Journal of Monetary Economics 117 (2021) 354–376

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jmoneco

Learning, confidence, and business cycles[☆]

Cosmin Ilut[a,*], Hikaru Saijo[b]

[a] Duke University, and NBER, United States
[b] UC Santa Cruz, United States

ARTICLE INFO

Article history:
Received 15 May 2019
Revised 27 January 2020
Accepted 28 January 2020
Available online 30 January 2020

JEL classification:
E3

Keywords:
Business cycles
Learning
Ambiguity
Firm dynamics
Wedges

ABSTRACT

We argue that information accumulation provides a quantitatively successful propagation mechanism that challenges and empirically improves on the conventional New Keynesian models with many nominal and real rigidities. In particular, we build a tractable heterogeneous-firm business cycle model where firms face Knightian uncertainty about their profitability and learn it through production. The feedback between uncertainty and economic activity maps fundamental shocks into an as if procyclical equilibrium confidence process, generating co-movement driven by demand shocks, amplified and hump-shaped dynamics, countercyclical correlated wedges in the equilibrium conditions for labor, risk-free and risky assets, and countercyclical firm-level and aggregate dispersion of forecasts.

1. Introduction

Analysts and policy makers generally view aggregate fluctuations as triggered by impulses that vary across historical episodes, such as excessive monetary policy tightening, technological boom-busts, or disturbances in the financial markets. While these impulses differ in their source, business cycles have remarkably consistent patterns, leading to important restrictions on a theory of propagation of shocks. First, there is positive and persistent co-movement of key aggregate quantities, such as hours worked, consumption, and investment. Second, this co-movement jointly occurs with predictable cross-equation restrictions between quantities and returns, a pattern that the literature refers to as reduced-form countercyclical labor, consumption, and risk premium 'wedges'.[1]

☆ We would like to thank the Editor, two anonymous referees, our discussants Wei Cui, Stefano Eusepi, Christoph Grosse Steffen, as well as George-Marios Angeletos, Yan Bai, Nick Bloom, Fabrice Collard, Jesus Fernandez-Villaverde, Tatsuro Senga, Martin Schneider, Mathieu Taschereau-Dumouchel, Vincenzo Quadrini, Carl Walsh, and many other seminar and conference participants for helpful comments.

* Corresponding author.
E-mail addresses: cosmin.ilut@duke.edu (C. Ilut), hsaijo@ucsc.edu (H. Saijo).

[1] In particular, in a recession, a larger 'labor wedge' appears as hours worked are lower than predicted by the comparison of labor productivity to the marginal rate of substitution between consumption and labor, as analyzed through the lenses of standard preferences and technologies (see Shimer (2009) and Chari et al. (2007) for evidence and discussion). At the same time, a higher 'consumption wedge' manifests as the risk-free return is unusually low compared to realized future aggregate consumption growth (see Christiano et al. (2005) and Smets and Wouters (2007) as examples for a large literature that uses shocks to the discount factor). Finally, a 'risk premium wedge' increases as the excess return on risky assets over the return of risk-free assets is unusually large (see Cochrane (2011) for a review on countercyclical excess returns).

https://doi.org/10.1016/j.jmoneco.2020.01.010
0304-3932/

Conventional quantitative business cycle models typically approach these recurring patterns through the lenses of the New Keynesian (NK) paradigm.[2] This view of propagation relies on a wide array of frictions and shocks to explain the data. First, a large set of real and nominal rigidities, such as habit formation, investment adjustment costs, sticky prices and wages, is used to fit the observed persistence, amplification and co-movement of aggregate quantities. Second, even if endowed with a variety of frictions, quantitative NK models still typically appeal to latent 'wedges', taken as exogenous residuals in the optimality conditions for hours, consumption, and capital accumulation. These residuals appear to an analyst as correlated and countercyclical, since these models view recessions (booms) as periods of 'unusually' low (high) hours worked, real interest rates, and asset prices.

In this paper we argue that information accumulation provides a propagation mechanism that challenges and empirically improves on the conventional view. The friction is based on plausible inference difficulties faced by firms, which are uncertain about their own profitability and learn about it through production. We show how the friction maps fundamental shocks into procyclical movements in confidence about aggregate conditions. Through this mapping a variety of aggregate shocks are propagated into business cycle fluctuations that have similar patterns - persistent positive co-movement and measured countercyclical wedges.

Our theory is consistent with a view shared by analysts and policymakers that various impulses lead to a demand-driven propagation through which 'confidence' or 'uncertainty' affect the aggregate economy's desire to spend, hire and invest.[3] Moreover, our proposed narrative speaks to a recent econometric strategy developed in Angeletos et al. (2018a), in the form of extracting common patterns across business cycles, that supports the empirical need to build models of demand-driven co-movement that do not rely on nominal rigidities.[4]

Our structure of uncertainty is motivated by three factors regarding firm decision making under uncertainty. First, starting from Hymer and Pashigian (1962), a common view in the industrial organization literature is that a firm is a collection of production units, which are subject to a firm and unit-specific profitability shock. The key implication of this view is that the variance of the firm growth rate decreases with firm size. As summarized in Coad (2007), this scaling relationship has been documented by numerous studies, including Stanley et al. (1996) and Bottazzi and Secchi (2003). Second, experimental evidence motivated by Ellsberg (1961) and more recently documented by Bossaerts et al. (2010) and Asparouhova et al. (2015) argues that decision makers perceive uncertainty not only as risk but also as ambiguity. Third, firm-level evidence suggests that the above two factors are closely related. In particular, Senga (2018) shows that forecast errors and forecast dispersions decrease with firm size. As we explain below, we take the forecast dispersion as a measure of ambiguity.

We capture these factors through two key modeling assumptions. On the one hand, firms accumulate information about their unobserved profitability through production. We share this feature with models of learning in the firm dynamics literature, similar to Jovanovic (1982). On the other hand, we assume that facing a larger estimation uncertainty, an ambiguity-averse decision maker is less confident in his probability assessments and entertains a wider set of beliefs about the conditional mean of the persistent firm-specific component. In particular, we model ambiguity-aversion through the recursive multiple priors preferences of Epstein and Schneider (2003b). The preference representation makes an agent facing lower confidence behave *as if* the true unknown mean becomes worse.

We embed this structure of uncertainty into a standard business cycle model with heterogeneous firms and a representative agent. The structure generates a feedback loop at the firm level: Consistent with the scaling relationship, lower production leads to more estimation uncertainty, which in turn shrinks the optimal size of productive inputs. In our model, this firm-level feedback aggregates linearly so that recessions are periods of a high cross-sectional mean of firm-level uncertainty because firms *on average* invest and hire less. In turn, the higher uncertainty, and the implied lower confidence, further dampens aggregate activity.

In particular, when confidence is low, the *uncertainty-adjusted* return to working, to consuming and to investing are *jointly* perceived to be low. This leads to a high measured labor wedge since equilibrium hours worked are low even if consumption is low and the realized marginal product of labor is on average unchanged under the econometrician's data generating process. The endogenous countercyclical labor wedge is a crucial property that explains why labor and consumption can both fall following a contractionary supply or demand shock. The low confidence also leads to a high measured consumption wedge because the increased desire to save depresses the real risk-free rate more than the econometrician's measured growth rate of marginal utility. Finally, it makes capital less attractive to hold so investors are compensated in equilibrium by a higher measured excess return.

As a result of the feedback between production and uncertainty, our mechanism generates as if correlated wedges that respond to the endogenous aggregate states of the economy. As such, these wedges manifest *conditional on any type* of fundamental shock, as long as that shock affects the aggregate states. Importantly, these fundamental shocks can arise in any form, including standard productivity, demand or monetary policy shocks, as well as more recently proposed sources, such as disturbances in the financial sector, exogenous changes in beliefs, perceived volatility or confidence.

[2] Barro and King (1984) emphasize how in a standard RBC model hours and consumption co-move negatively unless there is a total factor productivity (TFP) or a preference shock to the disutility of working. NK models overturn this impossibility result through countercyclical markups.

[3] Baker et al. (2016) document how the word "uncertainty" in leading newspapers and the FOMC's Beige Book spikes up in recessions.

[4] Angeletos (2018) provides a critical analysis of the empirical and theoretical underpinnings of NK models.

We quantitatively evaluate our mechanism by focusing squarely on propagation and employ an impulse-response-matching estimation, using the local projection method of Jordá (2005). We recover responses to four shocks. First, we use the excess bond premium from Gilchrist and Zakrajšek (2012) to measure a financial shock. Second, the utilization-adjusted TFP from Fernald (2014) directly identifies a neutral technology shock. Third, the relative price of investment goods (Fisher (2006)) provides a measure of an investment-specific technology shock. Fourth, we use the Romer and Romer (2004) series, extended by Coibion et al. (2017), for our monetary policy shocks. In addition, we use the observables to construct empirical measures of labor, consumption, and excess return wedges.

Our analysis produces four specific results on the quantitative role of the learning friction. First, as a headline finding, the information friction *alone* provides a narrative for business cycle propagation that is *qualitatively different and quantitatively fits the data better* than the typical rational expectations (RE) NK model augmented with many rigidities. To establish this, we first estimate a model *only* with the information friction, without additional real or nominal rigidities, by fitting only the impulse responses to the financial shock. We do so because this shock is quantitatively important, accounting for a significant fraction of business cycle variation, and informative, as it provides a laboratory for the relevant empirical cross-equation restrictions. This parsimonious model matches the impulse response well. After a reduction in the credit spread faced by entrepreneurs, the model replicates the persistent and hump-shaped dynamics of aggregate quantities as *hours, investment and consumption jointly* rise. It also matches *price dynamics*: inflation is stable and the real rate increases. Finally, the model is consistent with the observed *countercyclical wedges* as the labor, consumption and excess return wedges *jointly* fall.

We then estimate two competing models. One is an RE model with ambiguity turned off but with *all* the remaining rigidities: habit formation, investment adjustment costs, sticky prices and wages. This workhorse RE model can match the hump-shaped increase in output, hours, and investment following an expansionary financial shock. However, that model predicts a reduction, rather than an increase, in consumption, as well as an overprediction of inflation in the medium run. The second comparison is to an otherwise frictionless RE model but augmented with learning-by-doing (LBD) as in Chang et al. (2002). We find that the LBD model generates a simultaneous increase in output, hours, investment, and consumption, but it lacks the endogenous variation in the labor, consumption, and excess return wedges. While in our proposed mechanism a larger scale of production lowers uncertainty about the unobserved firm-level productivity and it increases the *perceived* return to working and investing, in the LBD model a larger scale raises *actual* productivity through accumulation of skill. The two mechanisms therefore have similarities in that learning can generate co-movement of consumption and labor out of a range of shocks. The conceptual difference is that in the LBD model there are no predictable wedges since, by the rational expectations assumption, there is no systematic difference between agents' beliefs and the equilibrium objects recovered by an analyst under the true DGP. Overall, we find that both competing RE versions fit the data worse compared to the learning model.

Our second result is that when the learning friction is embedded into a workhorse NK model, it significantly helps fit the data and reduces the need for standard rigidities. To show this, we match impulse responses to *all four* structural shocks and compare the fit of the standard set of rigidities with a model that *also* includes the information friction. We find that the learning model matches the four sets of impulse responses well and is heavily favored by the marginal data density against a standard RE model.

Our third result is that the estimated information friction generates movements in confidence that are consistent with survey data, used here as external model validation. The key logic, building on Ilut and Schneider (2014), is to map the empirically observed *dispersion of experts' forecasts* into the intervals of forecasts implied by our decision makers' equilibrium *sets* of beliefs.[5] Here we use both firm-level dispersion of forecasts, employing a series constructed by Senga (2018), as well as aggregate-level, where we exploit the dispersion of forecasts from the Survey of Professional Forecasters (SPF).

Our fourth result is on policy implications. In our estimated learning model an interest rate rule that reacts to the credit spread would significantly lower output variability because it stabilizes the variation in endogenous uncertainty. Instead, if the confidence process would be *counterfactually* held fixed at its pre-policy change path, the output variability would be largely unaffected. We also find a significantly larger government spending multiplier arising from the endogenous confidence process. Therefore, modeling confidence as a response to the state of the economy matters crucially for how policies affect equilibrium outcomes.

The endogenous correlation between fundamental shocks and the resulting *as if* confidence process that sustains the equilibrium allocations connects two literatures that suggest the empirical and theoretical appeal of information-driven business cycles. First, the low activity-high uncertainty feedback implied by the friction has been analyzed, in different forms, as a source for business cycle asymmetries, non-linearities, persistence or amplification, in a related learning literature (Caplin and Leahy (1993); Fajgelbaum et al. (2017); van Nieuwerburgh and Veldkamp (2006); Ordoñez (2013); Straub and Ulbricht (2016) and Saijo (2017)). While there the feedback typically matters through non-linear dynamics and learning occurs from aggregate market outcomes, we study an endogenous uncertainty mechanism driven by linear dynamics and learning about firm-level profitability. These two properties lead to a tractable characterization and evaluation of the feedback mechanism even within linear, workhorse quantitative models, as well as to novel policy implications.

[5] The idea is that the representative agent samples experts' opinions and aggregates them, according to his preferences, when making decisions. Since the decision maker is ambiguity averse, stronger disagreement among these professional forecasters generates lower confidence in the probability assessments of the future.

Second, recent work proposes movements in agents' beliefs, typically modeled as exogenous confidence shocks, as important drivers of business cycles (Angeletos et al. (2018b); Angeletos and La'O (2009, 2013); Barsky and Sims (2012); Benhabib and Spiegel (2018); Benhabib et al. (2015); Bhandari et al. (2016); Ilut and Schneider (2014); Milani (2011, 2017), and Huo and Takayama (2015)). Our analysis provides a theory disciplined by micro and macro moments of the formation of those beliefs, in which the confidence process changes endogenously as a response to the state of the economy.[6]

In Section 2 we introduce our learning model and discuss the solution method. We describe the potential of endogenous uncertainty as a parsimonious propagation mechanism in Section 3. In Section 4 we estimate our model on US aggregate data.

2. The model

Our baseline model is a real business cycle model in which, as in the standard framework, firms are owned by a representative household and maximize shareholder value. We augment the standard framework along two key features: the infinitely-lived representative household is ambiguity averse and that ambiguity is about the firm-level profitability processes.

2.1. Technology

There is a continuum of firms, indexed by $l \in [0, 1]$, which act in a monopolistically competitive manner. They rent capital $K_{l,t-1}$ and hire labor $H_{l,t}$ to operate $J_{l,t}$ number of production units, where each unit is indexed by j. The firm decides how many production units to operate, where $J_{l,t}$ is given by

$$J_{l,t} = \overline{N} F_{l,t}. \tag{1}$$

We define $F_{l,t} \equiv K_{l,t-1}^{\alpha} H_{l,t}^{1-\alpha}$ and $\overline{N}$ is a normalization parameter that controls the level of disaggregation inside a firm. As analyzed below, in our model the uncertainty faced by a firm is invariant to the level of disaggregation.

