds with equality in all periods. Each household acts as a representative family, in which all members share all types of income. Therefore, all members enjoy the same level of consumption regardless of their employment status. Patient households' preferences are the same as those of impatient households, except for the discount factor. Their period t budget constraint is

$$C_{H,t} + a_{H,t+1} = (1-\tau_t)\left(\frac{W_t}{P_t} n_t + b^u \frac{W_t}{P_t}(1-n_t)\right) + \frac{R_{t-1}}{\pi_t} a_{H,t} + \frac{D_t}{P_t}, \tag{2.2}$$

where $C_{H,t}$ and $a_{H,t+1}$ are consumption and savings of liquid assets by a patient household, respectively. n_t is the employment rate, and D_t denotes the sum of dividends collected from wholesale and intermediate goods firms. Werning (2015) and Bilbiie (2019, 2020) argue that the cyclicality and distribution of dividends affect the amplification in HANK models. In particular, in New Keynesian models, it is well known that dividends are countercyclical, which is at odds with the data. If dividend payments are given to households that have high MPCs and are exposed to unemployment risk, the amplification that arises from high MPCs and precautionary motives dampens. Following most studies that stress amplification in HANK, I assume that dividends are given to households that have low MPCs and are not exposed to unemployment risk.[5] Therefore, the amplification engendered by MPC heterogeneity and unemployment risk is operative in my model.

2.2 *Final goods firm*

A representative final goods firm combines differentiated wholesale goods and produces a final good according to a Dixit–Stiglitz aggregator,

$$Y_t = \left[\int_0^1 Y_{h,t}^{\frac{1}{1+\eta_t^p}} \, \mathrm{d}h \right]^{1+\eta_t^p}, \tag{2.3}$$

[5] Galí, López-Salido, and Vallés (2007) assume patient households receive dividends. Ravn and Sterk (2017), Challe et al. (2017), and Challe (2020) assume that dividends are given to agents that do not precautionary-save.

where $\eta_t^p > 0$ denotes the price markup in the wholesale goods market and evolves according to

$$\log(1+\eta_t^p) = (1-\rho_{\eta^p})\log(1+\eta^p) + \rho_{\eta^p}\log(1+\eta_{t-1}^p) + \epsilon_t^{\eta^p}, \quad \epsilon_t^{\eta^p} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{\eta^p}{}^2).$$

The final goods firm minimizes the expenditures on wholesale goods, taking the prices as given, subject to the aggregator (2.3). Its optimal choices imply the demand specification for wholesale good h,

$$Y_{h,t} = \left(\frac{P_{h,t}}{P_t}\right)^{-\frac{1+\eta_t^p}{\eta_t^p}} Y_t, \tag{2.4}$$

where $P_{h,t}$ is the price of wholesale good h in period t. P_t denotes the aggregate price index, which is given by

$$P_t = \left(\int_0^1 P_{h,t}{}^{-\frac{1}{\eta_t^p}}\,\mathrm{d}h\right)^{-\eta_t^p}. \tag{2.5}$$

2.3 *Wholesale firms*

Wholesale firm $h \in [0, 1]$ converts each intermediate good into a specialized good according to

$$Y_{h,t} = Y_t^I - A_t F,$$

where F represents the fixed cost of production, and Y_t^I denotes the intermediate good. A_t is included to ensure the existence of a balanced growth path and is defined below. Firm h's real profits are given by

$$\frac{D_{h,t}}{P_t} = \frac{P_{h,t}}{P_t}Y_{h,t} - \frac{MC_t}{P_t}Y_t^I = \left(\frac{P_{h,t}}{P_t} - \frac{MC_t}{P_t}\right)Y_{h,t} - \frac{MC_t}{P_t}A_t F,$$

where MC_t is the price of intermediate goods and is interpreted as the nominal marginal cost for wholesale firms.

Wholesale firms are subject to nominal price rigidity. I introduce nominal price rigidity following Calvo (1983), so that, in every period, a fraction ξ_p of the firms index their prices to lagged inflation according to

$$P_{h,t} = \pi_{t-1}^{\iota_p}\pi^{1-\iota_p}P_{h,t-1}.$$

The remaining fraction of the firms choose their period t optimal price P_t^* by maximizing the present discounted value of expected future real profits. Formally,

$$\max_{P_t^*} \mathbb{E}_t \sum_{s=0}^{\infty} (\xi_p)^s \left(\frac{1}{\prod_{k=1}^{s} \frac{R_{t+k-1}}{\pi_{t+k}}}\right)\left\{\left[\frac{P_t^*\chi_{t,t+s}}{P_{t+s}} - \frac{MC_{t+s}}{P_{t+s}}\right]Y_{h,t+s} - \frac{MC_t}{P_t}A_t F\right\},$$

subject to the demand constraint (2.4), where $\chi_{t,t+s} = \prod_{k=1}^{s}\pi_{t+k-1}^{\iota_p}\pi^{1-\iota_p}$ if $s \geq 1$, and $\chi_{t,t} = 1$.

2.4 *Intermediate goods firms*

A representative competitive intermediate goods firm produces with technology

$$Y_t^I = (A_t \widetilde{A}_t)^{1-\alpha}(u_{k,t}K_{t-1})^{\alpha} n_t^{1-\alpha}, \tag{2.6}$$

where K_{t-1} denotes the installed capital, and $u_{k,t}$ is the capital utilization rate. A_t is the nonstationary aggregate technology, and its growth rate $\mu_t = \frac{A_t}{A_{t-1}}$ evolves according to

$$\log\mu_t = (1-\rho_\mu)\log\mu + \rho_\mu \log\mu_{t-1} + \epsilon_t^\mu, \quad \epsilon_t^\mu \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_\mu^2).$$

$\widetilde{A}_t$ is the stationary technology, and its process is

$$\log\widetilde{A}_t = \rho_{\widetilde{A}} \log\widetilde{A}_{t-1} + \epsilon_t^{\widetilde{A}}, \quad \epsilon_t^{\widetilde{A}} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{\widetilde{A}}^2).$$

In every period, the firm posts vacancies. Vacancies and households that seek jobs are randomly matched according to the aggregate matching function

$$M(\widetilde{u}_t, \mathrm{v}_t) = \overline{M}(\widetilde{u}_t)^{\gamma}(\mathrm{v}_t)^{1-\gamma}, \tag{2.7}$$

where $M(\widetilde{u}_t, \mathrm{v}_t)$ is the number of matches in period t, $\widetilde{u}_t$ is the mass of job seekers, and v_t is the measure of vacancies posted by the firm. $\overline{M}$ is the matching efficiency, and γ represents the elasticity of matches with respect to job seekers. The mass of job seekers in period t consists of the mass of unemployment carried over from the previous period and the mass of existing employment relationships that are severed with probability $\rho_{x,t}$ at the beginning of period t. Formally,

$$\widetilde{u}_t = u_{t-1} + \rho_{x,t} n_{t-1}.$$

The job separation rate $\rho_{x,t}$ evolves according to[6]

$$\rho_{x,t} = \frac{1}{1 + \exp(\overline{\rho}_x - \widetilde{\rho}_{x,t})}, \tag{2.8}$$

where

$$\widetilde{\rho}_{x,t} = \rho_{\rho_x} \widetilde{\rho}_{x,t-1} + \epsilon_t^{\rho_x}, \quad \epsilon_t^{\rho_x} \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_{\rho_x}^2).$$

Given the matching function, the probability that a vacant job is filled and the probability that a job seeker finds a job are

$$\lambda_t = \overline{M}(\mathrm{v}_t/\widetilde{u}_t)^{-\gamma} \quad \text{and} \quad f_t = \overline{M}(\mathrm{v}_t/\widetilde{u}_t)^{1-\gamma},$$

respectively. In period t, a household that is severed from an employment relationship is assumed to find a job immediately with probability f_t. Therefore, the transition rate from period $t-1$ employment to period t unemployment is $\rho_{x,t}(1-f_t)$, which I label as the job-loss rate. This probability measures the degree of unemployment risk faced by

[6]This functional form ensures $\rho_{x,t} \in [0, 1]$.

employed households. Moreover, $1-f_t$ is the transition rate from period $t-1$ unemployment to period t unemployment and measures the degree of unemployment risk faced by unemployment households. Using these transition rates, I obtain the law of motion for the unemployment rate

$$u_t = (1-f_t)u_{t-1} + \rho_{x,t}(1-f_t)n_{t-1}. \tag{2.9}$$

In every period, vacancies are filled with probability λ_t. Therefore, the evolution of employment that the representative intermediate goods firm faces is

$$n_t = (1-\rho_{x,t})n_{t-1} + \lambda_t \mathrm{v}_t. \tag{2.10}$$

The firm owns capital, invests, and chooses the capital utilization rate $u_{k,t}$. The cost of capital utilization is $\Psi(u_{k,t})$ per unit of physical capital, where $\Psi(u_{k,t}) = \rho^{u_k}\frac{u_{k,t}^{\frac{1}{1-\psi}}-1}{\frac{1}{1-\psi}}$. In the steady state, $u_k = 1$, $\Psi(1) = 0$, and $\frac{\Psi''(1)}{\Psi'(1)} = \frac{\psi}{1-\psi}$, where $\psi \in (0,1)$. Aggregate physical capital K_t accumulates according to

$$K_t = v_t I_t\left[1 - \frac{s''}{2}\left(\frac{I_t}{I_{t-1}} - \mu\right)^2\right] + (1-\delta)K_{t-1}, \tag{2.11}$$

where δ denotes the depreciation rate, I_t denotes investment, and s'' captures the convex investment adjustment cost. v_t is the marginal efficiency of investment, which evolves according to

$$\log v_t = \rho_v \log v_{t-1} + \epsilon_t^v, \quad \epsilon_t^v \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_v^2).$$

Taking into account (2.10) and (2.11), the firm maximizes the present discounted stream of profits. Formally,

$$\max_{n_t, \mathrm{v}_t, I_t, u_{k,t}, K_t} \mathbb{E}_t \sum_{s=0}^{\infty}\left(\frac{1}{\prod_{k=1}^{s}\frac{R_{t+k-1}}{\pi_{t+k}}}\right)\frac{D_{t+s}^I}{P_{t+s}},$$

where

$$\frac{D_{t+s}^I}{P_{t+s}} = \left(\frac{MC_{t+s}}{P_{t+s}}\right)Y_{t+s}^I - \frac{W_{t+s}}{P_{t+s}}n_{t+s} - A_t\kappa \mathrm{v}_{t+s} - I_t - \Psi(u_{k,t})K_{t-1}.$$

The costs for the firm are the wage bill paid to all employees, expenditures on investment goods, and forgone resources from searching for new employees and utilizing capital. κ is the cost associated with posting a vacancy.

Nominal wages In the presence of frictional labor markets, there is a surplus in the employment relationship. An intermediate goods firm's surplus is the expected profits from hiring a new employee net of the searching costs of finding a new employee. A household's surplus comes from the wage income net of the cost of unemployment. To maintain the employment relationship, a wage-setting rule must be bilaterally efficient so that both household and firm surpluses are positive. Moreover, a wage-setting rule should incorporate the possibility of nominal wage stickiness.[7] As in Den Haan, Rendahl, and Riegler (2018), nominal wage stickiness affects the cyclicality of unemployment, and thus unemployment risk. Accordingly, it affects the evaluation of unemployment risk on aggregate consumption volatility, which is a central contribution of the present paper.

