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

Cosmin Ilut[†], Ralph Luetticke[‡] and Martin Schneider[§]

December 24, 2024

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 both the steady state and the linearized dynamics, allowing for fast computation and estimation with linear methods. The estimated model jointly fits cyclical variation in US macro aggregates and asset prices. In the presence of portfolio frictions, aggregate uncertainty shocks are a powerful driver of the business cycle, more so than idiosyncratic uncertainty shocks. Their effect is much stronger than in a representative agent model: portfolio substitution by rich capital owners helps fit investment and return dynamics.

JEL codes: C11, D31, E32, E63

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

*We thank seminar and conference participants at Bundesbank, ECB, Lancaster ETM Workshop, NBER Summer Institute, NYU, 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. Thus, 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. There is also a first-order effect of aggregate uncertainty on intertemporal choices by firms. All intertemporal decisions – savings and portfolio choice by households as well as price and wage setting by firms – are made more cautiously after an aggregate uncertainty shock.

Our main quantitative result is that aggregate uncertainty shocks generate powerful comovement of macroeconomic aggregates and asset prices over the business cycle. In our baseline estimation, a single shock to ambiguity about TFP jointly explains more than 70% of cyclical variation in key macroeconomic aggregates as well as in the excess return on capital and the real interest rate. 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. We also infer an important effect of aggregate uncertainty on asset prices: the average excess return on capital of 5.5% reflects a sizable uncertainty premium of 3.2% and a smaller illiquidity premium of 2.3%.

[1] See the review by Kaplan and Violante (2018) or more recently Auclert et al. (2024).

Our results are driven by the *interaction* of aggregate uncertainty and the portfolio frictions of the two-asset HANK model. Higher uncertainty about TFP generates a flight to safety: rich households who own most capital substitute away from capital towards bonds, while poor households want bonds for precautionary savings. The resulting decline in investment and increase in the capital premium is much stronger than in a representative agent model or even a one-asset HANK model. Indeed, when we shut down heterogeneity or liquidity frictions in our model, the effect of aggregate uncertainty on investment as well as the capital premium essentially disappear. Intuitively, when all households own capital, the precautionary motive tends to stabilize capital demand after an uncertainty shock. Similarly, higher labor income volatility generates relatively more precautionary savings and less of a flight away from capital; this makes volatility less suitable as drivers 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 enters linearly, we can selectively shut off one or more of the correlated wedges it introduces into model equations. 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. (2024). 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 because of 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 total factor productivity (TFP) and to the volatility of labor productivity. We also allow for shocks to monetary policy and 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 similar.[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.

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. 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 lower than actual mean TFP. An uncertainty shock corresponds to an even wider interval and hence a lower worst case mean. 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.

Our estimation strategy is thus designed to identify parameters jointly from long run

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

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 we 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 25% of variation in investment 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 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 magnitude of shocks 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.

Our paper contributes to a recent literature on estimating HANK models, which requires in particular solving the models sufficiently quickly. Auclert et al. (2021) solve HANK models in sequence space and use an MA-∞ representation for estimating their dynamics. Bayer et al. (2024) propose full-information Bayesian techniques for estimation using the models' state-space representation. Both methods assume that agents' actions do not respond to aggregate uncertainty – instead future movements in aggregate variables are treated as certainty-equivalent. This is what allows linear representation of the dynamics with standard risk preferences. Our approach extends Bayer et al. (2024), showing that their optimal dimensionality reduction and speed can also handle aggregate uncertainty.

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 highlighted the role of aggregate uncertainty as a demand shock in New Keynesian models with a representative agent (for example, 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 as multiple priors. 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.[3]

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

[3]The result is consistent with Berger et al. (2023), who show that labor and discount factor wedges due to HANK frictions do not contribute much to aggregate fluctuations. Chang et al. (2021) develop a complementary empirical approach and find that adding cross-sectional data to a VAR model does not affect the decomposition of US business cycles. Bayer et al. (2024) also find this in a structural HANK model.

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)).

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. Section 4 presents quantitative results for our baseline HANK model and compares them to a RANK model.

2 Model

Our setup shares technology and asset market frictions as well as the modeling of household heterogeneity with the two asset HANK environment in Bayer et al. (2024). However, we replace expected utility preferences by recursive multiple priors utility and add aggregate uncertainty shocks. In this section, we first lay out the familiar parts of the environment and then explain how ambiguity aversion affects the objectives of households and firms.

2.1 Technology

There are four sectors of production in this economy, for final goods, intermediate goods, capital goods and labor services. We now introduce their technologies.

Final goods

Output of a final good Y_t is made by competitive firms combining a continuum of intermediate goods $Y_{j,t}$ according to the Dixit-Stiglitz aggregator

$$Y_t = \left(\int Y_{jt}^{\frac{\eta-1}{\eta}} dj \right)^{\frac{\eta}{\eta-1}} ,$$

with elasticity of substitution η. Each of these differentiated goods is offered at price p_{jt}, so profit maximization and the zero profit condition imply a demand for intermediate good j

$$Y_{jt} = \left(\frac{P_{jt}}{P_t} \right)^{-\eta} Y_t; \qquad P_t = \left[\int_0^1 P_{j,t}^{1-\eta} dj \right]^{\frac{1}{1-\eta}} \tag{1}$$

where P_t the aggregate price level.

[4] Panageas (2020) provides an overview of this mechanism, and how it emerges in models of heterogeneous beliefs, attitude towards uncertainty or access to financial markets.

Intermediate goods

Each intermediate good is made using capital services $u_{jt}K_{jt}$ and labor N_{jt} with a constant returns to scale production function

$$Y_{jt} = Z_t N_{jt}^{\alpha} (u_{jt} K_{jt})^{(1-\alpha)} \tag{2}$$

Uncertainty enters via total factor productivity shocks Z_t, described in detail later. Here u_{jt} is the intensity with which the capital stock K_{jt} is used. An intensity higher than normal results in increased depreciation of capital $\delta(u_{jt})$.[5]

The producer's per-period profit function is

$$\Pi_{j,t}^F = \left(\frac{P_{jt}}{P_t} - mc_{j,t} \right) y_{jt} \tag{3}$$

where mc_{jt} is the real marginal cost of firm j. Taking as given demand in equation (1), the producer minimizes costs, $w_t^F N_{jt} - [r_t + q_t \delta(u_{jt})] K_{jt}$, where r_t and q_t are the rental rate and the (producer) price of capital goods, respectively, and w_t^F is the real wage the firm faces. Factor markets are perfectly competitive. Hence, the first-order conditions for labor and effective capital read

$$w_t^F = \alpha mc_{jt} Z_t \left(\frac{u_{jt} K_{jt}}{N_{jt}} \right)^{1-\alpha} \tag{4}$$

$$\text{and} \qquad r_t + q_t \delta(u_{jt}) = u_{jt}(1-\alpha) mc_{jt} Z_t \left(\frac{N_{jt}}{u_{jt} K_{jt}} \right)^{\alpha}, \tag{5}$$

where mc_{jt} is the real marginal cost of firm j. The optimal utilization is given by

$$q_t \delta'(u_{jt}) = (1-\alpha) mc_{jt} Z_t \left(\frac{N_{jt}}{u_{jt} K_{jt}} \right)^{\alpha}.$$

By combining these conditions, marginal costs are constant across firms, $mc_{jt} = mc_t$.

The technology to adjust nominal prices P_{jt} is subject to a standard Calvo (1983) friction with indexation - nominal prices are indexed to the steady-state level of inflation, where the latter is defined as $\pi_t = P_t / P_{t-1}$, and can be discretionally adjusted with probability $1 - \lambda_Y$.

Capital goods

Perfectly competitive capital goods producers take the relative price of capital goods, q_t,

[5] In particular, the depreciation takes the form $\delta(u_{jt}) = \delta_0 + \delta_1 (u_{jt} - 1) + \delta_2/2 (u_{jt} - 1)^2$, which, assuming $\delta_1, \delta_2 > 0$, is an increasing and convex function of utilization. Without loss of generality, capital utilization in steady state is normalized to 1, so that δ_0 denotes the steady-state depreciation rate of capital goods.

as given and choose investment I_t. Their per-period profit function is

$$\Pi_t^Q = I_t \left\{ q_t \left[1 - \frac{\phi}{2} \left(\log \frac{I_t}{I_{t-1}} \right)^2 \right] - 1 \right\} \tag{6}$$

where ϕ captures the adjustment costs, which equal to 0 in the steady state. Since capital goods producers are symmetric, for a given path of I_t aggregate capital evolves as:

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

Labor services

Workers sell their labor services to a mass-one continuum of unions indexed by j, each of whom offers a different variety of labor to labor packers who then provide labor services to intermediate goods producers. Labor packers produce final labor services according to

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

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

$$\hat{n}_{jt} = \left(\frac{W_{jt}}{W_t^F} \right)^{-\zeta} N_t, \qquad W_t^F = \left[\int_0^1 W_{jt}^{1-\zeta} dj \right]^{\frac{1}{1-\zeta}} \tag{7}$$

where W_{jt} is the nominal wage set by union j and W_t^F is the nominal wage at which labor packers sell labor services to intermediate goods producers.

Given the demand curve in equation (7) and the real wage w_t at which they buy labor from households, the per-period profit function of a union j is

$$\Pi_{j,t}^U = \left(\frac{W_{jt}}{P_t} - w_t \right) \hat{n}_{jt} \tag{8}$$

Similar to price setting, the technology to adjust nominal wages W_{jt} is subject to a Calvo (1983) friction - nominal wages are indexed to the steady-state level of wage inflation, defined as $\pi_t^W = W_t^F / W_{t-1}^F$, and can be discretionally adjusted with probability $1 - \lambda_W$.

Due to adjustments costs in investment, pricing and wage setting, the producers of intermediate and capital goods, as well as the unions, will face a dynamic maximization problem. In particular, as in standard New Keynesian models (eg. Smets and Wouters (2007), Christiano et al. (2005)), firms choose actions to maximize the present value of their corresponding profits. An important input in that valuation is the stochastic discount factor (SDF), which is the channel through which precautionary behavior might affect firms' optimal choices. We detail the SDF construction and its implications in section 2.4.

Technology shocks

The source of aggregate uncertainty in this economy is TFP, through the Z_t shocks in equation (2). We follow common practice in business cycle analysis by assuming that TFP follows a persistent AR(1) process and let the true data generating process (DGP) be

$$\log Z_{t+1} = (1-\rho_z)\log \bar{Z} + \rho_z \log Z_t + \epsilon^Z_{t+1} \tag{9}$$

where ϵ^Z_{t+1} is an iid normal sequence of innovations with mean zero and variance σ_z^2, and $\bar{Z}$ is the long-run mean. As we detail below, agents observe the state Z_t but, different than in a standard Rational Expectations (RE) model, they are not confident about the conditional mean of the distribution of the innovations in (9).[6]

2.2 Household preferences

In modeling the economic structure of the household sector, we first build on the incomplete markets environment proposed by Bayer et al. (2024). Our key modeling contribution is then to enrich this economic environment by letting households perceive uncertainty over aggregate TFP not only just as risk but also as ambiguity (Knightian uncertainty).

Multiple priors utility

Households have preferences over consumption goods and labor. At date t, household i obtains utility from a composite good

$$x_{it} = c_{it} - G(h_{it}, n_i^t)$$

where c_{it} is consumption, n_{it} is labor, G captures the disutility of labor and h_{it} is household i's idiosyncratic labor productivity shock. The quasilinear functional form eliminates wealth effects of labor supply, following Greenwood et al. (1988).[7]

[6]See Ilut and Schneider (2014) for a detailed statistical argument underlying such a possible lack of confidence. In particular, Ilut and Schneider (2014) generalizes the true DGP to non-stationary models where the innovation $\left(\log Z_{t+1} - (1-\rho_z)\log \bar{Z} - \rho_z \log Z_t\right)$ has a conditional mean μ_t^* that is deterministic and unknown to agents. To an observer who samples TFP data from that process, the sequence μ_t^* thus cannot be learned: even with a large amount of data, it is impossible to disentangle the parameter sequence μ^* and the shock sequence ϵ^Z.