Each unit j produces output, which is driven by three components: an economy-wide shock, a firm-specific shock and a unit-specific shock. This view of the firm is common in the industrial organization literature (see Coad (2007) for a survey) and has been motivated by the observed negative relationship between the size of a firm and its growth rate variance. The output at the unit level equals

$$x_{l,j,t} = e^{A_t + z_{l,t} + \tilde{\nu}_{l,j,t}} / \overline{N}, \tag{2}$$

where A_t is an economy-wide technology shock that follows $A_t = \rho_A A_{t-1} + \epsilon_{A,t}$, with $\epsilon_{A,t} \sim N(0, \sigma_A^2)$, $z_{l,t}$ is a firm-specific shock that follows

$$z_{l,t} = \rho_z z_{l,t-1} + \epsilon_{z,l,t}, \quad \epsilon_{z,l,t} \sim N(0, \sigma_z^2), \tag{3}$$

and the unit-specific shock follows $\tilde{\nu}_{l,j,t} \sim N(0, \overline{N}\sigma_\nu^2)$. The variance of a unit-specific shock is proportionally increasing in $\overline{N}$. Intuitively, as each production unit becomes smaller (i.e., as the level of disaggregation increases), the unit-specific component becomes larger compared to the firm-level component.[7]

Since the firm operates $J_{l,t}$ number of production units given by (1) and each unit produces according to (2), the firm's total output equals $Y_{l,t} = \sum_{j=1}^{J_{l,t}} x_{l,j,t}$.

Perfectively competitive final-goods firms produce aggregate output Y_t by combining goods produced by each firm l:

$$Y_t = \left[\int_0^1 Y_{l,t}^{\frac{\theta-1}{\theta}} dl \right]^{\frac{\theta}{\theta-1}}, \tag{4}$$

where θ determines the elasticity of substitution across goods. The demand function for intermediate goods l is $P_{l,t} = \left(\frac{Y_{l,t}}{Y_t}\right)^{-\frac{1}{\theta}}$, where we normalize the price of final goods $P_t = 1$. The revenue for firm l is then given by $P_{l,t}Y_{l,t} = Y_t^{\frac{1}{\theta}} Y_{l,t}^{1-\frac{1}{\theta}}$. Because the idiosyncratic shocks $z_{l,t}$ and $\tilde{\nu}_{l,j,t}$ can be equivalently interpreted as productivity or demand disturbances by adjusting the relative price $P_{l,t}$, we simply refer to $z_{l,t}$ and $\tilde{\nu}_{l,j,t}$ as profitability shocks. Note also that the firm-level returns to scale in terms of revenue, $1 - \frac{1}{\theta}$, is less than one, which gives us a notion of firm size that is well-defined.

Given production outcomes and its associated costs, firms pay out dividends

$$D_{l,t} = Y_t^{\frac{1}{\theta}} Y_{l,t}^{1-\frac{1}{\theta}} - W_t H_{l,t} - r_t^K K_{l,t-1}, \tag{5}$$

where W_t is the real wage and r_t^K is the rental rate for capital.

---

[6] Angeletos and Lian (2016) offer a distinct but complementary theory of propagation through endogenous confidence based on a lack of common knowledge, whose main effect is to attenuate general equilibrium effects.

[7] The assumption prevents output to be fully-revealing about firm-specific shocks even as we take the limit $\overline{N} \to \infty$. See Fajgelbaum et al. (2017) for a similar approach; in their model, the precision of a signal regarding an aggregate fundamental is decreasing in the number of total firms in the economy.

2.2. Imperfect information

We assume that agents cannot directly observe the realizations of idiosyncratic shocks $z_{l,t}$ and $\widetilde{\nu}_{l,j,t}$. Instead, every agent in the economy observes the economy-wide shocks A_t, the inputs used for operating production units $F_{l,t}$, as well as output $Y_{l,t}$ and $x_{l,j,t}$ of each firm l and production unit j. The imperfect observability assumption leads to a non-invertibility problem. Agents cannot tell whether an unexpectedly high realization of a production unit's output $x_{l,j,t}$ is due to the firm being 'better' (an increase in the persistent firm's specific profitability $z_{l,t}$) or just 'lucky' (an increase in the unit-specific shocks $\widetilde{\nu}_{l,j,t}$).

Faced with this uncertainty, agents use the available information, including the path of output and inputs, to form estimates on the underlying source of profitability $z_{l,t}$. 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.[8]

The measurement equation of the Kalman filter is given by the following sufficient statistic $s_{l,t}$ that summarizes observations from all production units within a firm l:

$$s_{l,t} = z_{l,t} + \nu_{l,t}, \tag{6}$$

where the average realization of the unit-specific shock is

$$\nu_{l,t} \equiv \frac{1}{J_{l,t}} \sum_{j=1}^{J_{l,t}} \widetilde{\nu}_{l,j,t} \sim N\left(0, \frac{\sigma_\nu^2}{F_{l,t}}\right),$$

and the transition equation for $z_{l,t}$ is given by (3).

The solution to the filtering problem is standard. The one-step-ahead prediction from the period $t-1$ estimate $\tilde{z}_{l,t-1|t-1}$ and its associated error variance $\Sigma_{l,t-1|t-1}$ are given by

$$\tilde{z}_{l,t|t-1} = \rho_z \tilde{z}_{l,t-1|t-1}; \; \Sigma_{l,t|t-1} = \rho_z^2 \Sigma_{l,t-1|t-1} + \sigma_z^2.$$

Then, firms update their estimates according to

$$\tilde{z}_{l,t|t} = \tilde{z}_{l,t|t-1} + \frac{\Sigma_{l,t|t-1}}{\Sigma_{l,t|t-1} + F_{l,t}^{-1}\sigma_\nu^2} \cdot (s_{l,t} - \tilde{z}_{l,t|t-1}), \tag{7}$$

and the updating rule for variance is

$$\Sigma_{l,t|t} = \left[\frac{\sigma_\nu^2}{F_{l,t}\Sigma_{l,t|t-1} + \sigma_\nu^2}\right]\Sigma_{l,t|t-1}. \tag{8}$$

The dynamics according to the Kalman filter can thus be described as

$$z_{l,t+1} = \rho_z(\tilde{z}_{l,t|t} + u_{l,t}) + \epsilon_{z,l,t+1}, \tag{9}$$

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

For our purposes, the important feature of the updating formulas is that the variance of the 'luck' component, which acts as a noise in the measurement Eq. (6), is decreasing in scale $F_{l,t}$. Thus, holding $\Sigma_{l,t|t-1}$ constant, the posterior estimation uncertainty $\Sigma_{l,t|t}$ in Eq. (8) increases as the scale decreases. Firm-level output becomes more informative about the underlying profitability $z_{l,t}$ as more production units operate.

2.3. Household wealth

There is a representative agent whose budget constraint is given by

$$C_t + B_t + I_t + \int P_{l,t}^e \theta_{l,t} dl \leq W_t H_t + r_t^K K_{t-1} + R_{t-1}B_{t-1} + \int (D_{l,t} + P_{l,t}^e)\theta_{l,t-1} dl + T_t,$$

where C_t is consumption of the final good, H_t is the amount of labor supplied, I_t is investment into physical capital, B_t is the one-period riskless bond, R_t is the interest rate, and T_t is a transfer. $D_{l,t}$ and $P_{l,t}^e$ are the dividend and price of a unit of share $\theta_{l,t}$ of firm l, respectively. Capital stock depreciates at rate δ so that it evolves as $K_t = (1-\delta)K_{t-1} + I_t$.

The market clearing conditions for labor, bonds and shares are:

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

The resource constraint is given by

$$C_t + I_t + G_t = Y_t, \tag{10}$$

[8] In Jovanovic (1982), firms use the observed outcome of production to learn about some unobserved technological parameter. In our model, firms learn about their time-varying, persistent profitability. The learning problem of the model with growth is in the online appendix, along with other equilibrium conditions.

where G_t is government spending and we assume a balanced budget each period ($G_t = -T_t$). For most of the analysis, we assume that government spending is a constant share of output, $\bar{g} = G_t/Y_t$.

Notice that our model is one with a typical infinitely-lived representative agent. Therefore, this agent is the relevant decision maker for the firms that operate the technology described in the previous section, since this agent owns in equilibrium the firms, i.e. $\theta_{l,t} = 1, \forall(t,l)$. The difference from a standard expected utility model, in which uncertainty is modeled only as risk, is that the decision maker faces ambiguity over the distribution of firm-level productivities, an issue that we take on next.[9]

2.4. Optimization

The representative household perceives ambiguity (Knightian uncertainty) about the vector of firm-level productivities $\{z_{l,t}\}_{l\in[0,1]}$. We now describe how that ambiguity process evolves. The agent uses observed data to learn about the hidden profitability through the Kalman filter to obtain a benchmark probability distribution. The Kalman filter problem has been described in Section 2.2 Ambiguity is modeled as a one-step-ahead *set* of conditional beliefs that consists of alternative probability distributions surrounding the benchmark Kalman filter estimate $\tilde{z}_{l,t}$ in (9) of the form

$$z_{l,t+1} = \rho_z \tilde{z}_{l,t|t} + \mu_{l,t} + \rho_z u_{l,t} + \epsilon_{z,l,t+1}, \quad \mu_{l,t} \in [-a_{l,t}, a_{l,t}] \tag{11}$$

In particular, the agent considers a set controlled by a bound on the relative entropy distance. More precisely, the agent only considers the conditional means $\mu_{l,t}$ that are sufficiently close to the long run average of zero in the sense of relative entropy:

$$\frac{\mu_{l,t}^2}{2\rho_z^2 \Sigma_{l,t|t}} \le \frac{1}{2}\eta^2, \tag{12}$$

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 ($\mu_{l,t}$ and zero), and η is a parameter that controls the size of the entropy constraint. The entropy constraint (12) results in a set $[-a_{l,t}, a_{l,t}]$ for $\mu_{l,t}$ in (11) that is given by

$$a_{l,t} = \eta \rho_z \sqrt{\Sigma_{l,t|t}}. \tag{13}$$

The interpretation of the entropy constraint is that the agent is less confident, i.e. the set of beliefs is larger, when there is more estimation uncertainty. 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, the agent becomes less confident and contemplates wider sets of conditional probabilities.[10] In our context, the entropy constraint allows us to account for the firm-level evidence in Senga (2018) that both forecast errors and dispersions decrease with firm size.

We model the household's aversion to ambiguity through recursive multiple priors preferences (Epstein and Schneider (2003b)), which capture an agent's lack of confidence in probability assessments. To define probability distributions over exogenous state variables, let us collect them in a vector $\vartheta_t \in \Theta$. Since $z_{l,t}$ and $\nu_{l,t}$ are not separately observed, the exogenous firm-level state is the observed $s_{l,t}$, i.e. the sum of $z_{l,t}$ and $\nu_{l,t}$, in Eq. (6), so that ϑ_t is comprised of A_t and the vector of idiosyncratic $\{s_{l,t}\}_{l\in[0,1]}$. The agent's lack of confidence about the persistent component $z_{l,t+1}$, showing up in Eqs. (11) and (13), therefore maps directly into a set one-step-ahead conditional beliefs about $s_{l,t+1}$. A household consumption plan C gives, for every history ϑ^t, the consumption of the final good $C_t(\vartheta^t)$ and the amount of hours worked $H_t(\vartheta^t)$. For a given consumption plan C, the household recursive multiple priors utility is defined by

$$U_t(C;\vartheta^t) = \ln C_t - \frac{H_t^{1+\phi}}{1+\phi} + \beta \min_{\mu_{l,t}\in[-a_{l,t},a_{l,t}],\forall l} E^{\mu}[U_{t+1}(C;\vartheta^t,\vartheta_{t+1})], \tag{14}$$

where β is the subjective discount factor and ϕ is the inverse of Frisch labor supply elasticity. We use the expectation operator $E^{\mu}[\,\cdot\,]$ to make explicit the dependence of expected continuation utility on the conditional means $\mu_{l,t}$.

Notice that there is a *cross-sectional distribution of sets* of beliefs over the future $\{s_{l,t+1}\}_{l\in[0,1]}$. Indeed, for each firm l, the agent entertains a set of conditional means $\mu_{l,t} \in [-a_{l,t}, a_{l,t}]$. If each set is singleton we obtain the standard expected utility case of separable log utility with those conditional beliefs. When the set is not a singleton, it reflects the assumption that the agent perceives Knightian uncertainty, in addition to the standard risk embedded in the conditional variances about $z_{l,t+1}$. As instructed by their preferences, in response to the aversion to that Knightian uncertainty, households take a cautious approach to decision making and act *as if* the true data generating process (DGP) is given by the worst-case conditional belief, which we will denote by $E_t^{\star}[\cdot]$.

[9] Note that while agents have ambiguity over the distribution of exogenous variables, they know the equilibrium relationships governing the evolution of the endogenous variables. See Eusepi and Preston (2011) as an example of learning about those laws of motions.

[10] Bianchi et al. (2018) also connect time-varying volatility, as measured by an econometrician, to perceived ambiguity to explain low frequency dynamics of asset prices in a business cycle model. In Ilut et al. (2016) ambiguity-averse firms learn about their demand curve and there a lower estimation uncertainty increases confidence in their estimate of demand, making them less willing to change their prices away from certainty.

Worst-case belief and the law of large numbers

Modeling idiosyncratic uncertainty also as ambiguity matters crucially for its effect on the decision maker's beliefs of continuation utility. Both risk and ambiguity share similar grounds: the sources of uncertainty are independent and identical and the rational decision maker — here the representative agent that owns the firms — does not evaluate the firms comprising the portfolio in isolation. In particular, in both cases, uncertainty over their idiosyncratic profitability matters only if it lowers the agent's continuation utility. That utility is a function of the wealth obtained through the average dividend from the portfolio.

The difference between risk and ambiguity is how it affects continuation utility. With risk only, uncertainty lowers that continuation utility by increasing the volatility of consumption. With purely idiosyncratic risk, uncertainty is diversified away since the law of large numbers (LLN) implies that the variance of consumption tends to zero as the number of firms becomes large. When uncertainty consists also of ambiguity, it affects utility by making the worst-case probability less favorable to the agent, through its effect on continuation utility in Eq. (14). Since ambiguity is over the conditional means of firm-level profitability, which in equilibrium affects dividends paid out to the agent, uncertainty affects utility by lowering the worst-case mean of firm-level profitability, i.e. $E_t^* z_{l,t+1}$. The agent faces independent and identical sources of uncertainty, represented here by the sets of distributions indexed by $\mu_{l,t}$, and therefore acts as if the mean on each source is lower. Therefore, in contrast to the risk case, the average dividend obtained on the portfolio, which is the equilibrium object that the agent cares about, does not become less uncertain, which here means being characterized by a narrower set of beliefs, as the number of firms increases.[11]

Put differently, the assumption of the sources of perceived uncertainty being independent and identical means that the agent is not willing to view a new firm added to the portfolio as 'hedging' out any ambiguity already perceived on that portfolio. Therefore, the agent ends up lacking confidence about the cross-sectional average (i.e. 'uncertainty over the size of the pie') as opposed to fully trusting that average but lacking confidence only about its composition (i.e. 'uncertainty over the shares of the pie'). It is this lack of confidence about the cross-sectional average that makes firm-level uncertainty not disappear through the LLN.

