Learning, Confidence, and Business Cycles*

Cosmin Ilut
Duke University & NBER

Hikaru Saijo
UC Santa Cruz

February 2016

Abstract

We construct and estimate a heterogeneous-firm business cycle model where firms face Knightian uncertainty about their profitability and learn it through production. The cross-sectional mean of firm-level uncertainty is high in recessions because firms invest and hire less. The higher uncertainty reduces agents' confidence and further discourages economic activity. This feedback mechanism endogenously generates properties traditionally explained through additional shocks or rigidities: countercyclical labor and financial wedges, co-movement driven by demand shocks, and amplified and hump-shaped dynamics. We find that endogenous idiosyncratic confidence reduces the empirical role of standard rigidities and changes inference about sources of fluctuations and policy experiments.

1 Introduction

Is firms' confidence about their business conditions important for understanding aggregate fluctuations? To answer this question, we construct a heterogeneous-firm business cycle model in which firms face uncertainty about their own profitability and need to learn it through production. The learning process generates countercyclical uncertainty that feeds from, and back into, economic activity. We find that this feedback loop generates properties that make our model behave as a standard business cycle model with (i) countercyclical labor wedge and financial spread, (ii) positive co-movement of aggregate variables in response to either supply or demand shocks, and (iii) strong internal propagation with amplified and hump-shaped dynamics.

*Email addresses: Ilut cosmin.ilut@duke.edu, Saijo hsaijo@ucsc.edu. We would like to thank Yan Bai, Nick Bloom, Jesus Fernandez-Villaverde, Tatsuro Senga, Martin Schneider, Mathieu Taschereau-Dumouchel, Carl Walsh, as well as to seminar and conference participants at the "Ambiguity and its Implications in Finance and Macroeconomics" Workshop, Bank of England, NBER Summer Institute, Econometric Society World Congress, San Francisco Fed, Stanford, UC Santa Cruz, SED Annual Meeting, the SITE conference on Macroeconomics of Uncertainty and Volatility, and Wharton for helpful comments.

Traditionally, macroeconomists have employed a variety of shocks and frictions to produce these key business cycle patterns. For example, labor wedge shocks are commonly cited as the explanation for the measured countercyclical "wedge" between the marginal rate of substitution of consumption for labor and the marginal product of labor[1] and aggregate uncertainty shocks are often used to explain the countercyclical excess return of uncertain assets, such as capital, over the risk free rate.[2] Additional rigidities, such as sticky prices and wages, are usually required to break the Barro and King (1984) critique and make other types of shocks, besides productivity or intratemporal labor supply shocks, generate positive co-movements of macro aggregates. To generate persistent and hump-shaped dynamics, arbitrary frictions such as consumption habit and investment adjustment costs are often added to the model. That our model can generate these key properties without relying on additional shocks or rigidities suggests that endogenous uncertainty could be a parsimonious mechanism for business cycle models that would at least partially replace these features. The goal of this paper is to theoretically and quantitatively evaluate this hypothesis.

There are two key ingredients for our mechanism. First, learning occurs through production. In particular, by producing at a larger scale, the firm gets to see more signals about the business prospects of the firm. We introduce two unobservable shocks into a firm's production. The first shock is a standard, persistent productivity shock that affects the marginal return. The second is a transitory shock that does not scale up with the level of inputs. While what matters for optimal investing and hiring decisions is the realization of the productivity shock, the path of firm's output and inputs is not perfectly revealing about its productivity because it is confounded by the transitory disturbance.

In the model, the level of inputs endogenously determines the informativeness of output about the idiosyncratic productivity level. Intuitively, when a firm allocates less resources into production, its estimate about its persistent productivity is imprecise because the level of output is largely determined by the realization of the transitory shock. Conversely, its estimate becomes more accurate when it uses more resources because output mostly reflects the realization of productivity. This results in a procyclical signal-to-noise ratio at the firm level. It follows that recessions are periods of a high cross-sectional mean of firm-level uncertainty because firms *on average* invest and hire less.

The second key ingredient is that uncertainty should matter for the productive inputs. When firms do not observe their profitability, the optimal choice of inputs is naturally made under an uncertain evaluation of their return. We model uncertainty as ambiguity, or Knightian uncertainty. In contrast to models of risk, we can study the effect of Knigh-

[1] See Shimer (2009) and Chari et al. (2007) for evidence and discussion of labor wedges.
[2] See Cochrane (2011) for a review of the evidence on countercyclical excess returns.

tian uncertainty on inputs using linear methods, which in turn facilitates aggregation and estimation. Ambiguity aversion is described by the recursive multiple priors preferences, axiomatized in Epstein and Schneider (2003b), that capture agents' lack of confidence in probability assessments. This preference representation makes agents act *as if* they evaluate plans according to a worst case scenario drawn from a set of multiple beliefs. A wider set of beliefs corresponds to a loss of confidence.

In our model, ambiguity averse agents cope with estimation uncertainty by first estimating the underlying productivity process using a standard Kalman filter and then considering a set of probability distributions around the estimates. In particular, when facing a larger estimation uncertainty, the firm is less confident about the conditional mean of the underlying persistent profitability. The lower confidence makes the firm behave as if the worst-case mean becomes worse. This is simply a manifestation of precautionary behavior, which lowers the certainty equivalent of the return to production, but, compared to risk, has the advantage that it allows for first-order effects of uncertainty on decisions.

The two ingredients generate a feedback loop at the firm level: lower production leads to more estimation uncertainty, which in turn shrinks the optimal size of productive inputs. Importantly, this feedback arises from any shock that starts to move the economic activity. It may be for example a shock to productivity, to demand or a government spending change.

We show how countercyclical uncertainty shows up as 'excess volatility', or wedges, at the firm level. Specifically, in periods of low production, the estimation uncertainty is larger and this lowers the worst-case estimate of profitability. As a consequence, the firm lowers its demand for inputs. From the perspective of the econometrician that measures, under the true data generating process, equilibrium objects such as the marginal rate of substitution between consumption and labor, marginal product of labor or realized return on capital, these factor demands are excessively low. They can be instead rationalized as labor and investment wedges, or taxes on inputs, which, in a reduced form, seem higher in recessions.

Importantly, when uncertainty is modeled as ambiguity with multiple priors that differ in their mean, the law of large numbers imply that sets of beliefs converge, but that idiosyncratic uncertainty cannot be diversified away. Intuitively, the agents who do not know the distribution of firms' productivity (and hence profit) evaluate choices as if mean productivity is low. An agent who owns a portfolio of many firms then acts *as if* the entire portfolio has a low mean payoff. This means that the confidence about idiosyncratic conditions does not vanish in the aggregate.

The linearity of decision rule is crucial for a tractable aggregation. Indeed, once aggregated, these reduced-form taxes will survive as aggregate countercyclical taxes. The emergence of the labor wedge is particularly important for the business cycle dynamics.

Through it the model is able to generate positive co-movement in aggregate variables, such as consumption, investment and hours, not only out of productivity shocks, but in general out of any type of aggregate shock, including demand shocks.

We quantitatively evaluate the role of endogenous uncertainty using a calibrated version of the baseline model. Our main findings are as follows. First, we find that endogenous uncertainty is a powerful propagation mechanism. A positive shock that raises economic activity increases the level of confidence, which in turn further affects economic activity, leading to an amplified and hump-shaped impulse response.

Second, consistent with the data, our model generates countercyclical labor wedge and ex-post excess return. During recessions, ambiguity increases endogenously and thus leads to an unusually low equilibrium labor supply. The increase in ambiguity also makes capital less attractive to hold. Investors holding an ambiguous asset are thus compensated by the higher measured excess return. The countercyclical labor and financial wedges do not arise from separate labor supply or premia shocks but instead are generated by any underlying shock that moves economic activity.

Third, the model generates positive co-movements in response to demand shocks. Consider, for example, a positive shock to government spending. In standard RBC models, consumption sharply declines and hours worked rise due to the negative income effect. In our model, since firms produce more, they become more confident and are willing to hire and invest more. The higher confidence leads to a positive wealth effect for the representative agent, who owns the portfolio of firms. This positive effect can be strong enough to overturn the initial negative one and to make consumption actually rise.

Fourth, firms have an incentive to actively use their productive inputs to learn about their profitability. We show that this experimentation incentive has a first-order effect in our setup and we characterize its cyclical properties. We find that experimentation exerts a procyclical influence on aggregates in the short run, turning experimentation into an amplifying factor, and a countercyclical one in the medium term, where it dampens fluctuations.

Due to easy aggregation, we can estimate using likelihood based methods quantitative models with a large state space, such as Christiano et al. (2005). Our estimated model includes shocks to aggregate TFP and government spending, as well as a financial wedge shock, analyzed for example in Christiano et al. (2015). We use standard observables for US aggregate data: the growth rate of output, hours worked, investment and consumption, as well as inflation and nominal interest rate. In addition, we use the Baa corporate spread as an observable proxy, assumed to be measured with error, for the financial wedge shock.

The estimated model provides evidence that endogenous confidence changes significantly the inference on the role of rigidities and shocks driving fluctuations. In terms of rigidities, learning reduces the need of additional frictions for fitting the data. In particular, compared to the rational expectations (RE) version, the habit formation parameter is lowered by 25%, the investment adjustment cost becomes negligible, the average Calvo adjustment period of prices falls from 4 to 2 quarters, and for wages from 25 quarters to 1 quarter. The reason for these smaller estimated frictions is that the learning mechanism provides strong internal propagation and induces by itself co-movement in response to the financial wedge shock.

In terms of model comparison, adding endogenous uncertainty allows the model to fit the data significantly better. The reason is twofold. First, as described above, learning emerges as a parsimonious friction, as opposed to the array of other rigidities. Second, while in the RE version the observed spread is mostly accounted as measurement error, the model-implied spread with learning tracks well the empirical one. This happens because the observed countercyclical spread is now consistent with co-movement and thus the estimation prefers observed contractionary wedge shocks during recessions. At the same time, agents lose confidence during recessions, which in our model endogenously contributes further to the countercyclicality of the spread and improves the model fit.

In the estimated model, endogenous uncertainty changes the propagation mechanism to the point that the financial wedge shock, disciplined by the observed spread, becomes an important driver of fluctuations. This illustrates that endogenous confidence can make demand shocks a relevant source of fluctuations even in the absence of traditionally used rigidities. At the same time, generating time-variation in uncertainty from an endogenous mechanism means that policy matters for its evolution. We compare some counterfactual models in which a Taylor rule responds to the financial spread to underscore the importance of modeling a policy-variant uncertainty process.

The paper is organized as follows. In Section 2, we discuss relation to literature. In Section 3 we introduce our heterogeneous-firm quantitative model. In Section 4 we present the model's equilibrium characterization and solution. We discuss the main results based on a calibrated version in Section 5. In Section 6 we add additional rigidities to estimate a model using US aggregate data and also also provide some firm-level supportive evidence.

2 Relation to literature

Our paper is related to a rapidly growing literature on time-varying uncertainty. There are models that use exogenous increases in the volatility of firm-level productivity to generate recessions, such as Bloom (2009), Arellano et al. (2012), Bachmann and Bayer (2013), Christiano et al. (2014), Bloom et al. (2014) and Schaal (2015). Another strand of literature analyzes the effects of aggregate uncertainty shocks, such as Fernández-Villaverde et al. (2011), Basu and Bundick (2012), Bidder and Smith (2012), Born and Pfeifer (2014), Ilut and Schneider (2014), Fernández-Villaverde et al. (2015) and Leduc and Liu (2015). There are two main differences to this literature. First, we do not rely on additional real or nominal rigidities to generate contractionary effects of higher uncertainty and positive co-movement.[3] The reason is that in our model the level of input (such as labor supply) is chosen under imperfect information, and thus under uncertainty, about the underlying productivity process. Second, in our economy changes in uncertainty are endogenous and hence affect the propagation of any aggregate shocks and outcomes of policy experiments.

We differ from other papers in the literature on endogenous uncertainty and learning in business cycles, such as Veldkamp (2005), van Nieuwerburgh and Veldkamp (2006), Fajgelbaum et al. (2015), Ordoñez (2013), Saijo (2014), in two major aspects. One is that we focus attention on imperfect information at the idiosyncratic level. We are motivated by the fact that firm-level volatilities are empirically much larger than fluctuations at the aggregate. This fact suggests that uncertainty about idiosyncratic fundamentals may be larger and hence quantitatively more important than uncertainty about the macro-level process. In addition, because firms learn about their individual-specific characteristics, the impact of imperfect information is unlikely to be substantially affected by an introduction of market for information or a release of official statistics.[4]

Our formulation of uncertainty at the firm level has important modeling implications. First, the additive noise shock has reasonable interpretations at the firm-level. This shock may arise either as a noisy signal about firm demand or from aggregating productive units at the firm level.[5] Second, our competitive equilibrium is constrained Pareto optimal. Differently from recent work such as Fajgelbaum et al. (2015) or Ordoñez (2013), this is a model without information externalities since learning occurs at the individual firm level and not from observing the aggregate economy.[6] Third, in our model the firm has an incentive to experiment since it controls the action that produces information. We thus

[3]In Angeletos et al. (2014) co-movement is possible, out of correlated higher-order beliefs shocks, without additional rigidities because labor is assumed to be chosen before observing profitability. Arellano et al. (2012) also make a similar timing assumption and generate countercyclical labor wedge through an uncertainty shock. In our model the endogenous propagation mechanism makes any type of shock, including various types of demand shocks, potentially generate labor and financial wedges as well as co-movement.

[4]At the same time, our approach can incorporate learning not only about idiosyncratic but also aggregate conditions, and as such, our results can be viewed as a lower bound on the overall business cycle effects of endogenous uncertainty.

[5]Applying this logic to additive noise about aggregate productivity requires instead additional assumptions on the observability of consumption bundles or the size of the units that are aggregated (see van Nieuwerburgh and Veldkamp (2006)).

[6]For example, the increased economic activity, and the associated average reduction in uncertainty, produced by a government spending increase can not be welfare increasing in our model.

further contribute to the literature on learning by showing how experimentation at the firm level affects the equilibrium dynamics of a business cycle model.

The second major difference from the learning literature is that we show that endogenous confidence not only leads to more persistence and amplification, as in Fajgelbaum et al. (2015) and Saijo (2014), but that in the process it produces countercyclical endogenous labor and financial wedges. As a consequence, we contribute to the business cycle literature by proposing a mechanism that alters the Barro and King (1984) critique of positive co-movement of aggregates.[7] In addition, this co-movement may now be accompanied by countercyclical premia, which is a key insight of the asset pricing literature.

We also build on the literature on recursive multiple priors, introduced by Gilboa and Schmeidler (1989), and extended to intertemporal choice by Epstein and Wang (1994) and Epstein and Schneider (2003b). Ilut and Schneider (2014) and Bianchi et al. (2014) study tractable business cycle model where time-variation in confidence about aggregate conditions arises from exogenous ambiguity or volatility shocks. The present paper instead considers time-variation in confidence about idiosyncratic shocks that emerges entirely endogenously and which thus changes the propagation mechanism of standard shocks.

Methodologically, this paper develops tractable linear methods to study the effects of such endogenous uncertainty about in a heterogeneous-firm model. An important property of the model is that this idiosyncratic uncertainty does not vanish in the aggregate. As such, this paper is connected to the decision-theoretical literature on the law of large numbers developed in Marinacci (1999) and Epstein and Schneider (2003a).

3 The model

Our baseline model is a real business cycle model augmented with two key features: Agents are ambiguity averse and face Knightian uncertainty about the firm-level profitability processes. After presenting the environment, we discuss in detail the information friction that gives rise to equilibrium fluctuations in confidence.

[7]Recent line of attacks to break Barro and King (1984) critique include: heterogeneity in labor supply and consumption across employed and non-employed (Eusepi and Preston (2015)), variable capacity utilization together with a large preference complementarity between consumption and hours (Jaimovich and Rebelo (2009)), and the large literature on countercyclical markups through nominal rigidities (such as in the quantitative models of Justiniano et al. (2010) or Christiano et al. (2014)).

3.1 Environment

Households

There is a representative household that has recursive multiple priors utility. Collect the exogenous state variables, to be described later, in a vector $s_t \in S$. A household consumption plan C gives, for every history s^t, the consumption of the final good $C_t(s^t)$ and the amount of hours worked $H_t(s^t)$. For a given consumption plan C, utility is defined recursively by

$$U_t(C; s^t) = \ln C_t - \varphi \frac{H_t^{1+\eta}}{1+\eta} + \beta \min_{p \in \mathcal{P}_t(s^t)} E^p[U_{t+1}(C; s^t, s_{t+1})], \tag{3.1}$$

where β is the subjective discount factor, φ is a scaling parameter that determines hours worked, and η is the inverse of Frisch labor supply elasticity. $\mathcal{P}_t(s^t)$ is a set of conditional probabilities about next period's state $s_{t+1} \in S_{t+1}$. We specify the evolution of this set in section 3.2.2.[8]

The household maximizes utility subject to the budget constraint

$$C_t + B_t + \int P^e_{l,t}\theta_{l,t} dl \leq W_t H_t + R_{t-1}B_{t-1} + \int (D_{l,t} + P^e_{l,t})\theta_{l,t-1} dl + T_t,$$

where B_t is the one-period risk-free bond, W_t is the real wage, R_t is the risk-free interest rate, and T_t is a transfer. $D_{l,t}$ and $P^e_{l,t}$ are the dividend payout and the price of a unit of share $\theta_{l,t}$ of firm l, respectively.

Firms

There is a continuum of firms, indexed by $l \in [0, 1]$, which act in a perfectly competitive manner. They use capital $K_{l,t-1}$, which is utilized at rate $U_{l,t}$, and hire labor $H_{l,t}$ to produce goods $Y_{l,t}$ according to the production function

$$Y_{l,t} = A_t\{z_{l,t}(U_{l,t}K_{l,t-1})^\alpha(\gamma^t H_{l,t})^{1-\alpha} + \gamma^t \nu_{l,t}\}, \qquad \nu_{l,t} \sim N(0, \sigma^2_\nu), \tag{3.2}$$

where γ is the growth rate of labor augmenting technical progress. The scale of the idiosyncratic i.i.d. shock $\nu_{l,t}$ grows at rate γ, which ensures that the shock does not vanish

[8] The recursive formulation ensures that preferences are dynamically consistent. Details and axiomatic foundations are in Epstein and Schneider (2003b). If the set is singleton we obtain standard separable log utility with those conditional beliefs. Otherwise, agents are not willing to integrate over the beliefs and narrow down the set to a singleton. In response, households take a cautious approach to decision making and act *as if* the true data generating process is given by the worst-case conditional belief.

along the balanced growth path. $z_{l,t}$ is an idiosyncratic technology shock that follows[9]

$$z_{l,t} = (1-\rho_z)\bar{z} + \rho_z z_{l,t-1} + \epsilon_{z,l,t}, \qquad \epsilon_{z,l,t} \sim N(0, \sigma_z^2), \tag{3.3}$$

and A_t is an aggregate technology shock that follows

$$\ln A_t = \rho_A \ln A_{t-1} + \epsilon_{A,t}, \qquad \epsilon_{A,t} \sim N(0, \sigma_A^2).$$

Firms cannot directly observe the realizations of idiosyncratic shocks $z_{l,t}$ and $\nu_{l,t}$. This informational assumption leads to a non-invertibility problem: Firms cannot tell whether an unexpectedly high realization of output is due to an increase in individual technology or a favorable transitory disturbance. Instead, they need to form the estimates using all other available information, including the path of output and inputs. In contrast, they perfectly observe the aggregate shock A_t.

