HANK's Response to Aggregate Uncertainty in an Estimated Business Cycle Model*

Cosmin Ilut†, Ralph Luetticke‡and Martin Schneider§

February 10, 2025

Abstract

This paper studies a HANK model with agents who respond to both idiosyncratic and aggregate uncertainty. Since aggregate uncertainty is modeled as ambiguity, it affects the steady state and linearized dynamics, allowing for fast computation and estimation. The interaction of aggregate uncertainty shocks and portfolio frictions generates a high capital premium as well as most cyclical comovement in macroeconomic aggregates. Heterogeneity in portfolios is crucial: when it is shut down, the model fails to explain investment dynamics and the capital premium disappears. Cautious price and wage setting by firms in anticipation of aggregate uncertainty shapes employment and inflation dynamics.

JEL codes: C11, D31, E32, E63

Keywords: Aggregate Uncertainty, Bayesian Estimation, Business Cycles, Incomplete Markets, Monetary and Fiscal Policy

*We thank Andy Atkeson, Markus Brunnermeier, Axelle Ferriere, Mark Gertler, Martin Holm, Karin Kinnerud, Kieran Larkin, Moritz Lenel, Tom Sargent, Karthik Sastry, Frank Schorfheide, Ludwig Straub, Gianluca Violante, Tom Winberry and Christian Wolf as well as seminar and conference participants at the Bundesbank, the ECB, Lancaster ETM Workshop, NBER Summer Institute, NYU, the Oslo Macro Conference, Paris School of Economics, Pavia Macro Forum, Princeton, SITE, and Wharton for helpful comments. Ralph Luetticke gratefully acknowledges funding by the European Union (ERC, AIRMAC, 101114991). Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or the European Research Council. Neither the European Union nor the granting authority can be held responsible for them.

†Duke University and NBER, *email:* cosmin.ilut@duke.edu

‡Universität Tübingen and CEPR, *email:* ralph.luetticke@uni-tuebingen.de

§Stanford University, NBER and CEPR, *email:* schneidr@stanford.edu

1 Introduction

A rapidly growing literature studies models with nominal rigidities and rich household heterogeneity.[1] At the heart of such HANK models are frictions in financial markets, in particular incomplete insurance against idiosyncratic income shocks, and differences in liquidity across assets. As a result, households' *responses to uncertainty about the future* matter for savings and portfolio choice. For example, higher uncertainty encourages precautionary savings and makes illiquid assets less attractive. In equilibrium, an increase in uncertainty lowers the risk-free rate and raises the premium on illiquid assets.

However, most work on HANK models to date restrict attention to responses to *idiosyncratic* uncertainty. This is for technical reasons: with expected utility preferences, uncertainty has only second-order effects on utility and choice that are not captured by popular linear solution methods. As a result, most HANK models do not quantify precautionary savings or asset premia due to *aggregate* uncertainty. They also abstract from the effect of aggregate uncertainty on firm decisions. These technical features create a disconnect between the HANK literature and large bodies of work in macroeconomics and finance that emphasize time-varying aggregate uncertainty.

This paper develops and estimates a two-asset HANK model with agents who respond to both idiosyncratic and aggregate uncertainty. We show that such a model is very tractable when aggregate uncertainty is modeled as ambiguity using multiple priors preferences. Aggregate uncertainty then has first-order effects on utility and is reflected in the equations for the steady state and linear dynamics, so we can use standard methods to characterize and estimate our model. As would be the case for aggregate risk, the anticipation of aggregate ambiguity introduces wedges in households' and firms' intertemporal first order conditions that jointly fluctuate with uncertainty shocks. All intertemporal decisions – savings and portfolio choice by households as well as price and wage setting by firms – are therefore taken more cautiously when uncertainty is higher.

Our main quantitative result is that aggregate uncertainty shocks not only allow the model to generate large premia in asset markets, but also powerful comovement of macroeconomic aggregates. In our baseline estimation, the average excess return on capital of 5.5% consists of a premium for aggregate uncertainty of 3.2% and a smaller illiquidity premium of only 2.3%. At the same time, a single shock to ambiguity about total factor productivity (TFP) jointly explains more than two thirds of cyclical variation in key macroeconomic aggregates. In contrast, our estimation infers only a modest role for shocks to idiosyncratic volatility of labor income. Identification comes from the dynamics of investment: only an aggregate uncertainty shock generates a recession with a strong protracted investment slump.

Our results are driven by the *interaction* of aggregate uncertainty and the portfolio

[1] See the review by Auclert et al. (2024a).

frictions of the two-asset HANK model. Higher uncertainty about TFP generates a flight to safety: rich households who own most capital try to substitute away from capital towards bonds, while poor households want bonds for precautionary savings. The resulting decline in investment is much stronger than in a representative agent model. Indeed, when we shut down heterogeneity in our model, the effect of aggregate uncertainty on investment essentially disappears. Intuitively, when all households are identical and hence equally exposed to capital, the precautionary motive tends to stabilize capital demand after uncertainty shocks. The capital premium also disappears when heterogeneity is shut down: uncertainty without frictions does not call for large premia. Higher labor income volatility similarly generates relatively more precautionary savings and less of a flight away from capital; this makes labor income volatility less suitable as driver of recessions.

A second key mechanism in our model is cautious price and wage setting by firms and unions, respectively. Higher uncertainty about TFP means that firm owners worry about future cost and unions worry about the future marginal product of labor. This worry is a force that pushes up both prices and wages after an aggregate uncertainty shock. It thus works against deflationary pressure from lower demand for goods due to precautionary savings. As a result, recessions triggered by aggregate uncertainty shocks exhibit less deflation than those triggered by a typical New Keynesian demand shifter. Since uncertainty wedges appear in the linear approximation of the model, we can selectively shut off one or more of them to assess their contributions. When we do so for the price and wage Phillips curves, the model produces a shallow recession with substantial deflation. Interaction of uncertainty with nominal rigidities is thus another important channel.

Our model builds on the two-asset HANK setup in Bayer et al. (2024a). Households experience uninsurable shocks to labor productivity. They save in liquid, safe nominal bonds and in illiquid capital with uncertain payoffs. Firms' price and wage setting and households' trading of capital are subject to Calvo frictions. The price of capital moves due to capital adjustment costs. Government policy determines the net supply of nominal bonds and sets a rule for the nominal short rate. There are two aggregate real shocks, to TFP and to the volatility of labor productivity. We also allow for shocks to monetary policy and to the inflation target.

Households in our model are averse to both risk (uncertainty with known probabilities) and ambiguity (uncertainty with "unknown odds"). The Ellsberg (1961) paradox established a behaviorally meaningful distinction between the two. It motivated multiple priors preferences (Gilboa and Schmeidler, 1989) that capture ambiguity via sets of beliefs: agents evaluate plans *as if* they hold a worst-case belief that minimizes expected utility. In many macro and finance applications, responses to both types of uncertainty are qualitatively sim-

ilar.[2] In particular, ambiguity averse agents may take precautionary actions when the future is more uncertain. They also dislike assets with uncertain payoffs and will hold them only if compensated by an uncertainty premium. For the purposes of this paper, the advantage of an ambiguity approach is that equilibria are easier to compute.[3]

Our solution strategy relies on multiple priors preferences with ambiguity *in means*. We parameterize belief sets by the mean TFP innovation: it lives in a symmetric interval around zero, the mean innovation under the true data generating process (DGP). The width of the interval, a measure of ambiguity, is an exogenous stochastic process. It has a positive mean, so households act *as if* long run mean TFP is lower than actual mean TFP. An uncertainty shock widens the interval and implies a lower worst case. Since households like to smooth consumption, they save more – precautionary savings here requires only diminishing marginal utility, not prudence. Moreover, households invest *as if* the expected payoff on capital is low, driving down prices. Both effects are about means, and hence have first order effects.

We discipline the exogenous ambiguity process by a priori bounding its mean to be no larger than one standard deviation of the TFP innovation under the true DGP. In the long run, the worst case mean is therefore no worse than a bad scenario that occurs relatively often along any sample path. The bound serves as a consistency criterion that connects the true DGP measured by the econometrician to the size of agents' belief sets. It weakens the strict criterion of rational expectations – belief sets contain only one belief, namely the true DGP – but shares the idea that sets should be close to the true DGP, and more so when observed volatility is low. As discussed in more detail in Ilut and Schneider (2014), such a consistency criterion is sensible when we want ambiguity aversion to capture cautious behavior in a world where agents observe repeated regular patterns, such as business cycles.

Our computational approach leverages the fact that the model is observationally equivalent to one with pessimistic expected utility agents. We make the standard assumption that agents understand the law of motion of the economy, so that any reasoning about endogenous variables follows from that law of motion together with pessimism about exogenous TFP. In particular, households always behave *as if* endogenous variables are on a transition path towards a worst case steady state with low mean TFP. In the ergodic steady state of the model, in contrast, TFP is constant at its higher true mean. Endogenous variables nevertheless reflect cautious behavior, since households anticipate the bad transition path. It follows that steady state and dynamics must always be studied jointly, since dynamics are crucial to describe worst case beliefs in steady state.

[2]See Ilut and Schneider (2022) for a survey of ambiguity models and their applications.

[3]We also note that recursive multiple priors utility shares with standard expected utility preferences the property that choice is indifferent to the timing of the resolution of uncertainty. Quantitative models with risk aversion instead often work with Epstein-Zin utility in order to generate large risk premia. Common parametrizations imply a large willingness to pay for early resolution, as clarified by Epstein et al. (2014).

Our estimation strategy is thus designed to identify parameters jointly from long run moments and business cycle dynamics. Standard HANK models allow a sequential two step approach. The first step calibrates a steady state without aggregate uncertainty to match long run moments of household portfolios and the wealth distribution. The second then uses linearized dynamics around that steady state to build a likelihood for estimation. In our model, long run moments reflect cautious behavior not only due to idiosyncratic risk, as in the standard case, but also to aggregate uncertainty, as captured by the anticipated transition to the worst case. We thus propose a new procedure that iterates between the steps: in every iteration we first choose a subset of parameters to match long run moments and then estimate the remaining parameters from the linearized dynamics.

Our estimation exercise uses six observables: the growth rates of consumption, investment and hours, the inflation rate, the short nominal interest rate and the return on capital constructed by Gomme et al. (2011). We work with five shock series and allow for measurement error on the observables. The aggregate uncertainty shock drives the bulk of business cycle variation in both quantities and real asset prices. Nominal shocks to monetary policy and the inflation target are important for inflation as well as the nominal interest rate, but less relevant for other variables. TFP shocks contribute about 20% of variation in inflation and the nominal rate, and much less to movement in other variables.

Households' response to an aggregate uncertainty shock works along both the savings and the portfolio choice margin. When households worry about bad times ahead, they save more for precautionary reasons. Moreover, they substitute away from assets with uncertain payoff towards safe bonds. In a two-asset HANK model, the relative strength of these forces differs across the wealth distribution. Since uncertainty about TFP affects both labor and capital income everyone likes to save more. Portfolio substitution, in contrast, is relevant only for rich households who actually invest in capital, and not for the poor who hold only safe bonds. The former households are also affected negatively by a decline in the price of capital, a feedback effect that amplifies an aggregate uncertainty shock.

Portfolio substitution by rich capital owners implies that investment as well as the capital premium are more responsive to an aggregate uncertainty shock in a HANK model than in a representative agent (RANK) model. When we consider a RANK version of our model with otherwise identical parameters, we find that aggregate uncertainty has essentially no effect on investment and the capital premium. For a representative agent who receives all labor and capital income, precautionary savings and portfolio choice effects effectively cancel each other out. With liquidity frictions, in contrast, rich households, substitute away from capital whereas precautionary savings of the poor flow to bonds. The price of capital falls to open up a premium on capital to induce rich households to invest.

The relative magnitude of shocks we infer is identified by the patterns of comovement they generate. The special feature of the aggregate uncertainty shock is that it not only works as a typical (negative) demand shock in a New Keynesian model that jointly lowers quantities and the interest rate, but also affects the relative attractiveness of capital and bonds. In other words, it also activates a strong "investment wedge", due to portfolio substitution. A shock to idiosyncratic risk, in contrast, mostly encourages precautionary savings and thus has weak effects on investment. It plays only a negligible role for business cycle dynamics in our estimation. Finally, nominal and TFP shocks cannot generate comovement of aggregate quantities and prices for familiar reasons. As a result, their volatility and impact on the business cycle is estimated to be relatively small.

There is a large literature on uncertainty shocks in business cycle models (see Fernandez-Villaverde and Guerron-Quintana (2020) for an overview). In particular, several papers have quantified the role of aggregate uncertainty as a demand shock in New Keynesian models with a representative agent (e.g. Ilut and Schneider (2014), Leduc and Liu (2016), Basu and Bundick (2017), Bhandari et al. (2024)). Ilut and Schneider (2014) is closest to our paper since they also model aggregate uncertainty using multiple priors utility. They emphasize that an aggregate uncertainty shock generates countercyclical labor and discount factor wedges and can thereby generate comovement of hours, consumption and output. However, their model struggles to fit the dynamics of investment and they do not match the capital return. Our HANK model with aggregate uncertainty and liquidity frictions fits investment and the capital return well since it introduces a quantitatively important investment wedge.

Relative to the large and rapidly growing literature on HANK models, our paper makes three contributions. First, it clarifies that frictions matter for aggregate dynamics. This result is not obvious: Berger et al. (2023) have shown that micro frictions matter little for aggregate fluctuations in a large class of HANK models. Their conclusion comes from a novel sufficient statistic approach: they measure labor and discount factor wedges in household first order conditions for labor and savings, respectively, relative to a RANK benchmark, building on insights of Nakajima (2005), Werning (2015) and Debortoli and Galí (2017). Our result falls outside their setting since it is due to an investment wedge in households' first order condition for the portfolio choice of capital, not their choice of savings. Second, we jointly estimate the role of aggregate and idiosyncratic uncertainty shocks. We obtain responses to idiosyncratic risk driven by precautionary savings, as in Bayer et al. (2019) and Bilbiie et al. (2023). At the same time, our estimation assigns a relatively small role to such shocks as a driver of business cycles, consistent with the VAR analysis in Chang et al. (2021) and the structural estimation in Bayer et al. (2024a).

A third contribution to the HANK literature is methodological: modeling aggregate

uncertainty as ambiguity means first order approximations work well. Computing the law of motion of our economy involves the same steps as under certainty equivalence - it just uses pessimistic expectations – and is therefore just as fast. It is also straightforward to allow for wedges in firm Euler equations. When uncertainty is described by risk, in contrast, solution techniques need to take into account nonlinearities. One approach is to use global methods that have been popular in the finance literature (see Schaab (2020), Fernández-Villaverde et al. (2024) or Kase et al. (2022) for recent applications to HANK models). An alternative is to work out perturbations to higher order (e.g. Bhandari et al. (2023), Bayer et al. (2024b), Auclert et al. (2024b)). In any case, estimation requires a method that is sufficiently fast so the model can be evaluated many times. Our characterization of linear dynamics can be an input to estimation techniques already popular in the HANK literature. We derive a state space representation of our model and use the Bayesian techniques of Bayer et al. (2024a).

In our model, the allocation of aggregate uncertainty among heterogeneous investors matters for asset prices and real activity. This broad theme, which we study in an estimated HANK model, permeates several other active literatures. First, in many models of financial markets, bad shocks lower an asset price because they affect precisely those agents who like the asset most, including by lowering their wealth.[4] Second, in the wake of the financial crisis, a new class of models studied the role of financial intermediaries as "specialist" investors (for example, Gertler (2010), Jermann and Quadrini (2012), Brunnermeier and Sannikov (2014), Brunnermeier and Sannikov (2014), Bocola and Lorenzoni (2023)). More recently, there has been growing interest in linking valuation and inequality (for example, Kacperczyk et al. (2019), Gomez (2024), Fernández-Villaverde et al. (2023) or Ilut et al. (2022)). Our contribution is to study these mechanisms using the frictions of the two-asset HANK model.

The rest of the paper is structured as follows. Section 2 introduces the model and its solution method. Section 3 describes how we estimate the model. Sections 4 and 5 present quantitative results for our baseline HANK model and compare them to a RANK model.

2 Model

Our setup shares technology and asset market frictions as well as the treatment of household heterogeneity with the two-asset HANK model in Bayer et al. (2024a). The difference is that we replace expected utility preferences by recursive multiple priors utility (common for all agents) and add aggregate uncertainty shocks. Our description of the physical environment

[4]Panageas et al. (2020) provides an overview of this mechanism, and how it emerges in models of heterogeneous beliefs, attitude towards uncertainty or access to financial markets. Kekre and Lenel (2022) show that it helps understand the transmission of monetary policy in a model with heterogeneous agents and nominal rigidities.

in Section 2.1 and of markets and institutions in Section 2.2 thus spends less time on the familiar ingredients. To ease notation, we do not carry around in those sections explicit notation for the state space. It is understood that all endogenous variables – choices and prices – are functions of the exogenous shocks. We make explicit the state once we define recursive equilibrium in Section 2.3 below.