Therefore, once the representative agent correctly understands the effect of firm-level profitability on the continuation utility in Eq. (14), the worst-case belief can be easily solved for at the equilibrium consumption plan. Given the bound in Eq. (13), the worst-case conditional mean for each firm's $z_{l,t+1}$ is therefore given by

$$E_t^* z_{l,t+1} = \rho_z \tilde{z}_{l,t|t} - \eta \rho_z \sqrt{\Sigma_{l,t|t}} \tag{15}$$

where $\tilde{z}_{l,t|t}$ is the Kalman filter estimate of the mean obtained in Eq. (7). Thus, the worst-case conditional distribution of each firm's productivity is

$$z_{l,t+1} \sim N\left(E_t^* z_{l,t+1}, \rho_z^2 \Sigma_{l,t|t} + \sigma_z^2\right). \tag{16}$$

Once the worst-case distribution is determined, it is easy to compute the cross-sectional average realization $\int z_{l,t+1} dl$. By the law of large numbers (LLN) this average converges to

$$\int E_t^* z_{l,t+1} dl = -\eta \rho_z \int \sqrt{\Sigma_{l,t|t}} dl. \tag{17}$$

where we have used that $\int \tilde{z}_{l,t|t} dl = 0$.[12]

Eq. (17) is a manifestation in this model of the LLN for ambiguous random variables analyzed by Marinacci (1999) or Epstein and Schneider (2003a). In particular, now the *average* idiosyncratic uncertainty, $\int \sqrt{\Sigma_{l,t|t}} dl$, matters for the *average* worst-case expected $z_{l,t+1}$. That formula shows that once ambiguity is taken into account by the agent, the LLN implies that risk itself does not matter anymore for beliefs since the volatility of consumption converges to zero even under the worst-case conditional beliefs.

Firms' problem

Each firm chooses $H_{l,t}$ and $K_{l,t-1}$ to maximize

$$E_0^* \sum_{t=0}^{\infty} M_0^t D_{l,t}, \tag{18}$$

where E_0^* denotes expectation under the representative agent's worst-case probability and $D_{l,t}$ is given by Eq. (5). In equilibrium the representative household is the owner of the firm, and thus the present discounted value in (18) is evaluated under the shareholder equilibrium belief, here given by E_0^*. The random variables M_0^t denote state prices of t-period ahead contingent claims based on conditional worst-case probabilities, 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.

[11] 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 cross-sectional averages must (almost surely) lie in an interval bounded by the highest and lowest possible cross-sectional mean, and these bounds are tight in the sense that convergence to a narrower interval does not occur. See also Epstein and Schneider (2008) for an application of this argument to pricing a portfolio of firms with ambiguous dividends.

[12] Indeed, since under the true DGP the cross-sectional mean of $z_{l,t}$ is constant, the cross-sectional mean of the Kalman posterior mean estimate is a constant as well.

Compared to a standard model of full information and expected utility, the problem in (18) has two important characteristics. First, as described above, unlike the case of expected utility, the firm-level uncertainty that shows up in these state prices does not vanish under diversification. The second concerns experimentation. Under incomplete information but Bayesian decision making, experimentation is valuable as it raises expected utility by improving posterior precision. Here, ambiguity-averse agents also value experimentation as it affects utility by tightening the set of conditional probability considered. Therefore, in (18) firms take into account the impact of the level of input on the worst-case mean.

We summarize the timing of events within a period t as follows. In a first, pre-production, stage agents observe the realization of economy-wide shocks (here A_t). Given forecasts about $z_{l,t}$ (estimate of $\tilde{z}_{l,t-1|t-1}$ and uncertainty $\Sigma_{l,t-1|t-1}$) and its associated worst-case scenario $\eta\rho_z\sqrt{\Sigma_{l,t-1|t-1}}$, firms hire labor $H_{l,t}$ and rent capital $K_{l,t-1}$. The household supplies labor H_t and capital K_{t-1}, so that the labor and capital rental markets clear at the wage rate W_t and capital rental rate r_t^K. In a second, post-production, stage idiosyncratic shocks $z_{l,t}$ and $\nu_{l,t}$ realize (but are unobserved) and production occurs. Then, given the already chosen productive inputs, firms update uncertainty $\Sigma_{l,t|t}$. They also use the observed signal $s_{l,t}$ to update the estimate $\tilde{z}_{l,t|t}$ and use it together with the associated worst-case scenario $\eta\rho_z\sqrt{\Sigma_{l,t|t}}$ to form forecasts for production for next period. Given the realized production, firms pay out to the representative agent dividends $D_{l,t}$. The household then makes an intertemporal savings decision, by choosing current consumption, investment, and asset purchase decisions (C_t, I_t, B_t, and $\theta_{l,t}$) under the worst-case conditional probability distribution for $s_{l,t+1}$.

2.5. 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. 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. Once decision rules are log-linearized, aggregation occurs exactly and the cross-sectional mean becomes the only relevant statistic for tracking aggregate dynamics. More precisely, controlling for its mean, the cross-sectional distribution of firms is not a part of the state space anymore.

We log-linearize equilibrium conditions around the steady state based on the worst-case beliefs.[13] Given the equilibrium laws of motion we then characterize the dynamics of the economy under the true DGP. Our solution method extends the one in Ilut and Schneider (2014) by endogenizing the process of ambiguity perceived by the representative household. More substantially, the methodology allows for a tractable aggregation of the endogenous uncertainty faced by heterogeneous firms.

We present additional details in the online appendix, including the recursive representation of the model, the optimality conditions (which will be a subset of those characterizing the estimated model with additional rigidities introduced in Section 4.1), a general description of the solution method, and an example that illustrates the log-linearization logic and the feedback between the average level of activity and the cross-sectional average of the worst-case mean.

3. Propagation mechanism: Countercyclical correlated wedges

In this section we characterize the main properties of the information driven propagation mechanism by exploring how the model generates *as if correlated* wedges that respond to the aggregate state of the economy.

3.1. Co-movement and the labor wedge

Of particular importance for aggregate dynamics is the implied correlation between the fundamental shock and a labor wedge. This endogenous correlation provides the potential for a wide class of fundamental shocks to produce the basic business cycle pattern of co-movement between hours, consumption and investment, without additional rigidities.

The optimal labor tradeoff of equating the marginal cost to the expected marginal benefit under the worst-case belief E_t^* is given by:

$$H_t^{\phi} = E_t^*(\lambda_t MPL_t) \tag{19}$$

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 preference shock to hours worked in agent's utility (14).

Instead, in our model, there can be such co-movement. Suppose that there is a period of low confidence. From the negative wealth effect current consumption is low and marginal utility λ_t is high, so the standard effect would be to see high

[13] 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, i.e. low (high) output and high (low) uncertainty, similar to the analysis in Fajgelbaum et al. (2017). We circumvent this multiplicity by treating the posterior variance of the level of idiosyncratic TFP as a parameter and by focusing on the unique steady state implied that posterior variance.

labor supply as a result. However, because the firm chooses hours *as if* productivity is low, there is a counter substitution incentive for hours to be low. To see how the model generates a countercyclical labor wedge, note that a decrease in hours worked due to an increase in ambiguity 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^H as

$$H_t^{\phi} = (1 - \tau_t^H)\lambda_t MPL_t \tag{20}$$

Using the optimality condition in (19), the labor tax is

$$\tau_t^H = 1 - \frac{E_t^*(\lambda_t MPL_t)}{\lambda_t MPL_t} \tag{21}$$

Consider first the linear rational expectations case. There the role of firm-level uncertainty disappears and the labor tax in equation (21) is constant and equal to zero. To see this, note our timing assumption that labor is chosen after the economy-wide shocks are realized and observed at the beginning of the period. This makes the optimality condition in (19) take the usual form of an intratemporal labor decision.[14]

Consider now the econometrician that measures realized H_t, C_t and MPL_t in our model. The ratio in Eq. (21) between the expected benefit to working under the worst-case belief compared to the econometrician's measure, which uses the average $\mu = 0$, is not equal to one. This ratio is affected by standard wealth and substitution effects. Take for example a period of low confidence. On the one hand, since the agent is now more worried about low consumption, the agent's expected marginal utility λ_t is larger than measured by the econometrician's. On the other hand, now the expected marginal product of labor MPL_t is lower than measured by the econometrician. When the latter substitution effect dominates, the econometrician rationalizes the 'surprisingly low' labor supply by a high labor tax τ_t^H.[15]

In turn, periods of low confidence are 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 by Eq. (17). Therefore, when the substitution effect on the labor choice dominates, the econometrician finds 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 et al. (2007)).

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 may add an additional reason why labor moves 'excessively', from the perspective of an observer that only uses equation (19) to understand labor movements.

3.2. The intertemporal consumption wedge

Uncertainty also affects the consumption-savings decision of the household. This is reflected in the Euler condition for the risk-free asset:

$$1 = \beta R_t E_t^*(\lambda_{t+1}/\lambda_t) \tag{22}$$

As with the labor wedge, let us implicitly define an intertemporal consumption wedge:

$$1 = (1 + \tau_t^B)\beta R_t E_t(\lambda_{t+1}/\lambda_t) \tag{23}$$

Importantly, this wedge is time varying, since the bond is priced under the uncertainty adjusted distribution, E_t^*, which differs from the econometrician's DGP, given by E_t. By substituting the optimality condition for the interest rate from (22), the wedge becomes:

$$1 + \tau_t^B = \frac{E_t^* \lambda_{t+1}}{E_t \lambda_{t+1}} \tag{24}$$

Eq. (24) makes transparent the predictable nature of the wedge. In particular, during low confidence times, the representative household acts *as if* future marginal utility is high. This heightened concern about future resources drives up demand for safe assets and leads to a low interest rate R_t. However, from the perspective of the econometrician, the measured average marginal utility at $t+1$ is not particularly high. To rationalize the low interest rate without observing large changes in the growth rates of marginal utility, the econometrician recovers a high consumption wedge τ_t^B, or a high 'tax' on consumption. Therefore, the model offers a mechanism to generate movements in the relevant stochastic discount factor that arise endogenously as a countercyclical desire to save in risk-free assets.

[14] If we would assume that labor is chosen before the aggregate shocks are realized, there would be a fluctuating labor tax in (21) even in the rational expectations model. In that model, the wedge is $\tau_t^H = 1 - \frac{E_{t-1}(\lambda_t MPL_t)}{\lambda_t MPL_t}$, 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^H will not be predictable using information at time $t-1$, including the labor choice, such that $E_{t-1}\tau_t^H = 0$. In contrast, our model with learning produces predictable, countercyclical, labor wedges.

[15] Given the equilibrium confidence process, which determines the worst-case belief E_t^*, the economic reasoning behind the effects of distorted beliefs on labor choice has been well developed by existing work, such as Angeletos and La'O (2009, 2013). There they describe the key income and substitution forces through which correlated higher-order beliefs, a form of confidence shocks, show up as labor wedges in a model where hiring occurs under imperfect information on its return.

3.3. The excess return wedge

Conditional beliefs matter also for the Euler condition for capital:

$$\lambda_t = \beta E_t^*[\lambda_{t+1} R_{t+1}^K]. \tag{25}$$

Under our linearized solution, using Eq. (22), we get $E_t^* R_{t+1}^K = R_t$, where $E_t^* R_{t+1}^K$ is the expected return on capital under the worst-case belief. As with the intertemporal consumption wedge, let us define the measured excess return wedge as

$$E_t R_{t+1}^K = R_t(1 + \tau_t^K) \tag{26}$$

As with bond pricing, this wedge is time-varying and takes the form

$$1 + \tau_t^K = \frac{E_t R_{t+1}^K}{E_t^* R_{t+1}^K} \tag{27}$$

During low confidence times demand for capital is 'surprisingly low'. This is rationalized by the econometrician, measuring R_{t+1}^K under the true DGP, as a high ex-post excess return $R_{t+1}^K - R_t$, or as a high wedge τ_t^K in Eq. (27). In the linearized solution, the excess return, similarly to the labor tax and the discount factor wedge, is inversely proportional to the time-varying confidence. In times of low economic activity, when confidence is low, the measured excess return is high.

Putting together the consumption wedge and the excess return we can characterize the linearized version of the Euler equation for capital in (25) as

$$\lambda_t = \frac{(1 + \tau_t^B)}{(1 + \tau_t^K)} \beta E_t[\lambda_{t+1} R_{t+1}^K]. \tag{28}$$

Eq. (28) and the emergence of *both* τ_t^B and τ_t^K provide cross-equation restrictions that connects our model to three interpretations of shocks to the Euler equations present in the literature. First, it clarifies that the τ_t^B wedge does not simply take the form of an 'as if' shock to β. If that would be the case, then τ_t^K would be zero since the desire to save through a higher β would show up equally in the Euler equations for bonds in (22) and capital in (25).[16] Second, it clarifies that the friction generates more than just an 'as if' tax in the capital market. If that would be the case, then τ_t^B would be zero since the desire of the representative agent to save would not be affected.[17] Third, the simultaneous presence of the two wedges relates the friction to a large DSGE literature that uses reduced-form 'risk-premium' shocks. Such shocks are introduced as a stochastic preference for risk-free over risky assets, by distorting the Euler equation for bonds but not for capital, which can be interpreted in our model as $\tau_t^B = \tau_t^K$.[18]

Therefore, the model predicts that in a recession an outside econometrician should observe 'excessively low' hours worked, at the same time when prices of riskless assets and excess returns for risky assets are 'excessively inflated'. These correlations arise from any type of shock that moves the economic activity.

3.4. Learning from aggregate market outcomes

We conclude the description of the model's qualitative properties by discussing the generality of the proposed economic forces. In particular, there are two basic features of uncertainty that were crucial in our proposed mechanism for business cycle dynamics. First, the cross-sectional average estimation uncertainty is lower in times when the cross-sectional average production is larger. Second, this state-dependent estimation uncertainty affects consumption and production decisions, including the labor choice.

In the context of these general forces, an alternative approach to generate the negative feedback loop between estimation uncertainty and aggregate economic activity is to modify two of our basic features by the following assumptions. First, firms learn about the aggregate-level productivity A_t. Second, lower aggregate output corresponds to fewer signals available to the firms. This approach of learning from market outcomes is present, in different forms, in the existing macroeconomic literature on endogenous uncertainty, such as Caplin and Leahy (1993); Fajgelbaum et al. (2017); van Nieuwerburgh and Veldkamp (2006); Ordoñez (2013), and Saijo (2017).

