UNCERTAINTY TRAPS*

Pablo D. Fajgelbaum Edouard Schaal Mathieu Taschereau-Dumouchel

We develop a theory of endogenous uncertainty and business cycles in which short-lived shocks can generate long-lasting recessions. In the model, higher uncertainty about fundamentals discourages investment. Since agents learn from the actions of others, information flows slowly in times of low activity and uncertainty remains high, further discouraging investment. The economy displays uncertainty traps: self-reinforcing episodes of high uncertainty and low activity. Although the economy recovers quickly after small shocks, large temporary shocks may have long-lasting effects on the level of activity. The economy is subject to an information externality but uncertainty traps may remain in the efficient allocation. Embedding the mechanism in a standard business cycle framework, we find that endogenous uncertainty increases the persistence of large recessions and improves the performance of the model in accounting for the Great Recession. *JEL Codes*: E32, D80.

I. INTRODUCTION

We develop a theory of endogenous uncertainty and business cycles. The theory combines two forces: higher uncertainty about economic fundamentals deters investment, and uncertainty evolves endogenously because agents learn from the actions of others. The unique rational expectations equilibrium of the economy features uncertainty traps: self-reinforcing episodes of high uncertainty and low economic activity that cause recessions to persist. Because of uncertainty traps, large but short-lived shocks can generate long-lasting recessions. We first build and characterize a model that only includes the essential features that give rise to uncertainty traps. Then we embed these features into a standard

*We thank the editor, Robert Barro, and two anonymous referees for valuable suggestions. Liyan Shi and Chunzan Wu provided superb research assistance. We are grateful to Andrew Abel, Andrew Atkeson, Michelle Alexopoulos, Jess Benhabib, Harold Cole, Michael Evers, João Gomes, William Hawkins, Patrick Kehoe, Lars-Alexander Kuehn, Ali Shourideh, Laura Veldkamp, Pierre-Olivier Weill, and seminar participants for useful comments. Edouard Schaal acknowledges financial support from the Spanish Ministry of Economy and Competitiveness, through the Severo Ochoa Programme for Centres of Excellence in R&D (SEV2015-0563).

The Quarterly Journal of Economics (2017), 1641–1692. doi:10.1093/qje/qjx021. Advance Access publication on May 29, 2017.

[@] The Author(s) 2017. Published by Oxford University Press on behalf of the President and Fellows of Harvard College. All rights reserved. For Permissions, please email: journals.permissions@oup.com

real business cycle model and quantify the impact of endogenous uncertainty during the Great Recession.

In the model, firms decide whether to undertake an irreversible investment whose return depends on an imperfectly observed fundamental that evolves randomly according to a persistent process. Firms are heterogeneous in the cost of undertaking this investment and hold common beliefs about the fundamental. Beliefs are regularly updated with new information, as firms learn by observing the return on the investment of other producers. We define uncertainty as the variance of these beliefs.

This environment naturally produces an interaction between beliefs and economic activity. Firms are more likely to invest if their beliefs about the fundamental have higher mean and if they have smaller variance (lower uncertainty). At the same time, the laws of motion for the mean and variance of beliefs depend on the investment rate. In particular, when few firms invest, little information is released, so uncertainty rises.

The key feature of the model is that this interaction between information and investment leads to uncertainty traps, formally defined as the coexistence of multiple stationary points in the dynamics of uncertainty and economic activity. Without shocks, the economy converges to either a high regime (with high economic activity and low uncertainty) if the current level of uncertainty is sufficiently low, or to a low regime (with low activity and high uncertainty) if the current level of uncertainty is sufficiently high. Because of the presence of multiple stationary points, the economy exhibits nonlinearities in its response to shocks: starting from the high regime, it quickly recovers after small temporary shocks, but it may shift to the low-activity regime after a large temporary shock. Once it has fallen into the low regime, only a large enough positive shock can push the economy back to the high-activity regime.

An important feature of the model is that despite the presence of uncertainty traps, there is a unique recursive competitive equilibrium. That is, multiplicity of stationary points does not mean multiplicity of equilibria. Therefore, unlike other macro models with complementarities, there is no room in our model for multiple equilibria or sunspots.¹

^{1.} For recent examples of business cycle models with multiple equilibria see Farmer (2013), Benhabib, Wang, and Wen (2015), Schaal and Taschereau–Dumouchel (2015, 2016), and Kaplan and Menzio (2016).

The model features an inefficiently low level of investment because agents do not internalize the effect of their actions on public information. This inefficiency naturally creates room for welfareenhancing policy interventions. We therefore study the problem of a constrained planner that is subject to the same informational constraints as private agents. The socially constrained-efficient allocation can be implemented with a subsidy to investment, but the optimal policy does not necessarily eliminate uncertainty traps. Therefore, while policy interventions are desirable, they do not eliminate the adverse feedback loop between uncertainty and economic activity.

To evaluate the quantitative implications of uncertainty traps, we embed the key features of the baseline model into a standard general equilibrium framework and then compare its predictions with a real business cycle (RBC) model and the data. To isolate the impact of endogenous movements in uncertainty, we also compare our full model to a "fixed θ -uncertainty" version in which uncertainty about the fundamental productivity θ is fixed over time. We discipline the key parameters of the model, those that determine option-value effects and the evolution of uncertainty, by targeting moments from the distribution of uncertainty about real GDP growth from the Survey of Professional Forecasters (SPF).

We first show that our calibrated model performs as well as the RBC and fixed θ -uncertainty models in terms of traditional business cycle moments. Therefore, incorporating endogenous uncertainty in a standard business cycle model does not impair its ability to predict well-known patterns of business cycle data.

Then we demonstrate that the nonlinearities generated by uncertainty traps, studied in the baseline theory, are active in the calibrated model. Specifically, we compute the economy's response to one-period negative shocks to beliefs of different magnitudes. We find that (i) recessions are longer and deeper under the full model than under the fixed θ -uncertainty model, and (ii) the difference between both models is more important for large shocks than for small ones. In response to a -1% shock, the ensuing recession is 22% deeper (in terms of the peak-to-trough fall in output) and 40% longer (in terms of quarters until the economy has recovered half of the peak-to-trough fall in output) in the full model than in the fixed θ -uncertainty model. However, in response to a larger -5% shock, the recession is 35% deeper and 66% longer in the full model than in the fixed θ -uncertainty model. Therefore, in the calibrated model, the endogenous uncertainty mechanism, whose impact is captured by the difference between the full and fixed θ -uncertainty models, makes recessions deeper and longer for shocks of any magnitude but relatively more so for larger shocks.

Finally, our main quantitative exercise evaluates the predictions of our calibrated model for past U.S. recessions. Because our mechanism provides amplification and persistence to large shocks, we expect that it might help explain particularly severe recessions observed in the data. We therefore investigate the largest recession in our sample, the Great Recession. To do so, we feed each of the three models (our full model, the RBC model, and the fixed θ -uncertainty model) with the observed total factor productivity (TFP) series and signals such that each model replicates the time series of forecasts about output growth from the SPF during the first part of the recession. We then contrast each model's response with the data.

Our main quantitative finding is that during the Great Recession, our model generates declines in output, consumption, employment, and investment which are clearly more protracted. and closer to patterns observed in the data, than what the alternative models predict. Endogenous uncertainty adds 1.8 percentage points in terms of the recession's depth and slows the recovery by about two years relative to the fixed θ -uncertainty model.² The corresponding numbers are 5.2 percentage points and five years when comparing to the RBC model. We also evaluate the performance of the three models against the data in terms of one key statistic that summarizes both the depth and the length of the recession: the cumulative output loss between the start of the recession and 2015. We find that our model generates 93% of the Great Recession's cumulative output loss, relative to the 70% and 30% generated by the fixed θ -uncertainty and RBC models, respectively. Reassuringly, the model also generates patterns for the evolution of uncertainty about output growth that are roughly consistent with the data.

We demonstrate the robustness of these conclusions by replicating this exercise under alternative assumptions about

^{2.} This measure refers to the time the economy takes to recover 20% of its peak-to-trough decline. We use the 20% threshold instead of the usual half-life since in the data, detrended output has only recovered about 20% of its peak-to-trough decline by the end of our sample in 2015.

the shocks hitting the economy, the source of the TFP data, the detrending strategy, and the preferences of the household. In each case, we find that the model with endogenous uncertainty performs better than its alternatives. To make sure that the full model does not generate counterfactual amounts of persistence for milder recessions, we also replicate the second largest recession in our sample, the 1981–1982 recession, which was characterized by a relatively rapid recovery. We find that our model behaves similarly to the RBC and to the fixed θ -uncertainty models in that case. We conclude that the inclusion of uncertainty traps in a standard macroeconomic model of business cycles improves its performance during the Great Recession and leads to predictions similar to standard models for smaller recessions.

The remainder of the introduction contains the literature review followed by a discussion of our notion of uncertainty and its business cycle properties. The article is then structured as follows. Section II presents the baseline model and the definition of the recursive equilibrium. Section III characterizes the investment decision of an individual firm and demonstrates the existence and uniqueness of the equilibrium. Section IV shows the existence of uncertainty traps, examines the nonlinearities that they generate, and characterizes the planner's problem. Section V describes the quantitative model, shows how uncertainty traps influence the response of the economy to shocks, and compares the dynamic properties of our model to an RBC model, a fixed θ -uncertainty model, and the data over the Great Recession. Section VI concludes. The full statement of the proposition and the proofs can be found in the Online Appendix.

I.A. Relation to the Literature

The theory is motivated by an empirical literature that investigates the impact of uncertainty on economic activity using vector autoregressions (VARs), as in Bloom (2009) and Bachmann, Elstner, and Sims (2013), or using instrumental variables, as in Carlsson (2007), and finds that increases in uncertainty typically slow down economic activity. It also relates to the uncertaintydriven business cycle literature that analyzes the impact of uncertainty through real option effects as in Bloom (2009), Bloom et al. (2012), Bachmann and Bayer (2013) and Schaal (2015), or through financial frictions as in Arellano, Bai, and Kehoe (2012) and Gilchrist, Sim, and Zakrajšek (2014).³

Our analysis also relates to a theoretical literature in macroeconomics that studies environments with learning from market outcomes, such as Veldkamp (2005), Ordoñez (2009), and Amador and Weill (2010). Closely related to our article is the analysis of Van Nieuwerburgh and Veldkamp (2006). They focus on explaining business-cycle asymmetries in an RBC model with incomplete information in which agents receive signals with procyclical precision about the economy's fundamental. During recessions, agents discount new information more heavily and the mean of their beliefs recovers slowly. Their paper provides a theory of endogenous pessimism that can explain business cycle asymmetries. Our model introduces a similar learning environment in a model of irreversible investment under uncertainty in the spirit of Dixit and Pindyck (1994) and Stokey (2008). The resulting feedback loop between endogenous uncertainty and real option effects, specific to our approach, offers a novel propagation mechanism that can lead to persistent episodes of high uncertainty and low economic activity.

The interaction of endogenous uncertainty and real option effects in our model is also reminiscent of the literature on learning and strategic delays as in Lang and Nakamura (1990), Rob (1991), Caplin and Leahy (1993), Chamley and Gale (1994), Zeira (1994), and Chamley (2004). This article differs from this literature in its attempt to evaluate and quantify the role of uncertainty and delays in a standard business cycle framework. In a recent paper in which learning and economic activity interact, Straub and Ulbricht (2015) propose a theory of endogenous uncertainty in which financial constraints impede learning about firm-level fundamentals. Financial crises cause uncertainty to rise, leading to a further tightening of financial constraints that amplifies and propagates recessions.⁴ In another recent paper considering the role of learning during the Great Recession, Kozlowski, Veldkamp, and Venkateswaran (2017) suggest that the Great Recession was

^{3.} Another literature studying time-varying risk is the literature on rare disasters (Barro 2006) and time-varying disaster risk as in Gabaix (2012) and Gourio (2012) and surveyed in Barro and Ursúa (2012).

^{4.} Some recent papers discuss alternative channels that give rise to endogenous volatility over the business cycle. See Bachmann and Moscarini (2011) and Decker and D'Erasmo (2016).

the result of an unlikely shock that caused agents to substantially revise their beliefs about the probability of lower-tail events. They find that the resulting increase in pessimism may account for part of the long-lasting downturn.

This study is also related to the literature on fads and herding in the tradition of Banerjee (1992) and Bikhchandani, Hirshleifer, and Welch (1992). Articles in that tradition consider economies with an unknown fixed fundamental and study a one-shot evolution toward a stable state, whereas we study the full cyclical dynamics of an economy that fluctuates between regimes.