2.1 Physical environment

Technology. There are four types of goods: final, intermediate, capital and labor. Final output Y_t is a CES aggregate of a continuum of intermediates, each of which is made from capital and labor. Production functions are

$$Y_t = \left(\int Y_{jt}^{\frac{\eta-1}{\eta}} dj \right)^{\frac{\eta}{\eta-1}}; \qquad Y_{jt} = Z_t N_{jt}^{\alpha} (u_{jt} K_{jt})^{(1-\alpha)}, \tag{1}$$

where η is the elasticity of substitution across intermediates, Z_t, is an exogenous TFP shock common to all intermediates, and u_{jt} is the utilization rate of capital K_{jt} used in the production of intermediate j.

Capital is made from final goods subject to investment rate adjustment costs: if investment was I_{t-1} at date $t-1$, then investing I_t units of the final good at date t makes

$$I_t \left[1 - \frac{\phi}{2} \left(\log \frac{I_t}{I_{t-1}} \right)^2 \right] \tag{2}$$

units of capital, available for use in production at date $t+1$. Capital used to produce intermediate j depreciates at the rate $\delta(u_{jt})$, with δ an increasing convex function.

Household-supplied *labor* is transformed into a continuum of labor varieties that are in turn bundled into aggregate *labor services* used in production according to

$$N_t = \left(\int \hat{n}_{jt}^{\frac{\zeta-1}{\zeta}} dj \right)^{\frac{\zeta}{\zeta-1}}, \tag{3}$$

where ζ is the elasticity of substitution across varieties. Introducing intermediate goods as well as a distinction between labor, labor services and labor varieties helps tractably introduce sticky prices and wages, as discussed below.

The TFP shock Z_t is a key source of aggregate uncertainty in our model. The true data generating process (DGP) is a standard persistent AR(1)

$$\log Z_{t+1} = (1-\rho_z) \log \bar{Z} + \rho_z \log Z_t + \epsilon_{t+1}^Z, \tag{4}$$

where ϵ_{t+1}^Z is an iid normal sequence of innovations with mean zero and variance σ_z^2, and $\bar{Z}$ is the long run mean.

Preferences. Households care about consumption goods and labor. At date t, household i obtains utility from a composite good

$$x_{it} = c_{it} - h_{it}\frac{n_{it}^{1+\sigma_f}}{1+\sigma_f},$$

where c_{it} is consumption, n_{it} is labor, h_{it} is household i's idiosyncratic labor productivity shock and σ_f is the Frisch elasticity of labor supply. The quasilinear functional form eliminates wealth effects of labor supply, following Greenwood et al. (1988).[5]

Household i's information set includes all aggregate exogenous variables of the model as well as the idiosyncratic shock h_{it}. We introduce ambiguity via *sets* of one-step-ahead conditional distributions. Let $\mathcal{P}_t$ denote the set of probabilities relevant at date t for computing conditional moments of random variables at $t+1$. The set $\mathcal{P}_t$ is itself a random variable. A larger set after some history describes an agent who is less confident about assigning probabilities to events at date $t+1$, perhaps because she has only poor information. The evolution of ambiguity is thus described by an entire stochastic process of belief sets. In particular, uncertainty shocks correspond to expansions of the set.

Fix a specific consumption plan C, that is, a collection of stochastic processes (c_{it}, n_{it}) and hence also x_{it}. We would like to describe continuation utility after any history under ambiguity. We write E_t^p for the conditional expectation taken under the one-step-ahead probability $p \in \mathcal{P}_t$. The utility process for the consumption plan C is then defined as the solution to the stochastic difference equation

$$U_t^C = \frac{x_{it}^{1-\gamma}}{1-\gamma} + \beta \min_{p\in\mathcal{P}_t} E_t^p \left[U_{t+1}^C\right], \tag{5}$$

where the discount factor β is between zero and one. Utility at date t is the sum of felicity from the current composite $u(x_{it})$ and discounted expected continuation utility, where the expectation is taken under the worst case conditional distribution *for that plan* C. For an agent who perceives more ambiguity, the worst case for each plan is more pessimistic – this is how the model captures cautious evaluation of plans.[6]

The multiple priors functional form (5) captures a strict preference for knowing probabilities, and is thus consistent with behavior exhibited in the Ellsberg (1961) experiments. The

[5]The GHH preferences assumption is motivated by estimated DSGE models typically finding small aggregate wealth effects (see e.g. Schmitt-Grohé and Uribe (2012); Born and Pfeifer (2014)). Using GHH is not important for introducing ambiguity - for that, we can alternatively use King et al. (1988) preferences.

[6]See Ilut and Schneider (2022) for a recent review of the multiple priors model and its applications.

key feature is that the worst case belief endogenously varies with the consumption plan C. In the special case where every $\mathcal{P}_t$ contains only a single conditional probability, the difference equation can be solved forward and the solution is standard time-separable expected utility. The recursive definition ensures that preferences share the dynamic consistency property of expected utility even under ambiguity (see Epstein and Schneider (2003) for axiomatic foundations). The primitives of the utility representation are the function u, the discount factor β, and the process $\mathcal{P}_t(s^t)$. We assume that there is no ex-ante preference heterogeneity, and thus let these utility primitives be common across households.

Ambiguity about mean TFPs and aggregate uncertainty shocks. Our modeling approach relies on a particular tractable form for the process of one-step-ahead belief sets $\mathcal{P}_t$: we parametrize each set by an interval of means for the innovation to TFP

$$\log Z_{t+1} = (1-\rho_z)\log \bar{Z} + \rho_z \log Z_t + \mu_t + \epsilon^Z_{t+1}; \quad \mu_t \in [-a_t, a_t], \tag{6}$$

where ϵ^Z is normal with mean zero and variance σ_z^2. When a_t is higher, there is more ambiguity and the set of belief is larger – in particular, a wider interval implies a lower worst case mean. We also note that the one-step-ahead conditional of the true DGP (4) is always contained in the set.

Ambiguity a_t is an exogenous stochastic process, common to all households. It captures the cumulative effect of news that affect confidence about future TFP and evolves as

$$\log a_t = (1-\rho_a)\log \bar{a} + \rho_a \log a_{t-1} + \epsilon^a_t, \tag{7}$$

with long run mean $\bar{a} > 0$, persistence $0 < \rho_a < 1$, and $\epsilon^a_t \sim i.i.d\ N(0, \sigma_a)$. Intuitively, agents sometimes obtain news that makes them relatively more confident that the correct forecast of future TFP $\log Z_{t+1}$ is $(1-\rho_z)\log \bar{Z} + \rho_z \log Z_t$. In other periods, they receive contradictory information, and end up less confident about their forecast. Periods of low $a_t < \bar{a}$ represent unusually low uncertainty about future productivity, whereas $a_t > \bar{a}$ describes periods of high uncertainty.

Endowments and shocks to idiosyncratic risk. An individual household switches between two states: worker or an entrepreneur. A worker becomes an entrepreneur with probability ζ. Entrepreneurs do not work but instead receive a stream of pure rents from firms, described further below.[7] With probability ι, an entrepreneur becomes a worker with

[7] In our quantitative exercise below, the entrepreneur state serves as a "superstar" state that helps generate a positively-skewed income and wealth distribution.

idiosyncratic labor productivity $h_{it} = 1$. Worker productivity evolves as

$$\log h_{it} + \log \bar{h}_t = \rho_h(\log h_{it-1} + \log \bar{h}_{t-1}) + \epsilon^h_{it}, \tag{8}$$

Here the innovations ϵ^h_{it} are independent both over time and across individual workers, and are normally distributed with mean zero and variance $\bar{\sigma}^2_{h,t}$. Moreover, $\bar{h}_t$ adjusts every period so the cross sectional average productivity over all workers remains constant at one.

The variance $\bar{\sigma}^2_{h,t}$ is itself stochastic: it captures changes in the idiosyncratic income risk faced by households. Income risk risk follows a log-AR(1) process

$$\log \bar{\sigma}^2_{h,t+1} = (1 - \rho_h) \log \bar{\sigma}^2_h + \rho_h \log \bar{\sigma}^2_{h,t} + \epsilon^\sigma_t, \tag{9}$$

where $\bar{\sigma}^2_h$ is steady state income risk, and $\epsilon^\sigma_t \sim i.i.d\ N(0, \sigma_\sigma)$. Thus, at a given time t households know that there is an aggregate change in the variance of shocks that drive their idiosyncratic productivity next period. This type of variation in income risk allows us to study within the same model changes to aggregate uncertainty, through a_t, and to idiosyncratic uncertainty, through $\bar{\sigma}^2_{h,t}$.

2.2 Markets and institutions

Firms and unions. Final goods producers are competitive firms that operate the final goods technology in (1), sell the final good at the nominal price P_t and buy individual varieties at nominal prices P_{jt}. Those firms solve static problems. Intermediate goods producers operate the intermediates technology in (1): they rent capital at the real rental rate r^k_t and labor services at the nominal wage W^F_t. They solve intertemporal problems, since they are monopolistic competitors subject to Calvo (1983) pricing frictions. In particular, nominal prices P_{jt} are usually indexed to the steady-state level of inflation and can be discretionally adjusted only with probability $1-\lambda_Y$. Let Π^F_t denote profits made by intermediate producers – we assume that they are earned by entrepreneurs.

Labor services are bundled by competitive firms called "labor packers", according to (3). Labor packers solve static problems: they sell labor services at the nominal wage W^F_t and buy varieties at wages W_{jt}. Calvo frictions enter again in the pricing of varieties, sold by monopolistically competitive "unions" that buy labor from households at the real wage w_t and transform it one-for-one. Here the wage is per unit of effective labor – a household i who works n_{it} hours supplies $n_{it}h_{it}$ units of labor. Nominal wages of varieties W_{jt} are usually indexed to the steady-state level of wage growth and can be discretionally adjusted with probability $1 - \lambda_w$. Let Π^U_t denote profits made by unions - they are proportionately rebated to workers.

Finally, capital goods producers are competitive firms that use final goods to make capital according to (2) and sell the capital at the real price q_t. Their problem is intertemporal because of time-to-build and adjustment costs. Let M_t^S denote the common sequence of stochastic discount factors (SDF) used to discount profits of all firms that solve intertemporal problems, that is, intermediate goods firms, unions and capital producers. For now we fix the SDF exogenously so the firm problem is well-defined. Below we derive it from the marginal utility of ambiguity averse firm owners and union members.

Assets and household budget constraint. Households have access to two assets: nominal bonds and capital. A household i who invests $b_{it} \geq 0$ units of the final good in bonds at date $t-1$ earns the nominal interest rate $R_{it} = R_t$ between $t-1$ and t. A household i may also borrow up to a limit, $b_{it} \in [\bar{B}, 0)$ and then pays the higher borrowing rate $R_{it} = R_t + \bar{R}$. Capital $k_{it} \geq 0$ cannot be sold short. It is also illiquid: only a fraction λ of households may adjust their capital holdings in a given period. Households that do not participate in the capital market must choose $k_{it+1} = k_{it}$, but still obtain dividends and can adjust their bond holdings.

The household budget constraint equates expenditure on goods and assets to income:

$$c_{it} + b_{it+1} + q_t k_{it+1} = b_{it}\frac{R_{it}}{\pi_t} + (q_t + r_t^k)k_{it} + (1-\tau)(w_t h_{it} n_{it} + \mathbb{I}_{h_{it}=0}\Pi_t^F + \mathbb{I}_{h_{it}\neq 0}\Pi_t^U) + L_t, \quad (10)$$

where $\pi_t = P_t/P_{t-1}$ is the inflation rate, so R_t/π_t is the real return on bonds. Income on the right hand side consists of non-transfer income, taxed at the rate τ, and lump sum transfers L_t. If household i is an entrepreneur ($h_{it} = 0$), non-transfer income consists of profits from firms Π_t^F. For a worker household, it consists instead of wages plus profits from unions Π_t^U.

Government. The government follows rules for monetary and fiscal policy. It sets the nominal interest rate on bonds according to a Taylor-type (1993) rule:

$$\frac{R_{t+1}}{\bar{R}} = \left(\frac{R_t}{\bar{R}}\right)^{\rho_R} \left(\frac{\pi_t}{\pi_t^*}\right)^{(1-\rho_R)\theta_\pi} \left(\frac{Y_t}{Y_{t-1}}\right)^{(1-\rho_R)\theta_Y} \varepsilon_t^R. \quad (11)$$

The coefficient $\bar{R} \geq 0$ determines the nominal interest rate in the steady state. The coefficients $\theta_\pi, \theta_Y \geq 0$ govern the extent to which the central bank attempts to stabilize inflation and output growth, $\frac{Y_t}{Y_{t-1}}$. The parameter $\rho_R \geq 0$ captures interest rate smoothing.

Two standard sources of shocks enter the Taylor rule: a monetary policy shock and an inflation target shock. They follow AR(1) processes

$$\log \varepsilon_t^R = \rho_R^\epsilon \log \varepsilon_{t-1}^R + \epsilon_t^R, \quad (12)$$

$$\log \pi_t^* = (1-\rho_\pi)\log \bar{\pi} + \rho_\pi \log \pi_{t-1}^* + \epsilon_t^\pi, \quad (13)$$

with innovations $\epsilon_t^R \sim i.i.d.\ N(0, \sigma_R)$ and $\epsilon_t^\pi \sim i.i.d.\ N(0, \sigma_\pi)$, respectively. While the monetary policy shock captures short run surprises, the inflation target shock accounts for the low frequency behavior of inflation possibly generated by the slow moving beliefs and resulting conduct of the monetary authority (see, for example Justiniano et al. (2013)).

The government further implements rules for government spending G_t as well as per capita lump sum transfers L_t:

$$\frac{G_t}{\bar{G}} = \left(\frac{G_{t-1}}{\bar{G}}\right)^{\rho_G} \left(\frac{B_t}{\bar{B}}\right)^{-(1-\rho_G)\gamma_B^G} \left(\frac{Y_t}{\bar{Y}}\right)^{(1-\rho_G)\gamma_Y^G}, \tag{14}$$

$$\frac{L_t}{\bar{L}} = \left(\frac{L_{t-1}}{\bar{L}}\right)^{\rho_L} \left(\frac{B_t}{\bar{B}}\right)^{-(1-\rho_L)\gamma_B^L} \left(\frac{Y_t}{\bar{Y}}\right)^{(1-\rho_L)\gamma_Y^L}. \tag{15}$$

Here policy parameters γ_Y^G and γ_Y^L control the cyclicality of the two fiscal instruments, γ_B^G and γ_B^L their adjustment to government debt to ensure debt stability, and ρ_G, ρ_L their mean reversion to the steady state values $\bar{G}$ and $\bar{L}$, respectively.

The government issues short term bonds that are perfect substitutes for bonds issued by households. The government budget constraint equates new debt B_{t+1} issued at date t to current spending and transfers less taxes, as well as principal plus interest on debt issued the previous period

$$B_{t+1} = G_t + L_t - \tau\left(w_t n_{it} h_{it} + \mathbb{I}_{h_{it}\neq 0}\Pi_t^U + \mathbb{I}_{h_{it}=0}\Pi_t^F\right) + R_t B_t/\pi_t. \tag{16}$$

Given rules for spending and transfers, as well as a tax rate, government debt adjusts endogenously to ensure that the budget constraint holds.

Market clearing. In intermediate goods markets, producer j adjusts the price P_{jt}, taking as given final goods producers' demand function for variety j. Market clearing for final goods requires that consumption, investment and government spending add up to total output. Since the market for final goods and labor services are competitive, the aggregate price level follow as the ideal price index for the final goods production function in (1).

Similarly, in markets for labor varieties, union j adjusts the wage W_{jt}, taking as given labor packers' demand for variety j. The two competitive labor markets clear at two economy-wide wages. The market for labor clears at the real wage w_t paid by unions to households. The market for labor services clears at the real wage $w_t^F = W_t^F/P_t$ paid by intermediate goods producers to labor packers. The two wages are linked since w_t is unions' marginal cost and is lower than w_t^F because of union market power.

All capital is owned by households. The rental market for capital clears at the rate r_t^k where demand from intermediate goods producers equals supply from households. The

market for owned capital clears at the price q_t where demand from households equals undepreciated capital that household bring into the period plus new supply provided by capital goods producers. Bond market clearing requires that the aggregate net bond holdings of households are equal to government debt. Since some households are borrowers, gross outstanding debt consists of both private and public claims.