In a setup with ambiguity like ours, where uncertainty changes the decision maker's plausible set of conditional means, this alternative approach of learning from market outcomes generates a propagation mechanism for the aggregate dynamics that is qualitatively similar to our benchmark model. The reason is that in both approaches the cross-sectional average estimation uncertainty is countercyclical and that uncertainty affects beliefs about aggregate conditions. Indeed, as discussed

[16] See Christiano et al. (2005) and Smets and Wouters (2007) as examples of a large literature of DSGE models that use shocks to β. Recent work, such as Eggertsson and Woodford (2003) and Christiano et al. (2015), also models the heightened desire to save as an independent stochastic shock that is responsible for the economy hitting the zero lower bound on the nominal interest rate.

[17] Quantitative DSGE models typically employ these as if taxes when modeling financial frictions. See for example Christiano et al. (2014); Gilchrist and Zakrajšek (2011) and Del Negro and Schorfheide (2013).

[18] Reduced-form risk premium shocks have typically emerged as a key business cycle driver in quantitative DSGE models, starting with Smets and Wouters (2007). See Gust et al. (2017) for a recent contribution emphasizing the quantitative role of these shocks. See Fisher (2015) for an interpretation of these shocks as time-varying preference for liquidity.

in Section 2.4, even when ambiguity is solely about the mean of each firm's productivity, the law of large numbers still preserves an effect of firm-level uncertainty on the worst-case beliefs of the cross-sectional average productivity.

While qualitatively similar to learning from aggregate market outcomes in its implications for aggregate dynamics, the friction present in our benchmark model, namely learning about firm-level profitability, also has some properties that are different. First, the competitive equilibrium of our economy is constrained Pareto optimal. Indeed, in this world there are no information externalities since learning occurs at the individual firm level and not from observing the aggregate economy. This stands in contrast to the case of learning from aggregate market outcomes, where an individual firm does not take into account the positive externality of generating signals that are useful for the rest of the economy.[19] Thus, even if policy interventions affect the aggregate dynamics similarly in the two cases, the welfare properties are different. For example, the increased economic activity, and the associated increase in the signal-to-noise ratio, produced by a government spending increase is not welfare increasing in our model. Second, extending the sources of imperfect information to firm-level shock offers a new way of disciplining endogenous uncertainty process through micro data. These include, as we will discuss in our quantitative model, firm-level technological or informational parameters.

4. Quantitative analysis

We now bring our endogenous uncertainty mechanism to the data in order to quantify the potential of the proposed information friction as a propagation mechanism and contrast it to other frictions. Our analysis consists of four steps. First, we embed the friction into a standard medium-scale business cycle model by allowing for an array of real and nominal rigidities.[20] Second, we employ an estimation procedure that focuses squarely on propagation. Since our friction predicts that we should observe regular patterns of co-movement and correlated wedges *conditional* on different types of shocks, our estimation consists of matching the model-implied and empirical impulse responses recovered using the local projection method (Jordá (2005)). Third, we run monetary and fiscal policy experiments to evaluate the impact of our friction on policies. Fourth, we use observable dispersion of beliefs to externally test the model's implications.

4.1. A quantitative business cycle model with rigidities

We add several standard features to the estimated model. The production function with capital utilization is $F_{l,t} = (U_{l,t}K_{l,t-1})^{\alpha}(\gamma^t H_{l,t})^{1-\alpha}$, where γ is the deterministic growth rate and $a(U_{l,t})K_{l,t-1}$ is an utilization cost that reduces dividends in Eq. (5).[21]

We modify the representative household's utility (14) to allow for habit persistence in consumption. In particular, the per-period utility from consumption is now $\ln(C_t - bC_{t-1})$, where $b > 0$ is a habit parameter. We also introduce an investment adjustment cost and an investment-specific technology (IST) shock:

$$K_t = (1-\delta)K_{t-1} + \left\{1 - \frac{\kappa}{2}\left(\frac{I_t}{I_{t-1}} - \gamma\right)^2\right\}\zeta_t I_t, \tag{29}$$

where $\kappa > 0$ is a parameter and ζ_t is the level of investment-specific technology that follows $\ln\zeta_t = \rho_\zeta \ln\zeta_{t-1} + \epsilon_{\zeta,t}$, where the innovation $\epsilon_{\zeta,t}$ is iid Gaussian with a standard deviation σ_ζ. For nominal rigidities we consider standard Calvo-type price and wage stickiness, along with monopolistic competition.[22]

We follow Bernanke et al. (1999) and introduce ex-ante homogeneous entrepreneurs that purchase capital from households and use it to produce output. The purchases of capital are financed by two sources: their own net worth and borrowing from financial intermediaries. The financial intermediaries provide external finance to entrepreneurs using funds obtained from households. The agency problem between entrepreneurs and financial intermediaries gives rise to an external finance premium. We introduce a financial shock, Δ_t^K, in the form of a time-varying difference between the financial intermediary's revenue and its opportunity cost of its funds (the risk-free return). We assume Δ_t^K follows an AR(1) process $\ln\Delta_t^K = \rho_\Delta \ln\Delta_{t-1}^K + \epsilon_{\Delta,t}$, where the innovation $\epsilon_{\Delta,t}$ is iid Gaussian with a standard deviation σ_Δ. An increase in Δ_t^K raises the credit spread (the difference between the loan rate to entrepreneurs and the risk-free rate) and drives up the cost of external finance. The interpretation and identification of this financial shock follows the standard literature, along the lines

[19] See Caplin and Leahy (1993), Ordoñez (2013), and Fajgelbaum et al. (2017) for a discussion of the information externalities arising in models based on learning from aggregate market outcomes.

[20] We follow the standard approach and include nominal rigidities as the main friction to generate co-movement. Other directions that address the Barro and King (1984) critique include: strategic complementary in a model with dispersed information (Angeletos and La'O (2013)), heterogeneity in labor supply and consumption across employed and non-employed (Eusepi and Preston (2015)), variable capacity utilization and a large preference complementarity between consumption and hours (Jaimovich and Rebelo (2009)).

[21] 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 is one. The cost $a(U)$ is increasing in U and χ_1 determines the degree of the convexity of utilization costs. In a linearized equilibrium, the dynamics are controlled by χ_1.

[22] 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 $1 - \xi_w$.

of Gilchrist and Zakrajšek (2012). In the online appendix we provide a complete exposition of the financial friction and the financial shock.

The central bank follows a Taylor-type rule. We consider a general form and allow the monetary authority to respond to current and lagged endogenous variables:

$$\hat{R}_t = \rho_R \hat{R}_{t-1} + \sum_{i=0}^{2} \phi_\pi^i \hat{\pi}_{t-i} + \sum_{i=0}^{2} \phi_Y^i \Delta \hat{Y}_{t-i} + \epsilon_{R,t}, \qquad \epsilon_{R,t} \sim N(0, \sigma_R^2),$$

where ρ_R, ϕ_π^i, and ϕ_Y^i are parameters and $\epsilon_{R,t}$ is a monetary policy shock.

4.2. Recovering empirical impulse responses using local projections

The starting point of our investigation is a local projection estimation of empirical impulse responses using U.S. quarterly macroeconomic data over the sample period 1973Q1–2008Q3. We trim the observation after 2008Q4 to avoid complications arising from the zero lower bound.[23] Specifically, we estimate the following regressions

$$x_{t+h} = c^h + \tau_1^h t + \tau_2^h t^2 + \sum_{j=1}^{J} \alpha_j^h x_{t-j} + \sum_{i=0}^{I} \sum_{k=1}^{K} \beta_i^{k,h} e_{t-i}^k + \varepsilon_{t+h}, \quad h = 0, \ldots, H \tag{30}$$

where x_t is the variable of interest and e_t^k are the identified structural shocks.[24] Our coefficient of interest is $\{\beta_0^{k,h}\}_{h=0}^H$.[25] We set $J = 8$ and $I = 2$ and compute the impulse response for $H = 15$ horizons for four structural shocks — neutral technology, investment-specific technology, financial and monetary policy shocks — and thus $K = 4$. We directly use the log-growth rate of utilization-adjusted TFP from Fernald (2014) to identify neutral technology shocks. Similarly, we use the log-growth rate of relative price of investment goods[26] for the investment-specific technology shocks. For the financial shock, we use the excess bond premium from Gilchrist and Zakrajšek (2012). Finally, to measure the monetary policy shock, we use the Romer and Romer (2004) series, extended by Coibion et al. (2017).[27]

Fig. 1 reports the fraction of variance of each endogenous variable explained by the identified shocks for different horizons.[28] First, the identified shocks account for a significant fraction of fluctuations. For example, the shocks account for 40 to 60% of output and investment variations. While the identified shocks account for less variation of hours, they explain up to 80% of the observed variance in consumption. For the nominal variables such as inflation and the federal funds rate, the shocks account for over half of the observed fluctuations depending on the horizons. Second, the financial shock is the most important among the identified shocks. This shock accounts for a significant fraction of the contribution to fluctuations in hours and investment and more than a quarter of the total variations in output and consumption explained by the identified shocks.

4.3. Bayesian impulse response matching estimation

We fix a small number of parameters before the estimation. The growth rate of technology γ, the discount factor β, the depreciation rate of capital δ, and the share of government spending to output $\bar{g}$ are set to 1.004, 0.998, 0.025, and 0.2 respectively. We set θ, θ_p, θ_w to 11, which imply steady-state firm-level markups, price markups, and wage markups of 10%. The survival rate of the entrepreneurs is set to $\zeta = 0.98$ and the steady-state capital to net worth ratio is set to 1.7, which are in line with the values used in Bernanke et al. (1999). The remaining set of parameters is estimated using a Bayesian version of a impulse-response-matching method, developed by Christiano et al. (2010). The description of the methodology is contained in the online appendix.

We conduct two main estimation experiments. In the first, we estimate our model using only the impulse responses to the financial shock. This shock is particularly informative for our objective for two reasons. First, it is quantitatively important, as it accounts for a significant fraction of business cycle variation. Second, the shock is characterized by cross-equation restrictions, in the form of positive co-movement of aggregate variables as well as correlated wedges, that provide stark identification of the underlying propagation mechanisms. In the second experiment, we estimate the model using impulse responses to all four identified shocks. This allows us to examine the quantitative robustness of the conclusion from the first experiment, explore the implications of endogenous uncertainty for other structural shocks, and more generally evaluate the role played by the additional cross-equations restrictions in the identification of the model.

[23] The sample starting point coincides with the initial availability of the credit spread series from Gilchrist and Zakrajšek (2012).

[24] In terms of x_t, output, investment, consumption, and real wages enter in log-growth rates. All other variables enter in log-levels.

[25] We cumulate $\beta_0^{k,h}$s when presenting impulse responses for variables that enter in growth rates.

[26] The series is retrieved from FRED (https://fred.stlouisfed.org/series/PIRIC).

[27] Following Romer and Romer (2004) and Coibion et al. (2017), we set the contemporaneous impact for all variables (except for the federal fund rate) to the monetary policy shock to zero. We modify the timing of the quantitative model so that it is consistent with the timing assumption of the empirical impulse response. In the appendix, we list the equilibrium conditions following this timing convention.

[28] We compute the variance decomposition using the local projection based method described in Gorodnichenko and Lee (2017). Output, investment, and consumption are in log-growth rates and other variables are in log-levels.

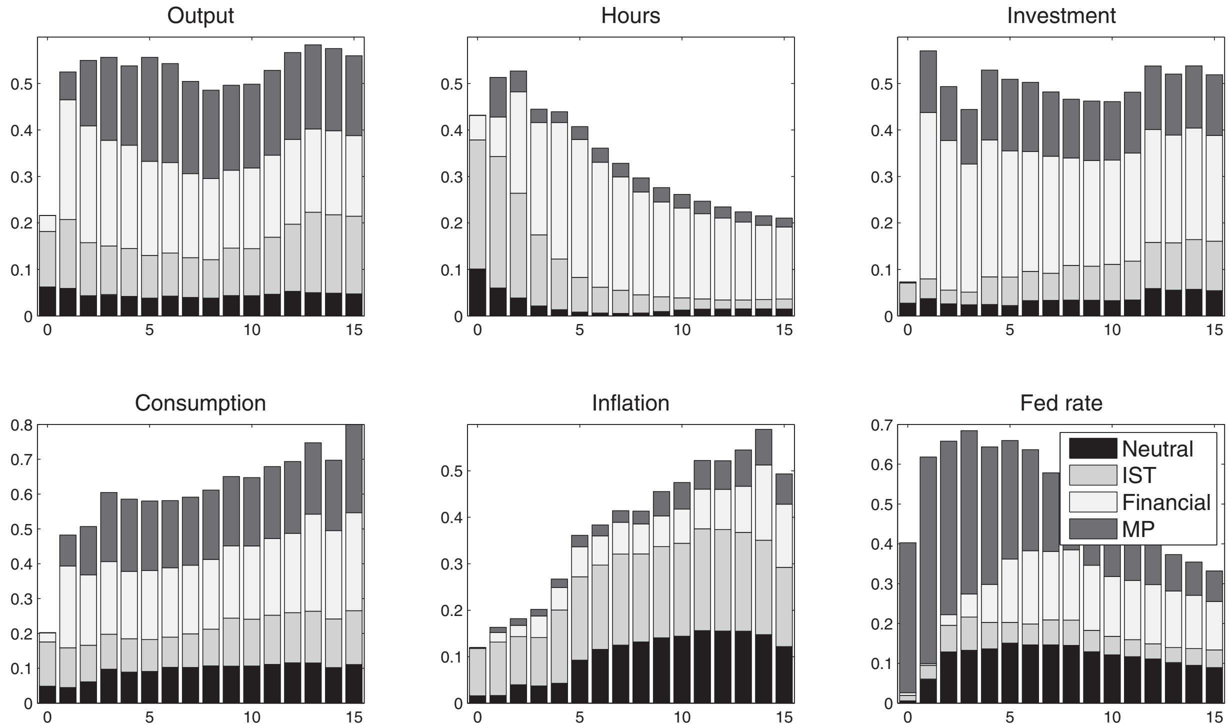

Fig. 1. Variance decomposition. *Notes*: We report the fraction of variance that can be explained due to the indicated shocks in the local projection for different horizons.

For both experiments, we stack the current and 15 lagged values of impulse response functions for the following variables — real GDP, hours worked, real investment, real consumption, real wages, inflation, fed rate as well as the Gilchrist and Zakrajšek (2012) credit spread, TFP growth rate, growth rate of relative price of investment — in the vector of responses to be matched. As additional discipline coming from the empirical cross-equation restrictions, we also incorporate the responses of labor and consumption wedges and excess return implicitly computed from the local projection, using the log-linearized first-order conditions from (20), (23), and (26). To calculate these wedges from the data, we need to take stand on some parameter values. We assume $\phi = 0.25$ and $b = 0$. When we calculate the wedges implied by the models, we use the same log-linearized conditions and parameter values and the expectations are computed under the econometrician's DGP.[29] Thus, in computing the wedges, the data and the model are treated symmetrically.