Firms choose $\{U_{l,t}, K_{l,t}, H_{l,t}, I_{l,t}\}$ to maximize shareholder value

$$E_0^* \sum_{t=0}^{\infty} M_0^t D_{l,t}, \tag{3.4}$$

where we use E_0^* to denote expectation under the worst case probability. Random variables M_0^t denote prices of t-period ahead contingent claims based on conditional worst case probabilities and is given by

$$M_0^t = \beta^t \lambda_t,$$

where λ_t is the marginal utility of consumption at time t by the representative household. $D_{l,t}$ is the dividend payout given by

$$D_{l,t} = Y_{l,t} - W_t H_{l,t} - I_{l,t} - a(U_{l,t}) K_{l,t-1},$$

where $I_{l,t}$ is investment and $a(U_{l,t})$ is the cost of utilization.[10] Their capital stock follows the law of motion

$$K_{l,t} = (1-\delta) K_{l,t-1} + I_{l,t}.$$

[9]We assume that the idiosyncratic shock $z_{l,t}$ follows a normal process, a technical assumption that is useful in solving the learning problem because it makes the information friction linear.

[10]We specify: $a(U) = 0.5\chi_1\chi_2 U^2 + \chi_2(1-\chi_1)U + \chi_2(0.5\chi_1 - 1)$, where χ_1 and χ_2 are parameters. We set χ_2 so that the steady-state utilization rate is one. The cost $a(U)$ is increasing in the rate of utilization and χ_1 determines the degree of the convexity of utilization costs. In a linearized equilibrium, the dynamics are controlled by the value of χ_1.

Interpretation of the additive shock

We generate a procyclical signal-to-noise ratio by adding an unobservable additive shock $\nu_{l,t}$ to the production function. In Appendix 8.1, we offer an additional interpretation of this shock based on noisy demand signals, in which firms learn more about the demand of their goods when they produce and sell more. In that version of the model, firms are subject to unobservable idiosyncratic shocks to the weight attached to their goods in the CES aggregator for final goods. It is natural to interpret the shock as a shock to the quality or demand of goods produced by an individual firm l. The additive shock, which vanishes in the aggregate due to the law of large numbers, is replaced with an i.i.d. observation error of the underlying idiosyncratic shock; agents observe noisy signals about the demand, whose precision is increasing in the level of individual production.

Market clearing and resource constraint

We impose the market clearing conditions for the labor market and the bond market:

$$H_t = \int_0^1 H_{l,t} dl, \qquad B_t = 0.$$

The resource constraint is given by

$$C_t + I_t + G_t + \int_0^1 a(U_{l,t}) K_{l,t-1} dl = Y_t$$

where $I_t \equiv \int_0^1 I_{l,t} dl$ and G_t is the government spending. We assume that the government balances budget each period ($G_t = -T_t$). We also assume $G_t + \int_0^1 a(U_{l,t}) K_{l,t-1} dl = g_t Y_t$ where g_t follows

$$\ln g_t = (1 - \rho_g)\bar{g} + \rho_g \ln g_{t-1} + \epsilon_{g,t}, \qquad \epsilon_{g,t} \sim N(0, \sigma_g^2).$$

Timing

The timing of events within a period t is as follows:

Stage 1 : Pre-production stage

- Agents observe the realization of aggregate shocks (A_t and g_t).
- Given forecasts about the idiosyncratic technology and its associated worst-case scenario, firms make utilization decisions and hire labor ($U_{l,t}$ and $H_{l,t}$). The household supply labor H_t and the labor market clears at the wage rate W_t.

Stage 2 : Post-production stage

- Idiosyncratic shocks $z_{l,t}$ and $\nu_{l,t}$ realize (but are unobservable) and production takes place.
- Given output and input, firms update estimates about their idiosyncratic technology and use it to form forecasts for production next period.
- Firms make investment $I_{l,t}$ and pay out dividends $D_{l,t}$. The household makes consumption and asset purchase decisions (C_t, B_t, and $\theta_{l,t}$).

3.2 Uncertainty and preferences

3.2.1 Learning about idiosyncratic productivity

Firms form estimates about the idiosyncratic shock $z_{l,t}$ from the observables. Since the problem is linear and Gaussian, Bayesian updating using Kalman filter is optimal from the statistical perspective of minimizing the mean square error of the estimates. To ease notation, we set the trend growth rate γ to zero.[11] We denote $F_{l,t} \equiv (U_{l,t}K_{l,t-1})^{\alpha}H_{l,t}^{1-\alpha}$. After production at period t, the measurement equation of the Kalman filter is given by

$$Y_{l,t}/A_t = F_{l,t}z_{l,t} + \nu_{l,t},$$

and the transition equation is given by

$$z_{l,t} = (1-\rho_z)\bar{z} + \rho_z z_{l,t-1} + \epsilon_{z,l,t}. \tag{3.5}$$

Note that, unlike the standard time-invariant Kalman filter, the coefficient in the measurement equation, $F_{l,t}$, is time-varying.[12] The key property of our filtering system is that the signal-to-noise ratio is procyclical, which follows from the fact that the input $F_{l,t}$ is procyclical. The flip side implication of this property is that uncertainty is countercyclical: the posterior variance of idiosyncratic technology $z_{l,t}$ rises during recessions. Intuitively, when a firm puts less resources into production, its estimate about its productivity is imprecise because the level of output is largely determined by the realization of the transitory shock. Conversely, its estimate is accurate when it uses more resources because output mostly reflects the realization of productivity.

[11]The learning problem of the model with positive growth is provided in Appendix 8.2 along with other equilibrium conditions.

[12]However, also note that, after production, the coefficient is pre-determined, which allows us to use Kalman filter.

To characterize the filtering problem, we start by deriving the one-step-ahead prediction from the period $t-1$ estimate $\tilde{z}_{l,t-1|t-1}$ and its associated error variance $\Sigma_{l,t-1|t-1}$. We have

$$\begin{aligned}\tilde{z}_{l,t|t-1} &= (1-\rho_z)\bar{z} + \rho_z \tilde{z}_{l,t-1|t-1},\\ \Sigma_{l,t|t-1} &= \rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2.\end{aligned}$$

Then, given observables (output $Y_{l,t}$ and aggregate productivity A_t) firms update their estimates according to

$$\tilde{z}_{l,t|t} = \tilde{z}_{l,t|t-1} + Gain_{l,t}(Y_{l,t}/A_t - \tilde{z}_{l,t|t-1}F_{l,t}), \tag{3.6}$$

where $Gain_{l,t}$ is the Kalman gain and is given by

$$Gain_{l,t} = \left[\frac{F_{l,t}^2 \Sigma_{l,t|t-1}}{F_{l,t}^2 \Sigma_{l,t|t-1} + \sigma_\nu^2}\right] F_{l,t}^{-1}. \tag{3.7}$$

The updating rule for variance is

$$\Sigma_{l,t|t} = \left[\frac{\sigma_\nu^2}{F_{l,t}^2 \Sigma_{l,t|t-1} + \sigma_\nu^2}\right] \Sigma_{l,t|t-1}. \tag{3.8}$$

Intuitively, the error variance is increasing in the un-informativeness of the observation, which is the variance of noise divided by the total variance. We can see that, holding $\Sigma_{l,t|t-1}$ constant, the posterior variance $\Sigma_{l,t|t}$ increases as input $F_{l,t}$ decreases.

The dynamics of the idiosyncratic technology $z_{l,t}$ according to the Kalman filter can thus be described as

$$z_{l,t+1} = (1-\rho_z)\bar{z} + \rho_z(\tilde{z}_{l,t|t} + u_{l,t}) + \epsilon_{z,l,t+1}, \tag{3.9}$$

where $u_{l,t}$ is the estimation error of $z_{l,t}$ and $u_{l,t} \sim N(0, \Sigma_{l,t|t})$.

3.2.2 Ambiguity

We assume that the representative agent that owns the firms perceives ambiguity about the idiosyncratic shock $z_{l,t}$. The agent uses observed data to learn about the hidden technology by using the Kalman filter to obtain a benchmark probability distribution. Ambiguity is modeled as a one-step ahead set of conditional beliefs $\mathcal{P}_t(s^t)$ in (3.1), which here consists of alternative probability distribution surrounding the benchmark controlled by a bound on the relative entropy distance. Thus, our ambiguity-averse agents continue to use the ordinary Kalman filter to estimate the latent technology and evaluate plans according to the worst-case means that are implied by the posterior estimates.

The agent is not confident in the benchmark Kalman filter estimate $\tilde{z}_{l,t}$ in (3.9) and considers a set of probability distributions, of the form

$$z_{l,t+1} = (1-\rho_z)\bar{z} + \rho_z \tilde{z}_{l,t|t} + \mu^*_{l,t+1} + \rho_z u_{l,t} + \epsilon_{z,l,t+1}, \tag{3.10}$$

where $\mu_{l,t+1} \in [-a_{l,t}, a_{l,t}]$. From the perspective of the agent, a change in posterior variance translates into a change in uncertainty about the one-step-ahead realization of technology. As in Bianchi et al. (2014), the change in uncertainty, in turn, affects the set of possible $\mu_{l,t+1}$ and thus the worst-case mean.

More precisely, agents only consider the conditional means $\mu^*_{l,t+1}$ that are sufficiently close to the long run average of zero in the sense of relative entropy:

$$\frac{(\mu^*_{l,t+1})^2}{2\rho_z^2 \Sigma_{l,t|t}} \leq \frac{1}{2}\eta_a^2, \tag{3.11}$$

where the left hand side is the relative entropy between two normal distributions that share the same variance $\rho_z^2\Sigma_{l,t|t}$, but have different means ($\mu^*_{l,t+1}$ and zero), and η_a is a parameter that controls the size of the entropy constraint. Agents compare the normal distributions with variance $\rho_z^2\Sigma_{l,t|t}$ because we assume that they only treat the estimation error u_t as ambiguous. They are fully confident in the law of motion in equation (3.5) and treat the technology shock $\epsilon_{z,l,t+1}$ as risk. The relative entropy can be thought of as a measure of distance between the two distributions. When uncertainty $\Sigma_{l,t|t}$ is high, it becomes difficult to distinguish between different processes. As a result, agents become less confident and contemplate wider sets $\mu_{l,t+1}$ of conditional probabilities.

The worst-case belief can be easily solved for at the equilibrium consumption plan: the worst-case expected idiosyncratic productivity is low. In particular, equation (3.11) implies that the worst-case mean is given by

$$-a_{l,t} = -\eta_a \rho_z \sqrt{\Sigma_{l,t|t}}. \tag{3.12}$$

Thus, the agent's cautious behavior faced with the set $\mathcal{P}_t(s^t)$ of beliefs results in acting *as if* the conditional mean of each firm's idiosyncratic technology is given by the worst-case mean in (3.12). We denote conditional moments under these worst case belief by stars.

We make two additional remarks. The first concerns the role of experimentation. In Bayesian decision making, experimentation is valuable because it raises expected utility by improving posterior precision. Ambiguity-averse agents also value experimentation since it affects utility by tightening the set of conditional probability considered. In our model, firms take into account the impact of the level of input on worst-case mean when they make decisions. Although we allow active learning by firms, our model can still be solved using standard linear methods. When we present our quantitative results, we assess the contribution of experimentation by comparing our baseline results with those under passive learning, i.e. where agents do not actively experiment.

The second remark is to note that the shareholder value, under which firms take optimal decisions in equation (3.4), depends on the worst case expectations E_0^*. This is because state prices reflect the representative household's ambiguity. An important feature is that, unlike the case of risk, the idiosyncratic uncertainty that shows up in these state prices does not vanish under diversification. Uncertainty affects ambiguity-averse household's utility by lowering the worst-case mean and hence the household acts as if the mean of each individual firm's technology is lower. As a result, ambiguity is not diversified away in the aggregate and uncertainty lowers the mean of the expected aggregate technology.[13]

4 Equilibrium characterization and solution

We start by discussing the recursive representation of the model. We then build on the framework to describe the solution method we use to solve for the equilibrium law of motion. Finally, we characterize the endogenous wedges that arise from equilibrium fluctuations in uncertainty.

4.1 Recursive competitive equilibrium

As in Angeletos et al. (2014), it is useful to divide the agents' problem into two stages; stage 1 (pre-production stage) and stage 2 (post-production stage). To ease exposition, we abstract from utilization momentarily. We collect exogenous aggregate state variables (such as aggregate TFP) in a vector X with a cumulative transition function $F(X'|X)$. The endogenous aggregate state is the distribution of firm-level variables. A firm's type is identified by the posterior mean estimate of productivity $\tilde{z}_l$, the posterior variance Σ_l, and its capital stock K_l. The worst-case TFP is not included because it is implied by the posterior mean and variance. We denote the cross-sectional distribution of firms' type by ξ_1 and ξ_2. ξ_1 is a stage 1 distribution over $(\tilde{z}_l, \Sigma_l, K_l)$ and ξ_2 is a stage 2 distribution over $(\tilde{z}_l', \Sigma_l', K_l)$.

[13] See Marinacci (1999) or Epstein and Schneider (2003a) for formal treatments of the law of large numbers for i.i.d. ambiguous random variables. There they show that sample averages must (almost surely) lie in an interval bounded by the highest and lowest possible mean, and these bounds are tight in the sense that convergence to a narrower interval does not occur.

ξ_1', in turn, is a distribution over $(\tilde{z}_l', \Sigma_l', K_l')$ at stage 1 in the next period.[14]

First, consider the household's problem. The household's wealth can be summarized by a portfolio $\overrightarrow{\theta_l}$ which consists of share θ_l for each firm and the risk-less bond holdings B. We use V_1^h and V_2^h to denote the household's value function at stage 1 and stage 2, respectively. We use m to summarize the income available to the household at stage 2. The household's problem at stage 1 is

$$\begin{aligned} V_1^h(\overrightarrow{\theta_l}, B; \xi_1, X) &= \max_H \left\{ -\varphi \frac{H^{1+\eta}}{1+\eta} + E^*[V_2^h(\hat{m}; \hat{\xi}_2, X)] \right\} \\ \text{s.t.} \quad \hat{m} &= WH + RB + \int (\hat{D}_l + \hat{P}_l)\theta_l dl \end{aligned} \tag{4.1}$$

where we momentarily use the *hat* symbol to indicate random variables that will be resolved at stage 2. The household's problem at stage 2 is

$$\begin{aligned} V_2^h(m; \xi_2, X) &= \max_{C, \overrightarrow{\theta_l}', B'} \left\{ \ln C + \beta \int V_1^h(\overrightarrow{\theta_l}', B'; \xi_1', X') dF(X'|X) \right\} \\ \text{s.t.} \quad & C + B' + \int P_l \theta_l' dl \leq m \\ & \xi_1' = \Gamma(\xi_2, X) \end{aligned} \tag{4.2}$$

In problem (4.1), households choose labor supply based on the worst-case stage 2 value (recall that we use E^* to denote worst-case conditional expectations). The problem (4.2), in turn, describes the household's consumption and asset allocation problem given the realization of income and aggregate states. In particular, they take as given the law of motion of the next period's distribution $\xi_1' = \Gamma(\xi_2, X)$, which in equilibrium is consistent with the firm's policy function. Importantly, in contrast to the stage 2 problem, a law of motion that describes the evolution of ξ_2 from (ξ_1, X) is absent in the stage 1 problem. Indeed, if there is no ambiguity in the model, agents take as given the law of motion $\xi_2 = \Upsilon(\xi_1, X)$, which in equilibrium is consistent with the firm's policy function and the true data generating process of the firm-level TFP. Since agents are ambiguous about each firm's TFP process, they cannot settle on a single law of motion about the distribution of firms. Finally, the continuation value at stage 2 is governed by the transition density of aggregate exogenous states X.

Next, consider the firms' problem. We use v_1^f and v_2^f to denote the firm's value function

[14] See also Senga (2015) for a recursive representation of an imperfect information heterogeneous-firm model with time-varying uncertainty.

at stage 1 and stage 2, respectively. Firm l's problem at stage 1 is

$$\begin{aligned} v_1^f(\tilde{z}_l, \Sigma_l, K_l; \xi_1, X) &= \max_{H_l} E^*[v_2^f(\hat{\tilde{z}}_l', \Sigma_l', K_l; \hat{\xi}_2, X)] \\ &\text{s.t.} \quad \text{Updating rules (3.6) and (3.8)} \end{aligned} \tag{4.3}$$

and firm l's problem at stage 2 is

$$\begin{aligned} v_2^f(\tilde{z}_l', \Sigma_l', K_l; \xi_2, X) = \max_{I_l} &\left\{ \lambda(Y_l - WH_l - I_l) + \beta \int v_1^f(\tilde{z}_l', \Sigma_l', K_l'; \xi_1', X') dF(X'|X) \right\} \\ \text{s.t.} \quad & K_l' = (1-\delta)K_l + I_l \\ & \xi_1' = \Gamma(\xi_2, X) \end{aligned} \tag{4.4}$$

where we simplify the exposition by expressing a firm's value in terms of the marginal utility λ of the representative household. Similar to the household's problem, a firm's problem at stage 1 is to choose the labor demand so as to maximize the worst-case stage 2 value. Note that the posterior mean $\tilde{z}_l'$ will be determined by the realization of output Y_l at stage 2 while the posterior variance Σ_l' is determined by Σ_l and the input level at stage 1. In problem (4.4), the firm then chooses investment taking as given the realization of output and the updated estimates of its productivity. Note that, as in the household's problem, firms take as given the (equilibrium) law of motion of the distribution of firms in the stage 2 problem but not in the stage 1 problem.

The discussion above highlights one of the key features of our model; the level of labor input is chosen before the realization of firm-level productivity and that this timing arises naturally from imperfect information about the underlying productivity process. This labor-in-advance feature allows us to circumvent the Barro and King (1984) critique and hence generate feedback effects of time-varying uncertainty consistent with business cycle co-movement without additional rigidities.