One popular bilaterally efficient wage-setting approach is Nash bargaining. However, it assumes that wages are renegotiated every period, and so wages become too flexible. To ensure bilateral efficiency and to incorporate wage stickiness in a flexible manner, I choose a wage rule adopted by Challe et al. (2017),[8]

$$W_t = W_{t-1}^{\iota_w}\left(P_t A_t \overline{w}\left(\frac{n_t}{n}\right)^{\xi_w}\right)^{1-\iota_w}, \tag{2.12}$$

where $P_t A_t$ is a scaling factor that ensures the existence of a balanced growth path, and $\overline{w}$ is a constant that ensures the existence of a steady-state real wage in the detrended equilibrium. $\xi_w \in [0, 1]$ is the elasticity of the nominal wage with respect to the deviation of employment from its steady state, and ι_w is the degree of nominal wage indexation. In the model, nominal wage stickiness, together with price stickiness, determines real wage stickiness.[9]

2.5 *Government*

In the model, the government is the only provider of liquid assets. The government raises tax revenue and issues liquid assets to finance government expenditures on unemployment insurances, government purchases, and interest payments on liquid assets. The government budget constraint at each date is

$$B_{t+1}^g + \tau\left(\frac{W_t}{P_t}n_t + b^u\frac{W_t}{P_t}u_t\right) = b^u\frac{W_t}{P_t}u_t + G_t + \frac{R_{t-1}}{\pi_t}B_t^g,$$
$$B_{t+1}^g = A_t\overline{B}^g,$$

where B_{t+1}^g and G_t denote the supply of liquid assets available in the economy and government purchases, respectively. $\overline{B}^g$ is the detrended stock of liquid assets, premultiplied by scaling factor A_t to ensure the existence of a balanced growth path. I assume that the government cannot adjust the supply of liquid assets, and hence $\overline{B}^g$ is constant. Although governments issue debt to finance spending in practice, especially during recessions, government debt accounts for only a very small fraction of household liquid assets. Kaplan, Moll, and Violante (2018) document that most of the liquid assets are

[7] See Den Haan, Rendahl, and Riegler (2018) for ample evidence on nominal wage stickiness.

[8] It is not unrealistic to assume that a wage does not respond to changes in cash-on-hand (e.g., unemployment insurance and asset holdings). Card, Chetty, and Weber (2007), Lalive (2007), van Ours and Vodopivec (2008), and Le Barbanchon (2016) find that unemployment insurances do not have a significant effect on wages.

[9] Challe et al. (2017) show that a sequence of real wages predicted in their estimated model lies within the bargaining set over their sample period. This is also the case in my estimated model.

held as deposits in financial institutions and that the share of government bonds in households' liquid assets is less than 10%. Therefore, interpreting $\overline{B}^g$ as public debt and making it countercyclical in my model might overstate the stock of liquid assets that households use to self-insure.

As in Justiniano, Primiceri, and Tambalotti (2010), government purchases are determined exogenously as a time-varying fraction of output, $G_t = (1 - \frac{1}{g_t})Y_t$, where the government purchase shock g_t follows

$$\log\left(\frac{g_t}{g}\right) = \rho_g \log\left(\frac{g_{t-1}}{g}\right) + \epsilon_t^g, \quad \epsilon_t^g \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_g^2).$$

I assume the monetary policy follows a feedback rule, following the tradition in the New Keynesian literature. The nominal interest rate reacts to the previous nominal interest rate, deviations of inflation from its steady state, and deviations of stationary GDP level from its steady state

$$\log\left(\frac{R_t}{R}\right) = \rho_R \log\left(\frac{R_{t-1}}{R}\right) + (1 - \rho_R)\left[\phi_\pi \log\left(\frac{\pi_t}{\pi}\right) + \phi_X \log\left(\frac{x_t}{x}\right)\right] + \epsilon_t^R, \tag{2.13}$$

where ϵ_t^R is the monetary policy shock with $\epsilon_t^R \overset{\text{iid}}{\sim} \mathcal{N}(0, \sigma_R^2)$. $x_t = \frac{C_t + I_t + G_t}{A_t}$ is stationary GDP, where $C_t = \int_{[0,1-\Omega]} C_{i,t}\, \mathrm{d}i + \Omega C_{H,t}$ is aggregate consumption.

Lastly, by combining the government's and the households' budget constraints, one can obtain the aggregate resource constraint

$$C_t + I_t + G_t = Y_t - A_t \kappa \mathrm{v}_t - \Psi(u_{k,t})K_{t-1}.$$

Clearly, the defined GDP ($\equiv C_t + I_t + G_t$) is not equal to output Y_t, but numerically, the dynamics of GDP and output are similar. I express the model's equilibrium conditions in terms of stationary variables by applying a standard detrending technique. Appendix B describes the full set of equilibrium conditions.

3. Solution method

3.1 *Discretization*

The equilibrium conditions are infinite-dimensional due to the presence of infinite-dimensional objects: decision rules of impatient households and the distribution of households over liquid wealth. These objects are known to pose a challenge in solving incomplete markets models. I use the method proposed by Winberry (2018) to overcome the computational hurdle of this class of models. In particular, I approximate the distribution using a parametric function so that the distribution is summarized by the parameters of the function and the finite number of moments of the distribution. Moreover, I approximate the conditional expectation of the future consumption function using a linear combination of polynomials so that the household decision rule is represented by the coefficients of these polynomials. For further details and the definition of the approximated equilibrium of the model, see Appendix C.

3.2 *Aggregate dynamics*

To solve for aggregate dynamics, I apply a standard technique to the approximated model. That is, I compute the model's stationary equilibrium, an equilibrium with no aggregate shocks. I then linearize the model's equilibrium conditions around their stationary values. Finally, I solve for the dynamics of all variables using a standard method that solves the linear rational expectation models.[10]

4. Estimation

I solve the model using the method outlined above and estimate it using a Bayesian method. This section discusses the calibration, the property of the stationary equilibrium, the data, and the parameter estimates.

4.1 *Calibration*

In this subsection, I describe parameters that are not subject to estimation. The model period is one quarter. The capital share in the production function α is 0.33. The capital depreciation rate δ is 0.02, implying 8% annual depreciation of physical capital. I set the utility function's curvature parameter σ to 2, as in McKay, Nakamura, and Steinsson (2016). I fix the steady-state markup η_p to 0.2, in line with Basu and Fernald (1997). The matching function elasticity to job seekers γ is 0.5, as suggested by Petrongolo and Pissarides (2001).

Unless otherwise noted, the steady-state values of aggregate variables described below correspond to their empirical average over the sample period, from 1964Q2 to 2008Q3. I choose the steady-state nonstationary technology growth rate to match the average GDP growth rate of 0.47% per quarter. The steady-state inflation rate is the average growth rate of the GDP deflator, which is 0.96% per quarter. I set the discount factor of patient households β_H equal to 1.0032 to match the average federal funds rate of 1.58% per quarter.[11] The steady-state ratio of government purchases to GDP is 0.2, which is the average government purchases of consumption and investment goods over GDP. The fixed cost F ensures that the steady-state profits of monopolistic competitive wholesale firms are zero. The steady- state unemployment rate, which corresponds to the average unemployment rate, is 6%. The steady-state job-finding rate f and the job-separation rate ρ_x are determined as follows. I use the quarterly job finding rates constructed by Challe et al. (2017) who follow the approach of Shimer (2005). The average job-finding rate is 0.78, which corresponds to f. Using equation (2.9), I then compute ρ_x. For the matching efficiency $\overline{M}$, I target a quarterly vacancy-filling rate of 0.71, calculated by Den Haan, Ramey, and Watson (2000). The expected cost of hiring an employee κ/λ is calibrated to match 4.5% of quarterly wages, following Hagedorn and Manovskii

[10] Linearization, solving for the dynamics, and estimation were executed using `Dynare`.

[11] One might wonder how β_H can be larger than 1. Note that in the detrended equilibrium, in which all nonstationary variables are divided by nonstationary technology level, the steady-state Euler equation implies $1 = \beta_H \mu^{-\sigma}(R/\pi)$, where μ is the steady-state nonstationary technology growth. Because of the term $\mu^{-\sigma}$, β_H is larger than 1.

(2008), whose calculation is based on the time spent hiring one worker. I then obtain the steady-state real wage from the optimal vacancy-posting condition under the free entry assumption. Under these parameterizations, the model produces the ratio of capital to (annual) output of 3.15 in steady state.

One of the main objectives of this paper is to decompose the aggregate consumption responses into the part that is attributable to precautionary savings against unemployment risk and the part that results from high average MPCs. Accordingly, it is crucial to match the extent of consumption insurance upon unemployment and the average MPC observed in the data. As for the extent of consumption insurance, I target the average consumption difference between employed and unemployed households of 23%, an estimate obtained by Eusepi and Preston (2015) using data from the Consumer Expenditure Survey (CEX).[12] The value implies that unemployed households consume an average of 23% less than employed households. Recent estimates on the decline in consumption during unemployment are smaller. For example, Ganong and Noel (2019) estimate that the spending of unemployed households falls by 9% during the receipt of unemployment insurance and a further 10% after its exhaustion using the JPMorgan Chase panel. Therefore, my chosen value of 23% implies that my model does not understate the role of unemployment risk.

Recall that patient households are fully insured, so their members experience no decrease in consumption when they become unemployed. Hence, the model-based average employed–unemployed (E-U hereafter) consumption difference is $\Omega \times 0 + (1 - \Omega) \times \frac{c_L^e - c_L^u}{c_L^e}$, where 0 is the average E-U consumption difference for the patient group. $\frac{c_L^e - c_L^u}{c_L^e}$ is the average E-U consumption difference for the impatient group, where c_L^e and c_L^u are the steady-state consumption level of employed and unemployed households in the impatient group, respectively. For model-based average MPCs, I compute $\int_{[0,1-\Omega]} MPC_i \, \mathrm{d}i + \Omega MPC_H$, where $MPC_H = 1 - \beta_H \mu^{1-\sigma}$ is the MPC for patient households. The MPC for impatient household i, MPC_i, is calculated based on the slope of the consumption function. I choose the values of the replacement rate b^u and the discount factor of impatient households β_L so that the model-based average MPC and E-U consumption difference jointly match their empirical counterparts.

An important parameter for the model to capture both the MPC heterogeneity and unemployment risk channels is the share of patient households Ω. Ω must be smaller than 0.8 to capture these two channels. If $\Omega = 0.8$, to match the average MPC of 0.2, all impatient households must be at the borrowing limit, living hand-to-mouth. In this scenario, no impatient households operate on the Euler equation, and thus precautionary savings against unemployment risk do not exist. If $\Omega > 0.8$, the model cannot match the average MPC of 0.2, even if all impatient households are at the borrowing limit. I set the share of liquid-wealthy patient households Ω to 0.1, based on the evidence that the top

[12] Note that the average consumption difference between employed and unemployed households is different from temporary consumption losses upon an unemployment shock. See Den Haan, Rendahl, and Riegler (2018) for a discussion of the evidence on the latter. I target the former, as this is a more relevant empirical counterpart to the steady-state consumption difference between employed and unemployed households.

Table 1. Parameters that are not estimated.

