Creative Destruction and Uncertainty*

Petr Sedláček

University of Bonn[†]

September 13, 2016

Abstract

Uncertainty rises in recessions. But does uncertainty cause downturns or vice versa? This paper argues that rather than one causing the other, the coincidence of both events is to a large extent driven by a third common force: technology growth. Indeed, uncertainty and technology growth correlate positively in the data, a novel stylized fact. I show analytically that when technology adoption is gradual, expansions of the technological frontier widen the dispersion of firm-level productivity shocks, a benchmark measure of uncertainty. In a general equilibrium model of endogenous firm dynamics and innovation, technological expansions spur a process of creative destruction, generate a temporary downturn and render uncertainty counter-cyclical. The model's predictions find support in U.S. data showing that up to 40 percent of uncertainty fluctuations are driven by technology growth.

JEL codes: D22, E32, D80

*I thank Christian Bayer, Nicholas Bloom, Steven Davis, Wouter den Haan, Greg Kaplan, Matthias Kehrig, Benjamin Pugsley, Markus Riegler, Moritz Schularick, Emily Sedlacek-Swift, Vincent Sterk, Robert Swift and participants in the 5th Ifo Conference on Macroeconomics and Survey data, the SITE workshop and in seminars at the Australian National University, CERGE-EI, Monash University, Norwegian Business School, University College London, University of Adelaide, University of Cologne, University of Oxford and University of Queensland 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 relationship has been extensively documented in recent years and is robust to many refinements, the question of whether uncertainty is an exogenous driving force of businesses cycles or an endogenous response to them is not well understood. This paper argues that the coincidence of high uncertainty and aggregate downturns is, to a large extent, driven by a third factor: technology growth. Therefore, rather than spikes in uncertainty causing recessions or vice versa, both events are temporary symptoms of the process of creative destruction.

This paper begins by documenting a novel stylized fact. While uncertainty is robustly counter-cyclical, it also correlates positively with measures of technology growth.[1] To understand why uncertainty and technology growth are related, I use a stylized model of firm growth through gradual technology adoption and show analytically that the dispersion of firm-level productivity shocks, i.e. firm-level uncertainty, depends on the speed of technology growth. The intuition is that faster technology growth leads to larger productivity gains when firms successfully innovate and relatively larger productivity losses when they do not.

In order to study the aggregate implications of this novel link between uncertainty and technology growth, I build a tractable general equilibrium model of endogenous firm entry, exit and growth through investment into research and development (R&D). A calibrated version of the model, which matches salient features of U.S. firm dynamics and R&D patterns, is able to replicate both that firm-level uncertainty correlates positively with technology growth and that it is counter-cyclical, as in the data. An increase in technology growth spurs a process of creative destruction as innovative firms start up and grow, while businesses that failed in their R&D efforts shrink or shut down. Because firm entry is costly and firm growth is gradual, it takes time to replace the (creatively) destroyed jobs

[1] I consider three distinct concepts for measuring technology growth: patent application growth, estimates based on structural vector autoregressions and estimates taken directly from the literature. The positive co-movement of uncertainty and technology growth is also robust to other refinements discussed in the Appendix.

generating a short-run contraction. Expansions of the technological frontier are hence responsible for both a wider dispersion of firm-level productivity shocks and a temporary downturn, i.e. uncertainty is counter-cyclical.

Importantly, the model predicts that increases in uncertainty are associated with *more* firm and worker churn. Therefore, in line with the newly documented positive correlation between uncertainty and technology growth, the model provides an alternative interpretation of the observed counter-cyclical nature of uncertainty as a temporary symptom of Schumpeterian reallocation. To the extent that these model predictions can be found in the data, this limits the quantitative relevance of the "wait-and-see" view on uncertainty shocks. In the latter, an exogenous rise in uncertainty generates an aggregate downturn because firms *freeze* their hiring and investment decisions.

To test the model's predictions, I use industry-level data on uncertainty and R&D expenditures and aggregated firm-level data on reallocation. First, the model predicts that faster technology adoption is associated with both a higher average level of uncertainty and with more volatile uncertainty fluctuations. I show that both these predictions hold using 4-digit industry-level data for the manufacturing sector where the speed of technological adoption is proxied by the share of R&D in total sales. Second, to explore reallocation dynamics associated with the process of creative destruction, I make use of Business Dynamics Statistics data on firm entry, exit and employment of continuing businesses. I estimate a series of mixed frequency structural vector autoregressions (VAR) and identify technology shocks via long-run restrictions as in e.g. Blanchard and Quah (1989); Gali (1999); Fisher (2006); Canova, Lopez-Salido, and Michelacci (2013).[2] The estimation reveals that following technological expansions, job creation of entrants, job reallocation among continuing firms and job destruction of firms shutting down all increase significantly and persistently.[3]

[2]The Appendix shows that similar results are obtained when considering alternative uncertainty measures, such as those constructed by Jurado, Ludvigson, and Ng (2015). The results are robust to an alternative empirical strategy based on local projections of technology shocks developed by Basu, Fernald, Fisher, and Kimball (2013).

[3]Using worker and job flow data, Mecikovsky and Meier (2015) and Riegler (2014) provide empirical

Finally, having established the empirical relevance of the model mechanism, I investigate to what extent observed uncertainty fluctuations are indeed driven by technology growth in the data. Towards this end, I employ the uncertainty proxies developed by Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) and estimate their responses to technology shocks using the same strucutral VAR methodology as above. The results suggest that uncertainty measures indeed increase significantly and persistently in response to technology improvements. Quantitatively, technology shocks explain between 20 and 40 percent of the overall variation in uncertainty.

This paper is related to several strands of the literature. First, it connects to studies investigating the business cycle properties and the aggregate implications of exogenous uncertainty shocks (see e.g. Bloom, 2009; Christiano, Motto, and Rostagno, 2010; Schaal, 2012; Bachmann and Bayer, 2014; Bachmann, Elstner, and Hristov, 2014; Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry, 2014; Gilchrist, Sim, and Zakrajšek, 2014; Vavra, 2014; Jurado, Ludvigson, and Ng, 2015; Kehrig, 2015). A handful of papers explore models with positive growth effects of uncertainty or models with causality running from downturns to uncertainty increases (Oi, 1961; Hartman, 1972; Abel, 1983; Bar-Ilan and Strange, 1996; Bachmann and Moscarini, 2012; Gourio, 2014; Orlik and Veldkamp, 2015; Boedo, Decker, and D'Erasmo, 2016).[4] A recent study by Ludvigson, Ma, and Ng (2016) closely connects the some of the empirical results in this paper. While the authors use a different empirical strategy based on instrumental variables, they also conclude that a large part of uncertainty fluctuations are endogenous responses to other structural shocks. In contrast to the above studies, the current paper not only provides novel evidence that uncertainty is positively correlated with technology growth, it also shows the underlying mechanism analytically and it provides a testable theory in which such growth-driven uncertainty fluctuations are counter-cyclical, as in the data.