We conclude this subsection by providing a brief definition of the recursive competitive equilibrium of our model. The recursive competitive equilibrium is a collection of value functions, policy functions, and prices such that

1. Households and firms optimize; (4.1) – (4.4).

2. The labor market, goods market, and asset markets clear.

3. The law of motion $\xi_1' = \Gamma(\xi_2, X)$ is induced by the firms' policy function $I_l(\tilde{z}_l', \Sigma_l', K_l; \xi_2, X)$.

4.2 Log-linearized solution

We solve for the equilibrium law of motion using standard log-linear methods. This is possible for two reasons. First, since the filtering problem firms face is linear, the law of motion of the posterior variance can be characterized analytically (Saijo (2014)). Because the level of inputs has first-order effects on the level of posterior variance, linearization captures the impact of economic activity on confidence. Second, we consider ambiguity about the mean and hence the feedback from confidence to economic activity can be also approximated by linearization. In turn, log-linear decision rules facilitate aggregation because the cross-sectional mean becomes a sufficient statistic for tracking aggregate dynamics.

We follow Ilut and Schneider (2014) and solve for the equilibrium law of motion using a guess-and-verify approach:

(a) guess the worst case beliefs p^0.

(b) solve the model assuming that agents have agents have expected utility and beliefs p^0.

(c) compute the agent's value function V.

(d) verify that the guess p^0 indeed achieves the minima.

In what follows we explain step (b) by deriving log-linearized expressions for the expected worst-case output at stage 1 and the realized output at stage 2.[15] We use the example to illustrate that uncertainty about the firm-level TFP has a first-order effect at the aggregate.

We first find the worst-case steady state by evaluating a deterministic version of the filtering problem and standard first-order conditions under the guessed worst-case belief. Potential complications arise because the worst-case TFP depends on the level of economic activity. Since the worst-case TFP, in turn, determines the level of economic activity, there could be multiple steady states. We circumvent this multiplicity by treating the posterior variance of the level of idiosyncratic TFP as a parameter and by focusing on the steady state that is implied by that posterior variance.

Next, we log-linearize the model around the worst-case steady state. To do this, we first log-linearize the expected worst-case output of individual firm l at stage 1:

$$E_t^* \hat{Y}_{l,t}^0 = \hat{Z}_t^0 + E_t^* \hat{z}_{l,t}^0 + \hat{F}_{l,t}^0, \tag{4.5}$$

and the realized output of individual firm l at stage 2:

$$\hat{Y}_{l,t}^0 = \hat{Z}_t^0 + \hat{z}_{l,t}^0 + \hat{F}_{l,t}^0, \tag{4.6}$$

[15] We provide a general description of the procedure in Appendix 8.3.

where we use $\hat{x}_t^0 = x_t - \bar{x}^0$ to denote log-deviations from the worst-case steady state and set the trend growth rate γ to zero to ease notation. The worst-case individual output (4.5) is the sum of three components: the current level of aggregate TFP, the worst-case individual TFP, and the input level. The realized individual output (4.6), in turn, is the sum of aggregate TFP, the *realized* individual TFP, and the input level.

We then aggregate the log-linearized individual conditions (4.5) and (4.6) to obtain the cross-sectional mean of worst-case individual output:

$$E_t^* \hat{Y}_t^0 = \hat{A}_t^0 + E_t^* \hat{z}_t^0 + \hat{F}_t^0, \tag{4.7}$$

and the cross-sectional mean of realized individual output:

$$\hat{Y}_t^0 = \hat{A}_t^0 + \hat{z}_t^0 + \hat{F}_t^0, \tag{4.8}$$

where we simply eliminate subscript l to denote the cross-sectional mean, i.e., $\hat{x}_t^0 \equiv \int_0^1 \hat{x}_{l,t}^0 dl$.

So far we have characterized the dynamics of output under the worst-case scenario. Our final step is to characterize the dynamics under the true data generating process (DGP). To do this, we feed in the cross-sectional mean of individual TFP, which is constant under the true DGP, into (4.7) and (4.8). Using (4.7), the cross-sectional mean of worst-case output is given by

$$E_t^* \hat{Y}_t = \hat{A}_t + E_t^* \hat{z}_t + \hat{F}_t, \tag{4.9}$$

where we use $\hat{x}_t = x_t - \bar{x}$ to denote log-deviations from the steady-state under the true DGP. Using (4.8), the realized aggregate output is given by

$$\hat{Y}_t = \hat{A}_t + \hat{F}_t, \tag{4.10}$$

where we used $\hat{z}_t = 0$ under the true DGP. Importantly, $E_t^* \hat{z}_t$ in (4.10) is not necessarily zero outside the steady state. To see this, combine (3.10) and (3.12) and log-linearize to obtain an expression for $E_t^* \hat{z}_{l,t}$:

$$E_t^* \hat{z}_{l,t} = \varepsilon_{z,z} \hat{\hat{z}}_{l,t-1|t-1} - \varepsilon_{z,\Sigma} \hat{\Sigma}_{l,t-1|t-1}. \tag{4.11}$$

From (3.8), the posterior variance is negatively related to the level of input F:

$$\hat{\Sigma}_{l,t-1|t-1} = \varepsilon_{\Sigma,\Sigma} \hat{\Sigma}_{l,t-2|t-2} - \varepsilon_{\Sigma,F} \hat{F}_{l,t-1}, \tag{4.12}$$

The elasticities $\varepsilon_{z,z}$, $\varepsilon_{z,\Sigma}$, $\varepsilon_{\Sigma,\Sigma}$, and $\varepsilon_{\Sigma,F}$ are functions of structural parameters and are all positive. We combine (4.11) and (4.12) to obtain

$$E_t^* \hat{z}_{l,t} = \varepsilon_{z,z} \hat{\hat{z}}_{l,t-1|t-1} - \varepsilon_{z,\Sigma} \varepsilon_{\Sigma,\Sigma} \hat{\Sigma}_{l,t-2|t-2} + \varepsilon_{z,\Sigma} \varepsilon_{\Sigma,F} \hat{F}_{l,t-1}. \tag{4.13}$$

Finally, we aggregate (4.13) across all firms:

$$E_t^* \hat{z}_t = -\varepsilon_{z,\Sigma} \varepsilon_{\Sigma,\Sigma} \hat{\Sigma}_{t-2|t-2} + \varepsilon_{z,\Sigma} \varepsilon_{\Sigma,F} \hat{F}_{t-1}, \tag{4.14}$$

where we used $\int_0^1 \hat{\hat{z}}_{l,t-1|t-1} dl = 0.$[16]

Notice again that the worst-case conditional cross-sectional mean simply aggregates linearly the worst-case conditional mean, $-a_{l,t}$, of each firm. Since the firm-specific worst-case means are a function of idiosyncratic uncertainty, which in turn depend on the firms' scale, equation (4.14) shows that the average level of economic activity, $\hat{F}_{t-1}$, has a first-order effect on the cross-sectional average of the worst-case mean. For example, this means that during recessions, firms on average produce less, which leads to lower confidence about their firm-level TFP. This endogenous reduction in confidence further reduces equilibrium hours worked and other economic activity.

4.3 Wedges from uncertainty

Co-movement and the labor wedge

Endogenous uncertainty leads to co-movement and a countercyclical labor wedge. This can be analyzed by considering the optimal labor tradeoff of equating the marginal cost to the expected marginal benefit under the worst-case belief E_t^*

$$\varphi H_t^{\eta} = E_t^* \left(\lambda_t MPL_t \right) \tag{4.15}$$

In the standard model, there is no expectation on the right-hand side. As emphasized by Barro and King (1984), there hours and consumption move in opposite direction unless there is a TFP or a labor supply shock (a shock to 'φ').

Instead, in our model, there can be such co-movement. Suppose that there is a period of low confidence. From the negative wealth effect there is a low consumption, so the standard effect would be to see high labor supply as a result of the high marginal utility of consumption λ_t. However, because the firm chooses hours *as if* productivity is low, there is a counter substitution incentive for hours to be low.

[16]This follows from aggregating the log-linearized version of (3.6) and evaluating the equation under the true DGP. Intuitively, since the cross-sectional mean of idiosyncratic TFP is constant, the cross-sectional mean of the Kalman posterior mean estimate is a constant as well.

To see how the model generates countercyclical labor wedge, note that an increase in ambiguity due to a reduction in labor supply looks, from the perspective of an econometrician, like an increase in the labor income tax. The labor wedge can now be easily explained by implicitly defining the labor tax τ_t as

$$\varphi H_t^\eta = (1 - \tau_t)\lambda_t MPL_t$$

Using the optimality condition in (4.15), the labor tax is

$$\tau_t = 1 - \frac{E_t^*\left(\lambda_t MPL_t\right)}{\lambda_t MPL_t} \tag{4.16}$$

Consider first the linear rational expectations case. There the role of idiosyncratic uncertainty disappears and the labor tax in equation (4.16) is constant and equal to zero. The reason is that our timing assumption that labor is chosen after the aggregate shocks are realized and observed at the beginning of the period makes the optimality condition in (4.15) take the usual form of an intratemporal labor decision.[17]

In our model, the role of idiosyncratic uncertainty does not vanish and instead it shows up in the *as if* expected return to working, formed under the worst-case belief E_t^*. Thus, even if labor is chosen after aggregate shocks are realized and observed, the average idiosyncratic uncertainty has a first-order effect on the cross-sectional average of the worst-case mean, as detailed in section 4.2.

Consider now the econometrician that measures realized hours, consumption and the marginal product of labor as of time t. While agents take the labor decision under E_t^*, the econometrician measures these equilibrium objects under the data generating process which uses the average $\mu = 0$. The difference between the worst-case distribution and the average realization under the econometrician's data generating process produces a labor wedge, which, in log-linear deviations, is inversely proportional to the time-varying confidence.

In a period of low confidence, the ratio between the expected benefit to working under the worst-case belief compared to the econometrician's measure of $\lambda_t MPL_t$ is typically lower. Thus the econometrician rationalizes the 'surprisingly low' labor supply by a high labor tax τ_t. In turn, the low confidence is generated endogenously from a low level of average economic activity, as reflected in the lower cross-sectional average of the worst-case mean, as given

[17] If we would assume that labor is chosen before the aggregate shocks are realized, there would be a fluctuating labor tax in (4.16) even in the rational expectations model. In that model, the wedge is $\tau_t = 1 - \frac{E_{t-1}(\lambda_t MPL_t)}{\lambda_t MPL_t}$, where, by the rational expectations assumptions, E_{t-1} reflects that agents form expectations using the econometrician's data generating process. Crucially, in such a model, the labor wedge τ_t will not be predictable using information at time $t-1$, including the labor choice, such that $E_{t-1}\tau_t = 0$. In contrast, our model with learning produces predictable, countercyclical, labor wedges.

by equation (4.14). Thus, the econometrician will find a systematic negative relationship between economic activity and the labor income tax. This relationship is consistent with empirical studies that suggest that in recessions labor falls by more than what can be explained by the marginal rate of substitution between labor and consumption and the measured marginal product of labor (see for example Shimer (2009) and Chari, Kehoe, and McGrattan (2007)). This countercyclical labor wedge does not arise from separate labor supply shocks but instead it is generated by any underlying shock that moves labor supply.

Finally, for an ease of exposition, we have described here the behavior of the labor wedge by ignoring the potential effect of experimentation on the optimal labor choice. This effect adds an additional reason why labor moves 'excessively', from the perspective of an observer that only uses equation (4.15) to understand labor movements. As discussed in section 5.2, experimentation amplifies the effects of uncertainty during the short-run, and thus leads to even more variable labor wedges, while it dampens fluctuations in the medium-run.

Excess return

A similar logic applies to the countercyclical ex-post excess return. The Euler conditions for capital and risk-free assets state

$$\lambda_t = \beta E_t^*[\lambda_{t+1} R_{t+1}^K] = \beta E_t^*[\lambda_{t+1} R_t]$$

Under our linearized solution, $E_t^* R_{t+1}^K = R_t$, where $E_t^* R_{t+1}^K$ is the expected return on capital under the worst-case belief. During low confidence times, demand for capital will be 'surprisingly low'. This is rationalized by the econometrician, measuring R_{t+1}^K under the true DGP, as a high ex-post excess return $R_{t+1}^K - R_t$. In the linearized solution, the excess return, similarly to the labor tax, is inversely proportional to the time-varying confidence. In times of low economic activity, when confidence is low, the measured excess return is high.

Thus, the model can generate a countercyclical labor wedge at times when the measured premia on uncertain assets is high.[18] This arises from any type of shock that moves the economic activity.

5 Quantitative results

We are interested in studying the role of endogenous firm-level uncertainty in business cycles. In this section we evaluate the empirical performance of the calibrated version of our model

[18]The pricing logic can be extended to defaultable corporate bonds. This will likely generate countercyclical excess bond premia, as documented for example by Gilchrist and Zakrajšek (2012).

Table 1: Parameters for the calibration exercise

γ	α	β	η	δ	χ_1	η_a	$\bar{\Sigma}$	$\bar{g}$	ρ_z	σ_z	ρ_A	ρ_g
1.004	0.3	0.99	0	0.025	0.15	0.4	0.1	0.2	0.5	0.4	0.95	0.95

and contrast its quantitative implications with those of a standard RBC model. We find the parsimony of this model a useful laboratory for the purpose of facilitating comparison with the standard RBC paradigm. We perform a more ambitious quantitative exercise in Section 6, where we introduce further nominal and real rigidities and estimate the model using a likelihood-based method.

5.1 Parameterization

Table 1 summarizes the parameters used in our exercise. In order to facilitate comparison of our model with a standard RBC model, we use common values used in the literature whenever possible.

The magnitude of the feedback loop between uncertainty and economic activity is determined by three factors. The first factor is the variability of inputs which is determined by the elasticities of labor supply and capital utilization. Regarding the labor supply elasticity, it is well known that standard real business cycles models understate the volatility of hours compared to the data. Motivated by this tension, we set the inverse Frisch elasticity η equal to zero following the indivisible labor model by Hansen (1985) and Rogerson (1988). We set the parameter that relates utilization rate to depreciation (χ_1) so that in equilibrium utilization rate is slightly less volatile than output.[19]

Second, the parameters that are related to the idiosyncratic processes control how changes in inputs translate to changes in the posterior variance. We choose $\rho_z = 0.5$ and $\sigma_z = 0.4$ for the idiosyncratic TFP process. Idiosyncratic TFP is less persistent than the aggregate, which is in line with the finding in Kehrig (2015). The values imply a cross-sectional standard deviation of TFP of 0.46 and is in line with the estimates found in Bloom et al. (2014) using the establishment-level data. Recall from the Discussion in section 4.2 that we re-parameterize the model so that we take the worst-case steady state posterior variance $\bar{\Sigma}^0$ of idiosyncratic TFP as a parameter. This posterior variance, together with ρ_z and σ_z, will pin down the standard deviation of the additive shock σ_ν. David et al. (2015) estimate the posterior variance of a firm-specific shock (in the context of our model, a TFP shock) to

[19] In the data, the standard deviation of the capacity utilization index, provided by the Federal Reserve Board of Governors, is more than twice as the standard deviation of GDP.

be around 8–13%. We choose the worst-case steady state posterior variance so that at the zero-risk steady state the posterior variance $\bar{\Sigma}$ is 10%.[20]

Finally, the size of the entropy constraint η_a determines how changes in the posterior variance translate into changes in confidence. To gauge the size of ambiguity that agents face, we use the cross-sectional mean of the dispersions of firm-level capital return forecasts across analysts from the Institutional Brokers' Estimate System (I/B/E/S) data. Specifically, we compute the min-max range of forecasts, defined as (max-min)/mean, for each firms and take the cross-sectional average.[21] For example, if there are three analysts for a firm X and the analyst 1's return on capital forecast is 8%, the analyst 2's forecast is 10%, and the analyst's forecast is 12%, then this will deliver a min-max range of 40%.[22] Averaged across the period 1985–2014, the cross-sectional mean of min-max range of forecasts is 43%. The model counterpart of the min-max range by analysts is the cross-sectional mean of min-max range of the capital return forecast implied by the set of productivity process (3.10) at the zero-risk steady state, where the minimum forecast is based on the worst-case mean $\mu_l^* = -a_l$ and the maximum forecast is based on the best-case mean $\mu_l^* = a_l$. We set $\eta_a = 0.4$, which generates the steady-state min-max range of forecasts of 39%.[23]

Our parameterization implies that the cross-sectional mean of the worst-case individual TFP at the zero-risk steady state is about 94 percent of the actual realized level. The Kalman gain at the zero-risk steady state, normalizing the level of input to one, is 0.47. To put this in perspective, the gain implies that an observation from quarters ago will receive a weight $(1 - 0.47)^4 \approx 0.08$. Thus, learning is fairly precise and quick under our parameterization.

5.2 Impulse response analysis

Figure 1 plots the impulse response to a positive TFP shock.[24] In addition to the response from the baseline model (labeled 'Baseline'), we also report the responses from the model with passive learning (labeled 'Passive'), in which the firm does not internalize the effect of its input choice on its future uncertainty, and the standard rational expectation (RE) RBC model (labeled 'RE'). The solution to the RE model is obtained by simply setting

[20]The zero-risk steady state is the ergodic steady state of the economy where optimality conditions take into account uncertainty and the data is generated under the econometrician's DGP. Appendix 8.3 provides additional details.

[21]We are grateful to Tatsuro Senga for providing us the statistics.

[22]During 1985–2014, the average number of firms sampled each year is 1460 and the average number of analysts surveyed for each firm is 9.8 in the I/B/E/S data set.

[23]Larger ambiguity leads to a higher min-max range of the capital return forecasts. For example, setting $\eta_a = 0.3$, holding other parameters fixed, generates a min-max range of 29% and $\eta_a = 0.5$ generates a 50% range. Ilut and Schneider (2014) argue that a reasonable upper bound for η_a is 2, based on the view that agents' ambiguity should not be "too large", in a statistical sense, compared to the variability of the data.

[24]We show the response to 0.1% increase in TFP.