Symbol	Description	Value	Target (Source)
δ	Capital depreciation rate	0.02	8% annual depreciation rate
α	Capital share	0.33	
σ	Risk aversion coeff.	2	McKay et al. (2016)
η_p	Markup	0.2	Basu and Fernald (1997)
γ	Matching elasticity	0.5	Petrongolo and Pissarides (2001)
μ	Technology growth	1.0047	GDP growth
β_H	Pat. households discount factor	1.0032	1.58% quarterly federal funds rate
Ω	Share of pat. households	0.1	See the text
ρ_x	Job separation rate	0.23	Average job-finding rate
$\overline{M}$	Matching efficiency	0.74	Average vacancy-filling rate
κ	Cost of posting vacancy	0.06	4.5% of quarterly wages
b^u	Replacement rate	0.58	Average E-U consumption difference
β_L	Imp. households discount factor	0.981	Average quarterly MPC of 0.2

10% hold most of the total liquid wealth (Kaplan, Moll, and Violante (2018)). Table 1 lists the parameters discussed in this subsection.

The total quantity of liquid assets $\overline{B}^g$ is 26% of the annual GDP following Kaplan, Moll, and Violante (2018). Targeting realistic average MPCs leads impatient households to hold less liquid wealth than the total amount of liquid assets. The difference between the total amount of liquid assets and the amount held by impatient households gives the wealth holdings of patient households.[13]

4.2 *Stationary equilibrium*

There are two individual states: unemployed and employed. The distribution for each state is approximated using a smoothed parametric function with the degree of approximation of 4.[14] Figure 1 represents the stationary consumption policy functions and the distribution of impatient households over liquid wealth. The slope of the consumption function is large for households whose liquid asset position is close to the borrowing limit. Consumption of these households responds very strongly to an additional increase in transitory income. As these households live close to hand-to-mouth, they barely respond to changes in future employment prospects. In the next section, I show that the relative contribution of unemployment risk on aggregate consumption fluctuations crucially depends on the fraction of households close to the borrowing limit or the average MPCs. Moreover, given the equal asset position, while consumption is higher for employed households than unemployed households, the slope is higher for unemployed households. Therefore, the difference in average MPCs between employed and

[13]In my setup, the total amount of liquid assets is not essential in determining the amount of liquid assets that impatient households can use for self-insurance. The amount of liquid assets that impatient households hold is primarily disciplined by average MPCs, given the return on liquid assets. The main reason I target the total amount of liquid assets is realism.

[14]Appendix C.4 provides the robustness of the aggregate dynamics to an alternative degree of approximation.

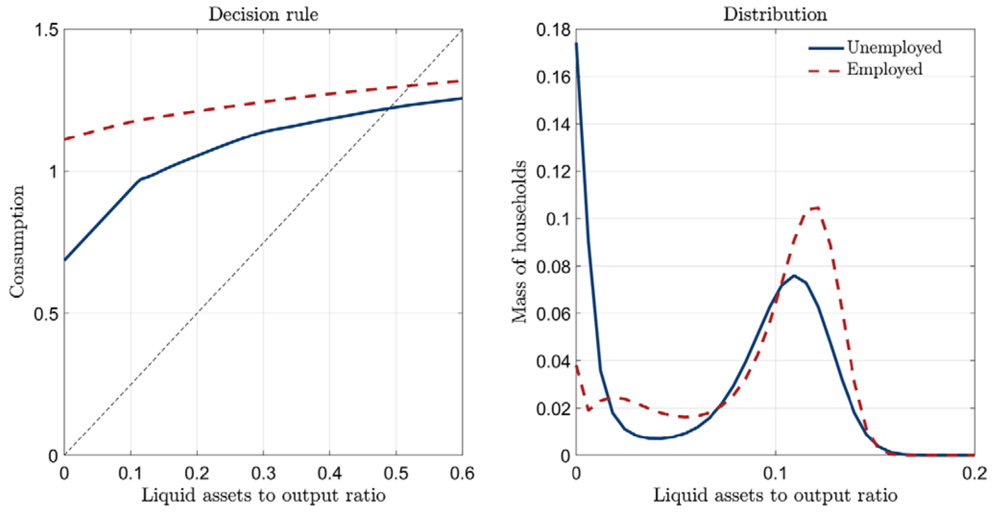

FIGURE 1. Stationary decision rules and liquid wealth distribution.

unemployed households depends on how the employed and unemployed households are distributed over the liquid asset positions. The right panel of the figure indicates that unemployed households are more likely to hit the borrowing constraint, implying a higher MPC for the unemployed households on average.

Table 2 reports the population share, MPCs, liquid wealth share, and consumption share for three groups of households: (1) impatient and unemployed, (2) impatient and employed, and (3) patient. Among the impatient group, unemployed households have a much higher MPC than employed households, as expected from the shape of the wealth distribution. Patient households, which consume based on the permanent income hypothesis, react little to a transitory income change, and thus their MPC is close to zero. Observing the liquid wealth share, it is evident that MPCs are negatively correlated with the liquid wealth share: the group holding the least amount of liquid wealth, that is, the impatient and unemployed group has the largest MPCs. The consumption share reveals that consumption by the impatient and employed group constitutes the largest share of total consumption in the economy, because it accounts for the largest share of the population.

TABLE 2. MPC, population, wealth, and consumption for each group.

Group (Dis. Factor)	Population	MPC	Liquid Wealth Share (Gini)	Cons. Share
Impatient U (β_L)	5.4%	0.8366	0.4% (0.48)	3.9%
Impatient E (β_L)	84.6%	0.1951	7.3% (0.22)	81.6%
Patient (β_H)	10%	0.0014	92.3% (0)	14.5%

Note: Impatient U (E) denotes unemployed (employed) households within the impatient group. Liquid wealth share is the share of total liquid wealth held by each group, while consumption share (“Cons. share” column) is the share of total consumption held by each group. The MPC for patient households is $1 - \beta_H \mu^{1-\sigma}$.

Although the model only targets the average MPC, the model's prediction of MPCs by employment status and liquid wealth position is qualitatively consistent with the data. Using the 2010 Survey of Household Income and Wealth, Kekre (2019) finds that the self-reported annual MPC is higher for unemployed than for employed households. In addition, Broda and Parker (2014) find much stronger consumption responses to the 2008 fiscal stimulus payments among households with low liquid funds, implying a pattern of MPCs declining in liquid wealth. Lastly, the model's prediction of patient households' wealth share of 92% is close to that reported by Kaplan, Moll, and Violante (2018), who show that the top 10% of households holds 86% of the total liquid wealth using data from the Survey of Consumer Finances (SCF).

4.3 *Data and estimation results*

I estimate the remaining structural parameters using the following quarterly US series:[15] the inflation rate, the federal funds rate, the log difference of real per-capita consumption, investment, and government purchases, the wage inflation rate, the job-finding rate, and the job-loss rate. The job-finding rate and the job-loss rate are included because the cyclicality of these two rates directly measures unemployment risk over the business cycle. The inflation rate is the growth rate of the GDP deflator, while the wage inflation rate is the growth rate of average hourly earnings of production and nonsupervisory employees. Real per-capita consumption is the nominal consumption divided by the civilian noninstitutional population (16 years and older) and the GDP deflator. The real series for per-capita investment and government purchases are obtained in the same manner. Consumption corresponds to the sum of nondurables and services, while investment is the sum of consumer durables and total private investment. Government purchases are constructed by adding government consumption and investment. The sample starts from 1964Q2 due to the limited availability of the wage data and ends in 2008Q3, which is the quarter right before the nominal interest rate hits the zero lower bound. As the job-finding and the job-loss rates exhibit trends, I detrend these rates using the filtering method proposed by Hamilton (2018). The remaining series are demeaned before estimation. Appendix D.1 provides more details on the data.

The data series on wage inflation imperfectly match the model's concept of wage inflation due to well-known difficulties in measuring aggregate nominal wages. To address the absence of wage markup shocks and the mismatch between the data and the model in nominal wages, I augment the model wage inflation rates with measurement errors, as in Boivin and Giannoni (2006) and Justiniano, Primiceri, and Tambalotti (2013).[16] I obtain 700,000 draws of parameters to recover the posterior distribution by relying on the Random Walk Metropolis–Hastings algorithm. Table 3 lists the parameter names and the prior and posterior distributions of the estimated parameters. Appendix D.2 includes the convergence diagnostics for the posterior.

[15] The labor market data I use in estimation are job-finding and job-loss rates. In Appendix E.1, I compare the model-based unemployment rate and vacancies with the data.

[16] Justiniano, Primiceri, and Tambalotti (2013) find that the wage inflation dynamics are largely attributable to wage measurement errors.

TABLE 3. Prior and posterior distribution.

Parameter	Description	Prior Dist.			Posterior Dist. (HANK)			Posterior Dist. (RANK)		
		Distribution	Mean	SD	Mean	5%	95%	Mean	5%	95%
s''	Invest. adjustment cost	Gamma	4	1	0.93	0.51	1.33	3.15	2.07	4.25
ψ	Capital utilization cost	Beta	0.5	0.15	0.90	0.83	0.97	0.89	0.82	0.96
ξ_p	Price stickiness	Beta	0.5	0.1	0.52	0.46	0.57	0.69	0.64	0.74
ι_p	Price indexation	Beta	0.5	0.15	0.50	0.27	0.73	0.25	0.09	0.39
ξ_w	Wage flexibility	Gamma	1	0.2	0.73	0.58	0.88	0.63	0.48	0.77
ι_w	Wage indexation	Beta	0.5	0.15	0.74	0.69	0.79	0.77	0.71	0.83
ρ_R	Taylor rule: smoothing	Beta	0.6	0.1	0.75	0.71	0.79	0.70	0.66	0.74
ϕ_π	Taylor rule: inflation	Norm	1.7	0.3	2.89	2.58	3.18	2.09	1.83	2.35
ϕ_X	Taylor rule: GDP	Norm	0	0.3	0.30	0.23	0.38	0.11	0.07	0.15
ρ_{η_p}	Auto. price markup	Beta	0.6	0.1	0.64	0.55	0.74	0.73	0.64	0.83
ρ_μ	Auto. nonstat. tech.	Beta	0.4	0.1	0.26	0.16	0.37	0.35	0.22	0.46
ρ_υ	Auto. MEI	Beta	0.6	0.1	0.75	0.66	0.84	0.59	0.51	0.68
ρ_ζ	Auto. preference	Beta	0.6	0.1	0.97	0.96	0.98	0.94	0.91	0.96
ρ_{ρ_x}	Auto. job-separation	Beta	0.6	0.1	0.81	0.75	0.86	0.81	0.75	0.87
ρ_g	Auto. gov. purchase	Beta	0.6	0.1	0.94	0.92	0.96	0.96	0.95	0.98
$\rho_{\tilde{A}}$	Auto. stat. tech.	Beta	0.6	0.1	0.90	0.88	0.94	0.94	0.92	0.96
$100\sigma_{\eta_p}$	Std price markup	Inv. Gamma	0.15	1	0.88	0.71	1.05	1.34	0.92	1.73
$100\sigma_\mu$	Std nonstat. tech.	Inv. Gamma	1	1	0.42	0.35	0.49	0.42	0.34	0.50
$100\sigma_\upsilon$	Std MEI	Inv. Gamma	0.5	1	2.53	1.92	3.12	4.40	3.03	5.74
$100\sigma_R$	Std mon. policy	Inv. Gamma	0.15	1	0.29	0.25	0.32	0.28	0.25	0.31
$100\sigma_\zeta$	Std preference	Inv. Gamma	1	1	6.52	5.44	7.58	1.45	1.26	1.63
$100\sigma_{\rho_x}$	Std job-separation	Inv. Gamma	1	1	12.6	11.5	13.7	12.7	11.5	13.7
$100\sigma_g$	Std gov. purchase	Inv. Gamma	0.5	1	0.28	0.25	0.30	0.27	0.25	0.30
$100\sigma_{\tilde{A}}$	Std stat. tech.	Inv. Gamma	0.5	1	0.74	0.65	0.83	0.76	0.67	0.85
$100\sigma_w$	Std wage measurement	Inv. Gamma	0.5	0.5	0.27	0.24	0.30	0.24	0.21	0.27

All parameters except for two are standard, and their prior distributions are in line with the literature. The two exceptions are the wage elasticity with respect to employment and the wage indexation, which are embedded in the wage rule (2.12). Because this rule is borrowed from Challe et al. (2017), I adopt their prior distribution for these two parameters. The covariance matrix of the vector of shocks is diagonal. For posterior estimates, I only comment on the parameters that govern price stickiness, nominal wage stickiness, and the responsiveness of the policy rate. These New Keynesian ingredients determine the strength of aggregate demand externality, and thus the degree of amplification in any New Keynesian models.

