Creative Destruction and Uncertainty*

Petr Sedláček

University of Bonn†

July 14, 2017

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 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, shocks to technology growth explain 1/4 of the cyclical variation in uncertainty, and up to 2/3 around the "dot-com" bubble.

JEL codes: D22, E32, D80

Keywords: creative destruction, uncertainty, business cycles, growth

*I thank Christian Bayer, Nicholas Bloom, 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.

†Address: Department of Economics, University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany. Email: sedlacek@uni-bonn.de

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 businesses 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 at 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.

[1] See Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). The Appendix shows that the results are robust to alternative firm-level uncertainty measures.

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.

To see whether this channel is also empirically relevant I test the model's predictions in the data. First, while uncertainty is counter-cyclical, the model predicts that uncertainty co-moves positively with changes in technology growth. I use information on patent applications and R&D expenditures to proxy for technology growth and show that, as the model predicts, the correlation between these proxies and uncertainty fluctuations is positive and statistically significant.

Next, 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).[2] The identifying assumption, consistent with the structural model, is that only technology shocks affect productivity in the long-run.

The estimated structural VARs 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 percent of the business cycle variation in uncertainty is driven by technology

[2]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, Fernald, Fisher, and Kimball (2013).

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, Ludvigson, and Ng, 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 (2015); Boedo, Decker, and D'Erasmo (2016); Berger and Vavra (2016).[3] Ludvigson, Ma, and Ng (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 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, Fernald, and Kimball, 2006; Lopez-Salido and Michelacci, 2007;

[3] Oi (1961); Hartman (1972); Abel (1983); Bar-Ilan and Strange (1996) entertain the possibility that increases in uncertainty come with positive effects and Ilut, Kehrig, and Schneider (2016) argue that counter-cyclical cross-sectional volatility is a natural result of concave hiring rules.

Canova, Lopez-Salido, and Michelacci, 2013).[4] 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.

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 are uncertainty fluctuations 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, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014), 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 is described in detail in Section 2.2.

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

[4] 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.

incumbent and entering firms.

2.1.1 Household preferences and choices

The representative household chooses consumption, C_t, and supplies labor, N_t, on a perfectly competitive labor market. 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

$$\ln C_t - \upsilon N_t,$$

where $\upsilon > 0$ 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

$$C_t = N_t W_t + \Pi_t, \tag{1}$$

which states that total income stems from employment (with W_t being the competitive wage rate) and from the ownership of firms, where Π_t are aggregate profits. This total income is entirely spent on consumption. The resulting optimal labor supply condition is given by

$$W_t = \upsilon C_t \tag{2}$$

2.1.2 Technology adoption, firm-specific productivity and growth

It is assumed that the frontier technology evolves exogenously according to the follow process

$$\ln Z_t = \overline{Z} + \ln Z_{t-1} + \epsilon_{Z,t}, \tag{3}$$

where $\overline{Z} > 0$ is a positive drift term and $\epsilon_{Z,t}$ are iid innovations distributed according to a Normal distribution with zero mean and standard deviation σ_Z. Individual firms are characterized by a particular vintage of technology $z_{j,t} = Z_{t-j}$ with $j \geq 0$. For future reference, it is useful to also define the technology gap, $\gamma_{j,t} = \ln z_{j,t} - \ln Z_t$.

Because the frontier is growing over time, an individual firm which fails to adopt newer technology vintages will experience a gradual decline in relative productivity. At some point the firm will become so unproductive that it will no longer be profitable to remain in operation. To prevent this, incumbent firms can undertake investment in order to improve their prevailing productivity levels. These investments are interpreted broadly, not only as costs of adopting a well-defined technology. For instance, they also represent the costs of identifying best practices and attempts at personalizing and implementing such practices at a specific firm. Therefore, the outcome of such an investment is inherently uncertain.[5]

Following Klette and Kortum (2004) a firm investing r units of the final good has a probability p of adopting a newer technology vintage, where

$$p = \left(\frac{r}{\chi}\right)^{\frac{1}{\eta}} \gamma^{1-\frac{1}{\eta}}.$$

In the above expression, χ is a scaling factor, γ is the above-defined technology gap (or "stock of knowledge") and $1/\eta$ is a curvature parameter. The associated cost function can be written as

$$R(p, \gamma) = \chi\gamma \left(\frac{p}{\gamma}\right)^{\eta}. \tag{4}$$

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, 2016). 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, letting j denote age of a particular vintage of technology, firm-specific productivity evolves

[5] 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.

according to

$$\ln z_{j,t} \rightarrow \begin{cases} \ln z_{j+1,t+1} & \text{with probability } 1 - p_{j,t}, \\ \ln z_{j,t+1} & \text{with probability } p_{j,t}(1-\theta), \\ \ln Z_{t+1} & \text{with probability } p_{j,t}\theta. \end{cases} \tag{5}$$

Finally, it is assumed that the process of technology adoption is the same for potential startups as it is for incumbents 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, $\overline{\gamma}$. Startups are assumed to enter the economy only if they manage to adopt the latest technology vintage.[6]

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, n, as the only production factor 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, Haltiwanger, and Syverson, 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.

[6]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.

In particular, it is assumed that firms accrue efficiency gains, ψ, through learning-by-doing (as in e.g. Stein, 1997). 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. This makes more mature businesses, which do not necessarily operate cutting-edge technologies, competitive and able to fend off more innovative newcomers.[7]

Finally, in addition to variable costs, firms must also pay stochastic fixed costs of operation, ϕ. Firms endogenously shut down when the realization of the fixed cost is too high rendering them unprofitable.[8]

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 ($\mathbb{V}_a(z_{i,t}, \mathcal{F}_t)$) by choosing employment ($n_{i,a,t}$), a technology adoption probability ($p_{i,a,t}$) and by deciding whether or not to remain in operation

$$\mathbb{V}_a(z_{i,t}, \mathcal{F}_t) = \max_{n_{i,a,t}, p_{i,a,t}} \int_\phi \max\left[0, \widetilde{\mathbb{V}}_a(z_{i,t}, \phi, \mathcal{F}_t)\right] dH_t(\phi), \tag{6}$$

where $\widetilde{\mathbb{V}}_a(z_{i,t}, \phi, \mathcal{F}_t)$ is the value of a firm conditional on a particular draw of operation costs defined as

$$\widetilde{\mathbb{V}}_a(z_{i,t}, \phi, \mathcal{F}_t) = y_{i,a,t} - W_t n_{i,a,t} - R(p_{i,a,t}, \gamma_{i,t}) + \psi_{a,t} n_{i,a,t} - \phi + \mathbb{E}_t \beta_{t+1} \mathbb{V}_{a+1}(z_{i,t+1}, \mathcal{F}_{t+1}), \tag{7}$$

where $\mathcal{F}_t$ is the aggregate state, $\beta_t = \beta C_t / C_{t+1}$ is the household's stochastic discount factor with $\beta < 1$ and $y_{i,a,t}$ is firm-level production given by

$$y_{i,a,t} = A_t z_{i,t} n_{i,a,t}^{\alpha}.$$

In the above, α controls the returns to scale in production, $z_{i,t}$ is firm-specific productivity evolving according to (5) and A_t represents an aggregate total factor productivity shock.

[7] 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 adjustment costs where the entire firm size distribution becomes a state variable.

[8] 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.

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

$$\ln A_t = \rho_A \ln A_{t-1} + \epsilon_{A,t},$$

where ρ_A is the autocorrelation coefficient and $\epsilon_{A,t} \sim N(0, \sigma_A^2)$.

Given the perfectly competitive nature of the labor market, the optimal firm-specific employment decision boils down to workers' wages being equal to the marginal product of labor and the firm's efficiency gains from learning-by-doing

$$W_t = \alpha y_{i,a,t}/n_{i,a,t} + \psi_{a,t}. \tag{8}$$

The point at which firms decide to shut down, $\widetilde{\phi}_{i,a,t}$, is defined by (7) equaling zero

$$0 = y_{i,a,t} - W_t n_{i,a,t} - R(p_{i,a,t}, \gamma_{i,t}) - \psi_{a,t} n_{i,a,t} - \widetilde{\phi}_{i,a,t} + \mathbb{E}_t \beta_t \mathbb{V}_{a+1}(z_{i,t+1}, \mathcal{F}_{t+1}).$$

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

$$\chi\eta \left(\frac{p_{i,a,t}}{\gamma_{i,t}}\right)^{\eta-1} = \frac{\partial \mathbb{E}_t \beta_t \mathbb{V}_{a+1}(z_{i,t+1}, \mathcal{F}_{t+1})}{\partial p_{i,a,t}},$$

$$\chi\eta \left(\frac{p_{e,t}}{\overline{\gamma}_t}\right)^{\eta-1} = \theta \mathbb{V}_0(Z_t, \mathcal{F}_t).$$

In the above, $p_{e,t}$ is the probability a potential entrant successfully adopts a newer technology and $\mathbb{V}_0$ represents the firm value of startups.[9]

2.1.4 The firm distribution, market clearing and balanced growth

Letting j denote the age of a particular vintage of technology, $z_{j,t} = Z_{t-j}$, we can then define $\omega_{j,a,t}$ as the beginning-of-period mass of firms of age a and productivity $z_{j,t}$. In addition, let there be a fixed mass $\overline{E}$ of potential startups attempting to enter the economy in each period. The mass of startups entering the economy in each period is given by

[9]Note that in the above expressions $\mathbb{V}_{a+1}(z_{i,t+1}, \mathcal{F}_{t+1})$ incorporates the endogenous evolution of firm-specific productivity as described by (5). To ease the exposition, formulas making this explicit are presented only in the Appendix.

$$\omega_{0,0,t} = \overline{E} p_{e,t} \theta.$$

The mass of firms older than one year, but nevertheless at the frontier, is given by

$$\omega_{0,a+1,t+1} = \begin{array}{ll} \sum_j \sum_a \int^{\widetilde{\phi}_{j,a,t}} p_{j,a,t} \theta \omega_{j,a,t} dH_t(\phi) & j = 0, 1, 2, ..., \\ & a = 0, 1, 2, ..., \\ \int^{\widetilde{\phi}_{0,a,t}} p_{0,a,t} (1-\theta) \omega_{0,a,t} dH_t(\phi) & j \leq a, \end{array}$$

where firms at the technological frontier are either last period's surviving adopters with radical improvements from any part of the firm distribution (top line) or last period's surviving frontier firms which managed to adopt the next younger vintage enabling them to keep up with technology growth (bottom line). The distribution of firm masses at productivity levels below the frontier is given by

$$\omega_{j+1,a+1,t+1} = \sum_j \sum_a \left[\begin{array}{l} \int^{\widetilde{\phi}_{j,a,t}} (1 - p_{j,a,t}) \omega_{j,a,t} dH_t(\phi) + \\ \\ \int^{\widetilde{\phi}_{j+1,a,t}} p_{j+1,a,t} (1-\theta) \omega_{j+1,a,t} dH_t(\phi) \end{array} \right] \begin{array}{l} j = 0, 1, 2, ..., \\ a = 0, 1, 2, ..., \\ j \leq a, \end{array}$$

where the mass of firms with productivity z_{j+1} is given by the mass of last period's surviving firms with productivity z_j which did not adopt newer technologies (top line) and the mass of last period's surviving firms with productivity z_j which adopted the next younger technology vintage enabling them to keep up with technology growth (bottom line).

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

$$\begin{array}{ll} N_t = \sum_j \sum_a \int^{\widetilde{\phi}_{j,a,t}} \omega_{j,a,t} n_{j,a,t} dH_t(\phi), & j = 0, 1, 2, ... \\ & a = 0, 1, 2, ... \\ Y_t = C_t + \Xi_t, & j \leq a \end{array}$$

where aggregate production, $Y_t = \sum_a \sum_j \int^{\widetilde{\phi}_{j,a,t}} \omega_{j,a,t} (y_{j,a,t} + n_{j,a,t} \psi_{a,t}) dH_t(\phi)$, which includes efficiency gains from learning-by-doing is spent on consumption and aggregate costs $\Xi_t = \sum_a \sum_j \int^{\widetilde{\phi}_{j,a,t}} \omega_{j,a,t} \left(\phi + R(p_{j,a,t}, \gamma_{j,t}) + \right) dH_t(\phi)$. The latter include operational costs and technology adoption expenditures. Aggregate profits are then defined as $\Pi_t = Y_t - W_t N_t - \Xi_t$.

Note that the frontier technology is the only source of growth and therefore the economy fluctuates around the stochastic trend Z_t. The following aggregate and firm-specific variables are stationary

$$\frac{C_t}{Z_t}, \frac{W_t}{Z_t}, \frac{R(p_{e,t}, \overline{\gamma}_t)}{Z_t}, \frac{\Pi_t}{Z_t}, N_t, \frac{z_{i,t}}{Z_t}, \frac{\widetilde{\phi}_{i,a,t}}{Z_t}, \frac{\psi_{a,t}}{Z_t}, \frac{\Phi_{i,a,t}}{Z_t}, \frac{R(p_{i,a,t}, \gamma_{i,t})}{Z_t}, n_{i,a,t}, \quad \forall i, a.$$

Finally, the aggregate state $\mathcal{F}_t$ consists of not only the two aggregate shocks, but also of the entire joint distribution of firm-specific productivity and employment levels. The reason for the latter is the perfectly competitive labor market where the aggregate wage rate depends on the distribution of workers across the heterogeneous firms.