Figure 1: Impulse response to an aggregate TFP shock

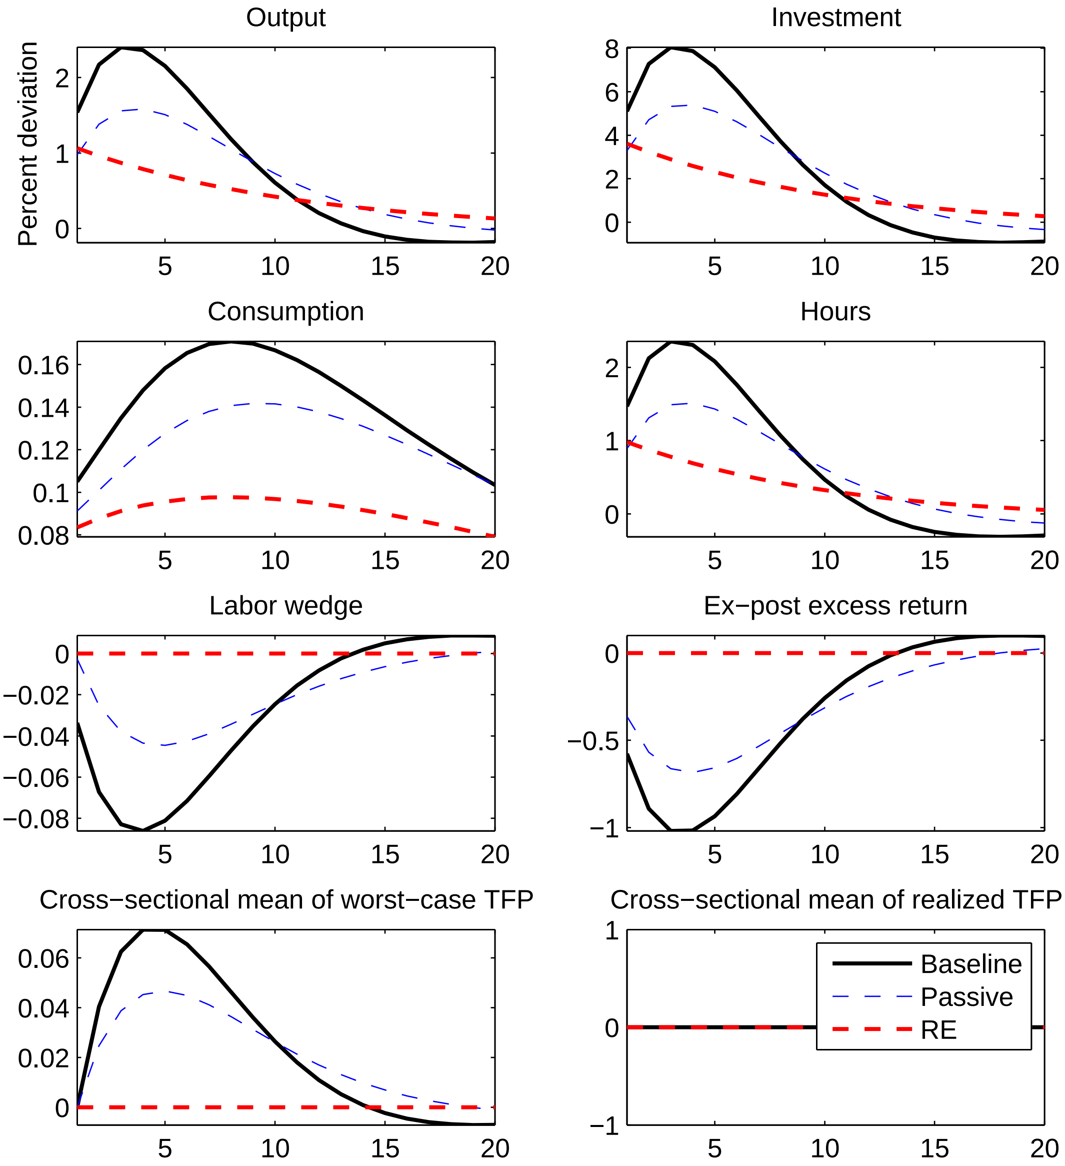

Notes: Thick black solid line is our baseline model with active learning, thin blue dashed line is the model with passive learning, and thick red dashed line is the frictionless, rational expectations, RBC model.

Figure 2: Impulse response to a government spending shock

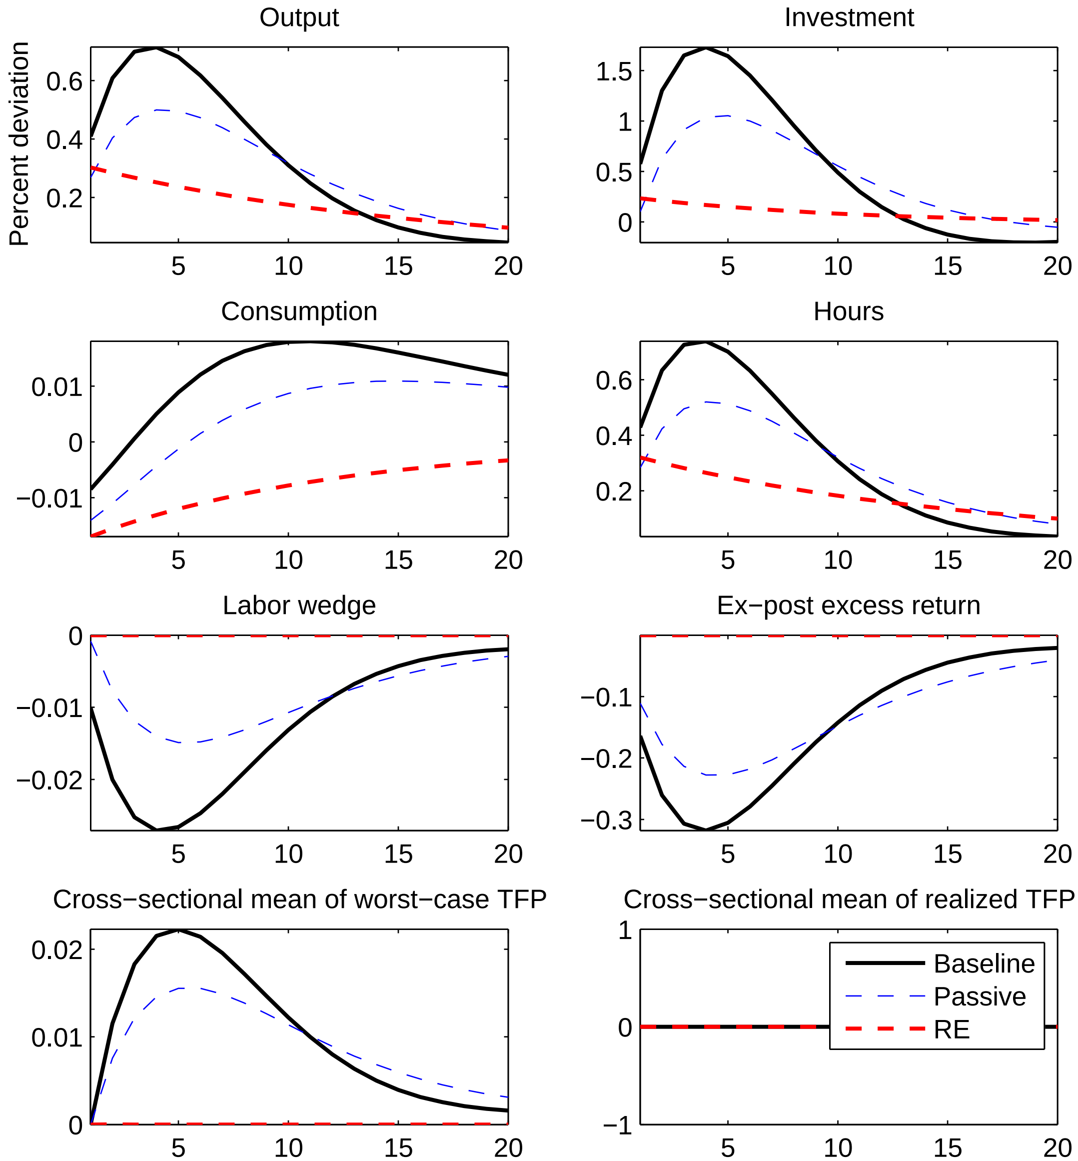

Notes: Thick black solid line is our baseline model with active learning, thin blue dashed line is the model with passive learning, and thick red dashed line is the frictionless, rational expectations, RBC model.

the entropy constraint η_a to zero. In this case, agents think in terms of single probabilities and the model reduces to a rational expectation model. Note that when $\eta_a = 0$, firm-level learning cancels out in the aggregate due to linearization and the law of large numbers.

Compared to the RE version, our model generates amplified and hump-shaped response in output, investment, and hours. These dynamics are due to the endogenous variation in firms' confidence. In response to a positive TFP shock, firms (on average) increase their inputs, such as hours and the capital utilization rate. The increase in inputs lowers uncertainty which implies that firms contemplate a narrower set of conditional probabilities; the worst-case scenarios become less worse. As a result, the agent acts *as if* the mean idiosyncratic productivities are higher and this may further stimulate economic activity. At the same time, from the perspective of the econometrician, labor supply and the demand for capital are surprisingly high. Thus, both labor wedge and ex-post excess return on capital decline.

Finally, we compare our baseline impulse response with the response from the passive learning model. Initially the output and hours responses of the baseline model with active learning are larger than those of passive learning. In the medium run, however, the responses of passive learning become larger. This is due to a dynamic interaction of two opposing forces. On one hand, higher production during booms increases the value of experimentation because it raises the marginal benefit of an increase in the expected worst-case technology. On the other hand, there is an offsetting effect coming from a reduction in posterior variance. Since the level of posterior variance is downward convex in the level of inputs, the marginal reduction in posterior variance due to an increase in inputs is smaller during booms. During the initial period of a positive technology shock, the first effect dominates the second. As the economy slows down, the second effect becomes more important.

Figure 2 shows the impulse response to a 1% increase in government spending. In the standard RBC model, due to the negative wealth effect hours increase but consumption decline. In our model, an increase in government spending, due to its effects of raising hours worked, also raises firms' confidence, which further stimulates economic activity. As a result, output, hours, and investment increases are larger. The negative response of consumption is overturned after five periods due to the positive income effect generated by the increase in confidence. Thus, a demand shock, such as the government spending shock may lead to positive comovement in consumption and labor.

5.3 Business cycle moments

Table 2 reports the HP-filtered second moments. To facilitate comparison with the standard RBC model, we assume that the only source of aggregate disturbance is the TFP shock.

Table 2: HP-filtered moments

	Data	Our model	RE
Standard deviations			
$\sigma(y)$	1.11	1.11	0.41
$\sigma(c)/\sigma(y)$	0.72	0.05	0.09
$\sigma(i)/\sigma(y)$	3.57	3.40	3.40
$\sigma(h)/\sigma(y)$	1.64	1.00	0.92
Correlations with labor wedge			
$\sigma(y, \tau_l)$	-0.83	-0.98	0
$\sigma(h, \tau_l)$	-0.97	-0.98	0
Autocorrelations			
$\sigma(y_t, y_{t-1})$	0.89	0.88	0.68
$\sigma(\Delta y_t, \Delta y_{t-1})$	0.39	0.49	-0.05

Notes: Both data and model moments are in logs, HP-filtered ($\lambda = 1600$) if the variables are in levels, and multiplied by 100 to express them in percentage terms. The model moments are the median values from 200 replications of simulations of 120 periods (after throwing away the initial 50 periods). The sample period for the data is 1985Q1–2014Q4. We choose the standard deviation of the aggregate TFP shock so that the output standard deviation in the baseline model matches the data. The government spending shock is set to zero.

First, with endogenous uncertainty the output standard deviation is more than twice larger than the RE version. The baseline model is also successful in generating a larger standard deviation of hours relative to output. The low volatility of hours has been a major shortcoming of RBC theories. Our model is less successful in reproducing the volatilities of consumption and investment. Second, our model can replicate the strong negative correlation of economic activity and the labor wedge. Third, our model gives a closer match in terms of autocorrelations. The baseline model generates higher autocorrelations in levels and, more importantly, positive autocorrelations in growth rates of output. As pointed out by Cogley and Nason (1995) and Rotemberg and Woodford (1996), a standard RBC model cannot generate persistence in output growth due to its weak internal propagation mechanism.

6 Empirical investigation: Bayesian estimation and firm-level evidence

In the previous section, we investigated the quantitative potential of endogenous uncertainty using a calibrated version of the model. In this section, we embed the learning mechanism into two DSGE models and conduct Bayesian likelihood estimations on US data. We ask two main questions. First, does endogenous uncertainty still play a key role when conventional rigidities are introduced? In turn, how important are those conventional rigidities when our learning mechanism is in place? Second, how does endogenous uncertainty affect inference about the source of business cycles and outcomes of policy experiments? We conclude the section by providing some firm-level data that supports the key implication of our model.

6.1 Models with conventional rigidities

We estimate two DSGE models: a flexible-price and a sticky-price model. In the flexible-price model, we augment the baseline model introduced in the previous section with two real rigidities. Specifically, we modify the representative household's utility (3.1) to allow for habit persistence in consumption:

$$U_t(C; s^t) = \ln(C_t - bC_{t-1}) - \varphi \frac{H_t^{1+\eta}}{1+\eta} + \beta \min_{p \in \mathcal{P}_t(s^t)} E^p[U_{t+1}(C; s^t, s_{t+1})],$$

where $b > 0$ is a parameter, and introduce an investment adjustment cost into the capital accumulation equation:

$$K_{l,t} = (1-\delta)K_{l,t-1} + \left\{1 - \frac{\kappa}{2}\left(\frac{I_{l,t}}{I_{l,t-1}} - \gamma\right)^2\right\} I_{l,t}, \tag{6.1}$$

where $\kappa > 0$ is a parameter.

In the sticky-price model, in addition to the consumption habit and investment adjustment cost, we consider Calvo-type price and wage stickiness along with monopolistic competition.[25] We assume a Taylor-type reaction function by the central bank:

$$\frac{R_t}{\bar{R}} = \left(\frac{R_{t-1}}{\bar{R}}\right)^{\rho_R} \left[\left(\frac{\pi_t}{\bar{\pi}}\right)^{\phi_\pi} \left(\frac{Y_t}{\bar{Y}}\right)^{\phi_Y}\right]^{1-\rho_R} \epsilon_{R,t}, \qquad \epsilon_{R,t} \sim N(0, \sigma_{R,t}^2),$$

where $\bar{\pi}$ is the inflation target, $\bar{R} = \bar{\pi}\gamma/\beta$, and $\bar{Y}$ is output along the balanced growth path. ρ_R, ϕ_π, and ϕ_Y are parameters and $\epsilon_{R,t}$ is a monetary policy shock.

[25]To avoid complications arising from directly embedding infrequent price adjustment into firms, we follow Bernanke et al. (1999) and assume that the monopolistic competition happens at the "retail" level. Retailers purchase output from firms in a perfectly competitive market, differentiate them, and sell them to final-goods producers, who aggregate retail goods using the conventional CES aggregator. The retailers are subject to the Calvo friction and thus can adjust their prices in a given period with probability $1 - \xi_p$. To introduce sticky wages, we assume that households supply differentiated labor services to the labor packer with a CES technology who sells the aggregated labor service to firms. Households can only adjust their wages in a given period with probability $1 - \xi_w$.

Finally, we introduce a "financial wedge" shock Δ_t^k to agents' Euler equation for capital accumulation (Christiano et al. (2015)):

$$1 = (1 - \Delta_t^k) E_t^* M_{t+1} R_{t+1}^k,$$

where $M_{t+1} \equiv M_0^{t+1}/M_0^t$ and R_{t+1}^k is the return on capital. We assume the process for Δ_t^k:

$$\Delta_t^k = (1 - \rho_\Delta)\bar{\Delta}^k + \rho_\Delta \Delta_{t-1}^k + \epsilon_{\Delta,t}, \qquad \epsilon_{\Delta,t} \sim N(0, \sigma_\Delta^2).$$

Combining the Euler equation for the risk-free asset, $1 = E_t^* M_{t+1} R_t$, with that for capital accumulation and rearranging, we obtain $\Delta_t^k \simeq E_t^* R_{t+1}^k - R_t$. As discussed in Christiano et al. (2015), the financial wedge could reflect variations in costs of financial intermediation such as bankruptcy costs or changes in the desirability of corporate bonds due to, for example, liquidity concern. In addition, our model generates the wedge through endogenous variation in confidence. We include the financial wedge shock because of a wide perception that financial frictions played a key role during the Great Recession. Our main focus is to explore how endogenous uncertainty affects the propagation of the financial wedge shock.

6.2 Bayesian estimation

Because the aggregate law of motion is linear, we can use standard Bayesian techniques as described in An and Schorfheide (2007) to estimate the model. The sample period is 1985Q1–2014Q4. The data is described in Appendix 8.5. For the flexible-price model, the vector of observables are

$$[\Delta \ln Y_t, \Delta \ln H_t, \Delta \ln I_t, \Delta \ln C_t, Spread_t],$$

where $Spread_t$ is the Baa corporate bond yield relative to the yield of Treasury bond with ten-year maturity.[26] We assume that the model counterpart of $Spread_t$ is the excess-return on capital so that $Spread_t = R_t^k - R_{t-1}$.[27] For the sticky-price model, we add inflation and

[26] We also used the spread constructed in Gilchrist and Zakrajšek (2012) and obtained similar results.

[27] To understand what causes variations in the spread, it is useful to decompose it into three components:

$$Spread_t = (E_{t-1}^* R_t^k - R_{t-1}) + (R_t^k - E_t^{'} R_t^k) + (E_t^{'} R_t^k - E_{t-1}^* R_t^k), \tag{6.2}$$

where we use $E_t^{'}$ to use worst-case expectations at the end of stage 1 (after the realization of period t aggregate shock but before the realization of idiosyncratic shocks). The financial wedge shock Δ_t^k causes variations in the first component, changes in confidence cause variations in the second component, and innovations to aggregate shocks (such as TFP shocks) cause variations in the third component.

the Federal funds rate to the observables:

$$[\Delta \ln Y_t, \Delta \ln H_t, \Delta \ln I_t, \Delta \ln C_t, \ln \pi_t, \ln R_t, Spread_t].$$

We fix a small number of parameters prior to estimation. The depreciation rate of capital δ is set to 0.025. The steady-state share of government spending to output is 0.2. We set θ_p and θ_w so that the steady-state price and wage markups are both 10%. The prior distributions for other structural parameters are collected in Table 3 and 4. For most parameters, we choose relatively loose priors. We reparametrize the parameter that determines the size of ambiguity (η_a) and instead estimate $0.5\eta_a$. We set a Beta prior for $0.5\eta_a$, centered around the calibrated value in the previous section, so that the lowest value corresponds to rational expectations and the highest value corresponds to the upper bound $\eta_a = 2$ suggested by Ilut and Schneider (2014). Since all our estimation exercises have less structural shocks than observables, we add i.i.d. measurement error to observables except for the interest rate. At the posterior mean, measurement error explains 1 percent of variation of a particular observable while at one standard deviation it explains 5 percent.[28]

6.3 Results

Parameter estimates

Table 3 and 4 report the posterior distributions of the two models (labeled 'Baseline'). For comparison, we also report estimates of the rational expectations versions of the models (labeled 'RE'). Two key results emerge. First, aggregate data prefers sizable amount of ambiguity. The posterior estimates of the parameter that determines the size of the entropy constraint, η_a, are 1.24 and 1.32 for flexible-price and sticky-price models, respectively. The posteriors imply that agents contemplate a set of conditional means whose bounds are given by roughly ± 1 standard deviations around the benchmark Kalman filter estimate.