The posterior mean of the Calvo price stickiness parameter ξ_p implies that prices are adjusted approximately every two quarters, which are lower than what most estimated RANK models imply. In Section 6, I discuss why HANK delivers a lower estimate of price stickiness. The posterior mean of price indexation ι_p is quite high relative to most RANK models, perhaps due to the absence of habit formations in my model. With habit formations, aggregate demand and inflation become persistent, and thus there is less need to rely on price indexation for models to generate the persistent inflation seen in the data. The posterior mean of the wage flexibility parameter ξ_w is lower than its prior mean, and the posterior mean of wage indexation ι_w is higher than its prior mean. These estimates suggest that there is some degree of nominal wage stickiness. The posterior mean of the inflation and GDP coefficients in the monetary policy rule, that is, ϕ_π and ϕ_X, are 2.89 and 0.3, respectively. These coefficient estimates suggest that the monetary policy authority reacts strongly to inflation and economic activity. The value of these coefficients is larger than those in most existing calibrated HANK models, in which inflation and GDP coefficients are around 1.5 and 0, respectively.

5. Unemployment risk and MPC heterogeneity

In this section, I first explore whether aggregate consumption volatilities in HANK and RANK differ. Then I study how much of the difference between consumption volatilities in the two models is attributable to unemployment risk and high average MPCs. The RANK benchmark is obtained from the HANK model by setting the share of patient households to 1. Therefore, the production side, government, monetary authority, and shock processes in RANK are the same as those in HANK. Steady-state prices and aggregate quantities are also the same between the two models. The only difference in RANK compared with HANK is the absence of heterogeneity in the discount factors and consumption levels.

The experiment works as follows. Conditional on the sample information, I use the Kalman smoother to recover historical shocks and state variables from the estimated HANK model. I then exclude the preference shocks and feed the remaining shocks and state variables into the RANK model to generate a counterfactual path of aggregate variables.[17] I compare the volatility of the aggregate variables from the two economies.

[17]The reason for excluding the preference shocks is that these shocks only affect consumption by some households under HANK. Households at the borrowing limit live hand-to-mouth, so their consumption

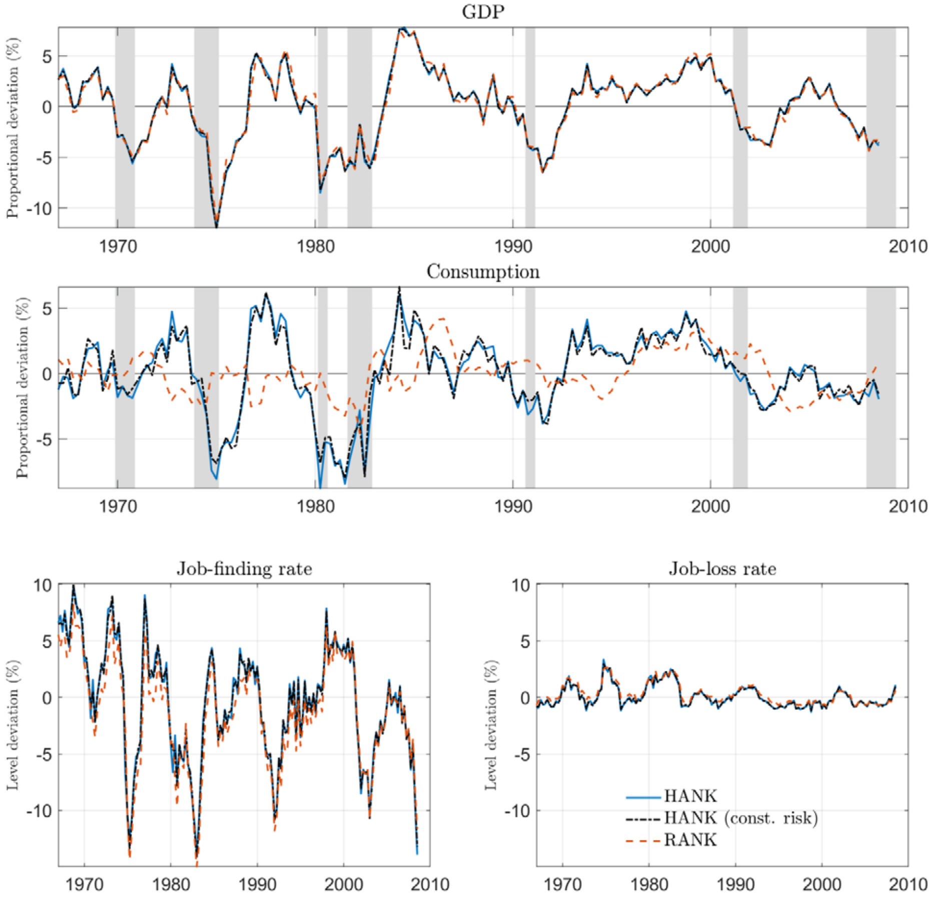

FIGURE 2. Historical dynamics in response to nonpreference shocks. *Note*: For GDP and consumption, the y-axis represents the proportional deviation from the trend produced from the Hamilton filter. For the job-finding and job-loss rates, the y-axis represents the level deviation from the mean (steady state).

Figure 2 visually highlights the comparison of GDP, consumption, the job-finding rate, and the job-loss rate in HANK and RANK. Table 4 compares the standard deviations of these variables over the sample period across models. The parameter values in the two models are identical and are fixed at the posterior mean of the HANK model. In Figure 2, GDP and consumption are expressed in proportional deviation from their respective trends, applying the filtering method proposed by Hamilton (2018). For job-finding and job-loss rates, they are expressed in level deviation from their respective means.

is not affected by a direct shift in time preferences. By contrast, under RANK, preference shocks affect all households, which are all identical. Thus, it may be difficult to interpret that the difference between consumption responses in the two models, conditional on preference shocks, arises from unemployment risk or high average MPCs. Rather, the difference arises largely due to the existence of households that are unresponsive to these shocks in the HANK model. Appendix E.2 compares the impulse responses to a preference

Table 4. Volatility subject to nonpreference shocks.

Model	Standard Deviation				
	GDP	Consumption	Investment	Job-Find. Rate	Job-Loss Rate
HANK	3.539	2.990	6.809	5.212	0.933
HANK (c. risk)	3.521	2.832	6.973	5.199	0.912
RANK	3.448	1.646	10.220	4.886	0.878

Note: For GDP, consumption, and investment, logs are taken and then detrended using the Hamilton filter. HANK (c. risk) denotes the constant risk model in which households perceive a constant probability of becoming unemployed.

Two results stand out. As noted in the figure, the job-loss rate in the HANK model moves little over the business cycle, showing results consistent with Challe et al. (2017) and implying that unemployment risk primarily arises from variations in the job-finding rates. This observation implies that the job-separation rate, $\rho_{x,t} = \frac{s_t}{1-f_t}$, moves little and plays a minor role in unemployment fluctuations, consistent with Shimer (2005). In addition, as noted in the figure and table, consumption is clearly more volatile in HANK than in RANK.

Two channels lead to a more volatile consumption in HANK. Consider recessions as an example. First, aggregate household income falls due to two effects: a fall in real wages and an increase in unemployment. As the average MPC in HANK is higher than in RANK, aggregate consumption drops more in HANK. This is the MPC heterogeneity channel. Second, even if households' employment status does not change, households cut consumption due to a strong precautionary savings motive as their perceived probability of becoming unemployed increases. This is the unemployment risk channel. The table shows that the volatilities of GDP, job-finding rate, and job-loss rate are higher for HANK than for RANK. However, the differences in the volatility of these variables are small, given the large differences in consumption volatility. This is because more volatile consumption in HANK than in RANK requires a larger response of real interest rates to clear the goods market, thus dampening the volatility of investment.

I now assess the importance of each of the two channels in explaining the difference between aggregate consumption volatility in HANK and RANK. To control for the effect of time-varying unemployment risk, I introduce the constant risk model. This model is equal to the baseline HANK model with job-separation and job-finding rates in the Euler equation set to constant, that is, steady state.[18] The other parts remain the same as in the baseline model and, therefore, prices/wages and (un)employment rate move due to general equilibrium forces triggered by aggregate shocks. Given that households' expected income is affected by movements in both wages and labor market variables,

shock in HANK and RANK and shows substantially weaker consumption under HANK due to the existence of constrained households.

[18]Put differently, the constant risk benchmark is a model in which the transition matrix across labor endowments is time invariant. Here, the endowment of labor is 1 if employed and b^u if unemployed. HANK models that adopt the time invariant transition matrix are Kaplan, Moll, and Violante (2018), Hagedorn, Manovskii, and Mitman (2019), and Auclert, Rognlie, and Straub (2020), among many others.

shutting down the time-varying unemployment risk dampens fluctuations in households' expected income relative to those in the baseline HANK model. Comparing the aggregate consumption dynamics in the baseline HANK model with those in the constant risk benchmark allows me to gauge the effect of time-varying unemployment risk on aggregate consumption. I interpret the difference between the consumption dynamics in the constant risk benchmark and the RANK model as the contribution of MPC heterogeneity.

The black dash-dotted line in Figure 2 corresponds to the dynamics produced from the constant risk benchmark in response to the same set of shocks used to simulate HANK and RANK. The parameter values of the constant risk benchmark are equal to those of HANK and RANK. Consumption responses in this benchmark are very similar to those in HANK, implying that closing down a precautionary motive from unemployment risk does not alter the consumption responses much. I use the standard deviations of consumption in Table 4 to answer what fraction of the aggregate consumption response is driven by a precautionary motive from unemployment risk or MPC heterogeneity. A simple calculation indicates that the contribution of unemployment risk to the difference between consumption volatilities in HANK and RANK is $\frac{(2.990-2.832)}{(2.990-1.646)} \times 100 \approx 12\%$. 88% of the difference is attributable to MPC heterogeneity.

It is noteworthy to compare my results with those in the existing papers such as Ravn and Sterk (2017), Challe et al. (2017), and Challe (2020), among many others. The key exercise in my work is to decompose the differences in aggregate consumption fluctuations between HANK and RANK into the part that arises from time-varying unemployment risk and the part that arises from MPC heterogeneity. Of these two parts, I find that the one contributed by time-varying unemployment risk is smaller. The existing papers do not decompose the differences between HANK and RANK into these two parts; in fact, they present a HANK model in which the average MPCs are counterfactually low and interpret the differences between HANK and RANK as the effect of time-varying unemployment risk. In contrast, these differences in my work capture the summed effects of time-varying unemployment risk and MPC heterogeneity. As discussed below, I show that the effect of time-varying unemployment risk crucially depends on average MPCs.