2.3 Equilibrium

In general, an equilibrium is a collection of production and consumption plans together with asset holdings and prices of goods and assets such that households and firms optimize, markets clear and the government budget constraint and policy rules are satisfied. Since the endogenous variables – agents' choices and prices – are random variables that depend on exogenous shocks, the standard definition implicitly assumes that agents understand the correct mapping from exogenous to endogenous variables; in other words, they have *structural knowledge* of the economy.

In addition, the definition of equilibrium may restrict the stochastic discount factor used to compute the objective of firms and unions. For example, when financial markets are complete, it is usually set to households' common intertemporal marginal rate of substitution (IMRS). With frictional markets, IMRSs need not be equated and there is no generally agreed-upon way to formulate the objective of the firm. In the New Keynesian literature, however, this issue is not central since models are characterized via first order approximations. The only component of the SDF that matters for firm choice is then the steady state mean. One convention is to set the SDF equal to households' common discount factor, or $M_t^S = \beta$. In other words, all firms are run as if owned by a risk neutral agent.

Worst-case belief. In a model with multiple priors utility, household optimization determines not only optimal choices, but also worst-case beliefs supporting those choices. From the definition of utility (5), ambiguity averse agents take the same equilibrium actions as pessimistic expected utility agents whose subjective belief is the worst case belief. Computation of multiple priors modes thus typically proceeds in two steps. The first conjectures the worst case beliefs of all agents after every history. Given the conjecture, the law of motion of the economy can be characterized using observational equivalence with an economy with pessimistic agents. A second step then verifies that the conjecture is indeed correct.

Since belief sets in our model are parametrized by intervals of means for TFP innovations as in (6), worst case beliefs are characterized by a stochastic process for the worst case conditional mean. A natural conjecture is that, for all households, at all dates and in all states, the worst case mean for the next TFP innovation is the lowest possible mean, or

$\mu_t^* = -a_t$. Intuitively, a negative aggregate TFP innovation decreases the available resources and thus overall surplus in the economy. Individually, through equilibrium prices, it lowers labor and capital income as well as rents. We thus expect the continuation value of any household, whether worker or entrepreneur, and whether able to adjust capital or not, to be increasing in aggregate TFP state, so low TFP is the worst case.

In what follows, we indicate conditional expectations taken under the worst case belief by stars to distinguish them from expectations under the true DGP. We also denote by ξ_t the change of measure, or Radon-Nikodym derivative, from the true DGP to the worst case belief, that is, a stochastic process with a conditional mean of one (under the true DGP) such that for any random variable $E_t^*[v_{t+1}] = E_t[\xi_{t+1} v_{t+1}]$. Relative to a rational expectations benchmark, the process ξ_t thus introduces a wedge in households' intertemporal first order conditions. For example, the FOC for bonds of an unconstrained net lender household is

$$x_{it}^{-\gamma} = \beta E_t \left[\xi_{t+1} x_{it+1}^{-\gamma} \frac{R_{t+1}}{\pi_{t+1}} \right]. \tag{17}$$

Agents thus choose bonds *as if* future consumption is expected to be low. This is how the model captures precautionary savings due to aggregate uncertainty.

The wedge ξ_t has two key properties. First, its conditional distribution varies over time. In particular, when ambiguity a_t is higher, the worst case belief is relatively more pessimistic. The model therefore captures, for example, time variation in the strength of the precautionary savings motive and its effect on interest rates. Second, it shows up in first order approximations of conditions such as (17). This is because of ambiguity in means: since utility is concave, the right hand side depends on consumption, which is increasing in aggregate TFP. In a first order approximation, the right hand side thus contains a linear term in log TFP, so ambiguity a_t enters linearly.

Firms and unions in equilibrium. Under ambiguity, all households value the future *as if* they held the worst-case belief. Other than that, their outlook is as in the standard model. The natural generalization of the usual convention is thus that firms and unions evaluate production plans by computing *worst case* profits and discounting them by the factor β. Formally, we assume that the SDF relevant for computing the present value of profits is

$$M_t^S = \beta \xi_t.$$

It follows that ξ_t also introduces wedges in all firm first order conditions that are correlated with those for households: all of them capture cautious behavior under ambiguity.

With this convention in place, we can follow standard steps in the New Keynesian literature to summarize aggregate consequences of firm and union behavior in a few difference

equations. In particular, the optimal behavior of intermediate goods producers and unions can be solved out in closed form. Aggregation properties of CES production functions as well as Calvo price and wage setting then deliver Phillips curves for final goods price inflation as well as wage inflation $\pi_t^W = W_t/W_{t-1}$. Let mc_t denote the unit cost of making intermediate goods according to the production function (1), an explicit function of the rental rate of capital r_t and the real wage for labor services w_t^F. We note that through those prices, marginal cost reflects TFP and hence is perceived as ambiguous in equilibrium.

We write both Phillips curves directly without terms that are not relevant for a first order approximation:

$$\log\left(\frac{\pi_t}{\bar{\pi}}\right) = \beta E_t^* \log\left(\frac{\pi_{t+1}}{\bar{\pi}}\right) + \kappa_Y \left(mc_t - \tfrac{1}{\mu^Y}\right), \tag{18}$$

$$\log\left(\tfrac{\pi_t^W}{\bar{\pi}_W}\right) = \beta E_t^* \log\left(\tfrac{\pi_{t+1}^W}{\bar{\pi}_W}\right) + \kappa_w \left(\tfrac{w_t}{w_t^F} - \tfrac{\zeta-1}{\zeta}\right), \tag{19}$$

where the slope parameters $\kappa_j = \frac{(1-\lambda_j)(1-\lambda_j\beta)}{\lambda_j}$, for $j = Y, w$, reflect the frequencies of price and wage adjustment. The only difference to standard New Keynesian models is that expectations are taken under the worst case belief. Prices and wages are adjusted in response to news about future *worst case* marginal cost mc_t and $\frac{w_t}{w_t^F}$, respectively. This cautious pricing and wage setting importantly shapes firms' response to uncertainty shocks.

The first order condition for capital goods producers is also intertemporal, due to the form of adjustment costs, and links the price of capital and investment:

$$q_t \left[1 - \phi \log \frac{I_t}{I_{t-1}}\right] = 1 - \beta E_t^* \left[q_{t+1} \phi \log\left(\frac{I_{t+1}}{I_t}\right)\right]. \tag{20}$$

Even without the forward looking term, however, it provides a link between ambiguity that affects the price of capital and investment. Indeed, households require compensation for ambiguity via a higher *capital premium*, i.e. an excess return,

$$Prem_t = \frac{q_t + r_t^k}{q_{t-1}} - \frac{R_t}{\pi_t}, \tag{21}$$

and more so after an uncertainty shock. Higher excess returns come about because of low prices, which makes investment - the production of capital goods – less attractive.

In addition, optimal firm behavior implies that factor prices equal marginal factor products for intermediates and that utilization optimally trades off the current marginal product versus a higher depreciation. The symmetry of shocks across intermediates delivers the

aggregate conditions

$$w_t^F = \alpha mc_t Z_t \left(\frac{u_t K_t}{N_t}\right)^{1-\alpha} \tag{22}$$

$$\text{and} \qquad r_t + q_t \delta(u_t) = u_t (1-\alpha) mc_t Z_t \left(\frac{N_t}{u_t K_t}\right)^{\alpha}, \tag{23}$$

$$q_t \delta'(u_t) = (1-\alpha) mc_t Z_t \left(\frac{N_t}{u_t K_t}\right)^{\alpha}. \tag{24}$$

Since costs and benefits for these decisions are within-period, the presence of ambiguity does not alter them.

Finally, we have equations for profits earned by unions and firms - in both cases rents come from a markup over marginal cost:

$$\Pi_t^U = \left(w_t^F - w_t\right) N_t, \qquad \Pi_t^F = \left(1 - mc_t\right) Y_t. \tag{25}$$

Unions use their market power to charge labor packers a higher wage than what they pay households in the spot market – although they rebate the rents to worker households. Firms use their market power to charge final goods producers a higher price than the factor costs they pay for labor and capital – rents that accrue to entrepreneurs in equilibrium.

Recursive household problem. Our functional form for utility implies that optimal hours satisfy a simple static first order condition that equates the marginal disutility of labor to the hourly wage:

$$h_{it} n_{it}^{\sigma_f} = (1-\tau) w_t h_{it}. \tag{26}$$

The absence of wealth effects means that labor supply is determined separately from savings and portfolio choice. The constant Frisch elasticity further implies that labor productivity cancels so all workers work the same number of hours. We denote by S the aggregate state of the economy, described further below. For now, what matters is that all aggregate endogenous variables, such as the wage, can be written as functions of S.

To describe savings and portfolio choice, we summarize an individual household's state by a tuple (b, k, h, χ, S). The first four state variables are idiosyncratic: b and k are bond and capital positions at the beginning of the period, h is labor productivity, $\chi \in \{A, N\}$ indicates whether the household is allowed to adjust capital ($\chi = A$) or not. Let $C(b, k, h, \chi; S)$ denote the household within-period constraint set for the choice vector (x, b', k'), that is, consumption, bonds and capital. It comprises the budget and borrowing constraints, as well as how trading constraints depend on χ. We note that the state variables capture in particular whether or not the agent is an entrepreneur, indicated by $h = 0$.

The household Bellman equation under the worst case belief is

$$V(b,k,h,\chi;S) = \max_{(x,b',k')\in C(b,k,h,\chi;S)} u(x) + \beta E^* \left[V(b',k',h',\chi';S')|h,S\right]. \tag{27}$$

Here the expectation conditions on the agent's idiosyncratic productivity state as well as the aggregate state. In particular, the aggregate state S contains current TFP and ambiguity, whereas the aggregate state S' contains TFP next period. The expected continuation value is computed under the worst case transition density from (6). We denote the optimal bond and capital policies by $b^*(b,k,h,\chi;S)$ and $k^*(b,k,h,\chi;S)$, respectively.

An important endogenous state variable is the joint distribution of asset holdings, labor productivity and adjustment states, denoted as $\Theta(b,k,h,\chi)$. It evolves in equilibrium as

$$\Theta'(b',k',h',\chi') = \int_{b'=b^*(b,k,h,\chi,S),k'=k^*(b,k,h,\chi,S)} \Phi(h,\chi,h',\chi';S^x)\, d\Theta(b,k,h,\chi), \tag{28}$$

where Φ is the exogenous transition density of the labor productivity and adjustment states and S^x is the exogenous component of S. This transition density depends on S^x, for example, because labor productivity is time-varying.

Recursive equilibrium. The exogenous state of the model is $S_t^x = (Z_t, a_t, \bar{\sigma}_{h,t}^2, \pi_t^*, \varepsilon_t^R)$. The aggregate state $S_t = (S_t^x, K_t, I_t, \Theta_t), Y_{t-1}, R_t, G_{t-1}, L_{t-1}, B_t$ further contains the capital stock, investment, the joint distribution of income and asset holdings as well as past variables that enter government reaction functions. A recursive equilibrium consists of stochastic processes for the aggregate state given by (4), (7), (9), (12) and (13), as well as prices $(w_t, w_t^F, q_t, r_t, R_{t+1}, \pi_t, \pi_t^w)$, profits Π_t^U, Π_t^F and government policy (G_t, L_t, B_{t+1}) that are all functions of the state S_t, such that

1. equations for optimal firm behavior (18)-(24) hold,
2. union and firm profits are defined by (25),
3. government reaction functions (11), (14) and (15) as well as the government budget constraint (16) hold,
4. markets for final goods, bonds and capital clear, with household sector demand computed as the integral over individual households' optimal choices (x, b', k') derived from the recursive household problem (27) using the cross-sectional distribution Θ_t over individual household states,
5. the market for labor clears, that is, firms' demand equals household sector supply from (26), so that $N_t = ((1-\tau)w_t)^{\frac{1}{\sigma_f}}$,

6. capital evolves according to $K_t = (1 - \delta(u_t))K_{t-1} + I_t$,

7. the distribution of household states evolves according to (28),

8. the household value function in (27) is strictly increasing in aggregate TFP Z.

Conditions 1-7 are standard: the only difference to the typical definition of a recursive rational expectations equilibrium is the worst case expectations in the intertemporal firm conditions (18)-(24) and the Bellman Equation (27). Since TFP is the only ambiguous shock in our model, this only affects the conditional distribution of TFP - expectations for all other variables, whether aggregate or idiosyncratic, are taken under the true DGP also in the equilibrium with ambiguity. Condition 8 ensures that selecting the lowest conditional mean for TFP, $\mu_t^* = -a_t$, is indeed the worst case belief: this is the case if all households are better off when TFP is higher. We emphasize again that worst case beliefs are a device to capture cautious behavior that comes from nondegenerate belief sets - rational expectations equilibrium instead assume that each belief set in (5) is singleton.

Representative agent benchmark. A second useful benchmark assumes frictionless complete markets. Equilibrium then works as if there is a representative household endowed with total labor who receives all rents from firms and unions. The household constraint set C changes to capture that this household can always freely adjust capital and bonds. The representative household perceives ambiguity about TFP, however, so wedges in the household Bellman equation and intertemporal firm conditions remain in place. The only change to the equilibrium condition is therefore that the representative household problem is used for market clearing in condition (4) and the evolution of the distribution in condition (7) is no longer relevant.

2.4 Solution method

We follow Reiter (2009) in solving the household problem globally while approximating aggregate dynamics by a first-order perturbation, including for the distribution Θ_t. In particular, we build on Bayer et al. (2024a), whose implementation of the Reiter approach allows for full information estimation of the model dynamics. The fact that ambiguity matters to first order for model allows us to adapt this approach with little extra complication. In addition, Bayer et al. (2024a) gives an upper bound on the achievable dimensionality reduction for first-order solutions and shows how to choose it optimally. These results apply to our model even though it captures effects of aggregate uncertainty.

The key new feature with ambiguity is that we need to distinguish between agents' perceived, or worst-case, law of motion for the economy – that includes the worst-case

beliefs about TFP, the ambiguous shock – and the actual law of motion that includes the true DGP for TFP. This is in contrast to rational expectations equilibrium (REE) where the actual and perceived laws of motion always coincide. However, a common denominator of both concepts is that the response of endogenous variables to shocks is determined entirely by the perceived law of motion. For example, precautionary savings are chosen *as if* future TFP is low and interest rates adjust to this behavior to clear the bond market. We can therefore proceed in two steps: we first use standard methods familiar from REE analysis to approximate the perceived (worst-case) law of motion and hence in particular compute all responses of endogenous variables. We then describe the actual model-implied law of motion taking into account the true TFP process.

For the first step, we modify the evolution of the exogenous state S_t to include the worst case TFP process that always sets $\mu_t = -a_t$ in (6). With this modification, the definition of recursive equilibrium describes a rational expectations equilibrium: perceived and actual laws of motion agree, albeit on the "wrong" dynamics. We define the *worst-case steady state* as the deterministic steady state of this law of motion. It can be found in standard fashion using global methods. We note that the worst case steady state TFP, Z^* say, is given by

$$Z^* = \bar{Z} \exp\left(\frac{-\bar{a}}{1-\rho_z}\right), \tag{29}$$

where $\bar{a}$ is the long run value of ambiguity in (7) and $\bar{Z}$ is the long run value under the true process in (4). We collect all other state variables in a vector $\mathbf{Y}$. We approximate the worst-case law of motion by log-linearizing around $(Z^*, \mathbf{Y}^*)$, applying directly the approach of Bayer et al. (2024a).

The second step characterizes equilibrium dynamics by combining the worst case law of motion, derived in the first step, and the true DGP. While endogenous variables respond as they would in a world with low TFP where the economy converges to the worst-case steady state, the actual process of TFP shocks (4) satisfies

$$\log Z_t = E^*_{t-1} \log Z_t + a_{t-1} + \epsilon^Z_t. \tag{30}$$

In other words, from the perspective of the worst-case belief that shapes action, agents always act *as if* they are positively surprised. From the perspective of the econometrician, agent responses instead look cautious: savings is higher than it would be under certainty equivalence, and interest rates are lower.

Our approach implies that aggregate uncertainty affects both the model's steady state and its dynamics. On the one hand, consider the *ergodic* steady state, $(\bar{Z}, \bar{\mathbf{Y}})$ say, as measured

by the econometrician. It is the long run value of variables when (i) the long run value of actual TFP in (4) equals $\log Z = \log \bar{Z}$, as in a deterministic steady state, but at the same time (ii) variables follow their log-linear equilibrium response to states under the worst-case law of motion, that is, they evolve *as if* TFP is declining towards its long run worst-case value of $\log Z^* = \log \bar{Z} - \frac{\bar{a}}{1-\rho_z}$. The anticipation of bad times (ii) makes steady state choices and prices incorporate caution. On the other hand, the log-linearized equilibrium responses derived under the worst-case belief also matter for dynamics around the ergodic steady state $(\bar{Z}, \bar{\mathbf{Y}})$, for example, for how the economy responds to times of higher/lower ambiguity than its steady state value of $a_t = \bar{a}$.