Second, when we allow for endogenous ambiguity, the estimated degrees of real and nominal rigidities are smaller than in the rational expectations versions. In the sticky-price model, the Calvo probability of not being able to adjust the price decreases from $\xi_p = 0.76$ to $\xi_p = 0.56$ and the Calvo wage probability decreases from $\xi_w = 0.96$ to $\xi_w = 0.02$. To put these into perspective, in the RE version of the model, prices are adjusted on average every $1/(1-0.76) \approx 4$ quarters and wages are adjusted every $1/(1-0.96) = 25$ quarters. In contrast, in our baseline model prices are adjusted every $1/(1-0.56) \approx 2$ quarters and wages are adjusted roughly every quarter. Compared to the RE version, the estimated frequency of

[28] The only exception to this is a measurement error for the spread; at one standard deviation it explains 20 percent of the variation.

Table 3: Estimated parameters: preference, technology, and policy

		Prior			Posterior mode			
					Flexible price		Sticky price	
		Type	Mean	Std	Baseline	RE	Baseline	RE
$100(\gamma-1)$	Growth rate	N	0.5	0.3	0.48	0.47	0.57	0.34
					[0.42, 0.55]	[0.39, 0.56]	[0.50, 0.62]	[0.26, 0.43]
$100(\beta^{-1}-1)$	Discount factor	G	0.25	0.05	1.01	1.01	0.27	0.19
							[0.15, 0.34]	[0.13, 0.26]
$100(\bar{\pi}-1)$	Net inflation	N	0.5	0.3	–	–	0.72	0.45
							[0.58, 0.82]	[0.17, 0.80]
α	Capital share	B	0.3	0.05	0.19	0.25	0.41	0.33
					[0.14, 0.31]	[0.21, 0.33]	[0.33, 0.45]	[0.28, 0.38]
η	Inverse Frisch elasticity	G	1	0.4	0.38	0.03	0.26	0.79
					[0.23, 0.55]	[0.01, 0.08]	[0.14, 0.39]	[0.89, 2.13]
χ_1	Utilization cost	G	0.7	0.3	0.46	0.08	0.22	1.27
					[0.24, 0.85]	[0.04, 0.23]	[0.16, 0.31]	[0.42, 2.00]
b	Consumption habit	B	0.4	0.1	0.53	0.40	0.47	0.62
					[0.34, 0.65]	[0.27, 0.53]	[0.38, 0.62]	[0.49, 0.71]
κ	Investment adj. cost	G	3	2	0.00	0.09	0.06	0.30
					[0.00, 0.00]	[0.07, 0.21]	[0.05, 0.08]	[0.20, 0.49]
ξ_p	Calvo price	B	0.5	0.2	–	–	0.56	0.76
							[0.52, 0.59]	[0.71, 0.80]
ξ_w	Calvo wage	B	0.5	0.2	–	–	0.02	0.96
							[0.00, 0.03]	[0.94, 0.97]
ρ_R	Interest smoothing	B	0.7	0.1	–	–	0.57	0.35
							[0.53, 0.64]	[0.28, 0.45]
ϕ_π	Inflation response	N	1.5	0.15	–	–	1.91	2.10
							[1.77, 2.12]	[1.88, 2.28]
ϕ_Y	Output response	N	0.1	0.05	–	–	0.07	0.00
							[0.06, 0.08]	[-0.01, 0.00]
ρ_z	Idiosyncratic TFP	B	0.5	0.05	0.62	–	0.62	–
					[0.46, 0.72]		[0.57, 0.65]	
σ_z	Idiosyncratic TFP	G	0.4	0.1	0.80	–	0.83	–
					[0.68, 1.11]		[0.72, 0.92]	
$0.5\eta_a$	Entropy constraint	B	0.2	0.1	0.62	0	0.66	0
					[0.46, 0.72]		[0.60, 0.80]	
$\bar{\Sigma}$	SS posterior variance	G	0.1	0.05	0.05	–	0.12	–
					[0.03, 0.11]		[0.03, 0.24]	

Notes: B refers to the Beta distribution, N to the Normal distribution, G to the Gamma distribution, IG to the Inverse-gamma distribution. 95% posterior intervals are in brackets and are obtained from 200,000 draws using the random-walk Metropolis-Hasting algorithm.

Table 4: Estimated parameters: shocks and measurement errors

		Prior			Posterior mode			
					Flexible price		Sticky price	
		Type	Mean	Std	Baseline	RE	Baseline	RE
ρ_A	Aggregate TFP	B	0.6	0.2	0.98	0.99	0.93	0.93
					[0.93, 1.00]	[0.99, 1.00]	[0.89, 0.95]	[0.90, 0.96]
ρ_g	Government spending	B	0.6	0.2	0.70	0.28	0.94	0.70
					[0.21, 0.95]	[0.12, 0.88]	[0.72, 0.99]	[0.26, 0.96]
ρ_Δ	Financial wedge	B	0.6	0.2	0.53	0.27	0.98	0.94
					[0.44, 0.62]	[0.09, 0.59]	[0.94, 1.00]	[0.91, 0.95]
$100\sigma_A$	Aggregate TFP	IG	1	1	0.28	0.35	0.22	0.26
					[0.23, 0.35]	[0.26, 0.43]	[0.19, 0.27]	[0.22, 0.34]
$100\sigma_g$	Government spending	IG	1	1	0.50	0.78	0.90	0.44
					[0.34, 0.94]	[0.37, 1.01]	[0.72, 1.10]	[0.29, 0.67]
$100\sigma_\Delta$	Financial wedge	IG	0.1	1	0.10	0.03	0.05	0.06
					[0.09, 0.12]	[0.02, 0.04]	[0.04, 0.06]	[0.05, 0.07]
$100\sigma_R$	Monetary policy	IG	0.1	1	–	–	0.04	0.03
							[0.03, 0.05]	[0.02, 0.04]
$100\sigma_Y$	Meas. error output	IG	0.060	0.13	0.29	0.03	0.45	0.32
					[0.11, 0.40]	[0.02, 0.32]	[0.40, 0.54]	[0.22, 0.37]
$100\sigma_H$	Meas. error hours	IG	0.073	0.16	0.44	0.56	0.58	0.39
					[0.38, 0.53]	[0.46, 0.62]	[0.47, 0.62]	[0.33, 0.47]
$100\sigma_I$	Meas. error investm.	IG	0.18	0.40	1.07	0.99	1.21	1.06
					[0.86, 1.32]	[0.79, 1.16]	[1.08, 1.39]	[0.91, 1.26]
$100\sigma_C$	Meas. error consum.	IG	0.050	0.11	0.41	0.35	0.39	0.33
					[0.35, 0.47]	[0.31, 0.42]	[0.33, 0.44]	[0.29, 0.40]
$100\sigma_{spread}$	Meas. error spread	IG	0.018	0.08	0.01	0.20	0.01	0.28
					[0.00, 0.03]	[0.18, 0.24]	[0.01, 0.05]	[0.26, 0.34]
$100\sigma_\pi$	Meas. error inflation	IG	0.024	0.11	–	–	0.23	0.28
							[0.19, 0.25]	[0.25, 0.33]
Log marginal likelihood					2336	2262	3489	3413

Notes: See notes for Table 3.

Figure 3: Estimated impulse response to a financial wedge shock in the flexible-price model

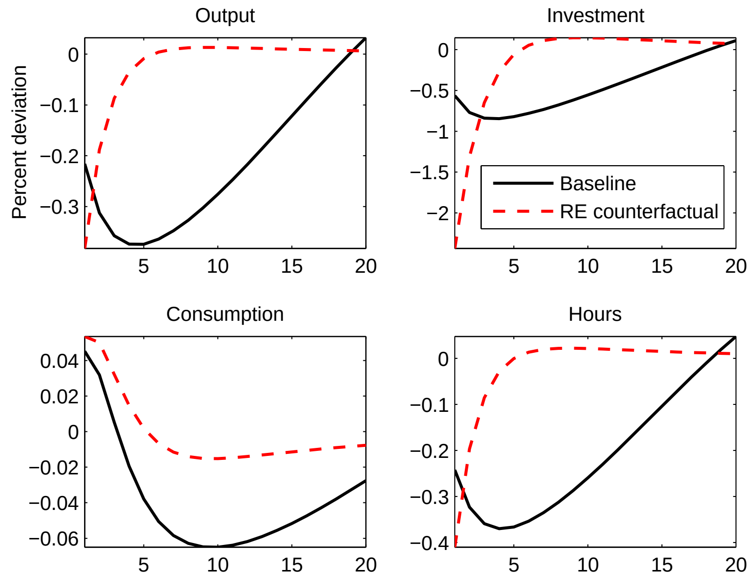

Notes: Black solid line is the baseline impulse response and the red dashed line is the response where we set the entropy constraint $\eta_a = 0$ while holding other parameters at their estimated values.

price change is thus more in line with the micro evidence presented in Bils and Klenow (2004). The degree of consumption habit decreases by 25% ($b = 0.62$ to $b = 0.47$) and the investment adjustment cost falls from $\kappa = 0.3$ to $\kappa = 0.06$. In the flexible-price model, while consumption habit slightly increases with endogenous uncertainty, the investment adjustment cost falls from $\kappa = 0.09$ to zero.[29]

Propagation of the financial wedge shock

We argue that the reason why our model requires smaller rigidities than the RE version is because our learning mechanism provides strong internal propagation and induces co-movement in response to the financial wedge shock.

To understand this, Figure 3 shows the estimated impulse response to a financial wedge shock in the flexible-price model. In addition to the baseline impulse response (labeled 'Baseline'), we also display the response for a version of the model in which we set $\eta_a = 0$ and freeze

[29]The intuition for the increase in the habit parameter in the flexible-price model is as follows. As we will see below, while endogenous uncertainty induces co-movement in response to a financial wedge shock in the medium run, consumption moves in the opposite direction in the short run. Because in our baseline flexible-price model the financial wedge shock accounts for a non-trivial variation of consumption, the estimation tries to mitigate this initial negative co-movement by increasing the habit. In the sticky-price model, such force is absent because all quantities move in the same direction from the initial period.

Figure 4: Estimated impulse response to a financial wedge shock in the sticky-price model

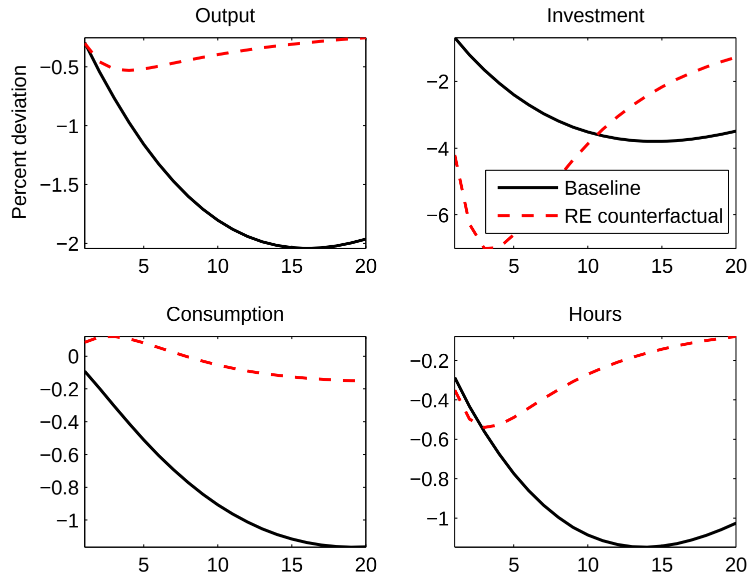

Notes: See notes for Figure 3.

the remaining parameters to the baseline estimated values (labeled 'RE counterfactual'). In the baseline, output, investment, and hours falls but consumption increase in the short run. In the medium-run, however, consumption also falls because of the decline in confidence. At the same time, the model generates hump-shaped dynamics even though the estimated real rigidities are small. In the RE counterfactual, on the other hand, consumption and other variables move in the opposite direction and the impulse response is monotonic.[30]

Figure 4 shows that similar forces are at work in the sticky-price model. However, all quantities fall immediately because countercyclical markups provide an additional wedge that breaks the Barro and King (1984) critique. At the same time, our model requires less nominal rigidities because endogenous uncertainty partly acts as a substitute to the countercyclical markups in generating co-movement.

Model fit

The final row of Table 4 reports the log marginal likelihoods, computed by the Geweke's modified harmonic mean estimator, across different model specifications. For both flexible-

[30]A different way to evaluate the impact of our learning mechanism is to examine the impulse response in the RE model, in which the parameters are re-estimated subject to the constraint $\eta_a = 0$. We find that the model is still unable to produce co-movement.

Table 5: Fit of observables

Model	Statistic	Output	Hours	Investment	Consumption	Inflation	Spread
	σ(data)	0.60	0.73	1.79	0.50	0.24	0.18
Panel A: Flexible price							
Baseline	σ(model)	0.60	0.45	1.21	0.21	–	0.18
	ρ(data, model)	0.94	0.83	0.86	0.62	–	1.00
RE	σ(model)	0.60	0.30	1.43	0.28	–	0.04
	ρ(data, model)	1.00	0.76	0.87	0.77	–	-0.37
Panel B: Sticky price							
Baseline	σ(model)	0.60	0.40	1.20	0.33	0.19	0.18
	ρ(data, model)	0.75	0.69	0.82	0.67	0.53	1.00
RE	σ(model)	0.57	0.58	1.17	0.35	0.32	0.25
	ρ(data, model)	0.92	0.88	0.86	0.80	0.52	0.19

Notes: σ(data) and σ(model) refer to standard deviations of the observables and the model counterparts, respectively, expressed in percentage points. ρ(data, model) is the correlation between data and the model.

price and sticky-price models, the model fit, measured in terms of log marginal likelihoods, is improved when our endogenous uncertainty mechanism is activated. To understand this result, in Table 5 we report the summary statistics of our observables and their model counterparts at the posterior mean, computed by the Kalman smoothing algorithm. The fit of endogenous variables except for the spread in the baseline flexible-price and sticky-price models is similar to the RE versions. The last column of Panel A shows that while the baseline flexible-price model perfectly matches the movement of the spread, in the RE version the model generates less than a quarter of the variation and it is negatively correlated with data. Panel B shows that the baseline sticky-price model also perfectly reproduces the movement of the spread. In contrast, the fit is much worse in the RE version. Indeed, Table 4 shows that measurement errors for spreads are much larger in the RE versions for both flexible-price (0.20 vs. 0.01) and sticky-price (0.28 vs. 0.01) models.

Figure 5 visualizes this result by plotting the spread along with the flexible-price model counterpart computed from the Kalman smoother. The model-implied spread, labeled 'Baseline', tracks the data while the spread from the RE model, labeled 'RE', does not. There are two reasons for the success of our baseline model. First, with endogenous uncertainty the financial wedge shock can now generate co-movement and thus the estimation prefers contractionary wedge shocks during recessions. Second, agents lose confidence during recessions, which further contribute to the countercyclicality of the spread. This endogenous

Figure 5: Credit spread: data and model

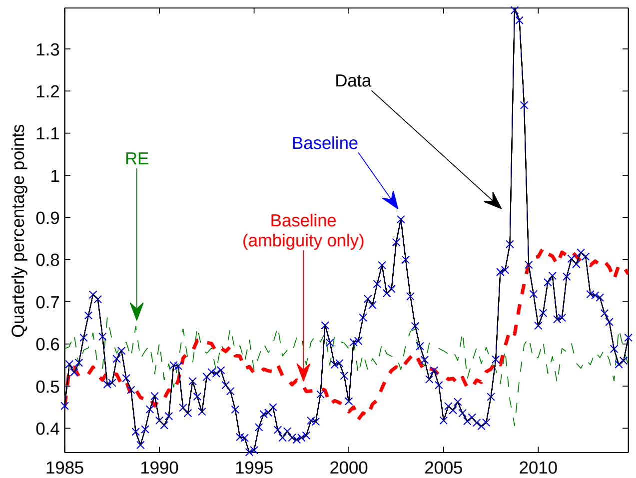

Notes: We plot the Baa corporate bond spread along with the flexible-price model counterpart computed from the Kalman smoother. Thick black solid line is the data, thin blue line (marked with x) is the spread from the baseline model, thick red dashed line is the spread generated by endogenous variation in confidence, and thin green dashed line is the spread from the rational expectations model.

variation in the spread is labeled 'Baseline (ambiguity only)'.[31] The plot of the spreads from the sticky-price model (not shown) also shows a similar pattern. We conclude that the ability of our model to generate countercyclical spreads, due to the financial wedge shock being countercyclical and the endogenous variation in confidence, is the key reason the estimation favors our learning mechanism.

Variance decomposition

To see how the introduction of endogenous uncertainty alters inference about the source of business cycles, Table 6 reports the contributions of TFP and financial wedge shocks. Panel A shows the variance decomposition for the flexible-price model. In the RE version, the TFP shock accounts for almost all variations of macro variables. In contrast, in our baseline model, the financial wedge shock explains a non-trivial fraction of business cycles since the shock can now generate co-movement; it accounts for 14 and 16 percents of volatilities in output and hours, respectively. Panel B shows that, with nominal rigidities, the financial wedge shock accounts for a sizable fraction of variations in quantities even in the RE version. For example,

[31] This endogenous variation in the spread corresponds to the second component in the previous decomposition of the spread (6.2).

Table 6: Theoretical variance decomposition

Model	Shock	Output	Hours	Investm.	Consum.	Inflation	Interest rate
Panel A: Flexible price							
Baseline	TFP	0.86	0.84	0.83	0.92	–	–
	Financial wedge	0.14	0.16	0.16	0.08	–	–
RE	TFP	0.99	0.98	0.99	1.00	–	–
	Financial wedge	0.00	0.01	0.01	0.00	–	–
Panel B: Sticky price							
Baseline	TFP	0.27	0.18	0.20	0.38	0.09	0.16
	Financial wedge	0.73	0.81	0.76	0.61	0.88	0.84
RE	TFP	0.84	0.77	0.88	0.78	0.12	0.10
	Financial wedge	0.15	0.23	0.12	0.22	0.88	0.90

Notes: We report the percent of variance at the business cycle frequency (6–32 quarters) that can be explained by each shock. For the interest of space, we omit the percent of variance explained by government spending and monetary policy shocks.

the shock explains 15 and 23 percents of variations in output and hours, respectively. With endogenous uncertainty, however, the financial wedge shock becomes the dominant source of the business cycle. For example, the shock now explains 73 and 81 percents of variations in output and hours, respectively. The financial wedge shock accounts for a large share of variation in inflation and interest rates for both the baseline and the RE models.