The dynamics generated by the model, with endogenous fluctuations between regimes, is reminiscent of the literature on static coordination games such as Cooper and John (1988), Morris and Shin (1998, 1999), and the dynamic coordination games literature as Angeletos, Hellwig, and Pavan (2007) and Chamley (1999). These publications study games in which a complementarity in payoffs leads to multiple equilibria under complete information. The introduction of strategic uncertainty through noisy observation of the fundamental leads to a departure from common knowledge that eliminates the multiplicity. In contrast, the complete-information version of our model does not feature multiplicity, and complementarity only arises under incomplete information through social learning. Uniqueness is obtained not through strategic uncertainty but by limiting the strength of the complementarities.

I.B. Bayesian Uncertainty and the Business Cycle

Throughout the article, we adopt the concept of Bayesian uncertainty: in our theory, all agents have access to the same information \mathcal{I}_t at time t and use Bayes' rule to form beliefs about the fundamental of the economy θ_t , which is, in our context, the aggregate productivity process. We define uncertainty as the variance $\operatorname{Var}(\theta_t \mid \mathcal{I}_t)$ of the probability distribution that describes these common beliefs. In contrast, the uncertainty-driven business cycle literature that we referenced above defines uncertainty as time-varying volatility in exogenous aggregate or idiosyncratic variables.

These definitions of uncertainty are related, but they are not identical. They are related because time-varying volatility may generate uncertainty about the future fundamentals of the economy, giving rise to Bayesian uncertainty. However, they are



Various Measures of Subjective Uncertainty

Panel A: The CBOE's VXO series is a measure of market expectations of stock market volatility over the next 30 days constructed from S&P100 option prices. We present monthly averages of the series over 1986–2014. Panel B: Jurado, Ludvigson, and Ng (2015) estimate a large-scale structural model with time-varying volatility on the U.S. economy and use it to compute an implied measure of ex ante forecast error. The series we present corresponds to the H12 measure, that is, an equal-weighted average of the 12-month ahead standard deviations over 132 macroeconomic series. Panel C: The SPF series is the standard deviation of the "mean probability forecast": an average of the probability distribution provided over forecasters, of one-year-ahead output growth in percentage terms. Panel D: The Michigan Survey series correspond to the percent fraction of all respondents that reply "uncertain future" to the question why people are not buying large household items. Shaded areas correspond to NBER recessions.

different because Bayesian uncertainty may fluctuate without the presence of time-varying volatility. In our model, the variance of beliefs varies over time through learning, whereas the volatility of exogenous variables is constant.

A basic and well-known feature of the data that motivates our theory is that uncertainty increases during recessions. Instead of direct measures of time-varying volatility, we present in Figure I the evolution of four measures that capture our notion of Bayesian uncertainty to the extent that they reflect uncertainty in subjective beliefs.⁵ Panel A shows the VXO, a measure of stock market volatility as perceived by market participants: Panel B shows the uncertainty measure proposed by Jurado, Ludvigson, and Ng (2015), which captures a Bayesian notion of ex ante forecast error in a statistical model of the macroeconomy: Panel C shows the standard deviation of the average perceived distribution of output growth from the SPF: and Panel D shows the fraction of respondents who answer "uncertain future" as a reason for why it is a bad time to buy major household goods from the Michigan Survey of Consumers. Although these series attempt to measure distinct objects, they all capture the notion of subjective uncertainty. These measures support the key implication of our mechanism, that uncertainty rises during recessions. In the quantitative section of the article, we use the SPF measure to calibrate and evaluate the performance of our model because it has a natural counterpart in our framework

II. BASELINE MODEL

We begin by presenting a stylized model that only features the necessary ingredients to generate uncertainty traps. The intuitions from this simple model as well as the laws of motion governing the dynamics of uncertainty carry through to the extended model that we use for numerical analysis.

II.A. Population and Technology

Time is discrete. There are a fixed number of firms \overline{N} , chosen large enough that firms behave atomistically. Each firm $j \in \{1, \ldots, \overline{N}\}$ holds a single investment opportunity that produces output θ , common to all firms. We refer to θ as the economy's

5. The uncertainty-driven business cycle literature measures aggregate uncertainty by the conditional heteroskedasticity of various aggregates such as TFP (Bloom et al. 2012). Time-varying volatility in idiosyncratic variables is typically proxied by cross-sectional dispersions in sales growth rates (Bloom 2009), output and productivity (Kehrig 2011), prices (Vavra 2013), employment growth (Bachmann and Bayer 2014), or business forecasts (Bachmann, Elstner, and Sims 2013). All these measures have been shown to be countercyclical. Since all agents have the same beliefs about θ , these cross-sectional measures are uninformative about uncertainty in our model.

fundamental and assume that it follows the autoregressive process

(1)
$$\theta' = \rho_{\theta}\theta + \varepsilon^{\theta}, \ \varepsilon^{\theta} \sim \text{i.i.d. } \mathcal{N}\left(0, \left(1 - \rho_{\theta}^{2}\right)\sigma_{\theta}^{2}\right),$$

where $0 < \rho_{\theta} < 1$ is the persistence of the process and σ_{θ}^2 the variance of its ergodic distribution. To produce, a firm must pay a fixed cost f, drawn each period from the continuous cumulative distribution F with mean μ_f and standard deviation σ_f . Once production has taken place, the firm exits the economy and is immediately replaced by a new firm holding an investment opportunity. This assumption ensures that the mass of firms in the economy remains constant.⁶

Upon investment, the firm receives the payoff θ . Firms have constant absolute risk aversion,⁷

$$u(\theta) = \frac{1}{a} \left(1 - e^{-a\theta} \right),$$

where a > 0 is the coefficient of absolute risk aversion.

II.B. Timing and Information

Firms do not know the true value of the fundamental θ and decide whether to invest based on their beliefs. As time unfolds, they learn about θ in various ways. First, they learn from a public signal Z with precision $\gamma_z > 0$ observed at the end of each period,

(2)
$$Z = \theta + \varepsilon^{Z}, \ \varepsilon^{Z} \sim \text{i.i.d. } \mathcal{N}(0, \gamma_{z}^{-1}).$$

This signal captures the information released by statistical agencies or the media. Second, agents acquire information through social learning. When firm j invests, a noisy signal about its return, $x_j = \theta + \varepsilon_j^x$, is sent to all firms.⁸ The noise ε_j^x is

6. This assumption is made for tractability and is relaxed in the quantitative section.

7. Here, agents can be thought of as entrepreneurs with risk-averse preferences. In our quantitative model, firms use the representative household's stochastic discount factor.

8. Social learning captures the idea that firms learn from each other about various common components that affect their revenues such as productivity, demand, or regulations. Social learning has been found to influence economic decisions in



Timing of Events

normally distributed with precision $\frac{\gamma_{\kappa}}{\overline{N}} > 0$, independent over time and across investors, but common to all observers.⁹ We denote by $N \in \{0, \ldots, \overline{N}\}$ the endogenous number of firms that invest and $n = \frac{N}{\overline{N}}$ the fraction of investing firms. Because of the normality assumption, a sufficient statistic for the information provided by investing firms is the public signal

(3)
$$X \equiv \frac{1}{N} \sum_{j \in I} x_j = \theta + \varepsilon_N^X,$$

where I is the set of such firms, and

$$arepsilon_N^X \equiv rac{1}{N}\sum_{j\in I}arepsilon_j^x\sim \mathcal{N}\left(0, (n\gamma_x)^{-1}
ight).$$

Importantly, the precision $n\gamma_x$ of this signal increases with the fraction of investing firms *n*.

The timing of events is summarized in Figure II.

II.C. Beliefs

Under the assumption of a common initial prior, and because all information is public, beliefs are common across firms. In particular, there is no cross-sectional dispersion in beliefs. The

various contexts. Foster and Rosenzweig (1995) estimate a model of the adoption of high-yielding seeds in India and find it consistent with social learning. Guiso and Schivardi (2007) find that peer-learning effects matter for the behavior of Italian industrial firms. Bikhchandani, Hirshleifer, and Welch (1998) survey the empirical social learning literature.

^{9.} We assume that the precision of each individual signal x_j is inversely proportional to \overline{N} to prevent the signals to be fully revealing when we take the limit $\overline{N} \to \infty$, while preserving the positive relationship between economic activity and the amount of information. This captures the idea that uncertainty may subsist even when the number of firms is large, either because their information is correlated and arises from the same sources or because large economies are more complex and subject to more shocks, preventing the learning problem from becoming trivial.

normality assumptions about the signals and the fundamental imply that beliefs are also normally distributed

$$\theta \mid \mathcal{I} \sim \mathcal{N}\left(\mu, \gamma^{-1}\right),$$

where \mathcal{I} is the information set at the beginning of the period. The mean of the distribution μ captures the optimism of agents about the state of the economy, while γ represents the precision of their beliefs about the fundamental. Precision γ is inversely related to the amount of uncertainty. As γ increases, the variance of beliefs decreases and uncertainty declines.

Firms start each period with beliefs (μ, γ) and use all the information available to update their beliefs. By the end of the period, they have observed the public signals X and Z. Therefore, using Bayes' rule, the beliefs about next period's fundamental θ' are normally distributed with mean and precision equal to

(4)
$$\mu' = \rho_{\theta} \frac{\gamma \mu + \gamma_z Z + n \gamma_x X}{\gamma + \gamma_z + n \gamma_x},$$

(5)
$$\gamma' = \left(\frac{\rho_{\theta}^2}{\gamma + \gamma_z + n\gamma_x} + \left(1 - \rho_{\theta}^2\right)\sigma_{\theta}^2\right)^{-1} \equiv \Gamma(n,\gamma).$$

These standard updating rules have straightforward interpretations: the mean of future beliefs μ' is a precision-weighted average of the present belief μ and the new signals, X and Z, whereas γ' depends on the precision of current beliefs, the precision of the signals, and the variance of the shock to θ . Importantly, the precision of future beliefs depends not on the realization of the public signals but only on n and γ . The higher is n, the more precise is the public signal X, and the lower is uncertainty in the next period. We define $\Gamma(n, \gamma)$ in equation (5) as the law of motion of the precision of information.

II.D. Firm Problem

We now describe the problem of a firm. In each period, given its individual fixed cost f and the common beliefs about the fundamental, a firm can either wait or invest. It solves the Bellman equation

(6)
$$V(\mu,\gamma,f) = \max\left\{V^{W}(\mu,\gamma), V^{I}(\mu,\gamma) - f\right\},$$

where $V^{W}(\mu, \gamma)$ is the value of waiting and $V^{I}(\mu, \gamma)$ is the value of investing after incurring the investment cost *f*.

If a firm waits, it starts the next period with updated beliefs (μ', γ') about the fundamental and a new draw of the fixed cost f'. Therefore, the value of waiting is

(7)
$$V^{W}(\mu,\gamma) = \beta \mathbb{E}\left[\int V\left(\mu',\gamma',f'\right) dF\left(f'\right) \mid \mu,\gamma\right].$$

In turn, upon investment, a firm receives output θ and exits the economy. Therefore,

(8)
$$V^{I}(\mu,\gamma) = \mathbb{E}\left[u(\theta) \mid \mu,\gamma\right] = \frac{1}{a} \left(1 - e^{-a\mu + \frac{a^{2}}{2\gamma}}\right).$$

The firm's optimal investment decision takes the form of a cutoff rule $f^c(\mu, \gamma)$ such that a firm invests if and only if $f \leq f^c(\mu, \gamma)$. The cutoff is defined by the following indifference condition

(9)
$$f^{c}(\mu,\gamma) = V^{I}(\mu,\gamma) - V^{W}(\mu,\gamma).$$

II.E. Law of Motion for the Number of Investing Firms N

We now aggregate the individual decisions of the firms. As the investment decision follows the cutoff rule $f^c(\mu, \gamma)$, the process for the number of investing firms N satisfies

(10)
$$N\left(\mu,\gamma,\left\{f_{j}\right\}_{1\leqslant j\leqslant \overline{N}}\right) = \sum_{j=1}^{\overline{N}} \mathrm{I\!I}\left(f_{j}\leqslant f^{c}\left(\mu,\gamma\right)\right).$$

Because investment depends on a random fixed cost, the number of investing firms is a random variable that depends on the realization of the shocks $\{f_j\}_{1 \leq j \leq \overline{N}}$. As these costs are i.i.d., the ex ante probability of investment is identical across firms and equal to $F(f^c(\mu, \gamma))$. Therefore, the ex ante distribution of N, as perceived by firms, is binomial,

(11)
$$N \mid \mu, \gamma \sim \operatorname{Bin}\left(\overline{N}, F(f^{c}(\mu, \gamma))\right).$$

Note that N is only a function of the beliefs (μ, γ) and the individual shocks $\{f_j\}_{1 \le j \le \overline{N}}$. Since these shocks are independent from the fundamental θ and the investment decisions are made

Downloaded from https://academic.oup.com/qje/article-abstract/132/4/1641/3857745 by UNIVERSITAT POMPEU FABRA user on 04 December 2017 before the observation of returns, there is nothing to learn from the noninvestment of firms, nor from the realization of N itself.