3 Estimation

The frequency of the model is quarterly. We estimate it using quarterly US data from 1985Q1 to 2019Q4, based on a two-pronged approach. We target average moments over this period that speak to key portfolio choice and asset pricing moments of the ergodic wealth distribution. At the same time, we also fit the model to business cycle and asset price dynamics. For the latter, we leverage the linearity of the model's state space representation and the normality of the shocks to fit the dynamics using standard full-information Bayesian likelihood methods, as discussed in e.g. Smets and Wouters (2007), Fernández-Villaverde (2010) and Bayer et al. (2024a). In addition to the estimated parameters, we also follow standard practice in the literature, and set some parameters based on external evidence.

Data used for estimation. We use five average moments that are common in the HANK literature and pin down the distribution of wealth, as well as the size of the government: the average ratio of capital to output, capital to government debt, top 10 wealth share, the share of borrowers, and government spending to output. For the full information likelihood estimation we use six observable time series. In particular, we include the growth rates of per capita hours, private consumption, investment (all in real terms), the log difference of the GDP deflator and the (shadow) federal funds rate. Our model is stationary so all growth rates are demeaned. These data are standard in the estimation of typical DSGE models. In addition, we use the non-demeaned capital premium from Gomme et al. (2011) as observable. Since our model has fewer structural shocks than observables, we allow for iid measurement error in the observation equation of the state-space representation that links all six observed variables to their model counterparts. Appendix A.1 lists the data sources.

3.1 Estimation approach

Our goal is to fit two features of the data. First, the ergodic steady state of the model should be consistent with a number of long run averages from the data, introduced above and listed in Table 3. They include the mean capital premium as well as moments of the wealth distribution commonly used to quantify HANK models. Second, movement around the steady state should resemble the fluctuations around the mean in our six observables. In our model, as in other models where anticipation of aggregate uncertainty matters for decisions and prices, the two features of the data are closely connected, in the sense that a common set of parameters shapes both. For example, parameters that govern the dynamics of uncertainty (be it ambiguity or risk) and the model's response to it matter not only for the business cycle, but also for the mean capital premium, and hence the average relative rate of return earned by rich capital owners and poorer bondholders, an important determinant of the long run wealth distribution.

It is common practice when estimating linearized business cycle models to divide parameters into two groups. A first group affects the model-implied distribution of the data primarily by shifting the perturbation point, usually the deterministic steady state. The second group is instead responsible for how the model captures fluctuations around that point. The first group is then typically calibrated up front by matching the deterministic steady state of the model to the data, whereas the second group is found via Bayesian estimation of the dynamics. In HANK models this approach is particularly attractive since finding the steady state involves solving nonlinearly a heterogeneous agent model and is therefore relatively more costly. This is why parameters such as the discount factor β or the probability of trading capital λ are usually chosen up front as part of the first group of parameters.

In our context, the standard two-step approach is not applicable because aggregate uncertainty matters for both the steady state and the dynamics. The linear solution method described in Section 2.4 nevertheless allows a tractable estimation approach. We describe the perceived law of motion of the data via perturbation around the worst-case steady state $(Z^*, \mathbf{Y}^*)$. This means that there is still a set of parameters, θ_{SS} say, that affects the solution of the model primarily by shifting the perturbation point. As discussed in Section 2.4, the perceived law of motion works like the law of motion for a rational expectations equilibrium. It thus makes sense to have the θ_{SS} contain the same parameters as in a standard HANK model, including the discount factor and the probability of trading capital.

Further simplification comes from parameterizing the model such that average ambiguity $\bar{a}$ does *not* affect the perturbation point and hence need not be included in θ_{SS}. Our model has the standard property that the (true) level of TFP $\bar{Z}$ is not identified. If we were to estimate it under rational expectations, we would thus set $\bar{Z} = 1$. With this normalization,

the worst-case mean TFP defined in Equation (29) would depend on $\bar{a}$. We thus instead normalize $Z^* = 1$. As a result, the parameter $\bar{a}$ is not relevant for the worst-case steady state. Of course, it is critical for steady state moments, for example the mean capital premium. It affects the model-implied distribution of the data because that distribution depends on the true DGP for TFP. For example, the value of $\bar{a}$ determines true mean TFP $\bar{Z}$ in Equation (29) that helps fit the data best given the normalization of $Z^* = 1$.

These observations motivate an estimation procedure with an outer loop over θ_{SS} and an inner loop over the other parameters. Conditional on values for θ_{SS}, the inner loop runs a full information Bayesian maximum likelihood estimation based on our six observables for the rest of the parameters. For each parameter draw, we thus avoid the costly re-solving of the approximation point $(Z^*, \mathbf{Y}^*)$, which involves the rich wealth distribution, fixed at this step. The estimation converges quickly, leveraging the linearity of the state-space representation. Given the posterior mode obtained in the Bayesian estimation, we compute the ergodic steady state $(\bar{Z}, \bar{\mathbf{Y}})$, and compare moments from Table 3 to their counterparts. The outer loop minimizes the (equally weighted) moment squared distance for the five targeted average moments. We note that since parameters affect jointly the model's distance to the targeted data moments *and* indirectly the Bayesian maximum likelihood score, there is no a-priori reason why under the best fitting parameters the former distance becomes zero.

3.2 Parameters

Our first set of parameters, in Table 1, is set based on external evidence. In particular, we take the idiosyncratic income process from Storesletten et al. (2004), which gives us $\rho_h = 0.98$ and $\bar{\sigma}_h = 0.12$. Guvenen et al. (2014) gives the probability of a household falling out of the top one percent of the income distribution in a given year, which we take to be the transition probability from entrepreneur to worker, $\iota = 6.25\%$. We set the relative risk aversion, γ, to 4, which is common in the incomplete markets literature; see Kaplan and Violante (2014). We set the Frisch elasticity to 0.5; see Chetty et al. (2011). The steady-state price and wage mark-ups are both fixed at 10%, following Born and Pfeifer (2014). The labor share of production, $\alpha = 0.68$, is determined by the average labor income share (given by η). The average quarterly depreciation is $\delta_0 = 0.025$.

Table 2 presents our *estimated* parameters. Its first three rows refer to the previously-introduced notation of set θ_{SS} of parameters, which shift the perturbation point and get adjusted in the outer estimation loop of moment-matching. We have five such parameters: the discount factor β, the trading friction λ, the probability to enter the entrepreneur state ζ, the borrowing wedge $\bar{R}$, and the tax rate τ. The rest of Table 2 comprises the remaining

Table 1: Parameters set externally

Par.	Value	Description	Target
Households			
ρ_h	0.980	Persistence income	Storesletten et al. (2004)
σ_h	0.120	Std. income	Storesletten et al. (2004)
ι	0.063	Trans. prob. E. to W.	Guvenen et al. (2014)
γ	4.000	Relative risk aversion	Kaplan and Violante (2014)
σ_f	0.500	Frisch elasticity	Chetty et al. (2011)
Firms			
α	0.680	Share of labor	Standard value
δ_0	0.025	Depreciation rate	Standard value
$\bar{\eta}$	11.000	Elasticity of sub.	Born and Pfeifer (2014)
$\bar{\zeta}$	11.000	Elasticity of sub.	Born and Pfeifer (2014)

parameters, which get estimated in the inner estimation loop of Bayesian maximum likelihood and for which we report the prior and posterior credible intervals.

We now detail the *prior* construction of our estimated parameters. For the real and nominal frictions, standard in the literature, we follow Justiniano et al. (2011) and impose a *Gamma* distribution with prior mean 5.0 and standard deviation 2.0 for δ_2/δ_1, the elasticity of marginal depreciation with respect to capacity utilization, and a *Gamma* prior with mean 4.0 and standard deviation 2.0 for the investment adjustment costs parameter, ϕ. For the slopes of the price and wage Phillips curves, κ_Y and κ_w, we assume *Gamma* priors with mean 0.10 and standard deviation 0.03, corresponding to contracts with an average length of four quarters.

For monetary policy, we estimate the Taylor rule responses to inflation and output growth, θ_π and θ_Y. We impose $Normal$ distributions with prior means of 1.7 and 0.13, respectively. We allow for interest rate smoothing with the parameter ρ_R. We assume a *Beta* distribution with parameters $(0.5, 0.2)$. For fiscal policy, we estimate the response of government spending and transfers to government debt deviations and output growth. We impose *Gamma* distributions to ensure debt stabilization and countercyclical responses of both rules.

Following Smets and Wouters (2007), the autoregressive parameters of the shock processes and policy rules are assumed to follow a *Beta* distribution with mean 0.5 and standard deviation 0.2. The standard deviations of the shocks follow $Inverse-Gamma$ distributions. We assume a prior mean of 0.1% and a standard deviation of 2% for the nominal shocks. For shocks to ambiguity and idiosyncratic risk the prior mean is 50% and the standard deviation is 25% as they capture relative shifts in uncertainty.

An important question in the estimation is how do we interpret and *discipline* the new

Table 2: Estimated parameters

Parameters θ_{SS} estimated by moment-matching					
Parameter		Value	Parameter		Value
β		0.977	λ		0.073
ζ		6.0E-4	$\bar{R}$		0.044
τ		0.260			
Parameters estimated by Bayesian estimation					
Parameter	Prior	Posterior	Parameter	Prior	Posterior
	Frictions and Ambiguity			Shocks	
δ_s	Gamma(5, 2)	6.036 (6.006, 6.064)	ρ_Z	Beta(0.5, 0.2)	0.945 (0.930, 0.959)
ϕ	Gamma(4, 2)	0.561 (0.366, 0.806)	ρ_A	Beta(0.5, 0.2)	0.963 (0.948, 0.975)
κ	Gamma(0.1, 0.03)	0.046 (0.031, 0.065)	σ_A	Inv.-Gamma(0.5, 0.25)	90.879 (68.496, 120.78)
κ_w	Gamma(0.1, 0.03)	0.101 (0.069, 0.137)	ρ_S	Beta(0.5, 0.2)	0.496 (0.325, 0.670)
$\tilde{z}$	Beta(0.99, 0.01)	0.972 (0.966, 0.978)	σ_S	Inv.-Gamma(0.5, 0.25)	22.182 (15.366, 30.374)
	Monetary Policy			Fiscal Policy	
ρ_R	Beta(0.5, 0.2)	0.171 (0.060, 0.316)	ρ_G	Beta(0.5, 0.2)	0.181 (0.042, 0.422)
θ_π	Normal(1.7, 0.3)	2.202 (1.996, 2.435)	γ_{GB}	Gamma(1.0, 0.2)	0.806 (0.611, 1.030)
θ_Y	Normal(0.13, 0.05)	0.075 (0.048, 0.102)	γ_{GY}	Gamma(1.0, 0.2)	0.954 (0.684, 1.253)
ρ_R^ϵ	Beta(0.5, 0.2)	0.553 (0.344, 0.686)	ρ_L	Beta(0.5, 0.2)	0.361 (0.102, 0.687)
σ_R^ϵ	Inv.-Gamma(0.1, 2.0)	0.136 (0.113, 0.175)	γ_{LB}	Gamma(0.2, 0.2)	0.206 (0.155, 0.263)
ρ_π^ϵ	Beta(0.5, 0.2)	0.948 (0.929, 0.985)	γ_{LY}	Gamma(0.2, 0.2)	1.044 (0.865, 1.222)
σ_π^ϵ	Inv.-Gamma(0.1, 2.0)	0.045 (0.045, 0.045)			

Notes: The table displays the set of parameters estimated through moment-matching and Bayesian likelihood estimation, respectively. For the latter we show their prior distribution and posterior means. The 90% credible intervals are shown in parentheses. Posteriors are obtained by an MCMC method. The standard deviations have been multiplied by 100 for better readability.

parameter $\bar{a}$ governing long run ambiguity. Here we build on Ilut and Schneider (2014), which discuss what sets of models are consistent with a sample of iid innovations measured by an econometrician. They propose a bound on the set of one-step ahead mean beliefs that is proportional to the standard deviation of the innovation measured by an econometrician. The idea is that if an econometrician estimates a more volatile process, there is more room

for agents' concern about ambiguity and hence the interval of means can be wider.

We thus follow the estimation approach in Ilut and Schneider (2014) and look to bound upwards the parameter $\bar{a}$ by one standard deviation σ_z of the TFP innovation. We find that in our estimated model this upper bound is tight. In particular, when we estimate the model without imposing this bound on $\bar{a}$, its resulting estimated value is significantly larger than σ_z. Disciplined by the bound constraint, the reported results are for a restricted estimation version where the upper bound is tight and so $\bar{a} = \sigma_z$.

In terms of a specific prior distribution for long run ambiguity, we find it convenient and informative to estimate the ratio $\tilde{z} \equiv \frac{Z^*}{Z}$, of the long run value of TFP under the worst-case belief (in turn normalized to $Z^* = 1$) relative to the true process, by using a prior *Beta* distribution. When $\tilde{z}$ is small, agents worry about a path towards a low long run TFP value, or put differently, the true one is large compared to that worst-case belief. Specifically, in Table 2 the prior mean for the ratio is 0.99, with a standard deviation of 0.01, consistent with the priors in Ilut and Schneider (2014). Given an estimated value for the persistence ρ_z of TFP, the primitive long run value $\bar{a}$ of the one-step ahead ambiguity is then implicitly estimated by the ratio $\tilde{z}$ of long run values. Indeed, by Equation (29) we have $\bar{a} = -(1-\rho_z)\log\tilde{z}$, which also gives the value σ_z given the tight disciplining bound on $\bar{a}$. The parameter σ_z thus does not show up as an additional free parameter in Table 2.

The *posteriors* for the Bayesian estimation in Table 2 get reported using a single RWMH chain after an extensive mode search. After a long burn-in, 150,000 draws from the posterior are used to compute the posterior statistics. Appendix A.2 provides details on convergence. We only briefly comment here on the estimated values for the standard set of parameters. These parameters influence the model's dynamics, which we discuss in detail in Section 5. For now, we note that the parameter estimates for the nominal and real frictions and for the policy rules are broadly consistent with the literature. We find sizable countercyclical fiscal policy, a strong reaction of the Taylor rule to inflation, nominal stickiness of around 4 quarters for wages and 5 quarters for prices, and higher frictions in capital utilization than in investment adjustment. In addition, there are persistent shocks to the inflation target and less persistent shocks to the Taylor rule, which we estimate quite accurately from observations of the policy rate and inflation.

4 Ergodic steady state and aggregate uncertainty

Aggregate uncertainty matters for our ergodic steady state. We now discuss its implications and mechanisms. Table 3 reports model-implied moments, based on the parameters set externally in Table 1 and posterior estimates of Table 2. The middle column contains the

Table 3: Ergodic moments

	Data	Ergodic SS	Worst-case SS
	Estimation moments		
Capital to output	9.88	9.87	9.41
Liquid to illiquid	0.20	0.20	0.32
Top 10 wealth	0.67	0.71	0.60
Share of borrowers	0.16	0.16	0.08
Gov. spending	0.22	0.22	0.18
Capital premium (%)	6.06	5.55	2.31
	Non-estimation moments		
Share of zero-liquidity	0.20-0.30	0.22	0.10

model-implied moments in the ergodic steady state and for comparison the same moments in the worst-case deterministic steady-state are reported in the last column. The difference between these moments can be used to identify the role of aggregate uncertainty in driving the long run behavior of our model.

Consider the first five rows of moments. These are the targeted moments in the moment-matching part of our estimation. The values of the five parameters, θ_{SS} in Table 2, are primarily responsible for fitting these moments, given the posterior mode of the Bayesian estimation. These moments get fit very well in the ergodic SS. The sixth moment, the average capital premium, is part of the Bayesian likelihood estimation since capital premium is an observable. The model-implied premium, at 5.55%, is close to the sample average of 6.06%. Finally, the last row presents a moment that was not part of the estimation at all, namely the share of agents with zero-liquidity, which the model also gets close to.

We emphasize two key parameters in shaping the ergodic moments. The first is the estimated value of the trading friction, as a probability $\lambda = 7.3\%$ of accessing the capital market. This parameter is of particular importance for the two-asset HANK literature. The trading probability comes in actually higher (i.e. the friction is lower) than the corresponding value of 6.2% in the earlier work of Bayer et al. (2024a), which does not feature aggregate uncertainty. The second is the steady-state amount of ambiguity, which gets estimated within the Bayesian estimation step. From Table 2, the posterior value of the ratio $\tilde{z}$ is 0.973, meaning that under the worst-case belief the long run TFP is about 2.7% lower than under the true process. The one-step ahead estimated ambiguity, by Equation (29), can be read as $\bar{a} = 0.0015$. Interestingly, this value is about half of the corresponding estimate of $\bar{a}$ in the representative agent business cycle model of Ilut and Schneider (2014).