We use the data on the dispersion of forecasts available from the Survey of Professional Forecasters (SPF) as an external model validation. Similar to Ilut and Schneider (2014), we relate the set of forecasts about real GDP growth and inflation in the model to the observed dispersion of forecasts in the data. The logic is as follows: the representative agent samples experts' opinions and aggregates them, according to his preferences, when making decisions. Since the decision maker is ambiguity averse, stronger disagreement among these professional forecasters generates lower confidence in the probability assessments of the future. Using the bounds in Eq. (11), the dispersions of forecasts ranging from the maximum to the minimum forecast about one quarter ahead real GDP growth and inflation are given by $D_t^y = \varepsilon_{yz} 2|\int_0^1 a_{l,t} dl|$ and $D_t^\pi = \varepsilon_{\pi z} 2|\int_0^1 a_{l,t} dl|$, where ε_{yz} and $\varepsilon_{\pi z}$ is the equilibrium elasticity of output growth and inflation, respectively, with respect to the innovation to average firm-level profitability in the linear decision rule.[30]

While the ambiguity model produces such dispersions, the RE does not since in that case the set of forecasts collapses by assumption to a singleton. Excluding SPF dispersion from the estimation criterion allows us to keep the number of observables between our model and its RE counterpart the same and thus facilitates the comparison between the two models. Nevertheless, when we report the estimated impulse response from our model with ambiguity, we plot the implied range of growth forecasts against that from the SPF, as an outside of model check on the plausibility of the mechanism.

[29] To be precise, we use the following equations to calculate the wedges:

$$\hat{\tau}_t^H = -(1+\phi)\hat{H}_t - \hat{C}_t + \hat{Y}_t;\ \hat{\tau}_t^B = -\hat{C}_t + E_t\hat{C}_{t+1} - \hat{R}_t + E_t\hat{\pi}_{t+1};\ \hat{\tau}_t^K = E_t\hat{R}_{t+1}^K - \hat{R}_t,$$

where $\hat{R}_{t+1}^K$ is measured using return on assets.

[30] For an alternative approach to discipline the equilibrium set of beliefs see Bhandari et al. (2016), who use the Michigan survey of consumers to interpret the reported unique subjective belief as the equilibrium worst-case forecast made by the typical household. Instead, we use the experts' forecasts as discipline on the belief sets that the decision maker entertains as plausible.

Table 1
Estimated parameters: preference and technology.

		Prior			Single shock		All shocks	
		Type	Mean	Std	Ambiguity	RE	Ambiguity	RE
α	Capital share	B	0.3	0.02	0.38 (0.02)	0.38 (0.02)	0.46 (0.02)	0.48 (0.02)
ϕ	Inv. Frisch elasticity	G	1	0.15	0.17 (0.02)	1.42 (0.14)	0.10 (0.01)	0.06 (0.01)
χ_1	Utilization cost	IG	0.01	0.25	0.004 (0.003)	0.005 (0.17)	0.004 (0.003)	0.004 (0.002)
b	Consumption habit	B	0.3	0.1	0	0.99 (0.002)	**0.62** (0.02)	**0.73** (0.02)
κ	Investment adj. cost	G	1.5	0.2	0	0.99 (0.13)	**0.21** (0.02)	**1.90** (0.12)
$\frac{1}{1-\xi_p}$	Avg. freq. of price adjustment	G	3	0.5	1.0001	2.50 (0.34)	**1.20** (0.05)	**10.54** (0.62)
$\frac{1}{1-\xi_w}$	Avg. freq. of wage adjustment	G	3	0.5	1.0001	2.15 (0.16)	**1.00** (0.01)	**2.07** (0.11)
σ_ω	Std. of entrepreneur shock	G	0.5	0.15	0.19 (0.01)	0.58 (0.03)	0.36 (0.01)	0.37 (0.03)
μ	Monitoring cost	B	0.1	0.05	0.07 (0.05)	0.002 (0.002)	0.13 (0.01)	0.02 (0.02)
Δ^K	SS financial shock	B	0.015	0.01	0.013 (0.001)	0.002 (0.001)	0.002 (0.001)	0.012 (0.001)
ρ_z	Idiosyncratic shock	B	0.6	0.1	0.70 (0.02)	–	0.58 (0.02)	–
σ_z	Idiosyncratic shock	B	0.4	0.2	0.98 (0.02)	–	0.98 (0.01)	–
0.5η	Entropy constraint	B	0.5	0.2	0.98 (0.01)	0	0.99 (0.01)	0
$\bar{\Sigma}$	SS posterior variance	G	0.1	0.02	0.10 (0.02)	–	0.01 (0.004)	–

Notes: See notes from Table 2.

Tables 1 and 2 report the prior distributions. Since we use standard choices for priors whenever possible, our discussion focuses on the parameters that affect the strength of the feedback loop between economic activity and uncertainty, which are determined by three factors. The first factor is the variability of inputs which is determined by the elasticities of capital utilization and labor supply. χ_1, which controls the elasticity of utilization, is centered around 0.01, where lower values indicate more elastic utilization,[31] while the inverse Frisch elasticity ϕ is centered around 1, which is in line with the recommended value by Chetty et al. (2011).[32] Second, the parameters that are related to the firm-level processes control how changes in inputs translate to changes in posterior variance. For the persistence of the idiosyncratic shocks, we set a prior for ρ_z centered around 0.6 following the estimates in the literature. Guided by the establishments-level evidence by Bloom et al. (2018), we set the prior mean of the innovation σ_z to be 0.4. David et al. (2015) estimate the posterior variance of a firm-specific shock to be around 8–13%. We set the prior mean for the posterior variance at the zero-risk steady state $\bar{\Sigma}$ to 10%.[33] Finally, the size of the entropy constraint η determines how changes in the posterior standard deviation translate into changes in confidence. Ilut and Schneider (2014) argue that a reasonable upper bound for η 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. We re-parameterize the parameter and estimate 0.5η, for which we set a Beta prior.[34]

4.4. Impulse response for the financial shock only: Learning friction alone *fits data better*

Our first analysis is to estimate the model using *only* the impulse response to the financial shock. We highlight the properties of our endogenous uncertainty mechanism by keeping *only* the learning friction, so we otherwise shut down the standard real and nominal rigidities such as consumption habit, investment adjustment cost, sticky prices and wages. To understand the advantage of our mechanism compared to other propagation channels proposed in the literature, we

---

[31] The choice of the prior mean is motivated by Christiano et al. (2005), who use $\chi_1 = 0.01$.

[32] Christiano et al. (2014) also use $\phi = 1$ in their calibration.

[33] 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 unit-specific shock σ_ν. 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. We provide additional details in the online appendix.

[34] The priors for the standard deviation of a shock to the entrepreneurs σ_ω and the monitoring cost μ are centered around 0.5 and 0.1, respectively, in line with the values used in Bernanke et al. (1999). We set the prior mean of the steady-state financial shock Δ^K to 0.015, motivated by the finding by Phillipon (2015) that financial intermediation costs around 1.5 percent of intermediated assets.

Table 2
Estimated parameters: Monetary policy and structural shocks.

		Prior			Single shock		All shocks	
		Type	Mean	Std	Ambiguity	RE	Ambiguity	RE
ρ_R	Interest smoothing	B	0.8	0.1	0.87 (0.08)	0.16 (0.04)	0.76 (0.08)	0.99 (0.01)
ϕ_π^0	Inflation response	N	1.3	0.2	1.30 (0.19)	0.60 (0.10)	0.62 (0.08)	0.99 (0.16)
ϕ_π^1	Inflation response	N	1.3	0.2	1.27 (0.20)	0.62 (0.10)	1.22 (0.17)	0.89 (0.15)
ϕ_π^2	Inflation response	N	1.3	0.2	1.35 (0.20)	0.64 (0.11)	0.80 (0.11)	0.79 (0.13)
ϕ_Y^0	Output response	G	0.1	0.05	0.01 (0.01)	0.46 (0.08)	0.06 (0.02)	0.01 (0.003)
ϕ_Y^1	Output response	G	0.1	0.05	0.04 (0.02)	0.10 (0.06)	0.06 (0.02)	0.00 (0.003)
ϕ_Y^2	Output response	G	0.1	0.05	0.08 (0.05)	0.06 (0.03)	0.05 (0.02)	0.00 (0.002)
ρ_Δ	Financial	B	0.6	0.2	0.97 (0.006)	0.95 (0.009)	0.99 (0.004)	0.88 (0.008)
$100\sigma_\Delta$	Financial	IG	1	1	4.38 (0.25)	55.3 (22.8)	13.3 (10.8)	5.43 (0.63)
ρ_A	Neutral tech.	B	0.6	0.2	–	–	0.98 (0.002)	0.99 (0.0004)
$100\sigma_A$	Neutral tech.	IG	1	1	–	–	0.20 (0.01)	0.23 (0.01)
ρ_ζ	Investment-specific tech.	B	0.6	0.2	–	–	0.87 (0.02)	0.44 (0.11)
$100\sigma_\zeta$	Investment-specific tech.	IG	1	1	–	–	0.20 (0.01)	0.64 (0.16)
$100\sigma_R$	Monetary policy	IG	0.1	1	–	–	0.15 (0.01)	0.09 (0.01)
Log marginal likelihood					**-486**	**-514**	**-2705**	**-3092**

Notes: 'Single shock' refers to the posterior modes of the estimation using only the financial shock and 'All shocks' refers to the posterior modes from the estimation using all four shocks. 'Ambiguity' corresponds to the baseline model with endogenous uncertainty and 'RE' corresponds to its rational expectations version. *B* refers to the Beta distribution, *N* to the Normal distribution, *G* to the Gamma distribution, *IG* to the Inverse-gamma distribution. Posterior standard deviations are in parentheses and are obtained from draws using the random-walk Metropolis-Hasting algorithm. The marginal likelihood is calculated using Geweke's modified harmonic mean estimator.

compare this parsimonious estimated model with two alternatives. The first is a standard RE model augmented with all the features, except ambiguity, presented in Section 4.1 This model represents the conventional view for a quantitatively relevant business cycle model based on the New Keynesian paradigm.

The second alternative model that we consider is an otherwise frictionless RE model augmented with learning-by-doing (LBD) as in Chang et al. (2002). This version serves as a useful comparison for how learning affects dynamics differently in our proposed propagation mechanism based on uncertainty. In particular, in this LBD model, the level of inputs adjusted for the workers' skill is given by $F_{l,t} \equiv K_{l,t-1}^{\alpha}(X_t H_{l,t})^{1-\alpha}$, where, as in d'Alessandro et al. (2019), the skill level X_t depends on past per-capita labor supply and is given to the individual firm so that $X_t = X_{t-1}^{\rho_x} H_{t-1}^{\mu_n}$, where $0 \le \rho_x < 1$, $\mu_n \ge 0$.[35]

Fig. 2 reports the local projection impulse responses (labeled 'Local projection') as well as the estimated impulse responses from our model (labeled 'Ambiguity, no rigidities'), from the RE model (labeled 'RE, with rigidities'), and from the LBD model (labeled 'LBD, no rigidities') to a one-standard-deviation financial shock. Columns labeled 'Single shock' in Tables 1 and 2 report the posteriors.[36] According to the local projection, an expansionary financial shock raises output, hours, investment, and consumption in a hump-shaped manner. The federal funds rate rise but inflation does not move, translating into an increase in the real interest rate. Finally, all the three wedges — labor, consumption and the excess return — as well as the forecast dispersion of GDP growth fall.

Our model with endogenous ambiguity matches the local projection response well. First, our model generates persistent, hump-shaped dynamics and co-movement in real quantities. This property is due solely to the endogenous uncertainty mechanism. To see this, in Fig. 3 we calculate the responses when we turn off ambiguity (set the entropy constraint η to 0)

[35] For the prior distribution, we use the Beta distribution for ρ_x and the Gamma distribution for μ_n. Following Chang et al. (2002), the mean (standard deviation) of ρ_x and μ_n is 0.798 (0.012) and 0.111 (0.004), respectively. The posterior mode estimates are $\rho_x = 0.82$ and $\mu_n = 0.14$.

[36] We discipline the standard deviation of the financial shock by the response of the credit spread. Denoting $\varepsilon_{s\Delta}$ the equilibrium elasticity of the spread with respect to the financial shock, which in turn is a function of other structural parameters such as the standard deviation of the entrepreneurial shock and the monitoring cost, the reason we find a larger standard deviation of the financial shock in the ambiguity model is because $\varepsilon_{s\Delta}$ is higher in our model than in the RE model under the estimated parameters.

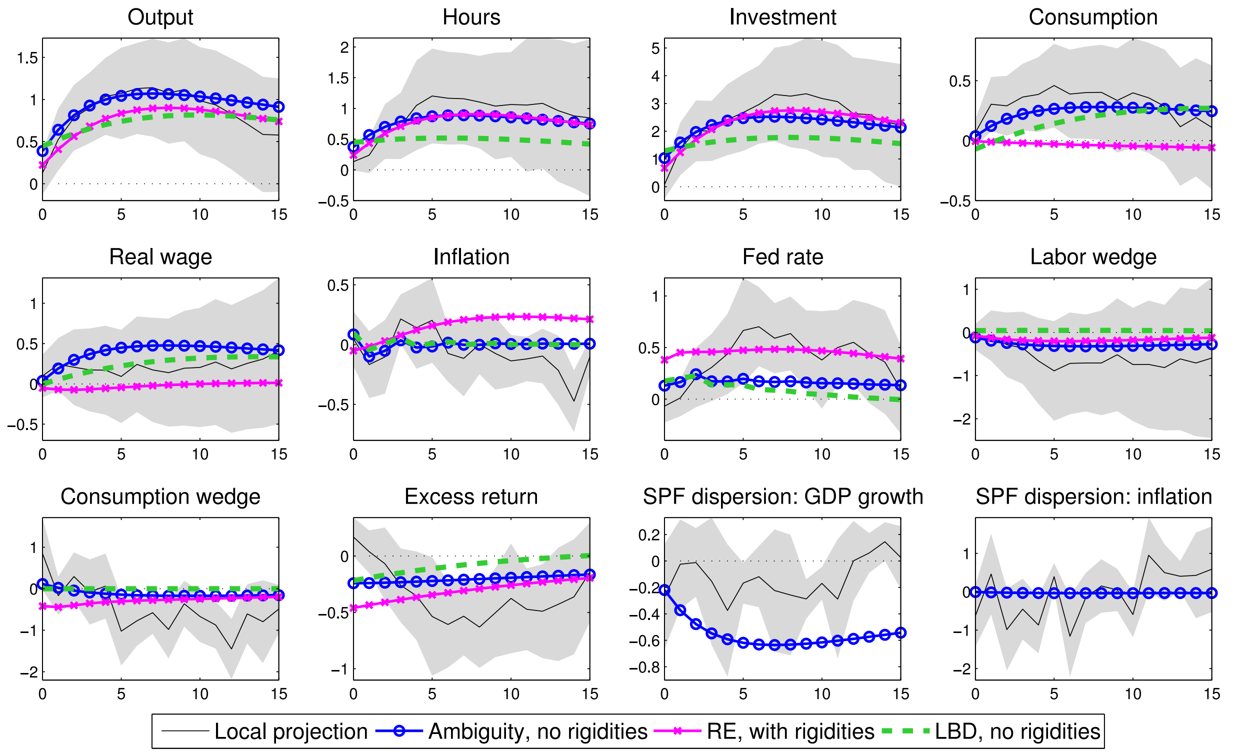

Fig. 2. Responses to a financial shock (single shock estimation). *Notes*: The black lines are the mean responses from the local projection and the shaded areas are the 95% confidence band. The blue circled lines are the impulse responses from the baseline model with ambiguity but without real and nominal rigidities. The purple lines are the impulse responses from the standard RE model featuring real and nominal rigidities. The green dashed lines are the impulse responses from the LBD model without ambiguity or rigidities. All impulse responses are estimated using only the local projection response to the financial shock. The responses of output, hours, investment, consumption, and real wages are in percentage deviations from the steady states while inflation, fed rate, and excess return are in annual percentage points. The rest are in quarterly percentage points.