II.F. Recursive Competitive Equilibrium

Focusing on the limiting case when $\overline{N} \to \infty$, the fraction of investing firms *n* becomes deterministic,

$$n = \frac{N}{\overline{N}} \xrightarrow{a.s} F(f^c(\mu, \gamma)).$$

We define a recursive competitive equilibrium as follows.¹⁰

- DEFINITION 1. A recursive competitive equilibrium consists of a cutoff rule $f^c(\mu, \gamma)$, value functions $V(\mu, \gamma, f)$, $V^W(\mu, \gamma)$, $V^I(\mu, \gamma)$, laws of motions for aggregate beliefs $\{\mu', \gamma'\}$, and a fraction of investing firms $n(\mu, \gamma)$, such that
 - i. The value function $V(\mu, \gamma, f)$ solves equation (6), with $V^{W}(\mu, \gamma)$ and $V^{I}(\mu, \gamma)$ defined according to equations (7) and (8), yielding the cutoff rule $f^{c}(\mu, \gamma)$ in equation (9);
 - ii. The aggregate beliefs (μ, γ) evolve according to equations (4) and (5), under the perceived fraction of investing firms $n(\mu, \gamma) = F(f^c(\mu, \gamma))$.

III. EQUILIBRIUM CHARACTERIZATION

We first characterize the evolution of beliefs. We then show the existence and uniqueness of an equilibrium and provide conditions under which firms are less likely to invest when uncertainty is high.

III.A. Evolution of Beliefs

The optimal investment rule $f^{c}(\mu, \gamma)$ depends on how beliefs evolve. We begin by establishing two simple lemmas about the dynamics of aggregate beliefs.

1. Evolution of the Mean of Beliefs. Using equation (4), we can characterize the stochastic process for the mean of beliefs as follows.

10. Fluctuations in N due to finite sampling are irrelevant for our purpose. Our results nonetheless carry over to the finite \overline{N} case. Our existence and uniqueness proof for that case is available on request.

LEMMA 1. For a given *n*, the mean of beliefs μ follows an autoregressive process with time-varying volatility *s*,

(12)
$$\mu' = \rho_{\theta} \mu + s(n, \gamma) \varepsilon,$$

where $s(n, \gamma) = \rho_{\theta} \left(\frac{1}{\gamma} - \frac{1}{\gamma + \gamma_{z} + n\gamma_{x}} \right)^{\frac{1}{2}}$ and $\varepsilon \sim \mathcal{N}(0, 1)$.

The mean of beliefs captures the optimism of agents about the fundamental and evolves stochastically due to the the arrival of new information. It inherits the autoregressive property of the fundamental, and its volatility $s(n, \gamma)$ is time-varying because the amount of information that firms collect over time is endogenous. The volatility is decreasing with γ and increasing with n. In times of low uncertainty (γ high) agents place more weight on their current information and less on new signals, making the mean of beliefs more stable. In contrast, in times of high activity (nhigh) more information is released, making beliefs more likely to fluctuate.

2. Evolution of Uncertainty. The precision of beliefs γ reflects the inverse of uncertainty about the fundamental and its dynamics play a key role for the existence of uncertainty traps. Its law of motion satisfies the following properties.

LEMMA 2. The law of motion $\Gamma(n, \gamma)$ increases with n and γ . For a given fraction of investing firms n, the law of motion for the precision of beliefs $\gamma' = \Gamma(n, \gamma)$ admits a unique stable stationary point in γ .

The thin solid curves in Figure III depict $\Gamma(n, \gamma)$ for different constant values of n. An increase in the level of activity raises the next period precision of information γ' for each level of γ in the current period. Since n is between 0 and 1, the support of the ergodic distribution of γ must lie between the bounds $\underline{\gamma}$ and $\overline{\gamma}$ defined by $\underline{\gamma} \equiv \Gamma(0, \underline{\gamma})$ and $\overline{\gamma} \equiv \Gamma(1, \overline{\gamma})$. In other words, $\underline{\gamma}$ is the stationary level of precision when no firm invests, while $\overline{\overline{\gamma}}$ is the one when all firms invest.

In equilibrium, *n* varies with μ and γ . Suppose, as an example, that *n* is an increasing step function of γ that takes the values 0, 0.5, or 1, and let us keep μ fixed for the moment. Figure III illustrates how the feedback from uncertainty to investment opens up the possibility of multiple stationary points in the dynamics of the precision of beliefs, and therefore uncertainty. In this example, the function $\gamma' = \Gamma(n(\mu, \gamma), \gamma)$, depicted by the thick solid curve,



FIGURE III

Example of Dynamics for Beliefs Precision γ

The thick continuous line shows the law of motion of uncertainty γ for an arbitrary policy $n(\mu, \gamma)$ that equals 0 for $\gamma \in [0, 0.45)$, 0.5 for $\gamma \in [0.45, 0.65)$, and 1 for $\gamma \in [0.65, 1]$.

has three fixed points. We formally establish in Section IV that this type of multiplicity can happen in equilibrium.

III.B. Existence and Uniqueness

We have described in Lemmas 1 and 2 how beliefs depend on the fraction of investing firms. We now characterize the equilibrium decision rule and provide existence and uniqueness conditions.

PROPOSITION 1. Under Assumptions 1–3, found in Online Appendix

D, and for γ_x sufficiently small, the equilibrium exists and is unique. Under some additional conditions satisfied when γ_x is small and risk aversion *a* is large enough, the equilibrium cutoff f^c is increasing in μ and γ .

This proposition establishes the monotonicity of the equilibrium cutoff rule. Anticipating higher returns, a more optimistic firm (higher μ) is more likely to invest. In turn, uncertainty

(lower γ) reduces the incentives to invest for two reasons. First, risk-averse firms dislike uncertain payoffs. Second, since investment is costly and irreversible, there is an option value of waiting: in the face of uncertainty, firms prefer to delay investment to gather additional information and avoid downside risk.

It is essential for our mechanism that uncertainty discourages investment, a feature typical of optimal stopping time models of investment. Assumption 3, satisfied if the persistence of the fundamental is high enough and its volatility is sufficiently low, ensures that the fundamental does not vary too much over time, so that firms have an incentive to wait to collect more information.¹¹ This condition alone, however, is not sufficient in our context. The monotonicity of the cutoff f^c in γ requires additional restrictions because of the endogeneity of beliefs, which gives rise to ambiguous feedback effects. For instance, the variations in *n* implied by fluctuations in μ or γ affect the volatility of next period's mean beliefs μ' (Lemma 1). This in turn can have ambiguous effects on firms' current incentives to invest. To ensure that the first-order effects of risk aversion and option value dominate, we must bound these feedback effects. Since they operate solely through social learning, we can do so by imposing an upper bound on the informativeness of this channel, γ_x . When γ_x is small enough, the equilibrium cutoff is guaranteed to be increasing in μ and γ , as one would expect in the absence of social learning.

Establishing the existence of an equilibrium is relatively straightforward, because the problem is continuous¹² and general fixed point theorems apply.¹³ Showing uniqueness is more

11. The law of motion equation (5) highlights the importance of the persistence of the fundamental ρ_{θ} for the dynamics of uncertainty. As ρ_{θ} declines, past observations contain less information about the current value of the fundamental and learning therefore becomes less relevant. At a result, the option value of waiting becomes smaller and the conditions for uncertainty traps to exist, provided in the next section, are less likely to be satisfied.

12. A key assumption for continuity to hold is Assumption 1, which imposes some regularity conditions on the fixed cost distribution F. In particular, it guarantees that the distribution is sufficiently smooth that the pattern of entry does not vary too much.

13. The contraction mapping theorem in the space of bounded continuous functions requires us to bound the belief process. The existence is established under Assumption 2 for a truncated version of equation (12), where the bounds are chosen sufficiently large to provide an arbitrarily good approximation to the true process.

QUARTERLY JOURNAL OF ECONOMICS



FIGURE IV Fraction of Investing Firms $n(\mu, \gamma)$

challenging because our economy features complementarities in information: the more firms invest, the more uncertainty declines, encouraging further investment. If these complementarities are strong enough, they can lead to multiple equilibria. To prevent this, we again use the insight that the magnitude of this feedback is governed by the precision γ_x of the social learning channel. We show in particular that the main fixed point problem that characterizes the optimal cutoff rule is a contraction, and it is therefore unique, when γ_x is small.¹⁴ The uniqueness of the equilibrium is an attractive feature, as it leads to unambiguous predictions and makes the model amenable to quantitative work. Despite the uniqueness of the equilibrium, the model features interesting nonlinear dynamics and multiple stationary points, as we show in Section IV.

Figure IV illustrates how the investment probability varies as a function of beliefs (μ, γ) when monotonicity is obtained. The fraction of investing firms increases as they are more optimistic $(\mu$ high) or less uncertain $(\gamma$ high) about the fundamental.

14. We show that the mapping that characterizes the optimal cutoff rule is a contraction in the space of Lipschitz continuous functions for some given moduli, which allows us to put a bound on these feedback effects. We cannot rule out the existence of equilibrium cutoffs that do not satisfy this property. We can, however, explicitly rule them out in the case of the planner's allocation, where Lipschitz continuity is necessarily satisfied.

IV. UNCERTAINTY TRAPS

We now examine the interaction between firms' behavior in the face of uncertainty and social learning. This interaction leads to episodes of self-sustaining uncertainty and low activity, which we call uncertainty traps. We provide sufficient conditions on the parameters that guarantee the existence of such traps and discuss the type of aggregate dynamics they imply. We find that the response of the economy to shocks is highly nonlinear: it quickly recovers after small shocks, but large, short-lived shocks may plunge the economy into long-lasting recessions. We also characterize the constrained planner's problem and discuss its policy implications.

IV.A. Definition and Existence

We define *uncertainty traps* as the coexistence of multiple stationary points in the dynamics of belief precision—a situation similar to the one depicted in Figure III.

DEFINITION 2. There is an *uncertainty trap* if there exists an interval (μ_l, μ_h) such that for every $\mu \in (\mu_l, \mu_h)$, there are at least two locally stable fixed points in the dynamics of the precision of beliefs $\gamma' = \Gamma(n(\mu, \gamma), \gamma)$.

We refer to these multiple stationary points as *regimes*. Note that multiplicity of regimes does not imply multiplicity of equilibria. This distinction is important because it highlights that the model is not subject to indeterminacy. Although multiple values of γ may satisfy the equation $\gamma = \Gamma(n(\mu, \gamma), \gamma)$ for a given μ , the regime that prevails at any given time is unambiguously determined by the history of past aggregate shocks, summarized by the current beliefs (μ, γ) . The definition of uncertainty traps also emphasizes the notion of stability, which is required for the type of self-sustaining dynamics that we describe. Notice, however, that we only require local stability along the dimension γ while μ keeps evolving according to its law of motion.

The following proposition formally establishes that uncertainty traps exist for a range of the mean of beliefs μ under some conditions on the dispersion of investment costs.

PROPOSITION 2. Under the conditions of Proposition 1 and one additional condition satisfied for σ_f small enough or risk aversion a high enough, the economy features an uncertainty trap with at least two regimes $\gamma_l(\mu) < \gamma_h(\mu)$ for $\mu \in (\mu_l, \mu_h)$. Regime



FIGURE V

Dynamics of the Precision of Beliefs $\gamma' = \Gamma(n(\mu, \gamma), \gamma)$ for Different Values of μ

 γ_l is characterized by high uncertainty and low investment, while regime γ_h is characterized by low uncertainty and high investment.

Figure V presents examples of the law of motion of γ when the investment costs f are normally distributed. The solid curves represent the function $\gamma' = \Gamma(n(\mu, \gamma), \gamma)$ evaluated at five different values of μ , with the thick solid curve corresponding to an intermediate value of μ . In all cases, for small γ uncertainty is high and firms do not invest. As a result, they do not learn from observing economic activity and the precision of beliefs γ' remains low. As the precision γ increases, uncertainty decreases and firms become sufficiently confident about the fundamental to start investing. As that happens, uncertainty decreases further.

In our example, the thick curve intersects the 45° line three times. The second intersection corresponds to an unstable regime, but the other two are locally stable. We denote these regimes by γ_l and γ_h . In regime γ_l , uncertainty is high and investment is low, while the opposite is true in regime γ_h .

Proposition 2 shows that this situation is a general feature of the equilibrium when the dispersion of investment costs σ_f is

small. This condition ensures that the feedback of investment on information is strong enough to sustain distinct stationary points.

IV.B. Dynamics: Nonlinearity and Persistence

We now describe the full dynamics of the economy by taking into account the evolution of μ in response to the arrival of new information. Figure V shows that as long as μ stays between the values μ_l and μ_h , defined in Proposition 2, the two regimes $\gamma_l(\mu)$ and $\gamma_h(\mu)$ preserve their stability. As a result, uncertainty and the fraction of active firms *n* are relatively unaffected by changes in μ . In contrast, for values of μ above μ_h , a large enough fraction of firms invest, so the dynamics of beliefs only admits the highactivity regime as a stationary point. Similarly, for values below μ_l , the economy only admits the low-activity regime. Therefore, sufficiently large shocks to μ can make one regime disappear and trigger a regime switch.