Alternative parameterizations and targets I investigate whether the small contribution of unemployment risk identified above is robust to alternative parameterizations and targets. I first consider parameters that are only relevant to the dynamics of my HANK model. Then I consider various steady-state targets.

Table 5 reports the standard deviations of GDP, consumption, and investment in HANK, RANK, and the constant risk model under three parameterizations. The same shocks that produced Table 4 are used to compute the standard deviations. The first case assumes a less responsive monetary policy rule than the estimated one. In particular, I set inflation and GDP coefficients, ϕ_π and ϕ_X, to 1.2 and 0, respectively, keeping the rest of the parameters equal to their estimates. The table shows that GDP and its components are more volatile in all three economies under the less responsive monetary policy than those under the estimated monetary policy. Regarding the decomposition of aggregate consumption volatility, the main result holds: unemployment risk accounts for 12% of the difference between consumption volatilities in HANK and RANK, which is small.

Table 5. Volatility under alternative parameterizations.

Parameterization	Model	Standard Deviation			U Risk Contribution
		GDP	Consumption	Investment	
Less responsive monetary policy	HANK	6.398	6.186	8.777	
	HANK (c. risk)	6.311	5.664	9.507	12%
	RANK	5.642	1.728	14.373	
High adjustment cost	HANK	2.473	3.574	3.171	
	HANK (c. risk)	2.471	3.478	3.199	14%
	RANK	2.437	2.899	4.191	
High wage rigidity	HANK	4.520	3.686	8.137	
	HANK (c. risk)	4.440	3.234	8.570	22%
	RANK	4.290	1.627	11.909	

Note: For GDP, consumption, and investment, logs are taken and then detrended using the Hamilton filter. HANK (c. risk) denotes the constant risk model in which households perceive a constant probability of becoming unemployed. The "U risk contribution" column represents the contribution of unemployment risk to the difference between the standard deviation of consumption in HANK and RANK.

The second parameterization considers the role of high investment adjustment costs. Because the investment adjustment costs parameter is quite low in the estimated HANK, I now impose much higher investment adjustment costs by letting $s'' = 20$, without changing the other parameters. As noted in the table, the high adjustment cost increases consumption variability in HANK, as aggregate resources are allocated more into consumption than costly investment. I find robust results regarding the decomposition: the relative contribution of unemployment risk is small, explaining 14% of the difference between consumption volatilities in HANK and RANK.

The third parameterization is related to highly sticky wages. I assume extreme nominal wage stickiness by letting the parameter of wage flexibility ξ_w equal to 0, keeping the other parameters equal to their estimates. The table shows that extremely high wage stickiness leads to more volatile GDP and its components than the estimated wage stickiness. As in Den Haan, Rendahl, and Riegler (2018), more rigid wages induce a more volatile unemployment rate and household income, making consumption more volatile via high MPCs and unemployment risk. It turns out that the relative contribution of unemployment risk increases more than under the estimated wage stickiness, accounting for 22% of the difference between consumption volatilities in HANK and RANK. However, MPC heterogeneity is still the dominant force that drives consumption fluctuations.

Next, I redo the consumption decomposition exercise under various steady-state targets. Table 6 reports the results associated with these targets. Again, I use the same shocks that produced Table 4 when computing the standard deviations. First, I consider two alternative values of the average employed–unemployed consumption differences: 10% and 30%. When targeting these values, I recalibrate the model so as to maintain

Table 6. Volatility under alternative steady-state targets.

Parameterization	Model	Standard Deviation			U Risk Contribution
		GDP	Consumption	Investment	
E-U consumption difference of 10%	HANK	3.528	2.863	6.972	
	HANK (c. risk)	3.524	2.824	7.019	3%
	RANK	3.448	1.646	10.220	
E-U consumption difference of 30%	HANK	3.556	3.048	6.745	
	HANK (c. risk)	3.514	2.749	7.098	21%
	RANK	3.448	1.646	10.220	
Liquid wealth to GDP ratio of 0.35	HANK	3.554	3.043	6.863	
	HANK (c. risk)	3.539	2.912	7.003	9%
	RANK	3.448	1.646	10.220	
Average MPC of 0.1	HANK	3.521	2.421	7.417	
	HANK (c. risk)	3.472	2.036	8.373	50%
	RANK	3.448	1.646	10.220	

Note: For GDP, consumption, and investment, logs are taken and then detrended using the Hamilton filter. HANK (c. risk) denotes the constant risk model in which households perceive a constant probability of becoming unemployed. The “U risk contribution” column represents the contribution of unemployment risk to the difference between the standard deviation of consumption in HANK and RANK.

the average quarterly MPC of 0.2. 10% is on the low side of empirical estimates and implies that the costs of unemployment measured in consumption losses are low. This case tends to reduce precautionary motives from unemployment risk, explaining only 3% of the difference between variations in consumption in HANK and RANK. Raising the value to 30% enhances the relative role of unemployment risk to 21%, which is still low.

Second, without changing the other targets, I increase the ratio of liquid assets to GDP to 0.35, which is the average government debt to GDP ratio over 1964Q2–2008Q3. This value is higher than the one measured by Kaplan, Moll, and Violante (2018), which I used for the baseline calibration. As the amount of liquid wealth held by impatient households must be small to match the empirical average MPC, patient households hold the additional supply of liquid wealth. In this case, the share of patient households' consumption in aggregate consumption grows larger. Hence, the effect of precautionary motive by impatient households on aggregate consumption is smaller, with unemployment risk explaining 9% of the differences between consumption volatilities in HANK and RANK, compared to 12% under the baseline calibration.

Lastly, I reduce the average MPC to 0.1. This experiment highlights how the relative contribution of unemployment risk changes depending on the average MPCs. The environment under the low MPC calibration resembles the environment presented in the existing studies such as Ravn and Sterk (2017), Challe et al. (2017), and Challe (2020) in that average MPCs are counterfactually low. Their models assume that all unemployed households live hand-to-mouth and that all employed households save in all periods.

This implies a low average MPC, which is close to the average unemployment rate. Table 6 shows that, under the average MPC of 0.1, unemployment risk explains 50% of the difference between consumption variations in HANK and RANK, compared with 12% under the baseline MPC.

Why is the role of unemployment risk in driving aggregate consumption so small when MPCs are high? In my model, households' wealth positions are heterogeneous, and so they differ in the degree of precautionary motive. By virtue of high average MPCs, a large mass of households is close to the borrowing limit. These households have little incentive to precautionary-save in the face of unemployment risk. For example, when the probability of becoming unemployed drops, precautionary savings motives for households with a relatively large amount of liquid wealth decrease, causing them to dissave and spend more today. Such an amount of wealth to deplete barely exists for households close to the borrowing limit. These wealth-poor households can spend more today if their income rise is realized today. This spending relates to a mechanical spending effect, that is, the MPC channel. Therefore, as the average MPC increases, the relative effect of precautionary motives from unemployment risk on aggregate consumption weakens.

In sum, I find that unemployment risk, the only source of countercyclical earnings risk in this model, plays a minor role in shaping aggregate consumption dynamics as long as realistic average MPCs are matched. This finding is robust to alternative parameterizations and steady-state targets considered in this section. However, this result does not necessarily imply that countercyclical earnings risk is not important. Recent empirical literature documents other sources of earnings risk that might be unrelated to unemployment (Guvenen, Ozkan, and Song (2014)). By omitting other sources of risk, my analysis could understate the role of countercyclical earnings risk in driving aggregate consumption fluctuations.[19]

6. Model fits: HANK vs. RANK

In this section, I evaluate how well the estimated HANK model explains the aggregate data compared with the estimated RANK model. To do so, I first compare the HANK and RANK models with respect to parameter estimates and aggregate shock decompositions.

Parameter estimates and shock decompositions I estimate the RANK model using the same aggregate series that I used to estimate the HANK model. Table 3 includes the posterior distributions of structural parameters in RANK. A notable difference between HANK and RANK comes from the price and wage stickiness and the stance of monetary policy. The posterior mean for the parameter of goods price stickiness ξ_p is 0.69 in RANK, implying that prices are adjusted approximately every three quarters, while 0.52 in HANK. The parameter of nominal wage flexibility ξ_w is 0.63 in RANK and 0.73 in

[19]Other important works that model cyclical earnings risk that is unrelated to unemployment risk include Acharya and Dogra (2020), Werning (2015), and Bilbiie (2019). They analytically study the implications of cyclical earnings risk on various issues in monetary economics such as determinacy and the forward guidance puzzle.

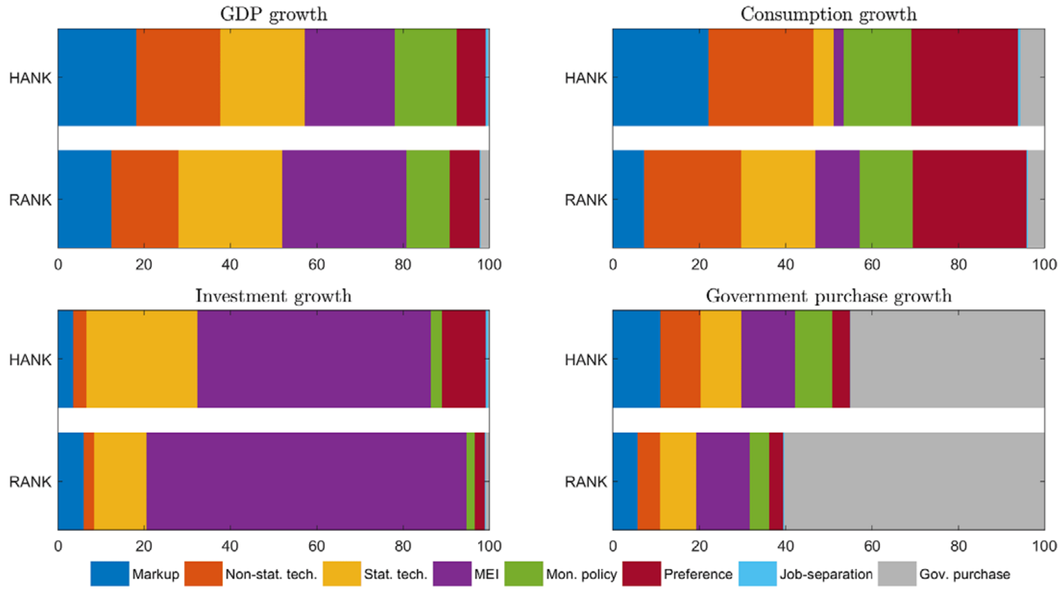

FIGURE 3. Variance decompositions of GDP growth and its components. *Note*: Unconditional variance decomposition in the estimated HANK and RANK. Decomposition is computed at the posterior mean and expressed in percentages.

HANK. The inflation and GDP coefficients in the monetary policy rule are $\phi_\pi = 2.09$ and $\phi_X = 0.11$ in RANK, while $\phi_\pi = 2.89$ and $\phi_X = 0.30$ in HANK. In summary, the RANK model delivers a higher degree of price and wage stickiness and a less responsive monetary policy than its HANK counterpart.