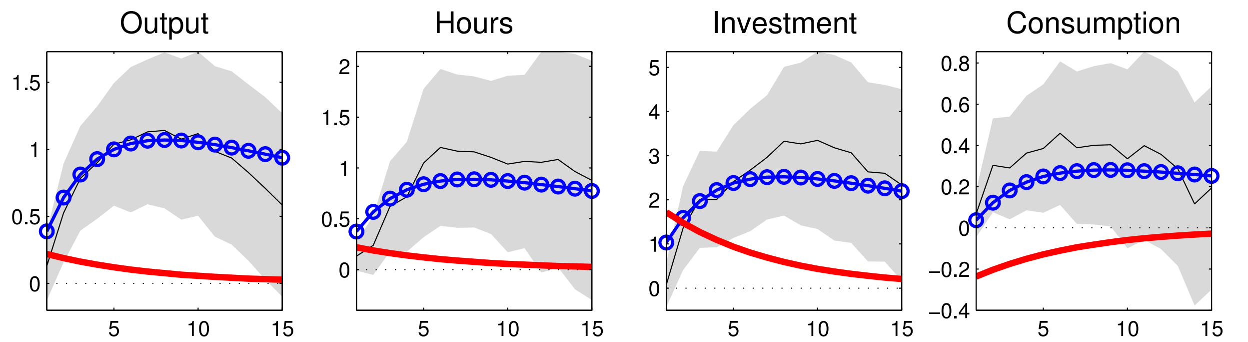

Fig. 3. Responses to a financial shock (single shock estimation): effect of confidence. *Notes*: The black lines are the mean responses from the local projection and the shaded areas are the 95% confidence bands. The blue circled lines are the impulse responses from the baseline model with ambiguity, estimated using only the local projection response to the financial shock. The red solid lines are the impulse responses when we turn off the effect of confidence by setting $\eta = 0$ and re-estimate the baseline model. In both versions, there is no real or nominal rigidity. Responses are in percentage deviations from the steady states.

and re-estimate the baseline model. In sharp contrast to the baseline model, output, hours, and investment all rise initially and then monotonically decrease while consumption declines, consistent with the Barro and King (1984) logic. Second, our model replicates the dynamics of inflation and of the nominal interest rate and hence of the real interest rate. Third, as a result of these successes, the labor and consumption wedges, as well as the excess return fall, as in the data. Fourth, even if not directly targeted in the estimation, the model implies a decline in the forecast range that is in line with the SPF. Finally, in Fig. 4 we report three perturbations of the baseline estimation: lower entropy constraint η, including SPF dispersions in the estimation criterion, and fixing the inverse Frisch elasticity at $\phi = 0.5$. We find that in all cases the estimated impulse responses are similar to the baseline estimation.

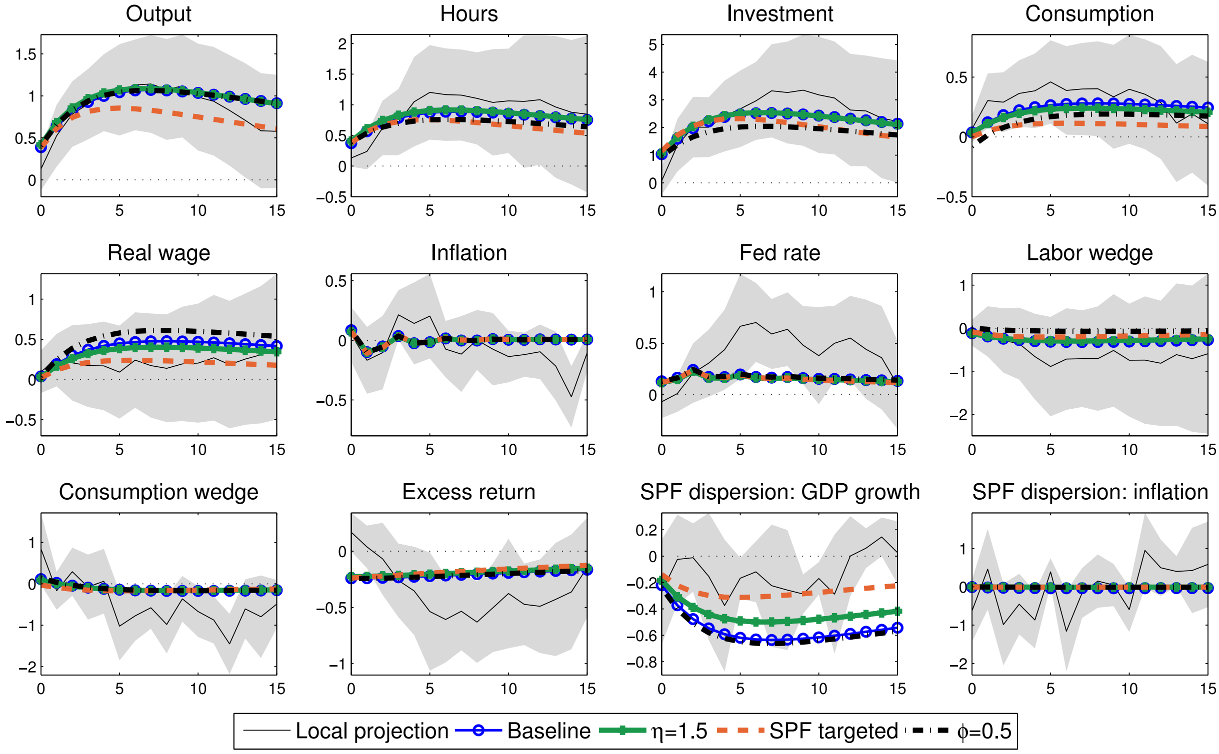

Fig. 4. Responses to a financial shock (single shock estimation): robustness. *Notes*: The black lines are the mean responses from the local projection and the shaded areas are the 95% confidence band. The blue circled lines are the impulse responses from the baseline model with ambiguity but without real and nominal rigidities. The green lines with crosses are the impulse responses from the baseline model with ambiguity but without real and nominal rigidities, where the entropy constraint is fixed at $\eta = 1.5$. The orange dashed lines are the impulse responses from the baseline model with ambiguity but without real and nominal rigidities, where we include the SPF dispersion on GDP growth and inflation in the estimation criterion. The black dashed lines are the impulse responses from the baseline model with ambiguity but without real and nominal rigidities, where the inverse Frisch elasticity is fixed at $\phi = 0.5$. All impulse responses are estimated using only the local projection response to the financial shock. The responses of output, hours, investment, consumption, and real wages are in percentage deviations from the steady states while inflation, fed rate, and excess return are in annual percentage points. The rest are in quarterly percentage points.

Consider now the RE model. The model is able to generate the hump-shaped increase in output, hours, and investment. This is due to the nominal rigidities, where at the posterior mode prices and wages are adjusted roughly every 2.5 and 2 quarters, respectively, and due to real rigidities, where at the posterior mode consumption habit $b = 0.99$ and the investment adjustment cost $\kappa = 0.99$. The RE model, however, cannot match several features of the data. First, the model overpredicts inflation in the medium run, which implies that the model understates the increase in the real interest rate relative to the data. Second, the model generates a flat consumption response.

To understand the real rate and consumption dynamics, consider a standard Euler equation for risk-free assets. Temporarily abstracting from consumption habits, in a first-order approximation, the Euler equation implies that expected consumption growth is equal to the real interest rate. A positive real interest rate thus implies that consumption has to grow over time. However, since the consumption response is hump-shaped in the local projection, a continuous growth in consumption becomes inconsistent with the data in the medium run. Thus, the estimation instead tries to generate a flat consumption path by forcing the consumption habit to be near the upper bound. In the short run, the high consumption habit allows the flat consumption and high real interest rate to coexist. In the medium run, however, the Euler equation implies that the low consumption growth should lead to a lower real interest rate, while in the data the interest rate remains persistently high.

Our model with ambiguity is able to break the counterfactual link between consumption growth and real interest rate through lowering the effective stochastic discount factor in the Euler equation, manifested as a reduction in the consumption wedge. The RE model, in contrast, fails to generate a decline in the consumption wedge in the medium run due to the aforementioned implication of the Euler equation. It also fails to predict a persistent drop in the excess return; instead, in the RE model the excess return traces the financial shock process and hence its fall is transitory. As a result, the data favors our model with ambiguity over the RE model: the marginal likelihood of our model is (-486-(-514)=) 28 log points higher than the RE model (Table 2).

Finally, consider the conceptual and empirical differences between our mechanism and the LBD model. In our proposed model a larger scale of production lowers uncertainty about the unobserved firm-level productivity and it increases the *perceived* return to working and investing. In contrast, in the LBD model a larger scale raises *actual* productivity through accumulation of skill. The two mechanisms have similarities in that co-movement of consumption and labor is now possible out of demand shocks, as in both cases an expansion of economic activity leads to improvements in the return to working. More precisely, Eq. (19) shows that the LBD model can overturn the Barro and King (1984) logic through the endogenous positive response of the realized marginal product of labor to the aggregate economic activity. The difference between the two mechanisms is that the LBD model remains a rational expectations economy so that, even if the realized marginal product of labor is affected by skill accumulation, there is no systematic difference between agents' beliefs and the equilibrium objects recovered by an analyst under the true DGP. Put differently, in the LBD model there are no predictable labor, consumption, and excess-return wedges in Eqs. (21), (24), and (27) respectively, since the belief operators E_t^* and E_t are the same.

Fig. 2 confirms that the LBD model has the potential to generate a simultaneous increase in output, hours, investment, and consumption. However, this alternative model fails to explain the data in two important dimensions. First, it does not replicate the empirical hump-shaped dynamics of hours, investment, and consumption. Second, for the reasons explained above, it generates no variation in the labor and consumption wedges, while the transitory excess return simply traces the exogenous financial shock process. The lack of movement in consumption wedge is reflected in the underprediction of the nominal interest rate. As a result, the marginal likelihood of the LBD model is -653 and is lower than the ambiguity and the RE models.[37]

4.5. Impulse responses for all four shocks: Learning friction helps fit data better

Our second experiment is to estimate the model using *all four* structural shocks. As in the first experiment, we estimate three models: our model with ambiguity, the RE model, and the LBD model. In order to produce real effects of monetary policy shocks, we incorporate nominal rigidities (sticky prices and wages) and for symmetry also real rigidities (consumption habit and investment adjustment cost) into all three models. Therefore, in this subsection we analyze how the standard RE model, with all its standard features, fares against *also* allowing for ambiguity into otherwise the same model. This permits asking to what extent our propagation mechanism quantitatively replaces standard rigidities used in medium-scale DSGE models.

Columns labeled 'All shocks' in Tables 1 and 2 report the posteriors. The ambiguity model requires smaller real and nominal frictions compared to the RE model. For example, at the posterior mode, adjustment of prices and wages occurs every 1.2 and 1.0 quarters, respectively, compared to 10.5 and 2.1 quarters in the RE model. The ambiguity model requires less frictions because the propagation of shocks and the co-movement are driven mostly by our learning mechanism rather than the traditional rigidities.

To further understand this, we first consider the performance of our model with endogenous uncertainty. Fig. 5 shows the impulse response to a one-standard-deviation IST shock. In the local projection, an expansionary IST shock leads to a persistent increase in output, hours, investment, and consumption. The shock also raises real wages and, in the short run, the nominal rate. Our model with ambiguity replicates data quite well; the model generates co-movement as well as the reduction in labor and consumption wedges and excess return. In addition, the fall in the forecast range of GDP growth is in line with the data. These properties are mainly due to learning. Indeed, when we set the entropy constraint η to zero while fixing other parameters at their estimated values, the impulse response is transitory, consumption mildly declines rather than increases, and the wedges stay constant.

Next, in Fig. 6, we report the impulse response to a monetary policy shock. In the data, an unexpected interest rate cut raises output, hours, investment, and consumption in a hump-shaped manner. In our model a surprise fall in interest rate also leads to a persistent expansion of output, hours, investment, and consumption. However, the model underpredicts the increase in consumption. The propagation of the monetary policy shock is driven by the confidence channel: When we turn off ambiguity, the real effect of a monetary policy shock is small and transitory.

We briefly discuss our model's fit regarding the financial and neutral technology shocks and report the figures in the online appendix. For the financial shock, as in the single shock estimation, our model is broadly successful in replicating the data. According to the local projection, a technology shock raises output and consumption but slightly reduces hours in the short-run, in line with the conventional finding in the literature such as Galí (1999). We find that our model fit the data reasonably well; in particular, the relatively moderate degree of estimated real and nominal rigidities in the ambiguity model is sufficient to generate the short-run decline in hours. In addition, the ambiguity model matches the empirical response of GDP forecast dispersion. Finally, our learning mechanism has a relatively moderate impact on the propagation of a technology shock. This is because both in the data and in our model a positive technology shock does not lead to a significant increase in inputs.

We find that the RE model does not fit the empirical impulse responses very well compared to our model. For example, while the RE model generates a persistent and hump-shaped responses thanks to real and nominal rigidities, in response to financial and IST shocks, consumption, inflation, and the nominal interest rate essentially remain constant. To understand this, note that the relatively high value of consumption habit ($b = 0.73$) allows the RE model to accommodate the main

[37] In the baseline version of the LBD model, we assume that the skill accumulated by learning-by-doing is reflected in the measured TFP. An alternative estimation where the skill level is not reflected in the measured TFP produces very similar results.