Ergodic steady state effects of aggregate uncertainty. There are two fundamental mechanisms through which aggregate uncertainty affects the ergodic moments. One is

precautionary savings and the other is an increase in the uncertainty-adjusted return on capital. These effects occur since agents act *as if* the economy is on a path towards a lower long run value of TFP. Therefore, agents (i) engage in more precautionary savings, and (ii) simultaneously require a higher equilibrium compensation for holding the uncertain capital.

We see these two forces at work in Table 3. First, the precautionary savings mechanism leads agents to invest in more capital, increasing capital to output. In addition, the portfolio choice between liquid government debt and the less liquid capital gets shifted in equilibrium towards the latter, reducing the ergodic liquid/illiquid asset ratio. The shift occurs as aggregate uncertainty leads agents to increase demand for both assets to save in, but the supply of capital is effectively more elastic in steady state than that of the government debt – the latter being determined by the government budget constraint of Equation (16). At the same time, the same precautionary savings force reduces the equilibrium real rate, which doubles the ergodic share of borrowers compared to the deterministic steady state.

Second, investors' exposure to aggregate uncertainty in capital, in the ergodic steady state, demands compensation through an excess return over the risk-free real rate.. Our headline result here is that aggregate uncertainty accounts for more than half of the model-implied ergodic capital premium of 5.55%. In particular, in the worst-case deterministic SS, where there is no compensation for uncertainty, the premium of 2.31% reflects *only* a financial compensation for illiquidity – the only source operating in standard linear RE HANK models. Instead, in our ergodic SS, the total premium reflects a liquidity *and* an uncertainty component. In particular, aggregate uncertainty opens up a premium that is larger by 3.21% than in the deterministic SS, to account for the total of 5.55%.

The large uncertainty premium also matters for the wealth distribution. In particular, following insights in the two-asset HANK literature (e.g. Kaplan and Violante (2014), Kaplan et al. (2018)), a higher premium also increases the share of wealth held by the top 10 percent by 11 percentage points and increases the share of agents with zero-liquidity, i.e. Hand-to-Mouth. Indeed, in Table 3 this latter share more than doubles from its deterministic SS to be 22%, in line with the otherwise untargeted data moment. This channel is important, as it shows that a model with aggregate uncertainty can produce an empirically relevant equilibrium premium, which is a key mechanism to generate a relevant share of Hand-to-Mouth agents, in particular of the wealthy type with illiquid assets.

The premium in a counterfactual RANK model. Our headline result that aggregate uncertainty generates a sizable premium in the ergodic steady state of our HANK model can be further compared to the premium obtained in a counterfactual RANK model that keeps the same parameters, including ambiguity, but eliminates the incomplete markets aspects of our economy (recall the discussion in Section 2.3 of how this variant is constructed).

In particular, we find that in the ergodic steady state of such a counterfactual RANK the premium is significantly smaller, at only 0.11%. Since the trading friction does not operate in this counterfactual, its premium is entirely a compensation for ambiguity. Thus, although the steady-state ambiguity is by construction the same in this counterfactual RANK, its equilibrium compensation for holding the uncertain capital is more than an order of magnitude *smaller* than in our HANK model, which we have reported above to be at 3.21%.

This counterfactual result thus also highlights that the heterogeneity in portfolios is crucial in producing a sizable equilibrium compensation for uncertainty and overall average premium. The mechanism has to do with the effective marginal investor in capital being different in the two economies. In particular, the rich households in the HANK model hold most of the capital. Their future income is thus more heavily exposed to the uncertain capital returns than a counterfactual representative agent which holds relatively more of their income as labor income. Due to this heightened exposure, a larger ambiguity premium opens up as an equilibrium compensation for the marginal investor in the HANK model.

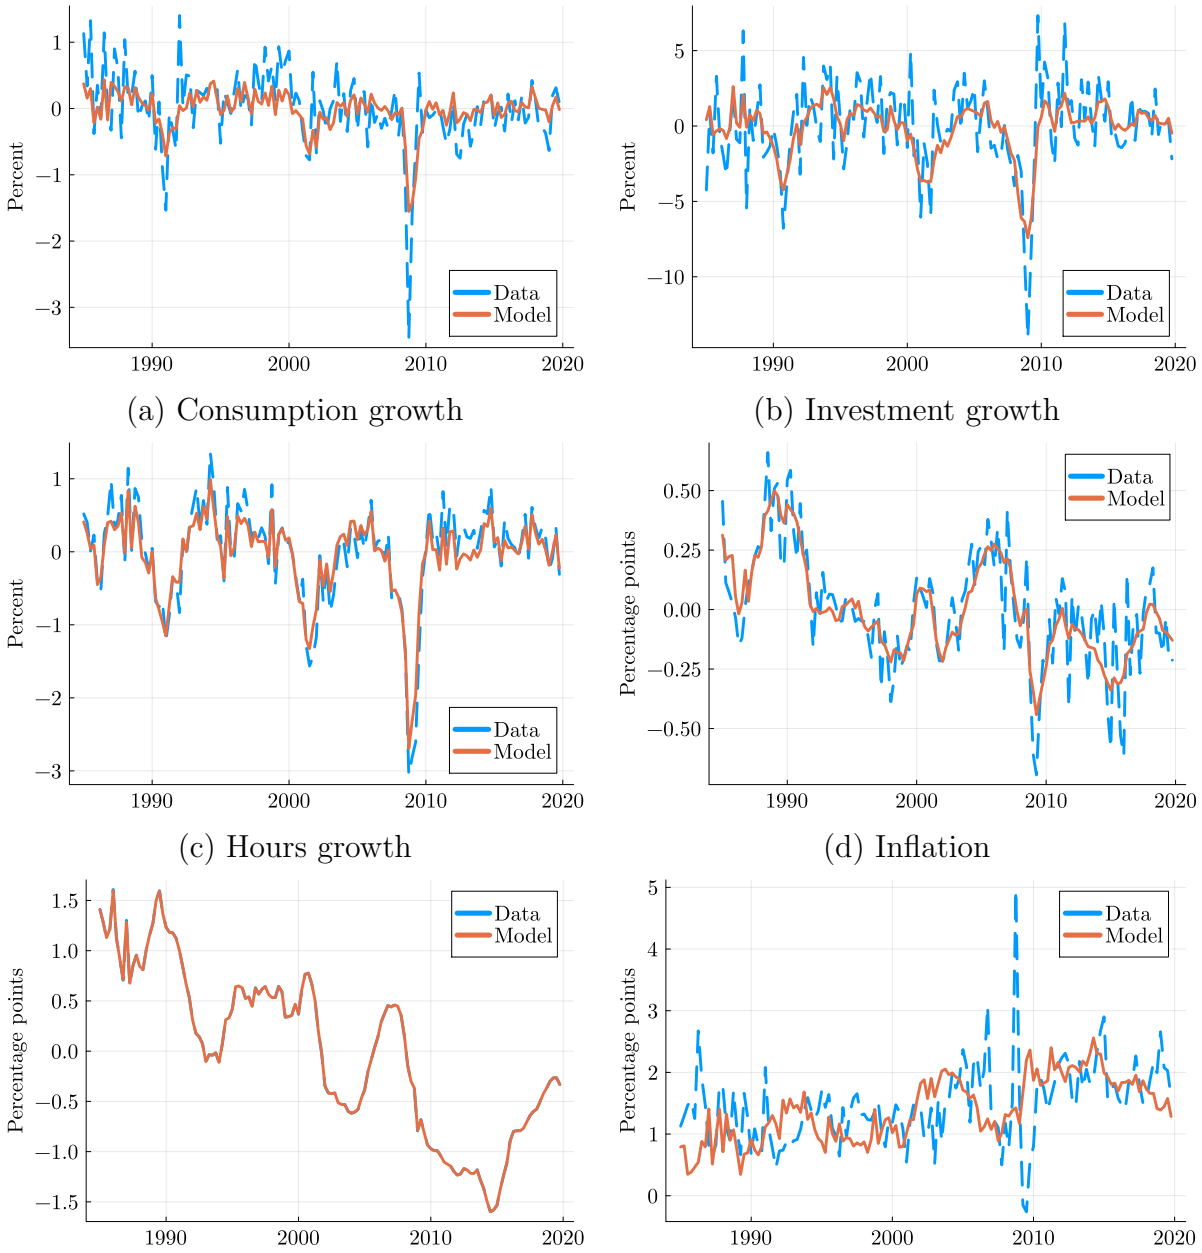
In our estimated model aggregate uncertainty thus provides a powerful economic force in driving the *long run* behavior of our model. In the next section we discuss how in turn time-variation in aggregate uncertainty emerges as the main quantitative driver of *short run*, business cycle dynamics. In that discussion, a common narrative for the quantitative significance of both short- and long run effects appears, namely the critical *interaction* of aggregate uncertainty and the trading frictions.

5 Business cycle dynamics

Aggregate uncertainty, modeled here as ambiguity, is the main business cycle driver in our estimated model. We discuss the model's empirical fit, response to shocks and mechanisms through a series of results.

First, Figure 1 plots the six observables (the 'Data' blue lines) against the corresponding historical path implied by our model estimates (the 'Model' red lines), computed by a Kalman smoother. The difference between the lines is the estimated measurement error, which we allowed for each observable. The model does a good job fitting the business cycle comovement of investment, consumption and hours growth. Out of these three real variables, the fit is closest for hours growth, since this series is the most persistent and thus less likely to be generated by measurement error. The model also closely tracks movements in the another persistent series, namely the nominal interest rate, and also matches well the business cycle and lower frequency movements in inflation and the capital premium.

Variance decomposition. To understand how our model generates a business cycle,



(a) Consumption growth

(b) Investment growth

(c) Hours growth

(d) Inflation

(e) Nominal rate

(f) Capital premium

Figure 1: Model vs Data: Kalman smoother for the estimated HANK model and the data used in estimation.

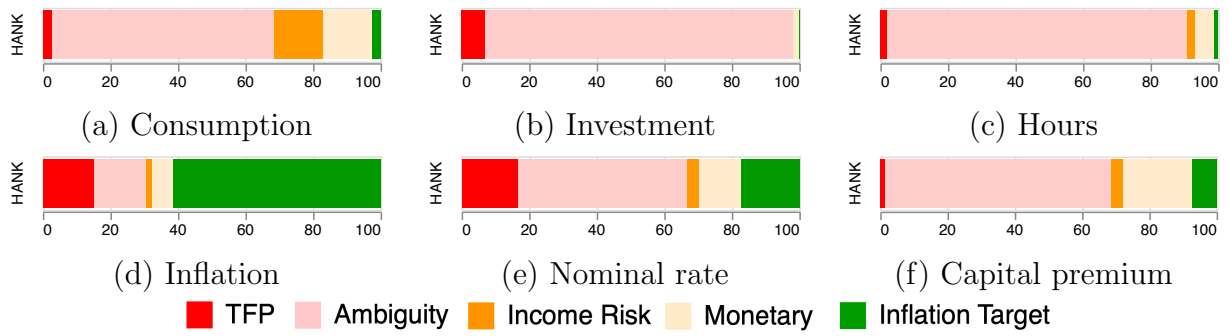


Figure 2: Variance decompositions for the estimated HANK model

with its corresponding movements in prices, we start by reporting in Figure 2 the role played by each shock in driving model-implied variation. In particular, we perform a variance decomposition at business cycle frequency following Uhlig (2001). The Figure shows that in our baseline HANK model the ambiguity shock accounts for the bulk of fluctuations in aggregate quantities, around 90% of the model-implied variability in investment and hours, and about 65% for consumption. The same aggregate uncertainty shock is also the most important shock for the nominal rate and the capital premium, driving more than half of their fluctuations, while being less significant for inflation dynamics.

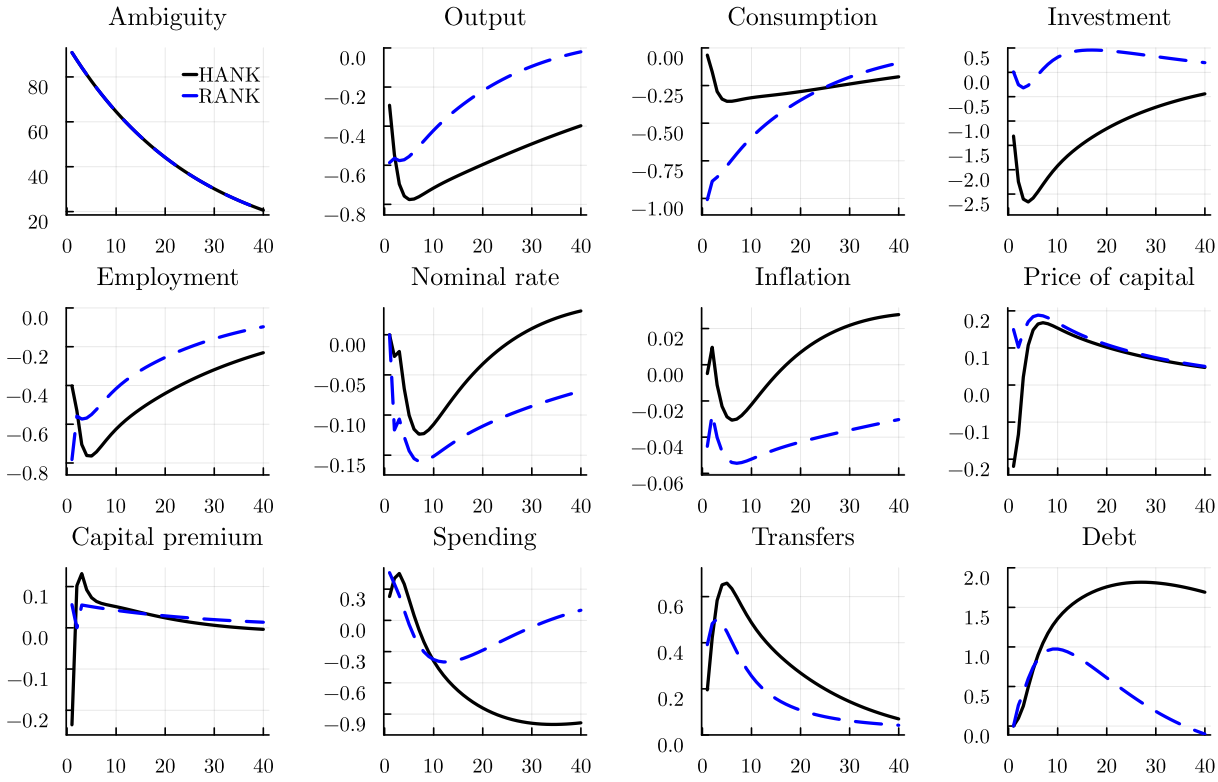
The other shocks play overall a more muted role. In particular, the idiosyncratic income risk shock has a negligible effect, except for consumption, where it accounts for about 15% of variation. The remaining shocks help the model generate more variation on the nominal side. Indeed, the inflation target shock accounts for most of the model-implied inflation fluctuations. The monetary policy shock moves the real interest rate, and through it affects the capital premium and consumption dynamics - the latter effect accounting for around 15% of variation. Finally, in the model TFP shocks matter little for quantities, but more so for the business cycle variation in inflation and nominal interest rate.

Next we use impulse response functions (IRFS) to understand why the ambiguity shock emerges as the prime business cycle factor driving the positive comovement of consumption, investment and hours worked, while also significantly contributing to movements in the capital premium.

The ambiguity shock in the baseline HANK. In particular, Figure 3 plots in solid dark lines the IRF to our ambiguity shock in the baseline HANK model. We will later also draw comparisons to the counterfactual RANK model, in dashed blue lines (recall its introduction in Section 4).

A loss of confidence over the conditional distribution of TFP, i.e. a negative aggregate

uncertainty shock, leads in the baseline HANK model of Figure 3 to a recession in which consumption, investment and employment all fall significantly on impact and remain persistently depressed. Intuitively, an increase in ambiguity acts like agents receive bad news about the conditional mean of aggregate TFP. We can decompose the economic effect of this anticipation along several margins, or, put differently, along several correlated 'wedges' that get activated by the ambiguity shock.



Y-axis: Percentage points for inflation, nominal rate and premium, otherwise percent. X-axis: Quarters.

Figure 3: Impulse responses to an ambiguity shock in estimated HANK model and counterfactual RANK (under the same parameters)

First, this lack of confidence affects the precautionary saving desire of all types of households – whether households are mostly exposed to aggregate TFP through their labor or capital income, they now worry that their respective future income streams are lower. This leads to precautionary saving and a desire to cut consumption and save. Overall, this precautionary effect is a type of 'wedge' in the Euler equation for saving that resembles the discount factor taken as a primitive shock in many NK models, with or without heterogeneity.