2.2 Firm-level uncertainty in the model

This subsection builds intuition as to how and why technology growth is linked to firm-level uncertainty. Before doing so, however, let us describe how the benchmark measure of firm-level uncertainty is constructed.

2.2.1 Measuring firm-level uncertainty

Throughout this paper the benchmark measure of uncertainty follows Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). The authors define uncertainty based on the cross-sectional dispersion of establishment-level total factor productivity shocks estimated from the following regression

$$\ln z_{i,t} = \mu_i + \rho \ln z_{i,t-1} + \lambda_t + \eta_{i,t}, \tag{9}$$

where $\ln z_{i,t}$ is the log of estimated establishment-level TFP, ρ is a persistence parameter, μ_i is an establishment fixed effect, λ_t are time fixed effects and $\eta_{i,t}$ are establishment-level TFP shocks.[10] To construct this uncertainty measure, the authors use the Census panel

[10] 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.

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, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014).

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 (5). 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

$$\ln z_{i,t} = (1-p)\ln z_{i,t-1} + p \ln Z_t + \upsilon_{i,t}, \tag{10}$$

where $\mathbb{E}[\upsilon_{i,t}] = 0$ 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 (9) used to estimate firm-level uncertainty shocks.[11] Recall that the latter is defined as the cross-sectional dispersion of the forecasting errors $\upsilon_{i,t}$. These errors can be be written in terms of structural parameters as

[11] To ease the exposition, establishment fixed effects are omitted here. However, all quantitative model exercises are based on (9) and thus include establishment fixed effects.

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

with their cross-sectional variance being

$$\mathrm{var}[\upsilon_{i,t}] = p(1+p)\sigma_\gamma^2 + p(1-p)\mu_\gamma^2 + p(1-p)\overline{Z}^2, \tag{11}$$

where $\mu_\gamma = \mathbb{E}[\gamma_{i,t-1}]$ and $\sigma_\gamma^2 = \mathrm{var}[\gamma_{i,t-1}]$ are the cross-sectional mean and variance of last period's distribution of productivity gaps, respectively.[12] The above expression shows that firm-level uncertainty is determined by three components: the distribution of (past) technology gaps, the probability of adopting the technological frontier and the growth rate of the frontier technology.

From (11) it is clear that time-variation in technology growth, $\overline{Z}_t = \overline{Z} + \epsilon_t$, directly translates into firm-level uncertainty fluctuations. 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.[13].

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.[14] 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.

[12]See the Appendix for a detailed derivation.

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

[14]It is straightforward to extend the model to include iid disturbances to firm-specific productivity such that $\mathrm{var}[\upsilon_{i,t}] > 0$ 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.[15]

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 1977-2013. Following the frequency of the BDS, the model period is therefore assumed to be one year.[16]

Finally, while firms in this model do not conduct research and development (R&D), but rather spend resources on adopting (exogenous) technology vintages, it is nevertheless interesting to confront some of the model predictions with data on R&D. I do so by proxying R&D expenditures in the model with technology adoption costs of firms which have adopted the leading technology vintage.

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 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 average probability of technology adoption. This, in turn, influences the persistence of establishment-level productivity as can be seen from (9). Therefore, χ is set such that the resulting establishment-level productivity persistence is consistent with that estimated in Foster, Haltiwanger, and Syverson

[15] The solution method follows Sedláček and Sterk (2016) and its description is deferred to the Appendix for brevity.

[16] When computing business cycle statistics, the data is logged and HP filtered with a smoothing coefficient 6.23.

Table 1: Model parameters

	parameter	value	target/source
β	discount factor	0.97	annual interest rate 3%
υ	disutility of worker labor	3.633	wage normalization, $W_{ss} = 1$
χ	R&D normalizing constant	0.280	firm productivity persistence 0.79, Foster et al. 2008
η	R&D cost curvature	2	patent elasticity w.r.t R&D 0.5, Acemoglu et al. 2013
θ	probability of radical improvements	0.1	Akcigit and Kerr (2016)
α	returns to scale	0.670	labor share 66%
ψ_s	learning-by-doing efficiency gains, startups	−0.545	rel. average size of startups 52%, BDS
ψ_y	learning-by-doing efficiency gains, young establishments	−0.344	rel. average size of young establishments 72%, BDS
ψ_m	learning-by-doing efficiency gains, medium-aged establishments	−0.195	rel. average size of medium-age establishments 90%, BDS
ψ_o	learning-by-doing efficiency gains, old establishments	0	normalization
μ_H	operational cost mean	0.093	0 average paid operational costs, normalization
σ_H	operational cost distribution, scale	0.238	average establishment exit rate of 11%, BDS
$\overline{E}$	mass of potential entrants	8.883	firm mass of 1, normalization
$\overline{Z}$	frontier technology drift	0.016	average patent application growth, USPTO
σ_Z	frontier technology shocks, standard deviation	0.011	corr(R&D,Y)=0.21, BEA
ρ_A	aggregate TFP shock, persistence	0.788	real GDP autocorrelation 0.79, BEA
σ_A	aggregate TFP shock, volatility	0.010	real GDP volatility 0.013, 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.

(2008).[17] 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. Proxying R&D expenditures with the adoption costs of firms adopting the leading technology, this is consistent with estimates in Acemoglu, Akcigit, Bloom, and Kerr (2013). Finally, θ is set to 0.1 following Akcigit and Kerr (2016) who estimate that roughly 10 percent of all innovations open up new technologies.

Next, returns to scale are set to a standard value of $\alpha = 0.67$.[18] 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).[19] Efficiency gains are then set in order to match average establishment size by age, relative to the economy's average (with efficiency gains of old establishments normalized to zero).[20] 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 aggregate paid operational costs are zero and the latter is set to match an average establishment exit rate of 11 percent observed in the BDS data.

Finally, aggregate TFP shocks, which are common across all establishments, are assumed to follow an AR(1) process. The persistence and the standard deviation of aggregate TFP shocks are set such that the model replicates the autocorrelation and volatility of aggregate output. Frontier technology shocks, on the other hand, are the only source of growth and are assumed to follow a random walk with a positive drift, $\overline{Z}$. The latter is set to 1.8 percent which is the average observed growth rate of labor productivity over

[17] The Appendix shows that similar results are obtained when targeting a lower persistence of establishment-level productivity.

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

[19] 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).

[20] In the data, on average about 30 percent 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.

the sample period. The standard deviation of innovations to frontier technology growth is set to match the correlation between aggregate output and R&D expenditures. The motivation for this last target will become clear in the next subsection which shows that while frontier technology and aggregate TFP shocks have the same qualitative effects on R&D expenditures, they have opposite effects on aggregate output. The relative size of the two aggregate shocks, which is important for the quantitative results, then determines to what extent R&D expenditures co-move with aggregate output over the business cycle.

2.3.2 Model performance

This subsection discusses the model's performance along several dimensions important for the quantitative results discussed next. In particular, it shows that the model predicts a realistic distribution of establishment and associated rates of job reallocation. In addition, the model also features empirically plausible R&D patterns at the micro-level even though only aggregate R&D patterns were used in the calibration process.

First, while average establishment sizes by age were a target of the calibration, the establishment and employment distribution were not. Nevertheless, the model does well in replicating these distributions (Table 2). The reason behind this fact is that the model correctly predicts a negative relation between average establishment exit rates and age. In particular, while the average exit rate of young (old) establishments is 17 (7) percent in the data, it is 15 (9) percent in the model.

Second, associated with the process of establishment churn is the overall degree of job reallocation. In the data, each year about 31 percent of all jobs are either created or destroyed. In the model, the overall reallocation rate is 24 percent. In addition, a large part of this reallocation process is due to the entry and exit of establishments. In the data, jobs created by entrants and destroyed by exiting establishments account for 12 percent of employment. In the model this fraction is equal to 15 percent.

Third, a crucial feature of the model is the technology adoption process. Again, using firms which adopt the leading technology as a proxy for innovating firms we find that the model does well in capturing the process of innovation at the micro-level. In partic-

Table 2: Establishment and employment distributions (in %)

	establishment age			
	0	1-5	6-10	11+
	establishment shares			
data	10.3	32.1	19.4	38.2
model	10.8	30.8	19.6	38.8
	employment shares			
data	5.3	23.2	17.4	54.1
model	5.6	21.5	18.0	54.9

Notes: The table reports the shares (in percent) of establishments and employment in the group of startups, young (1 to 5 years), medium-aged (6 to 10 years) and old (11 years and over) establishments as a share of all establishments and employment. Data is taken from the BDS.

ular, Akcigit and Kerr (2016) document that small firms innovate relatively more than larger businesses showing that patents per employee decrease with firm size. Using the probability of successful R&D (p) to proxy for patents, this negative relationship between innovation and firm size also holds in the model. In addition, in the data R&D 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 (4). At the same time, less productive firms that successfully innovate experience relatively larger productivity gains, see (5). Therefore, as in the data, also in the model R&D expenditures are positively correlated with firm-specific productivity but essentially uncorrelated with productivity growth.[21]

[21] 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 it's computation, in the steady state 20 percent of the firms exhibit an R&D intensity of less than 0.5 percent. The average R&D intensity is 3 percent in the model, which is close to the empirical value of 4 percent reported in Akcigit and Kerr (2016).

2.4 Model results

It was analytically shown that firm-level uncertainty is linked to technology growth in a simplified version of the structural model. This subsection begins by showing that also within the full structural model firm-level uncertainty increases in response to technology shocks. In addition, this subsection describes the economy's dynamics following the two structural shocks and explains why uncertainty is counter-cyclical, as in the data. The next section documents that these model-predicted dynamics are empirically plausible and it quantifies to what extent observed uncertainty fluctuations are growth-driven in U.S. data.

2.4.1 Model-predicted uncertainty fluctuations

In order to investigate the response of firm-level uncertainty to technology shocks, I construct exactly the same measure of uncertainty in the model as is estimated in the data according to (9). This is done by simulating the model $1,000$ times for $1,040$ periods with a cross-section of $100,000$ firms. All exogenous shocks are set to zero except for a positive one-standard-deviation innvoation to the frontier technology in period $1,001$. The first $1,000$ periods are discarded. The remaining time-periods are used to estimate equation (9) and construct the uncertainty measure.[22]

Figure 1 plots the resulting impulse response function of firm-level uncertainty averaged over the $1,000$ model simulations together with the associated one-standard-deviation confidence bands. It documents that firm-level uncertainty increases significantly and persistently following the positive one-time shock to the frontier technology. On impact, uncertainty rises above its steady state by about 5 percent and gradually converges back. The confidence bands suggest that this increase is relatively short-lived and lasts for only about two years.

In addition to the benchmark response, the dashed blue line depicts a counterfactual

[22]Following Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) 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.

Figure 1: Uncertainty impulse responses to a positive frontier technology shock

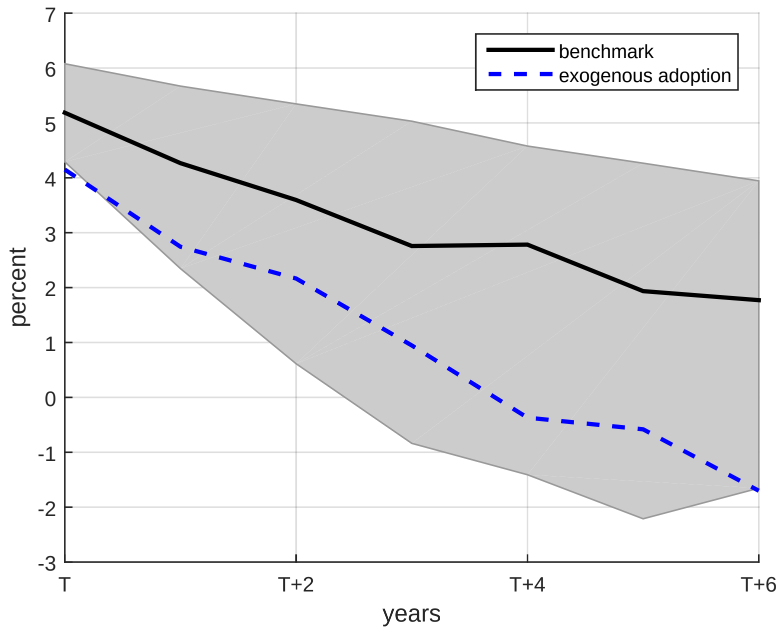

Notes: Impulse response function of the standard deviation of firm-level TFP shocks (computed according to (9)) to a positive one-standard-deviation shock to the frontier technology. The impulse response is generated by simulating a cross-section of firms $1,000$ times for $1,040$ periods. All exogenous 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 percent confidence bands (shaded areas) over the $1,000$ simulations. " Fixed adoption" referes to a counterfactual scenario when firms' probabilities of adopting newer technology vintages are held fixed at their respective steady state levels.

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. The difference between the benchmark and the counterfactual uncertainty response highlights that endogenous technology adoption serves to exacerbate the rise in uncertainty following a technology shock.