In Appendix 8.4, we show that the inclusion of the spread in the estimation is the key reason why in our model the financial wedge shock becomes an important driver of the business cycle. Intuitively, the financial wedge shock needs to be large since the spread is volatile and countercyclical. Without endogenous uncertainty, however, this comes at the cost of generating counterfactual predictions for other macro variables.

Policy implications of endogenous uncertainty

Finally, the fact that in our model uncertainty is endogenous has important policy implications. To illustrate this point, we evaluate the impact of modifying the Taylor rule to incorporate an adjustment to the credit spread.[32] Table 7 reports the standard deviation of output growth as we maintain all parameters as their baseline estimated values but change the Taylor rule coefficient on the credit spread ϕ_{spread} from the original value of zero in the estimated sticky-price model. The standard deviation decreases as the monetary policy responds more aggressively to the spread movements. For example, the standard deviation

[32] Such a policy was proposed by Taylor (2008), among others. Cúrdia and Woodford (2010) and Christiano et al. (2010) study the effects of similar policies using DSGE models.

Table 7: Changing the Taylor rule coefficient on the credit spread

	Std. of output growth	
ϕ_{spread}	Baseline	Fixed uncertainty
0	0.60	0.60
-0.5	0.59	0.60
-1.0	0.57	0.60
-1.5	0.52	0.63

Notes: The standard deviation of output growth as the Taylor rule coefficient on the credit spread ϕ_{spread} is changed from the original value ($\phi_{spread} = 0$). All other parameter values are fixed to the estimated values. "Baseline" refers to the baseline sticky-price model with endogenous uncertainty. "Fixed uncertainty" refers to a counterfactual economy where the path of uncertainty is fixed exogenously to the original one.

decreases from 0.60 to 0.52 when the coefficient ϕ_{spread} decreases from 0 to -1.5.

The reduction in output variability comes from stabilizing the endogenous variation in uncertainty. To see this, we also show the effects of policy changes in the counterfactual economy where the path of uncertainty is fixed to the original one (labeled "Fixed uncertainty"). In this economy, changes in ϕ_{spread} have smaller effects and even have a different sign. Indeed, the standard deviation increases from 0.60 to 0.63 when the coefficient ϕ_{spread} decreases from 0 to -1.5. The comparison of these counterfactual models underscores the importance for policy analysis of modeling time-variation in uncertainty as an endogenous response that in turn further affects economic decisions.

6.4 Firm-level evidence

The estimated business cycle model with conventional rigidities suggests that endogenous countercyclical idiosyncratic uncertainty may be an important propagation mechanism at the aggregate level. We conclude our empirical investigation by changing focus and report some cross-sectional evidence.

Our theory predicts that, through learning, a larger scale of production leads to lower uncertainty. We use here firm-level data to test this prediction. We use panel data of analysts' forecast error on each firm's return on capital from the I/B/E/S data. The data set spans annually from 1977–2014 and is constructed in Senga (2015).[33] We define each firm's forecast errors as (realized return - mean forecast)/mean forecast. We proxy the scale of production by the size of total assets held by each firm. Our hypothesis is that, cross-

[33]We thank Tatsuro Senga for conducting and sharing the analysis in this section and for numerous advice on the data set.

Table 8: Forecast errors are smaller for larger firms

Percentile	0–10	10–25	25–50	50–75	75–90	90–100
Mean error	0.93	0.62	0.51	0.38	0.32	0.25
Median error	0.35	0.21	0.16	0.12	0.09	0.08

Notes: The table reports the cross-sectional mean (second row) and the cross-sectional median (third row) of the forecast errors on capital returns, grouped according to the total asset held by each firm.

Table 9: Forecast errors are smaller for older firms

Age	1–10	11–20	21–30	31–
Mean error	0.65	0.40	0.30	0.27
Median error	0.22	0.13	0.10	0.08

Notes: The table reports the cross-sectional mean (second row) and the cross-sectional median (third row) of the forecast errors on capital returns, grouped according to the age. We proxy firm age by the number of years since the firm's first year of appearance in the I/B/E/S data set.

sectionally, the size of forecast errors are smaller the larger the scale of production. Table 8 documents that the data provides empirical support for our hypothesis.

We also consider a test that is more loosely connected to our theory. Our theory builds on the idea that firms learn about their fundamentals through production, suggesting that old firms face less uncertainty. Indeed, in Table 9, we show that the forecast error are smaller for older firms. While our model does not feature entry and exit and hence cannot directly speak to the relationship between firm age and uncertainty, the result in Table 9 is encouraging as it provides suggestive evidence for our learning mechanism.

Finally, we point out that the time-series implications of our model are supported by the data as well. Our theory predicts that the cross-sectional average of uncertainty increases during recessions. Senga (2015) shows that the average level of capital return forecast errors across firms is strongly countercyclical and rose sharply during the Great Recession.

7 Conclusion

In this paper we construct a tractable heterogeneous-firm business cycle model in which firms face Knightian uncertainty about their own profitability and need to learn it through production. The feedback loop between economic activity and confidence makes our model behave as a standard linear business cycle model with (i) countercyclical labor wedge and financial spread, (ii) positive co-movement of aggregate variables in response to either supply or demand shocks, and (iii) strong internal propagation with amplified and hump-shaped dynamics. When the model is estimated using US data, our learning process drives out the role of traditional rigidities and changes inference about the source of aggregate fluctuations and outcomes of policy experiments. We conclude that endogenous idiosyncratic uncertainty is a quantitatively important mechanism for understanding business cycles.

References

An, S. and F. Schorfheide (2007): "Bayesian Analysis of DSGE Models," *Econometric Reviews*, 26, 113–172.

Angeletos, G.-M., F. Collard, and H. Dellas (2014): "Quantifying Confidence," *Working Paper*.

Arellano, C., Y. Bai, and P. J. Kehoe (2012): "Financial Markets and Fluctuations in Uncertainty," *Federal Reserve Bank of Minneapolis Staff Report*.

Bachmann, R. and C. Bayer (2013): "'Wait-and-See' Business Cycles?" *Journal of Monetary Economics*, 60, 704–719.

Barro, R. J. and R. G. King (1984): "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *Quarterly Journal of Economics*, 99, 817–839.

Basu, S. and B. Bundick (2012): "Uncertainty Shocks in a Model of Effective Demand," *NBER Working Paper No. 18420*.

Bernanke, B. S., M. Gertler, and S. Gilchrist (1999): "The Financial Accelerator in a Quantitative Business Cycle Framework," in *Handbook of Macroeconomics*, ed. by J. B. Taylor and M. Woodford, Amsterdam: Elsevier, vol. 1, Part C, chap. 21, 1341–1393.

Bianchi, F., C. L. Ilut, and M. Schneider (2014): "Uncertainty Shocks, Asset Supply and Pricing over the Business Cycle," *NBER Working Paper No. 20081*.

Bidder, R. and M. E. Smith (2012): "Robust Animal Spirits," *Journal of Monetary Economics*, 59, 738–750.

Bils, M. and P. J. Klenow (2004): "Some Evidence on the Importance of Sticky Prices," *Journal of Political Economy*, 112, 947–985.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77, 623–685.

Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. Terry (2014): "Really Uncertain Business Cycles," *Working Paper*.

Born, B. and J. Pfeifer (2014): "Policy Risk and the Business Cycle," *Journal of Monetary Economics*, 68, 68–85.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2007): "Business Cycle Accounting," *Econometrica*, 75, 781–836.

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

Christiano, L. J., M. S. Eichenbaum, and M. Trabandt (2015): "Understanding the Great Recession," *American Economic Journal: Macroeconomics*, 7, 110–167.

Christiano, L. J., C. L. Ilut, R. Motto, and M. Rostagno (2010): "Monetary Policy and Stock Market Booms," NBER Working Papers 16402.

Christiano, L. J., R. Motto, and M. Rostagno (2014): "Risk Shocks," *American Economic Review*, 104, 27–65.

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

Cogley, T. and J. M. Nason (1995): "Output Dynamics in Real-Business-Cycle Models," *American Economic Review*, 85, 492–511.

Cúrdia, V. and M. Woodford (2010): "Credit Spreads and Monetary Policy," *Journal of Monetary, Credit, and Banking*, 42, 3–35.

David, J. M., H. A. Hopenhayn, and V. Venkateswaran (2015): "Information, Misallocation, and Aggregate Productivity," *Quarterly Journal of Economics*, forthcoming.

Epstein, L. G. and M. Schneider (2003a): "IID: independently and indistinguishably distributed," *Journal of Economic Theory*, 113, 32–50.

——— (2003b): "Recursive Multiple-Priors," *Journal of Economic Theory*, 113, 1–31.

Epstein, L. G. and T. Wang (1994): "Intertemporal Asset Pricing under Knightian Uncertainty," *Econometrica*, 62, 283–322.

Eusepi, S. and B. Preston (2015): "Consumption Heterogeneity, Employment Dynamics and Macroeconomic Co-movement," *Journal of Monetary Economics*, 71, 13–32.

Fajgelbaum, P., E. Schaal, and M. Taschereau-Dumouchel (2015): "Uncertainty Traps," *Working Paper*.

Fernández-Villaverde, J., P. Guerrón-Quintana, K. Kuester, and J. F. Rubio-Ramírez (2015): "Fiscal Volatility Shocks and Economic Activity," *American Economic Review*, forthcoming.

Fernández-Villaverde, J., P. Guerrón-Quintana, J. F. Rubio-Ramírez, and M. Uribe (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101, 2530–2561.

Gilboa, I. and D. Schmeidler (1989): "Maxmin Expected Utility with Non-unique Prior," *Journal of Mathematical Economics*, 18, 141–153.

Gilchrist, S. and E. Zakrajšek (2012): "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

Hansen, G. D. (1985): "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 16, 309–337.

Ilut, C. and M. Schneider (2014): "Ambiguous Business Cycles," *American Economic Review*, 104, 2368–2399.

Jaimovich, N. and S. Rebelo (2009): "Can News about the Future Drive the Business Cycle?" *American Economic Review*, 99, 1097–1118.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010): "Investment Shocks and Business Cycles," *Journal of Monetary Economics*, 57, 132–145.

Kehrig, M. (2015): "The Cyclical Nature of Productivity Dispersion," *Working Paper*.

Leduc, S. and Z. Liu (2015): "Uncertainty Shocks are Aggregate Demand Shocks," *Working Paper*.

Lorenzoni, G. (2009): "A Theory of Demand Shocks," *American Economic Review*, 99, 2050–2084.

Marinacci, M. (1999): "Limit Laws for Non-Additive Probabilities, and Their Frequentist Interpretation," *Journal of Economic Theory*, 84, 145–195.

Ordoñez, G. (2013): "The Asymmetric Effects of Financial Frictions," *Journal of Political Economy*, 5, 844–895.

Rogerson, R. (1988): "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics*, 21, 3–16.

Rotemberg, J. J. and M. Woodford (1996): "Real-Business-Cycle Models and the Forecastable Movements in Output, Hours, and Consumption," *American Economic Review*, 86, 71–89.

Saijo, H. (2014): "The Uncertainty Multiplier and Business Cycles," *Working Paper*.

Schaal, E. (2015): "Uncertainty and Unemployment," *Working Paper*.

Senga, T. (2015): "A New Look at Uncertainty Shock: Imperfect Information and Misallocation," *Working Paper*.

Shimer, R. (2009): "Convergence in Macroeconomics: The Labor Wedge," *American Economic Journal: Macroeconomics*, 1, 280–297.

Sims, C. A. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.

Taylor, J. B. (2008): "Monetary Policy and the State of the Economy," *Testimony before the Committee on Financial Services U.S. House of Representatives*.

van Nieuwerburgh, S. and L. Veldkamp (2006): "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53, 753–772.

Veldkamp, L. L. (2005): "Slow Boom, Sudden Crash," *Journal of Economic Theory*, 124, 230–257.

8 Appendix (For online publication)

8.1 A model of learning about firm-specific demand

In this section, we show that the baseline model with additive shock in the production function (3.2) can be reinterpreted as a model where firms learn about their demand from noisy signals.

There is a continuum of firms, indexed by $l \in [0,1]$, which produce intermediate goods and sell them to a large representative "conglomerate". The conglomerate, who holds shares of the intermediate firms, acts in a perfectively competitive manner and combine the intermediate goods to produce final goods. To ease exposition, we momentarily abstract from all aggregate shocks, labor-augmenting technological growth, and utilization.

The conglomerate combines intermediate output $Y_{l,t}$ according to the following CES aggregator:

$$Y_t = \left[\int_0^1 z_{l,t}^{\frac{\sigma-1}{\sigma}} Y_{l,t}^{\frac{\sigma-1}{\sigma}} dl \right]^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 1$$

where $z_{l,t}$ follows an AR(1) as in (3.3). In turn, intermediate goods $Y_{l,t}$ are produced according to

$$Y_{l,t} = K_{l,t-1}^{\alpha} H_{l,t}^{1-\alpha}.$$

We assume that all agents, including the conglomerate and households, cannot observe the realization of $z_{l,t}$. Instead, after the production of final and intermediate goods they observe noisy signals

$$s_{l,t} = Y_{l,t} z_{l,t} + \tilde{\nu}_{l,t}, \qquad \tilde{\nu}_{l,t} \sim N(0, \sigma_{\tilde{\nu}}^2) \tag{8.1}$$

where $\tilde{\nu}_{l,t}$ is an observation error. Agents use all available information, including the path of signals $s_{l,t}$ and intermediate output $Y_{l,t}$, to form estimates about the realization of $z_{l,t}$.

In this context, it is natural to think $z_{l,t}$ as a "quality" of intermediate good l that is difficult to observe. Alternatively, since the final good could directly be used for consumption, $z_{l,t}$ can be interpreted as an unobservable demand for a variety l. Crucially, (8.1) implies that the signal-to-noise ratio is increasing in the level of output $Y_{l,t}$. It is plausible that firms learn more about the quality or demand, $z_{l,t}$, of their goods when they produce and sell more. For example, when a restaurant serves more customers it generates more website reviews and hence people learn more about the quality of their meals.

Intermediate firms l choose price $P_{l,t}$ and inputs to maximize the shareholder value

$$E_0^* \sum_{t=0}^{\infty} M_0^t D_{l,t}.$$

$D_{l,t}$ is the dividend payout to the conglomerate given by

$$D_{l,t} = \frac{P_{l,t}}{P_t} Y_{l,t} - W_t H_{l,t} - I_{l,t},$$

where P_t is the price of the final good given by

$$P_t = \left[\int_0^1 P_{l,t}^{1-\sigma} dl \right]^{\frac{1}{1-\sigma}}.$$

The conglomerate, in turn, choose intermediate inputs $Y_{l,t}$ and shares $\theta_{l,t}$ to maximize the shareholder value

$$E_0^* \sum_{t=0}^{\infty} M_0^t D_t,$$

where D_t is the dividend payout to the households

$$D_t = Y_t + \int (D_{l,t} + P_{l,t}^e) \theta_{l,t-1} dl - \int P_{l,t}^e \theta_{l,t} dl.$$

The household side of the economy is the same as in the baseline model except that the households hold shares θ_t of conglomerates instead of shares $\theta_{l,t}$ of intermediate firms.

We now reintroduce aggregate shocks and utilization and describe the timing of the event at period t.

Stage 1 : Pre-production stage

- Agents observe the realization of aggregate shocks (A_t and g_t).
- Given forecasts about $z_{l,t}$ and its associated worst-case scenario, firms make utilization decision, hire labor, and choose price ($U_{l,t}$, $H_{l,t}$, and $P_{l,t}$). The household supply labor H_t and the labor market clears at the wage rate W_t.
- Firms produce intermediate output $Y_{l,t}$ and sells it to the conglomerates at price $P_{l,t}$.

Stage 2 : Post-production stage

- $z_{l,t}$ realize (but are unobservable) and production of the final goods Y_t takes place. Agents observe noisy signals $s_{l,t}$.
- Firms and conglomerates update estimates about $z_{l,t}$ and use it to form forecasts for production next period.

- Firms make investment $I_{l,t}$ and pay out dividends $D_{l,t}$ to the conglomerates. The conglomerates make asset purchase decisions $\theta_{l,t}$ and pay out dividends D_t to the households. Finally, households make consumption and asset purchase decisions (C_t, B_t, and θ_t).

In the perfect competition limit ($\sigma \to \infty$), this version of the model is observationally equivalent to the baseline model at the aggregate level. The introduction of the conglomerate is important for two reasons. First, it prohibits households from inferring $z_{l,t}$ from utility by directly consuming intermediate goods $Y_{l,t}$. Second, it generates countercyclical ex-post excess return on equity held by the household. This is because dividend payout by the conglomerate to the household (D_t) is based on the realized return on capital. Note that the dividend payout by the intermediate firms to the conglomerate ($D_{l,t}$) is not based on the realized return since the production and market clearing of $Y_{l,t}$ happens before the realization of $z_{l,t}$.

8.2 Equilibrium conditions for the estimated model

As we describe below in Appendix 8.3, we express equilibrium conditions from the perspective of agents at both stage 1 and stage 2. At stage 1, we need not only equilibrium conditions for variable determined before production (such as utilization and hours), but also those for variables determined after production (such as consumption and investment). At stage 2, we treat variables determined before production as pre-determined. To do this, we index period t variables determined at stage 1 by $t-1$ and period t variables determined at stage 2 by t. We then combine stage 1 and stage 2 equilibrium conditions by using the certainty equivalence property of linearized decision rules.

We scale the variables in order to introduce stationary:

$$c_t = \frac{C_t}{\gamma^t}, y_{l,t} = \frac{Y_{l,t}}{\gamma^t}, k_{l,t-1} = \frac{K_{l,t-1}}{\gamma^t}, i_{l,t} = \frac{I_{l,t}}{\gamma^t}, f_{l,t} = \frac{F_{l,t-1}}{\gamma^t}, w_t = \frac{W_t}{\gamma^t}, \tilde{\lambda}_t = \gamma^t \lambda_t, \tilde{\mu}_{l,t} = \gamma^t \mu_{l,t},$$

where $\mu_{l,t}$ is the Lagrangian multiplier on the capital accumulation equation of firm l.