By itself, this precautionary effect alone can generate comovement between consumption and labor for standard reasons present in NK models. Namely, due to nominal rigidities, equilibrium good prices and wages in this recession do not adjust sufficiently, and the mon-

etary policy through its Taylor rule does not lower sufficiently the real rate to undo those effects. Equilibrium markups in the good and labor markets rise, leading to a demand driven recession. In the absence of those rigidities, the typical Barro and King (1984) logic would prevail and labor and consumption would counterfactually move in opposite directions.

Second, what about aggregate investment? A pure precautionary saving effect would typically imply that aggregate investment would increase, as that is the equilibrium channel through which savings occurs.[8] The difference here is that an increase in aggregate uncertainty over TFP also decreases the *uncertainty-adjusted return* on investing in capital. This caution is formalized in our model as agents evaluate the future under the worst-case conditional belief for TFP. As a result of this worry, there is now also an *intertemporal substitution* away from capital. Another 'wedge' now simultaneously appears in households portfolio choice, making the uncertain capital particularly less attractive than the risk-free bond. As a result, there is a strong economic force that lowers the incentive to invest in capital. Put together, consumption, labor *and* investment significantly and persistently fall.

We note that this strong positive comovement of major aggregates is accompanied by two other dynamics that the data favors in its quantitative estimation. One is nominal price dynamics and arises from the property that higher aggregate uncertainty also affects firms' decisions. Of particular importance is the effect on goods price-setting, in the Phillips curve in Equation (18). On the one hand, a standard cost channel is at work: on impact, due to the lower household demand, marginal cost falls and pushes those firms who can adjust to lower prices. On the other hand, higher aggregate uncertainty also manifests as a novel wedge in the Phillips curve, since ambiguity shows up in the stochastic discount factor relevant for firms' intertemporal decisions (see the discussion in Section 2.3). In particular, through the *as if* risk neutral owner's worst-case belief of low future aggregate TFP, firms now worry that future *equilibrium* marginal costs will be higher. Due to nominal rigidities, firms that have an ability to reset prices anticipate that not increasing current prices would thus lead them exposed to sub-optimally low future markups. Therefore, this anticipation is a force that incentivizes firms to raise current prices. This precautionary effect is important in explaining the relatively small movements of inflation in an otherwise deep recession generated by the shock. As such, the ambiguity shock generates dynamics that speak to the challenge put forward by Angeletos et al. (2020) of having models of demand-driven business cycle that are consistent with stable inflation.

Second, a key important effect following an aggregate uncertainty shock is an ex-post *capital premium*, defined in Equation (21), which is persistently positive in this recession. The

[8]For this reason, in standard NK models, discount factor shocks, while typically leading to comovement between labor and consumption, do not simultaneously generate comovement with investment.

premium indicates that capital, an uncertain and illiquid asset, requires a higher equilibrium excess return compared to the risk-free and liquid asset.[9]

Decomposing the response of the capital premium. Figure 4 decomposes the sources underlying the predictability of a positive premium in response to the aggregate uncertainty shock. The solid lines in both panels plot the realized premium of the Figure 3 starting from the first period after the ambiguity shock. First, Panel (a) shows that over the first few quarters the predictably higher premium primarily comes from an increase in the capital return (dot-dashed blue line). The latter then stabilizes and the persistent fall in the real rate (dashed purple line) eventually accounts for the persistently higher capital premium. This decomposition, favored by the data in our Bayesian estimation, is further consistent with stylized facts documented in the asset pricing literature emphasizing not only that excess returns are predictable but that this predictability does not just reflect real rate movements (e.g. Cochrane (2011), Bianchi et al. (2018)).

Capital is both illiquid and uncertain. To understand the role of these two features in driving the premium response, recall that the IRF plots the premium as recovered by an econometrician belief, which measures realizations ex-post under the belief E_t. We can then leverage the linearity of the solution method to simply decompose the premium as

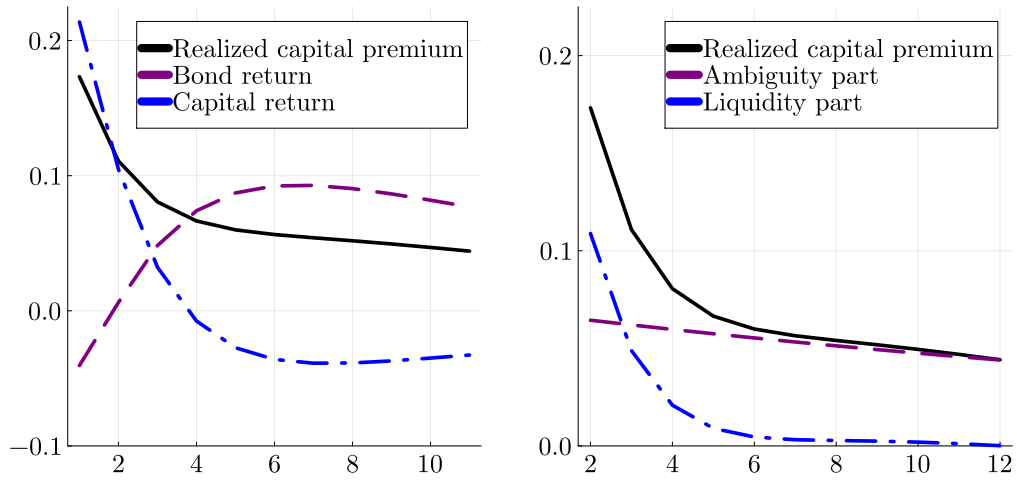
$$E_t Prem_{t+1} = \underbrace{E_t^* Prem_{t+1}}_{\substack{\text{liquidity} \\ \text{part}}} + \underbrace{E_t Prem_{t+1} - E_t^* Prem_{t+1}}_{\substack{\text{uncertainty} \\ \text{part}}}. \tag{31}$$

The liquidity part is the equilibrium compensation required to hold capital as an illiquid asset under the worst-case belief E_t^*, which is used in equilibrium in pricing assets. In the absence of a illiquidity friction, the expected premium under E_t^* would be zero in the impulse response, since the model is linearized. The uncertainty part is formally the result of the change of measure (i.e. a 'wedge') from the econometrician belief E_t, to the worst-case belief E_t^*. This part reflects the compensation in our linearized model for holding capital as an asset that is exposed to aggregate uncertainty. This ambiguity component would in turn be absent under Rational Expectations, as the econometrician and agents' worst-case belief would be assumed to coincide.[10]

Panel (b) of Figure 4 plots the decomposition of the premium in Equation (31) in response to the aggregate uncertainty shock into the liquidity (the dot-dashed blue line) and uncertainty part (the dashed purple line). The model implies that in the short-run the main

[9]The premium is on impact negative because of the surprise embedded in the ambiguity shock, which lowers dividends and the price of capital. After impact, the the premium is systematically positive.

[10]See Ilut and Schneider (2014) and Bianchi et al. (2018) for details of this argument in the context of representative agent models. For a model with liquidity and ambiguity premia see also Ilut et al. (2022).



(a) Return decomposition (b) Mechanism decomposition

Y-axis: Percentage points (quarterly). X-axis: Quarters.

Figure 4: Decomposing the response of capital premium following an increase in ambiguity

component is the compensation for trading frictions. Suddenly faced with higher aggregate uncertainty, the capital owners look to aggressively sell capital and shift away from its illiquidity property. This effect therefore arises from a strong *interaction* between the HANK friction and uncertainty. After a few quarters, this frictional component subsides and the capital premium becomes mainly a reflection of the compensation for aggregate uncertainty.

Comparing the response to ambiguity in HANK vs RANK. We can further evaluate the role of HANK frictions through the comparison in Figure 3 between our baseline HANK model and the counterfactual RANK. The key mechanism in this comparison has to do with the stronger incentive to move away from capital by the marginal investor in HANK compared to the RANK model. We have seen this mechanism at work already in Section 4 where we report that the *average level* of the equilibrium premium as a compensation for aggregate uncertainty is much larger than in the counterfactual RANK. Here, the same force of a differential marginal investor can explain the HANK model's stronger response to a *change* in uncertainty and the implied larger variability in premia and investment.

In particular, faced with more aggregate uncertainty about TFP, the rich agents in the HANK model, owning most capital and thus driving most of the investment, look to sell capital and shift their portfolio more towards the liquid asset. This shift and the higher demand for the liquid government debt is met in equilibrium by the increase in the supply of debt following the countercyclical government debt and fiscal transfers. Instead, the counterfactual representative agent worries relatively more about labor income, a larger

share of her future income in that case. She thus experiences a stronger precautionary savings demand which gets channeled in the RANK model more towards investment in capital. Thus, compared to its counterfactual RANK version, an increase in ambiguity interacts with the illiquidity friction to lower significantly more investment and the price of capital, leading to a capital premium that is larger and more persistent.

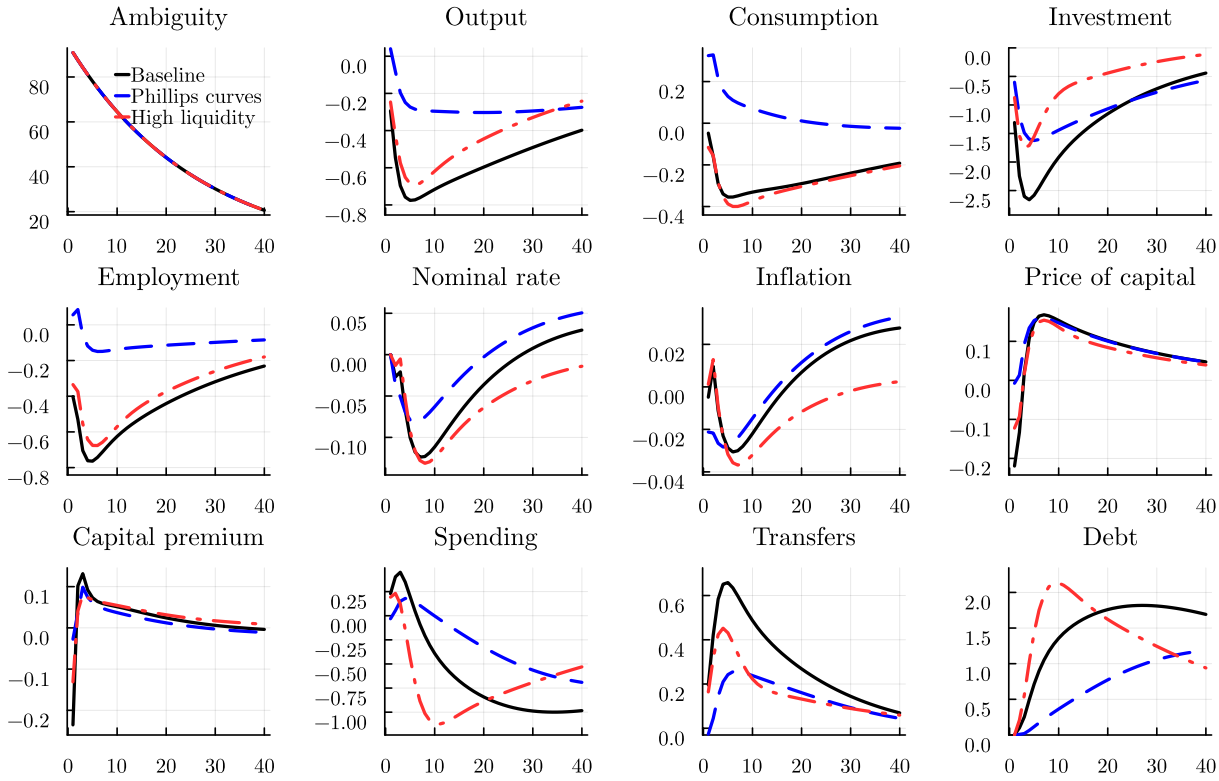
Beyond the IRF, a similar point can be made by analyzing the theoretical variance decomposition at business cycle frequency of the counterfactual RANK, where ambiguity shock accounts for about 50% of the model-implied variation in investment, compared to over 90% in HANK (recall panel (b) in Figure 2). The implications for the premium are also significantly different in HANK vs RANK. In Figure 3 the premium is essentially not moving in the RANK version. As shown in panel (f) of Figure 2, the ambiguity shock is the main model driver of premium in HANK, accounting for about 70% of the model-implied total variation. In contrast, in RANK this share is less than 10%.

Turning to consumption dynamics, these are also different across the two model versions. First, in the HANK model consumption falls by less on impact. This occurs for two reasons. On the one hand, as discussed earlier, there is relatively less precautionary savings demand than in the RANK model. On the other hand, the estimated fiscal policy is characterized by countercyclical lump-sum transfers. While these transfers have no effect in RANK due to its Ricardian equivalence nature, they help prop up consumption in the otherwise deep recession of the HANK model. Second, the consumption dynamic path features a hump-shape in our baseline. This stands in contrast to the monotonic mean-reversion from below in the RANK model, typical in models that lack habit-formation in consumption, like in ours. The hump shape in our baseline reflects the short-lived support from the countercyclical fiscal transfers in not letting consumption fall much on impact.

Counterfactual response to ambiguity: less illiquidity friction. To further diagnose mechanisms we now report results from a series of counterfactual experiments. Figure 5 plots in the red dot-dashed line a version where, keeping all the other parameters fixed as in the baseline, we weaken the illiquidity friction, by increasing the probability of trading the illiquid asset to $\lambda = 25\%$. Through this weakening of the trading friction, the resulting counterfactual model starts to resemble a one-asset HANK model.

We see three important effects in this counterfactual case compared to the baseline. First, consumption dynamics are very similar, indicating that the incomplete risk-sharing property of the model matters much more for consumption than the illiquidity friction. Second, investment falls by about 40% less than it does in the baseline. Third, the price of capital falls similarly by less and the premium is less volatile. Both of these latter effects confirm the key interaction between ambiguity and the illiquidity friction characterizing our

model's mechanism. When capital is less illiquid, its owners feel less of an urgency to shift away from it when aggregate uncertainty increases. As a result, the fall in investment and price of capital is significantly less dramatic than in the baseline.



Y-axis: Percentage points for inflation, nominal rate and premium, otherwise percent. X-axis: Quarters. Baseline refers to the estimated HANK model, Phillips curves is a counterfactual without ambiguity in the price and wage Phillips curves, High liquidity is counterfactual with trading probability of 25%.

Figure 5: Model counterfactuals

Counterfactual response to ambiguity: no effects through the Phillips curves. We can further diagnose mechanisms through counterfactuals where some decisions do not react to ambiguity. In particular, we can consider model versions where in steady-state all agents use the same worst-case belief but away from it, some decision-makers may not respond to time-varying ambiguity.

For example, an important property of our model discussed for the IRF in Figure 3 is that aggregate uncertainty matters for price setting through the expected inflation formed under the worst-case belief in the Phillips Curve of Equation (18). We can turn that effect off by leveraging the linearity of our solution method since

$$\mathbb{E}_t \widehat{\pi}_{t+1} = \mathbb{E}_t^* \widehat{\pi}_{t+1} + \varepsilon_{\pi z} a_t. \tag{32}$$

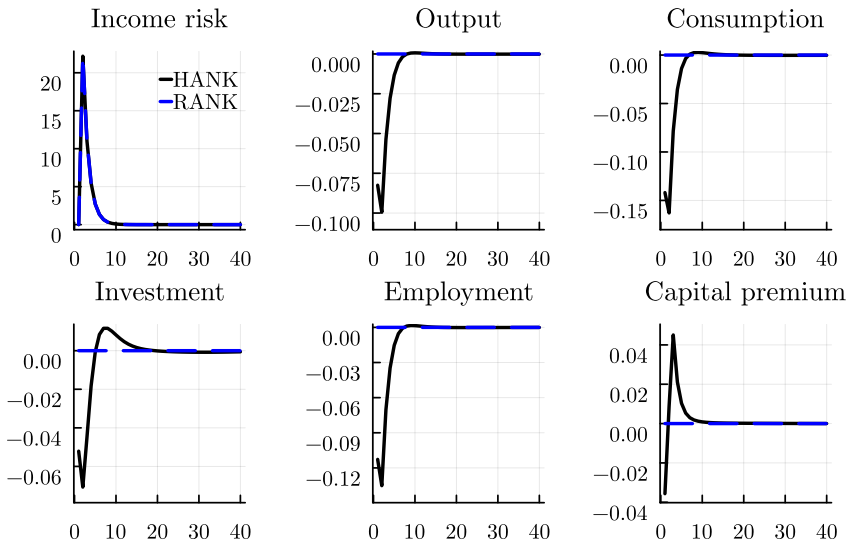
Computing the conditional expected inflation under the econometrician's belief means undoing the effect of the current worst-case belief about future TFP ($\mu_t^* = -a_t$) over future inflation, which occurs through $\varepsilon_{\pi z}$, the original equilibrium elasticity of inflation with respect to TFP. We can then compute a counterfactual economy where all forward-looking decisions are done under the worst-case belief except price-setting, where the expected inflation in Equation (18) is now given by $\mathbb{E}_t\widehat{\pi}_{t+1}$. A similar approach as in Equation (32) can turn off the effects of ambiguity on the nominal wage setting.