Second, this paper is related to models and empirical evidence on Schumpeterian creative destruction (see e.g. Aghion and Howitt, 1994; Caballero and Hammour, 1996;

evidence also suggesting that plants do not freeze upon exogenous uncertainty increases.

[4]A somewhat different literature links the average level of macroeconomic volatility to growth within endogenous growth models (see e.g. Ramey and Ramey, 1991; Barlevy, 2004).

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; Christiano, Eichenbaum, and Vigfusson, 2003; Francis and Ramey, 2005; Basu, Fernald, and Kimball, 2006; Canova, Lopez-Salido, and Michelacci, 2013).[5] To the best of my knowledge, the current paper is the first to document how firm dynamics and firm-level uncertainty measures respond to Schumpeterian technology shocks.

The rest of the paper is structured as follows. The next section documents that while uncertainty is robustly counter-cyclical, it co-moves positively with measures of technology growth. Section 3 uses a stylized model of firm growth through innovation to build intuition as to why technology growth and firm-level uncertainty are connected. Section 4 develops a richer general equilibrium model and provides the main quantitative results regarding the aggregate implications of such growth-driven uncertainty shocks. Section 5 shows supporting evidence of the underlying mechanism and quantifies to what extent uncertainty fluctuations are growth-driven in the data. Section 6 concludes.

2 Uncertainty fluctuations in the data

This section first describes the construction of the benchmark measure of uncertainty used throughout the rest of the paper. Next, using this measure, I document a novel stylized fact: while uncertainty is robustly counter-cyclical, it is positively correlated with proxies of technology growth.

I consider three distinct approaches to measuring technology growth. Following Hall, Jaffe, and Trajtenberg (2001) the first proxy is the number of patent applications taken from the United States Patent and Trademark Office. The second proxy is based on structural vector autoregressions (VARs) in which technology shocks are identified as the sole drivers of productivity growth in the long-run (see e.g. Gali, 1999; Fisher, 2006). The last proxy is taken from Basu, Fernald, Fisher, and Kimball (2013), who measure changes in technology as utilization-adjusted total factor productivity (TFP) shocks. A detailed

[5] Fisher (2006) stresses the importance of distinguishing between "neutral" and "investment-specific" technology shocks which typically have different qualitative effects.

description of all the data as well as robustness checks regarding alternative uncertainty and technology growth measures are presented in the Appendix.

2.1 Uncertainty, business cycles and technology growth

Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) develop one of the benchmark uncertainty measures based on the cross-sectional dispersion of establishment-level total factor productivity (TFP) shocks estimated from the following regression

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

where $z_{i,t}$ is the log of estimated establishment-level TFP, ρ is a persistence parameter, λ_t are time fixed effects and $\eta_{i,t}$ are establishment-level TFP shocks.[6] To construct this uncertainty measure, the authors use the Census panel of manufacturing establishments with annual data ranging from 1972 to 2009. As is documented in Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014), uncertainty is robustly counter-cyclical, rising sharply during recessions. This, by now well-established fact, is confirmed by the top panel of Table 1 which correlates the benchmark uncertainty measure with business cycle indicators.

The bottom panel of Table 1, however, shows that uncertainty is positively correlated with proxies of technology growth.[7] These two, seemingly contrasting, business cycle patterns therefore raise questions about the mechanism linking uncertainty to technology growth and about the macroeconomic consequences stemming from such a connection. The next two sections address these questions in turn.

[6] Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) also include establishment-level fixed effects in their regression. I omit them here to ease the exposition, but all the quantitative exercises in this paper exactly follow the methodology of Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). Furthermore, 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 the two.

[7] The Appendix shows that the proxies of 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.

Table 1: Co-movement of micro-level uncertainty with...

	...business cycle indicators		
	ΔY	ΔN	$\Delta Y/N$
corr(σ_t,x)	-0.49^{***}	-0.49^{***}	-0.23^{*}
	...technology growth proxies		
	patents	Gali	BFFK
corr(σ_t,x)	0.33^{**}	0.29^{**}	0.15

Notes: micro-level uncertainty, σ_t, is the cross-sectional dispersion of establishment-level TFP shocks taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). Δ indicates growth rates, Y refers to real GDP, N to civilian employment, "patents" refers to the total number of patent applications taken from the United States Patent and Trademark Office, "Gali" refers to technology shocks identified following the bivariate specification in Gali (1999) and "BFFK" refers to the utilization-adjusted TFP measure constructed by Basu, Fernald, Fisher, and Kimball (2013). One, two and three stars indicate that the correlation is significant at the 10, 5 and 1% level, respectively. Because the (growth rate of the) number of patent applications is characterized by an increasing trend, corr(σ_t,patents) is based on HP-filtered data with the number of patents in logs and using a smoothing coefficient of 6.23.

3 A stylized model of gradual technology adoption

This section builds intuition as to why firm-level uncertainty shocks go hand-in-hand with technology growth. Towards this end, I consider a partial equilibrium model of exogenous growth in which firms occasionally adopt a leading technology (as in e.g. Mortensen and Pissarides, 1998; Lopez-Salido and Michelacci, 2007). Within this framework, I show analytically that technology growth and firm-level uncertainty are inherently linked.

