

CREATIVE DESTRUCTION AND UNCERTAINTY

Petr Sedláček

University of Oxford

Abstract

Uncertainty rises in recessions. But does uncertainty cause downturns or vice versa? This paper argues that counter-cyclical uncertainty fluctuations are a by-product of technology growth. In a firm dynamics model with endogenous technology adoption, faster technology growth widens the dispersion of firm-level productivity shocks, a benchmark uncertainty measure. Moreover, faster technology growth spurs a creative destruction process, generates a temporary downturn and renders uncertainty counter-cyclical. Estimates from structural VARs on U.S. data confirm the model's predictions. On average, 1/4 of the cyclical variation in uncertainty is driven by technology shocks. This fraction rises to 2/3 around the “dot-com” bubble. (JEL: D22, E32, D80)

Keywords: creative destruction, uncertainty, business cycles, growth.

Acknowledgments: I thank Christian Bayer, Nicholas Bloom, Benjamin Born, Francisco Buera, Steven Davis, Wouter den Haan, Greg Kaplan, Matthias Kehrig, Zheng Liu, Emrah Mahmutoglu, Benjamin Pugsley, Markus Riegler, Moritz Schularick, Emily Sedlacek-Swift, Vincent Sterk, Robert Swift and participants of seminars and conferences at various institutions for helpful comments. I gratefully acknowledge the financial support of the Daimler and Benz foundation for this project.

E-mail: petr.sedlacek@economics.ox.ac.uk

1. Introduction

Uncertainty rises during recessions. While this stylized fact is robust to many refinements, the question of whether uncertainty is an exogenous source of business cycles or an endogenous response to them is not well understood. This paper argues that counter-cyclical fluctuations in uncertainty are a by-product of changes in technology growth. Moreover, such growth-driven uncertainty changes are found to be quantitatively important in U.S. data.

To study the link between technology growth, business cycles and firm-level uncertainty, I build a tractable general equilibrium model of endogenous firm dynamics and technology adoption. In this model, firms can improve their productivity by investing into the adoption of newer vintages of technology which grow stochastically over time. When the technological frontier expands, firms face relatively larger productivity gains if they successfully adopt newer vintages and relatively larger productivity losses if they do not. In other words, faster technology growth widens the dispersion of firm-level productivity shocks, a benchmark measure of uncertainty.¹ Endogenous technology adoption then serves as a strong source of magnification and propagation of uncertainty responses to technology shocks.

In addition, expansions of the technological frontier spur a process of creative destruction. A technological improvement raises productivity of firms utilizing the latest technology vintage. This leads to an increase in consumption and wages, but this rise is only gradual as consumption smoothing motives of the household direct some of the productivity gains into investment. Therefore, faster technology growth has opposing effects on firms, depending on their technology vintage. On the one hand, firms at the frontier enjoy productivity gains larger than the increase in wage costs prompting them to create jobs. On the other hand, firms which have not adopted the newer vintage of technology experience only a rise in labor costs and as a result they shed workers and shut down more often.

In the calibrated model, which matches salient features of U.S. firm dynamics, the initial surge in job destruction dominates and the economy undergoes a temporary Schumpeterian downturn. Over time, however, aggregate productivity increases as more firms adopt the leading technology and obsolete production units get weeded out. Therefore, in contrast to popular models of uncertainty-driven business cycles (see e.g. Bloom, 2009), in this model counter-cyclical increases in uncertainty are associated with positive long-run effects.²

To see whether this channel is also empirically relevant I test the model predictions in the data using two distinct approaches. First, while uncertainty is counter-cyclical, the model predicts that uncertainty co-moves positively with changes in technology growth. I use 4-digit industry level data between 1971 and 2009 and show that both

1. See Bloom et al. (2018). The Appendix shows that the results are robust to alternative firm-level uncertainty measures.

2. See Lester et al. (2014) for an RBC model in which uncertainty is welfare increasing under certain calibrations.

these predictions also hold empirically. Moreover, uncertainty in the data is also positively related to average R&D intensity of the industry. Interpreting the latter as a proxy for the speed of technology adoption, this is also consistent with the model in which endogenous technology adoption magnifies and propagates the effects of technology shocks on uncertainty.

Second, I test the causal implications of the model regarding the impact of technology shocks on uncertainty and the creative destruction process. The empirical strategy is to estimate a series of bi-variate structural vector autoregressions (VARs) identifying technology shocks using long-run restrictions as in e.g. Blanchard and Quah (1989) and Gali (1999).³ The identifying assumption, consistent with the structural model, is that only technology shocks affect productivity in the long-run.

The structural VARs, estimated on annual data between 1977 and 2014, support the model's predictions. Specifically, following positive technology shocks firm-level uncertainty, job creation and job destruction all increase, while at the same time aggregate employment falls temporarily. Moreover, the magnitudes of these empirical impulse responses are in line with the quantitative predictions of the calibrated model.

Finally, I use the estimated structural VARs to quantify to what extent observed firm-level uncertainty is growth-driven. Forecast error variance decompositions suggest that on average 27% of the business cycle variation in uncertainty is driven by technology shocks alone. Zooming in on uncertainty spikes around the four recessions in the sample shows that there are large differences in the degree to which changes in growth drive uncertainty. While more than two thirds of the uncertainty increase around the "dot-com" recession in 2001 were growth-driven, the Great Recession spike in uncertainty was essentially unrelated to technology shocks.

This paper is related to several strands of the literature. First, it is connected to the large set of studies analyzing uncertainty movements over the business cycle (see e.g. Bloom, 2009; Bachmann and Bayer, 2014; Jurado et al., 2015). The notion that uncertainty fluctuations may be endogenous to the business cycle features in e.g. Bachmann and Moscarini (2012), Gourio (2014), Orlik and Veldkamp (2013), Boedo et al. (2016) or Berger and Vavra (forthcoming).⁴ Ludvigson et al. (2017) use instrumental variables to estimate that indeed a large part of uncertainty fluctuations are endogenous responses to other structural shocks. Instead, the model in this paper shows that aggregate downturns and increases in uncertainty are both partly driven by a third common factor: changes in technology growth.

Second, the focus on a link between uncertainty and growth is related to Ramey and Ramey (1991) and Koren and Tenreyro (2007) who document a negative relationship between growth and the (constant) level of macroeconomic volatility. Baker and Bloom (2013) use natural disasters to estimate a negative relationship

3. The Appendix shows that the results are robust to an alternative empirical strategy based on local projections following Jorda (2005) and using the technology shocks estimated by Basu et al. (2013).

4. Oi (1961); Hartman (1972); Abel (1983); Bar-Ilan and Strange (1996) entertain the possibility that increases in uncertainty come with positive effects and Ilut et al. (2018) argue that counter-cyclical cross-sectional volatility is a natural result of concave hiring rules.

between changes in uncertainty and growth. In contrast to these studies, this paper provides a structural model of firm dynamics and growth in which it can be shown analytically that uncertainty fluctuations are a by-product of changes in technology growth.

Finally, this paper also relates to models and empirical evidence on Schumpeterian creative destruction (see e.g. Aghion and Howitt, 1994; Caballero and Hammour, 1996; Mortensen and Pissarides, 1998, for earlier contributions). Many studies have documented that such technology shocks are recessionary in the short-run (see e.g. Gali, 1999; Francis and Ramey, 2005; Basu et al., 2006).⁵ The model and VAR results in this paper are closely related to those in Lopez-Salido and Michelacci (2007) and Canova et al. (2013), who analyze the effects of technology shocks on labor market flows. In contrast to their model, the framework in this paper considers multi-worker firms and endogenous technology adoption. The latter turns out to be an important propagation force. Importantly, the theoretical and empirical focus of this paper is on the relation between the process of creative destruction and uncertainty. To the best of my knowledge, the current paper is the first to document how firm dynamics and firm-level uncertainty respond to Schumpeterian technology shocks and to quantify this nexus.

The rest of the paper is structured as follows. The next section describes the structural model, it explains its calibration and it provides the model-based results. Section 3 then tests the model results in the data and quantifies to what extent uncertainty fluctuations are growth-driven in the data. Section 4 concludes.

2. Structural model

This section builds a tractable general equilibrium growth model with endogenous firm dynamics, technology adoption and business cycle fluctuations. In this model firms endogenously enter, exit and conditional on survival they grow over their life-cycle. Throughout their life-cycles firms invest into adopting better production technologies which improve stochastically over time.

The main goal of the model is to understand the link between growth, business cycle fluctuations and firm-level uncertainty. Following Bloom et al. (2018), the benchmark measure of uncertainty used throughout the paper is the dispersion of firm-level total factor productivity (TFP) shocks. The construction of this measure, together with intuition about why it constitutes an uncertainty proxy, are described in detail in Section 2.2.

5. Fisher (2006) stresses the importance of distinguishing between “neutral” and “investment-specific” technology shocks which typically have different qualitative effects. The Appendix shows that the results are robust to accounting for investment-specific technology shocks.

2.1. Model environment

The economy is populated by a representative household with a continuum of members and by a continuum of heterogeneous firms which are owned by the household. To ease the exposition, aggregate variables are denoted by upper-case letters, while firm-specific variables are denoted by lower-case letters. Let us begin by describing household preferences and choices and then move on to the process of technology adoption and the behavior of incumbent and entering firms.