The blue dashed line of Figure 5 plots a counterfactual where we use this approach to turn off the effects of ambiguity in the Phillips curves for both price and wage setting. The key effect is that now the recession caused by the ambiguity increase is milder. The reason is that in contrast to the baseline version, in this counterfactual firms and unions now do not exhibit precautionary price-setting as they do not worry about future marginal costs being high. Therefore compared to the baseline, they set lower goods prices and nominal wages, leading to relatively higher demand for goods and employment. Thus, output, employment, investment and price of capital fall significantly less than the baseline. In fact, consumption even rises, still stimulated by the countercyclical fiscal transfers. Notably, inflation in this counterfactual is similar to the baseline despite the recession being much milder. Put differently, in our baseline model we obtain a deep recession without a correspondingly major deflation, since there firms do worry about high future marginal costs.

Impulse response to the idiosyncratic risk shock. We now discuss more briefly the impulse responses for the rest of the shocks in our baseline model. In particular, another source of time-varying uncertainty is the idiosyncratic income risk, i.e. an innovation ϵ_t^σ to the conditional volatility of labor income in Equation (9). Figure 6 plots the IRF to an increase in this risk. While in the counterfactual RANK model this shock would have clearly no effects, in the HANK model the increase in risk leads to a fall in consumption, labor and investment. However, these effects are short-lived and moreover aggregate investment overshoots soon after impact, and slowly returns to steady state from above.[11] Intuitively, this shock acts as a precautionary-savings inducing disturbance, leading on impact to a reduction in consumption, and through nominal rigidities to a demand-driven recession with lower employment and aggregate investment. In this respect, it is also similar to the precautionary saving property of the aggregate uncertainty increase, per the discussion around Figure 3.

A key contrast to the aggregate uncertainty increase is that the latter also implies a reduction on the uncertainty-adjusted return to capital, while the idiosyncratic uncertainty operates entirely through worries over labor income. That reduction in the perceived re-

[11]This type of down-and-up dynamic also resembles the IRF characterizing the rational expectations HANK model in Bayer et al. (2024a).



Y-axis: Percentage points for the premium, otherwise percent. X-axis: Quarters.

Figure 6: Impulse responses to an idiosyncratic income risk shock in estimated HANK model and counterfactual RANK (under the same parameters)

turn on capital pushes down significantly and persistently the desire to investment in the uncertain capital, a force that is absent here. Altogether, in contrast to the response to aggregate uncertainty, the short-lived recessionary effects and over-shooting response to the idiosyncratic income risk shock does not make it a promising source of systematic business cycle fluctuations. This is reflected in the variance decomposition of Figure 2, where shocks to idiosyncratic income risk play a small role.

The other shocks. We conclude the discussion of the model dynamics with a brief comment on the IRFs to the remaining shocks. These responses are consistent with standard findings in typical estimated NK models, so for brevity we relegate them to Appendix B.

Aggregate TFP shocks are not a sufficiently promising source of business cycles, for the standard reason of failing to generate in a quantitative NK model positive comovement between consumption, investment and hours. In particular, employment falls following a positive aggregate TFP shock (see Figure 9 in Appendix). The reason, in contrast to a typical response characterizing its RBC version, is standard in this class of models – it appears due to nominal rigidities, as price and wage markups become endogenously countercyclical.[12]

Finally, consider the responses to nominal shocks. A contractionary monetary policy shock induces a higher real rate, dampening demand for consumption and investment and

[12]The negative effects on hours appears in typical estimated NK models (like Smets and Wouters (2007), Justiniano et al. (2010)), including those that allow for incomplete markets (e.g. Bayer et al. (2024a)) or ambiguity with a representative agent (Ilut and Schneider (2014)).

leading to a relatively short recession with lower employment and persistently low inflation (see Figure 10 in the Appendix). Lastly, an increase in the inflation target lowers the real rate and produces a boom but one that is accompanied by a large and persistent increase in inflation (see Figure 11 in the Appendix). Quantitatively, as indicated earlier in the historical decomposition, the main role played by these nominal shocks is to improve the empirical fit of the model on the nominal side.

References

Angeletos, G.-M., Collard, F., and Dellas, H. (2020). Business-cycle anatomy. *American Economic Review*, 110(10):3030–70.

Auclert, A., Rognlie, M., and Straub, L. (2024a). Fiscal and monetary policy with heterogeneous agents. *NBER Working Paper*, 32991.

Auclert, A., Rognlie, M., Straub, L., and Tapak, T. (2024b). When do endogenous portfolios matter for hank? Unpublished manuscript.

Barro, R. J. and King, R. G. (1984). Time-separable preferences and intertemporal-substitution models of business cycles. *The Quarterly Journal of Economics*, 99(4):817–839.

Basu, S. and Bundick, B. (2017). Uncertainty shocks in a model of effective demand. *Econometrica*, 85(3):937–958.

Bayer, C., Born, B., and Luetticke, R. (2024a). Shocks, frictions, and inequality in us business cycles. *American Economic Review*, 114(5):1211–47.

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

Bayer, C., Luetticke, R., Weiss, M., and Winkelmann, Y. (2024b). An endogenous gridpoint method for distributional dynamics. CEPR Discussion Papers 19067.

Berger, D., Bocola, L., and Dovis, A. (2023). Imperfect risk sharing and the business cycle. *Quarterly Journal of Economics*, 138(3):1765–1815.

Bhandari, A., Borovička, J., and Ho, P. (2024). Survey data and subjective beliefs in business cycle models. *Review of Economic Studies*.

Bhandari, A., Bourany, T., Evans, D., and Golosov, M. (2023). A perturbational approach for approximating heterogeneous agent models. *NBER Working Paper*, 31744.

Bianchi, F., Ilut, C., and Schneider, M. (2018). Uncertainty shocks, asset supply and pricing over the business cycle. *The Review of Economic Studies*, 85:810–854.

Bilbiie, F., Primiceri, G. E., and Tambalotti, A. (2023). Inequality and business cycles. *NBER Working Paper*, 31729.

Bocola, L. and Lorenzoni, G. (2023). Risk-sharing externalities. *Journal of Political Economy*, 131(3):595–632.

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

Brunnermeier, M. K. and Sannikov, Y. (2014). A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421.

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

Chang, M., Chen, X., and Schorfheide, F. (2021). Heterogeneity and aggregate fluctuations. *NBER Working Paper*, 28853.

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

Cochrane, J. H. (2011). Presidential address: Discount rates. *The Journal of finance*, 66(4):1047–1108.

Debortoli, D. and Galí, J. (2017). Monetary policy with heterogeneous agents: Insights from tank models. Unpublished manuscript.

Ellsberg, D. (1961). Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics*, pages 643–669.

Epstein, L. G., Farhi, E., and Strzalecki, T. (2014). How much would you pay to resolve long-run risk? *American Economic Review*, 104(9):2680–97.

Epstein, L. G. and Schneider, M. (2003). Recursive multiple-priors. *Journal of Economic Theory*, 113(1):1–31.

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

Fernandez-Villaverde, J. and Guerron-Quintana, P. (2020). Uncertainty Shocks and Business Cycle Research. *Review of Economics Dynamics*, 37:S118–S1465.

Fernández-Villaverde, J., Hurtado, S., and Nuno, G. (2023). Financial frictions and the wealth distribution. *Econometrica*, 91(3):869–901.

Fernández-Villaverde, J., Marbet, J., Nuño, G., and Rachedi, O. (2024). Inequality and the zero lower bound. *Journal of Econometrics*.

Gertler, M. (2010). Financial intermediation and credit policy in business cycle analysis. *Handbook of Monetary Economics*.

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

Gilboa, I. and Schmeidler, D. (1989). Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153.

Gomez, M. (2024). Wealth Inequality and Asset Prices. *Columbia University working paper*.

Gomme, P., Ravikumar, B., and Rupert, P. (2011). The return to capital and the business cycle. *Review of Economic Dynamics*, 14(2):262–278.

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

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

Ilut, C. and Schneider, M. (2014). Ambiguous business cycles. *The American Economic Review*, 104(8):2368–2399.

Ilut, C. L., Krivenko, P., and Schneider, M. (2022). Uncertainty or frictions? a quantitative model of scarce safe assets. *NBER Working Paper*, 32198.

Ilut, C. L. and Schneider, M. (2022). Modeling uncertainty as ambiguity: A review. *NBER Working Paper*, 29915.

Jermann, U. and Quadrini, V. (2012). Macroeconomic effects of financial shocks. *American Economic Review*, 102(1):238–271.

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

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

Justiniano, A., Primiceri, G. E., and Tambalotti, A. (2013). Is there a trade-off between inflation and output stabilization? *American Economic Journal: Macroeconomics*, 5(2):1–31.

Kacperczyk, M., Nosal, J., and Stevens, L. (2019). Investor sophistication and capital income inequality. *Journal of Monetary Economics*, 107:18–31.

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

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

Kase, H., Melosi, L., and Rottner, M. (2022). Estimating nonlinear heterogeneous agents models with neural networks. Unpublished manuscript.

Kekre, R. and Lenel, M. (2022). Monetary policy, redistribution, and risk premia. *Econometrica*, 90(5):2249–2282.

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

Krishnamurthy, A. and Vissing-Jorgensen, A. (2012). The aggregate demand for Treasury debt. *Journal of Political Economy*, 120(2):233–267.

Leduc, S. and Liu, Z. (2016). Uncertainty shocks are aggregate demand shocks. *Journal of Monetary Economics*, 82:20–35.

Nakajima, T. (2005). A business cycle model with variable capacity utilization and demand disturbances. *European Economic Review*, 49(5):1331–1360.

Panageas, S. et al. (2020). The implications of heterogeneity and inequality for asset pricing. *Foundations and Trends® in Finance*, 12(3):199–275.

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

Schaab, A. (2020). Micro and macro uncertainty. Unpublished manuscript.

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

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

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

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

Uhlig, H. (2001). A toolkit for analysing nonlinear dynamic stochastic models easily. In Marimon, R. and Scott, A., editors, *Computational Methods for the Study of Dynamic Economies*, chapter 3, pages 30–61. Oxford University Press.

Werning, I. (2015). Incomplete markets and aggregate demand. *NBER Working Paper*, 21448.

World Inequality Database (2023). Top Income and Wealth Shares. Online Data Base. `https://wid.world/country/usa/` (accessed Feb 16, 2023).

Wu, J. C. and Xia, F. D. (2016). Measuring the macroeconomic impact of monetary policy at the zero lower bound: Dataset. *Journal of Money, Credit and Banking*, 48(2-3):253–291, `https://www.atlantafed.org/cqer/research/wu--xia--shadow--federal--funds--rate` (accessed: Feb 16, 2023).

Appendix

A Data and Estimation

A.1 Data: Sources and transformations

A.1.1 Data for moment-matching

The following list contains the data sources for the average data ratios we target in the calibration of the ergodic distribution. Unless otherwise noted, all series are available from 1985 to 2019 from the St.Louis FED - FRED database (mnemonics in parentheses).

Mean illiquid assets. Private fixed assets (K1PTOTL1ES000) over quarterly GDP (excluding net exports; see below), averaged over 1985 – 2019.

Mean government debt. Gross federal debt held by the public as percent of GDP (FYPUGDA188S), averaged over 1985 – 2019.

Average top 10 share of wealth. Source is the World Inequality Database (2023), averaged over 1985 – 2019.

A.1.2 Data for estimation

Formally, the vector of observable variables is given by:

$$OBS_t = \begin{bmatrix} \Delta \log(C_t) \\ \Delta \log(I_t) \\ \Delta \log(N_t) \\ \log\left(R_t^b\right) \\ \log(\pi_t) \\ \log(Prem_t) \end{bmatrix} - \begin{bmatrix} \overline{\Delta \log(C_t)} \\ \overline{\Delta \log(I_t)} \\ \overline{\Delta \log(N_t)} \\ \overline{\log\left(R_t^b\right)} \\ \overline{\log(\pi_t)} \\ 0.0 \end{bmatrix}$$

where Δ denotes the temporal difference operator and bars above variables denote time-series averages.

Unless otherwise noted, all series are available at quarterly frequency from 1985Q1 to 2019Q4 from the St.Louis FED - FRED database (mnemonics in parentheses).

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

Investment, I_t. Gross private domestic investment (GPDI) divided by the GDP deflator (GDPDEF) and the civilian noninstitutional population (CNP16OV).

Hours worked, N_t. Nonfarm business hours worked (HOANBS) divided by the civilian noninstitutional population (CNP16OV).

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

Nominal interest rate, R_t^b. Quarterly average of the effective federal funds rate (FEDFUNDS). From 2009Q1 to 2015Q4, we use the Wu and Xia (2016) shadow federal funds rate.

Capital Premium, $PREM_t$. We take the estimated time series for after-tax returns to all capital from Gomme et al. (2011) and substract the real yield on long-term U.S. government securities (LTGOVTBD) until June 2000 and 20-Year Treasury Constant Maturity Rate (GS20) afterwards (see Krishnamurthy and Vissing-Jorgensen, 2012). Available from 1985Q1 to 2019Q4.

A.2 MCMC diagnostics

We estimate the model using a single RWMH chain after an extensive mode search. After burn-in, 150,000 draws from the posterior distribution are used to compute the posterior statistics. The acceptance rate is close to 30%. We check Geweke (1992) convergence statistics for individual parameters as well as traceplots. Geweke (1992) tests the equality of means of the first 10% of draws and the last 50% of draws (after burn-in). If the samples are drawn from the stationary distribution of the chain, the two means are equal and Geweke's statistic has an asymptotically standard normal distribution. Taking the evidence from Geweke (1992) and the traceplotss together, we conclude that our RWMH chain has converged. There is still some movement in the parameter for the standard deviation of inflation, but only in the fourth decimal place. Otherwise no individual Geweke test rejects at the one percent level, and only a small number reject at the five percent level, which can be expected from the multiple-testing nature of the exercise.