3.1 Firm and technology growth

Consider an economy with a fixed mass of profit-maximizing firms which produce a homogenous final good using labor as the only factor of production. In particular, profits of an individual firm i are given by

$$\pi_i = y_i - W n_i = z_i n_i^{\alpha} - W n_i,$$

where z_i and n_i are firm-specific TFP and employment, respectively, α is the returns to scale parameter and W is a fixed aggregate wage rate. The optimal employment choice is static and results in

$$n_i = \left(\frac{\alpha z_i}{W}\right)^{\frac{1}{1-\alpha}}. \tag{2}$$

Assume further that firm-level productivity is given by the gradual adoption of a leading technology. In particular, consider the existence of a frontier technology level which grows (for now deterministically) over time according to

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

where $\overline{Z} > 0$ is average technology growth. In addition, assume that individual firms adopt the frontier technology with probability $p \in (0, 1)$, but otherwise they retain their previous productivity level

$$\ln z_{i,t} = \omega_{i,t} \ln Z_t + (1 - \omega_{i,t}) \ln z_{i,t-1}, \tag{4}$$

where $\omega_{i,t}$ is a random variable which is equal to one with probability p and zero otherwise. For future reference, let us define the "technology gap" as $\gamma_{i,t} = \ln z_{i,t} - \ln Z_t$.

3.2 Growth-driven uncertainty shocks

In a large enough cross-section a fraction p of firms will have adopted the leading technology Z_t, while productivity of all other firms would have remained fixed. This means that, on average, firm-level productivity can be described by the following law of motion

$$\ln z_{i,t} = (1 - p) \ln z_{i,t-1} + p \ln Z_t + \upsilon_{i,t}, \tag{5}$$

where $\mathbb{E}[\upsilon_{i,t}] = 0$ in the cross-section for every t. Notice, however, that defining $\rho = 1 - p$ and $\lambda_t = p \ln Z_t$ brings us back to the empirical regression (1) used to estimate firm-level uncertainty shocks. In particular, uncertainty is measured as the cross-sectional dispersion of the forecasting errors $\upsilon_{i,t}$, which can be written in terms of structural parameters as

$$\upsilon_{i,t} = \begin{cases} p\gamma_{i,t-1} - p\overline{Z} & \text{when firm } i \text{ does not adopt } Z_t, \\ (p-1)\gamma_{i,t-1} - (p-1)\overline{Z} & \text{when firm } i \text{ adopts } Z_t. \end{cases} \tag{6}$$

The cross-sectional variance of the forecasting errors can then be expressed as

$$\begin{aligned} \text{var}[\upsilon_{i,t}] =& \text{var}\left[\omega_{i,t}\{(p-1)\gamma_{i,t-1} - (p-1)\overline{Z}\} + (1-\omega_{i,t})\{p\gamma_{i,t-1} - p\overline{Z}\}\right] \\ =& \text{var}\left[p\gamma_{i,t-1} - \omega_{i,t}\gamma_{i,t-1} + \overline{Z}\omega_{i,t}\right] \\ =& p(1+p)\sigma_\gamma^2 + p(1-p)\mu_\gamma^2 + p(1-p)\overline{Z}^2, \end{aligned} \tag{7}$$

where $\mu_\gamma = \mathbb{E}[\gamma_{i,t-1}]$ and $\sigma_\gamma^2 = \text{var}[\gamma_{i,t-1}]$ are the cross-sectional mean and variance of last period's distribution of productivity gaps, respectively.[8] The above expression shows that firm-level uncertainty is determined by three components: the distribution of (past) productivity gaps, the probability of adopting the technological frontier and the growth rate of the frontier technology.

It is straightforward to see from (7) 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 aggregate growth are associated with more uncertainty, i.e. a wider dispersion of firm-level TFP shocks. Moreover, given the dependence of firm-level uncertainty on the *past* distribution of productivity gaps, even transitory changes in technology growth have the potential to generate persistent uncertainty effects. Finally, it is important that firm-level productivity be described by *gradual* technology adoption. In the extreme cases of no adoption (vintage technology) or full adoption (homogeneous technology) firm-level uncertainty shocks disappear.[9]

This analytical example shows that, when firm-level productivity is given by gradual adoption of a leading technology, firm-level uncertainty and technology growth are two

[8] The variance is derived using the fact that $\omega_{i,t}$ 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.

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

sides of the same coin. However, whether or not such growth-driven uncertainty fluctuations are empirically plausible, i.e. *counter-cyclical*, remains unclear.

The optimal employment choices in (2) suggest that aggregate employment rises upon an expansion of the leading technology, since firms at the frontier become more productive and the productivity of all other businesses remains unaffected. However, this does not take into account general equilibrium effects, firm selection and life-cycle patterns, nor the possibility that the adoption probability itself is endogenous. All these features, which will alter the reallocation process and in result the short run dynamics of the economy, are considered in the next section.

4 A quantitative model of firm growth through innovation

This section studies the link between technology growth and fluctuations in firm-level uncertainty in an environment in which firms endogenously enter, exit and grow via investment into research and development. The main goal is to quantitatively understand the uncertainty-growth nexus and its aggregate implications. A particular emphasis will be given to its implications for the process of reallocation over the business cycle. The latter will provide testable model predictions enabling the distinction of the presented model mechanism from that of the "wait-and-see" view on uncertainty shocks.

4.1 Model details

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 innovation and the behavior of incumbent and entering firms.

4.1.1 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

$$\frac{C_t^{1-\sigma}-1}{1-\sigma} - \upsilon N_t,$$

where $\sigma > 0$ is the coefficient of relative risk aversion and $\upsilon > 0$ is the disutility of labor. The representative household maximizes the expected present value of life-time utility, subject to its budget constraint

$$C_t = N_t W_t + \Pi_t, \tag{8}$$

which states that total income stems from employment 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

$$W_t C_t^{-\sigma} = \upsilon \tag{9}$$

4.1.2 Research and development and firm-specific productivity

Research and development is conducted by both incumbent and potential new firms. Incumbent firms undertake R&D in order to improve their prevailing firm-level productivity. In particular, following Klette and Kortum (2004) a firm investing r units of the final good into research and development has a probability p of successfully innovating, 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 firm-specific relative "stock of knowledge" and $1/\eta$ is the curvature of the R&D production function. The associated R&D cost function can be written as

$$R(p,\gamma) = \chi\gamma\left(\frac{p}{\gamma}\right)^{\eta}. \tag{10}$$

If an incumbent firm fails to innovate, it retains it's prevailing productivity level. Successful R&D attempts may lead to either radical or incremental innovations (as in e.g. Akcigit and Kerr, 2015). In particular, a fraction θ of innovations is radical leading firms to adopt the frontier technology. All other innovations are incremental in which case firms adopt the technology of the closest younger vintage. Therefore, letting j denote age of a particular vintage of technology, $z_{j,t} = Z_{t-j}$, firm-specific productivity evolves according to

$$\ln z_{j,t} \rightarrow \begin{cases} \ln z_{j+1,t+1} & \text{with probability } 1 - p_{j,t}, \\ \ln z_{j,t+1} & \text{with probability } p_{j,t}(1-\theta), \\ \ln Z_{t+1} & \text{with probability } p_{j,t}\theta. \end{cases} \tag{11}$$

As in the stylized example, frontier technology evolves according to the following process

$$\ln Z_t = \overline{Z} + \ln Z_{t-1} + \epsilon_{Z,t}, \tag{12}$$

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.