2.1.1. Household preferences and choices. The representative household chooses consumption, C_t , investment into physical capital I_t , and the supply of labor, N_t on perfectly competitive factor markets. Following the indivisible labor models (see e.g. Hansen, 1985; Rogerson, 1988), labor is assumed to enter linearly into the household's utility function and is interpreted as the employment rate. Formally, the per-period utility of the representative household is given by

$$\ln C - vN, \quad (1)$$

where $v > 0$ is the disutility of labor and the preference specification allows for balanced growth. The representative household maximizes the expected present value of life-time utility, subject to its budget constraint

$$C + I = NW + RK + \Pi, \quad (2)$$

which states that total income stems from employment (with W being the competitive wage rate), renting out of capital to firms (with R being the competitive interest rate) and from the ownership of firms, where Π are aggregate profits. This total income is spent on consumption and investment into physical capital, where $I = K' - (1 - \delta)K$ and where δ is the depreciation rate, K is the stock of capital and primes indicate next period's values. The resulting optimality conditions are given by

$$W = vC, \quad (3)$$

$$\frac{1}{C} = \beta \mathbb{E} \frac{1}{C'} (R' + 1 - \delta). \quad (4)$$

2.1.2. Technology adoption, firm-specific productivity and growth. It is assumed that the frontier technology evolves exogenously according to the following process

$$\ln Z' = \bar{Z} + \ln Z + \varepsilon'_Z, \quad (5)$$

where $\bar{Z} > 0$ is a positive drift term and ε_Z are iid innovations distributed according to a Normal distribution with zero mean and standard deviation σ_Z .

The productivity level of an individual firm i is denoted by z_i . Each individual firm owns a particular vintage of the frontier technology and therefore we can express individual firm productivity in terms of a particular lag of the frontier technology, $z_{i,t} = v_{j,t} = Z_{t-j}$ with $j \geq 0$. In addition, firms can attempt to improve their

prevailing productivity levels by investing into costly adoption of newer technology vintages. Because the frontier is growing over time, an individual firm which fails to adopt newer technologies will experience a gradual decline in relative productivity. At some point, such a firm will become so unproductive that it will no longer be profitable to remain in operation and it will shut down endogenously.

The investment into technology adoption is interpreted broadly, not only as direct costs of purchasing a well-defined technology. In particular, the costs of technology adoption, or implementation, include e.g. those related to identifying best practices, personalizing and implementing such practices at a specific firm, reorganizing existing business procedures or aligning of incentives to use the new technology efficiently (see e.g. Lientz and Rea, 1998; Comin and Hobijn, 2007; Atkin et al., 2017). Because any of the above implementation requirements may fail, investment into technology adoption is inherently uncertain.^{6,7}

Following Klette and Kortum (2004), the various costs and resulting benefits of technology adoption are summarized using the following simple function. In particular, a firm investing r units of the final good has a probability p of adopting a newer technology vintage, where

$$p_i = \left(\frac{r_i}{\chi} \right)^{\frac{1}{\eta}} \gamma_i^{1-\frac{1}{\eta}}. \quad (6)$$

In the above expression, χ is a scaling factor, γ_i is the technology gap (or “stock of knowledge”) defined as $\gamma_i = \ln z_i - \ln Z$ and $1/\eta$ is a curvature parameter.⁸ The associated cost function can be written as

$$T(p_i, \gamma_i) = \chi \gamma_i \left(\frac{p_i}{\gamma_i} \right)^{\eta}. \quad (7)$$

As explained, if an incumbent firm fails to adopt a newer technology vintage, it retains its prevailing productivity level. Successful adoption attempts may lead to either radical or incremental technological improvements (as in e.g. Akcigit and Kerr, 2018). In particular, a fraction θ of firms adopting newer vintages adopt the frontier technology, while all other adopting firms obtain the technology of the closest younger technology vintage. Formally, if a firm i in period t has a productivity level

6. The assumption of gradual adoption of (frontier) technology is related to Comin and Gertler (2006). In contrast to the latter study which assumes homogeneous firms (and competitive technology adopters), the primary focus of this paper is the time-varying *distribution* of technology vintages across firms.

7. An alternative interpretation of the uncertainty surrounding technology adoption is that in addition to expected adoption costs, firms face heterogeneous and stochastic implementation costs. As new information about these costs emerges, firms may choose to abandon adopting a technology that was profitable in expectation (see e.g. Jack et al., 2015).

8. While the baseline specification assumes that expenditures on technology adoption constitute a resource cost for the firm, the Appendix shows that similar results are obtained when technology adoption costs are specified in labor units instead.

of $z_{i,t} = v_{j,t} = Z_{t-j}$, then its productivity in the next period is described by

$$\ln z_{i,t+1} = \begin{cases} v_{j+1,t+1} & \text{with probability } 1 - p_{i,t}, \\ v_{j,t+1} & \text{with probability } p_{i,t}(1 - \theta), \\ Z_{t+1} & \text{with probability } p_{i,t}\theta. \end{cases} \quad (8)$$

Finally, it is assumed that the process of technology adoption is the same for potential startups as it is for incumbent firms. As a normalization, the stock of knowledge for potential entrants is assumed to be given by the average stock of knowledge in the economy, $\bar{\gamma}$. Startups are assumed to enter the economy only if they manage to adopt the latest technology vintage.⁹

2.1.3. Firm behavior. Firm dynamics play a key role in this model. They feature endogenous firm entry and exit, an endogenous firm productivity (and thus size) distribution and firm life-cycle growth. Let us first describe these individual features and then turn to the formal firm maximization problem.

Incumbent firms differ in terms of their productivity levels which they can improve as described in Section 2.1.2. Conditional on their productivity level, firms produce output using labor and capital in a decreasing-returns-to-scale production technology. The gradual nature of technology adoption together with the presence of decreasing returns to scale in production result in a non-degenerate endogenous firm-level productivity (and thus size) distribution.

In the data, however, productivity gaps alone cannot account for the observed average size differences between young and more mature firms (see e.g. Foster et al., 2016). Therefore, to generate a realistic firm size distribution, which will be quantitatively important for the aggregate dynamics of the economy, firms in this model also grow over their life-cycles independent of their productivity levels.

In particular, it is assumed that firms accrue efficiency gains, ψ , through learning-by-doing. These gains are proportional to firm size and can be rationalized by for instance established long-term relationships, well-developed distribution networks or better management practices (see e.g. Stein, 1997). This makes more mature businesses, which do not necessarily operate cutting-edge technologies, competitive and able to fend off more innovative newcomers.¹⁰

Finally, in addition to variable costs, firms must also pay stochastic fixed costs of operation, φ . Businesses endogenously shut down when the realization of the fixed cost is too high rendering them unprofitable.¹¹

9. The Appendix shows that a model in which startups are characterized by a distribution of different technology vintages, rather than all starting at the frontier, yields similar results.

10. In addition, modeling life-cycle growth using such deterministic efficiency gains greatly simplifies the computation of the model. The reason is that it does not introduce additional state variables as would be the case with e.g. labor or capital adjustment costs at the firm level.

11. Note that as with expenditures on technology adoption, also ψ and φ are assumed to be paid in units of the final good and therefore they grow at the same rate as the rest of the economy.

Formally, after observing aggregate shocks but prior to the realization of idiosyncratic operational costs, an incumbent firm i of age a maximizes its discounted stream of all future profits ($\mathbb{V}_a(z_i, \mathcal{F})$) by choosing employment ($n_{i,a}$), capital ($k_{i,a}$) a technology adoption probability ($p_{i,a}$) and by deciding whether or not to remain in operation

$$\mathbb{V}_a(z_i, \mathcal{F}) = \max_{n_{i,a}, k_{i,a}, p_{i,a}} \int_{\varphi} \max \left[0, \tilde{\mathbb{V}}_a(z_i, \varphi, \mathcal{F}) \right] dH(\varphi), \quad (9)$$

where \mathcal{F} is the aggregate state and $\tilde{\mathbb{V}}_a(z_i, \varphi, \mathcal{F})$ is the value of a firm conditional on a particular draw of operation costs defined as

$$\begin{aligned} \tilde{\mathbb{V}}_a(z_i, \varphi, \mathcal{F}) = & y_{i,a} - Wn_{i,a} - Rk_{i,a} - T(p_{i,a}, \gamma_i) + \psi_a n_{i,a} - \varphi \quad (10) \\ & + \mathbb{E} \beta \frac{C}{C'} \mathbb{V}_{a+1}(z'_i, \mathcal{F}'), \end{aligned}$$

and where $y_{i,a}$ is firm-level production given by

$$y_{i,a} = Az_i \left(k_{i,a}^{\alpha} n_{i,a}^{1-\alpha} \right)^{\kappa}.$$

In the above, α and κ lie between 0 and 1 with the latter controlling the returns to scale in production, z_i is firm-specific productivity evolving according to (8) and A represents an aggregate total factor productivity shock. Unlike individual firm productivity, aggregate TFP affects all firms symmetrically and as such allows for common movements in firm productivity. It is assumed to follow an AR(1) process

$$\ln A' = (1 - \rho_A) \bar{A} + \rho_A \ln A + \varepsilon'_A,$$

where ρ_A is the autocorrelation coefficient, \bar{A} is the mean of $\ln A$ and where $\varepsilon_A \sim N(0, \sigma_A^2)$.

Given the perfectly competitive nature of the factor markets, the optimal firm-specific employment and capital decision boil down to factor prices being equal to marginal products (and efficiency gains from learning-by-doing)

$$R = \alpha \kappa y_{i,a} / k_{i,a}, \quad (11)$$

$$W = (1 - \alpha) \kappa y_{i,a} / n_{i,a} + \psi_a. \quad (12)$$

The point at which firms decide to shut down, $\tilde{\varphi}_{i,a}$, is defined by (10) equalling zero

$$0 = y_{i,a} - Wn_{i,a} - Rk_{i,a} - T(p_{i,a}, \gamma_i) - \psi_a n_{i,a} - \tilde{\varphi}_{i,a} + \mathbb{E} \beta C / C' \mathbb{V}_{a+1}(z'_i, \mathcal{F}'). \quad (13)$$

Finally, optimal technology adoption, both for incumbent firms and potential new entrants, equates the marginal costs to the marginal benefits of investing into newer technology vintages

$$\chi \eta \left(\frac{p_{i,a}}{\gamma_i} \right)^{\eta-1} = \frac{\partial \mathbb{E} \beta C / C' \mathbb{V}_{a+1}(z'_i, \mathcal{F}')}{\partial p_{i,a}}, \quad (14)$$

$$\chi\eta \left(\frac{p_e}{\bar{\gamma}} \right)^{\eta-1} = \theta \mathbb{V}_0(Z, \mathcal{F}). \quad (15)$$

In the above, p_e is the probability a potential entrant successfully adopts a newer technology and \mathbb{V}_0 represents the firm value of startups. Note that $\mathbb{V}_{a+1}(z'_i, \mathcal{F}')$ incorporates the endogenous evolution of firm-specific productivity as described by (8). To ease the exposition, formulas making this explicit are presented only in the Appendix.