The economy displays nonlinear dynamics: it reacts very differently to large shocks in comparison to small ones. Figure VI shows various simulations to illustrate this feature using the example from Figure V. The top panel presents three different series of shocks to the mean of beliefs μ . The three series start from the high-activity/low-uncertainty regime. At t = 5, the economy is hit by a negative shock to μ , due to a bad realization of either the public signals or the fundamental. The mean of beliefs then returns to its initial value at t = 10. Across the three series, the magnitude of the shock is different.

The middle and bottom panels show the response of beliefs precision γ and the fraction of investing firms *n*. The solid gray line represents a small temporary shock, such that μ remains within (μ_l , μ_h). Despite the negative shocks to the mean of beliefs, all firms keep investing and the precision of beliefs is unaffected. When the economy is hit by a temporary shock of medium size (dashed line), some firms stop investing, leading to a gradual increase in uncertainty. As uncertainty rises, investment falls further and the economy starts to drift toward the low regime. However, when the mean of beliefs recovers, the precision of information and the number of active firms quickly return to the high-activity regime. In contrast, when the economy is hit by a large temporary shock (dotted line), the number of firms delaying investment is large enough to produce a self-sustaining



Persistent Effects of Temporary Shocks

The continuous line displays the response to a small negative transitory shock to the mean of beliefs μ , the dashed line to a medium-sized shock and the dotted line to a large shock.

increase in uncertainty. The economy quickly shifts to the lowactivity regime and remains there even after the mean of beliefs recovers.

We now discuss how the economy escapes from the trap in which it fell in Figure VI. Figure VII shows the effect of positive shocks when the economy starts from the low regime. The economy receives positive signals that lead to a temporary increase in mean beliefs between periods 20 and 25, possibly because of a recovery in the fundamental. When the temporary increase in average beliefs is not sufficiently strong, the recovery is interrupted as μ returns to its initial value. However, when the temporary increase is sufficiently large, the economy reverts back to the high-activity regime. Once again, temporary shocks of sufficient magnitude to the fundamental may lead to nearly permanent effects on the economy.



Escaping an Uncertainty Trap

The continuous line displays the response to a large positive transitory shock to the mean of beliefs μ , the dashed line to a medium-sized shock and the dotted line to a small shock.

IV.C. Additional Remarks

A number of additional lessons can be drawn from these simulations. First, in our framework uncertainty is a by-product of recessions. This result echoes the empirical findings of Bachmann, Elstner, and Sims (2013) who show that uncertainty is partly caused by recessions and conclude that it is of secondary importance for the business cycle. We show, however, that uncertainty may still have a large impact on the economy by affecting the persistence and depth of recessions, even if it is not what triggers them.

Second, as in models with learning in the spirit of Van Nieuwerburgh and Veldkamp (2006), our theory provides an explanation for asymmetries in business cycles. In good times, since agents receive a large flow of information, they react faster to shocks than they would in bad times.

Third, our economy may feature high uncertainty without volatility. For instance, in the low regime agents are highly uncertain about the fundamental but the volatility of economic aggregates is low. Therefore, according to our theory, subjective uncertainty may affect economic fluctuations even if no volatility is observed in the data. This distinguishes our approach from the existing uncertainty-driven business cycle literature in the spirit of Bloom (2009). In particular, direct measures of subjective uncertainty rather than measures of volatility are important to capture the full amount of uncertainty in the economy.

Finally, a recent body of literature (Bachmann, Elstner, and Sims 2013; Orlik and Veldkamp 2013) uses survey data to derive measures of uncertainty based on ex ante forecast errors. Our model highlights a potential shortcoming of this approach, as uncertainty about fundamentals differs from uncertainty about endogenous variables, such as output or investment. For example, when the economy is trapped in the low-activity regime, firms know that all firms are uncertain, and therefore that investment is likely to be low, so that the economy is less exposed to aggregate risk. As a result, their forecasts about economic aggregates are accurate, even though their uncertainty about the fundamental is high. As implied by the model, forecast errors about variables like output may not always be a good proxy for uncertainty about fundamentals.

IV.D. Policy Implications

The economy is subject to an information externality: in the decentralized equilibrium, firms invest less often than they should because they do not internalize the release of information to the rest of the economy caused by their investment. In Proposition 3, we solve the problem of a constrained planner subject to the same information technology as individual agents. We show that the decentralized economy is constrained inefficient, and that an investment subsidy is sufficient to restore constrained efficiency.

PROPOSITION 3. Under Assumptions 1–3, stated in the Online Appendix, the recursive competitive equilibrium is constrained inefficient. The efficient allocation can be implemented with positive investment subsidies $\tau(\mu, \gamma)$ and a uniform tax.

The subsidy that implements the optimal allocation takes a simple form to align social and private incentives. As shown in the proof of the proposition, it is simply the sum of the social value of releasing an additional signal to the economy and the private value of delaying investment.

The optimal policy being a subsidy, Proposition 3 implies that firms are more likely to invest in the efficient allocation than in the laissez-faire economy. However, uncertainty traps can still arise in the efficient allocation. Proposition 4 establishes the result.

PROPOSITION 4. Under Assumptions 1–2 and γ_x small enough, the planner's allocation is subject to an uncertainty trap for σ_f low enough and risk aversion *a* high enough.

The existence of uncertainty traps in the planner's allocation may be surprising if one thinks of the planner as a coordinator that should always prefer the high regime, as one might expect in a model with multiple equilibria. As it turns out, transitioning from one regime to the other is costly and risky. If the planner does not have more information than individual agents, it is still optimal to wait when uncertainty is high enough. Hence, there may still exist sufficiently strong feedback from beliefs to actions in the constrained-efficient allocation to generate uncertainty traps. However, although uncertainty traps remain present in the efficient allocation, they are less likely to arise than in the laissez-faire economy because firms have stronger incentives to invest.

V. QUANTITATIVE EVALUATION

To evaluate the quantitative importance of uncertainty traps, we embed the mechanism into a general equilibrium macroeconomic framework. We first describe the quantitative framework and its parameterization. We then compare the model's implications for standard business cycle moments with the data and two alternative models: the RBC model and a restricted version of our model in which uncertainty about the fundamental is fixed over time. Finally, we present our main quantitative exercise, in which we compare the behavior of economic aggregates in the data against the predictions from our model and alternative models in the context of the Great Recession.

V.A. Quantitative Model with Uncertainty Traps

1666

We extend the baseline model along several dimensions. First, firms are now long-lived, use both capital and labor to produce, and accumulate capital over time. They enter the economy endogenously depending on economic conditions and exit exogenously. Second, firms are owned by a risk-averse representative household that maximizes utility over consumption and leisure. Third, factor and goods prices are endogenously determined in general equilibrium. As in the baseline model, firms must pay an irreversible fixed cost to operate and social learning takes place when a firm begins to produce. As a result, the number of entering firms responds to uncertainty about the fundamental, and uncertainty depends on economic activity.

1. Preferences and Technology. The representative household chooses consumption C_t and labor L_t to maximize the expected discounted sum of future utility

(13)
$$\mathbb{E}\sum_{t=0}^{\infty}\beta^{t}U(C_{t},L_{t}),$$

where $0 < \beta < 1$ is the discount rate. The household supplies labor in a perfectly competitive market at a wage w_t . It also owns the firms in the form of claims to their dividends.

A single good used for consumption and investment is produced by a continuum of firms of measure *m*. Each firm $j \in [0, m]$ produces the final good by operating a Cobb-Douglas technology,

$$A(1+ heta)\left(k_{j}^{lpha}l_{j}^{1-lpha}
ight)^{\omega},\ 0$$

using l_j units of labor and k_j units of capital. The parameter α controls the capital intensity. The firm-level returns to scale, or span-of-control (Lucas 1978), parameter ω is assumed to be strictly less than 1 to deliver a well-defined notion of firm size. The fundamental θ follows the AR(1) process $\theta' = \rho_{\theta}\theta + \varepsilon^{\theta}$, with $\varepsilon^{\theta} \sim \text{i.i.d. } \mathcal{N} \left(0, \left(1 - \rho_{\theta}^2 \right) \sigma_{\theta}^2 \right).^{15}$ The total mass of firms m evolves endogenously. Each period, a mass Q > 0 of potential entrants

^{15.} The additive specification of TFP, $1 + \theta$, ensures that the variance of beliefs about θ does not affect expected output directly. As in our calibration the standard deviation of the ergodic distribution of θ is much smaller than 1, productivity is always positive in our simulations.

has the option to start production, but only an (endogenous) fraction *n* of them does so. The mass *Q* remains fixed over time. Firms exit at an exogenous rate $\delta_m > 0$.

Each period, firms pay a fixed cost f common across firms and denominated in units of the final good. We assume that $f \sim \mathcal{N}\left(\mu_f, \sigma_f^2\right)$ is drawn independently over time.¹⁶ Due to the irreversibilities created by these fixed costs, fewer firms enter in times of heightened uncertainty.

2. Information and Timing. As in the baseline model, agents do not observe the true value of the fundamental θ but learn about it from two sources. First, they learn from a public signal Z. In contrast to the baseline model, where this signal captured exogenous information released by media and statistical agencies, we now explicitly model the signal Z as a summary of the information collected through the observation of certain economic aggregates. As in any model with information frictions, restrictions about what agents observe must be imposed to avoid perfectly revealing the fundamental. Agents cannot, for instance, perfectly observe output as this would reveal θ . We assume instead that agents are able to observe the value added of each firm, as well as its aggregate counterpart, but are unable to perfectly distinguish between its individual components: revenue and fixed cost.¹⁷ As a result. a high level of value added may reflect either a high value of the fundamental θ or a low value of the fixed costs.¹⁸ Second, and more specific to the channel we study in this article, agents also learn from signals emanating from others. As in the $\overline{N} \to \infty$ case of the

16. The baseline theory included an idiosyncratic component to these fixed costs, which we ignore here for simplicity. Online Appendix C.4 performs sensitivity analysis on this assumption. We assume, however, that f is subject to aggregate shocks to be consistent with our information structure, as we explain in the next subsection.

17. This assumption is in the spirit of Lucas (1972), where firms cannot distinguish between real and nominal shocks. A previous version of the paper assumed that firms cannot distinguish between aggregate and idiosyncratic productivity shocks. The benefit of our current approach is to allow for a simple aggregation of the economy.

18. Despite this restriction on the observability of gross output and θ , all other economic aggregates are observed and agents use all the available information to make their decisions. However, as the timing will make clear, observing other variables such as the wage rate w, consumption C, the aggregate capital stock K, aggregate employment L, aggregate value added, the measure of entrants n, or the measure of incumbents m does not reveal any additional information.

baseline model, the entry of an infinitesimal measure of firms dj releases a normally distributed signal x_j about θ , observed by everyone, with a precision $\gamma_x dj$, proportional to the mass of entrants. Again, the information collected through this social learning channel can be summarized by a public signal X with precision $nQ\gamma_x$.¹⁹ With all signals being public, beliefs are common across firms and the representative household.

In each period, events unfold as follows:

- i. Incumbent firms, potential entrants and the household start with the same prior distribution over the fundamental, $\theta \mid \mathcal{I} \sim \mathcal{N}(\mu, \gamma^{-1})$. The fundamental θ and the fixed cost *f* are drawn but unobserved.
- ii. The Q potential entrants decide whether to enter. A fraction n of them enters and starts producing next period.
- iii. The *m* incumbent firms choose labor and investment. The household decides how much labor to supply and the labor market clears.²⁰
- iv. Fixed costs are paid, and production takes place. All agents observe the signal Z, which captures the information contained in value added, and the signal X from new entrants and update their beliefs. A fraction δ_m of firms exogenously exits.

3. *Firm-Level Problem*. The aggregate state space of the economy is (μ, γ, K, m) where $K = \int_0^m k_j dj$ is the aggregate capital stock. Realized individual profits for a firm operating with k units of capital and l units of labor are²¹

(14)

$$\pi(k,l;\mu,\gamma,K,m,\theta,f) = A(1+\theta)k^{\alpha\omega}l^{(1-\alpha)\omega} - w(\mu,\gamma,K,m)l - f.$$

19. A formal derivation of this information aggregation result is presented later in this section.

20. To simplify the inference problem from an endogenous variable like value added and avoid, in particular, a correlation between labor and the new information, we assume that the labor supply decision occurs at the beginning of the period, ahead of the consumption decision.

21. Note that the presence of fixed operating costs could lead to negative profits. Because f is small relative to output in our calibration, this virtually never happens.