Finally, the process of innovation is the same for potential startups as it is for incumbents firms. As a normalization, the accumulated stock of knowledge for potential entrants is assumed to be given by the average steady state productivity gap in the economy, $\overline{\gamma}$. Startups are assumed to enter the economy only if they succeed in radical innovation.

4.1.3 Firm behavior

Incumbent firms produce output using labor, n, as the only production factor. Apart from investing into firm-specific productivity through R&D expenditures, firms can be competitive in production also because of other factors. The latter is important especially for older firms which may not necessarily be on the cutting edge of technology but which nevertheless manage to survive under the pressure of more innovative newcomers.

In particular, it is assumed that firms undergo a learning-by-doing process in which they accrue efficiency gains in production as they age (as in e.g. Stein, 1997). Such learning-by-doing gains, denoted by ψ, can be rationalized by for instance established long-term relationships, well-developed distribution networks or better management practices.

Finally, all incumbent firms are also subject to idiosyncratic stochastic operational costs, denoted by ϕ and distributed according to a cumulative distribution function $H(\phi)$. If a particular realization of operational costs is large enough, firms will no longer find it profitable to operate and will decide to shut down.

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 by choosing employment ($n_{i,a,t}$), an innovation probability ($p_{i,a,t}$) and an exit threshold for operational costs ($\widetilde{\phi}_{i,a,t}$)

$$\mathbb{V}_a(z_{i,t}, \mathcal{F}_t) = \max_{n_{i,a,t}, p_{i,a,t}, \widetilde{\phi}_{i,a,t}} \int^{\widetilde{\phi}_{i,a,t}} \begin{bmatrix} y_{i,a,t} - W_t n_{i,a,t} - R(p_{i,a,t}, \gamma_{i,t}) \\ -\psi_{a,t} n_{i,a,t} - \Phi_{i,a,t} \\ +\mathbb{E}_t \beta_{t+1} \mathbb{V}_{a+1}(z_{i,t+1}, \mathcal{F}_{t+1}) \end{bmatrix} dH_t(\phi), \tag{13}$$

where $\mathbb{V}_a(z_{i,t}, \mathcal{F}_t)$ represents beginning-of-period firm value with $\mathcal{F}_t$ being the aggregate state described below. The first line on the right hand side of (13) represents production, $y_{i,a,t}$, net of wage and R&D costs, where W_t is the aggregate wage rate. The second line consists of age-dependent efficiency losses, which are assumed to be proportional to firm size, and expected paid operational costs. The former is assumed to be given by $\psi_{a,t} = \Psi\delta^a$, where $\Psi_t > 0$ and $\delta < 1$ resulting in firms accruing efficiency gains as they grow older. The latter is simply the expected value of operational costs conditional on survival $\Phi_{i,a,t} = \mathbb{E}[\phi|\phi < \widetilde{\phi}_{i,a,t}]$.[10] Finally, the last line is the firm's continuation value with $\beta_t = \beta(C_{t+1}/C_t)^{-\sigma}$ being the household's stochastic discount factor with $\beta < 1$.

In the above, firm-level production is given by

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

where α controls the returns to scale in production and A_t represents an aggregate total factor productivity shock affecting all firms symmetrically. The latter is assumed to follow an AR(1) process

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

[10] Note that as with R&D expenditures, 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.

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 the marginal product of labor going to workers' wages and to covering efficiency losses from (the lack of) learning-by-doing

$$W_t + \psi_{a,t} = \alpha y_{i,a,t}/n_{i,a,t}. \tag{14}$$

The point at which firms decide to shut down, $\widetilde{\phi}_{i,a,t}$, is implicitly defined by (current-period) firm value being equal to zero

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

And finally, the optimal expenditures on R&D equate the marginal costs of innovating to the marginal benefits both for incumbent and potential new entrants, respectively

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

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

where $p_{e,t}$ is the probability a potential entrant successfully innovates and $\mathbb{V}_0$ represents the firm value of startups.[11]

4.1.4 The distribution of firms, market clearing and balanced growth

Again, let j denote the age of a particular vintage of technology, $z_{j,t} = Z_{t-j}$, and let $\omega_{j,a,t}$ be the respective 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

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

[11]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 (11). To ease the exposition, formulas making this explicit are presented only in the Appendix.

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

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

where firms at the technological frontier are either last period's surviving radical innovators from any part of the firm distribution (top line) or last period's surviving firms which were at the frontier and managed to succeed with an incremental innovation enabling them to keep up with technology growth (bottom line). The distribution of firm masses at productivity levels below the frontier is given by

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

where the mass of firms with productivity z_j is given by the mass of last period's surviving firms which did not innovate and had the same level of productivity (top line) and the mass of last period's surviving firms which incrementally innovated their productivity 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

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

where aggregate production, $Y_t = \sum_a \sum_j \int^{\widetilde{\phi}_{j,a,t}} \omega_{j,a,t} y_{j,a,t} dH(\phi)$, is spent on consumption and aggregate costs $\Xi_t = \sum_a \sum_j \int^{\widetilde{\phi}_{j,a,t}} \omega_{j,a,t} \left(\phi + R(p_{j,a,t}, \gamma_{j,t}) + n_{j,a,t}\psi_{a,t}\right) dH(\phi)$ which include operational costs, R&D expenditures and efficiency losses from (the lack of) learning-by-doing. 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.