2.1.4. The firm distribution, market clearing, balanced growth and equilibrium. We can now define $\omega_{j,a}$ as the beginning-of-period mass of firms of age a and productivity vintage z_j . In addition, let there be a fixed mass \bar{E} of potential startups attempting to enter the economy in each period. The distribution of firm masses across frontier technology vintages ($j \geq 0$) and firm ages ($a \geq 0$) can be described by

$$\begin{aligned} \omega_{0,0} &= \bar{E} p_e \theta, \\ \omega'_{0,a+1} &= \sum_a \sum_j \int^{\tilde{\varphi}_{j,a}} p_{j,a} \theta \omega_{j,a} dH(\varphi) \\ &\quad + \sum_a \sum_{j \leq a} \int^{\tilde{\varphi}_{0,a}} p_{0,a} (1 - \theta) \omega_{0,a} dH(\varphi) \\ \omega'_{j+1,a+1} &= \sum_a \sum_j \int^{\tilde{\varphi}_{j,a}} (1 - p_{j,a}) \omega_{j,a} dH(\varphi) \\ &\quad + \sum_a \sum_{j \leq a} \int^{\tilde{\varphi}_{j+1,a}} p_{j+1,a} (1 - \theta) \omega_{j+1,a} dH_t(\varphi) \end{aligned} \quad (16)$$

In the above, the first expression describes the mass of startups entering the economy in each period. The second expression gives the mass of firms older than one year, but which are nevertheless at the technological frontier. Such firms are either last period's surviving adopters with radical improvements from any part of the firm distribution (first summation) or last period's surviving frontier firms which managed to adopt the next younger vintage enabling them to keep up with technology growth (second summation). Finally, the third expression defines the mass of firms at productivity levels below the frontier. These businesses are either last period's surviving firms with productivity z_j which did not adopt newer technologies (first summation) or the mass of last period's surviving firms with productivity z_{j+1} which adopted the next younger technology vintage enabling them to keep up with technology growth (second line).

The labor and capital market clearing conditions and the aggregate resource constraint can be written, respectively, as

$$N = \sum_j \sum_a \int^{\tilde{\varphi}_{j,a}} \omega_{j,a} n_{j,a} dH(\varphi), \quad (17)$$

$$K = \sum_j \sum_a \int^{\tilde{\varphi}_{j,a}} \omega_{j,a} k_{j,a} dH(\varphi), \quad (18)$$

$$Y = C + I + \Xi, \quad (19)$$

where aggregate production $Y = \sum_a \sum_j \int^{\tilde{\varphi}_{j,a}} \omega_{j,a} (y_{j,a} + n_{j,a} \psi_a) dH(\varphi)$, which includes efficiency gains from learning-by-doing, is spent on consumption, investment into physical capital and aggregate costs $\Xi = \sum_a \sum_j \int^{\tilde{\varphi}_{j,a}} \omega_{j,a} (\varphi + T(p_{j,a}, \gamma_j)) dH(\varphi)$. The latter include operational costs and technology adoption expenditures. Aggregate profits are then defined as $\Pi = Y - WN - RK - \Xi$.

Note that the frontier technology is the only source of growth and therefore the economy fluctuates around the stochastic trend Z .¹² The aggregate state \mathcal{F} consists of not only the aggregate capital stock K and the two aggregate shocks Z and A , but also of the entire joint distribution of firm age and productivity vintages $\omega_{j,a}$. The reason for the latter is that aggregate factor demands depend on the distribution of workers and capital across firms with different productivity vintages and efficiencies of operation (learning-by-doing gains) which are age-dependent.

Finally, we are now ready to define the model's equilibrium as

- individual firms' policy rules for employment ($n_{i,a}$), capital ($k_{i,a}$), technology adoption probabilities ($p_{i,a}$) and firm exit ($\tilde{\varphi}_{i,a}$),
- potential entrants' policy rules for technology adoption probabilities (p_e),
- household's policy rules for aggregate consumption (C), employment (N) and investment (I),
- the wage (W), interest rate (R) and the distribution of firms across technology vintages and ages ($\omega_{j,a}$),

which satisfy firms' optimal labor, capital, technology adoption and exit conditions (11) to (14), the free entry condition (15), household's optimal labor supply and Euler equations (3) and (4), and the aggregate resource constraint (19), which clear the labor and capital markets (17) to (18) and which are consistent with the law of motion for the distribution of firms across frontier technology vintages and ages (16).

2.2. Firm-level uncertainty in the model

Having described the theoretical framework, this subsection builds intuition as to how and why technology growth is linked to firm-level uncertainty. Before doing so, however, let us describe the construction of the benchmark measure of firm-level uncertainty.

2.2.1. Measuring firm-level uncertainty. Throughout this paper the benchmark measure of uncertainty follows Bloom et al. (2018). These authors consider business-level production functions similar to those in this paper, namely $y_i = Az_i f(k_i, n_i)$,

12. Only firm-level and aggregate employment are stationary. All other variables can be stationaryized by dividing them with Z .

where TFP has an aggregate (A) and a firm-specific (z) component and where firms use capital and labor inputs to produce output. In addition, they assume that firm-specific productivity evolves according to

$$\ln z_{i,t} = \rho_z \ln z_{i,t-1} + \eta_{i,t}, \quad (20)$$

where ρ is a persistence parameter and where, importantly, the variance of productivity innovations η_i is allowed to be time-varying. In other words, firms face time-varying idiosyncratic uncertainty since the magnitude of their productivity innovations fluctuates over time. Moreover, this also implies that the dispersion of productivity innovations *across* individual firms is time-varying.¹³

Using the above setting, the authors then define uncertainty based on the cross-sectional dispersion of establishment-level total factor productivity shocks estimated from the following regression

$$\ln z_{i,t} = \mu_i + \rho \ln z_{i,t-1} + \lambda_t + \eta_{i,t}, \quad (21)$$

where μ_i is an establishment fixed effect, λ_t are time fixed effects and $\eta_{i,t}$ are the associated establishment-level TFP shocks.¹⁴ To construct this uncertainty measure, the authors use the Census panel of manufacturing establishments with annual data ranging from 1972 to 2009. In order to avoid compositional changes, they focus only on a balanced panel of establishments which are at least 25 years old.

In what follows, all references to uncertainty are understood to be regarding the above-described concept of the cross-sectional dispersion of establishment-level TFP shocks. In addition, all empirical exercises concerning uncertainty will be conducted using the data constructed by Bloom et al. (2018).

2.2.2. The nexus between growth and uncertainty. To understand why growth is linked to firm-level uncertainty, we can consider a simplified version of the firm-level productivity process described in (8). In particular, consider that technology adoption is purely exogenous and that all firms face the same constant probability of adopting a newer technology vintage, i.e. $p_{i,t} = p$. In addition, let us assume that all firms adopting newer technologies obtain the latest technology vintage, i.e. $\theta = 1$.

Under these assumptions, in a large enough cross-section of firms, a fraction p of businesses will have adopted the leading technology Z_t , while productivity of all other firms would have remained fixed. The evolution of firm-level productivity can then be described, on average, by the following law of motion

$$\ln z_{i,t} = (1 - p) \ln z_{i,t-1} + p \ln Z_t + v_{i,t}, \quad (22)$$

13. Note that many other alternative uncertainty proxies correlate highly with cross-sectional TFP dispersion (see e.g. Bloom et al., 2018).

14. Despite that this particular measure is constructed with establishment-level data, I will use the term establishment- and firm-level uncertainty interchangeably because the structural model does not distinguish between firms and establishments.

where $\mathbb{E}[v_{i,t}] = 0$ in the cross-section for every t . Notice, however, that by defining $\rho = 1 - p$ and $\lambda_t = p \ln Z_t$ we obtain the empirical regression (21) used to estimate firm-level uncertainty shocks.¹⁵ Recall that the latter is defined as the cross-sectional dispersion of the forecasting errors $v_{i,t}$. These errors can be written in terms of structural parameters as

$$v_{i,t} = \begin{cases} p\gamma_{i,t-1} - p\bar{Z} & \text{when firm } i \text{ does not adopt } Z_t, \\ (p-1)\gamma_{i,t-1} - (p-1)\bar{Z} & \text{when firm } i \text{ adopts } Z_t, \end{cases} \quad (23)$$

with their cross-sectional variance being

$$\text{var}[v_{i,t}] = (p^2 - p + 1)\sigma_\gamma^2 + p(1-p)\mu_\gamma^2 + p(1-p)\bar{Z}^2, \quad (24)$$

where $\mu_\gamma = \mathbb{E}[\gamma_{i,t-1}]$ and $\sigma_\gamma^2 = \text{var}[\gamma_{i,t-1}]$ are the cross-sectional mean and variance of last period's distribution of productivity gaps, respectively.¹⁶ The above expression shows that firm-level uncertainty is determined by three components: the distribution of (past) technology gaps, the speed of adopting the technological frontier and the growth rate of the frontier technology.

From (24) it is clear that time-variation in technology growth, $\bar{Z}_t = \bar{Z} + \varepsilon_t$, directly translates into firm-level uncertainty fluctuations.¹⁷ In particular, periods of high growth are associated with more uncertainty. The intuition behind this result is simple. When the frontier technology expands, firms face larger productivity gains if they successfully adopt the leading technology and relatively larger productivity losses if they do not.¹⁸

Moreover, notice that firm-level productivity must be described by *gradual* technology adoption in order for growth to be linked to uncertainty. In the extreme cases of no adoption (purely vintage technology) or full adoption (homogeneous technology) growth-driven uncertainty fluctuations disappear.¹⁹ Finally, the dependence of firm-level uncertainty on the past distribution of productivity gaps allows for persistent uncertainty increases even following only transitory changes in technology growth. Similarly, endogenous changes in the probability of technology adoption, p , will also induce richer dynamics in the full structural model.