We first describe the stage 1 equilibrium conditions. An individual firm l's problem is to choose $\{U_{l,t}, K_{l,t}, H_{l,t}, I_{l,t}\}$ to maximize

$$E_t^* \sum_{s=0}^{\infty} \beta^{t+s} \lambda_{t+s} [P_{t+s}^W Y_{l,t+s} - W_{t+s} H_{l,t+s} - I_{l,t+s} - a(U_{l,t+s}) K_{l,t+s-1}],$$

where P_t^W is the price of whole-sale goods produced by firms and λ_t, and its detrended counterpart $\tilde{\lambda}_t$, is the marginal utility of the representative household:

$$\tilde{\lambda}_t = \frac{\gamma}{c_t - bc_{t-1}} - \beta b E_t^* \frac{1}{\gamma c_{t+1} - bc_t}, \tag{8.2}$$

subject to the following three constraints. The first constraint is the production function:

$$y_{l,t} = A_t \{E_{t-1}^* z_{l,t} f_{l,t} + \nu_{l,t}\}, \tag{8.3}$$

where $f_{l,t}$ is the input,

$$f_{l,t} = (U_{l,t} k_{l,t-1})^\alpha H_{l,t}^{1-\alpha}. \tag{8.4}$$

The worst case TFP $E_t^* z_{l,t+1|t+1}$ is given by

$$E_t^* z_{l,t+1|t+1} = (1-\rho_z)\bar{z} + \rho_z \tilde{z}_{l,t|t} - \eta_a \rho_z \sqrt{\Sigma_{l,t|t}}. \tag{8.5}$$

and the Kalman filter estimate $\tilde{z}_{l,t|t}$ evolves according to

$$\tilde{z}_{l,t|t} = \tilde{z}_{l,t|t-1} + Gain_{l,t}(y_{l,t}/A_t - \tilde{z}_{l,t|t-1} f_{l,t}), \tag{8.6}$$

where $Gain_{l,t}$ is the Kalman gain and is given by

$$Gain_{l,t} = \left[\frac{f_{l,t}^2 \Sigma_{l,t|t-1}}{f_{l,t}^2 \Sigma_{l,t|t-1} + \sigma_\nu^2}\right] f_{l,t}^{-1}. \tag{8.7}$$

The second constraint is the capital accumulation equation:

$$\gamma k_{l,t} = (1-\delta) k_{l,t-1} + \left\{1 - \frac{\kappa}{2}\left(\frac{\gamma i_{l,t}}{i_{l,t-1}} - \gamma\right)^2\right\} i_{l,t} \tag{8.8}$$

and the last constraint is the law of motion for posterior variance:

$$\Sigma_{l,t|t} = (1 - Gain_{l,t} f_{l,t}) \Sigma_{l,t|t-1}. \tag{8.9}$$

As described in the main text, firms take into account the impact of their input choice on worst-case probabilities.

The first-order necessary conditions for firms' input choices are as follows:

- FONC for $\Sigma_{l,t|t}$

$$\begin{aligned}\psi_{l,t} =&\beta E_t^* \Bigg[\frac{1}{2}\tilde{\lambda}_{t+1} P_{t+1}^W A_{t+1} \eta_a \rho_z \Sigma_{l,t|t}^{-\frac{1}{2}} f_{l,t+1} \\ &+ \psi_{l,t+1} \Bigg\{ \frac{\sigma_\nu^2 \rho_z^2}{f_{l,t+1}^2(\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2} - \frac{\sigma_\nu^2 \rho_z^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) f_{l,t+1}^2}{\{f_{l,t+1}^2(\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2\}^2} \Bigg\} \Bigg],\end{aligned} \tag{8.10}$$

where $\psi_{l,t}$ is the Lagrangian multiplier for the law of motion of posterior variance.

- FONC for $U_{l,t}$

$$\begin{aligned}&\tilde{\lambda}_t P_t^W \alpha \frac{y_{l,t}}{U_{l,t}} + \psi_{l,t} \frac{2\alpha \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) f_{l,t}^2}{\{f_{l,t}^2(\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) + \sigma_\nu^2\}^2 U_{l,t}} \\ =&\tilde{\lambda}_t \{\chi_1 \chi_2 U_{l,t} + \chi_2 (1 - \chi_1)\} k_{l,t-1}\end{aligned} \tag{8.11}$$

- FONC for $H_{l,t}$

$$\tilde{\lambda}_t (1-\alpha) P_t^W \frac{y_{l,t}}{H_{l,t}} + \psi_{l,t} \frac{2(1-\alpha)\sigma_\nu^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) f_{l,t}^2}{\{f_{l,t}^2(\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) + \sigma_\nu^2\}^2 H_{l,t}} = \tilde{\lambda}_t \tilde{w}_t, \tag{8.12}$$

where $\tilde{w}_t$ is the real wage: $\tilde{w}_t \equiv w_t / P_t$.

- FONC for $k_{l,t}$

$$\gamma \tilde{\lambda}_t = \beta E_t^* (1 - \Delta_t^k) \tilde{\lambda}_{t+1} \frac{R_{t+1}^k}{\pi_{t+1}}, \tag{8.13}$$

where the return on capital is defined as

$$\begin{aligned}R_t^k =&\Bigg[P_t^W \alpha \frac{y_{l,t}}{k_{l,t-1}} + q_{l,t}(1-\delta) - a(U_{l,t}) \\ &+ \psi_{l,t}^k \frac{2\alpha \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) f_{l,t}^2}{\{f_{l,t}^2(\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) + \sigma_\nu^2\}^2 k_{l,t-1}} \Bigg] \times \frac{\pi_t}{q_{l,t-1}},\end{aligned} \tag{8.14}$$

where

$$q_{l,t} = \tilde{\mu}_{l,t} / \tilde{\lambda}_t, \tag{8.15}$$

$$\psi_{l,t}^k = \psi_{l,t} / \tilde{\lambda}_t. \tag{8.16}$$

- FONC for $i_{l,t}$

$$\begin{aligned}\gamma\tilde{\lambda}_t =&\gamma\tilde{\mu}_{l,t}\left[1-\frac{\kappa}{2}\left(\frac{\gamma i_{l,t}}{i_{l,t-1}}-\gamma\right)^2-\kappa\left(\frac{\gamma i_{l,t}}{i_{l,t-1}}-\gamma\right)\frac{\gamma i_{l,t}}{i_{l,t-1}}\right]\\&+\beta E_t^*\left[\tilde{\mu}_{l,t+1}\kappa\left(\frac{\gamma i_{l,t+1}}{i_{l,t}}-\gamma\right)\left(\frac{\gamma i_{l,t+1}}{i_{l,t}}\right)^2\right]\end{aligned} \tag{8.17}$$

We eliminate l-subscripts to denote cross-sectional means (e.g., $y_t \equiv \int_0^1 y_{l,t}dl$).

Firms sell their wholesale goods to monopolistically competitive retailers. Conditions associated with Calvo sticky prices are

$$P_t^n = \tilde{\lambda}_t P_t^W y_t + \xi_p\beta E_t^*\left(\frac{\pi_{t+1}}{\bar{\pi}}\right)^{\theta_p} P_{t+1}^n \tag{8.18}$$

$$P_t^d = \tilde{\lambda}_t y_t + \xi_p\beta E_t^*\left(\frac{\pi_{t+1}}{\bar{\pi}}\right)^{\theta_p-1} P_{t+1}^d \tag{8.19}$$

$$p_t^* = \left(\frac{\theta_p}{\theta_p-1}\right)\frac{P_t^n}{P_t^d} \tag{8.20}$$

$$1 = (1-\xi_p)(p_t^*)^{1-\theta_p} + \xi_p\left(\frac{\bar{\pi}}{\pi_t}\right)^{1-\theta_p} \tag{8.21}$$

$$\tilde{y}_t = \tilde{p}_t^{-\theta_p} y_t \tag{8.22}$$

$$\tilde{p}_t = (1-\xi_p)(\tilde{p}_t)^{-\theta_p} + \xi_p\left(\frac{\bar{\pi}}{\pi_t}\right)^{-\theta_p} \tag{8.23}$$

Conditions associated with Calvo sticky wages are

$$v_t^1 = v_t^2 \tag{8.24}$$

$$v_t^1 = (w_t^*)^{1-\theta_w}\tilde{\lambda}_t H_t \tilde{w}_t + \xi_w\beta E_t^*\left(\frac{\pi_{t+1}^w w_{t+1}^*}{\bar{\pi}w_t^*}\right)^{\theta_w-1} v_{t+1}^1 \tag{8.25}$$

$$v_t^2 = \frac{\theta_w}{\theta_w-1}(w_t^*)^{-\theta_w(1+\eta)}H_t^{1+\eta} + \xi_w\beta E_t^*\left(\frac{\pi_{t+1}^w w_{t+1}^*}{\bar{\pi}w_t^*}\right)^{\theta_w(1+\eta)} v_{t+1}^2 \tag{8.26}$$

$$1 = (1-\xi_w)(w_t^*)^{1-\theta_w} + \xi_w E_t^*\left(\frac{\bar{\pi}}{\pi_t^w}\right)^{1-\theta_w} \tag{8.27}$$

$$\pi_t^w = \pi_t\tilde{w}_t/\tilde{w}_{t-1} \tag{8.28}$$

Households' Euler equation:

$$\gamma\tilde{\lambda}_t = \beta E_t^* \tilde{\lambda}_{t+1} R_t/\pi_{t+1} \tag{8.29}$$

Monetary policy rule:

$$\frac{R_t}{\bar{R}} = \left(\frac{R_{t-1}}{\bar{R}}\right)^{\rho_R} \left[\left(\frac{\pi_t}{\bar{\pi}}\right)^{\phi_\pi} \left(\frac{y_t}{\bar{y}}\right)^{\phi_Y}\right]^{1-\rho_R} \epsilon_{R,t} \tag{8.30}$$

Resource constraint:

$$c_t + i_t = (1 - g_t)y_t \tag{8.31}$$

The 30 endogenous variables we solve are:

$$k_t, y_t, i_t, c_t, H_t, U_t, f_t, \tilde{\lambda}_t, \tilde{\mu}_t, \psi_t, R_t, R_t^k, q_t, \psi_t^k, E_t^* z_{t+1}, \tilde{z}_{t|t}, \Sigma_{t|t}, Gain_t,$$
$$P_t^W, P_t^n, P_t^d, p_t^*, \pi_t, \tilde{y}_t, \tilde{p}_t, v_t^1, v_t^2, \tilde{w}_t, w_t^*, \pi_t^w$$

We have listed 30 conditions above, from (8.2) to (8.31). Of the above 30 endogenous variables, those that are determined at stage 1 are:

$$H_t, U_t, f_t, v_t^1, v_t^2, \tilde{w}_t, w_t^*, \pi_t^w$$

We now describe the state 2 equilibrium conditions. To avoid repetitions, we only list conditions that are different from the state 1 conditions.

(8.3):

$$y_{l,t} = A_t\{f_{l,t-1} + \nu_{l,t}\}, \quad y_{l,t+s} = A_{t+s}\{E_t^* z_{l,t+s} f_{l,t+s-1} + \nu_{l,t+s}\},$$

where $s \geq 1$.

(8.4):

$$f_{l,t} = (U_{l,t} k_{l,t})^\alpha H_{l,t}^{1-\alpha}$$

(8.6):

$$\tilde{z}_{l,t|t} = \tilde{z}_{l,t|t-1} + Gain_{l,t}(y_{l,t}/A_t - \tilde{z}_{l,t|t-1} f_{l,t-1})$$

(8.7):

$$Gain_{l,t} = \left[\frac{f_{l,t-1}^2 \Sigma_{l,t|t-1}}{f_{l,t-1}^2 \Sigma_{l,t|t-1} + \sigma_\nu^2}\right] f_{l,t-1}^{-1}$$

(8.9):

$$\Sigma_{l,t|t} = (1 - Gain_{l,t} f_{l,t-1}) \Sigma_{l,t|t-1}$$

(8.10):

$$\psi_{l,t} = \beta E_t^* \Bigg[\frac{1}{2} \tilde{\lambda}_{t+1} P_{t+1}^W A_{t+1} \eta_a \rho_z \Sigma_{l,t|t}^{-\frac{1}{2}} f_{l,t} \\ + \psi_{l,t+1} \left\{ \frac{\sigma_\nu^2 \rho_z^2}{f_{l,t}^2(\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2} - \frac{\sigma_\nu^2 \rho_z^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) f_{l,t}^2}{\{f_{l,t}^2(\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2\}^2} \right\} \Bigg]$$

(8.11):

$$E_t^* \left[\tilde{\lambda}_{t+1} P_{t+1}^W \alpha \frac{y_{l,t+1}}{U_{l,t}} + \psi_{l,t+1} \frac{2\alpha \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) f_{l,t}^2}{\{f_{l,t}^2(\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2\}^2 U_{l,t}} \right] \\ = E_t^* \tilde{\lambda}_{t+1} \{\chi_1 \chi_2 U_{l,t} + \chi_2(1 - \chi_1)\} k_{l,t}$$

(8.12):

$$E_t^* \left[\tilde{\lambda}_{t+1} (1-\alpha) P_{t+1}^W \frac{y_{l,t+1}}{H_{l,t}} + \psi_{l,t+1} \frac{2(1-\alpha) \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) f_{l,t}^2}{\{f_{l,t}^2(\rho_z^2 \Sigma_{l,t|t} + \sigma_z^2) + \sigma_\nu^2\}^2 H_{l,t}} \right] = E_t^* \tilde{\lambda}_{t+1} \tilde{w}_t$$

(8.14):

$$R_t^k = \Bigg[P_t^W \alpha \frac{y_{l,t}}{k_{l,t-1}} + q_{l,t}(1-\delta) - a(U_{l,t-1}) \\ + \psi_{l,t}^k \frac{2\alpha \sigma_\nu^2 (\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) f_{l,t-1}^2}{\{f_{l,t-1}^2(\rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2) + \sigma_\nu^2\}^2 k_{l,t-1}} \Bigg] \times \frac{\pi_t}{q_{l,t-1}}$$

(8.25):

$$v_t^1 = (w_t^*)^{1-\theta_w} E_t^* \tilde{\lambda}_{t+1} H_t \tilde{w}_t + \xi_w \beta E_t^* \left(\frac{\pi_{t+1}^w w_{t+1}^*}{\bar{\pi} w_t^*} \right)^{\theta_w - 1} v_{t+1}^1$$

(8.28):

$$\pi_t^w = E_t^* \pi_{t+1} \tilde{w}_t / \tilde{w}_{t-1}$$

8.3 Solution procedure

Here we describe the general solution procedure of the model. The procedure follows the method used in the example in Section 4.2. First, we derive the law of motion assuming that the model is a rational expectations model where the worst case expectations are on average correct. Second, we take the equilibrium law of motion formed under ambiguity and then evaluate the dynamics under the econometrician's data generating process. We provide a step-by-step description of the procedure:

1. Find the worst-case steady state.

   We first compute the steady state of the filtering problem (3.6), (3.7), (3.8), and (3.10), under the worst-case mean (3.12) to find the firm-level TFP at the worst-case steady state, $\bar{z}^0$. We then solve the steady state for other equilibrium conditions evaluated at $\bar{z}^0$.

2. Log-linearize the model around the worst-case steady state.

   We can solve for the dynamics using standard tools for linear rational expectation models. We base our discussion based on the method proposed by Sims (2002).

   We first need to deal with the issue that idiosyncratic shocks realize at the beginning of stage 2. Handling this issue correctly is important, since variables chosen at stage 1, such as input choice, should be based on the worst-case TFP, while variables chosen at stage 2, such as consumption and investment, would be based on the realized TFP (but also on the worst-case future TFP). To do this, we exploit the certainty equivalence property of linear decision rules. We first solve for decision rules *as if* both aggregate and idiosyncratic shocks realize at the beginning of the period. We call them "pre-production decision rules". We then solve for decision rules *as if* (i) both aggregate and idiosyncratic shocks realize at the beginning of the period and (ii) stage 1 variables are pre-determined. We call them "post-production decision rules". Finally, when we characterize the dynamics from the perspective of the econometrician, we combine the pre-production and post-production decision rules and obtain and equilibrium law of motion.

   To obtain pre-production decision rules, we collect the linearized equilibrium conditions, which include firm-level conditions, into the canonical form:

   $$\mathbf{\Gamma}_0^{pre}\hat{\mathbf{y}}_t^{pre,0} = \mathbf{\Gamma}_1^{pre}\hat{\mathbf{y}}_{t-1}^{pre,0} + \mathbf{\Psi}^{pre}\omega_t + \mathbf{\Upsilon}^{pre}\eta_t^{pre},$$

   where $\hat{\mathbf{y}}_t^{pre,0}$ is a column vector of size k that contains all variables and the conditional expectations. $\hat{\mathbf{y}}_t^{pre,0} = \mathbf{y}_t^{pre} - \bar{\mathbf{y}}^0$ denotes deviations from the worst-case steady state and η_t are expectation errors, which we define as $\eta_t^{pre} = \hat{\mathbf{y}}_t^{pre,0} - E_{t-1}^*\hat{\mathbf{y}}_t^{pre,0}$ such that $E_{t-1}^*\eta_t^{pre} = 0$. We define $\omega_t = [e_{l,t} \quad e_t]'$, where $e_{l,t} = [\epsilon_{z,l,t} \quad u_{l,t} \quad \nu_{l,t}]'$ is a vector of idiosyncratic shocks and e_t is a vector of aggregate shocks of size n. For example, for the baseline model introduced in Section 3, e_t is a 2×1 vector of aggregate TFP and government spending shocks.

The vector $\hat{\mathbf{y}}_t^{pre,0}$ contains firm-level variables such as firm l's labor input, $H_{l,t}$. In contrast to other linear heterogeneous-agent models with imperfect information such as Lorenzoni (2009), all agents share the same information set. Thus, to derive the aggregate law of motion, we simply aggregate over firm l's linearized conditions and replace firm-specific variables with their cross-sectional means (e.g., we replace $H_{l,t}$ with $H_t \equiv \int_0^1 H_{l,t} dl$) and set $e_{l,t} = \mathbf{0}$, which uses the law of large numbers for idiosyncratic shocks.