4.2 Calibration

This subsection discusses the calibration of the model. The solution method follows Sedláček and Sterk (2016) and its description is deferred to the Appendix for brevity. In order to ease the exposition, I discuss the calibrated parameters in relation to specific targets even though individual parameters typically influence the behavior of the entire model. Unless stated otherwise, the targeted moments are computed using U.S. data in the period from 1977-2013. The sample period is dictated by the availability of firm-level data taken from the Business Dynamics Statistics (BDS). Following the frequency of the BDS data, the model period is assumed to be one year. When computing business cycle statistics, the data is logged and HP filtered with a smoothing coefficient 6.23. All parameter values and the associated targets are presented in Table 2.

Let us start by discussing the parameters pertaining directly to the household. The discount factor, β, is set to 0.96 corresponding to an annual interest rate of 4%. The coefficient of relative risk aversion, σ, is set to one implying log-utility for consumption. The disutility of labor, ν, is backed out from the household's optimal labor supply condition (9) after normalizing the steady state wage rate to one.

The parameters governing the innovation process include the normalization constant χ, the curvature parameter η and the probability of radical innovation (conditional on successfully conducting R&D) θ. Akcigit and Kerr (2015) combine the Longitudinal Business Database with the NBER Patent Database and document that the average R&D intensity (R&D expenditures as a share of sales) is about 4.1 percent. Therefore, the normalization

Table 2: Model parameters

	parameter	value	target/source
β	discount factor	0.96	annual interest rate 4%
σ	relative risk aversion coefficient	1	log-utility
υ	disutility of worker labor	5.541	wage normalization, $W_{ss} = 1$
χ	R&D normalizing constant	0.538	R&D intensity 4.1%, Akcigit and Kerr (2015)
η	R&D cost curvature	2	patent elasticity w.r.t R&D 0.5, Acemoglu et al. 2013
θ	probability of radical innovation	0.1	Akcigit and Kerr (2015)
α	returns to scale	0.660	labor share 66%
ψ_s	learning-by-doing efficiency gains, startups	-1.156	rel. average size of startups 27.3%, BDS
ψ_y	learning-by-doing efficiency gains, young firms	-0.789	rel. average size of young firms 40.4%, BDS
ψ_m	learning-by-doing efficiency gains, medium-aged firms	-0.565	rel. average size of medium-age firms 54.2%, BDS
ψ_o	learning-by-doing efficiency gains, old firms	0	normalization
μ_H	operational cost mean	0.065	0 average paid operational costs, normalization
σ_H	operational cost distribution, scale	0.251	average firm exit rate of 8.7%, BDS
$\overline{E}$	mass of potential entrants	11.515	firm mass of 1, normalization
$\overline{Z}$	frontier technology drift	0.016	average technology growth, Basu et al. (2013)
σ_Z	frontier technology shocks, volatility	0.009	corr(R&D,Y)=0.21, BEA
ρ_A	aggregate TFP shock, persistence	0.552	real GDP autocorrelation 0.79, BEA
σ_A	aggregate TFP shock, volatility	0.008	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.

constant χ is set to match this target. Next, Acemoglu, Akcigit, Bloom, and Kerr (2013) use R&D survey data and estimate that the elasticity of patents with respect to R&D expenditures is roughly 0.5 implying $\eta = 2$ when patents proxy successful innovation in the model. Finally, Akcigit and Kerr (2015) estimate a probability with which innovations are significant enough to open new technologies of 10 percent. Based on this estimate, I set θ, the probability of radical innovation conditional on a successful R&D effort, to 0.1.

Turning towards firm-related parameters, returns to scale are set to a standard value of $\alpha = 0.66$. The age-dependent efficiency gains from learning-by-doing, ψ_a, control the extent to which incumbent firms grow (conditional on firm-specific productivity). Therefore, the associated parameters are chosen such that the model matches the growth patterns observed in the BDS data. To ease the computational burden, the model allows for four age categories: startups, young (one to five years), medium-aged (six to ten years) and old firms (11 years and more). While startups become young firms in the next period (conditional on survival), young (medium-aged) firms become medium-aged (old) firms with a probability $\delta = 1/5$ ensuring an "expected duration" of five years within these age categories (conditional on survival). Each of these categories is characterized by different age-specific parameters ψ_a with $a = s, y, m, o$ indicating the group of startups (s), young (y), medium-aged (m) and old (o) firms. The parameters ψ_a are then calibrated such that the model matches firm sizes by age relative to the economy-wide average (ψ_o is normalized to zero). The distribution of the operational costs, H, controls the extent to which firms exit the economy. It is assumed that H is logistic with mean μ_H such that paid operational costs are normalized to zero and with scaling parameter σ_H calibrated to match the average exit rate of 8.7 percent observed in the data.

Finally, persistence and the standard deviation of shocks to aggregate TFP are set such that the model replicates the autocorrelation and volatility of aggregate output. Average frontier technology growth is set to 1.6 percent which is the average observed growth rate of patent applications in the USPTO. The standard deviation of the innovations to frontier technology growth are set to match the pro-cyclical nature of R&D expenditures observed in the data. The motivation for this last target will become clear in the next subsection discussing the model's properties. In particular, it will be shown that frontier technology and aggregate TFP shocks have the same qualitative effects on R&D expenditures, but the opposite effects on aggregate output. The relative size of the two aggregate shocks then determines to what extent R&D expenditures co-move with aggregate output over the business cycle which is important for the implied reallocation dynamics.

4.3 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 firm distribution, empirically plausible R&D patterns *at the firm-level* as well as an extent of firm and worker churn consistent with the data.

First, while average firm sizes by age were a target of the calibration, the firm and employment distribution were not. Nevertheless, the model does well in replicating these distributions (Table 3). The reason behind this fact is that the model correctly predicts a negative relation between average firm exit rates and age. In particular, while the average exit rate of young (old) firms is 15 (6) percent in the data, it is 11 (6) percent in the model.

Second, a crucial feature of the model is the innovation process. Aggregate R&D expenditures were part of the calibration, but the model does well also in capturing the process of innovation at the firm level. In particular, Akcigit and Kerr (2015) document that small firms innovate relatively more than larger businesses. They estimate that a 10 percent increase in firm size is associated with a reduction of 0.02 standard deviations of patents per employee. Proxying obtained patents by the probability of successful R&D (p) in the model, the baseline framework predicts a value of 0.015 for this elasticity. Therefore, as in the data, also in the model R&D intensity falls with firm size in a quantitatively plausible manner.