15. To ease the exposition, establishment fixed effects are omitted here. However, all quantitative model exercises are based on (21) and thus include establishment fixed effects.

16. See the Appendix for a detailed derivation.

17. Similarly, it is possible to derive expressions for skewness and kurtosis of the firm-level TFP shocks. In particular, frontier growth enters as $\bar{Z}^3[1 - 2p]/\sqrt{p(1-p)}$ and $\bar{Z}^4[1 - 6p(1-p)]/[p(1-p)]$ in the skewness and kurtosis expressions, respectively. Therefore, the impact frontier growth has on these higher moments is quantitatively smaller and qualitatively ambiguous depending on the speed of technology adoption p .

18. This mechanism is similar to the growth option channel described in Bar-Ilan and Strange (1996) where higher uncertainty also has positive effects

19. It is straightforward to extend the model to include iid disturbances to firm-specific productivity such that $\text{var}[v_{i,t}] > 0$ even in the extreme cases of full or no technology adoption.

2.3. Calibration and model performance

The following paragraphs first describe the model's calibration and then evaluate its performance on dimensions not considered in the parametrization. In order to ease the exposition of the calibration strategy, I discuss the calibrated parameters in relation to specific targets even though individual parameters typically influence the behavior of the entire model. All parameter values and the associated targets are presented in Table 1.²⁰

In order to be consistent with the establishment-based uncertainty measure, the targeted moments are computed using U.S. establishment data taken from the Business Dynamics Statistics (BDS) for the available period of 1977-2014. Following the frequency of the BDS, the model period is therefore assumed to be one year.²¹

2.3.1. Calibration. Let us start by discussing the parameters pertaining directly to the household. The discount factor, β , is set to 0.97 corresponding to an annual interest rate of 3%. The disutility of labor, ν , is set such that the steady state wage rate is normalized to one. The depreciation rate of physical capital is set to 10%, consistent with values found in Cooper and Haltiwanger (2006).

The parameters governing the process of technology adoption include the normalization constant χ , the curvature parameter η and the probability of radical technology improvements θ . The normalization constant affects the level of technology adoption costs. Proxying adoption costs with expenditures spent on research and development, χ is set such that average adoption costs are 4% of output as in the data (see Akcigit and Kerr, 2018). The curvature parameter is set to 2 implying a 0.5 elasticity of the probability of adopting a new technology vintage with respect to the associated expenditures. This is consistent with estimates in Acemoglu et al. (2013). Finally, θ is set to 0.1 following Akcigit and Kerr (2018) who estimate that roughly 10% of all innovations open up new technologies.

Next, turning to the production function, $\alpha = 0.33$ while the returns to scale parameter is set to $\kappa = 0.8$ which falls within the values estimated in Basu and Fernald (1997).²² The efficiency gains from learning-by-doing, ψ_a , directly affect establishments' life-cycle growth. To ease the computational burden, I consider four age categories: startups, young (one to five years), medium-aged (six to ten years) and old establishments (11 years and more).²³ Efficiency gains are then set in order

20. The solution method follows Sedláček and Sterk (2017) and its description is deferred to the Appendix.

21. When computing business cycle statistics, the data is logged and HP filtered with a smoothing coefficient 100.

22. The Appendix provides robustness exercises with respect to this parameter which is quantitatively important for the resulting business cycle fluctuations of the economy.

23. While startups become young establishments in the next period (conditional on survival), young (medium-aged) establishments become medium-aged (old) establishments with a probability $\delta = 1/5$ ensuring an "expected duration" of five years within these age categories (conditional on survival).

TABLE 1. Model parameters

parameter	value	target/source
β	0.97	annual interest rate 3%
ν	0.17	wage normalization, $W_{ss} = 1$
δ_k	0.10	Cooper, Haltiwanger (2006)
χ	28.5	R&D intensity of 4%, Akcigit and Kerr (2016)
η	2	Acemoglu et al. 2013
θ	0.10	Akcigit and Kerr (2018)
κ	0.80	Basu and Fernald (1997)
ψ_s	-0.39	rel. average size of startups 52%, BDS
ψ_y	-0.22	rel. average size of young establishments 72%, BDS
ψ_m	-0.12	rel. average size of medium-age establishments 90%, BDS
ψ_o	0	normalization
μ_H	6.633	0 average paid operational costs, normalization
σ_H	4.896	average establishment exit rate of 11%, BDS
\bar{E}	8.883	firm mass of 1, normalization
\bar{Z}	0.018	average labor productivity growth, BEA
σ_Z	0.006	reallocation rate volatility 0.079, BDS
\bar{A}	2.664	average establishment size 20, BDS
ρ_A	0.625	real GDP autocorrelation 0.58, BEA
σ_A	0.004	real GDP volatility 0.019, BEA

Notes: The table reports model parameters and their respective targets or sources. Relative average size is defined as average size of the given firm category relative to the economy-wide average. The reallocation rate is the number of jobs created by startups and destroyed by exiting businesses relative to total employment.

to match average establishment size by age, relative to the economy's average (with efficiency gains of old establishments normalized to zero).²⁴ The distribution of the operational costs, H , controls the extent to which establishments exit the economy. It is assumed that H is logistic with mean μ_H and scaling parameter σ_H . The former is set such that the average establishment exit rate is 10%. The latter, which controls the dispersion of operational costs and which in turn shapes the relation between exit rates and firm-specific productivity, is pinned down by targeting the relative exit rate of startups and young establishments of 1.7 observed in the BDS data.

Finally, let us turn to the calibration of the two aggregate shocks, TFP (A) and frontier technology (Z). Aggregate TFP is assumed to follow an AR(1) process which is characterized by a mean TFP level \bar{A} , persistence parameter ρ_A and dispersion of TFP shocks σ_A . In contrast, frontier technology grows over time with a positive drift \bar{Z} and dispersion of frontier technology shocks σ_Z . These five parameters are chosen such that the model replicates the average establishment size, the persistence and volatility of real GDP, average labor productivity growth and the volatility of the reallocation rate related to entry and exit. The latter is defined as the number of jobs created by startups and destroyed by exiting businesses as a share of total employment.

2.3.2. Model performance. This subsection discusses the model's performance along several dimensions important for the quantitative results discussed next.

Table 2 shows that the model does well in replicating business cycle volatility and co-movement of (i) aggregate variables (top left panel), (ii) the reallocation process (top right panel), and (iii) the age distribution of establishments (bottom panels). Specifically, while the volatility of output was a calibration target, the model endogenously predicts that consumption and employment are less and that investment is more volatile than output and that all variables are highly correlated with output, consistent with the data.²⁵ Similarly, while the volatility of the employment reallocation rate Rn was a calibration target, the model does well in also replicating the volatility and business cycle co-movement of the number of establishments, the establishment reallocation rate and average establishment size. The exception is the employment reallocation rate which is more counter-cyclical in the model compared to the data. In addition, the model is also consistent with the magnitude of reallocation. In particular, in the data almost a third of all jobs are either created or destroyed each year and about 35% of this reallocation is because of establishment entry and exit. In the model, 30% of all jobs are created or destroyed every period and entry and exit account for 40% of this reallocation process. Finally, the model is also consistent with

24. In the data, on average about 30% of new establishments are created by existing firms. Such establishments may face different efficiency gains inherited from their parent firm. For simplicity, the calibration abstracts from such issues.

25. Volatility of model variables is computed using 1,000 simulations of 1,038 periods, where only the last 38 periods considered (consistent with the empirical sample length). Reported values are means of the respective standard deviations across the 1,000 simulations.

TABLE 2. Model performance

	<i>aggregates</i>				<i>reallocation process</i>			
	<i>Y</i>	<i>C</i>	<i>I</i>	<i>N</i>	Ω	<i>Rn</i>	<i>Re</i>	\bar{n}
	<i>standard deviation (in %)</i>							
data	1.9	1.4	8.3	1.8	1.2	7.9	4.3	1.8
model	1.8	1.0	6.7	1.3	1.5	7.1	2.2	1.7
	<i>correlation with output</i>							
data	1	0.92	0.84	0.89	0.65	-0.06	0.23	0.56
model	1	0.78	0.94	0.89	0.61	-0.66	0.19	0.14
	<i>establishment shares (%)</i>				<i>employment shares (%)</i>			
age (years)	0	1-5	6-10	11+	0	1-5	6-10	11+
data	10	32	19	39	5	23	17	54
model	10	27	17	46	5	19	15	61

Notes: The top panel shows the empirical and model-implied standard deviations (and correlations with output) of aggregate variables and variables related to the reallocation process. The data spans 1977-2014, consistent with the BDS sample. Y is real GDP, C is personal consumption expenditures, I is real gross private domestic investment and N is civilian employment, Ω is the number of establishments, Rn is the employment reallocation rate defined as the sum of job creation and destruction at entering and exiting establishments as a share of total employment, Re is the establishment reallocation rate defined as the sum of entering and exiting establishments as a share of all establishments and \bar{n} is average establishment size. All variables are logged and detrended with an HP filter with smoothing coefficient 100. The bottom panels of the table report the shares of establishments and employment in the respective age groups.

the empirical age distributions (of establishments and employment) with the share of old businesses being somewhat higher in the model compared to the data.

In addition to the above “macro” predictions, the model is also consistent with several “micro” patterns of the technology adoption process. In particular, Akcigit and Kerr (2018) document that small firms innovate relatively more than larger businesses showing that patents per employee decrease with firm size. Using the probability of successfully adopting newer technology vintages (p) to proxy for patents, this negative relationship between innovation and firm size also holds in the model. In addition, in the data innovation expenditures are positively correlated with firm productivity in the cross-section, but less so with productivity growth (see e.g. Klette and Kortum, 2004). In the model, lower firm-specific productivity impedes the innovation process, see (7). At the same time, less productive firms that successfully innovate experience relatively larger productivity gains, see (8). Therefore, as in the data, also in the model technology adoption expenditures are positively correlated with firm-specific productivity but essentially uncorrelated with productivity growth.²⁶

26. The data features a large share of firms reporting zero R&D expenditures (see e.g. Klette and Kortum, 2004). While this is not allowed in the model, which simplifies its computation, in the steady state 20% of the firms exhibit a technology adoption intensity of less than 0.5%.