The value of an incumbent firm that has accumulated k units of capital is then

(15)

$$V^{I}(k; \mu, \gamma, K, m) = \max_{k', l} \mathbb{E} \left\{ U_{c}(C, L) \left[\pi(k, l; \mu, \gamma, K, m, \theta, f) + (1 - \delta_{K})k - k' \right] + \delta_{m}U_{c}(C, L)k' + \beta(1 - \delta_{m})V^{I}(k'; \mu', \gamma', K', m') |\mu, \gamma \right\},$$

subject to the laws of motion for the aggregate state variables $\{\mu, \gamma, K, m\}$, described in the following sections. The parameter $0 < \delta_K < 1$ is the depreciation rate of capital. This firm chooses labor l and the next period capital k' to maximize the expected sum of profits, discounted by the marginal utility $U_c(C, L)$, which plays the role of the stochastic discount factor. When a firm exits, which happens with probability δ_m , its accumulated capital is scrapped and returned to the household at the end of the period.

Consider now the problem of a potential entrant. In each period, a potential entrant decides between waiting and entering. Its value is

(16)
$$V(\mu, \gamma, K, m) = \max \left\{ V^{W}(\mu, \gamma, K, m), V^{E}(\mu, \gamma, K, m) \right\}.$$

If a potential entrant waits, it preserves the option of entering next period; hence the value of waiting is

(17)
$$V^{W}(\mu, \gamma, K, m) = \beta \mathbb{E} \left[V \left(\mu', \gamma', K', m' \right) | \mu, \gamma \right].$$

If, instead, the potential entrant decides to enter, its value is

(18)
$$V^{E}(\mu, \gamma, K, m) = \max_{k'_{e}} (1 - \delta_{m}) \mathbb{E} \Big[-U_{c}(C, L) k'_{e} + \beta V^{I}(k'_{e}; \mu', \gamma', K', m') | \mu, \gamma \Big].$$

The definition of V^E indicates that a potential entrant chooses the amount of capital k'_e , carried into the next period if it survives the δ_m shock or returned to the household at the end of the period otherwise. Upon entry, its value next period is equal to the value of an incumbent firm that has accumulated k'_e units of capital, $V^I(k'_e; \mu', \gamma', K', m')$.

4. Aggregates. Incumbents and entrants face the same investment problem and choose the same next-period capital level $k'(\mu, \gamma, K, m)$. Therefore, the next period's aggregate capital stock is

(19)
$$K'(\mu, \gamma, K, m) = m'(\mu, \gamma, K, m)k'(\mu, \gamma, K, m),$$

where $m'(\mu, \gamma, K, m)$ is the mass of incumbents next period, given by

(20)
$$m'(\mu, \gamma, K, m) = (1 - \delta_m)(m + n(\mu, \gamma, K, m)Q).$$

The fraction of entering firms among the Q potential entrants $n(\mu, \gamma, K, m)$ must be consistent with individual entry decisions, in the sense that (21)

$$n(\mu, \gamma, K, m) = \begin{cases} 1 & \text{if } V^E(\mu, \gamma, K, m) > V^W(\mu, \gamma, K, m) \\ \in \begin{bmatrix} 0, 1 \end{bmatrix} & \text{if } V^E(\mu, \gamma, K, m) = V^W(\mu, \gamma, K, m) \\ 0 & \text{if } V^E(\mu, \gamma, K, m) < V^W(\mu, \gamma, K, m) . \end{cases}$$

In turn, aggregate labor demand is

(22)
$$L(\mu, \gamma, K, m) = m \times l\left(\frac{K}{m}; \mu, \gamma, K, m\right),$$

where $l(k; \mu, \gamma, K, m)$ is the firm-level labor demand resulting from equation (15).

5. Information and Beliefs. We now characterize the information contained in the signals observed by the agents. First, as in the baseline model, the information diffused through social learning can be aggregated into a single signal X which averages the individual signals released by entrants,²²

(23)
$$X = \frac{1}{nQ} \int_0^{nQ} x_j dj = \theta + \varepsilon^X, \ \varepsilon^X \sim \mathcal{N}\left(0, (nQ\gamma_x)^{-1}\right),$$

where $nQ\gamma_x$, the endogenous precision of the social learning channel, changes with economic activity. Second, the information

22. As in the infinite \overline{N} case in the baseline model, X in expression (23) is to be understood as the distributional limit of the average of N signals with precisions $\frac{\gamma_X}{\overline{N}}$, i.e., $\lim \frac{1}{N} \sum_{1}^{N} x_j \sim \mathcal{N}\left(0, \left(\frac{\gamma_X N}{\overline{N}}\right)^{-1}\right)$ as $\overline{N} \to \infty$ and $\frac{N}{\overline{N}} \to nQ$.

conveyed by observing value added is equivalent to the information conveyed by the signal $Z \sim \mathcal{N}\left(\theta, (\gamma_z(K, L, m))^{-1}\right)$ with precision²³

$$\gamma_{z}(K, L, m) = \left[A\left(\frac{K^{\alpha}L^{(1-\alpha)}}{m}\right)^{\omega}\frac{1}{\sigma_{f}}\right]^{2}.$$

In contrast to the benchmark model, the precision γ_z of this signal now changes with economic activity—a natural implication of assuming that agents observe economic aggregates instead of the fundamental directly. In our calibrated economy, we find that fluctuations in γ_z , which depends on the stock m of incumbent firms, are considerably smaller than fluctuations in the precision of X, which depends on the flow n of incumbent firms. Therefore, the endogenous uncertainty in the economy largely evolves as a function of the X signal.²⁴

As in the baseline model, agents are fully rational and use all information available to update their beliefs according to Bayes' law. The laws of motion for the mean and the precision of beliefs are

(24)
$$\mu' = \rho_{\theta} \frac{\gamma \mu + \gamma_z Z + n Q \gamma_x X}{\gamma + \gamma_z + n Q \gamma_x}$$

(25)
$$\gamma' = \left(\frac{\rho_{\theta}^2}{\gamma + \gamma_z + nQ\gamma_x} + \left(1 - \rho_{\theta}^2\right)\sigma_{\theta}^2\right)^{-1}.$$

6. *Recursive Competitive Equilibrium*. We are now ready to define a competitive equilibrium for this economy.

DEFINITION 3. A recursive competitive equilibrium is a collection of value functions $V(\mu, \gamma, K, m)$, $V^{W}(\mu, \gamma, K, m)$, $V^{E}(\mu, \gamma, K, m)$, and $V^{I}(k; \mu, \gamma, K, m)$; individual policy functions

23. Since all incumbent firms are identical, individual value added is $A(1+\theta)\left(\frac{K}{m}\right)^{\alpha\omega}\left(\frac{L}{m}\right)^{(1-\alpha)\omega} - f$. Since K, L, m, and the distribution of f are known, observing value added is equivalent to observing $\theta - \frac{m^{\omega}}{AK^{\alpha\omega}L^{(1-\alpha)\omega}}(f-\mu_f) \sim \mathcal{N}\left(\theta, \left[Am^{-\omega}K^{\alpha\omega}L^{(1-\alpha)\omega}\right]^{-2}\sigma_f^2\right)$.

24. In our calibrated economy, fluctuations in γ_z only accounts for 2.7% of the total fluctuation in uncertainty while social learning through *X* accounts for the rest.

 $k'(\mu, \gamma, K, m), k'_e(\mu, \gamma, K, m)$, and $l(k; \mu, \gamma, K, m)$; aggregate policy functions $K'(\mu, \gamma, K, m), m'(\mu, \gamma, K, m), \mu'(\mu, \gamma, K, m),$ $\gamma'(\mu, \gamma, K, m), n(\mu, \gamma, K, m), L(\mu, \gamma, K, m)$, and $C(\mu, \gamma, K, m, \theta, f)$; and wages $w(\mu, \gamma, K, m)$ such that

- i. The value functions V(μ, γ, K, m), V^W(μ, γ, K, m), V^E(μ, γ, K, m), and V^I(k; μ, γ, K, m), and the associated policy functions k'(μ, γ, K, m), k'_e(μ, γ, K, m), and l(k; μ, γ, K, m), solve the Bellman equations (15)–(18) under the entry schedule n(μ, γ, K, m), and the laws of motion K'(μ, γ, K, m), m'(μ, γ, K, m), μ'(μ, γ, K, m), and γ'(μ, γ, K, m) given by equations (19), (20), (24), and (25);
- ii. The fraction of entering firms $n(\mu, \gamma, K, m)$ satisfies the consistency equation (21);
- iii. The policy functions $L(\mu, \gamma, K, m)$ and $C(\mu, \gamma, K, m, \theta, f)$ solve the household's first-order condition on labor supply.²⁵

(26)

$$\frac{\mathbb{E}\left[U_{L}(C(\mu,\gamma,K,m,\theta,f),L(\mu,\gamma,K,m))\right]}{\mathbb{E}\left[U_{C}(C(\mu,\gamma,K,m,\theta,f),L(\mu,\gamma,K,m))\right]} = w(\mu,\gamma,K,m);$$

iv. The aggregate resource constraint is satisfied:

$$\begin{split} C\left(\mu,\gamma,K,m,\theta,f\right) + K'\left(\mu,\gamma,K,m\right) - \left(1-\delta_K\right)K + mf \\ &= A(1+\theta)m^{1-\omega}\left(K^{\alpha}L(\mu,\gamma,K,m)^{1-\alpha}\right)^{\omega}. \end{split}$$

V.B. Calibration

1. Standard Parameters. The time period is one quarter. Most of the moments that we target are computed starting in 1978:Q1, when the Longitudinal Business Database (LBD) that we use for firm-level moments begins, and stopping in 2007:Q3, at the onset of the 2007–2009 recession, allowing us to evaluate the out-of-sample properties of the model in our Great Recession exercise. In our benchmark specification, we assume GHH preferences, $U = \log \left(C - \frac{L^{1+\nu}}{1+\nu}\right)$, and we report the results of our main

25. Because the labor and consumption decisions do not occur simultaneously, the first-order condition that governs the labor supply decision features an expectation over consumption.

quantitative exercise under CRRA preferences in Online Appendix A as robustness.²⁶ We set the Frisch elasticity $\nu = 2.84$, which corresponds to the average aggregate Frisch elasticity of hours reported by Chetty et al. (2011). The discount rate β is chosen to match an annual value of 0.95. The depreciation rate is set to an annual value of 0.1.

For the production function parameters, we normalize A = 1 and set the returns-to-scale parameter ω to 0.89, which corresponds to the weighted average across two-digit SIC estimates of the returns to scale from Basu and Fernald (1997).²⁷ We set the capital intensity parameter α so that $(1 - \alpha)\omega = 0.645$ to match the average labor compensation over GDP from 1978–2007 according to annual data from the Penn World Tables (Feenstra, Inklaar, and Timmer 2015).²⁸

We set $\delta_m = 2.6\%$, which corresponds to the employmentweighted firm exit rate for all firms in the LBD between 1978 and 2007. The mass of potential entrants Q is normalized to 1.²⁹

The parameters $\{\rho_{\theta}, \sigma_{\theta}^2\}$ of the fundamental process θ are estimated using the quarterly utilization-adjusted TFP series from Fernald (2014) over 1978:Q1–2007:Q3 after removing a linear trend. This yields $\sigma_{\theta} = 0.028$ and $\rho_{\theta} = 0.964$.

2. Information and Fixed-Cost Distribution. With all the above parameters calibrated, we only need to set values for the precision of individual signals γ_x and the mean and variance of the distribution of fixed costs $\{\mu_f, \sigma_f^2\}$. These parameters govern the option-value effects and the evolution of Bayesian uncertainty about TFP. To the best of our knowledge, no direct empirical measure exists for this concept of uncertainty. The variance

26. GHH preferences are common in the information frictions literature. We adopt them in our benchmark specification because the usual CRRA preferences generate a counterfactual correlation between economic activity and positive signals (see Beaudry and Portier 2006 for a VAR estimation of the impact of news shocks). On receiving a positive signal about the economy, the wealth effect on the labor supply leads to a decline in output on impact. See Jaimovich and Rebelo (2009) for a discussion of the role of preferences in the news shocks literature.

27. Our parameterization corresponds to the γ^{v} parameter for the private economy estimated using OLS from Table 2 in Basu and Fernald (1997).

28. Specifically we use the series LABSHPUSA156NRUG from the FRED database.

29. For any given Q, we can replicate the aggregate allocation, albeit with a different measure of firms m, by rescaling A, γ_x , μ_f , and σ_f .

of beliefs about θ , however, is tightly related in our model to the ex ante forecast variance about endogenous variables like output. We thus target moments of the distribution of uncertainty in output growth forecasts provided by the SPF and use this series as our main empirical proxy for uncertainty.

The SPF asks a panel of forecasters to provide the distribution of their beliefs about the growth rate of real output in percentage terms between the current year and the last, and these distributions are averaged across forecasters.³⁰ We compute the standard deviation of this averaged distribution in every year, and use moments of its time series to calibrate the model. To fit the parameters, we compute the exact same object in a long-run simulation of our model. We pick the values of $\{\gamma_x, \mu_f, \sigma_f\}$ by targeting the mean, the 5th percentile, and the 95th percentile of the empirical distribution of uncertainty over the 1992–2007 period, corresponding to the time period over which data on the distribution of real GDP growth are available in the SPF up to the beginning of the Great Recession.