Finally, there is a large amount of job churn in the data. Each year about 28 percent of all jobs are either created or destroyed. The model predicts this share to be about 20 percent. Job creation by entrants and job destruction of firms shutting down accounts

Table 3: Firm and employment distributions (in %)

	firm age			
	0	1-5	6-10	11+
	firm shares			
data	10.2	30.7	18.7	40.4
model	8.2	24.7	17.5	49.6
	employment shares			
data	2.6	12.6	10.2	74.4
model	2.1	9.7	9.1	79.1

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

for about 8 percent of aggregate employment both in the data and in the model. Moreover, employment created by startups and jobs destroyed in firms shutting down are not only important for reallocation on average, they are important also for it's business cycle fluctuations (see e.g. Sedláček, 2015). In the data, job creation of entrants and job destruction of exiting firms are both about 3.5 times more volatile than aggregate employment. In the model these values are 3.8 and 3.9 for employment in entering and exiting firms, respectively.

4.4 Model results

This section presents the main quantitative results. It begins by showing that, even in a framework in which the innovation process and the resulting firm dynamics are endogenous, firm-level uncertainty is counter-cyclical and that it co-moves positively with technology growth as in the data. Next, I show that these growth-driven uncertainty increases lead to *more* worker and firm reallocation, contrary to the popular "wait-and-see" hypothesis in which firms freeze their hiring and investment decisions upon uncertainty increases. The next section provides empirical support for the models predictions and quantifies to what extent uncertainty fluctuations are growth-driven in the data.

4.4.1 Model-predicted uncertainty fluctuations

For convenience, I repeat that Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014) define firm-level uncertainty as the cross-section dispersion of residuals ($\eta_{i,t}$) from the following regression

$$\ln z_{i,t} = \mu_i + \rho \ln z_{i,t-1} + \lambda_t + \eta_{i,t}, \tag{15}$$

where firm-level productivity is regressed on its lagged value while allowing for firm (μ_i) and time (λ_t) fixed effects. In order to avoid compositional changes, the focus is only on a balanced panel of firms which are at least 25 years old. Figure 1 plots the impulse response function of the standard deviation of the the above-described firm-level TFP shocks in the model. The impulse response is generated by simulating the model $1,000$ times for $1,010$ 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 impulse response together with confidence bands across the $1,000$ simulations. It documents that firm-level uncertainty increases significantly and persistently following the positive one-time shock to the frontier technology. The intuition provided by the stylized example in Section 3 remains to hold.[12] Uncertainty rises upon impact because the adoption of the frontier technology entails a larger productivity gains while the failure of doing so results in larger relative productivity losses.

Next, to gage how such growth-driven uncertainty fluctuations move over the business cycle, I compute correlations of the standard deviation of firm-level TFP shocks with the cyclical components of aggregate output and employment (first two columns in Table 4). The generated data used in the table is based on $1,000$ simulations of length $1,040$. Unlike with the impulse response, however, the simulations are based on random draws of both exogenous shocks in line with the calibration discussed earlier. The first $1,000$ periods are

[12]The firm fixed effects, which are missing in the example of Section 3, merely capture the average distance of individual firms from the technological frontier over the sample period.

Figure 1: Impulse responses of firm-level uncertainty

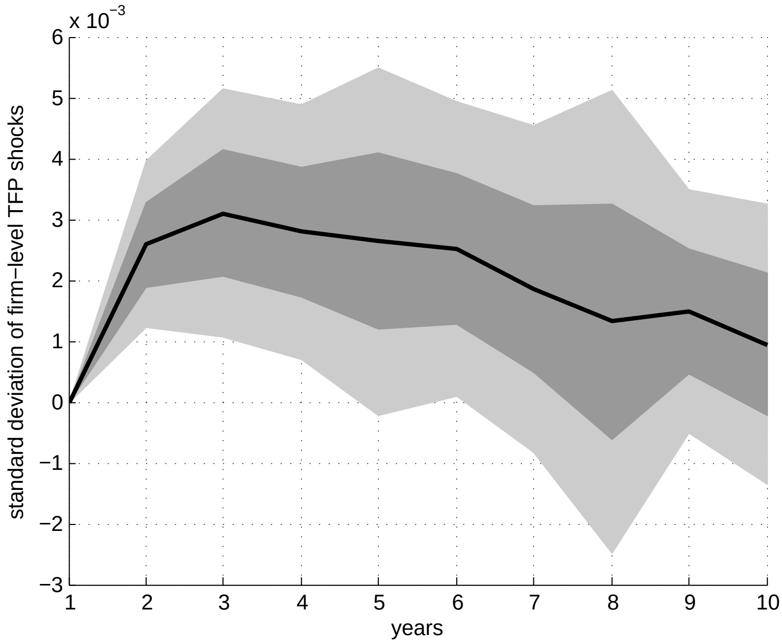

Notes: Impulse response function of the standard deviation of firm-level TFP shocks (computed according to (15)) to a positive one-standard-deviation shock to the frontier technology. The impulse response is generated by simulating a cross-section of firms $1,000$ times for $1,010$ periods. All exogenous 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 and 90% confidence bands (dark and light shaded areas, respectively) over the $1,000$ simulations.

discarded leaving a 40 year sample period consistent with that used in Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). The table reports average correlations and the respective one-standard-deviation confidence bands across the $1,000$ simulations.

Table 4 shows that measured firm-level uncertainty is counter-cyclical even unconditionally on frontier technology shocks, as in the data. The magnitude of the correlations is somewhat lower than in the data (see Table 1) indicating that firm-level uncertainty is likely fluctuating counter-cyclical also for other (potentially exogenous) reasons not present in the current model. Importantly, the model replicates the key business cycle features of firm-level uncertainty observed in the data: while it is robustly counter-cyclical, it

Table 4: Correlation of firm-level uncertainty in model with...

output	employment	frontier tech.
-0.15	-0.13	0.21
$[-0.31, 0.01]$	$[-0.28, 0.03]$	$[0.04, 0.38]$

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 (15)). "frontier tech." 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.

co-moves positively with technology growth (third column of Table 4).