2.4. Model results

This subsection presents the main model results. It begins by describing the model dynamics following the two aggregate shocks. It was analytically shown that firm-level uncertainty is linked to technology growth in a simplified version of the structural model. Importantly, this remains true in the full model. Moreover, unconditionally uncertainty fluctuations are counter-cyclical over the business cycle, as in the data. The next section presents empirical evidence for these model predictions using U.S. data.

2.4.1. Model dynamics following aggregate shocks. Figure 1 shows impulse response functions of several model variables in response to the aggregate TFP (A) and frontier technology (Z) shocks. All variables are expressed in percent deviations from their respective steady state growth paths.

Let us begin by inspecting the dynamics following a one-standard-deviation increase to aggregate TFP which affects all firms symmetrically. Because of consumption smoothing motives, the household ensures that part of this productivity increase gets invested (both into physical capital and greater technology adoption). The household's labor supply decision (3) shows that the aggregate wage rises along with consumption, undershooting the increase in firm productivity. Therefore, despite the temporary rise in labor costs, all firms are relatively more profitable and they expand leading to an increase in aggregate employment and output. For the same reason, job creation of startups rises and job destruction from firm exit drops. Finally, because all firms are affected symmetrically, average firm productivity mirrors the pattern of aggregate TFP without affecting the productivity dispersion across firms. Therefore, uncertainty is unaffected. In other words, the positive aggregate TFP shock generates a standard real business cycle.

The dynamics following shocks to the frontier technology are very different from those induced by an aggregate TFP shock. The main reason lies in the fact that changes in the frontier technology do not immediately affect the productivity of all firms, but rather they permeate through the economy only gradually. This is because it takes time and resources to adopt the new technology vintages. This asymmetry, which induces changes in the firm-level productivity distribution, generates a new force acting against the consumption smoothing channel. On impact, only the productivity of firms that have adopted the latest technology vintage rises. The productivity of all other businesses remains fixed until they also manage to successfully adopt newer vintages. Therefore, average firm productivity only gradually increases to its new long-run level. Importantly, uncertainty rises in response to the frontier technology shock. This is precisely because the frontier technology shock affects firms asymmetrically, widening the cross-sectional dispersion of firm-specific productivity shocks. Firms now face relatively larger productivity gains if they manage to successfully adopt the leading technology and relatively larger productivity losses if they do not. More details on the uncertainty response, including a description of its computation, are presented below.

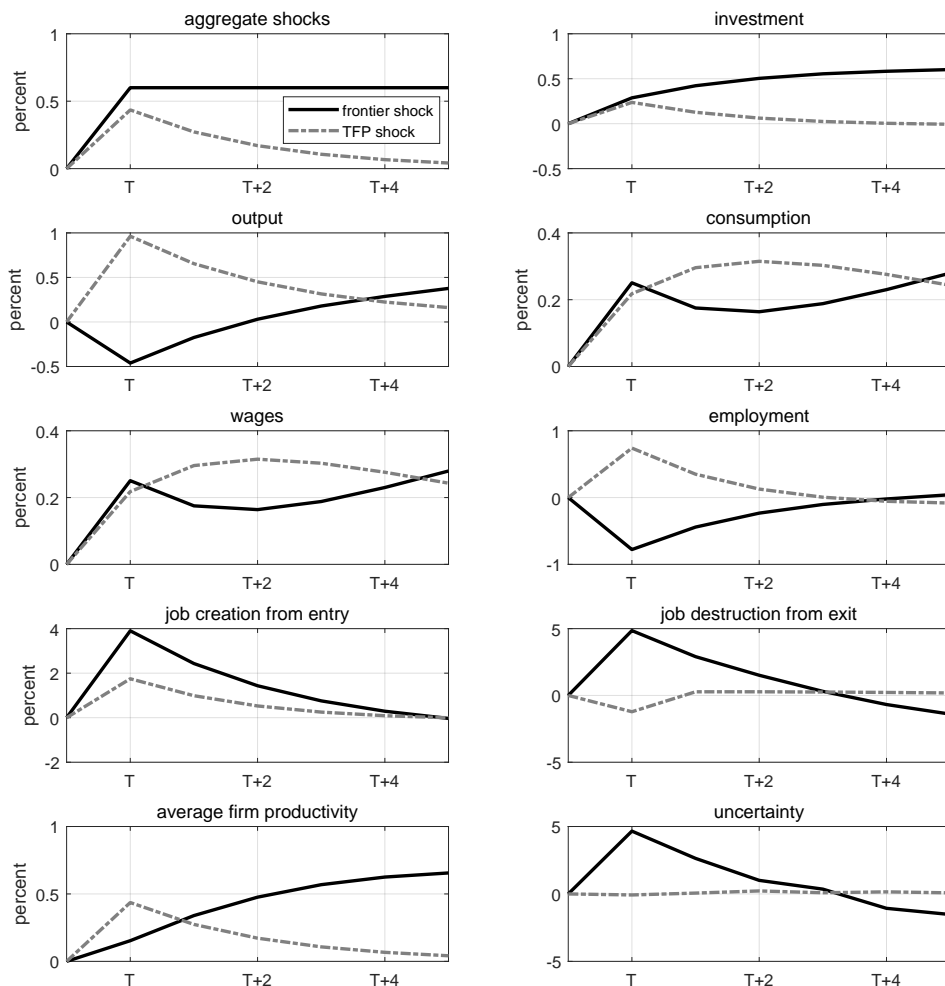


FIGURE 1. Impulse responses to aggregate shocks. Notes: Impulse response functions to positive one-standard-deviation shocks occurring in period T to aggregate TFP (A) and frontier technology (Z), respectively. All impulse responses are expressed in percentage deviations from the respective steady state growth paths.

However, not all the productivity gains are consumed. Instead, consumption smoothing motives ensure that part of these gains are invested and therefore consumption (and wages) rises by relatively less than productivity at the frontier. This, in turn, means that the economy experiences a simultaneous increase in job creation by firms utilizing the latest technology vintage and job destruction of all other firms for which productivity has remained fixed.

The precise nature of the creative destruction process is crucial for the aggregate employment response, which is in principle ambiguous. This process is described in detail below. However, under the present calibration, job destruction dominates initially and aggregate employment and output fall temporarily.

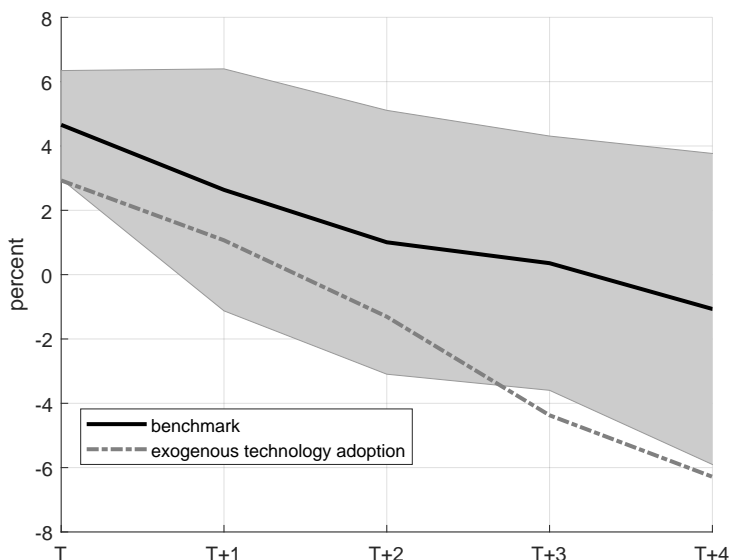


FIGURE 2. Uncertainty impulse responses to a positive frontier technology shock. Notes: Impulse response function of the standard deviation of firm-level TFP shocks, computed according to (21), to a positive one-standard-deviation shock to the frontier technology. The impulse response is generated by simulating a cross-section of firms 1,000 times for 1,038 periods. All exogenous shocks are set to zero except for a positive one-standard-deviation innovation to the frontier technology in period 1,001. The first 1,000 periods are discarded. The figure shows the average response and the respective 90% confidence bands (shaded areas) over the 1,000 simulations. “Fixed adoption” refers to a counterfactual scenario when firms’ probabilities of adopting newer technology vintages are held fixed at their respective steady state levels.

2.4.2. Model-predicted uncertainty fluctuations. Let us now turn back to analyzing the response of uncertainty to frontier technology shocks. The model-based uncertainty measure is exactly the same as is estimated in the data according to (21). Towards this end, the model is simulated 1,000 times for 1,038 periods with a cross-section of one million firms. All exogenous shocks are set to zero except for a positive one-standard-deviation innovation to the frontier technology in period 1,001. The first 1,000 periods are discarded. The remaining time-periods are used to estimate equation (21) and construct the uncertainty measure.²⁷ The impulse response function of firm-level uncertainty in Figure 1 is the average response over the 1,000 model simulations.

Figure 2 replicates the uncertainty impulse response, together with the associated 90% confidence bands stemming from the 1,000 simulations. While firm-level uncertainty increases on impact by about 5% and gradually converges back to its steady state, the confidence bands suggest that this increase is short-lived and lasts for at most one year.

27. Following Bloom et al. (2018) who construct their uncertainty measure using data on establishments with at least 25 years of observations, I restrict the sample in the model-counterpart to only old firms.

In addition to the benchmark response, the dash-dotted line depicts a counterfactual impulse response of uncertainty which excludes the effect endogenous technology adoption has on uncertainty, i.e. when the probabilities of updating firm-specific productivity are fixed to their respective steady state values p_i . Importantly, the magnitude of this effect is quantitatively important. On impact, an increase in technology adoption expenditures strengthens the uncertainty response by about 60%. In addition, it more than doubles the persistence of the uncertainty increase. While the average uncertainty response reverts back to its steady state after about four years in the benchmark economy, it dies out after about one and a half years in the counterfactual scenario.