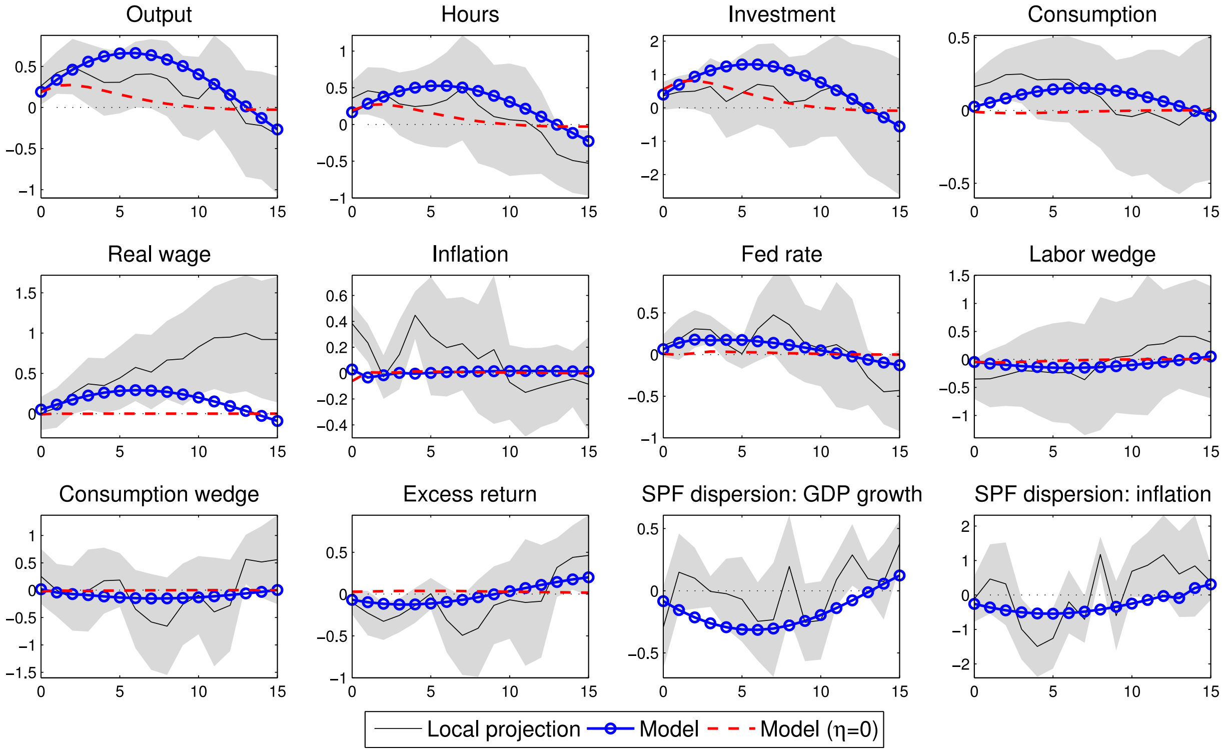

Fig. 5. Responses to an investment-specific technology shock. *Notes*: The black lines are the mean responses from the local projection and the shaded areas are the 95% confidence bands. The blue circled lines are the impulse responses from the baseline model with ambiguity, estimated using the local projection responses to all four structural shocks (technology, investment-specific, financial, and monetary policy). The red dashed lines are the counterfactual responses where we set the entropy constraint η to 0, while holding other parameters at the estimated values. The responses of output, hours, investment, consumption, and real wages are in percentage deviations from the steady states while inflation, fed rate, and excess return are in annual percentage points. The rest are in quarterly percentage points.

property of an expansionary monetary policy shock: consumption grows while the real interest rate is falling. While this negative co-movement between consumption and the real interest rate helps the RE model match the responses to a monetary policy shock, it becomes inconsistent with the responses to the financial shock. In order to strike a balance between matching consumption and the real interest rate, the four shocks estimation chooses parameter values so that both variables remain roughly constant in response to a financial shock.[38] The RE model also fails to generate a reduction in the consumption wedge and excess return in response to IST and financial shocks. Hence, our ambiguity model improves upon the RE model in terms of log marginal likelihood, which penalizes more parameters, by (-2705-(-3092)=) 387 log points.

We also find that the fit improves when we introduce learning-by-doing into the RE model with real and nominal rigidities. However, the marginal likelihood of that model is still lower than that of our model with ambiguity: log marginal likelihood of the LBD model is -2755 compared to -2705 for our model. This is because the LBD model, as in the single shock estimation, fails to generate countercyclical consumption wedge and excess return. We collect the complete impulse responses of the RE and LBD models in the online appendix.

We conclude by discussing two additional points regarding our estimation. First, as shown in Fig. 7, the model implies responses of the 1-year-ahead forecast range of GDP growth and inflation that are similar to the SPF. Second, we consider to what extent our model comparison exercises are affected by the fact that we include wedges in the estimation criterion. To do so, we repeat our estimations excluding wedges. In the single shock estimation, the marginal likelihood of the baseline model with ambiguity is -375 while for the RE and LBD models the marginal likelihoods are -389 and -496, respectively. In the four shocks estimation, the marginal likelihood of the baseline model with ambiguity is -2073 while for the RE and LBD models the marginal likelihoods are -2481 and -2176, respectively. Thus, while the difference between the baseline model and the RE/LBD models in terms of log points shrinks in the single shock estimation and widens in the four shocks

[38] Note that, in contrast to the single shock case, the estimation cannot resort to an extremely high consumption habit because it needs to match the consumption responses of neutral technology and monetary policy shocks. This is why the four shocks estimation not only requires consumption but also the interest rate (inflation and the nominal interest rate) to be flat in response to the financial shock.

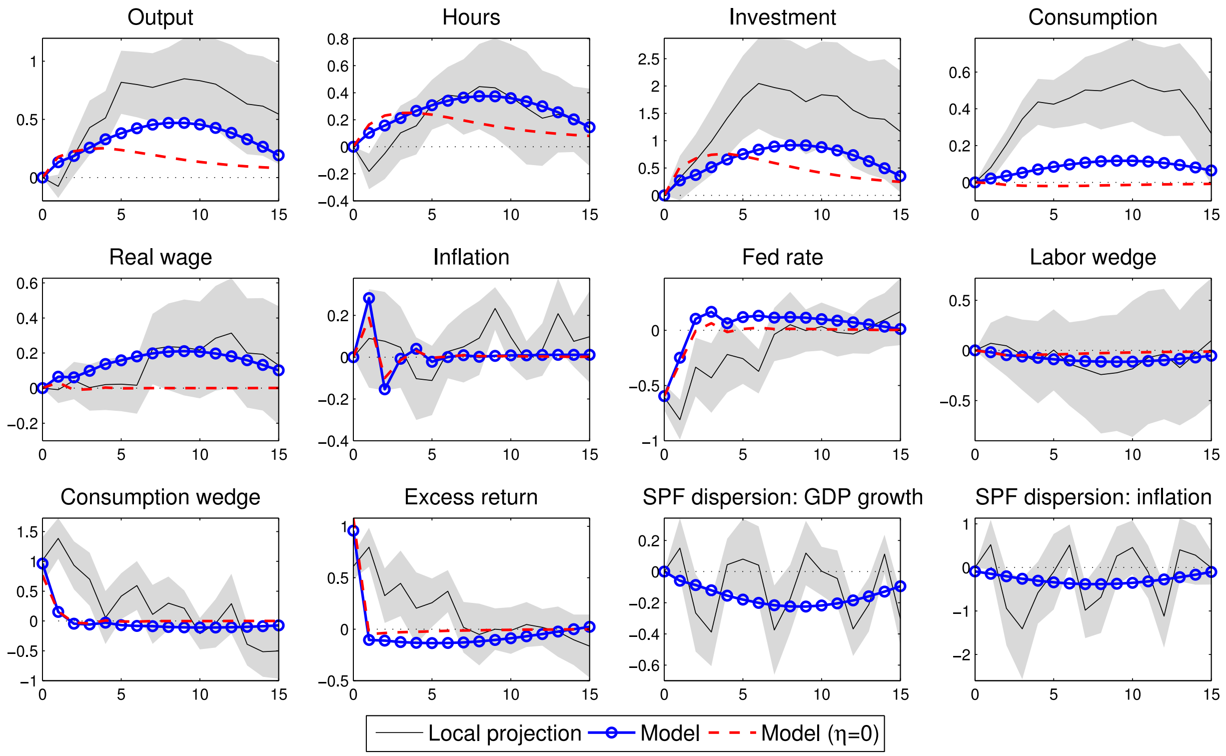

Fig. 6. Responses to a monetary policy shock. *Notes*: See notes from Fig. 5.

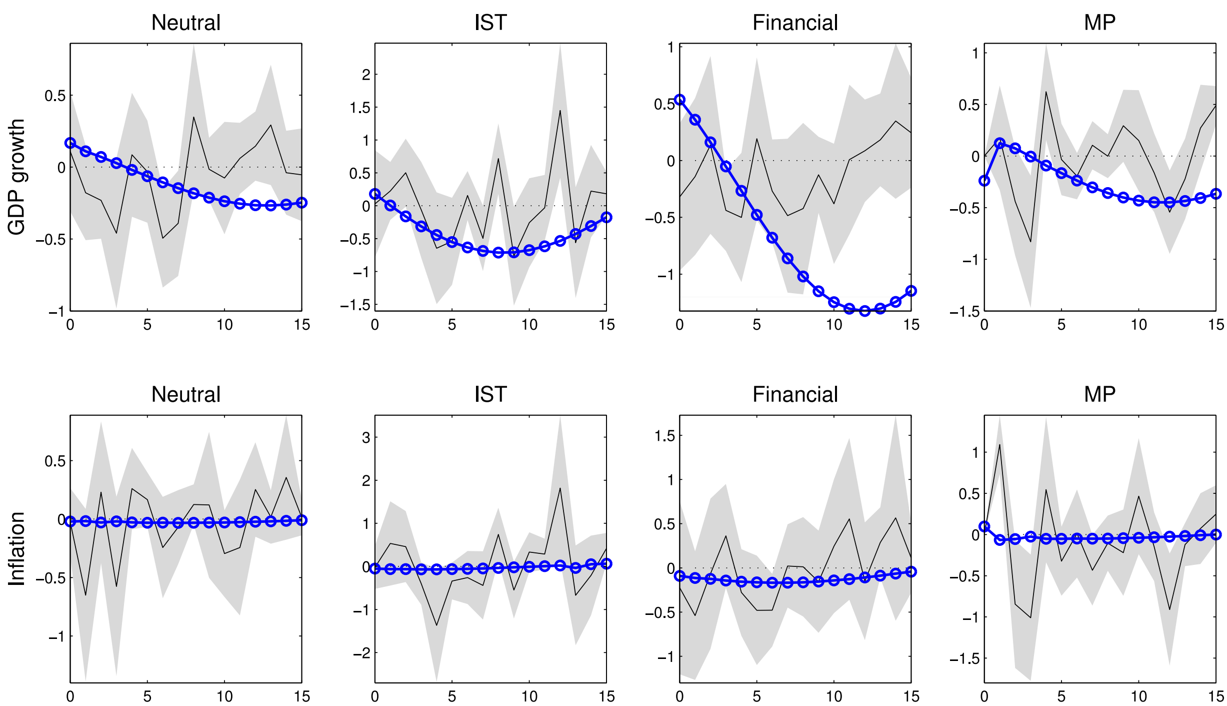

Fig. 7. 1-year ahead forecast range. *Notes*: The top panel plots the ranges of 1-year-ahead GDP forecast and the bottom panel plots the ranges of 1-year-ahead in inflation forecast for neutral technology, investment-specific technology, financial, and monetary policy shocks. The black lines are the mean responses from the local projection and the shaded areas are the 95% confidence band. The blue circled lines are the impulse responses from the baseline model with ambiguity.

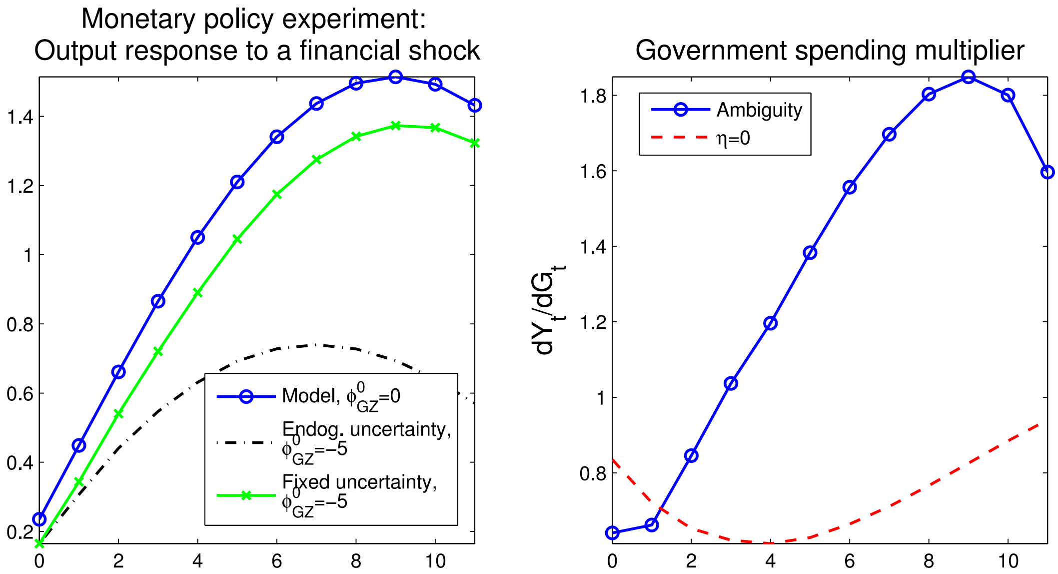

Fig. 8. Policy experiments. *Notes*: The left panel plots the output response to a financial shock. The blue circled line is the baseline model with ambiguity, estimated using the local projection responses to all four structural shocks. The black dashed line is the counterfactual where the Taylor rule coefficient on the GZ spread is $\phi^0_{GZ} = -5$. The green line is the response when $\phi^0_{GZ} = -5$ but with the path of uncertainty fixed at the original one. The right panel plots the government spending multiplier for output. The economy is hit by a positive spending shock at $t = 0$ and the path of government spending follows an AR(1) process. The blue circled line is the multiplier from the baseline model with ambiguity, estimated using the local projection responses to all four structural shocks. The red dashed line is the multiplier where $\eta = 0$, holding other parameters at the estimated values.

estimation,[39] the ranking between the models do not change in both the single and the four shocks estimation. We also find that the estimated parameters are very similar whether or not wedges are targeted in the estimation.

4.6. Policy implications

The fact that in our model uncertainty is endogenous has important policy implications. We illustrate this point by conducting two policy experiments, using parameter values that are based on the four shocks estimation throughout the analysis. First, we evaluate the impact of modifying the Taylor rule to incorporate an adjustment to the credit spread. In the left panel of Fig. 8 we report the impulse response of output to the financial shock in the ambiguity model as we keep all parameters at their baseline estimated values, but change the Taylor rule coefficient on the credit spread ϕ^0_{GZ} from its original value of zero to -5. The output effect decreases by roughly half when monetary policy responds aggressively to the spread movements (black dashed line). Much of the reduction in the output effect comes from stabilizing the endogenous variation in uncertainty. To see this, we show the effects of policy changes in the economy where the path of uncertainty is fixed to the original one. In this economy, a change in ϕ^0_{GZ} has a much smaller effect (green line).

Second, we consider fiscal policy effects. In standard models, an increase in government spending crowds out consumption and hence the government spending multiplier on output, dY_t/dG_t, tends to be modest and below one. In our model, however, an increase in hours worked triggered by an increase in government spending raises agents' confidence, which feeds back and raises the level of consumption and other economic activities. Because of this amplification effect, the government spending multiplier could be larger and above one. In the right panel of Fig. 8, we plot the multiplier in our estimated model after a one-time, positive shock to government spending at $t = 0$.[40] The model predicts a multiplier that becomes larger than one after four periods and stays persistently and significantly above one. In contrast, in a counterfactual economy where $\eta = 0$ and other parameters are at their estimated values, the multiplier stays persistently below or around one.