We order variables in $\hat{\mathbf{y}}_t^{pre,0}$ as

$$\hat{\mathbf{y}}_t^{pre,0} = \begin{bmatrix} \hat{\mathbf{y}}_{1,t}^{pre,0} \\ \hat{\mathbf{y}}_{2,t}^{pre,0} \\ \hat{\mathbf{s}}_t^{pre,0} \end{bmatrix},$$

where $\hat{\mathbf{y}}_{1,t}^{pre,0}$ is a column vector of size k_1 of variables determined at stage 1, $\hat{\mathbf{y}}_{2,t}^{pre,0}$ is a column vector of size k_2 of variables determined at stage 2, and $\hat{\mathbf{s}}_t^{pre,0} = [\hat{s}_{1,t}^{pre,0} \quad \hat{s}_{2,t}^{pre,0}]'$, where $s_{1,t} = \bar{z} - E_{t-1}^* z_t$ and $s_{2,t} = \bar{z} - \tilde{z}_{t|t}$.

The resulting solution of pre-production decision rules is obtained applying the method developed by Sims (2002):

$$\hat{\mathbf{y}}_t^{pre,0} = \mathbf{T}^{pre} \hat{\mathbf{y}}_{t-1}^{pre,0} + \mathbf{R}^{pre} [\mathbf{0}_{3\times 1} \quad e_t]', \tag{8.32}$$

where $\mathbf{T}^{pre}$ and $\mathbf{R}^{pre}$ are $k \times k$ and $k \times (n+3)$ matrices, respectively.

The solution of post-production decision rules can be obtained in a similar way by first collecting the equilibrium conditions into the canonical form

$$\mathbf{\Gamma}_0^{post} \hat{\mathbf{y}}_t^{post,0} = \mathbf{\Gamma}_1^{post} \hat{\mathbf{y}}_{t-1}^{post,0} + \mathbf{\Psi}^{post} \omega_t + \mathbf{\Upsilon}^{post} \eta_t^{post},$$

and is given by

$$\hat{\mathbf{y}}_t^{post,0} = \mathbf{T}^{post} \hat{\mathbf{y}}_{t-1}^{post,0} + \mathbf{R}^{post} [\mathbf{0}_{3\times 1} \quad e_t]', \tag{8.33}$$

where

$$\hat{\mathbf{y}}_t^{post,0} = \begin{bmatrix} \hat{\mathbf{y}}_{1,t}^{post,0} \\ \hat{\mathbf{y}}_{2,t}^{post,0} \\ \hat{\mathbf{s}}_t^{post,0} \end{bmatrix},$$

and $\mathbf{T}^{post}$ and $\mathbf{R}^{post}$ are $k \times k$ and $k \times (n+3)$ matrices, respectively.

3. Characterize the dynamics from the econometrician's perspective.

The above law of motion was based on the worst-case probabilities. We need to derive the equilibrium dynamics under the true DGP, where the cross-sectional mean of firm-level TFP is $\bar{z}$. We are interested in two objects: the zero-risk steady state and the dynamics around that zero-risk steady state.

(a) Find the zero-risk steady state.

This the fixed point $\bar{\mathbf{y}}$ where the decision rules (8.32) and (8.33) are evaluated at the realized cross-sectional mean of firm-level TFP $\bar{z}$:

$$\begin{aligned}\bar{\mathbf{y}}^{pre} - \bar{\mathbf{y}}^0 &= \mathbf{T}^{pre}(\bar{\mathbf{y}} - \bar{\mathbf{y}}^0),\\ \bar{\mathbf{y}}^{post} - \bar{\mathbf{y}}^0 &= \mathbf{T}^{post}(\bar{\mathbf{y}} - \bar{\mathbf{y}}^0) + \mathbf{R}^{post}[\bar{\mathbf{s}} \quad \mathbf{0}_{(n+1)\times 1}]',\end{aligned} \tag{8.34}$$

where

$$\bar{\mathbf{y}} = \begin{bmatrix} \bar{\mathbf{y}}_1^{pre} \\ \bar{\mathbf{y}}_2^{post} \\ \bar{\mathbf{s}}^{post} \end{bmatrix}.$$

Note that we do not feed in the realized firm-level TFP to the pre-production decision rules since idiosyncratic shocks realize at the beginning of stage 2.

We obtain $\bar{\mathbf{s}}$ from

$$\bar{\mathbf{s}} = [\mathbf{T}_{3,1}^{post} \quad \mathbf{T}_{3,2}^{post} \quad \mathbf{T}_{3,3}^{post}](\bar{\mathbf{y}} - \bar{\mathbf{y}}^0) + \bar{\mathbf{s}}^0,$$

where $\mathbf{T}_{3,1}^{post}$, $\mathbf{T}_{3,2}^{post}$, and $\mathbf{T}_{3,3}^{post}$ are $2 \times k_1$, $2 \times k_2$, and 2×2 submatrices of $\mathbf{T}^{post}$, respectively:

$$\mathbf{T}^{post} = \begin{bmatrix} \mathbf{T}_{1,1}^{post} & \mathbf{T}_{1,2}^{post} & \mathbf{T}_{1,3}^{post} \\ \mathbf{T}_{2,1}^{post} & \mathbf{T}_{2,2}^{post} & \mathbf{T}_{2,3}^{post} \\ \mathbf{T}_{3,1}^{post} & \mathbf{T}_{3,2}^{post} & \mathbf{T}_{3,3}^{post} \end{bmatrix},$$

where $\mathbf{T}_{1,1}^{post}$, $\mathbf{T}_{1,2}^{post}$, $\mathbf{T}_{1,3}^{post}$, $\mathbf{T}_{2,1}^{post}$, $\mathbf{T}_{2,2}^{post}$, and $\mathbf{T}_{2,3}^{post}$ are $k_1 \times k_1$, $k_1 \times k_2$, $k_1 \times 2$, $k_2 \times k_1$, $k_2 \times k_2$, and $k_2 \times 2$ matrices, respectively.

(b) Dynamics around the zero-risk steady state.

Denoting $\hat{\mathbf{y}}_t \equiv \mathbf{y}_t - \bar{\mathbf{y}}$ the deviations from the zero-risk steady state, we combine the decision rules (8.32) and (8.33) evaluated at the true DGP and the equations for the zero-risk steady state (8.34) to characterize the equilibrium law of motion:

$$\hat{\mathbf{y}}_t^{pre} = \mathbf{T}^{pre}\hat{\mathbf{y}}_{t-1} + \mathbf{R}^{pre}[\mathbf{0}_{3\times1} \quad e_t]', \tag{8.35}$$
$$\hat{\mathbf{y}}_t^{post} = \mathbf{T}^{post}[\hat{\mathbf{y}}_{1,t}^{pre} \quad \hat{\mathbf{y}}_{2,t-1} \quad \hat{\mathbf{s}}_{t-1}]' + \mathbf{R}^{post}[\hat{\hat{\mathbf{s}}}_t \quad 0 \quad e_t]', \tag{8.36}$$
$$\hat{\hat{\mathbf{s}}}_t = [\mathbf{T}_{3,1}^{post} \quad \mathbf{T}_{3,2}^{post} \quad \mathbf{T}_{3,3}^{post}][\hat{\mathbf{y}}_{1,t}^{pre} \quad \hat{\mathbf{y}}_{2,t-1} \quad \hat{\mathbf{s}}_{t-1}]' + \mathbf{R}_{3,3}^{post}[\mathbf{0}_{3\times1} \quad e_t]', \tag{8.37}$$

and

$$\hat{\mathbf{y}}_t = \begin{bmatrix} \hat{\mathbf{y}}_{1,t}^{pre} \\ \hat{\mathbf{y}}_{2,t}^{post} \\ \hat{\mathbf{s}}_t^{post} \end{bmatrix}. \tag{8.38}$$

$\mathbf{R}_{3,3}^{post}$ is a $2 \times n$ submatrix of $\mathbf{R}^{post}$:

$$\mathbf{R}^{post} = \begin{bmatrix} \mathbf{R}_{1,1}^{post} & \mathbf{R}_{1,2}^{post} & \mathbf{R}_{1,3}^{post} \\ \mathbf{R}_{2,1}^{post} & \mathbf{R}_{2,2}^{post} & \mathbf{R}_{2,3}^{post} \\ \mathbf{R}_{3,1}^{post} & \mathbf{R}_{3,2}^{post} & \mathbf{R}_{3,3}^{post} \end{bmatrix},$$

where $\mathbf{R}_{1,1}^{post}$, $\mathbf{R}_{1,2}^{post}$, $\mathbf{R}_{1,3}^{post}$, $\mathbf{R}_{2,1}^{post}$, $\mathbf{R}_{2,2}^{post}$, $\mathbf{R}_{2,3}^{post}$, $\mathbf{R}_{3,1}^{post}$, and $\mathbf{R}_{3,2}^{post}$ are $k_1 \times 2$, $k_1 \times 1$, $k_1 \times n$, $k_2 \times 2$, $k_2 \times 1$, $k_2 \times n$, 2×2, and 2×1 matrices, respectively.

We combine equations (8.35), (8.36), (8.37), and (8.38) to obtain the equilibrium law of motion. To do so, we first define submatrices of $\mathbf{T}^{pre}$ and $\mathbf{R}^{pre}$:

$$\mathbf{T}^{pre} = \begin{bmatrix} \mathbf{T}_1^{pre} \\ \mathbf{T}_2^{pre} \\ \mathbf{T}_3^{pre} \end{bmatrix},$$

where $\mathbf{T}_1^{pre}$, $\mathbf{T}_2^{pre}$, and $\mathbf{T}_3^{pre}$ are $k_1 \times k$, $k_2 \times k$, and $2 \times k$ matrices, respectively, and

$$\mathbf{R}^{pre} = \begin{bmatrix} \mathbf{R}_{1,1}^{pre} & \mathbf{R}_{1,2}^{pre} \\ \mathbf{R}_{2,1}^{pre} & \mathbf{R}_{2,2}^{pre} \\ \mathbf{R}_{3,1}^{pre} & \mathbf{R}_{3,2}^{pre} \end{bmatrix},$$

where $\mathbf{R}_{1,1}^{pre}$, $\mathbf{R}_{1,2}^{pre}$, $\mathbf{R}_{2,1}^{pre}$, $\mathbf{R}_{2,2}^{pre}$, $\mathbf{R}_{3,1}^{pre}$, and $\mathbf{R}_{3,2}^{pre}$ are $k_1 \times 3$, $k_1 \times n$, $k_2 \times 3$, $k_2 \times n$, 2×3, and $2 \times n$ matrices, respectively.

We then define matrices $\mathbf{T}$ and $\mathbf{R}$. A $k \times k$ matrix $\mathbf{T}$ is given by

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_1^{pre} \\ \mathbf{T}_2 \\ \mathbf{T}_3 \end{bmatrix},$$

where $\mathbf{T}_2$ and $\mathbf{T}_3$ are $k_2 \times k_2$ and 2×2 matrices, respectively, given by

$$\mathbf{T}_2 = [\mathbf{Q}_{2,1} \quad \mathbf{Q}_{2,2} + \mathbf{T}_{2,2}^{post} + \mathbf{R}_{2,1}^{post}\mathbf{T}_{3,2}^{post} \quad \mathbf{Q}_{2,3} + \mathbf{T}_{2,3}^{post} + \mathbf{R}_{2,1}^{post}\mathbf{T}_{3,3}^{post}],$$
$$\mathbf{T}_3 = [\mathbf{Q}_{3,1} \quad \mathbf{Q}_{3,2} + \mathbf{T}_{3,2}^{post} + \mathbf{R}_{3,1}^{post}\mathbf{T}_{3,2}^{post} \quad \mathbf{Q}_{3,3} + \mathbf{T}_{3,3}^{post} + \mathbf{R}_{3,1}^{post}\mathbf{T}_{3,3}^{post}],$$

and $\mathbf{Q}_{2,1}$, $\mathbf{Q}_{2,2}$, and $\mathbf{Q}_{2,3}$ are $k_2 \times k_1$, $k_2 \times k_2$, and $k_2 \times 2$ submatrices of $\mathbf{Q}_2$, where $\mathbf{Q}_2 \equiv (\mathbf{T}_{2,1}^{post} + \mathbf{R}_{2,1}^{post}\mathbf{T}_{3,1}^{post})\mathbf{T}_1^{pre}$, so that $\mathbf{Q}_2 = [\mathbf{Q}_{2,1} \quad \mathbf{Q}_{2,2} \quad \mathbf{Q}_{2,3}]$. Similarly, $\mathbf{Q}_{3,1}$, $\mathbf{Q}_{3,2}$, and $\mathbf{Q}_{3,3}$ are $k_3 \times k_1$, $k_3 \times k_2$, and $k_3 \times 2$ submatrices of $\mathbf{Q}_3$, where $\mathbf{Q}_3 \equiv (\mathbf{T}_{3,1}^{post} + \mathbf{R}_{3,1}^{post}\mathbf{T}_{3,1}^{post})\mathbf{T}_1^{pre}$, so that $\mathbf{Q}_3 = [\mathbf{Q}_{3,1} \quad \mathbf{Q}_{3,2} \quad \mathbf{Q}_{3,3}]$.

A $k \times n$ matrix $\mathbf{R}$ is given by

$$\mathbf{R} = \begin{bmatrix} \mathbf{R}_{1,2}^{pre} \\ \mathbf{R}_2 \\ \mathbf{R}_3 \end{bmatrix},$$

where

$$\mathbf{R}_2 = \mathbf{T}_{2,1}^{post}\mathbf{R}_{1,2}^{pre} + \mathbf{R}_{2,1}^{post}(\mathbf{T}_{3,1}^{post}\mathbf{R}_{1,2}^{pre} + \mathbf{R}_{3,3}^{post}) + \mathbf{R}_{2,3}^{post},$$
$$\mathbf{R}_3 = \mathbf{T}_{3,1}^{post}\mathbf{R}_{1,2}^{pre} + \mathbf{R}_{3,1}^{post}(\mathbf{T}_{3,1}^{post}\mathbf{R}_{1,2}^{pre} + \mathbf{R}_{3,3}^{post}) + \mathbf{R}_{3,3}^{post}.$$

The equilibrium law of motion is then given by

$$\hat{\mathbf{y}}_t = \mathbf{T}\hat{\mathbf{y}}_{t-1} + \mathbf{R}e_t.$$

8.4 Bayesian estimation: the role of the credit spread

We argue that the inclusion of the spread in the estimation is the key reason why in our model the financial wedge shock becomes an important driver of the business cycle. To see this, we first re-estimate our flexible-price and sticky-price models and their RE versions without using the credit spread in the estimation. The column labeled "No spread used" in Table 10 reports the contributions of the financial wedge shock in explaining the variance of output. First, consider the flexible-price model (Panel A). In line with the main estimation that uses the spread, in the RE version the financial wedge shock can only explain a negligible share of output variance. This is not surprising, since in the main estimation the measurement error for the spread was quite large. In contrast to the main estimation, however, the contribution of the financial wedge shock is negligible even with endogenous uncertainty. The intuition is as follows. In the flexible-price model, the model cannot generate co-movement

Table 10: Output variance explained by the financial wedge shock: the role of credit spread

Model	Spread not used	No meas. error in spread
Panel A: Flexible price		
Baseline	0.00	0.13
RE	0.00	0.04
Panel B: Sticky price		
Baseline	0.48	0.72
RE	0.15	0.15

Notes: We report the percent of variance of output at the business cycle frequency (6–32 quarters) that can be explained by the financial wedge shock. "Spread not used" refers to results from estimations where the credit spread is not used as an observable. "No meas. error in spread" refers to results from estimations where the credit spread is observed without a meaurement error.

of consumption and hours in response to the financial wedge shock in the short run (see Figure 3). The likelihood estimator thus tries to reduce the size of the financial wedge shock as much as possible.[34] This is feasible since the credit spread is not used in the estimation. Next, consider the sticky-price model (Panel B). As in the flexible-price model, the result of the RE version is in line with the main estimation. With endogenous uncertainty, the financial wedge shock becomes more important but its contribution to output variance is smaller than the main estimation (48 percent vs. 73 percent). In the data, the credit spread is strongly countercyclical, which pushes the estimation towards assigning a larger role to the financial wedge shock when the spread is used as an observable.

We also consider what happens when we estimate our models with no measurement error in the spread (Table 10, "No meas. error in spread"). For both flexible-price and sticky-price models, the results of the baseline models with endogenous uncertainty are very similar to those from the main estimation. This is because the estimated measurement errors for the spread for the baseline models are very small in the main estimation to begin with. For the RE version of the flexible-price model, the financial wedge shock explains a small fraction of output (4 percent) and it comes at the cost of countercyclical consumption; the correlation between output growth and consumption growth is -0.19. For the sticky-price RE model, the financial wedge shock explains 15 percent of output variance but the estimated real rigidities are large; the consumption habit is $b = 0.92$ and the investment adjustment cost is $\kappa = 2.42$.[35] The large real rigidities lead to macro quantities being too smooth. For example, the model

[34]Indeed, the estimated parameters for the financial wedge shock process are $\rho_\Delta = 0.36$ and $100\sigma_\Delta = 0.04$, compared to $\rho_\Delta = 0.53$ and $100\sigma_\Delta = 0.10$ in the main estimation.

[35]The estimated nominal rigidities are similar to those found in the main estimation.

substantially understates the standard deviations of output growth and consumption growth (0.50 and 0.23, respectively, in the model versus 0.60 and 0.50, respectively, in the data).

To summarize, we find that in our model the financial wedge shock becomes an important source of business cycle when the credit spread, measured with or without error, is used in the estimation. Intuitively, the financial wedge shock needs to be large since the spread is volatile and countercyclical. Without endogenous uncertainty, however, this comes at the cost of generating counterfactual predictions for other macro variables.

8.5 Data sources

We use the following data:

1. Real GDP in chained dollars, BEA, NIPA table 1.1.6, line 1.
2. GDP, BEA, NIPA table 1.1.5, line 1.
3. Personal consumption expenditures on nondurables, BEA, NIPA table 1.1.5, line 5.
4. Personal consumption expenditures on services, BEA, NIPA table 1.1.5, line 6.
5. Gross private domestic fixed investment (nonresidential and residential), BEA, NIPA table 1.1.5, line 8.
6. Personal consumption expenditures on durable goods, BEA, NIPA table 1.1.5, line 4.
7. Nonfarm business hours worked, BLS PRS85006033.
8. Civilian noninstitutional population (16 years and over), BLS LNU00000000.
9. Effective federal funds rate, Board of Governors of the Federal Reserve System.
10. Moody's Seasoned Baa Corporate Bond Yield Relative to Yield on 10-Year Treasury Constant Maturity (Baa spread), downloaded from Federal Reserve Economic Data, Federal Reserve Bank of St. Louis.

We then conduct the following transformations of the above data:

11. Real per capita GDP: (1)/(8)
12. GDP deflator: (2)/(1)
13. Real per capita consumption: $[(3)+(4)]/[(8)\times(12)]$
14. Real per capita investment: $[(5)+(6)]/[(8)\times(12)]$
15. Per capita hours: (7)/(8)