The above shows that uncertainty increases in response to higher growth and that endogenous technology adoption is an important magnification and propagation channel in this regard. To the extent that measured uncertainty is at least partly growth-driven and R&D expenditures can be proxied by the firms' costs of adopting the leading technology, this result is consistent with Stein and Stone (2013). The authors find that uncertainty increases have a positive effect on R&D expenditures while hampering many other forms of investment.

What the above does not show, however, is how growth-driven uncertainty moves over the business cycle, i.e. unconditional on frontier technology shocks. Therefore, I again simulate the model, but this time allowing for both aggregate shocks to vary in line with their calibration.²⁸

The top two rows of Table 3 report correlations of uncertainty (cross-sectional dispersion of firm-level productivity shocks) with several business cycle indicators. The first two columns show that firm-level uncertainty is unconditionally *counter-cyclical*, albeit with correlations that fall somewhat short of that observed in the data. For instance, Bloom et al. (2018) report a correlation between GDP and uncertainty of about -0.46 indicating that firm-level uncertainty is likely fluctuating in a counter-cyclical fashion also for other (potentially exogenous) reasons not present in the current model.

However, the table also shows that skewness and kurtosis of firm-level TFP shocks is essentially a-cyclical. Importantly, this is also consistent with the data, as Bloom et al. (2018) find no statistically significant correlation between these moments and the business cycle.

Finally, the last column depicts the correlation of the moments of the firm-level productivity shock distribution with changes in the frontier technology. In this case, uncertainty is strongly *positively* related, since the dispersion of firm-level TFP shocks widens as the technological frontier expands. However, skewness and kurtosis are again a-cyclical.

To gauge these relationships in the data, I consider three proxies for frontier technology growth: i) real expenditures on R&D, ii) patent applications and iii) patent

28. Once again, the model is simulated 1,000 times for 1,038 period with a cross-section of on million firms. The first 1,000 periods are discarded. The generated cross-section of firm-level productivity shocks is then used to compute the model-implied uncertainty (dispersion), skewness and kurtosis measures.

TABLE 3. Correlation of cross-sectional moments of firm-level productivity shocks with...

	output	employment	technology
dispersion	-0.27 [-0.44, -0.10]	-0.24 [-0.41, -0.07]	0.37 [0.19, 0.55]
skewness	-0.01 [-0.18, 0.16]	-0.01 [-0.18, 0.16]	0.01 [-0.16, 0.18]
kurtosis	0.07 [-0.09, 0.23]	0.06 [-0.11, 0.23]	-0.08 [-0.24, 0.08]

Notes: correlation coefficients between firm-level uncertainty and business cycle indicators. Uncertainty is measured as the standard deviation of firm-level TFP shocks (computed according to (22)). “Technology” refers to the stochastically growing frontier technology Z_t . All the data is logged and HP filtered. The reported values are averages over 1,000 model simulations of length 1,038 periods in which the first 1,000 periods are discarded. The respective one-standard-deviation intervals (across the 1,000 model simulations) are reported in brackets.

grants.²⁹ The empirical correlation between firm-level uncertainty and proxies of frontier technology growth ranges between 0.15 and 0.25. In the case of skewness and kurtosis, there is no significant relationship to the above proxies.

Clearly, the above model-implied business cycle properties crucially depend on the process of creative destruction which induces the Schumpeterian downturn following a speed-up in frontier technology growth. The following subsection, therefore, analyses this process in more detail. The next section then provides empirical evidence on the model-suggested link between technology growth, creative destruction and uncertainty.

2.4.3. The process of creative destruction. The aggregate employment response depends on the relative mass of created and destroyed jobs. This, in turn, depends on the shares of expanding and contracting firms in the economy and on the magnitude of their respective employment changes.

Let us begin by investigating the firm-level employment changes implied by the optimal hiring decision (11)

$$\hat{n}_i = 1/(1 - \kappa)(\hat{z}_i - (1 - \kappa\alpha)\hat{W} - \kappa\alpha\hat{R}), \quad (25)$$

where “hats” indicate percentage deviations from the respective steady state trends.³⁰ The above equation shows that on the one hand all firms which fail to innovate, i.e. for which $\hat{z}_i = 0$, experience the same percentage change in employment irrespective of their size. In particular, the percentage drop in employment among shrinking

29. Focusing on patents as a measure of frontier technology growth follows Hall et al. (2001). The patent data is taken from U.S. Patent and Trademark Office (USPTO), R&D expenditures are taken from the Bureau of Economic Analysis and the uncertainty measure is taken from Bloom et al. (2018). All variables are logged and detrended using an HP filter with smoothing coefficient 100.

30. For clarity, (25) ignores efficiency gains from learning-by-doing. Taking them into account introduces differences in the percentage responses between firms of different ages. Equation (25) also highlights the importance of the returns to scale parameter κ for the quantitative results. The Appendix provides sensitivity tests with respect to this parameter.

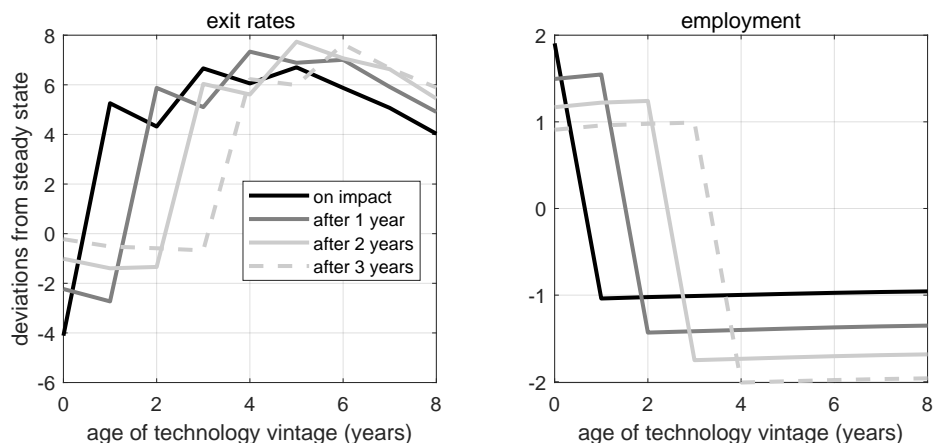


FIGURE 3. Impulse responses to a frontier technology shock: distributions. Notes: Impulse response functions (in percent) to a positive one-standard-deviation shock to the frontier technology (Z). The horizontal axis shows the age of the technology vintage of firm-specific productivity, i.e. $z_{j,t} = Z_{t-j}$. A zero year old technology refers to the frontier. The left panel depicts firm-specific exit rates deviations from their respective steady state values, as a function of the technology vintage of firm-specific productivity. The right panel depicts firm employment deviations from their respective steady state values, as a function of the technology vintage of firm-specific productivity. The different lines plot the impulse response on impact, 1, 2 and 3 years after the shock hits the economy, respectively.

firms is only a function of the change in factor prices and the production function parameters which are all common across firms. On the other hand, firms adopting newer technology vintages experience heterogeneous productivity gains depending on their prevailing productivity level. Specifically, firms far away from the frontier, but which nevertheless managed to adopt the leading technology, undergo relatively larger productivity (and thus size) increases compared to businesses with an initially higher productivity level. The shape of the firm size distribution is therefore key for the quantitative results.

In the calibrated model, which matches well the empirical firm size distribution, an average innovating firm creates more jobs than the average shrinking firm destroys. However, the overall impact on aggregate employment still depends on the relative shares of these two groups of firms. Figure 3 shows the impulse responses of the *distribution* of firm exit rates and employment levels as a function of firm-specific productivity ordered according to the age of the specific technology vintage, i.e. $z_{j,t} = Z_{t-j}$. Therefore, firms with a zero year old technology vintage possess the leading technology and older vintages are farther away from the growing frontier. The figure shows responses in the year of impact, 1, 2 and 3 years after the shock hit the economy. For clarity, the figure restricts the maximum age of a technology vintage to be 8 years.³¹ However, the vast majority (95%) of all firms are located

31. Vintages can be up to 30 years old as discussed in the solution method in the Appendix. The Appendix also shows more detailed impulse responses separated also by firm age.

within this range. In addition, the average exit rate and size of firms with technology vintages younger than 9 years are 8.5% and 18.3 employees, respectively. Therefore, Figure 3 summarizes changes in the (quantitatively) most important part of the firm distribution.

Upon impact, only firms at the technological frontier ($j = 0$) reap the benefits of improved technology. Such firms shut down relatively less often and they expand, compared to their steady state. All other firms, which now find themselves facing higher labor costs, shut down relatively more often and contract. The permanent nature of technology shocks means that one year after the improvement in frontier technology, also firms with a one year old technology vintage are relatively more productive and they expand and shut down less often. Similar logic applies to responses in later years, highlighting how the benefits of technology shocks only gradually permeate through the economy.

In the calibrated model, the mass of shrinking firms initially dominates that of expanding businesses. Therefore, following a positive frontier technology shock, the economy undergoes a temporary Schumpeterian downturn.

3. Creative destruction and uncertainty in the data

The structural model presents a theory in which positive growth options increase uncertainty at the firm-level and, at the same time, spur a process of creative destruction generating a temporary Schumpeterian downturn. This section presents empirical evidence in support of these model predictions and quantifies to what extent observed uncertainty fluctuations in the data are growth-driven.

First, using detailed industry-level data, I show that uncertainty is indeed positively correlated with productivity growth. Second, using structural VARs and aggregated firm-level data in the U.S., I document that a positive technology shock triggers a Schumpeterian downturn and an increase in firm-level uncertainty as predicted by the model.