[7]The assumption of GHH preferences is motivated by the fact that estimated DSGE models typically find small aggregate wealth effects in labor supply; see, e.g., Schmitt-Grohé and Uribe (2012); Born and Pfeifer

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 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

$$U_t^C = u(x_{it}) + \beta \min_{p \in \mathcal{P}_t} E_t^p \left[U_{t+1}^C \right], \tag{10}$$

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.

The multiple priors functional form (10) captures a strict preference for knowing probabilities, and is thus consistent with behavior exhibited in the Ellsberg (1961) experiments. The key feature is that the worst case belief endogenously varies with, the consumption plan C_i.[8] 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 functions u and G), 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.

(2014). Using GHH preferences is not important for introducing ambiguity - for that, we can alternatively use King et al. (1988) (KPR) preferences.

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

Belief sets and uncertainty shocks

Reflecting their lack of confidence in the conditional distribution of TFP in (9) as ambiguous, we let agents entertain a family of conditional distributions with different means. In particular, we now parameterize the one-step ahead belief sets $\mathcal{P}_t$ over TFP as an interval of means

$$\log Z_{t+1} = (1-\rho_z)\log \bar{Z} + \rho_z \log Z_t + \mu_t + \epsilon^Z_{t+1}; \quad \mu_t \in [-a_t, a_t] \tag{11}$$

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.

The stochastic process a_t captures agents' perceived ambiguity about TFP; it evolves as

$$\log a_t = (1-\rho_a)\log \bar{a} + \rho_a \log a_{t-1} + \epsilon^a_t \tag{12}$$

with long run mean $\bar{a} > 0$, persistence $0 < \rho_a < 1$, and $\epsilon^a_t \sim i.i.d\ N(0, \sigma_a)$.

We model here a_t as an exogenous persistent process, interpreted as the cumulative effect of news that affect confidence. In some periods, such information gathering leaves them relatively 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, various pieces of information might contradict each other, and agents are less confident about their forecast. Periods of low $a_t < \bar{a}$ represent unusually low ambiguity about future productivity, whereas $a_t > \bar{a}$ describes periods of high uncertainty.

As in a standard RE economy, in our model uncertainty over aggregate TFP matters for intertemporal decisions for two sets of economic agents. On the one hand, households' optimal choices of investment into bonds and into capital are made under uncertainty. On the other hand, due to the adjustments costs in investment, pricing and wage setting introduced in section 2.1, the firms' optimal choices become forward-looking and dynamic. Formally, those choices may be affected by uncertainty due to the stochastic discount factor used to evaluate the firms' future profits. Next we discuss how our assumed ambiguity, as a particular form of uncertainty, impacts both of these type of agents. Critically, these effects show up even in a log-linearized equilbrium.

2.3 Household heterogeneity and incomplete markets

We first turn attention to the household side.

Household heterogeneity

As in Bayer et al. (2024), the household sector is subdivided into two types of agents: workers and entrepreneurs. The transition between types is stochastic. Only workers supply labor and face idiosyncratic labor productivity risk. This productivity evolves as[9]

$$h_{it} = \begin{cases} \exp\left(\rho_h \log h_{it-1} + \epsilon_{it}^h\right) & \text{with probability } 1-\zeta \text{ if } h_{it-1} \neq 0, \\ 1 & \text{with probability } \iota \text{ if } h_{it-1} = 0, \\ 0 & \text{else.} \end{cases} \tag{13}$$

With probability ζ households become entrepreneurs ($h = 0$). Both workers and entrepreneurs rent out physical capital. Entrepreneurs earn all pure rents, except those of unions which are equally distributed across workers. With probability ι an entrepreneur returns to the labor force with median productivity.

Idiosyncratic shocks ϵ_{it}^h are normally distributed with mean zero and variance $\bar{\sigma}_{h,t}^2$. These variance shocks capture changes in the idiosyncratic risk faced by households. We let this this income risk follows a log-AR(1) process

$$\log \bar{\sigma}_{h,t+1}^2 = (1-\rho_h)\log \bar{\sigma}_h^2 + \rho_h \log \bar{\sigma}_{h,t}^2 + \epsilon_t^\sigma \tag{14}$$

where $\bar{\sigma}_h^2$ is the steady state income risk, and $\epsilon_t^\sigma \sim i.i.d\ N(0, \sigma_\sigma)$. Thus, that at a given time t households know that there is an aggregate change in the variance of shocks that drive their next period's idiosyncratic productivity. 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 $\bar{\sigma}_{h,t}^2$.

Budget constraint

All households self-insure against the income risks they face by saving in a liquid nominal asset (bonds) and a less liquid asset (capital). Given income, which we detail more in Appendix A, households optimize intertemporally subject to their budget constraint

$$c_{it} + b_{it+1} + q_t k_{it+1} = b_{it}\frac{R_{it}}{\pi_t} + (q_t + r_t)k_{it} + (1-\tau)(w_t h_{it} n_{it} + \mathbb{I}_{h_{it}=0}\Pi_t^F + \mathbb{I}_{h_{it}\neq 0}\Pi_t^U) + L_t \tag{15}$$

where τ is the tax rate and L_t are lump-sum transfers. The profits for firms $\Pi_t^F = (1-mc_t)Y_t$ go to entrepreneurs, and the profits of unions $\Pi_t^U = (w_t^F - w_t)N_t$ go to workers. Real liquid assets is b_{it}, while k_{it} is the amount of illiquid assets, q_t is the price of these assets and r_t is their rental rate. Here R_t is the nominal interest rate on liquid assets, while the borrowing rate, i.e. when $b_{it} < 0$, is higher by a constant $\bar{R}$. Borrowing constraints take the form of an exogenous debt limit $\underline{B}$ on bond holdings b_{it+1} and non-negativity on capital holdings k_{it+1}.

Households make their savings and portfolio choice between liquid bonds and illiquid

[9] We divide by the cross-sectional average, $\int h_{it} di$, so that average worker productivity is constant.

capital in light of a capital market friction that renders capital illiquid because participation in the capital market is random and i.i.d. In particular, only a fraction, λ, of households are selected to be able to adjust their capital holdings in a given period. Households that do not participate in the capital market ($k_{it+1} = k_{it}$) still obtain dividends and can adjust their liquid asset holdings. For the portfolio choice, an important factor is therefore the *capital premium*, defined as the ex-post real excess return between the illiquid and liquid asset as

$$Prem_t = \frac{r_t + q_t}{q_{t-1}} - \frac{R_t}{\pi_t} \tag{16}$$

Value functions

Since a household's saving decision—(b'_a, k') for the case of adjustment and (b'_n, k) for non-adjustment—will be some non-linear function of that household's wealth and productivity, inflation and all other prices will be functions of the joint distribution, Θ_t, of (b, k, h) in t. This makes Θ a state variable of the household's planning problem and this distribution evolves as a result of the economy's reaction to aggregate shocks. We can summarize all effects of aggregate state variables, including the distribution of wealth and income, by writing the dynamic planning problem with time-dependent continuation values. We also note in particular that the exogenous aggregate TFP shocks Z_t and time-varying ambiguity a_t are also part of what determines the time-dependency of these continuation values.

This leaves us with three functions that characterize the household's problem: value function V^A for the case where the household *adjusts* its capital holdings, the function V^N for the case in which it does *not* adjust, and the expected continuation value, $\mathbb{W}$, over both,

$$V_t^A(b, k, h) = \max_{b'_A, k'} \left\{ u[x(b, b'_A, k, k', h)] + \beta \min_{\mu \in [-a,a]} E_t^{\mu} \mathbb{W}_{t+1}(b'_A, k', h') \right\} \tag{17}$$

$$V_t^N(b, k, h) = \max_{b'_N} \left\{ u[x(b, b'_N, k, k, h)] + \beta \min_{\mu \in [-a,a]} E_t^{\mu} \mathbb{W}_{t+1}(b'_N, k, h') \right\} \tag{18}$$

Maximization is subject to the corresponding budget constraint in (15). The continuation value $\mathbb{W}$ averages over V^A and V^N, using the exogenous probability of adjusting or not:

$$\mathbb{W}_{t+1}(b', k', h') = \lambda V_{t+1}^A(b', k', h') + (1 - \lambda) V_{t+1}^N(b', k', h') \tag{19}$$

Worst-case belief

As in the generic multiple priors representation in (10), continuation utility is evaluated under a set of one-step ahead beliefs. Equation (11) parameterizes this set as an interval of means, $\mu \in [-a, a]$ over future Z', where the process for ambiguity, a, is given by (12).

Ambiguity is a state variable as agents internalize its evolution and effect on values.

Intuitively, a negative aggregate TFP innovation decreases the available resources and thus overall surplus in the economy. Individually, through equilibrium prices, this impacts negatively the available labor, capital and profit income sources of agents, to the extent that conditional on individual state variables, including whether the household is a worker or an entrepreneur, or whether she can adjust capital holdings or not, her current value function depends positively in equilibrium on the aggregate TFP state. It is then natural to conjecture that cautious agents worry by behaving *as if* innovations to future aggregate TFP are low. That is, we guess that the worst-case belief that supports the optimal choices in equations (17) and (18) after every history, given endogenous equilibrium prices, is that future aggregate TFP has the lowest conditional mean out of the set, i.e.

$$\mu_t^* = -a_t, \forall t \tag{20}$$

The worst-case conditional belief over future TFP is therefore:

$$E_t^* \log Z_{t+1} = (1 - \rho_z) \log \bar{Z} + \rho_z \log Z_t - a_t \tag{21}$$

where we denote the conditional expectation under the guessed worst case belief by E_t^*. If the guess is correct, then agents' equilibrium decision rules are the same as in a model with a fixed conditional belief given by E_t^*. We verify our guess by checking that the equilibrium value functions, i.e. implied by the optimal choices supported by that belief, are indeed increasing with the innovation to aggregate productivity.

Households also need to make forecasts about the cross-section distribution Θ_{t+1} over individual states (b', k', h'). Agents hold model consistent beliefs, by correctly understanding how this distributions evolves as

$$\begin{aligned} \Theta_{t+1}(b', k', h') = \lambda \int_{b'=b^*_{A,t}(b,k,h), k'=k^*_t(b,k,h)} \Phi_t(h, h') d\Theta_t(b, k, h) \\ + (1 - \lambda) \int_{b'=b^*_{N,t}(b,k,h), k'=k} \Phi_t(h, h') d\Theta_t(b, k, h) \end{aligned} \tag{22}$$

$\Phi(\cdot)$ is the exogenous transition probability for idiosyncratic productivity h in equation (13), and $b^*_{A/N,t}$ and k^*_t are the time-t optimal policies derived under the worst-case belief E_t^*.

Using standard insights in the literature, and in particular following Reiter (2009), the discretized version of equations (17), (18), (19) and (22) can then be viewed as a set of equations that pins down the dynamics of the value functions and optimal policy for each $b \times k \times h$ node as well as the transition of the mass of households at each of the nodes.

A model variant with a representative household and complete markets

We briefly introduce a model variant with the same structure of ambiguity but with complete markets, which we will use as a comparison in some of our quantification. There all households are homogeneous with equal and constant labor productivity $h_i = 1$ and equally obtain all profit incomes. We maintain the same multiple priors utility and primitive ambiguity over aggregate TFP. Since that was common to all agents in the baseline economy, it applies directly to the representative household in this model variant.

The planning problem can be described by the consumption Euler equation for bonds instead of the above mentioned set of equations. For an optimal consumption-savings policy,

$$u_c(x_t) = \beta E_t^* \frac{R_t}{\pi_{t+1}} u_c(x_{t+1})$$

needs to hold, replacing (17), (18) and (19). Here $x_t = c_t - G(n_t)$ is the composite consumption-leisure good for the representative houshold and E_t^* is the same conditional expectation based on the worst-case belief in (20).

The law of motion for the distribution (22) is replaced by the wealth accumulation equation given by the budget constraint

$$c_t + q_t K_{t+1} + B_{t+1} = \tfrac{R_t}{\pi_t} B_t + (q_t + r_t) K_t + (1 - \tau)(w_t h_{it} n_{it} + \Pi_t^U + \Pi_t^F) + L_t, \tag{23}$$

and the consumption Euler equation for capital

$$u_c(x_t) = \beta E_t^* \frac{q_{t+1} + r_{t+1}}{q_t} u_c(x_{t+1}), \tag{24}$$

which then yields the optimal portfolio combination of K and B given return expectations.