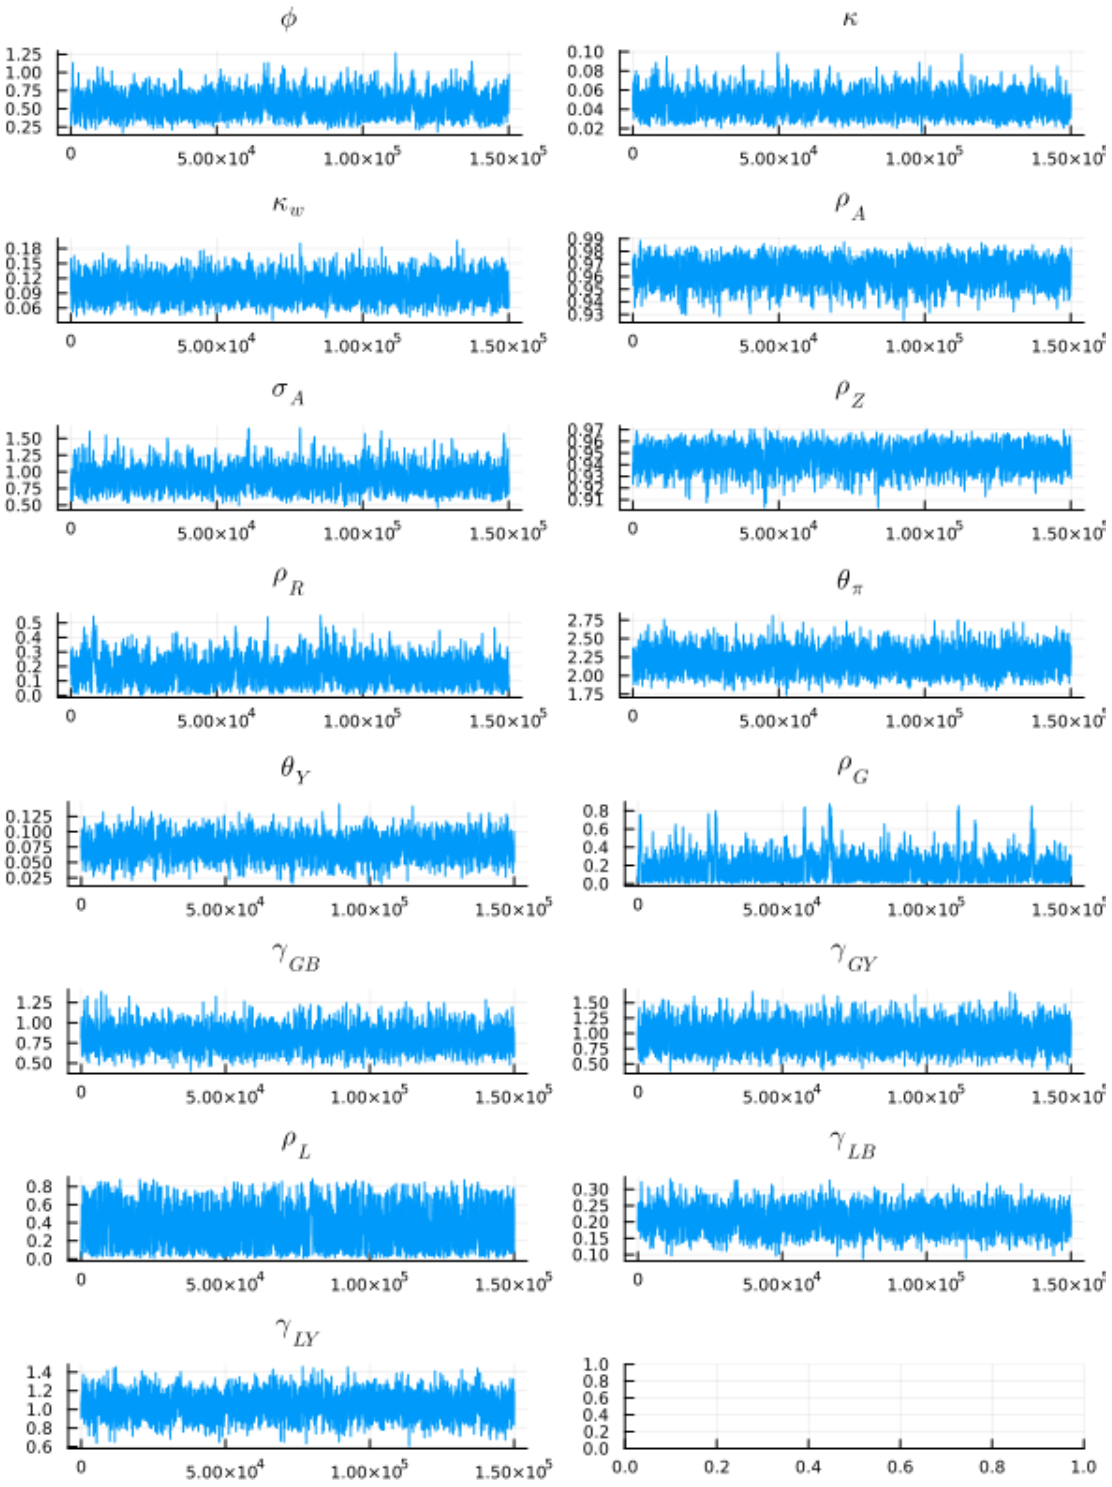


Figure 7: MCMC draws of HANK model

Table 4: Geweke convergence test results

Frictions and Ambiguity			Shocks		
Parameter	z-stat	p-value	Parameter	z-stat	p-value
δ_s	-1.045	0.296	ρ_Z	0.002	0.998
ϕ	0.602	0.547	ρ_A	0.369	0.712
κ	1.460	0.144	σ_A	0.177	0.860
κ_w	-0.309	0.757	ρ_S	1.988	0.047
$\tilde{z}$	-0.356	0.722	σ_S	-0.172	0.864
Monetary Policy			**Fiscal Policy**		
ρ_R	1.138	0.255	ρ_G	0.988	0.323
θ_π	0.094	0.925	γ_{GB}	1.879	0.060
θ_Y	-0.951	0.342	γ_{GY}	1.084	0.279
ρ_R^ϵ	-0.534	0.593	ρ_L	1.806	0.071
σ_R^ϵ	-0.065	0.948	γ_{LB}	2.205	0.027
ρ_π^ϵ	-1.985	0.047	γ_{LY}	1.150	0.250
σ_π^ϵ	8.970	0.000			

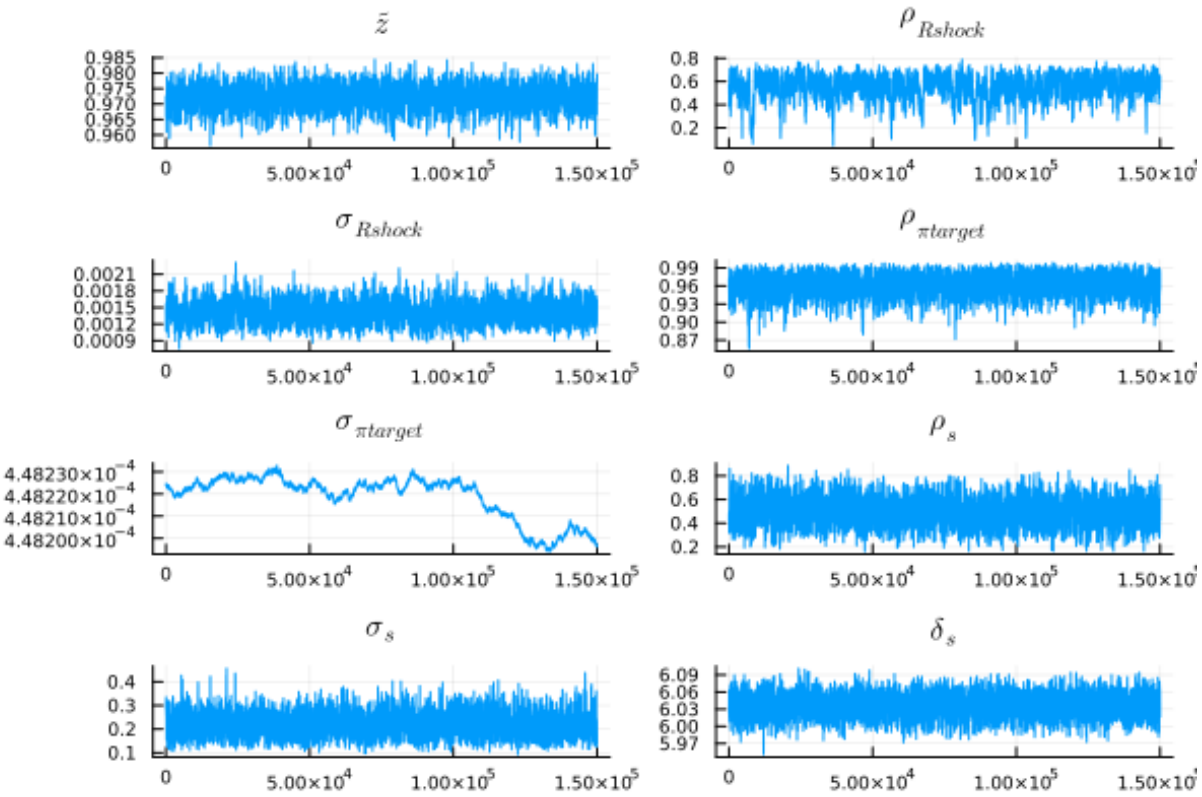
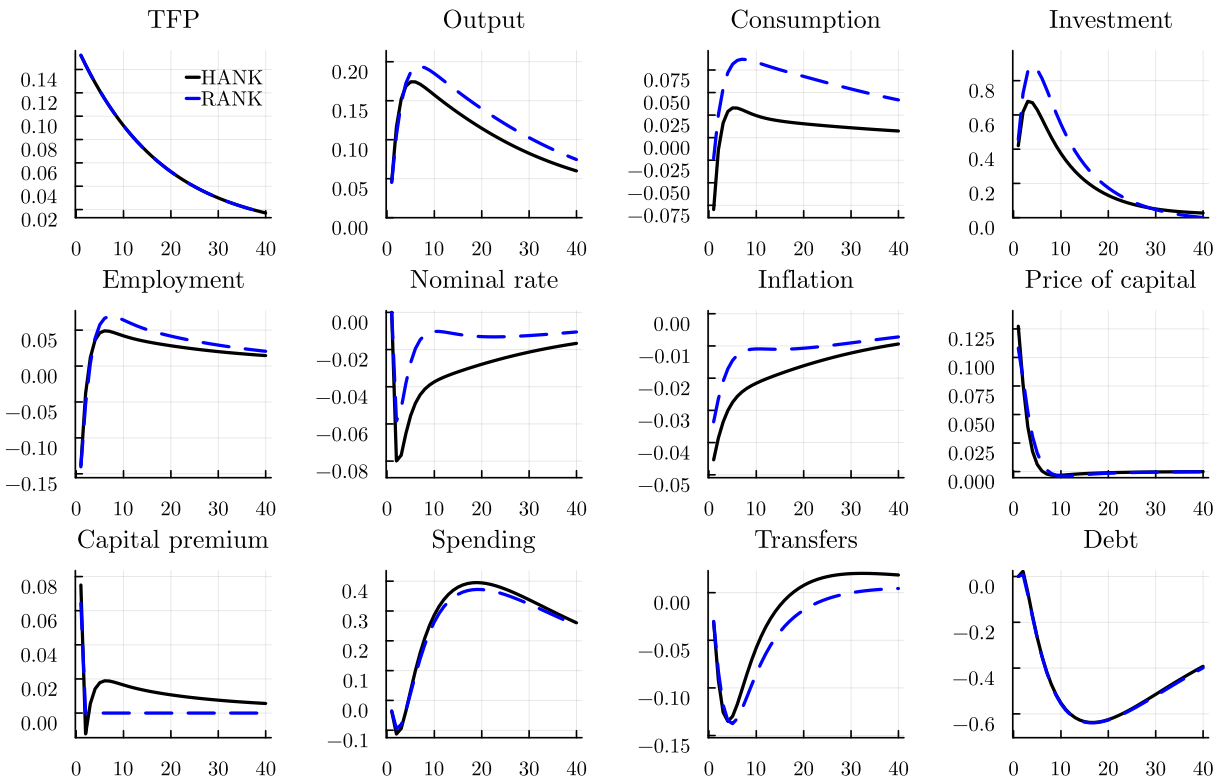


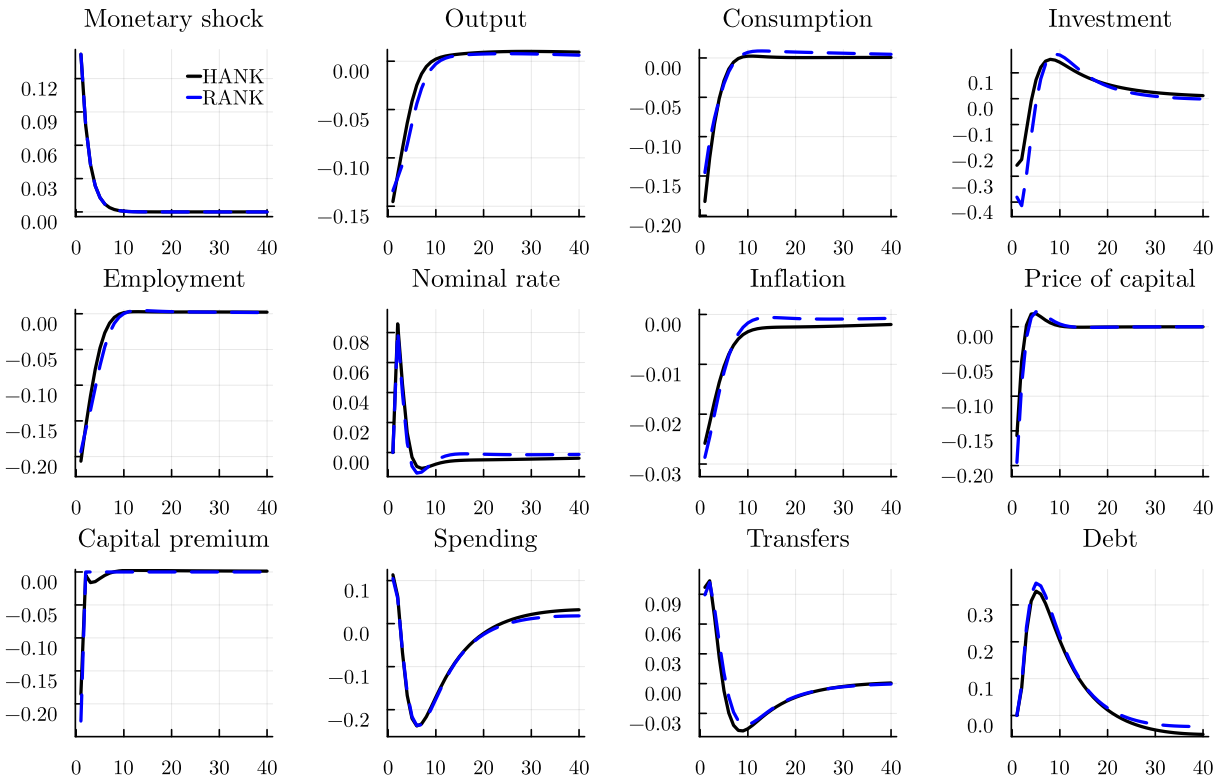
Figure 8: MCMC draws of HANK model

B Supplementary Figures



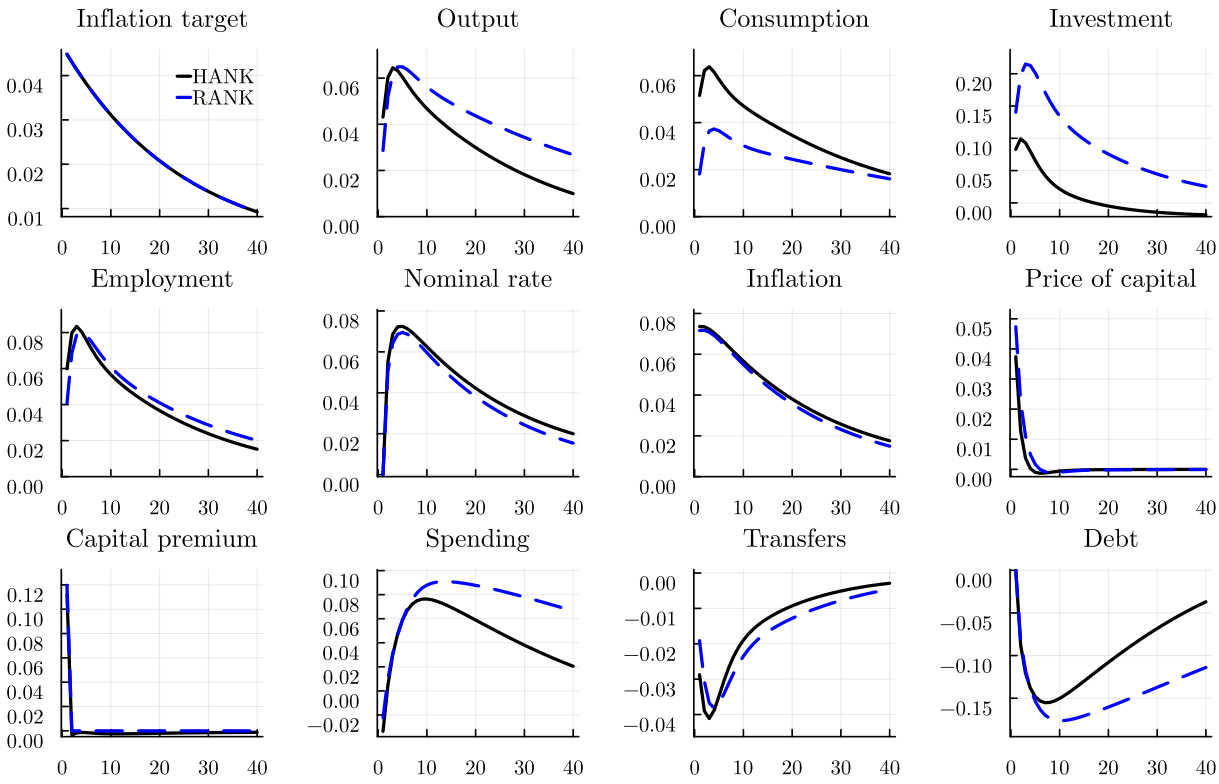
Impulse Responses to TFP shock in estimated HANK model and counterfactual RANK (under the same parameters). Y-axis: Percentage points for inflation, nominal rate and premium, otherwise percent. X-axis: Quarters.

Figure 9: Impulse responses to TFP



Impulse Responses to a monetary shock in estimated HANK model and counterfactual RANK (under the same parameters). Y-axis: Percentage points for inflation, nominal rate and premium, otherwise percent. X-axis: Quarters.

Figure 10: Impulse responses to monetary policy shocks



Impulse Responses to an inflation target shock in estimated HANK model and counterfactual RANK (under the same parameters). Y-axis: Percentage points for inflation, nominal rate and premium, otherwise percent. X-axis: Quarters.

Figure 11: Impulse responses to inflation target shocks