3.1. Technology growth and uncertainty at the industry-level

The model predicts that, while being counter-cyclical, uncertainty is *positively* related to productivity growth. Moreover, in the model this channel operates via endogenous technology adoption. In this subsection, I test these model predictions using industry data at the 4-digit level. Specifically, I consider the following regression

$$\sigma_{s,t} = \mu_s + \lambda_t + \alpha \Delta \ln VA_{s,t} + \beta \Delta \ln TFP_{s,t} + \gamma RD_s + \delta X_{s,t} + \varepsilon_{s,t}, \quad (26)$$

where s indicates an industry and t a time period. In (26), $\sigma_{s,t}$ is industry-specific uncertainty, μ_s is an industry fixed effect, λ_t is a time fixed effect, $\Delta \ln VA_{s,t}$ is the growth rate of industry-specific value added, $\Delta \ln TFP_{s,t}$ is the growth rate of industry-specific TFP, RD_s is the average industry-specific R&D intensity and $X_{s,t}$ include other control variables.

TABLE 4. Regression results: industry-level uncertainty

	I	II	III
ΔVA	-0.05***	-0.12***	-0.12***
ΔTFP		0.21**	0.22**
RD			0.01**

Notes: Regression results for (26). All regressions include industry and time fixed effects and control variables described in the main text. The three columns report results for different versions of (26) which, in addition to value added growth, sequentially include TFP growth and average R&D intensity within industries. Standard errors are clustered by industries.

The industry level uncertainty data is taken from Bloom et al. (2018), the R&D data is taken from Bloom et al. (2013) and all the other industry-level data is taken from the NBER-CES Manufacturing Industry Database. Uncertainty is measured again as the cross-sectional dispersion in establishment-level TFP shocks, TFP is the “5-factor-TFP” measure and R&D intensity is given by R&D expenditures relative to sales. Finally, $X_{s,t}$ includes a range of industry-specific control variables: growth in total payroll, material costs, capital investment and inventories and the average capital-labor ratio, share of energy costs in total costs and share of structures in total capital. The sample periods differ somewhat across industries, but overall they range between 1971 and 2009. After dropping industries with missing observations we are left with 77 industries and a total of 2,831 time-industry observations.

Table 4 shows the regression results. The top row indicates that uncertainty is robustly counter-cyclical as has been shown in many other studies. However, consistent with the model predictions, the second row shows that conditionally on value added growth uncertainty is in fact *positively* related to TFP growth. Moreover, uncertainty is shown to be positively related to R&D intensity of the given industries. Using R&D expenditures as a proxy for the speed of technology adoption, this is consistent with the model which can be directly seen from (24).³²

While one should be cautious in interpreting the above results as causal relationship, they do reflect conditional correlations which are consistent with the model predictions. The following subsection takes a step in the direction of identifying causal relationships in the data. In particular, it estimates technology *shocks* and investigates the responses of uncertainty, and other variables, to these structural innovations.

3.2. Technology growth and uncertainty in the aggregate

Let us now focus on the causal predictions of the model regarding technology shocks, the creative destruction process and uncertainty. Towards this end, I estimate a series of structural vector autoregressions (VARs) with long run restrictions as

32. As long as $p < 0.5$, higher values of p raise uncertainty (and its sensitivity to TFP shocks).

in e.g. Blanchard and Quah (1989), Gali (1999).³³ Consistent with the structural model, the identification is based on assuming that only technology shocks determine productivity in the long-run.

The estimation uses bi-variate VARs where the data vector is given by $Y_t = (\Delta a_t, x_t)'$, with Δa_t being productivity growth and x_t being the variable of interest.³⁴ Productivity is measured by output per hour in the non-farm business sector and x_t includes: firm-level uncertainty, job creation by new establishments, job destruction by exiting establishments and aggregate employment.³⁵ All variables are in annual frequencies, ranging from 1977-2014. In addition, following Fernald (2007), who documents that low-frequency movements in productivity impair the identification of technology shocks, the estimation allows for break points in the intercepts. Finally, all VAR specifications are estimated with two lags.³⁶

Figure 4 shows the responses of our variables of interest, both in the data and in the model. In the data, and as predicted by the model, a positive technology shock spurs a process of creative destruction associated with a simultaneous increase in job creation and job destruction. In the aggregate, this leads to a temporary Schumpeterian downturn with employment falling for several periods. Finally, the bottom right panel shows that, consistent with the model, also in the data firm-level uncertainty rises in response to a positive technology shock.

Therefore, the figure shows that the model predicts empirically plausible impulse responses, both qualitatively but also quantitatively. Only the model response of uncertainty is somewhat on the high side of its empirical counterpart. However, recall from Figure 2 that the confidence bands around the model-predicted uncertainty response suggest that it is different from zero only in the first year following the shock.

3.3. To what extent is uncertainty growth-driven in the data?

Finally, let us quantify to what extent observed fluctuations in firm-level uncertainty are growth-driven in the data. A variance decomposition of the forecast errors suggests that about 27% of the observed fluctuations in uncertainty are driven by technology shocks alone.

In addition, it is possible to gage which episodes of heightened uncertainty were predominantly driven by growth options and which were by driven by other factors. In what follows, I focus on the four NBER recessions in the sample, which were

33. The Appendix shows that very similar results are obtained with an alternative estimation strategy based on local projections following Jorda (2005) and using technology shocks estimated by Basu et al. (2013).

34. The identified technology shocks are nevertheless very similar with correlation coefficients around 0.8 across the different VARs.

35. Uncertainty is again measured as the cross-sectional variation in TFP shocks taken from Bloom et al. (2018), job creation and destruction data are taken from the Business Dynamics Statistics (both entering in logs) and employment is the growth rate of civilian employment taken from the Bureau of Labor Statistics.

36. The Appendix provides further details on the estimation procedure as well as several robustness checks.

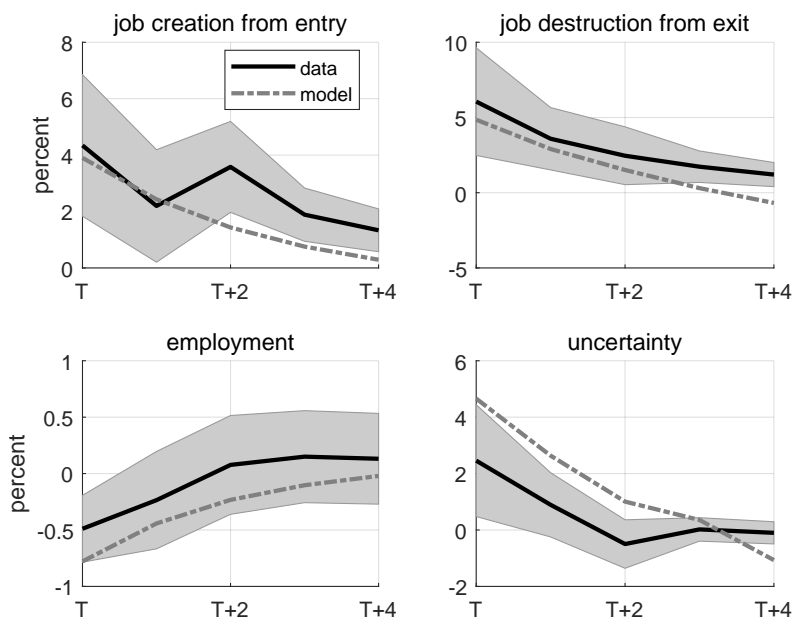


FIGURE 4. Aggregate responses to a positive technology shock: data and model. Notes: Impulse response functions to a positive one-standard-deviation technology shock in the “model” and the “data”. Job creation from entry and destruction from exit are taken from the Business Dynamics Statistics (both entering in logs), employment is the growth rate of civilian employment taken from the Bureau of Labor Statistics and R&D expenditures are measured as real R&D expenditures as a share in real GDP taken from the Bureau of Economic Analysis. Shaded areas depict one-standard-deviation confidence intervals.

all associated with a cyclical increase in uncertainty. To measure the latter, I de-mean the data using the estimated time-varying intercepts from the structural VAR.³⁷ The uncertainty run-ups are always measured from the respective trough prior to the recession up until one year after the official end of the downturn. In order to measure the “growth-driven” component of these uncertainty spikes, I use the estimated structural VAR to forecast uncertainty while allowing for only the identified technology shock to vary and fixing the second shock to zero.

Figure 5 plots the above-described uncertainty spikes around the four NBER recessions in the sample, together with the respective growth-driven components implied by the identified technology shocks alone. While technology shocks explain, on average, about a quarter of the uncertainty fluctuations, the patterns differ substantially across recessions.

37. Note that the level of uncertainty was highest in the Great Recession, but the time-varying intercepts soak up some of this (trend) increase which occurred already prior to the Great Recession. See the Appendix for more details on the estimated VARs.

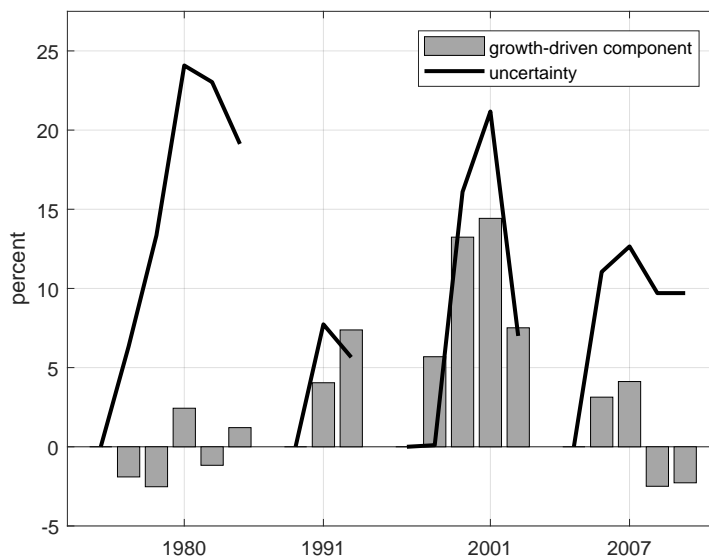


FIGURE 5. Uncertainty spikes around recessions and their growth-driven components. Notes: Cyclical uncertainty increases around NBER recessions measured as deviations from the respective prior troughs. “Growth-driven component” is based on the estimated structural VAR with the identified technology shocks being the only sources of variation. Both time-series are first demeaned using the time-varying intercepts from the estimated structural VAR.