2.4 Firms' forward-looking problems

As in standard RE models, firms in our economy solve intertemporal profit maximization problems due to the technological adjustment frictions introduced in section 2.1. Uncertainty can matter through the assumed stochastic discount factor used in these decisions. In particular, we let M_0^t denote the t-period ahead stochastic discount factor (SDF) as of date 0 used by a given firm to compute the expected value of its future profits.

In a representative agent model, or one with effectively complete markets, the SDF M_0^t used for any firm intertemporal decision would be unique, reflecting the (as if) representative agent's preferences. In an incomplete markets model however, it is generally not immediate how to pick M_0^t– it involves taking a stand on how different agents' preferences are aggregated to determine what a given firm does. A typical approach in the quantitative business cycle literature to deal with this difficulty is to observe that when uncertainty is assumed to equate risk (i.e. agents are SEU maximizers), fluctuations in the stochastic discount factors would be irrelevant up to log-linearization. Under linearity, that approach then makes the simplifying assumption that the firm sector is run by managers that are risk neutral. As with a SEU representative agent, risk would then not matter for firms' decisions.[10]

Our model addresses the challenge of SDF formation for firms' decision while maintaining a role for aggregate uncertainty in those decisions even under log-linearization. In particular, we show how under log-linearization the firms' intertemporal optimality conditions can be described *as if* they are owned by a risk-neutral owner evaluating the uncertain future under the worst-case conditional mean in equation (20).

As if risk neutral owner with the worst-case belief

Given the various types of producers in this economy, we develop the argument with more general notation. Consider a firm that faces states s_t and chooses actions x_t to solve

$$\max_x E\left[\sum_{t=0}^{\infty} M_0^t \Pi\left(x_t, x_{t-1}, s_t\right)\right]$$

where Π is profits and s collects aggregate state variables (both exogenous and endogenous) and M_0^t is a t-period ahead stochastic discount factor as of date 0 applicable to this firm.

Assuming usual regularity conditions, the optimal x_t choice follows from the firm's FOC

$$0 = \Pi_1\left(x_t, x_{t-1}, s_t\right) + E_t\left[M_{t+1}\Pi_2\left(x_{t+1}, x_t, s_{t+1}\right)\right]$$

where $M_{t+1} \equiv M_0^{t+1}/M_0^t$ is the one period ahead SDF and Π_j is the gradient with respect to the jth argument (possibly a vector). We note that in models where past actions matter only because of adjustment costs, we typically have $\Pi_2\left(x, x, s\right) = 0$ at the steady state. This indeed holds for the problems facing firms in section 2.1.

Without loss of generality, we assume that the firm's SDF takes the form

$$M_{t+1} = \beta\xi_{t+1}\widetilde{M}\left(s_t, s_{t+1}\right)$$

where $\widetilde{M}\left(s_t, s_{t+1}\right)$ is the weighted average of agents' risk-based component, capturing state price variation from reasons other than ambiguity. There are two key components of M_{t+1} that are common across agents: the discount factor β and ξ_{t+1}, the Radon-Nikodym-derivative

[10]See Carceles-Poveda and Coen-Pirani (2010) and Gornemann et al. (2021) for some discussion and approaches.

("change of measure") of the common worst-case-belief $p_t^*(s_t, s_{t+1})$ of the agents with respect to the econometrician's belief $p_t^0(s_t, s_{t+1})$ (given by (9))

$$\xi_{t+1} \equiv \frac{dp_t^*(s_t, s_{t+1})}{dp_t^0(s_t, s_{t+1})}$$

To be more precise, since the TFP innovation ϵ_t^Z is continuous, and the only difference between conditional distributions is the mean shifter, there exists a positive stochastic process ξ_{t+1} adapted to the agent's information set with $E_t[\xi_{t+1}] = 1$ such that $E_t^*[Y_{t+1}] = E_t[\xi_{t+1} Y_{t+1}]$ for any random variable Y_{t+1}. The adjustment ξ_{t+1} can therefore be thought of as a ratio of densities that shifts relatively more weight to states of the world that offer low utility value to the agent.

Let s^* denote the worst case steady state value of the state. The worst case steady state action of the firm thus solves

$$0 = \Pi_1(x^*, x^*, s^*) + \beta \widetilde{M}^* \Pi_2(x^*, x^*, s^*)$$

Consider the loglinearized FOC. Denote by Π_i^* the 1st and by Π_{ij}^* the 2nd derivative of profits evaluated at the worst case steady state, that is, $\Pi_{ij}^* = \Pi_{ij}(x^*, x^*, s^*)$. We get

$$\begin{aligned} 0 = &\Pi_{11}^* \widehat{x}_{t-1} + \Pi_{12}^* \widehat{x}_t + \left(\Pi_{13}^* + \beta \Pi_2^* \widetilde{M}_1\right) \widehat{s}_t + \\ &+ \beta \left[\Pi_{21}^* \widetilde{M}^* E_t \xi_{t+1} \widehat{x}_{t+1} + \Pi_{22}^* \widetilde{M}^* \widehat{x}_t + \left(\Pi_{23}^* \widetilde{M}^* + \Pi_2^* \widetilde{M}_2^*\right) E_t\left(\xi_{t+1} \widehat{s}_{t+1}\right)\right] \end{aligned}$$

where $\widetilde{M}^*$ and $\widetilde{M}_i^*$ are the level and i-th derivative of $\widetilde{M}$ at the worst case steady state.

We now have a difference equation that relates actions to shocks, but with expectations formed using the worst case belief. Like with households, this is what captures precautionary behavior by the firm.

Further simplification is possible when $\Pi_2^* = 0$. We then have

$$0 = \Pi_{11}^* \widehat{x}_{t-1} + \Pi_{12}^* \widehat{x}_t + \Pi_{13}^* \widehat{s}_t + \beta \widetilde{M}^* \left[\Pi_{21}^* E_t \xi_{t+1} \widehat{x}_{t+1} + \Pi_{22}^* \widehat{x}_t + \Pi_{23}^* E_t\left(\xi_{t+1} \widehat{s}_{t+1}\right)\right] \tag{25}$$

Equation (25) indicates that variation in uncertainty matters through ambiguity, while the standard risk-based $\widetilde{M}_{t+1}$ component does not matter up to first order. Instead, the latter component matters only through its steady-state level, by affecting the effective discount factor $\beta \widetilde{M}^*$. We follow a standard convention in the literature and set $\widetilde{M}^* = 1$ (see eg. Bayer et al. (2024)) so that the firms' effective discount factor is the primitive β.

Leveraging our earlier notation of conditional expectations under the worst-case beliefs as E_t^*, i.e. the change of measure ξ_{t+1}, means equation (25) then becomes

$$0 = \Pi_{11}^* \widehat{x}_{t-1} + \Pi_{12}^* \widehat{x}_t + \Pi_{13}^* \widehat{s}_t + \beta \left(\Pi_{21}^* E_t^* \widehat{x}_{t+1} + \Pi_{22}^* \widehat{x}_t + \Pi_{23}^* E_t^* \widehat{s}_{t+1} \right) \tag{26}$$

Equation (26) shows that the log-linearized optimality conditions are the same *as if* the firm is held by a risk-neutral manager with the same common worst-case belief as the agents in the economy.

Firms' log-linearized optimality conditions

In the economy described in section 2.1, there are three types of firms that make intertemporal decisions - the intermediate good producers, capital goods producers and unions, with per-period profit functions given by equations (3), (6) and (8), respectively. As detailed above (see equation (26)), the log-linearized optimality conditions for these intertemporal decisions maintain a role for precautionary behavior for firms through the worst-case belief embedded in firms' SDF.

In particular, consider first the intermediate goods sector, where a firm j chooses its nominal prices to maximize the present value of future profits

$$E_0 \sum_{t=0}^{\infty} M_{F,0}^t \lambda_Y^t (1-\tau^L) Y_t^{1-\tau^P} \left\{ \left(\frac{p_{jt}\bar{\pi}^t}{P_t} - mc_t \right) \left(\frac{p_{jt}\bar{\pi}^t}{P_t} \right)^{-\eta} \right\}^{1-\tau^P},$$

with indexation to steady-state inflation $\bar{\pi}$. Here $M_{F,0}^t$ is a SDF that can be specific to the firm sector. The corresponding first-order condition for price setting implies a Phillips curve

$$\log \left(\frac{\pi_t}{\bar{\pi}} \right) = \beta E_t^* \log \left(\frac{\pi_{t+1}}{\bar{\pi}} \right) + \kappa_Y \left(mc_t - \tfrac{1}{\mu^Y} \right), \tag{27}$$

where we dropped all terms irrelevant for a first-order approximation and defined $\kappa_Y = \frac{(1-\lambda_Y)(1-\lambda_Y\beta)}{\lambda_Y}$ where $\mu^Y = \frac{\eta}{\eta-1}$ is the target markup. This Phillips curve is an example of the more general equation (26), where we note again that expectations over a future endogenous variable, inflation, are formed under the worst-case belief E_t^*.

A similar Phillips curve under ambiguity obtains for wage inflation. In particular, given a stochastic discount factor $M_{U,0}^t$, a union j chooses nominal wages to maximize

$$E_0 \sum_{t=0}^{\infty} M_{U,0}^t \lambda_w^t \frac{W_t^F}{P_t} N_t \left\{ \left(\frac{W_{jt}\bar{\pi}_W^t}{W_t^F} - \frac{W_t}{W_t^F} \right) \left(\frac{W_{jt}\bar{\pi}_W^t}{W_t^F} \right)^{-\zeta_t} \right\},$$

where nominal wages that do not get re-optimized get indexed to steady-state $\bar{\pi}_W$. Since all unions are symmetric, we focus on a symmetric equilibrium and obtain the linearized wage

Phillips curve from the corresponding first-order condition as follows, leaving out all terms irrelevant at a first-order approximation around the stationary equilibrium

$$\log\left(\frac{\pi_t^W}{\bar{\pi}_W}\right) = \beta E_t^* \log\left(\frac{\pi_{t+1}^W}{\bar{\pi}_W}\right) + \kappa_w \left(mc_t^w - \frac{1}{\mu^W}\right), \tag{28}$$

with π_t^W being wage inflation, $mc_t^w = \frac{w_t}{w_t^F}$ is the actual and $\frac{1}{\mu^W} = \frac{\zeta-1}{\zeta}$ being the target markdown of wages the unions pay to households, W_t, relative to the wages charged to firms, W_t^F and $\kappa_w = \frac{(1-\lambda_w)(1-\lambda_w\beta)}{\lambda_w}$.

Finally, the representative capital goods producer chooses optimal investment I_t to maximize the expected present value of profits given by

$$E_0 \sum_{t=0}^{\infty} M_{Q,0}^t I_t \left\{ q_t \left[1 - \frac{\phi}{2}\left(\log \frac{I_t}{I_{t-1}}\right)^2\right] - 1 \right\}$$

where $M_{Q,0}^t$ is the stochastic discount factor relevant to the capital goods sector. Optimality of the capital goods production requires (again dropping all terms irrelevant up to first order)

$$q_t \left[1 - \phi \log \frac{I_t}{I_{t-1}}\right] = 1 - \beta E_t^* \left[q_{t+1}\phi \log\left(\frac{I_{t+1}}{I_t}\right)\right], \tag{29}$$

and each capital goods producer will adjust its production until (29) is fulfilled.

Put together, the firms' optimality conditions take the usual form familiar from standard linear RE models. The difference here is that aggregate uncertainty affects optimal choices and equilibrium outcomes through the conditional worst-case belief, common to all agents.

2.5 Government

We close the description of the economy with the government sector. The monetary authority controls the nominal interest rate on liquid assets, while the fiscal authority issues government bonds to finance deficits, chooses both the average tax rate in the economy and the tax progressivity, and makes expenditures.