Moreover, the magnitude of this effect is quantitatively important. On impact an increase in technology adoption expenditures strengthens the uncertainty response by about 25 percent. 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 ten years in the benchmark economy, it dies out after about four 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 is how growth-driven uncertainty moves over the business cycle, i.e. unconditional on technology shocks. Towards this end, the next subsections discuss the economy's dynamics in response to the two aggregate shocks and evaluate the cyclical properties of growth-driven uncertainty.

2.4.2 Business cycle dynamics following aggregate TFP shocks

Figure 2 characterizes the response of the economy to a positive one-standard-deviation shock to aggregate TFP (A). All variables in the figure are expressed in percent deviations from their respective steady state growth rates.[23]

By construction, the persistent increase in aggregate TFP affects all firms symmetrically raising average firm productivity. This symmetry is also the reason why firm-level uncertainty is unaffected by fluctuations in aggregte TFP.[24] Consumption smoothing motives on the side of the household ensure that part of this productivity increase gets channeled into greater technology adoption (both among incumbent firms and potential new startups).

The household's labor supply decision (2) 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

[23] In this case, average firm productivity and consumption grow together with the frontier technology, while all other variables in the figure are stationary.

[24] The impulse response of firm-level uncertainty to an aggregate TFP shock is provided in the Appendix.

leading to an increase in aggregate employment. For the same reason, job creation of startups rises and job destruction from firm exit drops. In other words, the positive aggregate TFP shock generates a standard real business cycle.

2.4.3 Business cycle dynamics following frontier technology shocks

The aggregate dynamics following shocks to the frontier technology are very different from those induced by an aggregate TFP shock, see Figure 3.[25] 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 as it takes time and resources to adopt 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. This, however, requires time and resources and average firm productivity only gradually rises to it's new long-run level. Consumption smoothing motives ensure that part of the productivity gains are "re-invested" into technology adoption and therefore consumption (and wages) rises by relatively less than average firm productivity. This means, however, 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. Under the present calibration, job destruction dominates initially and aggregate employment falls temporarily. The following paragraphs describe in detail how exactly creative destruction helps shape aggregate fluctuations.

[25] As before, all variables are expressed in percent deviations from their respective steady state trends.

Figure 2: Impulse responses to an aggregate TFP shock

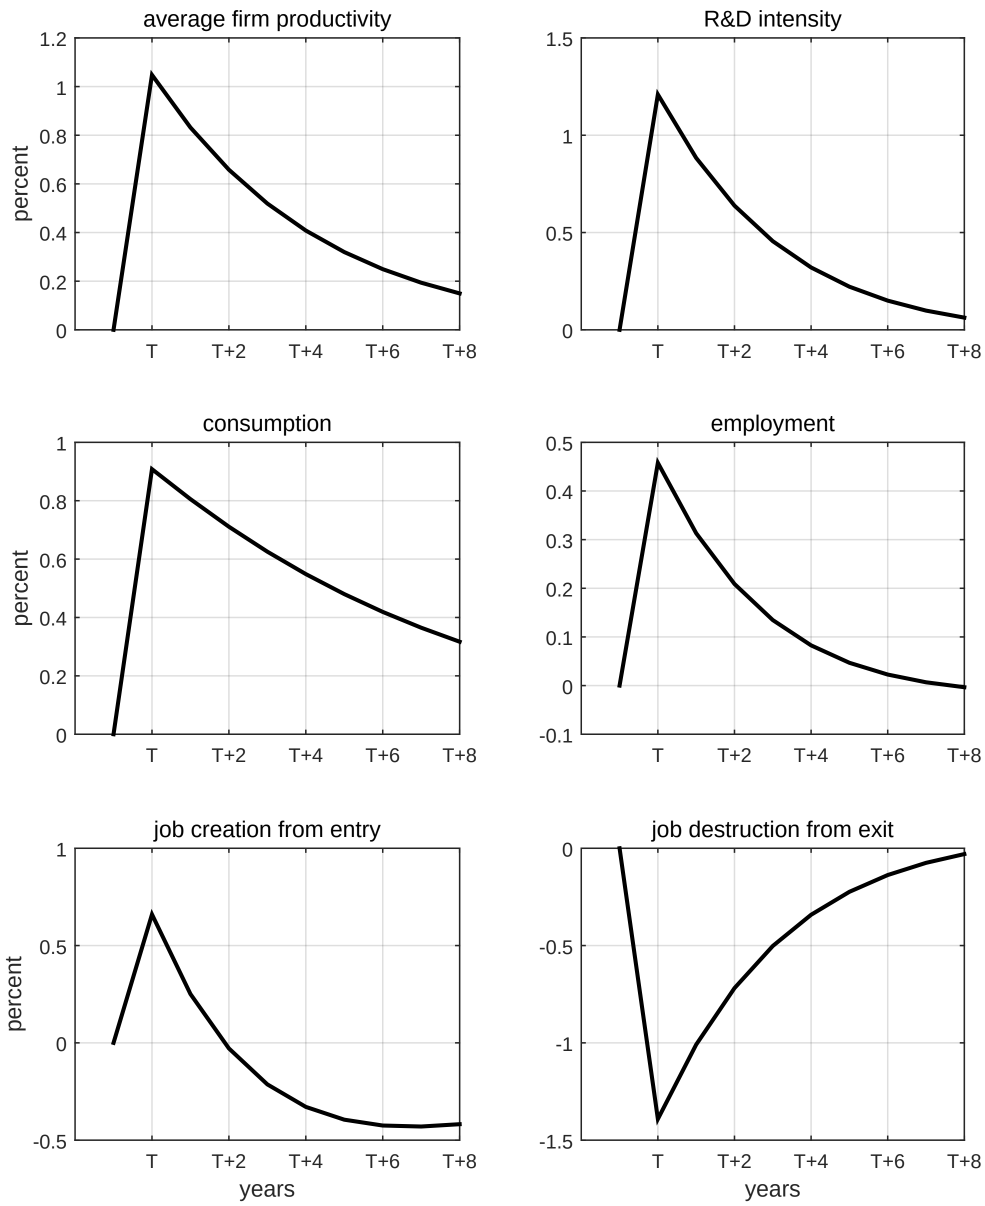

Notes: Impulse response functions to a positive one-standard-deviation shock to aggregate TFP (A) occurring in period T. "R&D intensity" is defined as the share of all technology adoption expenditures at firms adopting the leading technology relative to aggregate output. All other variables are self-explanatory. All impulse responses are expressed in percentage deviations from the respective steady state trends (average firm productivity and consumption grow at rate $\overline{Z}$ in steady state).

Figure 3: Impulse responses to a frontier technology shock: aggregates

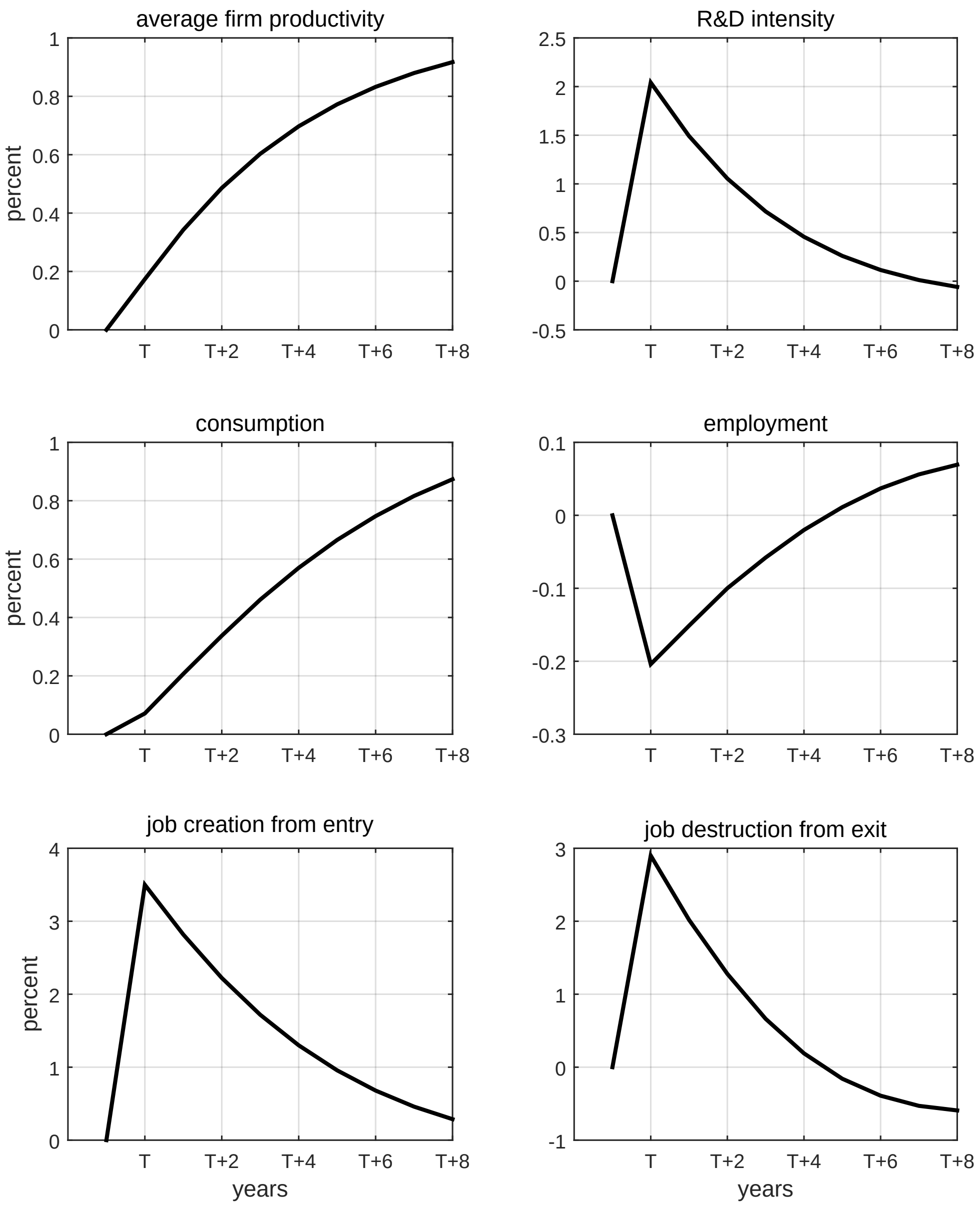

Notes: Impulse response functions to a positive one-standard-deviation shock to the frontier technology (Z) occurring in period T. "R&D intensity" is defined as the share of all technology adoption expenditures at firms adopting the leading technology relative to aggregate output. All other variables are self-explanatory. All impulse responses are expressed in percentage deviations from the respective steady state trends (average firm productivity and consumption grow at rate $\overline{Z}$ in steady state).

2.4.4 The process of creative destruction

Clearly, 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 (8)

$$\widehat{n}_{i,t} = 1/(1-\alpha)(\widehat{z}_{i,t} - \widehat{W}_t), \tag{12}$$

where "hats" indicate percentage deviations from the respective steady state trends.[26] The above equation shows that on the one hand all firms which fail to innovate, i.e. for which $\widehat{z}_{i,t} = 0$, experience the same percentage drop in employment irrespective of their size. In particular, the percentage drop in employment among shrinking firms is proportional to the wage increase. On the other hand, innovating firms experience heterogeneous productivity gains depending on their prevailing productivity level. Specifically, firms far away from the frontier, but which nevertheless managed to radically innovate, undergo relatively larger productivity (and thus size) increases compared to innovating 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 4 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. $z_{j,t} = Z_{t-j}$. Specifically, firms with a zero year old technology vintage posses the leading technology and older vintages are farther away from the growing frontier. The figure shows responses in the year of

[26] For clarity, (12) ignores efficiency gains from learning-by-doing. Taking them into account introduces differences in the percentage responses between firms of different ages. Equation (12) also highlights the importance of the returns to scale parameter α for the quantitative results. The Appendix provides sensitivity tests with respect to this parameter.

Figure 4: Impulse responses to a frontier technology shock: distributions

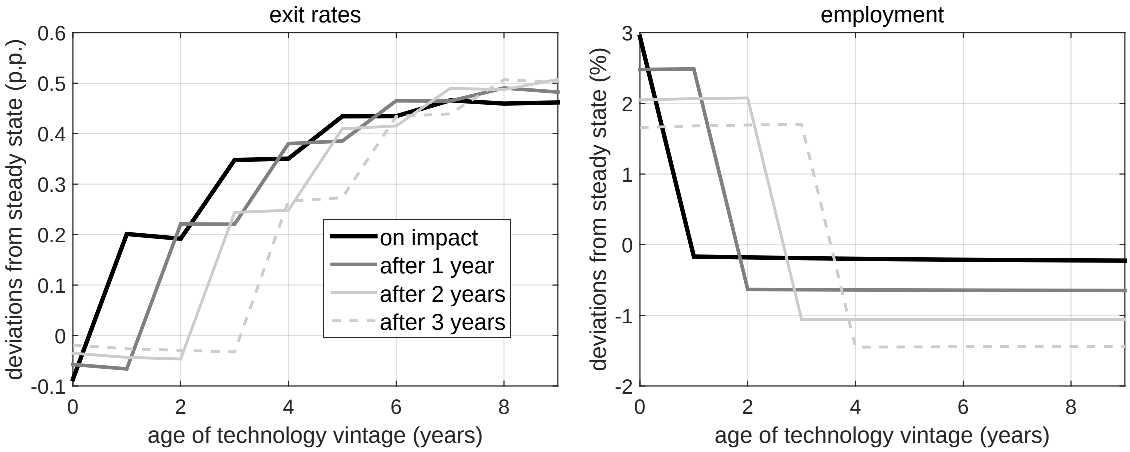

Notes: Impulse response functions to a positive one-standard-deviation shock to the frontier technology (Z). The horizontal axis shows the age of the technology vintage of firm-specific productivity, i.e. $z_{j,t} = Z_{t-j}$. A zero year old technology refers to the frontier. The left panel depicts firm-specific exit rates in percentage point deviations from their respective steady state values, as a function of the technology vintage of firm-specific productivity. The right panel depicts firm employment in percent 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.

impact, 1, 2 and 3 years after the shock hit the economy.[27]

Upon impact, only firms at the technological frontier (i.e. with $j = 0$) 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.