Next, I investigate the contributions of aggregate shocks to the volatility of the growth rates of GDP, consumption, investment, and government purchases in HANK and RANK. Figure 3 reports the unconditional variance decomposition in the two economies, evaluated at their respective posterior mean. As in Justiniano, Primiceri, and Tambalotti (2010), in RANK, the MEI shock accounts for the largest fraction of investment and GDP fluctuations. However, this shock accounts for only a modest fraction of consumption fluctuations, which are driven mainly by preference shocks. In HANK, the contribution of the MEI shock to GDP fluctuations decreases, while the contribution of supply shocks (markup and the two productivity shocks) to its fluctuations increases compared with RANK. The differences in the importance of supply shocks between HANK and RANK arise from nominal rigidities. As goods' prices, wages, and the monetary policy rate are estimated to be more flexible in HANK than in RANK, supply shocks matter more in explaining GDP and its components in HANK. In contrast, the MEI shock, which is a demand shock, matters less.

A natural question that arises from the large contribution of the MEI shock to the variance of GDP growth is whether this shock can generate a positive comovement of consumption with investment, employment, and GDP in HANK and RANK. It is well known that representative-agent macro models with standard preferences have hard times explaining this conditional comovement for the following reasons. When prices are flexible, in response to a shock that raises the demand for investment goods, the market forces work to drive up goods prices, thus depressing consumption. The comovement can arise if goods prices are very rigid, or the nominal interest rate is unresponsive. Extremely sticky goods prices dampen a fall in consumption and so allow the investment boom to induce more expansion in output, increasing households' permanent income and consumption. If the nominal interest rate is invariant, a rise in goods prices results in a fall in the real interest rate, causing consumption to increase. However, in most estimated RANK models, such an extreme degree of nominal rigidity and interest rate inertia does not emerge, and hence the comovement problem persists (Justiniano, Primiceri, and Tambalotti (2010)).

Departing from the representative-agent assumption gives more leeway to achieve the conditional comovement. For instance, Furlanetto, Natvik, and Seneca (2013) and Auclert, Rognlie, and Straub (2020) emphasize the investment-consumption complementarity when MPCs are high: as investment rises, the additional output leads to a rise in employment and labor income, increasing consumption by high-MPC households. It is natural to ask whether my version of the HANK model, which has high MPCs and unemployment risk, can also generate the conditional comovement. Figure 4 compares the impulse responses to a positive MEI shock produced from three models: HANK, RANK, and the constant risk model. Here, the responses in HANK and RANK are computed at their respective posterior mean. For the constant risk model, the parameter values are the same as those in HANK. As observed in the figure, under RANK, investment, employ-

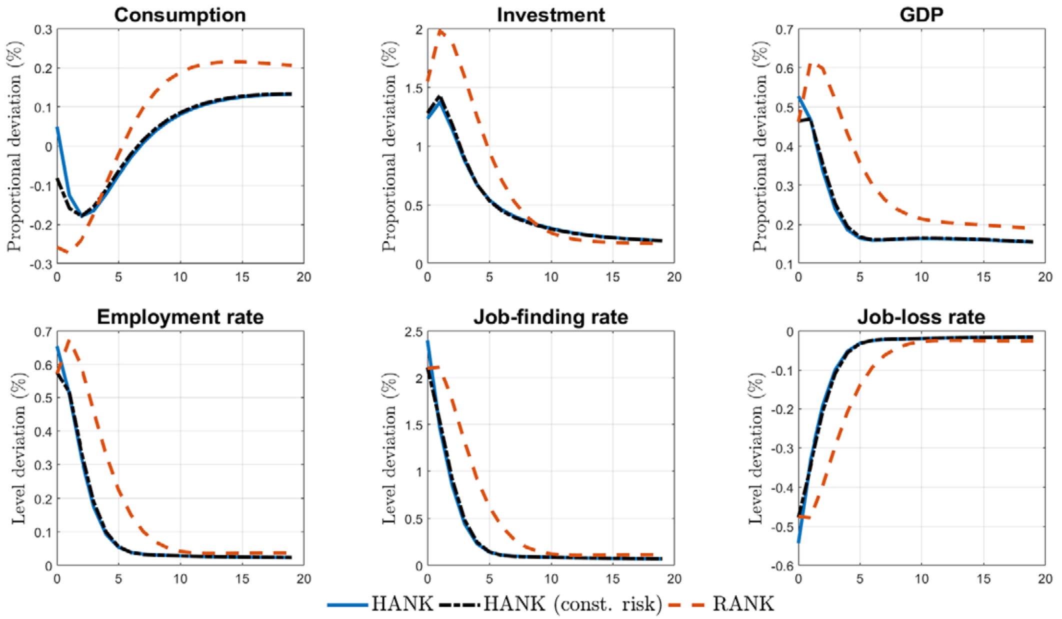

FIGURE 4. IRFs to a positive MEI shock. *Note*: The x-axis measures the time horizon in quarters. For the employment, job-finding, and job-loss rates, the y-axis represents the level deviation from a steady state. For the other variables, the y-axis represents the proportional deviation from a steady state.

ment, and GDP rise, while consumption falls, illustrating a pattern that appears in many estimated RANK models.[20]

In the estimated HANK economy, except for the impact period, consumption drops for six quarters, implying that the short-run comovement problem remains. Why do my results differ from those in Furlanetto, Natvik, and Seneca (2013) and Auclert, Rognlie, and Straub (2020)? The reason for this contrast stems from the more flexible prices and responsive monetary policy in my HANK model. In particular, Furlanetto, Natvik, and Seneca (2013) and Auclert, Rognlie, and Straub (2020) assume 1.5 for inflation coefficient and 0 for GDP coefficient in the Taylor rule, implying much stickier nominal interest rates than in my estimated model. In addition, the estimate of Calvo price stickiness by Auclert, Rognlie, and Straub (2020), which they obtain by matching the impulse responses from an identified monetary policy shock, is 0.93. Such a high degree of price stickiness is necessary for their model to match very persistent and small responses of inflation in the data. Thanks to highly sticky goods prices and nominal interest rates, Furlanetto, Natvik, and Seneca (2013) and Auclert, Rognlie, and Straub (2020) find strong feedback from investment to consumption via high MPCs. However, by estimating a HANK model using the full-information method, instead of calibrating or estimating parameters by matching impulse responses from a single shock, I find that goods prices and nominal interest rates are not sticky enough to generate the conditional comovement between consumption and investment.[21]

For completeness, I explore how much of consumption responses conditional on the MEI shock in HANK are attributable to unemployment risk. Because all parameter values in the constant risk model are the same as those in HANK, comparing the responses in both models enables one to assess the importance of unemployment risk. Figure 4 shows that the constant risk model (black dash-dotted line) predicts slightly more negative short-run consumption. However, it is hard to say that consumption responses in the two models are strikingly different, given that both models fail to deliver a positive comovement between consumption and investment in the short-run. Again, consistent with the analysis in the previous section, the precautionary savings motive

[20]The impulse responses in the figure are very persistent. This feature arises from the fact that capital has a high stock-flow ratio. Consider an unexpected one-time temporary MEI shock in period t for the sake of intuition. An increase in investment in period t raises the capital used in period $t+1$. Because capital has a high stock-flow ratio, capital remains high in subsequent periods, returning to a steady state very slowly. As capital is a state variable, its slow movement causes other variables to converge slowly as well. However, if there is a households' endogenous labor-leisure choice as in Justiniano, Primiceri, and Tambalotti (2010), capital would converge faster. This is because, in times of high capital stock, households reduce the hours worked due to the wealth effect. Thus, the reduced income causes households to invest less. Even in the absence of an endogenous labor-leisure choice, a higher capital depreciation rate lowers the stock-flow ratio, and thus leads to a less persistent movement of capital and employment.

[21]In the previous version, I estimated a HANK model in which the monetary policy rule includes four-quarter inflation and GDP growth as in Justiniano, Primiceri, and Tambalotti (2013) instead of one-quarter inflation and GDP level. In this case, the monetary policy is estimated to be less responsive than in the current version, generating the conditional comovement. Sala, Soderstrom, and Trigari (2010) (in their footnote 26) also argue that this type of monetary policy rule increases the role of the MEI shock in explaining GDP and its components. Regardless of the monetary policy rule, I find the estimates of price and interest rate stickiness in HANK to be lower than those in RANK.

against unemployment risk has a small effect on aggregate consumption responses. The high average MPC is the main determinant for aggregate consumption responses.

Model fits I now evaluate the fit of the HANK model in explaining the aggregate data. The data favors RANK over HANK, with the log marginal likelihoods being -1666.2 for HANK and -1546.6 for RANK. Why does the fit of the HANK model fall behind compared to the RANK model? The reason is that the empirical average MPC target in HANK generates counterfactually volatile aggregate consumption, making it more difficult for the estimated model to match the persistence of the aggregate data. I label this problem of targeting the average MPC as the MPC puzzle. To understand the problem further, it is useful to compare the model-based and empirical autocorrelations of the observable variables. These moments are displayed in Figure 5. Black dotted lines represent the moments predicted from HANK, while red dash-dotted lines indicate those predicted from RANK, computed at their respective posterior mean. The solid blue lines and the dashed blue lines correspond, respectively, to the empirical correlations and 95% confidence intervals.

From the figure, we see that both HANK and RANK do not perfectly capture the decaying autocorrelation structure in the data. One reason for this imperfect match is that, in this paper, I do not incorporate habit formations, which are known to improve the models to match the persistence of aggregate data (Christiano, Eichenbaum, and Evans (2005)). Between the two models, it is more difficult for HANK to capture the empirical autocorrelations than RANK. This is because of a lower degree of nominal rigidity

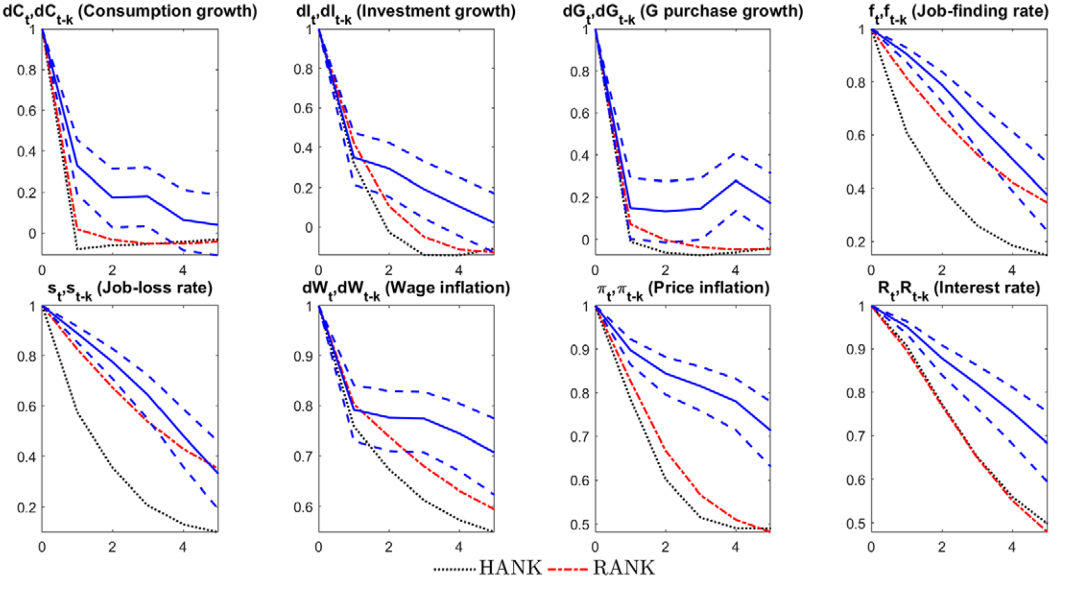

Figure 5. Autocorrelations of observable variables. *Note*: Solid lines are the empirical correlations, and the dashed lines are 95% confidence intervals centered on the empirical correlations. Dotted lines correspond to the HANK model's prediction evaluated at its posterior mean. Dash-dotted lines correspond to the RANK model's prediction evaluated at its posterior mean. k denotes the lag-order.