We assume that monetary policy sets the nominal interest rate following a Taylor-type (1993) rule with interest rate smoothing:

$$\frac{R_{t+1}}{\bar{R}} = \left(\frac{R_t}{\bar{R}}\right)^{\rho_R} \left(\frac{\pi_t}{\pi_t^*}\right)^{(1-\rho_R)\theta_\pi} \left(\frac{Y_t}{Y_{t-1}}\right)^{(1-\rho_R)\theta_Y} \varepsilon_t^R. \tag{30}$$

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. First, a monetary policy shock:

$$\log \varepsilon_t^R = \rho_R^\epsilon \log \varepsilon_{t-1}^R + \epsilon_t^R \tag{31}$$

where $\epsilon_t^R \sim i.i.d.\ N(0, \sigma_R)$. Second, an inflation target shock

$$\log \pi_t^* = (1 - \rho_\pi) \log \bar{\pi} + \rho_\pi \log \pi_{t-1}^* + \epsilon_t^\pi \tag{32}$$

where $\epsilon_t^\pi \sim i.i.d.\ N(0, \sigma_\pi)$. The two shocks capture disturbances in the Taylor rule that can be potentially different in their persistence, through the AR(1) parameters ρ_R^ϵ and ρ_π, respectively. Following for example Justiniano et al. (2013), the interpretation and purpose of the inflation target shock is to account for the low frequency behavior of inflation possibly generated by the slow moving beliefs and resulting conduct of the monetary authority.

Fiscal policy is implemented through a government spending rule

$$\frac{G_t}{\bar{G}} = \left(\frac{G_{t-1}}{\bar{G}}\right)^{\rho_G} \left(\frac{B_t}{\bar{B}}\right)^{-(1-\rho_G)\gamma_B^G} \left(\frac{Y_t}{\bar{Y}}\right)^{(1-\rho_G)\gamma_Y^G} \tag{33}$$

and a rule that governs lump-sum transfers, which get transferred uniformly to all agents:

$$\frac{\tilde{L}_t}{\bar{\tilde{L}}} = \left(\frac{\tilde{L}_{t-1}}{\bar{\tilde{L}}}\right)^{\rho_L} \left(\frac{B_t}{\bar{B}}\right)^{-(1-\rho_L)\gamma_B^L} \left(\frac{Y_t}{\bar{Y}}\right)^{(1-\rho_L)\gamma_Y^L} \tag{34}$$

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 $\bar{\tilde{L}}$, respectively. We formulate transfers as percent of output, $L_t = \log \tilde{L}_t\ \bar{Y}$, and set $\bar{\tilde{L}} = 1$, i.e. average transfers are zero as we later calibrate the idiosyncratic income process after transfers.

Total government tax revenues are $T_t = \tau \left(w_t n_{it} h_{it} + \mathbb{I}_{h_{it} \neq 0} \Pi_t^U + \mathbb{I}_{h_{it}=0} \Pi_t^F\right)$ and government debt is determined residually from the government budget constraint:

$$B_{t+1} = G_t + L_t - T_t + R_t B_t / \pi_t \tag{35}$$

2.6 Goods, asset, and labor market clearing

The labor market clears at the competitive wage given in (4). The liquid asset market clears whenever the following equation holds:

$$B_{t+1} = \mathbb{E}_t^{XS} \left[\lambda b_{a,t}^* + (1 - \lambda) b_{n,t}^*\right], \tag{36}$$

where $b^*_{a,t}, b^*_{n,t}$ are functions of the states (b, k, h), and depend on how households value asset holdings in the future, $\mathbb{W}_{t+1}(b, k, h)$, and the current set of prices $(w_t, w^F_t, \Pi^F_t, \Pi^U_t, q_t, r_t, R_t, \pi_t, \pi^W_t, \Theta_t, \mathbb{W}_{t+1})$. Future prices do not show up because we can express the value functions such that they summarize all relevant information on the expected future price paths.

Expectations $\mathbb{E}^{XS}_t$ in the right-hand-side expression are taken w.r.t. the *cross-sectional* distribution $\Theta_t(b, k, h)$. Equilibrium requires the total *net* amount of bonds the household sector demands, B^d, to equal the supply of government bonds. In gross terms there are more liquid assets in circulation; some households borrow up to $\underline{B}$.

Last, the market for capital has to clear, i.e.,

$$K_{t+1} = \mathbb{E}^{XS}_t[\lambda k_t{}^* + (1 - \lambda)k], \tag{37}$$

where the RHS defines the aggregate supply of funds from households—both those that trade capital, λk^*_t, and those that do not, $(1 - \lambda)k$. Again k^*_t is a function of the current prices and continuation values. The goods market then clears due to Walras' law, whenever labor, bonds, and capital markets clear.

When we consider the representative household model, we can think of the RHS of equations (36) and (37) as simply given by (23) and (24), respectively. In other words, the representative household model only changes equilibrium conditions in replacing the Bellman equation and the capital and bonds demand equations, but leaves the entire other model structure unchanged.

2.7 Equilibrium

A *sequential equilibrium with recursive planning under ambiguity* in our model is a sequence of policy functions $\{x^*_{a,t}, x^*_{n,t}, b^*_{a,t}, b^*_{n,t}, k^*_t\}$, a sequence of value functions $\{\mathbb{W}_t\}$, a sequence of prices $\{w_t, w^F_t, \Pi^E_t, \Pi^U_t, q_t, r_t, R_t, \pi_t, \pi^W_t, \tau_t\}$, a sequence of stochastic states $\{Z_t, a_t, \bar{\sigma}^2_{h,t}, \pi^*_t, \varepsilon^R_t\}$ and shocks $\{\epsilon^Z_t, \epsilon^a_t, \epsilon^\sigma_t, \epsilon^\pi_t, \epsilon^R_t\}$, aggregate capital and labor supplies $\{K_t, N_t\}$, distributions Θ_t over individual asset holdings and productivity, and expectations for the distribution of future prices, such that

1. Conditional beliefs E^*_t over exogenous variables are formed:

   (a) under the worst-case process for aggregate TFP, Z_{t+1}, as in equation (21);

   (b) under the true DGP processes for the other exogenous states and shocks.

2. Bellman equations are satisfied:

(a) Given the functionals $E_t^* \mathbb{W}_{t+1}$, formed over the continuation value under the conditional beliefs E_t^*, and period-t prices, policy functions $\{x_{a,t}^*, x_{n,t}^*, b_{a,t}^*, b_{n,t}^*, k_t^*\}$ solve the households' recurive multiple priors planning problem; and

(b) Given the policy functions $\{x_{a,t}^*, x_{n,t}^*, b_{a,t}^*, b_{n,t}^*, k_t^*\}$ and prices, the value functions $\{V_t^A\}, \{V_t^N\}$ and $\{\mathbb{W}_t\}$ are a solution to the Bellman equations (17), (18), 19).

3. Distributions of wealth and income evolve according to households' policy functions, per equation (22).

4. Market clearing: the labor, the final goods, the bond, the capital, and the intermediate goods markets clear in every period.

5. Policy: interest rates on bonds are set according to the Taylor rule in equation (30) and fiscal policies are set according to the fiscal rules in equations (33) and (34).

6. Consistency of beliefs:

   (a) Beliefs over functions mapping exogenous to endogenous variables are consistent with actual equilibrium functions.

   (b) The worst-case belief for TFP at each period t is the minimizing belief over the continuation value $\mathbb{W}_{t+1}$ out of the primitive set in (11).

7. Data simulated under the true DGP draws aggregate TFP Z_t from the process in (9).

A Rational Expectations version is a special case of this equilibrium, where the belief set over TFP is assumed to be a singleton that coincides with the true DGP. There are three observations to make here. First, under ambiguity, there is a sytematic difference between the conditional belief that supports optimal choices and the true DGP, as summarized by points (1.a) and (7). Second, the consistency of beliefs criterion in (6.a) imposes the same 'structural knowledge' on the part of agents as in RE models: agents are not confident in the probability distribution of exogenous variables (here aggregate TFP) but have full knowledge over the equilibrium mappings to endogenous variables. Third, while under RE there is a single belief, point (6.b) looks to verify that the equilibrium worst-case belief is indeed the conjectured one in point (1.a), as discussed earlier in section 2.3.

2.8 Solution method

We follow the tradition of Reiter (2009) and solve the household problem globally while approximating aggregate dynamics by a first-order perturbation. In particular, we build on

Bayer et al. (2024), whose implementation of the Reiter approach allows for full information estimation of the model's dynamics, and extend that framework to ambiguity along the lines of Ilut and Schneider (2014).

The key advantage of our method is to leverage the property characterizing our model that ambiguity allows aggregate uncertainty to have first-order effects on decisions, both in steady state and following changes in uncertainty. This first-order approach also means that the interaction between the global solution aspect of solving the household problem and (time-varying) aggregate uncertainty does not add significant computational complexity to the first-order perturbation solutions in Bayer et al. (2024) or Auclert et al. (2020). Additionally, Bayer et al. (2024) 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 general logic of the method mirrors the equilibrium definition in section 2.7. First, we solve the model *as if* it is a RE model in which the worst case scenario expectations in equation (21) are correct on average. In this step, we compute the equilibrium using a first-order perturbation. Second, we take the resulting log-linear equilibrium decision rules formed under ambiguity and then characterize the dynamics under the econometrician's law of motion for TFP described by the probabilistic evolution in equation (9).

Log-linearization around the worst-case steady state

The steady state of the log-linearized *as if* RE model under the worst-case belief can be computed as a deterministic steady state since agents act as if the economy converges there in the long run. In this deterministic steady state, TFP equals its long-run value Z^* under the worst-case belief in equation (21), given simply by

$$Z^* = \bar{Z} \exp\left(\frac{-\bar{a}}{1-\rho_z}\right), \tag{38}$$

where $\bar{a}$ is the long-run value of ambiguity in equation (12) and $\bar{Z}$ is the long-run value under the true process in equation (9). Under the deterministic steady state Z^*, we can then compute the worst case steady state of the endogenous variables, including the value functions $\mathbb{W}(b, k, h)$ and distributions of individual states $\Theta(b, k, h)$. For a simple notation, denote the collection of all other shocks and endogenous variables as $\mathbf{Y}$. The worst-case deterministic steady-state is then given by $(Z^*, \mathbf{Y}^*)$. With this steady state $(Z^*, \mathbf{Y}^*)$ at hand, we log-linearize the model around it.

At this stage, by solving the *as if* RE model under the worst-case belief, the computation builds on techniques developed for expected utility models. In particular, we follow the method in Bayer et al. (2024) to first compute the stationary distribution under Z^* and thus the collection $(Z^*, \mathbf{Y}^*)$, and then the model's linear dynamics around it. The resulting solution describes how endogenous variables respond to the state variables, including the current ambiguity a_t and current aggregate TFP Z_t. Importantly, this response does not depend on the true DGP in equation (9). Instead it is entirely determined by agents' (worst-case) beliefs. Finally, this computation also produces the agents continuation values $\mathbb{W}_{t+1}$. These functions can be used to verify, as in step 6 of the equilibrium definition of section 2.7, the consistency of the worst-case beliefs, given the equilibrium policies and prices derived under the log-linear approximation.

Ergodic steady state and equilibrium dynamics

To characterize equilibrium dynamics, we combine the equilibrium law of motion derived under ambiguity in the previous step, and the true DGP. The key step here is to account for the true DGP in equation (9) being different than agents' beliefs, by noting that

$$\log Z_t = E^*_{t-1} \log Z_t + a_{t-1} + \epsilon^Z_t \tag{39}$$

As described in sections 2.3 and 2.4, the fact that current TFP is on average higher than the worst-case belief by a_{t-1} is just a manifestation of the change of measure between the agent's cautious, worst case conditional beliefs $E^*_{t-1} \log Z_t$ and the true DGP.

Following the insights of Ilut and Schneider (2014), this change of measure makes aggregate uncertainty have a role in both the model's steady state and its dynamics around it. First, the *ergodic* steady state as measured by the econometrician is the long-run value of variables when (i) the long-run value of actual TFP in equation (9) equals $\log Z = \log \bar{Z}$, while (ii) variables follow their log-linear equilibrium response to states as determined under ambiguity, *as if* TFP is on a declining path towards its long-run worst-case value of $\log Z^* = \log \bar{Z} - \frac{\bar{a}}{1-\rho_z}$. Following the notation above, we can denote this ergodic steady state as the collection $(\bar{Z}, \bar{\mathbf{Y}})$. Second, given the log-linearized equilibrium responses derived under ambiguity and the law of motion in equation (39), we can then also characterize the model's dynamics around this ergodic steady state $(\bar{Z}, \bar{\mathbf{Y}})$. For example, to describe 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. (2024). 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 include the following observable time series: 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, here 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 and their model counterparts. Appendix B.1 lists the data sources.