Therefore, the aggregate employment response crucially depends on the speed of the

[27]For clarity, the figure restricts the maximum age of a technology vintage to be 9 years. In the full model, vintages can be up to 30 years old as discussed in the solution method in the Appendix.

technology adoption. In the calibrated model, which is consistent with several empirical R&D patterns at the firm-level and the average persistence of firm-level productivity, the mass of shrinking firms initially dominates that of expanding businesses and the economy undergoes a temporary Schumpeterian downturn.

2.4.5 Business cycle fluctuations of uncertainty

The above paragraphs make clear that conditional on technology shocks, uncertainty is *counter-cyclical*. However, it may still be the case that aggregate TFP shocks dominate the business cycle and uncertainty is acyclical unconditionally. Therefore, I simulate the model allowing for both aggregate shocks to vary in line with their calibration.[28]

Table 3 reports correlations of uncertainty with several business cycle indicators. The first two columns show that firm-level uncertainty is unconditionally *counter-cyclical*, albeit with correlations that falls somewhat short of that observed in the data. For instance, Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) 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. The last column of Table 3 highlights the growth-driven nature of uncertainty which is positively correlated with technology growth.

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.

[28] Once again, the model is simulated $1,000$ times for $1,040$ period with a cross-section of $100,000$ firms. The first $1,000$ periods are discarded. The generated data is used to estimate uncertainty according to (9) which I then correlate with the cyclical components of aggregate output and employment.

Table 3: Correlation of firm-level uncertainty in model with...

output	employment	technology
-0.33	-0.23	0.54
$[-0.50, -0.16]$	$[-0.42, -0.05]$	$[0.40, 0.68]$

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 (10)). "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 $1,040$ 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.

3.1 Cyclical pattern of firm-level uncertainty

Let us begin by the cyclical fluctuations of firm-level uncertainty. As is well known, uncertainty is robustly counter-cyclical and this is also true in the structural model. However, the growth option channel of uncertainty fluctuations of the model predicts that, while being counter-cyclical, uncertainty is *positively* correlated with technology growth.

In order to investigate this prediction in the data, I consider two distinct proxies for frontier technology growth. First, I compute the correlation between uncertainty and R&D expenditures as the latter are key in adoption process of the growing technology frontier in the model. The correlation between the cyclical components of R&D expenditures, taken from the Bureau of Economic Analysis, and uncertainty is 0.30 which is statistically significant at the 5 percent level. Second, I follow Hall, Jaffe, and Trajtenberg (2001) and use data on patent applications, taken from the U.S. Patent and Trademark Office (USPTO), to proxy for technology growth. The resulting correlation of patent applications with uncertainty is 0.33.[29]

[29] In both cases, the cyclical components were extracted using the HP-filter with a smoothing coefficient of 6.23 for annual data. Redoing the correlations between the level of uncertainty and growth rates in R&D expenditures (patent applications) gives correlation coefficients of 0.16 (0.15). However, both R&D expenditures and patent applications are characterized by medium-term swings in their growth rates which are taken out by the HP-filter.

Therefore, as in the model, also in the data firm-level uncertainty is robustly counter-cyclical but at the same time it correlates positively with proxies of frontier technology growth.

3.2 Technology shocks, uncertainty and creative destruction

Next, let us turn to 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 in e.g. Blanchard and Quah (1989), Gali (1999).[30] 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.[31] 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, aggregate employment and R&D expenditures (recall that in the model the proxy for R&D expenditures are costs of technology adoption of the leading technology).[32] In addition, following Fernald (2007), who documents that low-frequency movements in productivity may impair the identification of technology shocks, the estimation allows for break points in the intercepts. Finally, all VAR specifications are estimated with two lags.[33]

Figure 5 shows the impulse responses of firm-level uncertainty to a positive one-standard-deviation technology shock in the data and the model. The magnitude of the

[30] 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, Fernald, Fisher, and Kimball (2013).

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

[32] Uncertainty is again measured as the cross-sectional variation in TFP shocks taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014), job creation and destruction data 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.

[33] The Appendix provides further details on the estimation procedure as well as several robustness checks.

Figure 5: Uncertainty responses to a positive technology shock: data and model

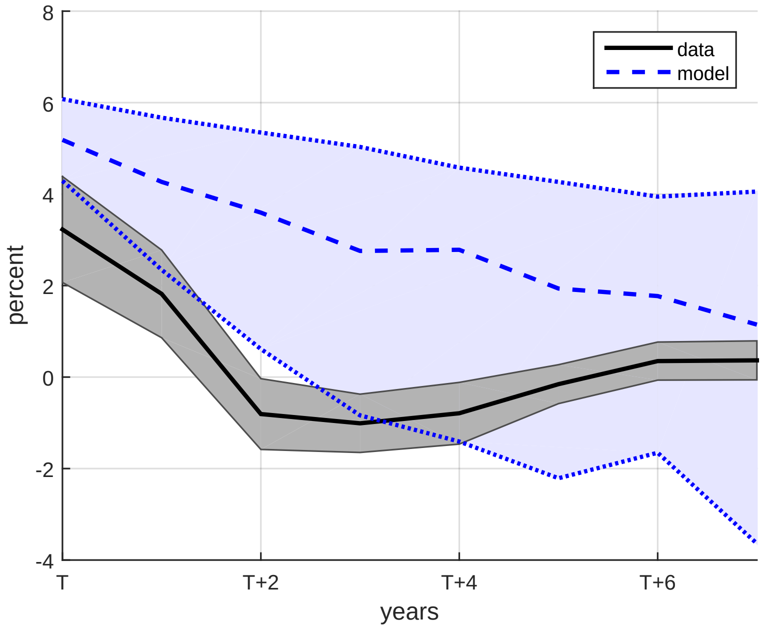

Notes: Impulse response functions to a positive one-standard-deviation technology shock in the "model" and the "data". Uncertainty is measured as cross-sectional variation in establishment-level TFP shocks taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). Shaded areas depict one-standard-deviation confidence intervals from the data. Thin dashed lines depict one-standard-deviation confidence intervals in the model.

uncertainty response in the model is on the high end of that identified in the data. The average model response is also somewhat more persistent, but the one-standard-deviation confidence bands suggest that even in the model the uncertainty increase is only short-lived.

Next, Figures 6 characterizes the technology adoption and creative destruction mechanism behind the counter-cyclical nature of growth-driven uncertainty fluctuations. As in the model, also in the data a positive technology shock raises the incentives to conduct R&D. In addition, an improvement in technology 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. Therefore, the model predictions are not only qualitatively, but also

Figure 6: Aggregate responses to a positive technology shock: data and model

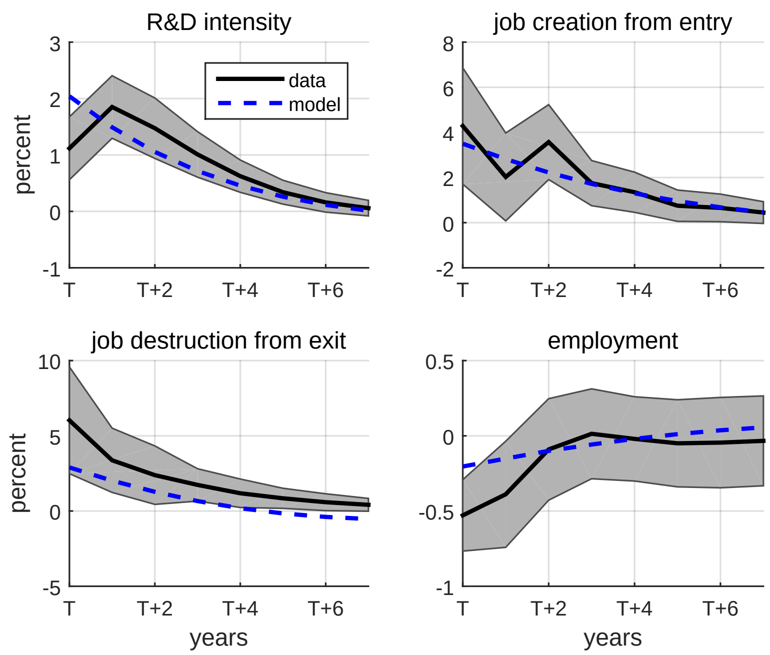

Notes: Impulse response functions to a positive one-standard-deviation technology shock in the "model" and the "data". Job creation from entry and destruction from exit are taken from the Business Dynamics Statistics (both entering in logs), employment is the growth rate of civilian employment taken 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.

quantitatively, consistent with U.S. data.

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 percent 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 all associated with

Figure 7: Uncertainty spikes around recessions and their growth-driven components

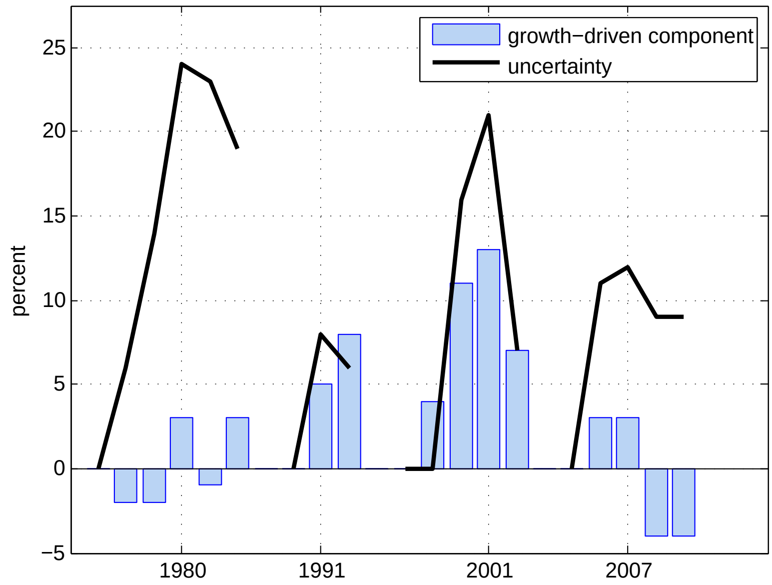

Notes: Cyclical uncertainty increases around NBER recessions measured as deviations from the respective prior troughs. "Growth-driven component" is based on the estimated structural VAR with the identified technology shocks being the only sources of variation. Both time-series are first detrended using the trend from the estimated structural VAR.

a cyclical increase in uncertainty. To measure the latter, I de-mean the data using the estimated intercepts from the structural VAR. 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 7 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.

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.[34] In fact, during both these downturns the growth-driven component contributed negatively to the overall increase in uncertainty.[35]

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 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.

[34]Note that the level of uncertainty was highest during the Great Recession. However a large part of this increase is attributed to the estimated trend. Similar results are obtained for alternative detrending methods.

[35]This is in line with Ludvigson, Ma, and Ng (2017), who argue that the Great Recession increase in uncertainty was primarily related to financial uncertainty.

Appendix to
Creative Destruction and Uncertainty

A Model results

This section of the Appendix provides details of the solution method used in the main text as well as further model details, robustness checks and additional results.

A.1 Solution method

The structural model is a general equilibrium framework with heterogeneous firms. Individual businesses must know the entire distribution of firm productivity and employment levels in order to be able to forecast the development of the wage rate, a key variable in their optimization decisions. In addition, the presence of two aggregate shocks makes these firm distributions time-varying rendering the solution of the model challenging.

The method employed in this paper follows that developed in Sedláček and Sterk (2016). The procedure is based on first-order perturbation along the stationary steady state life-cycle dynamics of individual firms, which depend on the evolution of their firm-specific productivity values. Notice that without persistent idiosyncratic shocks and without adjustment costs, all firms with the same productivity level will make the same decisions. Therefore, it is possible to treat a particular distance from the technological frontier as a separate "firm type". To economize on notation, we can express the model compactly as:

$$\mathbb{E}_t f\left(y_{t+1}, y_t, x_{t+1}, x_t; \Upsilon, \zeta\right) = 0$$

where x_t is a vector containing the state variables (all variables in $\mathcal{S}_t$) and y_t is a vector containing the non-preditermined variables, Υ is a vector containing all parameters of the model and ζ is a scalar parameter pre-multiplying the covariance matrix of the shock innovations, as in Schmitt-Grohé and Uribe (2004). Importantly, the above is system of a finite number of expectational difference equations.