4.4.2 Aggregate shocks and the reallocation process

The popular "wait-and-see" view on uncertainty shocks postulates that upon an uncertainty increase firms freeze their hiring and investment decisions leading to a temporary downturn (see e.g. Bloom, 2009). Based on the observed positive co-movement of uncertainty with technology growth, the presented model provides an alternative explanation of the counter-cyclical nature of uncertainty fluctuations. Specifically, aggregate downturns are temporary symptoms of the Schumpeterian reallocation process which, ultimately, leads to productivity and output growth. The following paragraphs explain this process of creative destruction in more detail. The next section then provides empirical evidence supporting the proposed model mechanism linking technology growth, creative destruction and uncertainty.

Figure 2 shows a set of graphs depicting the impulse response functions of the *distribution* of exit rates and firm sizes to a positive one-standard-deviation shock to frontier technology growth in period T. The horizontal axis represents firms with a particular productivity vintage, $z_{j,t} = Z_{t-j}$, with the most left point being the response of firms at the frontier and the most right point being the responses of firms with a thirty year old technology vintage. The top rows indicate the response of the distributions on impact, the middle row two years and the bottom row four years after the shock hit the economy.

The figure shows that upon impact only firms at the frontier have an easier time surviving in the economy and that they expand. All other firms become relatively less productive and they therefore shut down more often and those that survive shed labor (top panels). Two years after the shock hits the economy, the frontier technology has spread through the economy somewhat. Now even those firms who's technological level is not older than two years, are relatively more productive compared to the stationary distribution and they therefore shut down less frequently and expand their employment levels (middle panels). The same logic applies to the bottom panel indicating that after four years, all firms with a technological level older than four years have a harder time surviving and even if they do, they shrink.

The aggregate implications of these distributional changes are depicted in Figure 3. On impact, it becomes more attractive to startup new businesses and therefore job creation by entrants rises (top left). This effect is persistent as it takes time for the new technology to spread through the economy via innovation of incumbents. At the same time, however, we have seen that most of the economy experiences a decline in relative productivity since only a small fraction of firms is at the frontier. Therefore, job destruction generated by firms shutting down relatively more often also increases after a technological improvement (top right). The same logic also applies to continuing businesses, most of which do not posses the frontier technology and they therefore shrink (bottom left). In the short run, these reallocation dynamics generate an aggregate downturn (bottom right).[13]

Overall, periods of increased firm-level uncertainty are times when job reallocation *rises*. This is true not only because of reallocation through firm entry and exit, but also because among incumbent firms employment flows from relatively less productive firms behind the frontier to those that managed to innovate. The latter is true even though the innovation process is endogenous and expenditures on R&D increase following a frontier technology shock. As long as technology adoption is gradual, shocks to frontier growth will affect the distribution of firms asymmetrically. Technological improvements will therefore

[13] Aggregate output and consumption also fall.

Figure 2: Impulse responses to a frontier technology shock: distributions

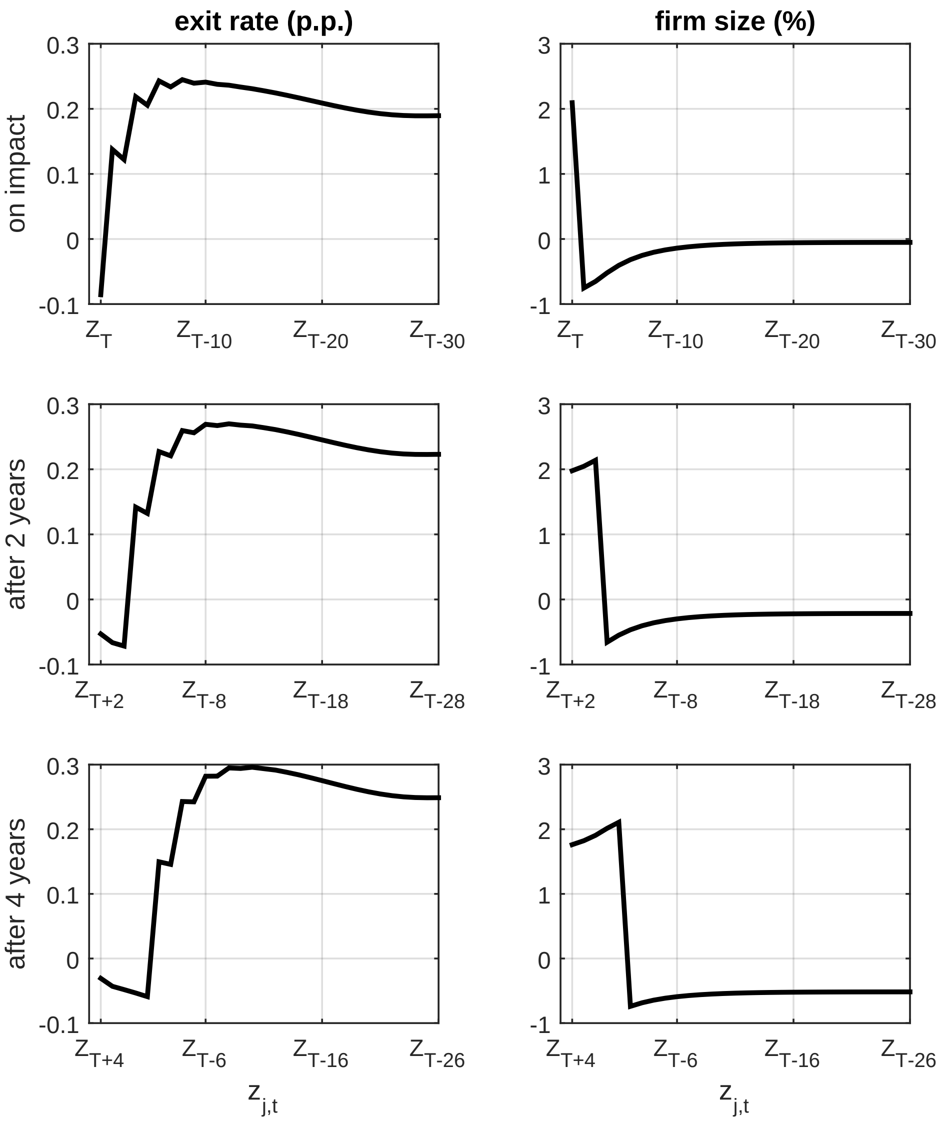

Notes: Impulse response functions to a positive one-standard-deviation shock to the growth rate of the frontier technology which realizes in period T. In all cases the horizontal axis is firm specific productivity expressed in lags of the frontier $z_{j,t} = Z_{t-j}$. The left panels show the response of exit rates in percentage point deviations from their respective steady states. The right panels show the responses of firm employment in percent. The top, middle and bottom rows show the responses upon impact, 2 and 4 years after the shock hit the economy, respectively.