3.1 Estimation approach

Our estimation approach exploits the tractability of our linear solution method described in Section 2.8. We use the model's ergodic distribution to fit the data, along two estimation objectives. To match average data moments, we use the ergodic steady state $(\bar{Z}, \bar{\mathbf{Y}})$. To match the historical path of the observed data in the Bayesian likelihood estimation, we use the model's dynamics. Critically, in our model, the ergodic steady-state is in itself a function of the dynamics. The model parameters generically influence the model's dynamics directly or indirectly, through the worst-case steady state approximation point $(Z^*, \mathbf{Y}^*)$. Through dynamics, parameters therefore also affect the ergodic steady state $(\bar{Z}, \bar{\mathbf{Y}})$. This makes estimation more challenging and the identification of parameters richer, since they affect jointly the moment-matching fit and the maximum likelihood score.

In contrast, in linear Rational Expectations DSGE models aggregate uncertainty has no effects on the long-run values of variables, therefore the ergodic steady state would coincide there with the deterministic one. In that approach, typical in the DSGE literature, the two estimation objectives (moment matching and likelihood maximization) can thus be maximized separately, since dynamics do not matter for the long-run model-implied values.[11]

Our estimation is nevertheless tractable and we implement it by an iterative procedure as follows. We isolate the parameters to be estimated that affect the approximation point, namely the worst-case steady state $(Z^*, \mathbf{Y}^*)$. Call this parameter set θ_{SS}. Besides standard parameters that would matter in a representative agent for this steady-state (like the discount factor or government spending etc.), in the HANK model parameters governing the incomplete markets aspect and the trading friction also matter since they affect the wealth distribution. We then implement a two-step procedure.

First, conditional on some initial parameter values in θ_{SS}, we run a Bayesian maximum likelihood estimation 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, and gets fixed at this step. The estimation converges quickly, leveraging the linearity of the state-space representation. Then, given the posterior mode obtained in the Bayesian estimation, we can compute the ergodic steady state $(\bar{Z}, \bar{\mathbf{Y}})$, which we compare with the counterpart data moments. The second step of the procedure is to adjust the parameters in θ_{SS} to minimize the (equally weighted) moment squared distance. The procedure then potentially restarts to improve the model fit. Finally, 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, ξ, 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

[11]The main exception and thus similar work to us is models that employ higher-order perturbation, like Fernández-Villaverde et al. (2011), Fernández-Villaverde et al. (2015). There, similar to us, the equilibrium is perturbed around the deterministic steady state, and the higher-order terms make aggregate uncertainty shift the ergodic steady state, which gets compared to data moments, away from the deterministic one. The difference for us is that this shift occurs in the linear approximation due to change of measure in the actual process from its worst-case belief to the true process.

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)
ξ	4.000	Relative risk aversion	Kaplan and Violante (2014)
γ	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)

average quarterly depreciation is $\delta_0 = 0.025$.

Estimated parameters

We first introduce the estimated parameters, shown in Table 2, and then in Section 3.3 discuss their estimated values. The first three rows of Table 2 refer to the aforementioned set θ_{SS} of parameters, which get adjusted in the moment-matching estimation step. 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. Table 3 lists the target moments and the model-implied ones. We discuss their interpretation in Section 3.3.

The rest of Table 2 comprises the remaining parameters, which get estimated via Bayesian likelihood and for which we report the prior and posterior credible intervals. We now detail the prior construction. Following Justiniano et al. (2011), we 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

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	ρ_Z	Beta(0.5, 0.2)	0.945
		(6.006, 6.064)			(0.930, 0.959)
ϕ	Gamma(4, 2)	0.561	ρ_A	Beta(0.5, 0.2)	0.963
		(0.366, 0.806)			(0.948, 0.975)
κ	Gamma(0.1, 0.03)	0.046	σ_A	Inv.-Gamma(0.5, 0.25)	90.879
		(0.031, 0.065)			(68.496, 120.78)
κ_w	Gamma(0.1, 0.03)	0.101	ρ_S	Beta(0.5, 0.2)	0.496
		(0.069, 0.137)			(0.325, 0.670)
$\tilde{z}$	Beta(0.99, 0.01)	0.972	σ_S	Inv.-Gamma(0.5, 0.25)	22.182
		(0.966, 0.978)			(15.366, 30.374)
	Monetary Policy			Fiscal Policy	
ρ_R	Beta(0.5, 0.2)	0.171	ρ_G	Beta(0.5, 0.2)	0.181
		(0.060, 0.316)			(0.042, 0.422)
θ_π	Normal(1.7, 0.3)	2.202	γ_{GB}	Gamma(1.0, 0.2)	0.806
		(1.996, 2.435)			(0.611, 1.030)
θ_Y	Normal(0.13, 0.05)	0.075	γ_{GY}	Gamma(1.0, 0.2)	0.954
		(0.048, 0.102)			(0.684, 1.253)
ρ_R^ϵ	Beta(0.5, 0.2)	0.553	ρ_L	Beta(0.5, 0.2)	0.361
		(0.344, 0.686)			(0.102, 0.687)
σ_R^ϵ	Inv.-Gamma(0.1, 2.0)	0.136	γ_{LB}	Gamma(0.2, 0.2)	0.206
		(0.113, 0.175)			(0.155, 0.263)
ρ_π^ϵ	Beta(0.5, 0.2)	0.948	γ_{LY}	Gamma(0.2, 0.2)	1.044
		(0.929, 0.985)			(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.

are assumed to follow a *Beta* distribution with mean 0.5 and standard deviation 0.2 for TFP and a standard deviation of 20 for ambiguity and idiosyncratic risk.[12] The standard deviations of the shocks follow $Inverse - Gamma$ distributions with prior mean 0.1% and standard deviation 2%, again 20% for ambiguity and idiosyncratic risk.

[12] The prior for ambiguity is consistent with estimates in Ilut and Schneider (2014).

Estimating steady-state ambiguity

For steady state ambiguity, we find it convenient and informative to estimate the ratio of the long run value of TFP under the worst-case belief relative to the true process. Using equation (38), this ratio, denoted by $\tilde{z}$, is

$$\tilde{z} \equiv \frac{Z^*}{\bar{Z}} = \exp\left(\frac{-\bar{a}}{1-\rho_z}\right) \tag{40}$$

In standard Rational Expectations models the long-run value $\bar{Z}$ is usually normalized to a simple constant, for example $\bar{Z} = 1$. We instead normalize the long-run value under the worst-case belief Z^* which controls the worst-case deterministic steady state. This allows us to avoid re-solving the approximation point $(Z^*, \mathbf{Y}^*)$ in the Bayesian likelihood. In particular, we can normalize $Z^* = 1$ and estimate the ratio $\tilde{z}$, using for example a *Beta* distribution like in Table 2, which imposes $\tilde{z} < 1$. For a given estimated value of $\tilde{z}$, the implied $\bar{Z}$ is then

$$\bar{Z} = \tilde{z}^{-1} > Z^* = 1 \tag{41}$$

The ratio $\tilde{z}$ controls how much the steady-state ambiguity shifts through its change of measure the steady state away from its worst-case deterministic value $(Z^*, \mathbf{Y}^*)$ to the ergodic steady state $(\bar{Z}, \bar{\mathbf{Y}})$. When $\tilde{z}$ is small, agents worry about a path towards a low long-run value of TFP, or put differently, the true one is large compared to that worst-case belief. The differences in long-run values get mapped in differences in endogenous variables ($\mathbf{Y}^*$ vs $\bar{\mathbf{Y}}$) as a function of the model's dynamics, which we discuss in Section 3.3.

In Table 2 we also estimate the persistence ρ_z of the TFP process. This parameter gets partly identified also from the shock dynamics and its effect on the likelihood. Given ρ_z, 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.[13] Indeed, using the definition in equation (40) we have

$$\bar{a} = -(1-\rho_z)\log\tilde{z} \tag{42}$$

The one-step ahead ambiguity $\bar{a}$ (or equivalently the long-run ratio $\tilde{z}$) is the only new steady-state parameter introduced by our multiple priors model. To interpret and discipline its magnitude we refer to 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

[13] This also makes clear that alternatively we could estimate $\bar{a}$ and ρ_z.

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

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 parameters $\tilde{z}, \rho_z$ and σ_z independently, the implied $\bar{a}$ is significantly larger than σ_z. Disciplined by the bound constraint, we the reported results are for a restricted estimation version where the upper bound is tight and so $\bar{a} = \sigma_z$. The latter parameter is thus also given by equation (42) and does not show up as an additional free parameter in Table 2.

3.3 Estimates and ergodic moments

Aggregate uncertainty matters for our ergodic steady state. We now discuss its implications and mechanisms. The last column in Table 3 reports model-implied moments in the worst-case deterministic SS $(Z^*, \mathbf{Y}^*)$, while the middle one for the ergodic SS $(\bar{Z}, \bar{\mathbf{Y}})$. The difference between these moments identifies the role of aggregate uncertainty.

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. (2024), 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, using equation (42), can be read as $\bar{a} = 0.0015$. Interestingly, this value is about half of the corresponding estimate of $\bar{a}$ in the representative 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 these 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 the 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 steady-state government budget constraint of equation (35). 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 SS.

Second, in the ergodic steady state the investors' exposure to the aggregate uncertainty characterizing capital needs to be compensated by a financial 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 (eg. 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.

Finally, consider the remaining parameters in Table 2. These are estimated in our Bayesian likelihood procedure, and the posteriors 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 B.2 provides details on convergence. We only briefly comment on the estimated values here. These parameters influence the model's dynamics, which we will discuss in detail in the next section. 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 both the Taylor rule and the inflation target, which we estimate quite accurately from observations of the policy rate and inflation.

4 Business cycle dynamics

Aggregate uncertainty, modeled here as ambiguity, emerges as 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.

4.1 Historical fit

We first show how well the model actually fits the data. 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