A.1.1 Solving for the steady state without aggregate uncertainty

One first solves for the equilibrium of a version of the model without aggregate uncertainty. That is, I find vectors $\overline{y}$ and $\overline{x}$ that solve $f(\overline{y}, \overline{y}, \overline{x}, \overline{x}; \Upsilon, 0) = 0$. As described in the main text, the calibration targets various parameters to match long-run statistics. The calibration procedure has the following steps:

1. given values for the technology types (i.e. technology gaps), the aggregate wage rate (W), the technology adoption probability (p) and the distribution of firm-specific operational and adjustment costs ($H(\mu_h, \sigma_H)$ and ψ), one can calculate the growth paths of firm-level employment, firm values and the endogenous exit rates.

2. given firm values and exit rates from (1.) and a normalization of the mass of entrants, it is possible to back out the entry cost and to compute the distribution of firm masses across technology types.

3. given the mass of firms in all technology types from (2.) and their optimal choices from (1.) and (2.), it is possible to compute all aggregate variables.

A.1.2 Solving for the equilibrium with aggregate uncertainty

Next, one can solve for the dynamic equilibrium using first-order perturbation around the stationary steady state (including the steady state life-cycle patterns of firms) found in the previous step. The first-order approximated solutions, denoted by hats, have the following form:

$$\begin{aligned} \widehat{x}_{t+1} &= \overline{x} + \Theta\left(\widehat{x}_t - \overline{x}\right) \\ \widehat{y}_{t+1} &= \overline{y} + \Phi\left(\widehat{x}_t - \overline{x}\right) \end{aligned}$$

where Θ and Φ are matrices containing the coefficients obtained from the approximation. The perturbation procedure is standard and carried out in one step.

An advantage of perturbation methods is that the computational speed is relatively high and many state variables can be handled. An important prerequisite for perturbations to be accurate, however, is that deviations from the steady-state are not too large. For firm dynamics models like the one in this paper it may seem problematic because differences in employment levels across firms may be very large. The solution method adopted here, however, overcomes this problem since the steady state we perturb around contains the entire life-cycle profiles of firms. These growth paths, captured by the constants in the above equations, are themselves non-linear functions of technology types.

Hence, the fact that most newborn firms starts off much below their eventual sizes does not involve large accuracy losses since the same is true for the steady-state sizes of newborn firms. Similarly, the fact that the equilibrium features various firm types with very different optimal sizes does not reduce accuracy since we perturb around the growth path for each individual firm type.

A.2 Details of firms' first order conditions

This subsection provides more detailed expressions for the firms' first order conditions presented in the main text. Specifically, it makes explicit the evolution of firm specific productivity. Let us rewrite the first order conditions in terms of firm-specific productivity levels defined by the age of the technology vintage operated by the firm, $z_{j,t} = Z_{t-j}$. The threshold level of operational costs is then defined by

$$\widetilde{\phi}_{j,a,t} = y_{j,a,t} - W_t n_{j,a,t} - R(p_{j,a,t}, \gamma_{j,t}) - \psi_{a,t} n_{j,a,t} + \mathbb{E}_t \beta_t \begin{pmatrix} p_{j,a,t} \theta \mathbb{V}_{a+1}(Z_{t+1}, \mathcal{F}_{t+1}) \\ +p_{j,a,t}(1-\theta)\mathbb{V}_{a+1}(z_{j,t+1}, \mathbb{F}_{t+1}) \\ +(1-p_{j,a,t})\mathbb{V}_{a+1}(z_{j+1,t+1}, \mathbb{F}_{t+1}) \end{pmatrix},$$

In a similar fashion, the optimal expenditures on R&D for incumbent and potential new firms, respectively, are given by

$$\chi(1+\eta)\left(\frac{p_{i,a,t}}{\gamma_{i,t}}\right)^{\eta} = \mathbb{E}_t \beta_t \begin{pmatrix} \theta \mathbb{V}_{a+1}(Z_{t+1}, \mathcal{F}_{t+1}) \\ +(1-\theta)\mathbb{V}_{a+1}(z_{j,t+1}, \mathbb{F}_{t+1}) \\ -\mathbb{V}_{a+1}(z_{j+1,t+1}, \mathbb{F}_{t+1}) \end{pmatrix},$$

$$\chi(1+\eta)\left(\frac{p_{e,t}}{\overline{\gamma}_t}\right)^{\eta} = \theta \mathbb{V}_0(Z_t, \mathcal{F}_t).$$

A.3 Derivation of growth-uncertainty link

The main text showed that the forecasting errors can be written as

$$\upsilon_{i,t} = \tau_{i,t}(p\gamma_{i,t-1} - p\overline{Z}) + (1 - \tau_{i,t})((p-1)\gamma_{i,t-1} - (p-1)\overline{Z}),$$

where $\tau_{i,t}$ is an indicator function which is equal to one with probability p and zero otherwise. Moreover, $\tau_{i,t}$ is independent from any other process, i.e. also from the distribution of past productivity gaps $\gamma_{i,t-1}$. I assume that the model parameters are such that these moments are well-defined. Using this and the fact that the mean and variance of $\tau_{i,t}$ in the cross-section is equal to p and $p(1-p)$, respectively, we can write the variance of the above forecasting errors as

$$\begin{aligned}
var[\upsilon_{i,t}] =& var[\tau_{i,t}(p\gamma_{i,t-1} - p\overline{Z}) + (1 - \tau_{i,t})((p-1)\gamma_{i,t-1} - (p-1)\overline{Z})] \\
=& var[p\gamma_{i,t} - p\overline{Z} + \tau_{i,t}\overline{Z} - \tau_{i,t}\gamma_{i,t}] \\
=& p^2\sigma_\gamma^2 + 0 + \overline{Z}^2 p(1-p) + p^2\sigma_\gamma^2 + \mu_\gamma p(1-p) + p(1-p)\sigma_\gamma^2 \\
=& p(1+p)\sigma_\gamma^2 + p(1-p)\mu_\gamma^2 + p(1-p)\overline{Z}^2,
\end{aligned}$$

where μ_γ and σ_γ are, respectively, the mean and variance of the cross-sectional distribution of the relative stocks of knowledge γ and where we have used the fact that the variance of the product of *independent* random variables is given by $var[XY] = \mu_X^2\sigma_Y^2 + \mu_Y^2\sigma_X^2 + \sigma_X^2\sigma_Y^2$, where again μ and σ^2 indicate the respective means and variances.

A.4 Uncertainty response to aggregate TFP shocks

The main text shows the model impulse response of firm-level uncertainty to a positive frontier technology shock. This subsection provides the same impulse response, but for a positive aggregate TFP shock. In contrast to frontier technology shocks, aggregate TFP shocks affect all firms symmetrically. Therefore, the mechanism through which growth-driven uncertainty fluctuations arise disappears. In particular, even conditional on a radical technology improvements, firms do not experience relatively larger productivity

Figure 8: Uncertainty impulse responses to a positive aggregate TFP shock (ϵ^A)

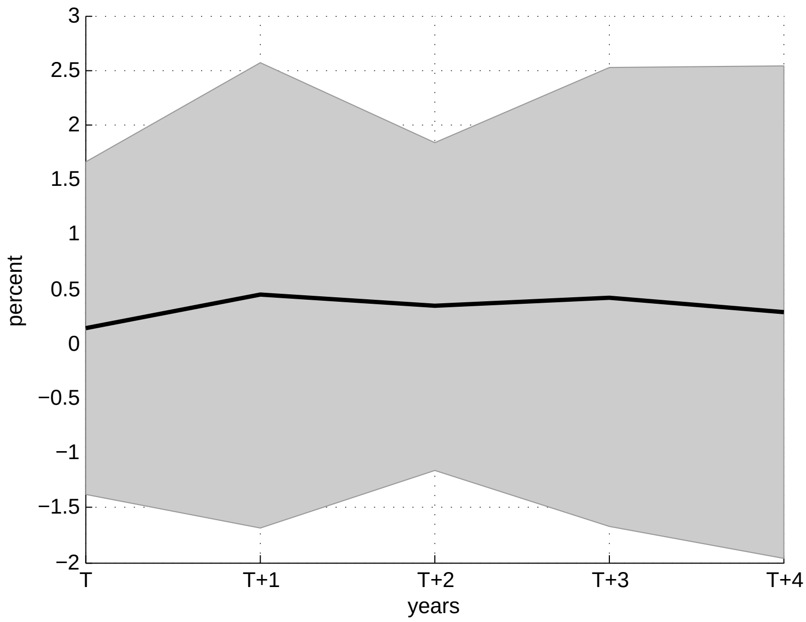

Notes: Impulse response function of the standard deviation of firm-level TFP shocks (computed according to (9)) to a positive one-standard-deviation shock to aggregate TFP. The impulse response is generated by simulating a cross-section of firms $1,000$ times for $1,040$ periods. All exogenous shocks 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 one-standard-deviation confidence bands (shaded areas) over the $1,000$ simulations.

gains (relative to steady state growth) because firm-specific productivity increases for *all* firms, not just those at the frontier.

Figure 8 shows the impulse response of firm-level uncertainty to a positive one-standard-deviation shock to aggregate TFP. The model is simulated $1,000$ times for $1,040$ periods with a cross-section of $100,000$ firms. All exogenous shocks are set to zero except for a positive one-standard-deviation innvoation to the aggregate TFP in period $1,001$. The first $1,000$ periods are discarded. The remaining time-periods are used to estimate equation (9). The figure plots the average response over the $1,000$ simulations together with one-standard-deviation confidence bands.

A.5 Returns-to-scale sensitivity

The benchmark calibration assumes a returns-to-scale parameter equal to $\alpha = 0.67$ delivering a realistic labor share in income. This subsection shows that similar results are obtained with a higher calibration for α. In particular, I consider a calibration in which α is 20 percent higher than it's benchmark value. I recalibrate the model to ensure that it still matches all the calibration targets. Figure 9 shows the impulse responses to a positive one-standard-deviation shock to the frontier technology under this alternative calibration together with the benchmark dynamics.

The figure documents that even with $\alpha = 0.8$, the results are very similar. In particular, the positive technology shock still generates a short-run Schumpeterian downturn with job creation due to entry and job destruction due to exit both rising simultaneously and aggregate employment dropping. In addition, it changes very little on the impulse response of firm-level uncertainty to technology shocks, which still remains to increase through the growth option channel.

A.6 A distribution of technology vintages for entrants

The benchmark calibration assumes that potential entrants can only start up if they manage to adopt the leading technology. This subsection shows that similar results are obtained when assuming a non-degenrate distribution of technology vintages for startups.

In particular, it is assumed that in addition to startups adopting the leading technology, potential entrants that fail to adopt the frontier technology are uniformly distributed between technology vintages of age $j = 0$ and $j = J$, where J is chosen such that the model matches the high exit rate of one year old establishments (21 percent).

Figure 10 shows the impulse responses to a positive one-standard-deviation shock to the frontier technology under this alternative calibration together with the benchmark dynamics. The two models are similar with the largest differences coming from a muted entry response (now that entrants potentially face less attractive technologies) and a more pronounced response of R&D intensity. However, the aggregate employment and uncertainty responses change very little.