These three moments are directly informative about the three parameters that we need to calibrate. The 95th percentile corresponds to periods of high uncertainty about real output growth, which in the model corresponds to periods with low firm entry when uncertainty is mostly driven by the aggregate public signal Z, as opposed to the social learning signal X. Therefore, the 95th percentile is useful to identify σ_f , which governs the informativeness of Z. Similarly, the 5th percentile corresponds to periods of low uncertainty, which in the model corresponds to periods of high firm entry, when γ_x is the main driver of uncertainty. The average fixed cost μ_f affects the average fraction of entrants n and relates to the average level of uncertainty.

The parameters are estimated by minimizing an equalweighted distance between the empirical and simulated moments. The numerical algorithm used to solve the model is described in Online Appendix F. Table I reports the fit. The calibrated parameters are $\gamma_x = 450$, $\mu_f = 0.0115$, and $\sigma_f = 0.0155$.³¹ As the table

30. Specifically, at the beginning of each quarter, the SPF asks each forecaster to report the probability that growth between the previous year and the current year will fall within each of several bins. We use the forecast reported in the fourth quarter, which represents a measure of uncertainty over just one quarter, because it maps easily into our model.

31. The average fixed cost amounts to close to 10% of steady-state firm level output. Bartelsman, Haltiwanger, and Scarpetta (2013) estimate overhead labor

TABLE 1

Uncertainty about GDP growth	Data (%)	Model (%)	
Mean	0.60	0.55	
5th percentile	0.45	0.50	
95th percentile	0.73	0.64	

CALIBRATED MOMENTS FROM THE SURVEY OF PROFESSIONAL FORECASTERS

Notes. Uncertainty is computed as the standard deviation of the SPF distribution over 1992:Q4 to 2007:Q4 of current year's annual over last year's annual real GDP stated in the last quarter of the current year. Growth rates and standard deviations are stated in percentage terms. We use the 5% and 95% percentiles, instead of the min and the max, for robustness against outliers. Uncertainty in our model is computed over a simulation of 50,000 periods using the same definition as in the data.

shows, the calibrated model cannot match all moments at the same time. In particular, because TFP is the only source of uncertainty while the SPF forecasters may worry about other shocks, we have difficulty matching the upper tail of uncertainty in the data and there is on average less uncertainty in our model. We thus view our results as conservative on the role of uncertainty in the economy.

V.C. Business Cycle Moments

We now evaluate the performance of our model in explaining standard business cycle moments. For that, we compare the benchmark model we have just described against the data and against two alternative models. The first alternative model (RBC) is the standard real business cycle model under complete information, identically parametrized. The second model (fixed θ -uncertainty) is a version of our model in which firms update their beliefs about the mean of the fundamental θ (i.e., they update μ using equation (24)) but uncertainty (i.e., the inverse of the precision of beliefs γ) remains constant at its long-run average in our benchmark model. Comparing our full framework to the fixed θ -uncertainty model highlights the specific role played by endogenous uncertainty.

Before comparing the models to the data, we must detrend the empirical time series. We note that the cyclical properties we are interested in, in particular the persistence and depth of the Great Recession, are sensitive to the specific detrending strategy. In Online Appendix E.2, we investigate the implications of three filters: the Hodrick and Prescott (1997) (HP) filter, a linear detrending, and a linear detrending allowing for a structural

costs on the order of 14% of U.S. employment, while Ramey (1991) reviews evidence according to which this fraction is on the order of 20%.

break in the trend. As shown in Online Appendix Figure E.12a, the HP-filtered data suggests that the Great Recession was a mild economic downturn and that the economy promptly recovered to its long-run trend after the trough. Both conclusions contradict essential features of the raw output data shown in Online Appendix Figure E.11a.³² We also find that a purely linear trend exaggerates the severity and persistence of the recession by ignoring low-frequency changes in the trend. Therefore, we choose a linear trend with a structural break estimated by least squares (Hansen 2000) as our benchmark, and we report the sensitivity of our results to using a standard linear trend.³³ See Online Appendix E for a more detailed discussion of the data and the detrending strategies.

We compute standard business cycle moments using the detrended data covering the period between the first quarter of 1978 and the last quarter of 2014. For the three models, we generate 10,000 simulations of the same length as the data (148 quarters), and then average each moment across all simulations. Panel A of Table II reports the results for the standard deviation of output (Y), consumption (C), employment (L), investment (I), the number of firms (m), and uncertainty about real output growth (U). Panel B reports the correlation between each of these variables and output, and Panel C reports their autocorrelation.

Because of the absence of learning, the RBC model generates more volatility in every variable than our full model. However, their performances are overall similar. Regarding the volatility of uncertainty about real output growth, our model is able to explain only a fraction of what is observed in the data, suggesting that other shocks—possibly exogenous uncertainty shocks—may be needed to fully account for the total fluctuations in uncertainty. Note that uncertainty about current-quarter real output growth is 0 in the RBC model under full information. Panels B and C show that the benchmark, the fixed θ -uncertainty model, and the

32. See King and Rebelo (1993) and Cogley and Nason (1995a) for a discussion of various drawbacks of the HP filter.

33. An additional advantage of the linear trends is that they allow us to cleanly interpret the (unfiltered) prediction of forecasters reported in the SPF data used to calibrate the model. For instance, a 2% growth forecast for this year relative to the previous year can be directly mapped to the model by removing the long-run growth rate in output from the linear filter. It is unclear how the analog exercise would be done if the data are HP filtered and detrended with a bandpass filter. See Online Appendix E.2 for a detailed discussion of the filters.

,,,						
	Y	С	L	Ι	т	U
Panel A: Standard deviation						
Data	0.039	0.031	0.066	0.132	0.038	0.253
Benchmark	0.035	0.030	0.027	0.072	0.020	0.053
Fixed θ -uncertainty	0.036	0.031	0.027	0.076	0.019	0.027
RBC	0.044	0.036	0.033	0.093	_	_
Panel B: Correlation w.r.t Y						
Data	1	0.588	0.921	0.523	0.646	-0.443
Benchmark	1	0.967	0.884	0.855	0.459	-0.644
Fixed θ -uncertainty	1	0.966	0.888	0.858	0.480	-0.619
RBC	1	0.987	1	0.944	_	_
Panel C: Autocorrelation (1st lag)						
Data	0.981	0.982	0.993	0.962	0.995	0.762
Benchmark	0.956	0.943	0.942	0.918	0.991	0.931
Fixed θ -uncertainty	0.957	0.943	0.941	0.915	0.989	0.869
RBC	0.934	0.949	0.934	0.907	_	_

TABLE II

BUSINESS CYCLE MOMENTS: DATA, RBC, AND BENCHMARK MODEL

Notes. All series are computed in log deviation from trend. Each of the 10,000 replications of the simulated series are linearly detrended. The annual series for m is interpolated to quarterly data. Uncertainty U is the standard deviation of current year real output growth expressed in the last quarter of the current year. Because it is only available annually and by definition cannot simply be interpolated, the moments we report about uncertainty are computed using observations from the fourth quarter of every year. Its annual autocorrelation is expressed in quarterly terms. C is gross of fixed costs for comparison across models. Employment and output are perfectly correlated in the RBC model as a consequence of the GHH preferences.

RBC model are roughly comparable in terms of the correlations with output and the autocorrelations.

The key conclusion from Table II is that a standard calibration of our model performs similarly to the RBC model in terms of standard business cycle moments. Therefore, incorporating endogenous uncertainty in a standard business cycle model does not impair its ability to predict well-known patterns of business cycle data. This result should not be surprising. As we saw in the baseline model, the uncertainty trap mechanism kicks in only for large shocks, which are rare in these simulations. In the next section, we show that the key difference between our model and standard models, and the value added of modeling uncertainty traps, lies in terms of predicting how the economy responds to large shocks.

V.D. Policy Function and Impulse Responses

We now examine the role of endogenous uncertainty in propagating shocks. First, we ask whether the key implication of the uncertainty trap mechanism identified in the baseline theory—generating protracted recessions out of sufficiently QUARTERLY JOURNAL OF ECONOMICS

1678



Investment Decision $n(\mu, \gamma, K, Q)$ for Constant K and m at Their Steady-State Levels

negative shocks—is also at work in the full calibrated model. In the next section, we ask whether this feature of the model can help explain the behavior of macro aggregates during the Great Recession.

1. Firm Entry and Beliefs. In Proposition 1 we showed that under certain conditions, the mass of producing firms increases with the mean of beliefs μ and decreases with uncertainty (increases with γ). We verify that these properties are inherited by the calibrated quantitative model. Figure VIII plots the fraction of entering firms $n(\mu, \gamma, K, m)$ as a function of mean beliefs μ , for three levels of uncertainty. Similarly to Figure IV in the baseline theory, we find that for any level of uncertainty, firms are more likely to enter when beliefs are more optimistic and that for any level of mean beliefs, firms are more likely to enter when uncertainty is lower. Therefore, the entry decision of the quantitative model inherits the key feature of the baseline theory leading to uncertainty traps.

2. Response to Small and Large Shocks. A central feature of uncertainty traps is the nonlinear response to shocks, shown in Figure VI in the context of the baseline theory. For shocks leading to large negative drops in the mean of beliefs, endogenous uncertainty leads to recessions which are relatively deeper and last relatively longer than for small shocks. We ask whether the full quantitative model inherits this feature. Figure IX shows the evolution of the mean of beliefs μ , the precision of beliefs γ , and output Y after a one-period drop in mean beliefs of 1% (left column) and 5% (right column) in both the full model and the fixed



FIGURE IX

Mean Beliefs, Precision of Beliefs, and Output in Response to a One-Period Shock

The left column shows the response of the economy to a -1% one-period shock to μ . The right column shows the response of the economy to a -5% one-period shock to μ . The solid curves show the evolution of the economy according to the full model, while the dashed curves show the evolution of a control economy in which the flow of public information is fixed at the steady-state level of the full model. Panels E and F are in log deviation from trend, the other figures are in levels.

 θ -uncertainty model.³⁴ In the fixed-information model, the economy starts to recover immediately after the initial shock regardless of the size of the shock. In contrast, in the full model, output continues to decline after the mean of beliefs has started to recover. Moreover, the duration of this decline is longer the larger is the shock. Reaching the trough in output takes 6 quarters in response to the small shock, and 15 quarters in response to the larger one. As in the baseline model, this nonlinearity is driven by the endogenous evolution of uncertainty. The fall in mean beliefs drives down the incentives to enter. As a result, fewer signals are released and the precision of beliefs falls (Panels C and D). After the shock, agents receive signals suggesting an improvement

34. Fundamental shocks affect the policy functions only through beliefs. Falls in beliefs may result from shocks to the fundamental θ or to the signals, *X* and *Z*.

TABLE III

THE IMPACT OF SHOCK SIZES ON THE DEPTH AND DURATION OF RECESSIONS ACROSS MODELS

	Small shock (-1%)		Large shock (-5%)	
	Depth (%)	50% recovery (qtrs.)	Depth (%)	50% recovery (qtrs.)
Full model	-1.1	42	-6.2	58
Fixed θ -uncertainty	-0.9	30	-4.6	35

Notes. The depth of the recession corresponds to the lowest value of output reached since the official beginning of the recession. The "50% recovery" column is the number of quarters before the economy recovers 50% of the peak-to-trough drop in output.

in the fundamental, and beliefs recover (Panels A and B). However, the recovery in output is delayed in the full model by the feedback between high uncertainty (Panels E and F) and slow entry.

Table III summarizes the properties of the recessions depicted in Figure IX. For both shocks, the table reports the ensuing recession's depth (the magnitude of the peak-to-trough fall in output) and the half-life of the recovery (the number of quarters for the economy to recover half of the peak-to-trough fall in output). The table highlights two key features of uncertainty traps: (i) endogenous uncertainty makes recessions deeper and longer for shocks of any magnitude (recessions' depth and duration are larger in the full than in the fixed θ -uncertainty model), and (ii) the differential effect of endogenous uncertainty is relatively larger for large than for small shocks. The recession is 22% deeper and 40% longer in the full model than in the fixed θ -uncertainty model under a negative 1% shock to beliefs, but 35% deeper and 66% longer under a negative 5% shock. Therefore, in the full calibrated model, endogenous uncertainty leads to amplification and persistence of shocks driving down beliefs.

V.E. Endogenous Uncertainty in U.S. Recessions

Our previous discussion established that, within our calibrated model, endogenous uncertainty amplifies and lengthens the decline in economic activity relatively more for large than for small shocks. We now evaluate the role of endogenous uncertainty in explaining the U.S. experience during past recessions. Specifically, we ask whether it can help explain the observed depth and persistence of a recession when there is a large enough shock and if, indeed, the mechanism only "kicks in" for sufficiently large shocks.