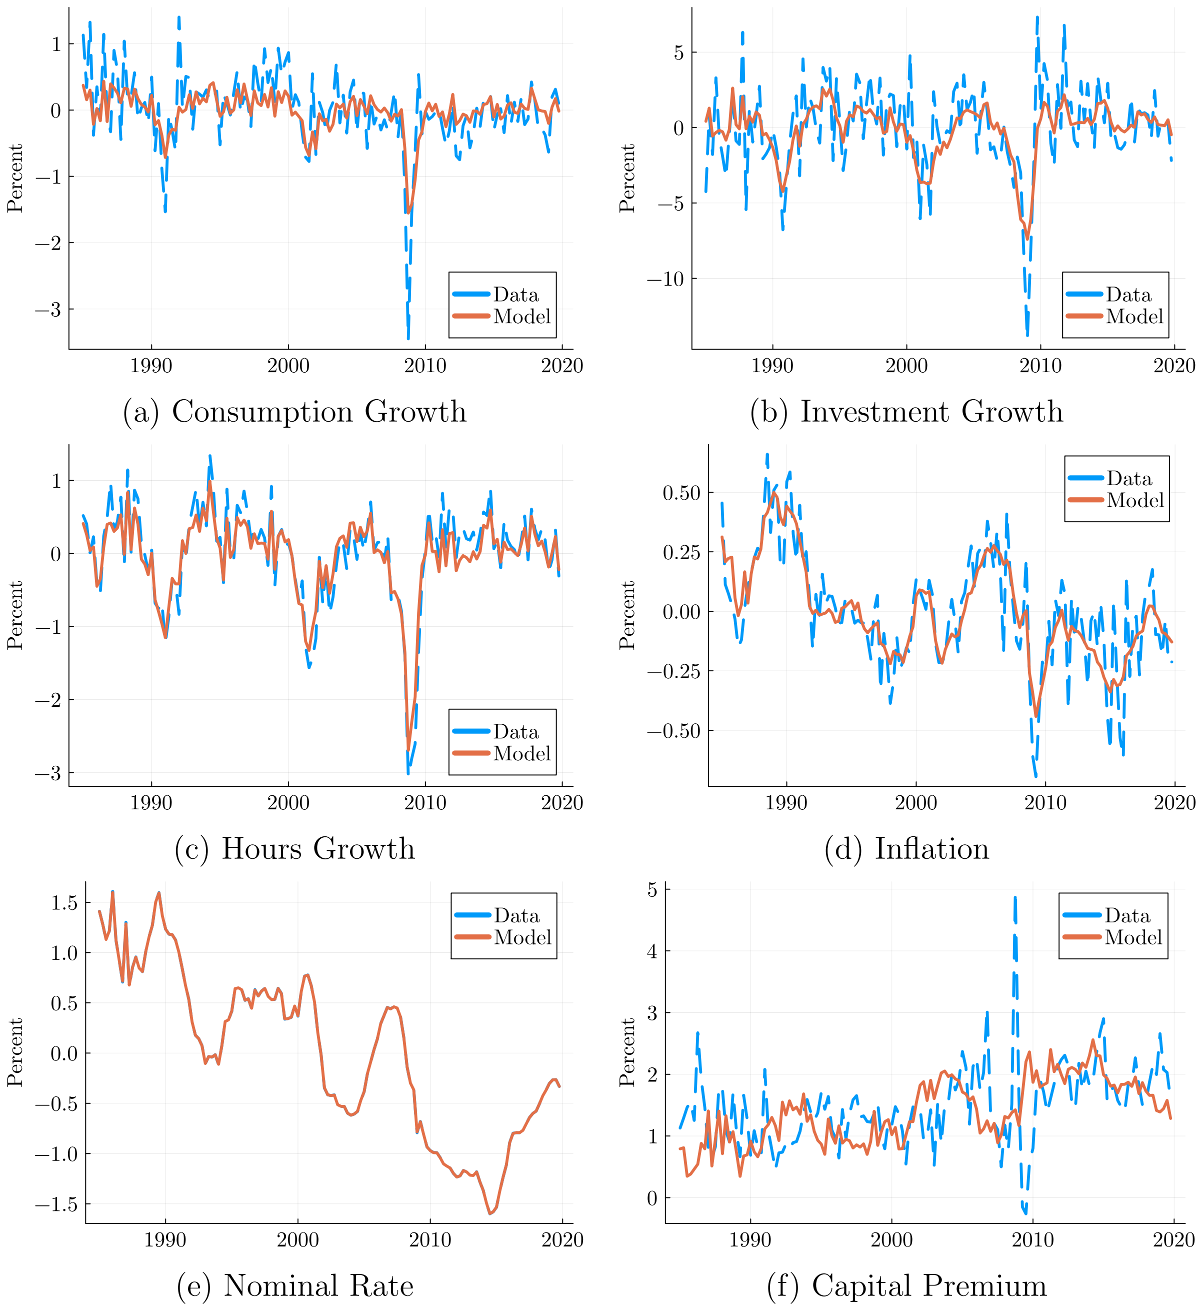

Figure 1: Model vs Data: Kalman smoother for the estimated HANK model and the data used in estimation.

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.

Historical variance decomposition

To understand how our estimated model delivers this fit, we start with investigating the role played by each shock in a historical variance decomposition exercise. In particular, given the linearity of the solution, Figure 2 decomposes the overall model fit (red line of Figure 1) into the specific historical contribution of each shock. In addition, the decomposition also accounts for the role played by the estimated initial condition (purple component), which is particularly important for the downward trend in the nominal rate and inflation.

Clearly, the most important shock in driving the historical fit across observables is the ambiguity one. First, this shock accounts for the historical booms and busts, with strong comovement in consumption, investment and employment. Second, it accounts especially for the lower frequency movements in the nominal rate that appear to be driven by the implied slower moving changes in the real rate. Third, aggregate uncertainty also helps explain business cycle and lower frequency movements in the historical capital premium. The other shocks play a more muted role. The exception is the inflation target shocks, which matters especially for the slow-moving dynamics of inflation and partly of the nominal rate.

4.2 Impulse responses

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. In particular, Figure (3) plots the IRF to our ambiguity shock in the baseline HANK model (in solid dark lines). We will later also draw comparisons to a counterfactual RANK model (in dashed blue lines) that keeps the same parameters but eliminates the incomplete markets aspects of our economy (recall the discussion in section 2.3 of how this variant is constructed).

The ambiguity shock in the baseline HANK

A loss of confidence over the conditional distribution of aggregate 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.

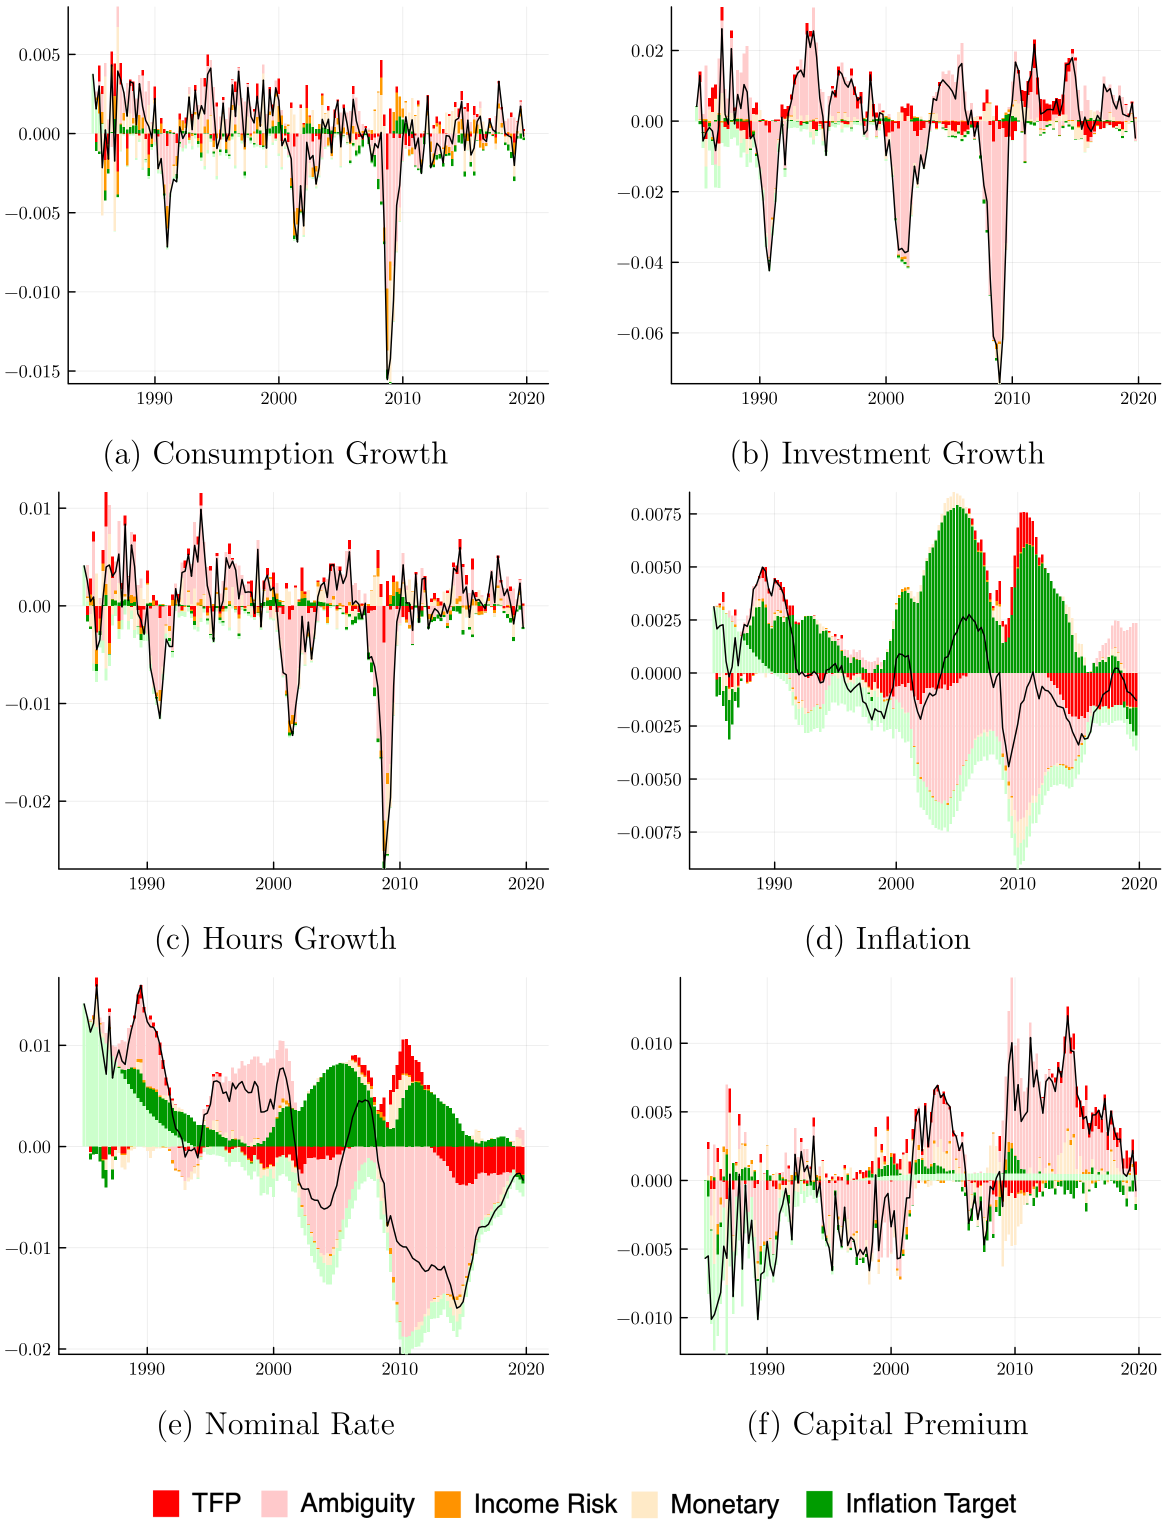

Figure 2: Historical decompositions for the estimated HANK model

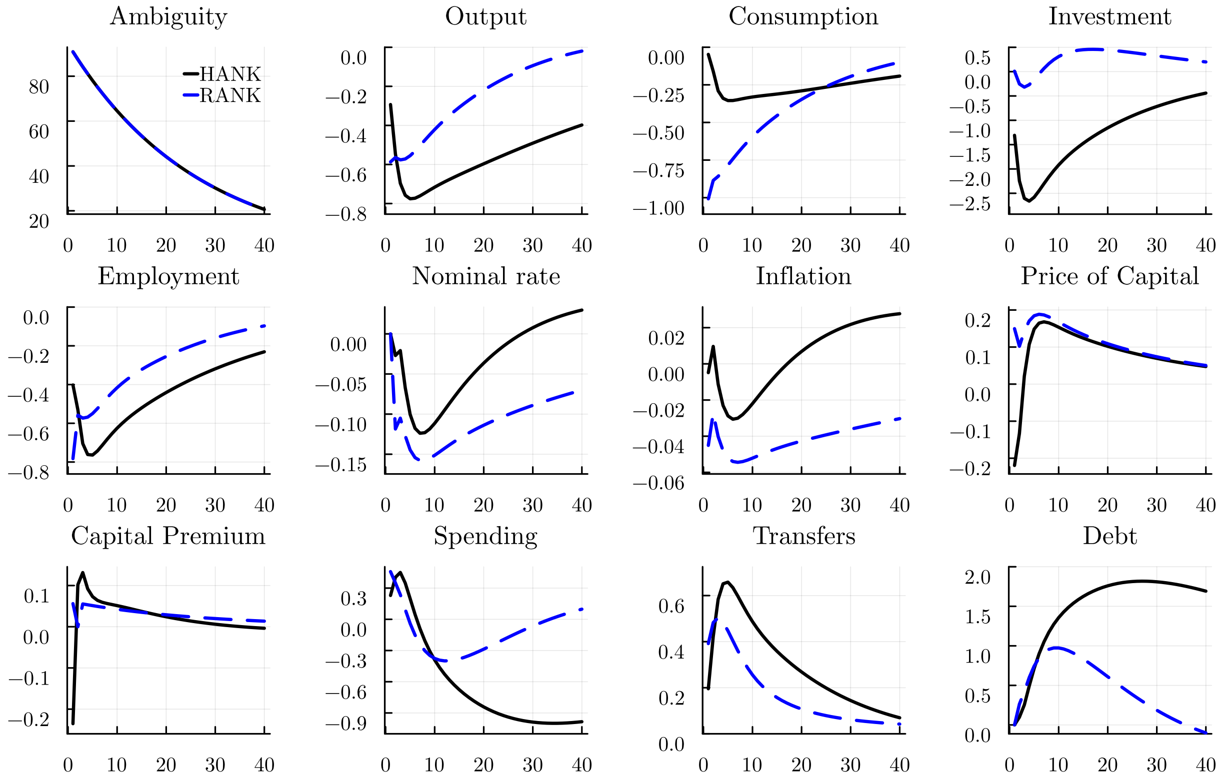

Figure 3: Impulse Responses to an ambiguity shock in estimated HANK model and counterfactual RANK (under the same parameters)

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 monetary 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-King 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.[14] The key 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 physical 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 (27). 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 2.4). 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 equation (16), which is persistently positive in this recession. The premium indicates that capital, an uncertain and illiquid asset, requires a higher equilibrium excess return compared to the risk-free and liquid asset.[15]

Decomposing the response of the capital premium

Figure 4 decomposes the sources underlying the predictability of a positive premium in

[14]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.