Figure 9: Responses to a positive technology shock: benchmark and alternative

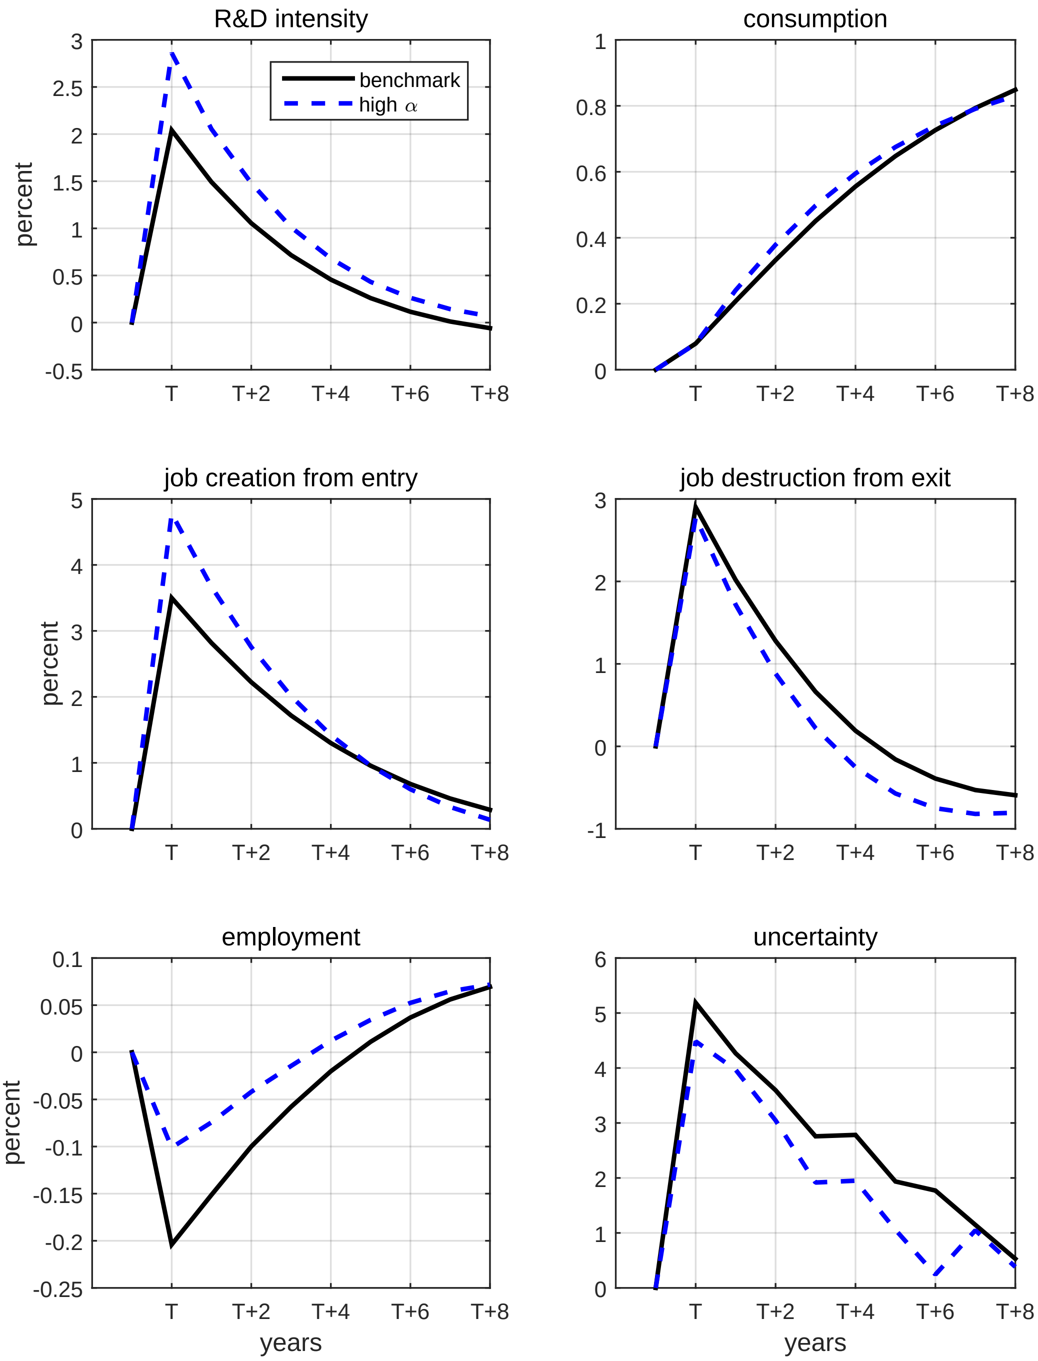

Notes: Impulse response functions to a positive one-standard-deviation technology shock in the "benchmark" and an alternative calibration with a "high alpha" value of 0.8.

Figure 10: Responses to a positive technology shock: benchmark and alternative

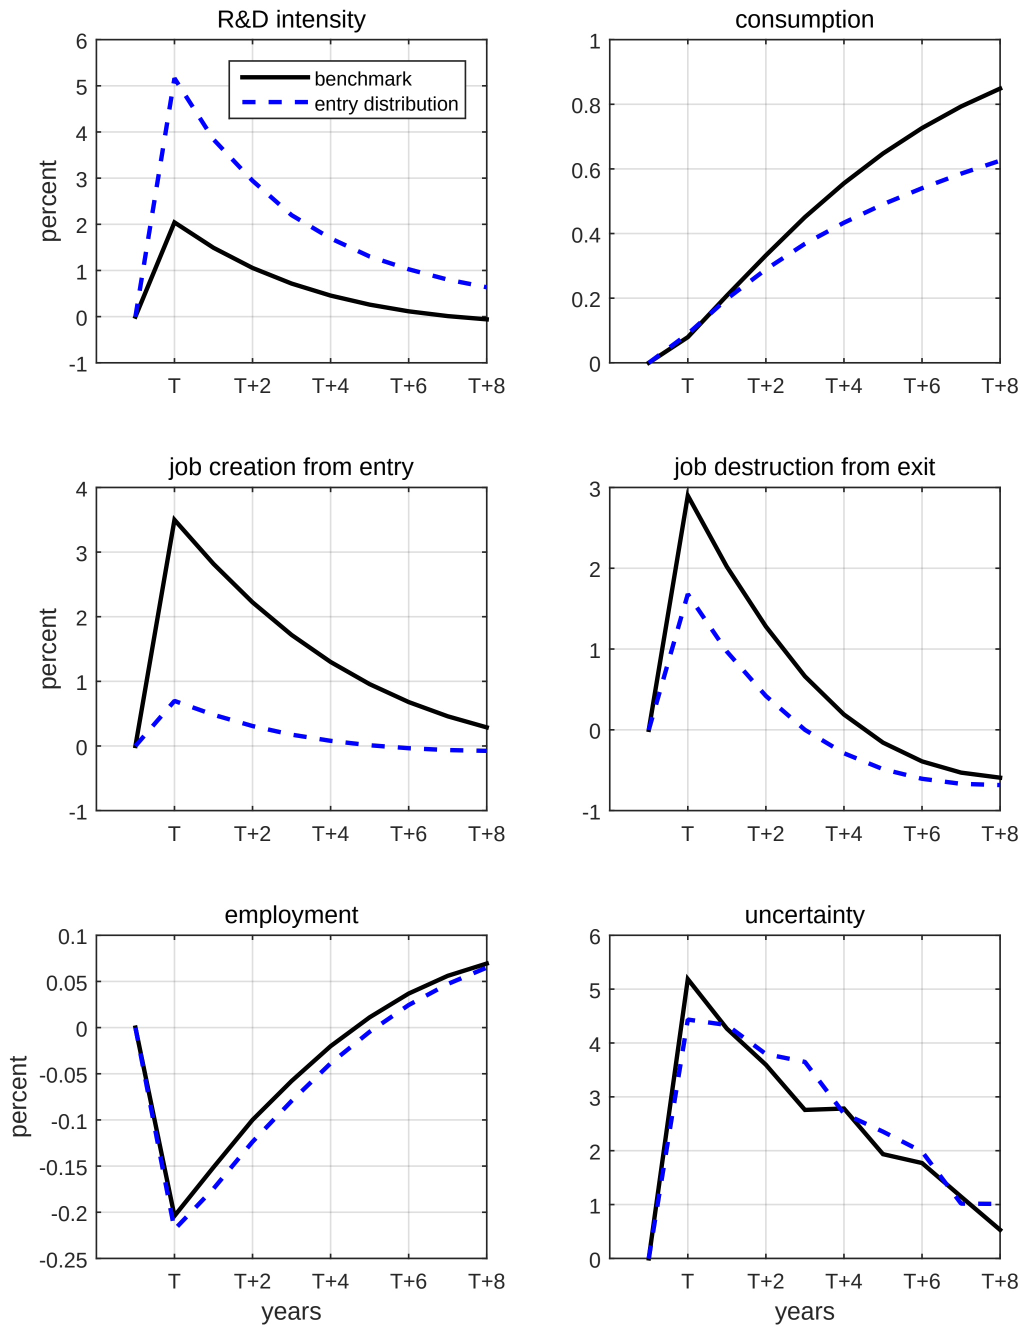

Notes: Impulse response functions to a positive one-standard-deviation technology shock in the "benchmark" and an alternative calibration with "entry distribution" referring to the case when startups face a non-degenrate technology distribution.

A.7 Cost of technology adoption sensitivity

The benchmark calibration calibrated the cost of technology adoption in order to match a establishment-level productivity persistence of 0.79. This subsection shows that similar results are obtained with a lower target of productivity persistence. Specifically, Decker, Haltiwanger, Jarmin, and Miranda (2017) estimate a persistence parameter in the range of 0.6 and 0.7.

Figure 11 shows the impulse responses to a positive one-standard-deviation shock to the frontier technology under this alternative calibration together with the benchmark dynamics. Overall the patterns are very similar with a somewhat muted R&D intensity response (as the average probability of technology adoption increases) and a stronger aggregate employment and uncertainty response.

B Empirical results: details, extensions and robustness

This section of the Appendix explains in more detail the estimation procedure and it provides further impulse responses not discussed in the main text. In addition, it shows results also for other uncertainty measures and results from an alternative estimation procedure.

B.1 Robustness of uncertainty co-movement

This section provides more detailed information on the data used for analyzing the co-movement of firm-level uncertainty and it conducts several robustness checks.

B.1.1 Data used for baseline results

The benchmark measure of firm-level uncertainty used throughout the paper is the cross-sectional dispersion in establishment-level TFP shocks constructed by Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). Specifically, it is the standard deviation of the cross-sectional dispersion estimated using a panel of establishments with at least 25 years of observations.

Figure 11: Responses to a positive technology shock: benchmark and alternative

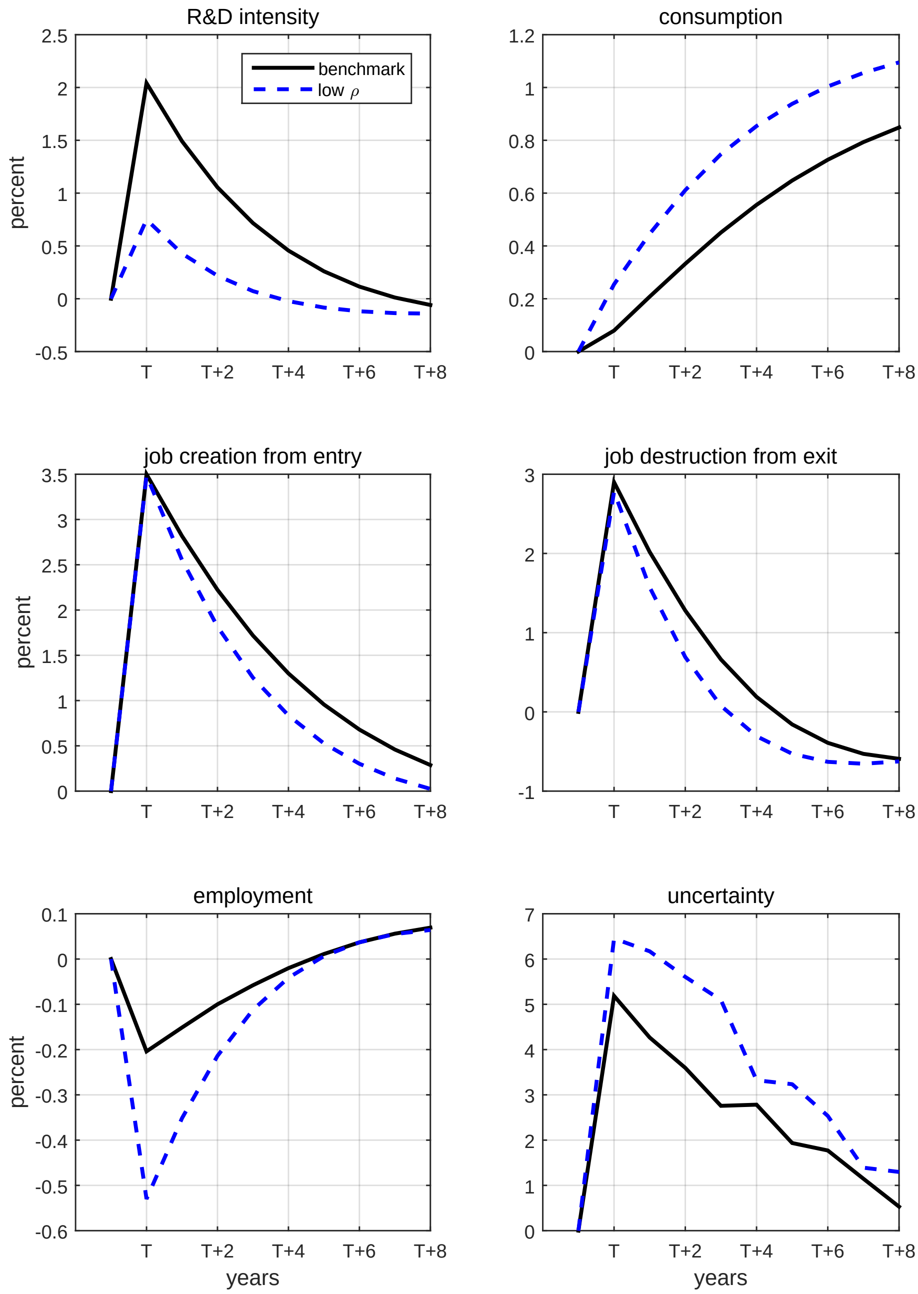

Notes: Impulse response functions to a positive one-standard-deviation technology shock in the "benchmark" and an alternative "low persistence" calibration with an implied persistence of firm-level productivity of 0.7.

The data used to proxy frontier technology growth in the main text comes from two distinct approaches: patent application growth, growth in real R&D expenditures. Patent applications are taken directly from the United States Patent and Trademark Office (USPTO).[36] The benchmark results correlate HP-filtered patent applications (in logs) with the above-mentioned uncertainty measure (also HP-filtered for comparability).

The two proxies for frontier technology growth are mildly pro-cyclical. This is reassuring in the sense that the positive co-movement between uncertainty and frontier technology growth is not merely a correlation between two highly counter-cyclical time series.