To implement this exercise we must first take a stand on the exogenous shocks that hit the U.S. economy in each recession. Within our model, the only exogenous shocks are the innovations to the signals $\{X_t, Z_t\}$ and to the fundamentals $\{\theta_t\}$. The shocks to the fundamentals $\{\theta_t\}$ are taken directly from the data. To discipline the evolution of $\{X_t, Z_t\}$, we choose their innovations such that our model and the fixed θ -uncertainty model exactly reproduce, as equilibrium outcomes, the forecast of current year's real output from the SPF during the first part of each recession.³⁵ We then compute the response of the three models we have introduced—our full quantitative model, the fixed θ -uncertainty model, and the RBC model-and contrast them with the data. Importantly, among the variables that we simulate and contrast against the data we include uncertainty about output growth from the SPF and the number of firms, both of which are directly related to the forces in the model.

We implement this exercise in the context of the largest and second-largest recessions in terms of decline in output within our sample, the Great Recession and the 1981–1982 recession. Since we expect uncertainty traps to be more important after large shocks, the Great Recession—the most important downturn since the Great Depression—is a historical episode where we would expect endogenous uncertainty to play a role. However, we also verify that endogenous uncertainty does not transform every downturn into a deeply protracted recession. For that, we also consider the impact of the mechanism during the 1981–1982 recession, which was characterized by a large but brief decline in productivity and was followed by a relatively rapid recovery in output. We expect the impact of endogenous uncertainty to be relatively weaker during this episode.

1. Great Recession. We start by considering the Great Recession. According to the NBER, the recession took place over the period 2007:Q4–2009:Q2. We first feed into the model the

^{35.} Note that we do not need to specify exactly the combination of X_t and Z_t shocks that is needed to match output forecasts. Only the time series of μ_t is relevant for our purpose. In the case of the RBC model we only introduce shocks to θ_t , as there is no degree of freedom to match the forecast about real output.

observed TFP shocks over the corresponding period.³⁶ We then feed each model a series of signals that allows them to replicate output forecasts from the SPF over that period. Specifically, the SPF surveys forecasters every quarter about their point expectations for the level of real output for the current year. Since the trough in output forecast from the SPF was reached in 2009:Q2, we focus on second quarter forecasts and match expectations about current year output in 2007:Q2, 2008:Q2, and 2009:Q2.³⁷

Figure X shows the evolution of TFP, real output forecasts, output, investment, employment, consumption, the number of firms, and uncertainty about real GDP growth. The solid lines in the left column represent the data. The right column shows the predictions of the three models: the full quantitative model (solid blue curves), the fixed θ -uncertainty model (dashed brown curves), and the RBC model (dot-dash pink curves).³⁸ The diamonds in Figure X, Panels B and D indicate the data that the models are calibrated to match. All shocks and innovations are set to 0 after 2009:Q2, and the economy is left to recover.

The central finding of this exercise is that for all the macro aggregates that we consider (output, investment, consumption, employment, number of firms, and uncertainty about GDP growth), our model generates dynamics that are closer to the data than the fixed θ -uncertainty model and the RBC model. Figure X, Panel F shows that the decline in output is clearly deeper and more protracted in our model than in the alternative models.

Table IV summarizes the key properties of the recession in the data and as predicted by the three models. The first column

36. We pick 2007:Q2 as our start date to be consistent with the dates at which we match output forecasts. However, we stop feeding the series in 2008:Q4 because the TFP series from Fernald (2014) displays a large increase at the beginning of 2009, which leads to a strong counterfactual expansion in the RBC economy (the other models fare better). For robustness, we fit the full TFP series in Figure A.1 of Online Appendix A and experiment with other TFP series in Figures A.3 and A.4. In all cases, the full model generates a deeper and longer recession than the fixed θ -uncertainty and RBC models.

37. Forecasts in the first, third, or fourth quarters in the SPF are about the current year and thus correspond to forecasts at different horizons. We focus on a single horizon with second-quarter forecasts to avoid using four different definitions of forecast in our model. Since we do not use forecast data in Q1, Q3 and Q4, we linearly interpolate the series of μ_t at these points. The empirical output forecast series stops in 2009:Q2 because of a change in variable definition.

38. A color rendition of Figure \boldsymbol{X} is available in the online version of this article.

UNCERTAINTY TRAPS





Great Recession: Data and Models

The left column shows the data while the right column shows the models. In the right column, the solid curves show the full model, the dashed curves show the fixed θ -uncertainty model and the dot-dash curves show the RBC model. All figures are in log deviation from trend. The scales are the same in both columns. A full color rendition of this figure is available in the online version of this article.

	Depth (%)	20% recovery	Cumulative loss (%)
Data	-7.9	2014:Q3	-184
Full model	-7.8	2015:Q1	-171
Fixed θ -uncertainty	-6.0	2012:Q4	-134
RBC	-2.6	2009:Q3	-56

TA	BL	Æ	IV
----	----	---	----

DEPTH, PERSISTENCE, AND CUMULATIVE OUTPUT LOSS OF THE GREAT RECESSION

Notes. We compute the depth, persistence, and cumulative loss of the recession from the official beginning of the recession in 2007:Q4 as determined by the NBER. The depth of the recession corresponds to the lowest value of output reached since the official beginning of the recession. The "20% recovery" column is the earliest date at which output recovers 20% of the depth. Cumulative loss is the sum of the percentage deviation of output from trend from 2007:Q4 to 2015:Q1.

shows the depth of the recession, measured as the percentage drop in output from the beginning of the recession to its lowest value. The full model is able to explain almost all of the depth. while the fixed θ -uncertainty model and the RBC model account for 76% and 33% of the fall, respectively. The second column of the table shows the earliest date at which output has recovered 20% of the depth.³⁹ In the full model, output recovers 20% of the lost ground in 2015:Q1, close to the corresponding date in the data (2014:Q3). In contrast, in the fixed θ -uncertainty and in the RBC models this event happens in 2012:Q4 and 2009:Q3, respectively. These results imply that within our model, endogenous uncertainty adds 1.8 (5.2) percentage points in terms of depth and about two (five) years until the 20% recovery relative to the fixed θ -uncertainty (RBC) model. Finally, the third column shows the cumulative loss in output, measured as the sum of the percentage deviations from trend from 2007:Q4 to 2015:Q1, the end of our sample. This measure combines the depth and the persistence of the recession into a single statistic. If we normalize this measure by the output produced in guarter 2007:Q4 at the onset of the recession, the U.S. economy lost a cumulated 1.84 quarters of output since that date. Our full model explains 93% of this loss, while the fixed θ -uncertainty and RBC models can account for 73% and 30%

39. We cannot use the standard half-life measure of persistence since output has not recovered half of the way from its trough in the data. We pick the 20% threshold because it roughly corresponds to the level attained by output in the data at the end of the sample. The full model also generates a longer recession measured in terms of the number of quarters from peak to trough relative to alternative models. The exact date for the trough is hard to identify in the data because of high-frequency noise.

of the loss, respectively. According to these various measures, our full model is better able to account for the overall severity of the recession.

Besides output, endogenous uncertainty can also help explain the behavior of other aggregates. As shown in Panels G and H, even though our model does not generate as pronounced a decline in investment, it does generate more persistence than the alternative models. Figure X. Panels M and N show that the number of firms also falls. By construction, the RBC model does not have predictions in this regard. Our full quantitative model does well in predicting the persistence and the amplitude of the fall in the number of active firms. Since entry and the number of firms govern the dynamics of uncertainty within the full model, the good performance of the model along that dimension is reassuring. Our model generates as large a fall in consumption as is observed in the data (Panels K and L). Relative to the alternative models, it also predicts a larger fall and a more protracted recovery in employment (Panels I and J), better reflecting the empirical patterns. Even though for most variables the fall in economic aggregates happens faster in the data, our model is closer to the empirical series than the alternative models for all the variables reported in the figure.

We also highlight that our model is relatively successful at explaining the shape of the uncertainty response. As observed in Figure X, Panels O and P, the uncertainty in output growth follows an inverted-U shape relationship that our model reproduces. although in our model the peak in uncertainty is delayed relative to the data. In contrast, the alternative models generate little variation in uncertainty in GDP forecasts.

We demonstrate the robustness of our findings by replicating the Great Recession exercise under alternative assumptions: (i) fitting the whole TFP series instead of only its beginning; (ii) using a simple linear detrending instead of allowing for a structural break; (iii) using TFP data from the Penn World Table, and the Bureau of Labor Statistics instead of the Fernald time series; and (iv) using standard CRRA preferences instead of the GHH preferences assumed in the main exercise. The full simulations from these robustness exercises are included in Online Appendix A. Table V summarizes these simulations by showing the cumulative loss in output generated by each model in each exercise and comparing it to the data. In all cases, we find that our conclusions are robust. The full model with endogenous uncertainty gets closer to

	Cumulative loss in output (%)			
	Data	Full model	Fixed θ -uncertainty	RBC
Benchmark	-184	-166	-129	-56
Full TFP fit	-184	-86	-38	+38
Linear detrending	-310	-173	-139	-72
PWT TFP	-184	-194	-175	-131
BLS TFP	-184	-171	-142	-85
CRRA Preferences	-184	-80	-65	-31

TABLE V

ROBUSTNESS: CUMULATIVE LOSS DURING THE GREAT RECESSION

Notes. Cumulative loss is the sum of the percentage deviation of output from trend from 2007:Q4 to 2015:Q1.

TABLE VI

DEPTH, PERSISTENCE, AND CUMULATIVE OUTPUT LOSS OF THE 1981-1982 RECESSION

	Depth (%)	20% recovery	Cumulative loss (%)
Data	-5.3	1983:Q2	-13
Full model	-4.9	1984:Q1	-32
Fixed θ -uncertainty	-4.7	1984:Q1	-32
RBC	-6.4	1984:Q2	-44

Notes. The table is constructed similarly to Table IV. The official beginning of the recession is in 1981:Q3, and the cumulative loss is computed between that quarter and 1985:Q1.

the data by generating more severe recessions than the alternative models. Table A.1 in the Online Appendix shows that the full model generates deeper and longer lasting recessions in all the robustness exercises.

2. 1981–1982 Recession. We consider the impact of endogenous uncertainty in the less severe recession of 1981:Q3–1982:Q4. The purpose of this section is to show that while endogenous uncertainty generates amplification and persistence in response to very large shocks, it does not necessarily increase the depth and persistence of every negative shock. As such, the model does not have counterfactual implications for the magnitude and persistence of milder recessions.

Table VI replicates Table IV for the 1981–1982 recession.⁴⁰ We see that the full and the fixed θ -uncertainty models generate shallower recessions than the data, while output drops more than the

^{40.} Since the output forecast series is fully available over the period 1981– 1985, we fit the whole time series of both beliefs and TFP. Online Appendix B shows the model predicted series.

data in the RBC model. In terms of persistence, the three models perform similarly, suggesting that the economy should have recovered 20% of the loss in output at the beginning of 1984—about a year after that threshold was reached in the data. In terms of cumulative loss, all models predict recessions that are more severe than the data, although the difference between model and data is larger for the RBC than for the full and fixed θ -uncertainty models. Overall, we find that for the 1981–1982 recession, characterized by a less severe combination of TFP and beliefs shocks, endogenous uncertainty does not generate additional persistence. This confirms that the full model generally behaves as a standard RBC model, except during large unusual downturns, such as the 2007–2009 recession, when the uncertainty trap mechanism kicks in.

V.F. Further Applications

In Online Appendix C we provide additional exercises that highlight various features of the model. First, we show that like other studies of the effects of uncertainty on the economy (Bloom 2009; Bloom et al. 2012), exogenous uncertainty shocks generate recessions in our framework. In our model, these exogenous shocks are propagated endogenously through the endogenous uncertainty channel.

Second, we consider the problem of a social planner in Online Appendix C.2.⁴¹ To do so, we compute the efficient allocation and consider how it responds to shocks relative to the competitive equilibrium we have studied so far. We find that the planner prevents the number of entering firms from falling too much during the recession to prevent uncertainty from rising too much. As a result, the recession is shorter and shallower in the planner's allocation, suggesting a potential role for government policies in mitigating the impact of recessions.

Third, we show in Online Appendix C.3 that our benchmark model improves on RBC in terms of propagation of shocks. As noted by Cogley and Nason (1995b), the RBC model features weak internal persistence so that the properties of GDP growth mimic

^{41.} The formulation of the constrained planner's problem is equivalent to that discussed in Proposition 4, with the difference that here, the planner also chooses the additional margins of the quantitative model (intensive margin of investment and labor supply) alongside the extensive margin of investment among potential entrants.

those of the exogenous TFP growth process. In our full model, however, the persistence in output growth is larger than in TFP growth, and more so at lower lags. Hence, although our model does not replicate the full amount of autocorrelation observed in the data, it does reproduce the qualitative feature of the data that the autocorrelation of GDP growth is larger than that of TFP growth at lower lags. We interpret this finding as evidence that endogenous uncertainty leads to a more persistent series for output growth relative to TFP growth, in line with the baseline theory.