[15]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.

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 (eg. 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

$$E_t Prem_{t+1} = \underbrace{E_t^* Prem_{t+1}}_{\substack{\text{liquidity} \\ \text{part}}} + \underbrace{E_t Prem_{t+1} - E_t^* Prem_{t+1}}_{\substack{\text{uncertainty} \\ \text{part}}} \tag{43}$$

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.[16]

Panel (b) of Figure 4 plots the decomposition of the premium in equation (43) 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 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.

[16]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).

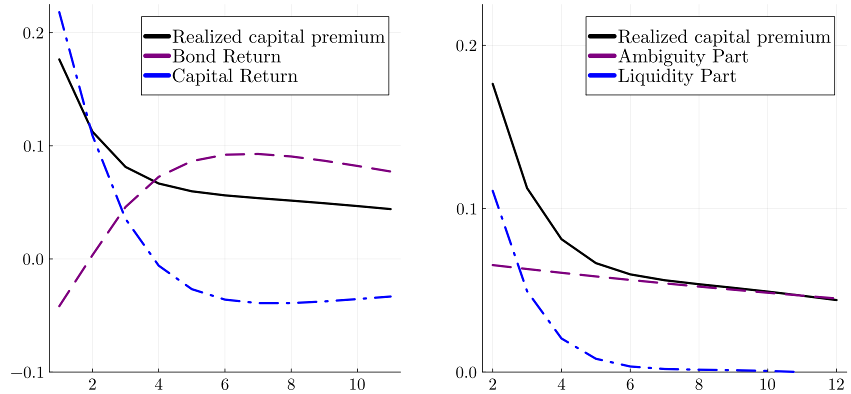

(a) Return Decomposition (b) Mechanism Decomposition

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

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. Following an increase in aggregate uncertainty, the presence of the illiquidity friction in the HANK model acts to amplify the households' incentive to move away from capital compared to the RANK version. 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 have relatively less labor income while holding most of the capital and thus driving most of the investment dynamics compared to a representative agent. Faced with more aggregate uncertainty about TFP, the rich agents in the HANK model 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, using a theoretical variance decomposition at business cycle frequency (following Uhlig (2001)), we can further establish that in our baseline HANK model the ambiguity shock accounts for about 90% of the model-implied variation in investment, compared to 50% in RANK. 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. 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%.

Importantly, the capital premium in the counterfactual RANK is not only less volatile, but we report that is also significantly smaller on average, at only 0.11%. Since the trading friction does not operate in this counterfactual, its premium is entirely a compensation for ambiguity. Earlier in Section 3.3 we discuss how the total ergodic premium in our baseline model is 5.55%, and that 3.21% can be decomposed as an ambiguity premium. The latter is thus more than an order of magnitude larger than its RANK counterpart. This difference showcases again that the same amount of ambiguity interacts with the HANK side of the model to dramatically increase the uncertainty compensation required by the marginal investor in our model.

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

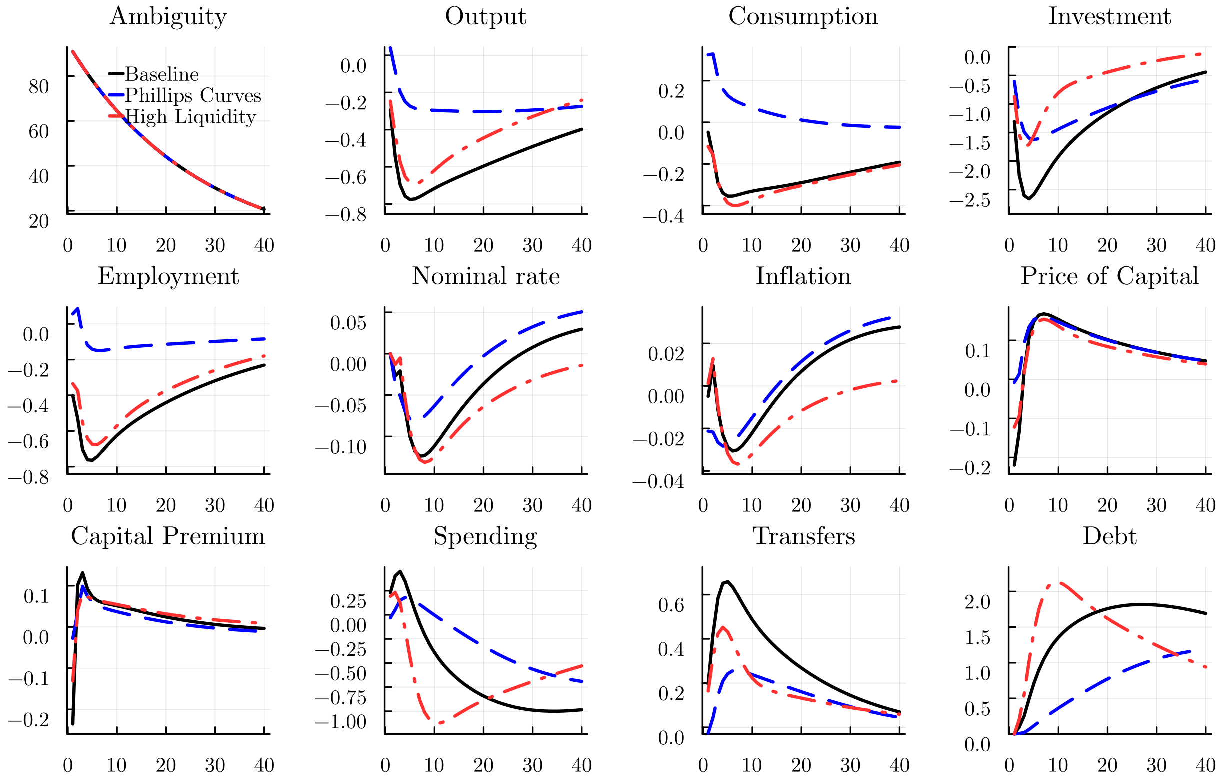

Figure 5: Model counterfactuals

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.

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 (27). We can turn that effect off by leveraging the linearity of our solution method since

$$\mathbb{E}_t \widehat{\pi}_{t+1} = \mathbb{E}_t^* \widehat{\pi}_{t+1} + \varepsilon_{\pi z} a_t \tag{44}$$

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 (27) is now given by $\mathbb{E}_t\widehat{\pi}_{t+1}$. A similar approach as in equation (44) 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 shock in the model is the idiosyncratic income risk, i.e. an innovation ϵ_t^σ to the conditional volatility of labor income in equation (14). 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 over-shoots soon after impact, by slowly returns to steady state from above.[17] 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 that regards it resembles the precautionary saving property of the aggregate uncertainty increase as well, 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 return on capital pushed down significantly and persistently the desire to investment in the

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

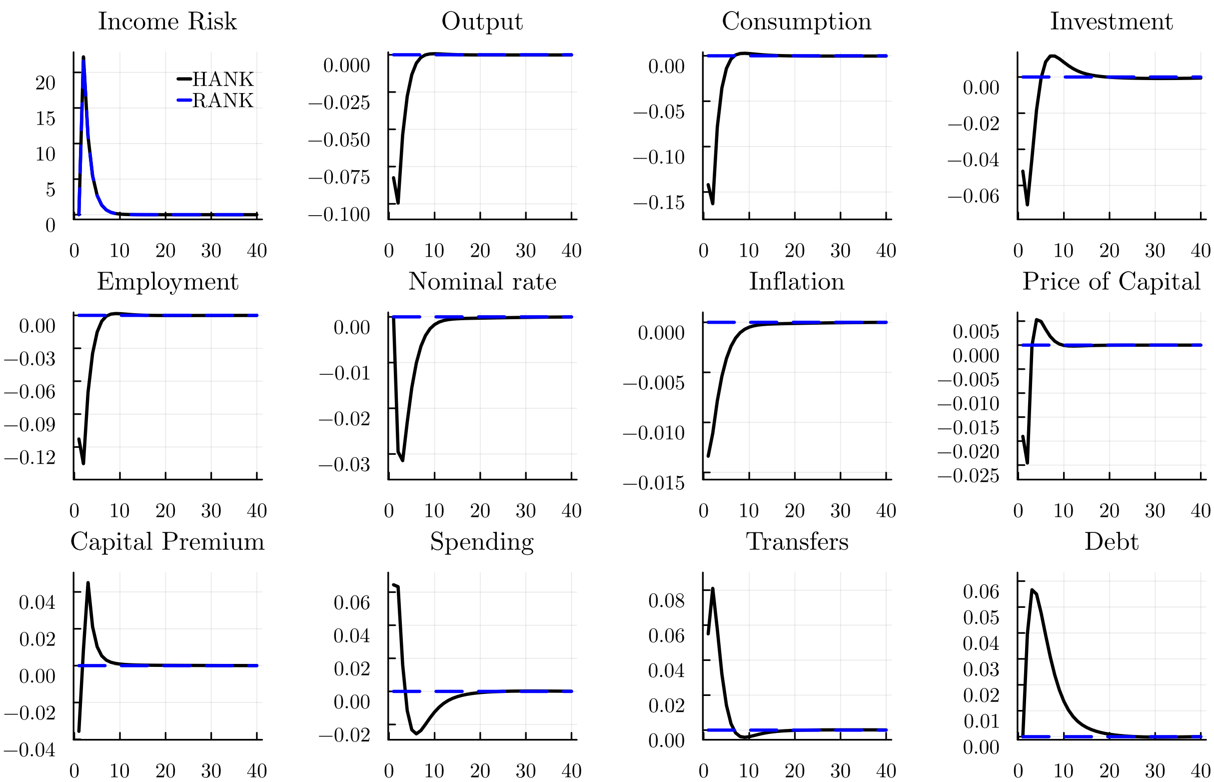

Figure 6: Impulse Responses to an Idiosyncratic income risk shock in estimated HANK model and counterfactual RANK (under the same parameters)

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 historical decomposition of Figure 2, where risk shocks play a small role.

The other shocks

We conclude the discussion of the model's response to shocks with a brief comment on the IRFS to the remaining shocks. These responses are consistent with standard findings in typical estimated NK models, and thus for brevity we relegate to the Appendix.

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 7 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.[18]

Finally, consider the responses to nominal shocks. A contractionary monetary policy shock induces a higher real rate, dampening demand for consumption and investment and leading to a relatively short recession with lower employment and persistently low inflation (see Figure 8 in the Appendix). Finally, 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 9 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., Bardóczy, B., Rognlie, M., and Straub, L. (2021). Using the sequence-space Jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408.

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

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

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. (2024). Shocks, frictions, and inequality in us business cycles. *American Economic Review*, 114(5):1211–47.

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*, page rdae054.

[18] 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 (eg. Bayer et al. (2024)) or ambiguity with a representative agent (Ilut and Schneider (2014)).

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.

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.

Carceles-Poveda, E. and Coen-Pirani, D. (2010). Owning capital or being shareholders: An equivalence result with incomplete markets. *Review of Economic Dynamics*, 13(3):537–558.

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.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.

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

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

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., Guerrón-Quintana, P., Kuester, K., and Rubio-Ramírez, J. (2015). Fiscal volatility shocks and economic activity. *American Economic Review*, 105(11):3352–3384.

Fernández-Villaverde, J., Guerrón-Quintana, P., Rubio-Ramirez, J. F., and Uribe, M. (2011). Risk matters: The real effects of volatility shocks. *American Economic Review*, 101(6):2530–2561.

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

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.