B.1.2 Alternative definitions of uncertainty

While the benchmark results are based on standard deviations of TFP shocks of a panel of establishments with at least 25 years of observations, Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) also provide alternative definitions. In particular, they consider a panel of establishments with at least 38 and 2 years of observations and they compute not only the standard deviations, but also inter-quartile ranges. The top panel of Table 4 shows that the results are robust to such alternative definitions.

B.1.3 Patterns after 1984

It is well established that several business cycle patterns underwent a change in the mid 80's (see e.g. Barnichon, 2010; Gali and van Rens, 2014). The second panel of Table 4 shows that the positive co-movement of frontier technology growth and micro-uncertainty does not vanish when considering a sample from 1984-2009.

B.1.4 Patent grants and growth rates

While Hall, Jaffe, and Trajtenberg (2001) argue for the use of patent applications as a measure of technology progress, rather than patent grants, the patterns are robust to this refinement. In addition, considering patent application growth, rather than HP-filtered log-levels does not change the results qualitatively (see third panel of Table 4).

[36] `http://www.uspto.gov/web/offices/ac/ido/oeip/taf/us_stat.htm`.

Table 4: Co-movement of micro-level uncertainty with frontier technology: robustness

	patents	R&D
alternative uncertainty definitions		
s.d. 38Y	0.32^{**}	0.21^{*}
s.d. 2Y	0.39^{**}	-0.08
IQR 25Y	0.08	0.31^{**}
IQR 38Y	0.08	0.27^{**}
IQR 2Y	0.31^{**}	-0.10
correlations after 1984		
corr(σ_t,x)	0.29^{**}	0.54^{***}
patent grants and application growth		
corr(σ_t,grants)	0.19	
corr(σ_t,Δ patents)	0.15	

Notes: micro-level uncertainty, σ_t, is the cross-sectional standard deviation of establishment-level TFP shocks of establishment with at least 25 years of observations taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). Δ indicates growth rates, "patents" refers to the total number of patent applications taken from the USPTO, "R&D" refers to real expenditures on R&D taken from the BEA. "s.d." refers to standard deviation, "IQR" is the interquartile range, "38Y", "25Y" and "2Y" refer to micro-uncertainty measures based on panels of establishments with at least 38, 25 and 2 years of observations, respectively. "grants" refers to the total number of patents granted, taken from the USPTO. One, two and three stars indicate that the correlation is significant at the 10, 5 and 1% level, respectively.

B.2 Baseline estimation procedure

The empirical results presented in the main text are based on a structural VAR with long-run restrictions.

B.2.1 Identification of technology shocks and specification

Let Y_t be a vector of variables with a moving average representation $Y_t = C(L)\epsilon_t$, where $C(L)$ is a matrix of lag polynomials and ϵ_t is a vector of (reduced-form) innovations with a variance-covariance matrix Σ. Furthermore, assume that the vector of variables also has a moving average representation linked to "structural" innovations υ_t given by $Y_t = A(L)\upsilon_t$, where the variance-covariance matrix of the structural innovations is normalized to the identity matrix. The structural and reduced form innovations are then related according to the following relation

$$\upsilon_t = A_0^{-1}\epsilon_t, \tag{13}$$

where A_0 is the coefficient matrix on the current values of υ_t. The variance-covariance matrix of the reduced-form innovations can then be expressed as

$$A_0 A_0' = \Sigma \tag{14}$$

Finally, let the first element of Y_t be the growth rate of productivity and assume, without loss of generality, that the first element of υ_t is a neutral technology shock. Following Gali (1999) the neutral technology shock can be identified using a long-run restriction. In particular, it is assumed that only a neutral technology shock can impact labor productivity in the long-run. This implies that only the first element in the first row of the matrix $\overline{A} = \sum_{i=0}^{\infty} A_i$ is non-zero and the rest are restricted to zero.

Finally, following Fernald (2007) and Canova, Lopez-Salido, and Michelacci (2013), I allow for intercept breaks to account for the low-frequency movements in the data.[37] All the bi-variate structural VAR specifications are conducted including two lags and for R&D intensity, which displays an increasing trend over the sample, I allow for a deterministic cubic time trend in the VAR.

B.3 Mixed-frequency estimation

The main text uses annual data to estimate the structural VARs because the uncertainty measures and establishment-level data are available only at this frequency. However, it is possible to utilize some higher frequency data, in particular on employment and labor productivity and estimate a mixed-frequency structural VAR. This subsection shows that the results remain the same even under this methodology.

[37]The break points are in 1973, 1997 and 2005.

Figure 12: Impulse responses of aggregates: mixed-frequency estimates

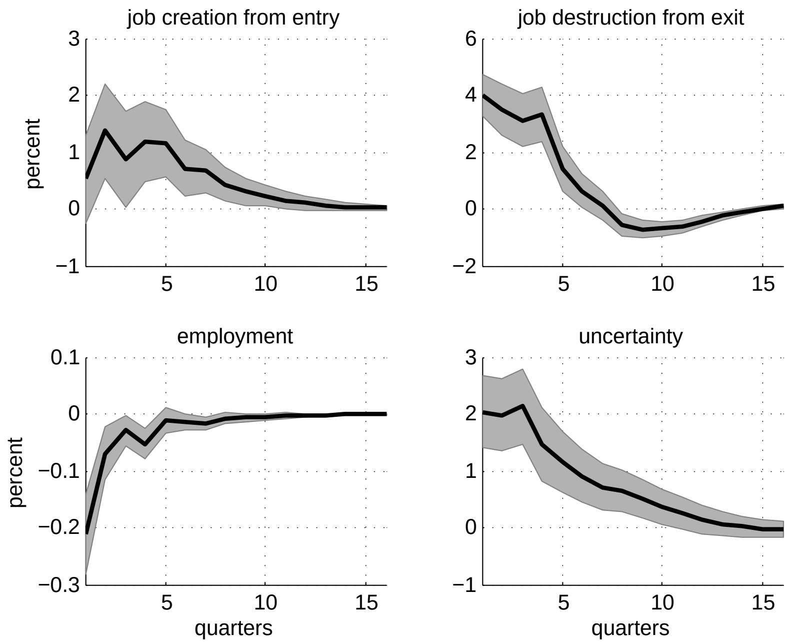

Notes: impulse response functions to a positive one-standard-deviation neutral technology shock estimated using a mixed-frequency structural VAR.

In particular, the used frequency of the VAR is quarterly. In case of the uncertainty measures, which are dispersions of annually observed establishment-level variables, it is assumed that their values in quarters two to four are unobserved. The structure of the VAR itself then serves as an imputation device for such missing observations.

In order to obtain good starting values for the Maximum Likelihood estimation, I first estimate the reduced-form VAR with OLS using the Kalman smoothed estimates of the annual data as one of the variables.[38] The resulting estimates are then used as starting values for the mixed frequency VAR. Once again, following Fernald (2007) and Canova, Lopez-Salido, and Michelacci (2013), I allow for intercept breaks to account for the low-frequency movements in the data.[39]

[38]The Kalman smoothed data is obtained by assuming an AR(1) process for the underlying, unobserved, quarterly variables. Note that this procedure does not utilize the additional information coming from the variation in the variables that are observed at higher frequencies.

[39]The break points are in 1973Q1, 1997Q1 and 2005Q1.

Figure 13: Impulse responses of aggregates: benchmark and when allowing for investment-specific shocks

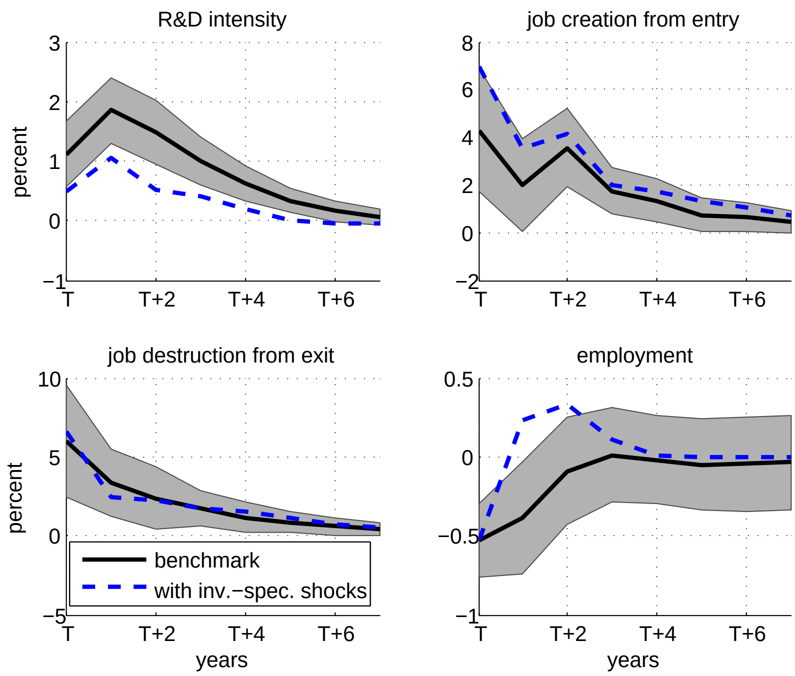

Notes: impulse response functions to a positive one-standard-deviation neutral technology shock in the benchmark and when including investment-specific shocks. The latter are identified by assuming that they can affect productivity and the relative price of investment in the long-run. Shaded areas represent the one-standard-deviation confidence bands from the benchmark estimation.

Figure 12 shows the impulse responses to a positive technology shock of annual variables estimated using the mixed-frequency methodology.[40] All the responses are very close qualitatively and quantitatively to the benchmark estimates in the main text which are based only on annual data.

B.4 Investment-specific technology shocks

The model presented in the main text does not include capital. Therefore, the empirical exercise focuses solely on so-called "neutral" technology shocks. However, Fisher (2006)

[40] Aggregate employment is available at higher frequencies, but is included here to show that the positive technology shock remains to be recessionary in the short run.

Figure 14: Impulse responses of uncertainty: benchmark and when allowing for investment-specific shocks

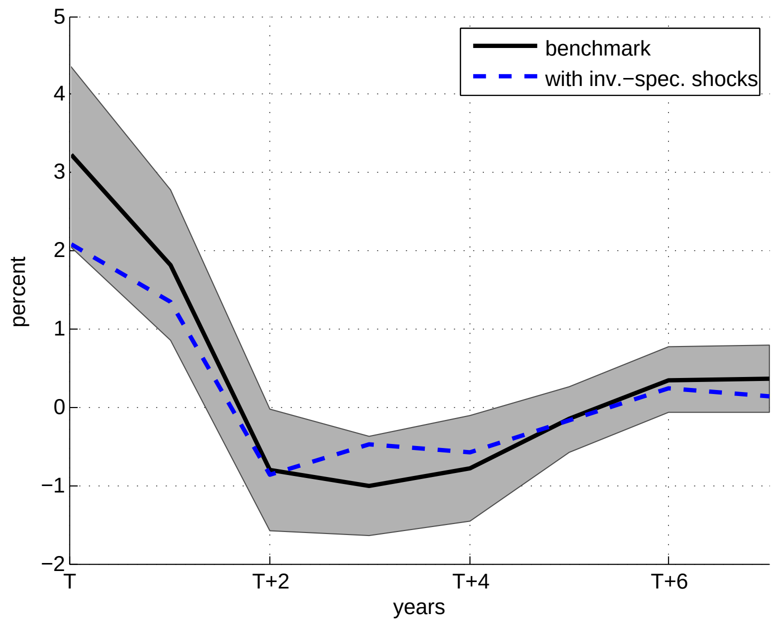

Notes: impulse response functions to a positive one-standard-deviation neutral technology shock in the benchmark and when including investment-specific shocks. The latter are identified by assuming that they can affect productivity and the relative price of investment in the long-run. Shaded areas represent the one-standard-deviation confidence bands from the benchmark estimation.

stresses the importance of distinguishing between neutral and "investment-specific" technology shocks which may have qualitatively different implications for the economy's dynamics. The latter can also affect productivity in the long-run. However, it is assumed that only investment-specific technology shocks can affect the relative price of investment goods in the long-run. The relative price of investment is defined as the investment deflator divided by the consumption deflator.

Figures 13 and 14 show that even when allowing for investment-specific technology shocks, the effects of neutral technology shock remain to be very similar to the benchmark specification.

Figure 15: Impulse responses of uncertainty: alternative uncertainty proxies

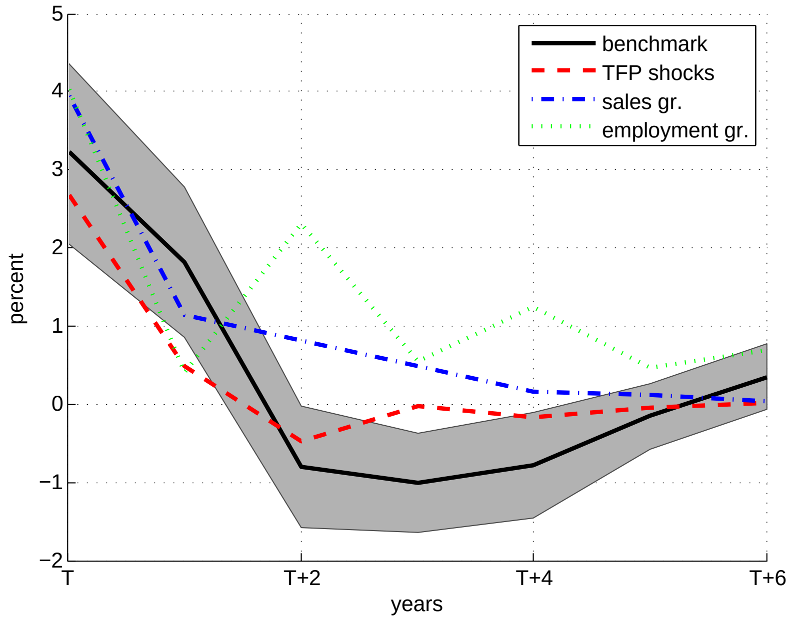

Notes: impulse response functions of firm-level uncertainty measures by Jurado, Ludvigson, and Ng (2015) to a positive one-standard-deviation neutral technology shock. Shaded areas represent one-standard deviation and 90% confidence bands, respectively.