In our model learning arises at firm-level and therefore there are no information externalities that the government can correct. This is an important contrast to models where learning occurs through observing the aggregate economy and highlights the relevance of modeling the underlying source of uncertainty for evaluating policies. At a more general level, the

[39] Recall that when wedges are included in the estimation criterion, the baseline model has a marginal likelihood that is larger than for the RE and LBD models by 28 and 167 log points, respectively, in the single shock estimation. In the four shocks estimation, that difference in marginal likelihoods is 387 and 50 log points, respectively. When wedges are excluded in the estimation criterion, the baseline model's marginal likelihood is larger than for the RE and LBD models by 14 and 121 log points, respectively, in the single shock estimation. In the four shocks estimation, that difference becomes 408 and 103 log points, respectively.

[40] We assume that the government spending G_t in the resource constraint (10) is given by $G_t = g_t Y_t$, where g_t follows $\ln g_t = (1 - 0.8)\ln \bar{g} + 0.8 \ln g_{t-1} + \epsilon_{g,t}$.

comparisons of these counterfactual models in the monetary and fiscal policy experiments underscore the importance for policy analysis of modeling time-variation in uncertainty as an endogenous response.

4.7. Evidence from firm-level survey data

We provide a further test of the model by comparing our model-implied firm-level confidence process with the time-series moments of uncertainty directly measured from the micro survey data. Our measure of confidence is the cross-sectional average dispersion of firm-level capital return forecasts. We use a series constructed by Senga (2018) using I/B/E/S and Compustat data.[41] For each firm, Senga (2018) measures the min-max range across analysts' forecasts of the return on capital for that firm. Taking the cross-sectional average across firms of that forecast range results in a time-series measure. For the model counterpart, we calculate the range D_t^R of capital return forecasts using the formula $D_t^R = \varepsilon_{Rz} 2|\int_0^1 a_{l,t} dl|$, where ε_{Rz} is the equilibrium elasticity of firm-level capital return with respect to the innovation to firm-level profitability in the linear decision rule.[42]

First, we find that the model matches the ratio of the standard deviation of the forecast range to that of real GDP reasonably well, at 4.4 in the model versus 4.9 in the data.[43] Second, the forecast dispersion and real GDP is negatively correlated, although the model overstates the negative correlation, at -0.96 in the model compared to -0.38 in the data. On the one hand, in the model uncertainty is driven solely by changes in economic activity, thus producing a strong negative co-movement. On the other hand, the firm-level data could be subject to measurement errors, which tend to bias the correlation between the forecast range and real GDP towards zero, while increasing the measured standard deviation of the range. Overall, we find that the time-series properties of uncertainty implied by our estimated model broadly matches firm-level data that was not used in the estimation.[44]

5. Conclusion

In this paper, we build a tractable heterogeneous-firm business cycle model where firms face Knightian uncertainty about their profitability and learn it through production. We show how, even in the absence of any other frictions, the feedback mechanism endogenously generates empirically desirable cross-equation restrictions such as: co-movement driven by demand shocks, amplified and hump-shaped dynamics, and countercyclical correlated wedges in the equilibrium conditions for labor, risk-free and risky assets. We embed our learning mechanism into a standard medium-scale model and estimate it by matching impulse responses of macroeconomic aggregates and asset prices to financial, monetary policy and neutral and investment-specific technology shocks. We find that our model improves on conventional models in replicating impulse responses, requires less real and nominal rigidities and predicts magnified responses of economic activity to monetary and fiscal policies, while at the same time producing a confidence process that is consistent with the survey data both at the macro and micro level.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.jmoneco.2020.01.010.

CRediT authorship contribution statement

Cosmin Ilut: Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing - original draft, Writing - review & editing, Visualization, Supervision, Project administration. **Hikaru Saijo:** Conceptualization, Methodology, Software, Validation, Formal analysis, Investigation, Resources, Data curation, Writing - original draft, Writing - review & editing, Visualization, Supervision, Project administration.

References

Angeletos, G.-M., 2018. Frictional coordination. J. Eur. Econ. Assoc. 16 (3), 563–603.
Angeletos, G.-M., Collard, F., Dellas, H., 2018a. Business cycle anatomy. NBER Working Paper 24875.
Angeletos, G.-M., Collard, F., Dellas, H., 2018. Quantifying confidence. Econometrica 86 (5), 1689–1726.
Angeletos, G.-M., La'O, J., 2009. Noisy business cycles. In: NBER Macroeconomics Annual 2009, Volume 24, pp. 319–378.

[41] We thank Tatsuro Senga for generously sharing his data.

[42] In terms of mapping the model to data, the idea is similar to the SPF dispersion on real GDP growth: stronger disagreement among experts about conditional firm-level mean returns is reflected in the agent's lower confidence in probability assessments.

[43] The sample period is 1977–2008, where the starting point corresponds to the initial availability of Senga (2018) data and we exclude the zero lower bound period by trimming the observations after 2008. The model moments are obtained by simulating the model from a four shock estimation, where we also include the SPF dispersion in the estimation criterion, and are annualized so that it matches the frequency of the data by Senga (2018). All variables are linearly de-trended.

[44] Senga (2018) additionally shows that the firms with higher uncertainty, measured in terms of analysts' forecast dispersion, tend to have lower sales and employment. This empirical observation is consistent with our theory that low activity is associated with high uncertainty at the firm level.

Angeletos, G.-M., La'O, J., 2013. Sentiments. Econometrica 81 (2), 739–779.
Angeletos, G.-M., Lian, C., 2016. A (real) theory of the keynesian multiplier. MIT mimeo.
Asparouhova, E., Bossaerts, P., Eguia, J., Zame, W., 2015. Asset pricing and asymmetric reasoning. J. Political Econ. 123 (1), 66–122.
Baker, S.R., Bloom, N., Davis, S.J., 2016. Measuring economic policy uncertainty. Q. J. Econ. 131 (4), 1593–1636.
Barro, R.J., King, R.G., 1984. Time-separable preferences and intertemporal-substitution models of business cycles. Q. J. Econ. 99 (4), 817–839.
Barsky, R.B., Sims, E.R., 2012. Information, animal spirits, and the meaning of innovations in consumer confidence. Am. Econ. Rev. 102 (4), 1343–1377.
Benhabib, J., Spiegel, M.M., 2018. Sentiments and economic activity: evidence from U.S. states. Econ. J. 129 (618), 715–733.
Benhabib, J., Wang, P., Wen, Y., 2015. Sentiments and aggregate demand fluctuations. Econometrica 83 (2), 549–585.
Bernanke, B.S., Gertler, M., Gilchrist, S., 1999. The financial accelerator in a quantitative business cycle framework. In: Taylor, J.B., Woodford, M. (Eds.), Handbook of Macroeconomics, 1, Part C. Elsevier, Amsterdam, pp. 1341–1393.
Bhandari, A., Borovička, J., Ho, P., 2016. Identifying ambiguity shocks in business cycle models using survey data. NBER WP 22225.
Bianchi, F., Ilut, C.L., Schneider, M., 2018. Uncertainty shocks, asset supply and pricing over the business cycle. Rev. Econ. Stud. 85, 810–854.
Bloom, N., Floetotto, M., Jaimovich, N., Saporta-Eksten, I., Terry, S., 2018. Really uncertain business cycles. Econometrica. Forthcoming
Bossaerts, P., Ghirardato, P., Guarnaschelli, S., Zame, W., 2010. Ambiguity in asset markets: Theory and experiment. Rev. Financ. Stud. 23 (4), 13–25.
Bottazzi, G., Secchi, A., 2003. Common properties and sectoral specificities in the dynamics of us manufacturing companies. Rev. Ind. Org. 23 (3-4), 217–232.
Caplin, A., Leahy, J., 1993. Sectoral shocks, learning, and aggregate fluctuations. Rev. Econ. Stud. 60 (4), 777–794.
Chang, Y., Gomes, J.F., Schorfheide, F., 2002. Learning-by-doing as a propagation mechanism. Am. Econ. Rev. 92 (5), 1498–1520.
Chari, V.V., Kehoe, P.J., McGrattan, E.R., 2007. Business cycle accounting. Econometrica 75 (3), 781–836.
Chetty, R., Guren, A., Manoli, D., Weber, A., 2011. Are micro and macro labor supply elasticities consistent? a review of evidence on the intensive and extensive margins. Am. Econ. Rev. 101 (3), 471–475.
Christiano, L.J., Eichenbaum, M., Evans, C.L., 2005. Nominal rigidities and the dynamic effects of a shock to monetary policy. J. Political Econ. 113 (1), 1–45.
Christiano, L.J., Eichenbaum, M., Trabandt, M., 2015. Understanding the great recession. Am. Econ. J.: Macroecon. 7 (1), 110–167.
Christiano, L.J., Motto, R., Rostagno, M., 2014. Risk shocks. Am. Econ. Rev. 104 (1), 27–65.
Christiano, L. J., Trabandt, M., Walentin, K., 2010. Involuntary unemployment and the business cycle. NBER Working Paper.
Coad, A., 2007. Firm growth: a survey. CES Working Papers.
Cochrane, J.H., 2011. Presidential address: Discount rates. J. Finance 66 (4), 1047–1108.
Coibion, O., Gorodnichenko, Y., Kueng, L., Silva, J., 2017. Innocent bystanders? monetary policy and inequality. J. Monet. Econ. 88, 70–89.
d'Alessandro, A., Fella, G., Melosi, L., 2019. Fiscal stimulus with learning-by-doing. Int. Econ. Rev.. Forthcoming
David, J.M., Hopenhayn, H.A., Venkateswaran, V., 2015. Information, misallocation, and aggregate productivity. Q. J. Econ. 131, 943–1005.
Del Negro, M., Schorfheide, F., 2013. Dsge model-based forecasting. Handbook of Economic Forecasting 2, 57–140.
Eggertsson, G.B., Woodford, M., 2003. Zero bound on interest rates and optimal monetary policy. Brookings Pap. Econ. Activity 2003 (1), 139–233.
Ellsberg, D., 1961. Risk, ambiguity, and the savage axioms. Q. J. Econ. 643–669.
Epstein, L.G., Schneider, M., 2003. Iid: Independently and indistinguishably distributed. J. Econ. Theory 113 (1), 32–50.
Epstein, L.G., Schneider, M., 2003. Recursive multiple-priors. J. Econ. Theory 113 (1), 1–31.
Epstein, L.G., Schneider, M., 2008. Ambiguity, information quality, and asset pricing. J. Finance 63 (1), 197–228.
Eusepi, S., Preston, B., 2011. Expectations, learning, and business cycle fluctuations. Am. Econ. Rev. 101 (6), 2844–2872.
Eusepi, S., Preston, B., 2015. Consumption heterogeneity, employment dynamics and macroeconomic co-movement. J. Monet. Econ. 71, 13–32.
Fajgelbaum, P., Schaal, E., Taschereau-Dumouchel, M., 2017. Uncertainty traps. Q. J. Econ. 132, 1641–1692.
Fernald, J.G., 2014. A quarterly, utilization-adjusted series on total factor productivity. Working Paper.
Fisher, J., 2006. The dynamic effects of neutral and investment-specific technology shocks. J. Political Econ. 114 (3), 413–451.
Fisher, J.D., 2015. On the structural interpretation of the smets-wouters risk premium shock. J. Money Credit Bank. 47 (2-3), 511–516.
Galí, J., 1999. Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? Am. Econ. Rev. 89 (1), 249–271.
Gilchrist, S., Zakrajšek, E., 2011. Monetary policy and credit supply shocks. IMF Econ. Rev. 59 (2), 195–232.
Gilchrist, S., Zakrajšek, E., 2012. Credit spreads and business cycle fluctuations. Am. Econ. Rev. 102 (3), 1692–1720.
Gorodnichenko, Y., Lee, B., 2017. A note on variance decomposition with local projection. Working Paper.
Gust, C., Herbst, E., López-Salido, D., Smith, M.E., 2017. The empirical implications of the interest-rate lower bound. Am. Econ. Rev. 107 (7), 1971–2006.
Huo, Z., Takayama, N., 2015. Higher order beliefs, confidence, and business cycles. University of Minnesota working paper.
Hymer, S., Pashigian, P., 1962. Firm size and rate of growth. J. Political Econ. 556–569.
Ilut, C., Schneider, M., 2014. Ambiguous business cycles. Am. Econ. Rev. 104 (8), 2368–2399.
Ilut, C., Valchev, R., Vincent, N., 2016. Paralyzed by fear: rigid and discrete pricing under demand uncertainty. NBER Working Paper 22490.
Jaimovich, N., Rebelo, S., 2009. Can News about the Future Drive the Business Cycle? Am. Econ. Rev. 99 (4), 1097–1118.
Jordá, O., 2005. Estimation and inference of impulse responses by local projections. Am. Econ. Rev. 95 (1), 161–182.
Jovanovic, B., 1982. Selection and the evolution of industry. Econometrica 649–670.
Marinacci, M., 1999. Limit laws for non-additive probabilities, and their frequentist interpretation. J. Econ. Theory 84 (2), 145–195.
Milani, F., 2011. Expectation shocks and learning as drivers of the business cycle. Econ. J. 121 (552), 379–401.
Milani, F., 2017. Sentiment and the u.s. business cycle. J. Econ. Dyn. Control 82, 289–311.
van Nieuwerburgh, S., Veldkamp, L., 2006. Learning asymmetries in real business cycles. J. Monet. Econ. 53 (4), 753–772.
Ordoñez, G., 2013. The asymmetric effects of financial frictions. J. Political Econ. 5 (121), 844–895.
Phillipon, T., 2015. Has the us finance industry become less efficient? on the theory and measurement of financial intermediation. Am. Econ. Rev. 105 (4), 1408–1438.
Romer, C.D., Romer, D.H., 2004. A new measure of monetary shocks: derivation and implications. Am. Econ. Rev. 94 (4), 1055–1084.
Saijo, H., 2017. The uncertainty multiplier and business cycles. J. Econ. Dyn. Control 78, 1–25.
Senga, T., 2018. A new look at uncertainty shocks: Imperfect information and misallocation. Working Paper.
Shimer, R., 2009. Convergence in macroeconomics: The labor wedge. Am. Econ. J.: Macroecon. 1 (1), 280–297.
Smets, F., Wouters, R., 2007. Shocks and frictions in U.S. business cycles: aBayesian DSGE approach. Am. Econ. Rev. 97 (3), 586–607.
Stanley, M.H., Amaral, L.A., Buldyrev, S.V., Havlin, S., Leschhorn, H., Maass, P., Salinger, M.A., Stanley, H.E., 1996. Scaling behaviour in the growth of companies. Nature 379 (6568), 804–806.
Straub, L., Ulbricht, R., 2016. Endogenous uncertainty and credit crunches. Manuscript.