TABLE 7. Nominal rigidities in HANK and aggregate volatility.

Model	Standard Deviation				
	C Growth	I Growth	P Inflation	Job-Find. Rate	Job-Loss Rate
Baseline	1.41	1.79	0.72	5.80	1.19
High rigidity	1.73	2.04	0.61	8.93	1.94
Data	0.49	2.07	0.59	5.02	0.88

Note: Standard deviations of consumption growth, investment growth, price inflation, the job-finding rate, and the job-loss rate from the estimated HANK model ("Baseline" row), HANK model with a higher degree of nominal rigidity ("High rigidity" row), and data.

and more responsive monetary policy in the estimated HANK than in RANK. The more flexible prices are, the faster the prices return to the steady state, causing quantities to converge faster. Therefore, it is more challenging for HANK than for RANK to match the persistent aggregate data.

The important ingredient that leads to a low degree of nominal rigidity and responsive monetary policy in the estimated HANK model is the high average MPC target. In order to understand this, it is useful to consider what happens if the degree of nominal rigidity and monetary policy coefficients are identical to those in RANK. Specifically, I set the values of parameters related to price and wage stickiness, that is, ξ_p, ι_p, ξ_w, and ι_w, and monetary policy coefficients, that is, ρ_R, ϕ_π, and ϕ_X, equal to those in the estimated RANK model without changing the other parameters. I refer to this as the high rigidity benchmark. Then I compute the theoretical standard deviations of consumption growth, investment growth, price inflation, the job-finding rate, and the job-loss rate from this benchmark and compare those computed from the estimated HANK model.

Table 7 shows that, except for investment growth, the volatility of the selected variables from the estimated HANK ("Baseline" row) is already higher than their empirical volatility ("Data" row). In the high rigidity case ("High rigidity" row), the volatility of inflation is close to its empirical volatility thanks to stickier goods prices, wages, and nominal interest rates. However, with stickier prices, the quantities—consumption growth, investment growth, job-finding rate, and job-loss rate—become more volatile due to the interaction between the average MPC target and a high degree of nominal rigidity. That is, an increase in household consumption can more likely be translated into a higher level of output and employment, which causes consumption to rise further via high MPCs. Highly volatile consumption is the price that must be paid for a high degree of nominal rigidity from a likelihood perspective. Therefore, the likelihood-based estimation leads to a low degree of nominal rigidity and responsive monetary policy in the HANK model to reduce the discrepancy between consumption volatility in the model and in the data.

In sum, estimation reveals that HANK with the empirical average MPC target underperforms RANK in explaining aggregate variables as evidenced from a likelihood evaluation. It is worth noting that this conclusion is drawn under the assumption that the share of patient households Ω is fixed to 0.1. If one ignores matching the empirical MPCs by including Ω in the set of estimated parameters, HANK with estimated Ω can lead to a higher value of marginal likelihoods than the estimated RANK.

7. Extensions: HANK vs. TANK

In the business cycle literature, two-agent New Keynesian (TANK hereafter) models are often used to generate a volatile consumption response because of their tractability in targeting high average MPCs (Galí, López-Salido, and Vallés (2007), Bilbiie (2008), Furlanetto, Natvik, and Seneca (2013)). This section evaluates how a HANK model compares with a TANK model regarding consumption volatility. I then discuss whether the estimated TANK fits the aggregate data better than the estimated RANK, resolving the MPC puzzle.

To construct a TANK model with search and matching frictions, I introduce hand-to-mouth households that coexist with patient households. Similar to patient households, each hand-to-mouth household acts as a family in which all members share the family income. Accordingly, family members do not suffer from unemployment risk. However, unlike patient households, hand-to-mouth households are borrowing constrained, and thus fully consume their income, $(1-\tau_t)(\frac{W_t}{P_t}n_t + b^u\frac{W_t}{P_t}(1-n_t))$, in every period. The supply side and fiscal/monetary policy are the same as in HANK. I calibrate TANK to have a 0.2 share of hand-to-mouth households to match the same quarterly MPC in HANK.

Figure 6 displays the GDP and consumption path predicted from TANK by feeding the same shocks used to produce Figure 2. For comparison, I reproduce the GDP and consumption paths from HANK and the constant risk model. The parameter values of all three models are set to the posterior estimates of HANK. Despite having the same average MPCs, TANK produces less volatile consumption than HANK and the constant risk model. There are two reasons for the difference between the consumption paths in TANK and HANK.

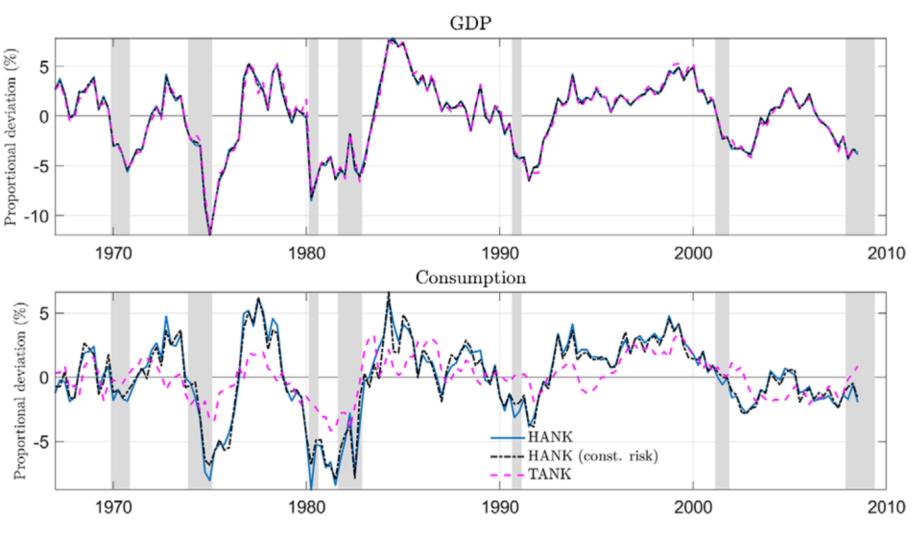

Figure 6. Historical dynamics in response to nonpreference shocks in TANK. *Note*: For GDP and consumption, the y-axis represents the proportional deviation from the trend produced from the Hamilton filter.

First, the HANK model in the current paper features compositional changes between employed and unemployed households. As the average consumption level is higher for employed versus unemployed households, the compositional changes over the business cycle directly affect aggregate consumption volatility; all else equal, during booms, aggregate consumption increases, as there are a higher number of employed households during such periods. This composition effect does not appear in my version of the TANK model, wherein all family members, whether employed or unemployed, consume the same level. The composition effect engendered by differences in the average consumption level across employment statuses is emphasized by Eusepi and Preston (2015), who show that this effect is quantitatively important, even in the absence of nominal rigidity.

Second, as Auclert, Rognlie, and Straub (2018, 2020) and Hagedorn, Manovskii, and Mitman (2019) emphasize, HANK models feature intertemporal MPCs, which are key moments that cause households' current consumption to respond to future income increases. These moments are crucial for HANK models to deliver more amplified consumption responses than TANK. The logic is as follows. Households immediately consume out of an increase in their current income, part of which is saved for future consumption. In the presence of nominal rigidity, this higher future consumption leads to a higher future income, allowing unconstrained households to consume even more today.[22] These intertemporal feedbacks from future income to current consumption are absent in TANK. As hand-to-mouth households respond to an increase in their current income only, TANK delivers smaller amplification. As shown in Auclert, Rognlie, and Straub (2018, 2020) and Hagedorn, Manovskii, and Mitman (2019), the effect of intertemporal MPCs on aggregate consumption is powerful when the degree of nominal rigidity is high. However, given that the degree of nominal rigidity is relatively low in my estimated HANK model, the effect of intertemporal MPCs might not be as strong as the composition effect discussed above.

Given the rather small amplified aggregate consumption in TANK, it is natural to ask whether the estimated TANK outperforms the estimated RANK based on a marginal likelihood evaluation. The log marginal likelihood for the estimated TANK is -1563.8, which is slightly lower than that for the estimated RANK (-1546.6), implying that TANK is not a solution to the MPC puzzle. The reason for the underperformance of TANK is similar to the reason for the underperformance of HANK. That is, the empirical average MPC target in TANK delivers a more responsive monetary policy rule than RANK in the estimation procedure. As observed in Table 8, focusing on the posterior mean, the inflation and GDP coefficients in the monetary policy rule (ϕ_π and ϕ_X) are higher for TANK relative to RANK, while the degree of price and wage stickiness (ξ_p and ξ_w) in TANK is similar to that in RANK. Intuitively, for any given inflation rate, a more responsive monetary policy will lead to a faster return of inflation to its steady state. Thus, it is more difficult for the estimated TANK to capture the highly persistent inflation than the estimated RANK.

The share of hand-to-mouth households in my version of TANK is lower than that used in the existing TANK models, which typically assume 0.3–0.5 (Galí, López-Salido,

[22] Constrained households do not respond to future income increases.

Table 8. Prior and posterior distribution in TANK.

Parameter	Description	Prior Dist.			Posterior Dist.		
		Distribution	Mean	SD	Mean	5 %	95 %
s''	Invest. adjustment cost	Gamma	4	1	3.02	1.91	4.13
ψ	Capital utilization cost	Beta	0.5	0.15	0.88	0.80	0.96
ξ_p	Price stickiness	Beta	0.5	0.1	0.69	0.64	0.74
ι_p	Price indexation	Beta	0.5	0.15	0.33	0.15	0.52
ξ_w	Wage flexibility	Gamma	1	0.2	0.68	0.52	0.83
ι_w	Wage indexation	Beta	0.5	0.15	0.78	0.73	0.83
ρ_R	Taylor rule: smoothing	Beta	0.6	0.1	0.74	0.70	0.77
ϕ_π	Taylor rule: inflation	Norm	1.7	0.3	2.40	2.12	2.68
ϕ_X	Taylor rule: GDP	Norm	0	0.3	0.14	0.10	0.18
ρ_{η_p}	Auto. price markup	Beta	0.6	0.1	0.71	0.61	0.82
ρ_μ	Auto. nonstat. tech.	Beta	0.4	0.1	0.37	0.25	0.50
ρ_υ	Auto. MEI	Beta	0.6	0.1	0.65	0.57	0.72
ρ_ζ	Auto. preference	Beta	0.6	0.1	0.95	0.93	0.97
ρ_{ρ_x}	Auto. job-separation	Beta	0.6	0.1	0.81	0.75	0.87
ρ_g	Auto. gov. purchase	Beta	0.6	0.1	0.95	0.93	0.97
$\rho_{\tilde{A}}$	Auto. stat. tech.	Beta	0.6	0.1	0.95	0.93	0.97
$100\sigma_{\eta_p}$	Std price markup	Inv. Gamma	0.15	1	1.50	1.05	1.92
$100\sigma_\mu$	Std nonstat. tech.	Inv. Gamma	1	1	0.37	0.29	0.44
$100\sigma_\upsilon$	Std MEI	Inv. Gamma	0.5	1	4.29	2.97	5.59
$100\sigma_R$	Std mon. policy	Inv. Gamma	0.15	1	0.28	0.25	0.31
$100\sigma_\zeta$	Std preference	Inv. Gamma	1	1	1.51	1.30	1.71
$100\sigma_{\rho_x}$	Std job-separation	Inv. Gamma	1	1	12.6	11.5	13.7
$100\sigma_g$	Std gov. purchase	Inv. Gamma	0.5	1	0.28	0.25	0.30
$100\sigma_{\tilde{A}}$	Std stat. tech.	Inv. Gamma	0.5	1	0.76	0.67	0.85
$100\sigma_w$	Std wage measurement	Inv. Gamma	0.5	0.5	0.25	0.22	0.27

and Vallés (2007), Furlanetto, Natvik, and Seneca (2013)). If my version of TANK were calibrated with a higher share of hand-to-mouth households to generate more volatile aggregate consumption as in the existing studies, estimation would deliver an even lower degree of nominal rigidity and more responsive monetary policy, causing more difficulty for TANK to capture the persistent aggregate data.