Gornemann, N., Kuester, K., and Nakajima, M. (2021). Doves for the rich, hawks for the poor? distributional consequences of systematic monetary policy. Technical report. Working Paper.

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. National Bureau of Economic Research WP 32198.

Ilut, C. L. and Schneider, M. (2022). Modeling uncertainty as ambiguity: A review. National Bureau of Economic Research WP 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.

Kaplan, G. and Violante, G. L. (2018). Microeconomic heterogeneity and macroeconomic shocks. *Journal of Economic Perspectives*, 32(3):167–194.

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.

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

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.

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 Some details on the household side

The household's felicity function u exhibits a constant relative risk aversion (CRRA) with risk aversion parameter ξ,

$$u(x_{it}) = \frac{x_{it}^{1-\xi} - 1}{1-\xi}, \tag{45}$$

where $x_{it} = c_{it} - G(h_{it}, n_{it})$ is household i's composite demand for goods consumption c_{it} and leisure and G measures the disutility from work.

Assuming a proportional income-tax, a household's net labor income, y_{it}, is given by

$$y_{it} = (1-\tau)(w_t h_{it} n_{it}), \tag{46}$$

where w_t is the aggregate real wage rate and τ the tax rate. Given net labor income, the first-order condition for labor supply is

$$\frac{\partial G(h_{it}, n_{it})}{\partial n_{it}} = (1-\tau)(w_t h_{it}). \tag{47}$$

Assuming that G has a constant elasticity w.r.t. n, $\frac{\partial G(h_{it},n_{it})}{\partial n_{it}} = (1+\gamma)\frac{G(h_{it},n_{it})}{n_{it}}$ with $\gamma > 0$, we can simplify the expression for the composite consumption good, x_{it}, making use of this first-order condition (47), and substitute $G(h_{it}, n_{it})$ out of the individual planning problem

$$x_{it} = c_{it} - G(h_{it}, n_{it}) = c_{it} - \frac{(1-\tau)w_t h_{it} n_{it}}{1+\gamma}. \tag{48}$$

When the Frisch elasticity of labor supply is constant and the tax schedule has the form (46), the disutility of labor is always a fraction of labor income and constant across households. Therefore, in both the household's budget constraint and felicity function, only after-tax income enters and neither hours worked nor productivity appear separately.

This implies that we can assume $G(h_{it}, n_{it}) = h_{it}\frac{n_{it}^{1+\gamma}}{1+\gamma}$ without further loss of generality as long as we treat the empirical distribution of income as a calibration target. This functional form simplifies the household problem as h_{it} drops out from the first-order condition and all households supply the same number of hours $n_{it} = N(w_t)$. Total effective labor input, $\int n_{it} h_{it} di$, is hence also equal to $N(w_t)$ because $\int h_{it} di = 1$.

B Data and Estimation

B.1 Data: Sources and transformations

B.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.

B.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.

B.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. 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.

C Supplementary Figures

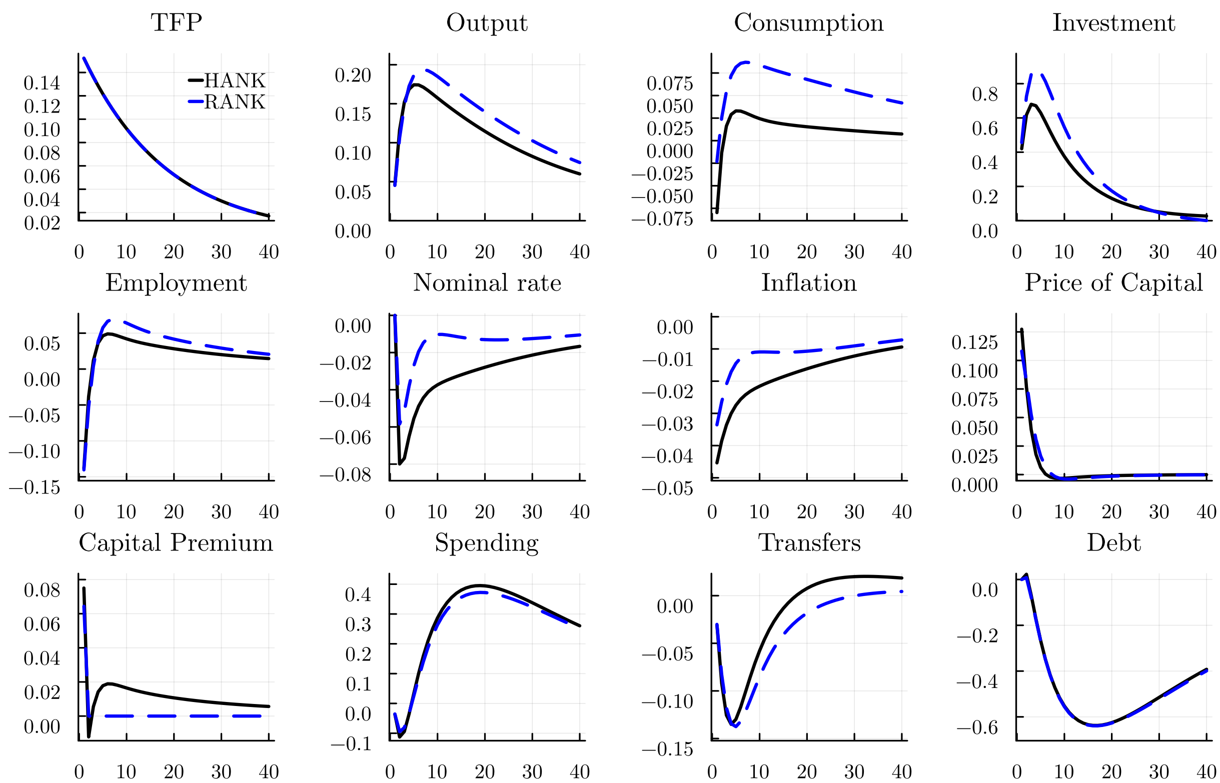

Impulse Responses to TFP shock in estimated HANK model and counterfactual RANK (under the same parameters).

Figure 7: Impulse Responses to TFP

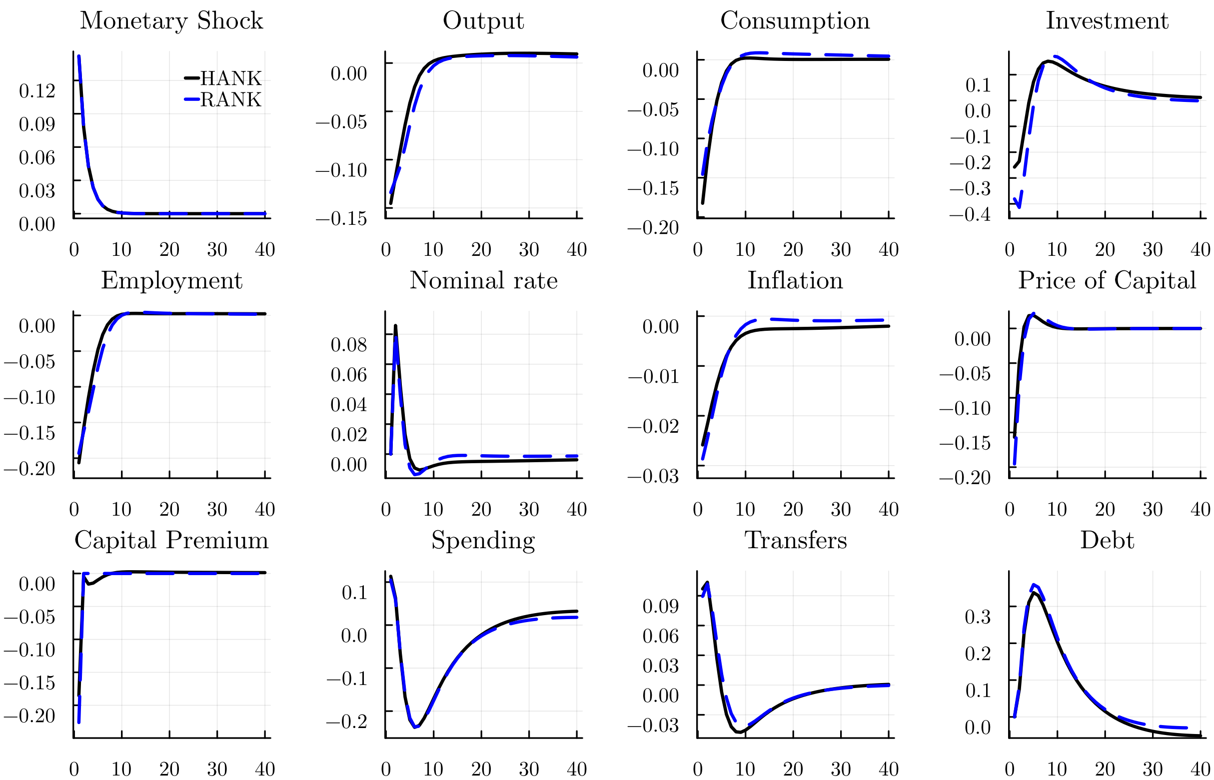

Impulse Responses to a monetary shock in estimated HANK model and counterfactual RANK (under the same parameters).

Figure 8: Impulse Responses to Monetary Policy

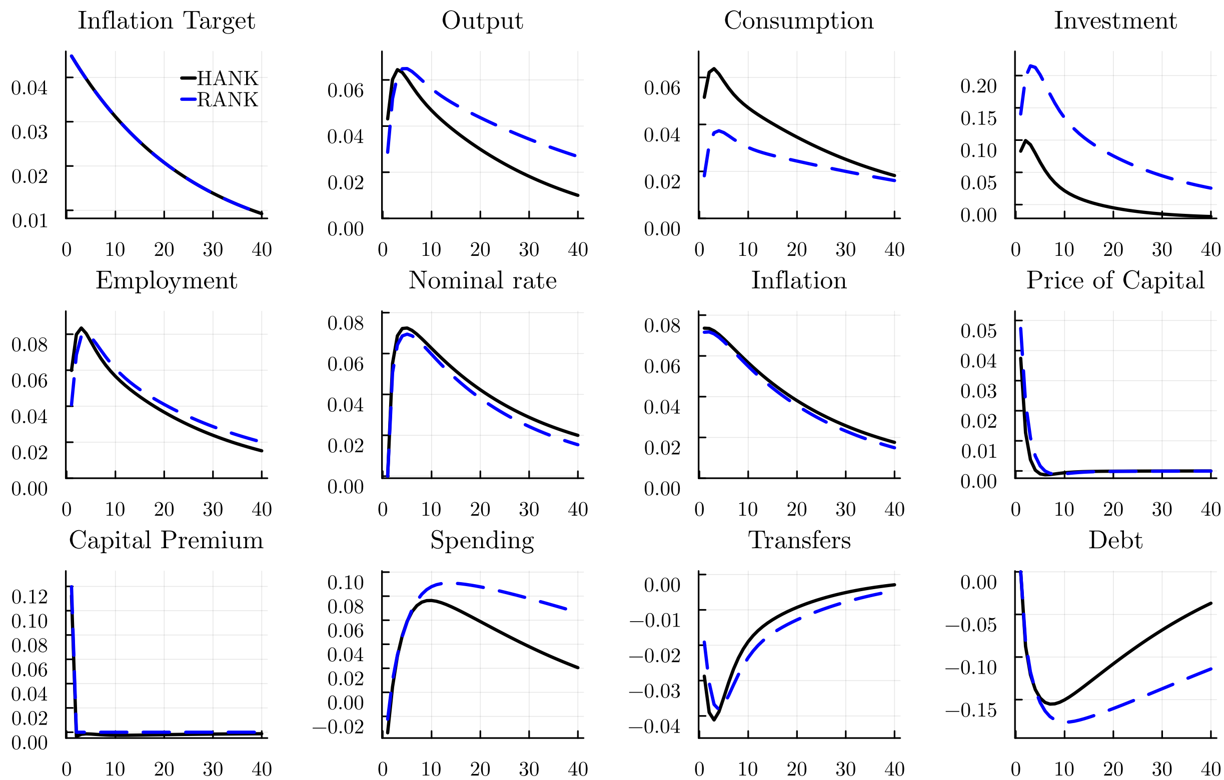

Impulse Responses to an inflation target shock in estimated HANK model and counterfactual RANK (under the same parameters).

Figure 9: Impulse Responses to Inflation Target