VI. CONCLUSION

We develop a theory of endogenous uncertainty and business cycles that combines two forces: higher uncertainty about economic fundamentals deters investment, and uncertainty evolves endogenously because agents learn from the actions of others. The interaction between investment and uncertainty leads to uncertainty traps: episodes in which high uncertainty leads firms to delay investment, further raising uncertainty. In the unique equilibrium of the model, the economy fluctuates between a high-activity/low-uncertainty regime and a low-activity/highuncertainty regime and is subject to strong nonlinear dynamics in which large shocks can have near permanent effects.

To quantify the importance of uncertainty traps, we embed the mechanism into a business cycle model. We calibrate the model to the U.S. economy and find that endogenous uncertainty increases the depth and duration of recessions, particularly after large shocks. We also find that our full model is able to better explain the behavior of standard macro aggregates during the Great Recession than are alternative models.

We believe that the novel channel proposed in this article is important for several reasons. First, the emphasis on subjective uncertainty—and beliefs about fundamentals in particular implies that not only exogenous volatility shocks but also other sources of uncertainty matter for the economy. Thus, we view recent empirical work using survey data on forecasts or consumer and business expectations as an important step toward a more complete understanding of the role of uncertainty in business cycles. Second, we believe that our framework may be useful as a theoretical benchmark for empirical and quantitative studies seeking to estimate the direct and feedback effects of uncertainty on economic activity. Despite the multiplicity of regimes and strong nonlinearities, the model features a single competitive equilibrium, which makes it amenable to applied work. Third, we have shown that allowing uncertainty to fluctuate endogenously leads to a significant propagation and amplification mechanism. The type of nonlinearities and the multiplicity in regimes that we obtain may be of broader interest for business cycle modeling in general and could also shed light on some particularly large historical downturns.

For the sake of clarity, we have exposited the mechanism in a purposely simple framework, but a number of generalizations may be worth investigating. In particular, it would be interesting to understand how uncertainty traps interact with frictions that could magnify their impact, such as financial frictions, demand externalities, or belief heterogeneity. We leave these questions to future research

UNIVERSITY OF CALIFORNIA, LOS ANGELES UNIVERSITAT POMPEU FABRA, CREI WHARTON SCHOOL OF THE UNIVERSITY OF PENNSYLVANIA

SUPPLEMENTARY MATERIAL

An Online Appendix for this article can be found at *The Quar*terly Journal of Economics online. Data and code replicating the tables and figures in this paper can be found in Fajgelbaum, Schaal, and Taschereau-Dumouchel (2017), in the Harvard Dataverse, doi:10.7910/DVN/53F81R.

REFERENCES

- Amador, Manuel, and Pierre-Olivier Weill, "Learning from Prices: Public Commu-
- nication and Welfare," *Journal of Political Economy*, 118 (2010), 866–907. Angeletos, George-Marios, Christian Hellwig, and Alessandro Pavan, "Dynamic Global Games of Regime Change: Learning, Multiplicity, and the Timing of Attacks," Econometrica, 75 (2007), 711-756.
- Arellano, Cristina, Yan Bai, and Patrick Kehoe, "Financial Markets and Fluctu-ations in Uncertainty," Working paper, Federal Reserve Bank of Minneapolis and NBER: Arizona State University: University of Minnesota and Federal Reserve Bank of Minneapolis, 2012.

Bachmann, Rudiger, and Christian Bayer, "Investment Dispersion and the Business Cycle," American Economic Review, 104 (2014), 1392-1416.

"Wait-and-See Business Cycles?," Journal of Monetary Economics, 60 (2013), 704-719.

Bachmann, Rudiger, and Giuseppe Moscarini, "Business Cycles and Endogenous Uncertainty," Mimeo, Yale University, July 2011.

- Bachmann, Rudiger, Steffen Elstner, and Eric R. Sims, "Uncertainty and Economic Activity: Evidence from Business Survey Data," American Economic Journal: Macroeconomics, 5 (2013), 217–249.
- Banerjee, Abhijit V., "A Simple Model of Herd Behavior," Quarterly Journal of Economics, 107 (1992), 797–817.
- Barro, Robert J., "Rare Disasters and Asset Markets in the Twentieth Century," Quarterly Journal of Economics, (2006), 823–866.
- Barro, Robert J., and José F. Ursúa, "Rare Macroeconomic Disasters," Annual Review of Economics, 4 (2012), 83–109.
- Bartelsman, Eric, John Haltiwanger, and Stefano Scarpetta, "Cross-Country Differences in Productivity: The Role of Allocation and Selection," American Economic Review, 103 (2013), 305–334.
- Basu, Susanto, and John G. Fernald, "Returns to Scale in US Production: Estimates and Implications," *Journal of Political Economy*, 105 (1997), 249–283.
- Beaudry, Paul, and Franck Portier, "Stock Prices, News, and Economic Fluctuations," American Economic Review, 96 (2006), 1293–1307.
- Benhabib, Jess, Pengfei Wang, and Yi Wen, "Sentiments and Aggregate Demand Fluctuations," *Econometrica*, 83 (2015), 549–585.
- Bikhchandani, Sushil, David Hirshleifer, and Ivo Welch, "Learning from the Behavior of Others: Conformity, Fads, and Informational Cascades," *Journal of Economic Perspectives* 12 (1998), 151–170.
- ——, "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, (1992), 992–1026.
- Bloom, Nicholas, "The Impact of Uncertainty Shocks," *Econometrica*, 77 (2009), 623–685.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry, "Really Uncertain Business Cycles," NBER Working Paper no. 18245, 2012.
- Caplin, Andrew, and John Leahy, "Sectoral Shocks, Learning, and Aggregate Fluctuations," *Review of Economic Studies*, 60 (1993), 777–794.
- Carlsson, Mikael, "Investment and Uncertainty: A Theory-Based Empirical Approach," Oxford Bulletin of Economics and Statistics, 69 (2007), 603–617.
- Chamley, Christophe, "Delays and Equilibria with Large and Small Information in Social Learning," *European Economic Review*, 48 (2004), 477–501.
- ——, "Coordinating Regime Switches," *Quarterly Journal of Economics*, 114 (1999), 869–905.
- Chamley, Christophe, and Douglas Gale, "Information Revelation and Strategic Delay in a Model of Investment," *Econometrica: Journal of the Econometric* Society (1994), 1065–1085.
- Chetty, Raj, Adam Guren, Day Manoli, and Andrea Weber, "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review*, 101 (2011), 471–475.
- Cogley, Timothy, and James M. Nason, "Effects of the Hodrick-Prescott Filter on Trend and Difference Stationary Time Series Implications for Business Cycle Research," Journal of Economic Dynamics and Control, 19 (1995a), 253–278.
 —, "Output Dynamics in Real-Business-Cycle Models," American Economic Review, 85 (1995b), 492–511.
- Cooper, Russell W., and Andrew John, "Coordinating Coordination Failures in Keynesian Models," *Quarterly Journal of Economics*, 103 (1988), 441–463.
- Decker, Ryan A., and Pablo N. D'Erasmo, "Market Exposure and Endogenous Firm Volatility over the Business Cycle," American Economic Journal: Macroeconomics, 8 (2016), 148–198.
- Dixit, Avinash K., and Robert S. Pindyck, *Investment under Uncertainty* (Princeton, NJ: Princeton University Press, 1994).
- Fajgelbaum, Pablo D., Edouard Schaal, and Mathieu Taschereau-Dumouchel, "Replication Data for: 'Uncertainty Traps'," Harvard Dataverse (2017), doi:10.7910/DVN/53F81R.
- Farmer, Roger, "Animal Spirits, Financial Crises and Persistent Unemployment," Economic Journal, 123 (2013), 317–340.

Feenstra, Robert C., Robert Inklaar, and Marcel P. Timmer, "The Next Generation of the Penn World Table," American Economic Review, 105 (2015), 3150-3182.

Fernald, John G., "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Working paper, Federal Reserve Bank of San Francisco, 2014.

- Foster, Andrew D., and Mark R. Rosenzweig, "Learning by Doing and Learning from Others: Human Capital and Technical Change in Agriculture," *Journal* of Political Economy, 103 (1995), 1176-1209.
- Gabaix, Xavier, "Variable Rare Disasters: An Exactly Solved Framework for Ten Puzzles in Macro-Finance," Quarterly Journal of Economics, 127 (2012), 645-700.
- Gilchrist, Simon, Jae W. Sim, and Egon Zakrajšek, "Uncertainty, Financial Frictions, and Investment Dynamics," Technical report, NBER 2014.
- Gourio, Fran, "Disaster Risk and Business Cycles," American Economic Review, 102 (2012), 2734-66.
- Guiso, Luigi, and Fabiano Schivardi, "Spillovers in Industrial Districts," Economic Journal, 117 (2007), 68-93.
- Hansen, Bruce E., "Testing for Structural Change in Conditional Models," Journal of Econometrics, 97 (2000), 93-115.
- Hodrick, Robert J., and Edward C. Prescott, "Postwar U.S. Business Cycles: An Empirical Investigation," Journal of Money, Credit and Banking, 29 (1997), 1 - 16.
- Jaimovich, Nir, and Sergio Rebelo, "Can News about the Future Drive the Business Cycle?," American Economic Review, 99 (2009), 1097–1118.
- Jurado, Kyle, Sydney C. Ludvigson, and Serena Ng, "Measuring Uncertainty," American Econonmic Review, 105 (2015), 1177–1216. Kaplan, Greg, and Guido Menzio, "Shopping Externalities and Self-Fulfilling Un-
- employment Fluctuations," Journal of Political Economy, 124 (2016), 771-825.
- Kehrig, Matthias, "The Cyclicality of Productivity Dispersion," U.S. Census Bureau Center for Economic Studies Paper No. CES-WP-11-15, 2011.
- King, Robert G., and Sergio T. Rebelo, "Low Frequency Filtering and Real Business Cycles," Journal of Economic Dynamics and Control, 17 (1993), 207–231.
- Kozlowski, Julian, Laura Veldkamp, and Venky Venkateswaran, "The Tail that Wags the Economy: Belief-Driven Business Cycles and Persistent Stagnation," Working paper, New York University, 2017.
- Lang, William W., and Leonard I. Nakamura, "The Dynamics of Credit Markets in a Model with Learning," Journal of Monetary Economics, 26 (1990), 305–318.
- Lucas, Robert E., "On the Size Distribution of Business Firms," Bell Journal of Economics (1978), 508-523.

"Expectations and the Neutrality of Money," Journal of Economic Theory, 4 (1972), 103–124.

- Morris, Stephen, and Hyun Song Shin, "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks," American Economic Review (1998), 587-597.
 - , "A Theory of the Onset of Currency Attacks," Asian Financial Crisis, 2 (1999), 230-255.
- Ordoñez, Guillermo L., "Larger Crises, Slower Recoveries: The Asymmetric Effects of Financial Frictions," Technical report, Federal Reserve Bank of Minneapolis, 2009.

Orlik, Anna, and Laura Veldkamp, "Understanding Uncertainty Shocks and the Role of Black Swans," Working paper, 2013.

Ramey, Valerie A., "[Markups and the Business Cycle]: Comment," NBER Macroeconomics Annual, 6 (1991), 134-139.

Rob, Rafael, "Learning and Capacity Expansion under Demand Uncertainty," Review of Economic Studies, 58 (1991), 655-675.

- Schaal, Edouard, "Uncertainty and Unemployment," Working paper, New York University, 2015.
- Schaal, Edouard, and Mathieu Taschereau-Dumouchel, "Coordinating Business Cycles," Working paper, 2015.
 - , "Aggregate Demand and the Dynamics of Unemployment," Working paper, 2016.

- Stokey, Nancy L., The Economics of Inaction: Stochastic Control Models with Fixed Costs (Princeton, NJ: Princeton University Press, 2008).
- Straub, Ludwig, and Robert Ulbricht, "Endogenous Uncertainty and Credit Crunches," manuscript, 2015.
- Van Nieuwerburgh, Stijn, and Laura Veldkamp, "Learning Asymmetries in Real Business Cycles," *Journal of Monetary Economics*, 53 (2006), 753–772.
 Vavra, Joseph, "Inflation Dynamics and Time-Varying Volatility: New Evidence
- Vavra, Joseph, "Inflation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *Quarterly Journal of Economics*, 129 (2013), 215– 258.
- Veldkamp, Laura L., "Slow Boom, Sudden Crash," Journal of Economic Theory, 124 (2005), 230–257.
- Zeira, Joseph, "Informational Cycles," Review of Economic Studies, 61 (1994), 31–44.