8. Conclusion

I have assessed the quantitative importance of (i) precautionary savings against unemployment risk and (ii) MPC heterogeneity for aggregate consumption dynamics. I did so by comparing the aggregate consumption volatility under the estimated HANK model to that under the RANK model. Most of the difference between consumption volatilities in HANK and RANK stems from MPC heterogeneity because of a large fraction of households that are close to the borrowing limit, which is required to target the empirical average MPC. I then assessed the model fit of the estimated HANK to the aggregate data and found that the estimated HANK underperforms the estimated RANK. The reason for the underperformance of the HANK model is its empirical average MPC target, which leads to counterfactually volatile aggregate consumption. Thus, the aggregate data favors a low degree of nominal rigidity and responsive monetary policy in HANK to reduce the model-implied consumption volatility, making it more challenging to match the estimated HANK to the persistence of the aggregate data.

In this paper, the asset structure is fairly simple in the sense that the model does not embed households' portfolio choice between liquid and illiquid assets. Bayer et al. (2019) investigate the effect of exogenous earnings risk on consumption and output in a two-asset HANK model and found that the interaction between precautionary savings and portfolio choices has a substantial effect on consumption and output. Future work could explore the impact of endogenous unemployment risk on aggregate fluctuations in a two-asset HANK model and see whether the results differ from those in the present paper. Therefore, my quantitative results should be regarded as benchmarks against which future models with a rich asset and labor market structure can be compared.

References

Acharya, Sushant and Keshav Dogra (2020), "Understanding HANK: Insights from a PRANK." *Econometrica*, 88 (3), 1113–1158. [739]

Auclert, Adrien (2019), "Monetary policy and the redistribution channel." *American Economic Review*, 109 (6), 2333–2367. [718]

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2018), "The intertemporal Keynesian cross." NBER Working Paper No. 25020. [746]

Auclert, Adrien, Matthew Rognlie, and Ludwig Straub (2020), "Micro jumps, macro humps: Monetary policy and business cycles in an estimated HANK model." NBER Working Paper No. 26647. [719, 735, 741, 742, 746]

Barro, Robert J. and Robert G. King (1984), "Time-separable preferences and intertemporal-substitution models of business cycles." *Quarterly Journal of Economics*, 99 (4), 817–839. [719]

Basu, Susanto and John G. Fernald (1997), "Returns to scale in U.S. production: Estimates and implications." *Quarterly Journal of Economics*, 105 (2), 249–283. [727]

Bayer, Christian, Benjamin Born, and Ralph Luetticke (2020), "Shocks, frictions, and inequality in US business cycles." Working Paper. [719]

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. [717, 748]

Bilbiie, Florin O. (2008), "Limited asset markets participation, monetary policy and (inverted) aggregate demand logic." *Journal of Economic Theory*, 140 (1), 162–196. [745]

Bilbiie, Florin O. (2019), "Monetary policy and heterogeneity: An analytical framework." Working Paper. [721, 739]

Bilbiie, Florin O. (2020), "The new Keynesian cross." *Journal of Monetary Economics*, 114, 90–108. [718, 721]

Boivin, Jean and Marc Giannoni (2006), "DSGE models in a data-rich environment." NBER Working Paper No. 12772. [731]

Broda, Christian and Jonathan A. Parker (2014), "The economic stimulus payments of 2008 and the aggregate demand for consumption." *Journal of Monetary Economics*, 68 (S), S20–S36. [731]

Calvo, Guillermo A. (1983), "Staggered prices in a utility-maximizing framework." *Journal of monetary Economics*, 12 (3), 383–398. [722]

Card, David, Raj Chetty, and Andrea Weber (2007), "Cash-on-hand and competing models of intertemporal behavior: New evidence from the labor market." *Quarterly Journal of Economics*, 122 (4), 1511–1560. [725]

Challe, Edouard (2020), "Uninsured unemployment risk and optimal monetary policy in a zero-liquidity economy." *American Economic Journal: Macroeconomics*, 12 (2), 241–283. [718, 721, 736, 738]

Challe, Edouard, Julien Matheron, Xavier Ragot, and Juan F. Rubio-Ramírez (2017), "Precautionary saving and aggregate demand." *Quantitative Economics*, 8 (2), 435–478. [718, 719, 720, 721, 725, 727, 733, 735, 736, 738]

Cho, Daeha (2023), "Supplement to 'Unemployment risk, MPC heterogeneity, and business cycles'." *Quantitative Economics Supplemental Material*, 14, https://doi.org/10.3982/QE1550. [719]

Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (2005), "Nominal rigidities and the dynamic effects of a shock to monetary policy." *Journal of Political Economy*, 113 (1), 1–45. [743]

Den Haan, Wouter, Pontus Rendahl, and Markus Riegler (2018), "Unemployment (fears) and deflationary spirals." *Journal of the European Economic Association*, 16 (5), 1281–1349. [718, 725, 728, 737]

Den Haan, Wouter J., Garey Ramey, and Joel Watson (2000), "Job destruction and propagation of shocks." *American Economic Review*, 90 (3), 482–498. [727]

Eusepi, Stefano and Bruce Preston (2015), "Consumption heterogeneity, employment dynamics and macroeconomic co-movement." *Journal of Monetary Economics*, 71 (C), 13–32. [719, 728, 746]

Furlanetto, Francesco, Gisle J. Natvik, and Martin Seneca (2013), "Investment shocks and macroeconomic co-movement." *Journal of Macroeconomics*, 37 (C), 208–216. [741, 742, 745, 747]

Galí, Jordi, David López-Salido, and Javier Vallés (2007), "Understanding the effects of government spending on consumption." *Journal of the European Economic Association*, 5 (1), 227–270. [718, 721, 745, 746, 747]

Ganong, Peter and Pascal Noel (2019), "Consumer spending during unemployment: Positive and normative implications." *American Economic Review*, 109 (7), 2383–2424. [728]

Guvenen, Fatih, Serdar Ozkan, and Jae Song (2014), "The nature of countercyclical income risk." *Journal of Political Economy*, 122 (3), 621–660. [739]

Hagedorn, Marcus and Iourii Manovskii (2008), "The cyclical behavior of equilibrium unemployment and vacancies revisited." *American Economic Review*, 98 (4), 1692–1706. [727, 728]

Hagedorn, Marcus, Iourii Manovskii, and Kurt Mitman (2019), "The fiscal multiplier." Working Paper, University of Oslo. [720, 735, 746]

Hamilton, James (2018), "Why you should never use the Hodrick–Prescott filter." *Review of Economics and Statistics*, 100 (5), 831–843. [731, 734]

Heathcote, Jonathan and Fabrizio Perri (2018), "Wealth and volatility." *Review of Economic Studies*, 85 (4), 2173–2213. [718]

Justiniano, Alejandro, Giorgio Primiceri, and Andrea Tambalotti (2010), "Investment shocks and business cycles." *Journal of Monetary Economics*, 57 (2), 132–145. [718, 726, 740, 741, 742]

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2013), "Is there a trade-off between inflation and output stabilization?" *American Economic Journal: Macroeconomics*, 5 (2), 1–31. [731, 742]

Kaplan, Greg, Benjamin Moll, and Giovanni L. Violante (2018), "Monetary policy according to HANK." *American Economic Review*, 108 (3), 697–743. [718, 725, 729, 731, 735, 738]

Kaplan, Greg and Giovanni L. Violante (2014), "A model of the consumption response to fiscal stimulus payments." *Econometrica*, 82 (4), 1199–1239. [720]

Kekre, Rohan (2019), "Unemployment insurance in macroeconomic stabilization." Working Paper. [718, 731]

Kreamer, Jonathan (2016), "Household debt, unemployment, and slow recoveries." Working Paper. [718]

Krusell, Per and Anthony A. Smith (1998), "Income and wealth heterogeneity in the macroeconomy." *Journal of Political Economy*, 106 (5), 867–896. [717, 720]

Lalive, Rafael (2007), "Unemployment benefits, unemployment duration, and post-unemployment jobs: A regression discontinuity approach." *American Economic Review*, 97 (2), 108–112. [725]

Le Barbanchon, Thomas (2016), "The effect of the potential duration of unemployment benefits on unemployment exits to work and match quality in France." *Labour Economics*, 42 (C), 16–29. [725]

McKay, Alisdair, Emi Nakamura, and Jón Steinsson (2016), "The power of forward guidance revisited." *American Economic Review*, 106 (10), 3133–3158. [727]

McKay, Alisdair and Ricardo Reis (2016), "The role of automatic stabilizers in the U.S. business cycle." *Econometrica*, 84 (1), 141–194. [720]

McKay, Alisdair and Ricardo Reis (2019), "Optimal automatic stabilizers." Working Paper. [718]

Oh, Hyunseung and Ricardo Reis (2012), "Targeted transfers and the fiscal response to the great recession." *Journal of Monetary Economics*, 59 (S), S50–S64. [718]

Oh, Joonseok and Anna Rogantini Picco (2019), "Macro uncertainty and unemployment risk." Working Paper, EUI. [718]

Petrongolo, Barbara and Christopher A. Pissarides (2001), "Looking into the black box: A survey of the matching function." *Journal of Economic Literature*, 39 (2), 390–431. [727]

Ravn, Morten O. and Vincent Sterk (2017), "Job uncertainty and deep recessions." *Journal of Monetary Economics*, 90 (C), 125–141. [718, 721, 736, 738]

Ravn, Morten O. and Vincent Sterk (2018), "Macroeconomic fluctuations with HANK & SAM: An analytical approach." Working Paper. [718]

Sala, Luca, Ulf Soderstrom, and Antonella Trigari (2010), "The output gap, the labor wedge and the dynamic behavior of hours." CEPR Discussion Paper No. 8005. [742]

Shimer, Robert (2005), "The cyclical behavior of equilibrium unemployment and vacancies." *American Economic Review*, 95 (1), 25–49. [727, 735]

van Ours, Jan C. and Milan Vodopivec (2008), "Does reducing unemployment insurance generosity reduce job match quality?" *Journal of Public Economics*, 92 (3–4), 684–695. [725]

Werning, Iván (2015), "Incomplete markets and aggregate demand." NBER Working Paper No. 21448. [721, 739]

Winberry, Thomas (2018), "A method for solving and estimating heterogeneous agent macro models." *Quantitative Economics*, 9 (3), 1123–1151. [718, 719, 726]

Co-editor Kjetil Storesletten handled this manuscript.

Manuscript received 10 February, 2020; final version accepted 2 March, 2023; available online 3 March, 2023.