B.5 Alternative uncertainty measures

The main text provided results for uncertainty measures based on the standard deviation of establishment-level TFP shocks. Not only are such measures commonly used in the literature, but the structural model in the main text can exactly replicate such an uncertainty proxy.

However, Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) consider also alternative uncertainty measures including the interquartile range of establishment-level TFP shocks, sales growth rates and employment growth rates. Figure 15 shows that the impulse responses of these alternative measures closely resemble that of the benchmark measure reported in the main text.

Figure 16: Impulse responses of aggregates: benchmark and local projections

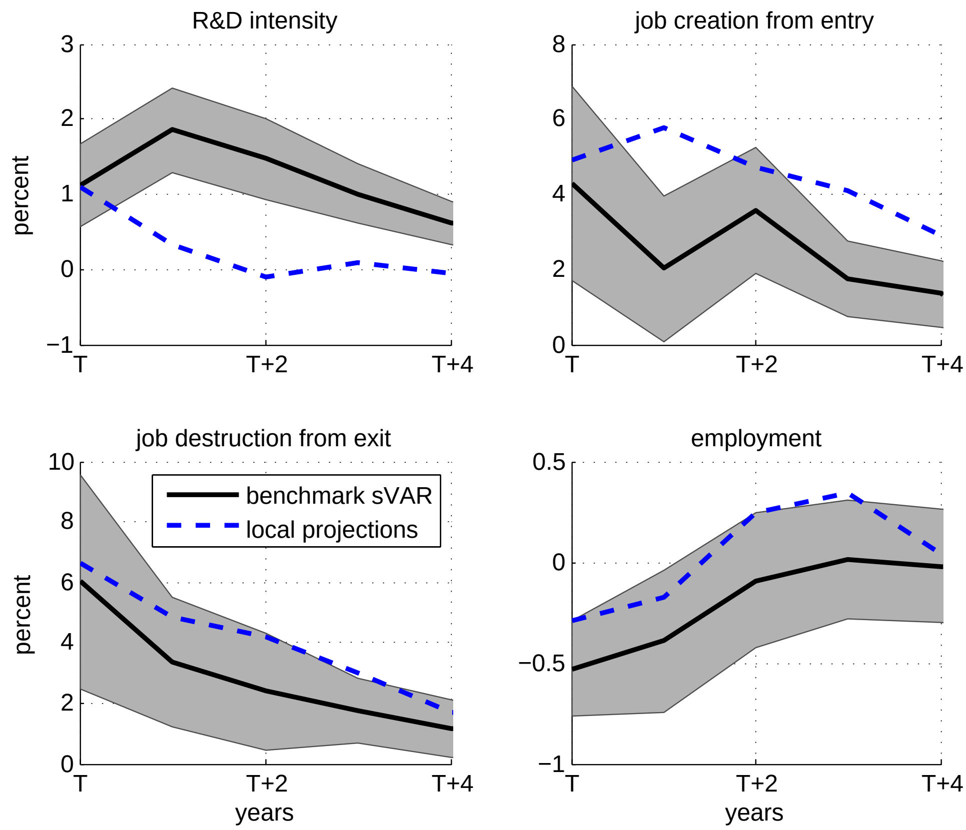

Notes: impulse response functions estimated using local projections and technology shocks estimated by Basu, Fernald, Fisher, and Kimball (2013) together with the associated one-standard-deviation confidence bands. The figure also shows the benchmark responses in the main text estimated using structural VARs.

B.6 Alternative empirical strategy

As an alternative empirical strategy I employ local projects following Jorda (2005) and an exogenous measure of technology shocks developed by Basu, Fernald, and Kimball (2006) and Basu, Fernald, Fisher, and Kimball (2013).[41]

Figure 16 shows the resulting impulse responses of R&D, job creation from entry, job destruction from exit and aggregate employment together with those reported in the

[41] The length of the data over the local projections is fixed to its maximum length at the longest horizon. For trending variables, such as R&D intensity and aggregate employment, the specification also includes a deterministic cubic time trend.

Figure 17: Impulse responses of uncertainty: benchmark and local projections

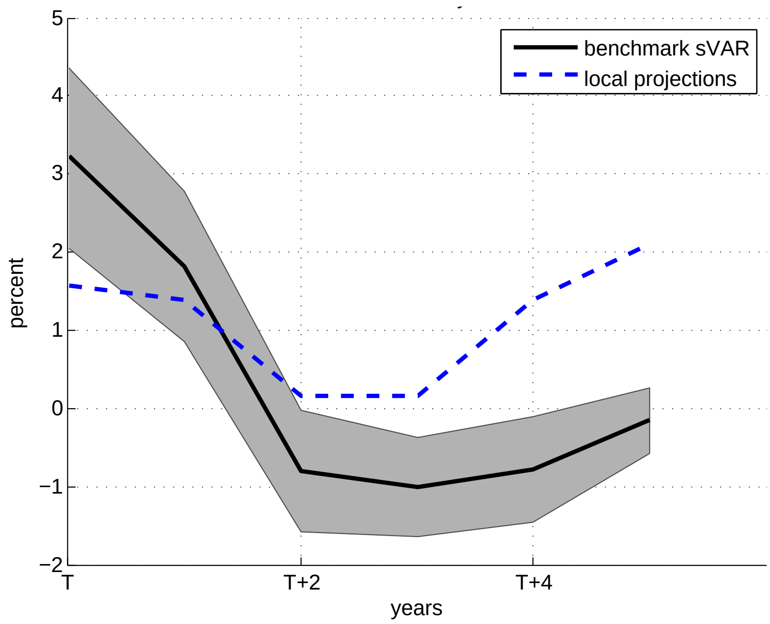

Notes: impulse response function estimated using local projections and and technology shocks estimated by Basu, Fernald, Fisher, and Kimball (2013) together with the associated one-standard-deviation confidence bands. The figure also shows the benchmark responses in the main text estimated using structural VARs.

main text based on the estimated structural VARs. Both the qualitative and quantitative patterns are very similar. The same is apparent from Figure 17, which depicts the impulse response for uncertainty. While the impact response is somewhat milder using local projections, uncertainty remains to increase significantly following a positive technology shock.

References

Abel, A. (1983): "Optimal Investment Under Uncertainty," *American Economic Review*, 73(1), 228–233.

Acemoglu, D., U. Akcigit, N. Bloom, and W. Kerr (2013): "Innovation, Reallocation and Growth," NBER Working Paper No. 18993.

Aghion, P., and P. Howitt (1994): "Growth and Unemployment," *Review of Economic Studies*, 61, 477–494.

Akcigit, U., and W. Kerr (2016): "Growth through Heterogeneous Innovations," mimeo.

Bachmann, R., and C. Bayer (2014): "Investment Dispersion and the Business Cycle," *American Economic Review*, 104(4), 1392–1416.

Bachmann, R., and G. Moscarini (2012): "Business Cycles and Endogenous Uncertainty," mimeo.

Baker, S., and N. Bloom (2013): "Does Uncertainty Reduce Growth? Using Disasters as Natural Experiments," NBER Working Paper 19475.

Bar-Ilan, A., and W. Strange (1996): "Investment Lags," *American Economic Review*, 86(3), 610–622.

Barnichon, R. (2010): "Productivity and Unemployment over the Business Cycle," *Journal of Monetary Economics*, 57(8), 1013–1025.

Basu, S., J. Fernald, J. Fisher, and M. Kimball (2013): "Sector-Specific Technical Change," mimeo.

Basu, S., J. Fernald, and M. Kimball (2006): "Are Technology Improvements Contractionary?," *American Economic Review*, 96(5), 1418–1447.

Berger, D., and J. Vavra (2016): "Shocks vs Responsiveness: What Drives Time-Varying Dispersion?," mimeo.

Blanchard, O., and D. Quah (1989): "The Dynamic Effects of Aggregate Demand and Supply Disturbances," *American Economic Review*, 79(4), 655–673.

Bloom, N. (2009): "The Impact of Uncertainty Shocks," *Econometrica*, 77(3), 623–685.

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

Boedo, M., R. Decker, and P. D'Erasmo (2016): "Market Exposure and Endogenous Firm Volatility over the Business Cycle," *American Economic Journal: Macroeconomics*, 8(1), 148–198.

Caballero, R., and M. Hammour (1996): "On the Timing and Effeciency of Creative Destruction," *Quarterly Journal of Economics*, 111, 805–852.

Canova, F., D. Lopez-Salido, and C. Michelacci (2013): "The Ins and Outs of Unemployment: an Analysis Conditional on Technology Shocks," *The Economic Journal*, 123, 515–539.

Comin, D., and M. Gertler (2006): "Medium-Term Business Cycles," *American Economic Review*, 96(3), 532–551.

Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2017): "Changing Business Dynamism and Productivity: Shocks vs. Responsiveness," mimeo.

Fernald, J. (2007): "Trend Breaks, Long-run Restrictions, and Contractionary Technology Improvements," *Journal of Monetary Economics*, 54, 2467–2485.

Fisher, J. (2006): "The Dynamic Effects of Neutral and Investment Specific Technology Shocks," *Journal of Political Economy*, 114(3), 413–451.

Foster, L., J. Haltiwanger, and C. Syverson (2008): "Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?," *American Economic Review*, 1(98), 394–425.

——— (2016): "The Slow Growth of New Plants: Learning about Demand?," *Econometrica*, 83(329), 91–129.

Francis, N., and V. Ramey (2005): "Is the Technology-Driven Real Business Cycle Hypothesis Dead? Shocks and Aggregate Fluctuations Revisited," *Journal of Monetary Economics*.

Gali, J. (1999): "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?," *American Economic Review*, 89(1), 249–271.

Gali, J., and T. van Rens (2014): "The Vanishing Procyclicality of Labor Productivity," CEPR Discussion Paper, DP9853.

Gourio, F. (2014): "Financial Distress and Endogenous Uncertainty," mimeo.

Hall, B., A. Jaffe, and M. Trajtenberg (2001): "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," NBER Working Paper 8498.

Hansen, G. (1985): "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics*, 16(3), 309–327.

Hartman, R. (1972): "The Effects of Price and Cost Uncertainty on Investment," *Journal of Economic Theory*, 5(2), 258–266.

Ilut, C., M. Kehrig, and M. Schneider (2016): "Slow to Hire, Quick to Fire: Employment Dynamics with Asymmetric Responses to News," mimeo.

Jorda, O. (2005): "Estimation and Inference of Impulse Responses by Local Projections," *American Economic Review*, 95(1), 161–182.

Jurado, K., S. Ludvigson, and S. Ng (2015): "Measuring Uncertainty," *American Economic Review*, 105(3), 1177–1216.

Klette, T. J., and S. Kortum (2004): "Innovating Firms and Aggregate Innovation," *Journal of Political Economy*, 112(5), 986–1018.

KOREN, M., AND S. TENREYRO (2007): "Volatility and Development," *Quarterly Journal of Economics*, 122(1), 243–287.

LOPEZ-SALIDO, D., AND C. MICHELACCI (2007): "Technology Shocks and Job Flows," *Review of Economic Studies*, 74, 1195–1227.

LUDVIGSON, S., S. MA, AND S. NG (2017): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?," mimeo.

MORTENSEN, D., AND C. PISSARIDES (1998): "Technological Progress, Job Creation, and Job Destruction," *Review of Economic Dynamics*, 1, 733–753.

OI, W. (1961): "The Desirability of Price Instability under Perfect Competition," *Econometrica*, 29(1), 58–64.

ORLIK, A., AND L. VELDKAMP (2015): "Understanding Uncertainty Shocks and the Role of Black Swans," mimeo.

RAMEY, G., AND V. RAMEY (1991): "Technology Commitment and the Cost of Economic Fluctuations," NBER Working Paper 3755.

ROGERSON, R. (1988): "Indivisible Labor, Lotteries, and Equilibrium," *Journal of Monetary Economics*, 21(1), 3–16.

SEDLÁČEK, P., AND V. STERK (2016): "The Growth Potential of Startups over the Business Cycle," mimeo.

STEIN, J. (1997): "Waves of Creative Destruction: Firm-Specific Learning-by-Doing and the Dynamics of Innovation," *Review of Economic Studies*, 64(2), 265–288.

STEIN, L., AND E. STONE (2013): "The Effect of Uncertainty on Investment, Hiring and R&D: Causal Evidence from Equity Options," mimeo.