Specifically, uncertainty increases during the milder downturn in 1991 and the run up towards the bursting of the “dot-com” bubble in 2001 were predominantly growth-driven. In these cases, growth-driven uncertainty accounts for about two thirds of the overall increase at its peak. On the contrary, the strong uncertainty increases during the double-dip recession in the early 1980’s and the Great Recession had little to do with growth. In fact, during both these downturns the growth-driven component contributed negatively to the overall increase in uncertainty.³⁸

The above empirical evidence therefore supports the predictions of the structural model that technology growth, creative destruction and firm-level uncertainty fluctuations go hand-in-hand. Not only does uncertainty respond to technology shocks in the data, it does so in a quantitatively important way.

4. Conclusion

This paper provides a theory and empirical evidence on how growth-options impact firm-level uncertainty and in turn the aggregate economy. The structural model of firm growth via endogenous technology adoption suggests that increases in

38. This is in line with Ludvigson et al. (2017), who argue that the Great Recession increase in uncertainty was primarily related to financial uncertainty.

technology growth go hand-in-hand with a process of creative destruction and with increases in firm-level uncertainty. Such growth-driven uncertainty fluctuations are therefore counter-cyclical, but co-move positively with technology growth. The model predictions are shown to hold in U.S. data not only qualitatively, but also quantitatively.

While the results show that technology growth is likely an important driver of uncertainty fluctuations, especially in certain periods, they also highlight the role of other factors in shaping uncertainty fluctuations. In particular, the Great Recession seems to be a period in which uncertainty increased dramatically for reasons unrelated to technology growth. In order to understand the aggregate implications of uncertainty fluctuations and the possible inefficiencies and associated policy implications related to them, it is important to further strive to understand the different sources of uncertainty variation.

References

- Abel, Andrew (1983). “Optimal Investment Under Uncertainty.” *American Economic Review*, 73(1), 228–233.
- Acemoglu, Daron, Ufuk Akcigit, Nicholas Bloom, and William Kerr (2013). “Innovation, Reallocation and Growth.” NBER Working Paper No. 18993.
- Aghion, Philippe and Peter Howitt (1994). “Growth and Unemployment.” *Review of Economic Studies*, 61, 477–494.
- Akcigit, Ufuk and William Kerr (2018). “Growth through Heterogeneous Innovations.” *Journal of Political Economy*, 126(4), 1374–1443.
- Atkin, David, Azam Chaudhry, Shamyala Chaudry, Amit Khandelwal, and Eric Verhoogen (2017). “Organizational Barriers to Technology Adoption: Evidence from Soccer-Ball Producers in Pakistan.” *Quarterly Journal of Economics*, 132(3), 1101–1164.
- Bachmann, Rüdiger and Christian Bayer (2014). “Investment Dispersion and the Business Cycle.” *American Economic Review*, 104(4), 1392–1416.
- Bachmann, Rüdiger and Giuseppe Moscarini (2012). “Business Cycles and Endogenous Uncertainty.” mimeo, University of Notre Dame.
- Baker, Scott and Nicholas Bloom (2013). “Does Uncertainty Reduce Growth? Using Disasters as Natural Experiments.” NBER Working Paper 19475.
- Bar-Ilan, Avner and William Strange (1996). “Investment Lags.” *American Economic Review*, 86(3), 610–622.
- Basu, Susanto and John Fernald (1997). “Returns to Scale in US Production: Estimates and Implications.” *Journal of Political Economy*, 105(2), 249–283.
- Basu, Susanto, John Fernald, Jonas Fisher, and Miles Kimball (2013). “Sector-Specific Technical Change.” Mimeo, Federal Reserve Bank of San Francisco.
- Basu, Susanto, John Fernald, and Miles Kimball (2006). “Are Technology Improvements Contractionary?” *American Economic Review*, 96(5), 1418–1447.
- Berger, David and Joseph Vavra (forthcoming). “Shocks vs Responsiveness: What Drives Time-Varying Dispersion?” *Journal of Political Economy*.
- Blanchard, Olivier and Daniel Quah (1989). “The Dynamic Effects of Aggregate Demand and Supply Disturbances.” *American Economic Review*, 79(4), 655–673.
- Bloom, Nicholas (2009). “The Impact of Uncertainty Shocks.” *Econometrica*, 77(3), 623–685.
- Bloom, Nicholas, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen Terry (2018). “Really Uncertain Business Cycles.” *Econometrica*, 86(3), 1031–1065.
- Bloom, Nicholas, Mark Schankerman, and John van Reenen (2013). “Identifying Technology Spillovers and Product Market Rivalry.” *Econometrica*, 81(4), 1347–1393.
- Boedo, Moscoso, Ryan Decker, and Pablo D’Erasmus (2016). “Market Exposure and Endogenous Firm Volatility over the Business Cycle.” *American Economic Journal: Macroeconomics*, 8(1), 148–198.

- Caballero, Ricardo and Mohamad Hammour (1996). “On the Timing and Efficiency of Creative Destruction.” *Quarterly Journal of Economics*, 111, 805–852.
- Canova, Fabio, David Lopez-Salido, and Claudio Michelacci (2013). “The Ins and Outs of Unemployment: an Analysis Conditional on Technology Shocks.” *The Economic Journal*, 123, 515–539.
- Comin, Diego and Mark Gertler (2006). “Medium-Term Business Cycles.” *American Economic Review*, 96(3), 532–551.
- Comin, Diego and Bart Hobijn (2007). “Implementing Technology.” NBER Working Paper 12886.
- Cooper, Russel and John Haltiwanger (2006). “On the Nature of Capital Adjustment Costs.” *Review of Economic Studies*, 73(3), 611–633.
- Fernald, John (2007). “Trend Breaks, Long-run Restrictions, and Contractionary Technology Improvements.” *Journal of Monetary Economics*, 54, 2467–2485.
- Fisher, Jonas (2006). “The Dynamic Effects of Neutral and Investment Specific Technology Shocks.” *Journal of Political Economy*, 114(3), 413–451.
- Foster, Lucia, John Haltiwanger, and Chad Syverson (2016). “The Slow Growth of New Plants: Learning about Demand?” *Economica*, 83(329), 91–129.
- Francis, Neville and Valerie Ramey (2005). “Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited.” *Journal of Monetary Economics*.
- Gali, Jordi (1999). “Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?” *American Economic Review*, 89(1), 249–271.
- Gourio, Francois (2014). “Financial Distress and Endogenous Uncertainty.” mimeo, Federal Reserve Bank of Chicago.
- Hall, Bronwyn, Adam Jaffe, and Manuel Trajtenberg (2001). “The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools.” NBER Working Paper 8498.
- Hansen, Gary (1985). “Indivisible Labor and the Business Cycle.” *Journal of Monetary Economics*, 16(3), 309–327.
- Hartman, Richard (1972). “The Effects of Price and Cost Uncertainty on Investment.” *Journal of Economic Theory*, 5(2), 258–266.
- Ilut, Cosmin, Matthias Kehrig, and Martin Schneider (2018). “Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News.” *Journal of Political Economy*, 126(5), 2011–2071.
- Jack, Kelsey, Paulina Oliva, Christopher Severen, Elizabeth Walker, and Samuel Bell (2015). “Technology Adoption Under Uncertainty: Take-Up and Subsequent Investment in Zambia.” NBER Working Paper 21414.
- Jorda, Oscar (2005). “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review*, 95(1), 161–182.
- Jurado, Kyle, Sydney Ludvigson, and Serena Ng (2015). “Measuring Uncertainty.” *American Economic Review*, 105(3), 1177–1216.
- Klette, Tor Jakob and Samuel Kortum (2004). “Innovating Firms and Aggregate Innovation.” *Journal of Political Economy*, 112(5), 986–1018.

- Koren, Miklos and Silvana Tenreyro (2007). “Volatility and Development.” *Quarterly Journal of Economics*, 122(1), 243–287.
- Lester, Robert, Michael Pries, and Eric Sims (2014). “Volatility and Welfare.” *Journal of Economic Dynamics and Control*, 38, 17–36.
- Lientz, Bennet and Kathryn Rea (1998). *Breakthrough Technology Project Management*. Academic Press.
- Lopez-Salido, David and Claudio Michelacci (2007). “Technology Shocks and Job Flows.” *Review of Economic Studies*, 74, 1195–1227.
- Ludvigson, Sydney, Sai Ma, and Serean Ng (2017). “Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?” mimeo, New York University.
- Mortensen, Dale and Christopher Pissarides (1998). “Technological Progress, Job Creation, and Job Destruction.” *Review of Economic Dynamics*, 1, 733–753.
- Oi, Walter (1961). “The Desirability of Price Instability under Perfect Competition.” *Econometrica*, 29(1), 58–64.
- Orlik, Anna and Laura Veldkamp (2013). “Understanding Uncertainty Shocks and the Role of Black Swans.” mimeo, Board of Governors of the Federal Reserve System.
- Ramey, Gary and Valerie Ramey (1991). “Technology Commitment and the Cost of Economic Fluctuations.” NBER Working Paper 3755.
- Rogerson, Richard (1988). “Indivisible Labor, Lotteries, and Equilibrium.” *Journal of Monetary Economics*, 21(1), 3–16.
- Sedláček, Petr and Vincent Sterk (2017). “The Growth Potential of Startups over the Business Cycle.” *American Economic Review*, 107(10), 3182–3210.
- Stein, Jeremy (1997). “Waves of Creative Destruction: Firm-Specific Learning-by-Doing and the Dynamics of Innovation.” *Review of Economic Studies*, 64(2), 265–288.
- Stein, Luke and Elizabeth Stone (2013). “The Effect of Uncertainty on Investment, Hiring and R&D: Causal Evidence from Equity Options.” mimeo, Arizona State University.