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

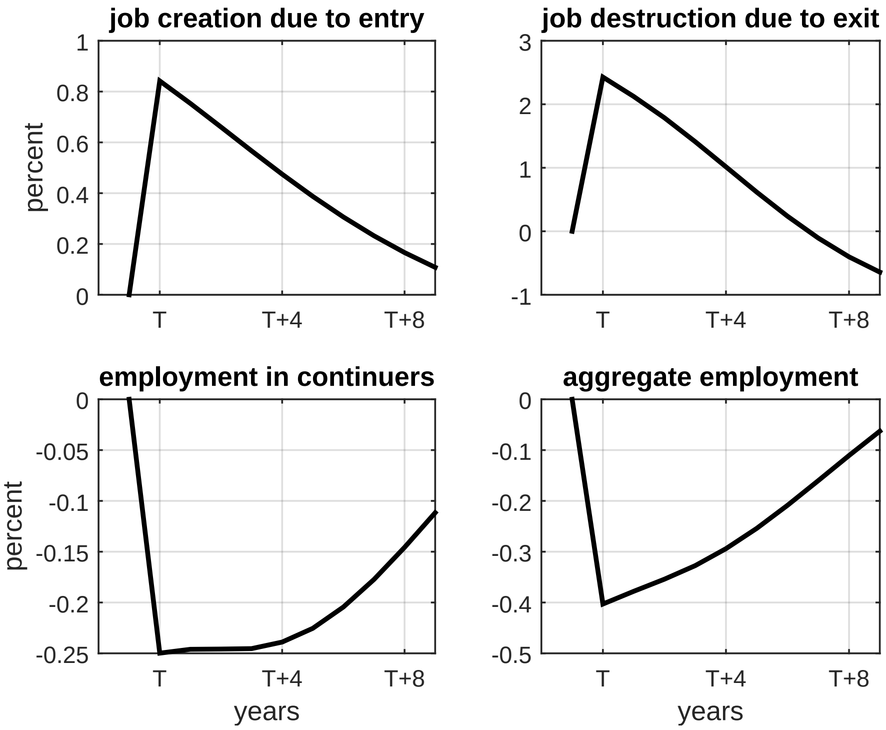

Notes: Impulse responses to a positive one-standard-deviation shock to the frontier technology which realizes in period T. The panels show the responses of employment (in percent) in startups (top left), firms that are shutting down (top right), continuing firms (bottom left) and the aggregate (bottom right).

be characterized by a pattern of *creative destruction*, as firms which are further away from the frontier shut down relatively more giving way to more productive firms at the frontier.[14]

The response to the aggregate TFP shock is standard as seen from Figure 4. Job creation by entrants and employment in incumbent firms both rise (left panels) as *all* firms become more productive. For the same reason, firms shut down relatively less often reducing job destruction from exit. All these effects lead to an aggregate boom (right

[14]Note that R&D expenditures increase after both a positive frontier technology shock and an aggregate TFP shock. However, as was explained, frontier technology improvements lead to short-run downturns which dampens the positive correlation between R&D and output. This explains why targeting this correlation in the calibration section helps pin down the relative size of the two aggregate shocks.

Figure 4: Impulse responses to an aggregate TFP shock: aggregate variables

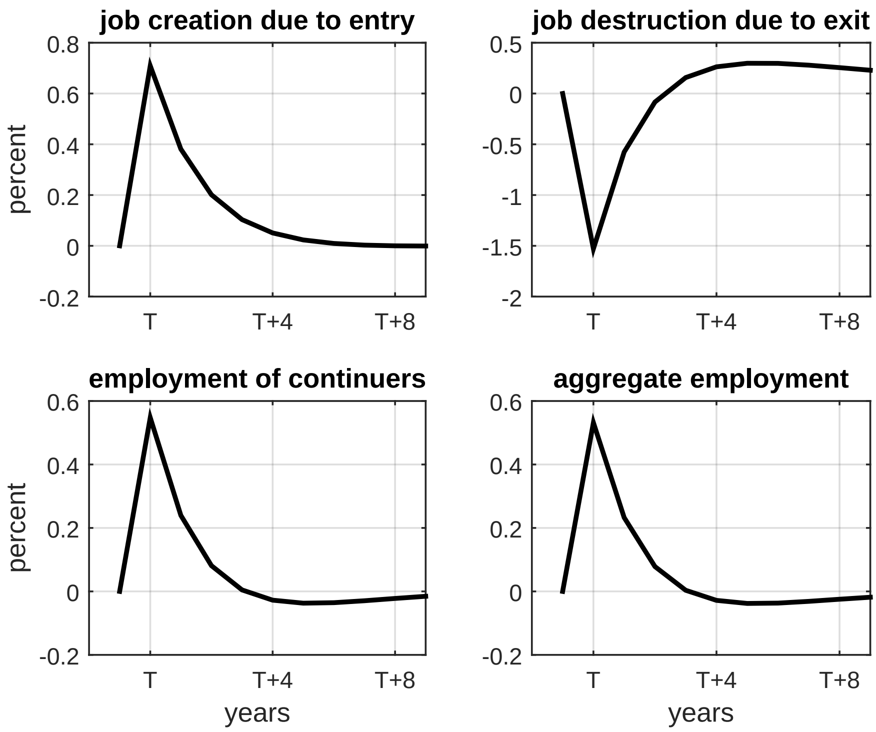

Notes: Impulse response functions to an expansionary (negative) one-standard-deviation aggregate TFP shock which realizes in period T. The panels show the responses of employment (in percent) in startups (top left), firms that are shutting down (top right), continuing firms (bottom left) and the aggregate (bottom right).

panels).

5 Supporting evidence and growth-driven uncertainty fluctuations in the data

In line with the newly documented positive co-movement of uncertainty and technology growth, the presented model explains counter-cyclical uncertainty as an epiphenomenon of creative destruction, thereby offering an alternative view to the "wait-and-see" hypothesis. This section presents empirical evidence supporting this alternative interpretation of uncertainty fluctuations. First, using industry-level data, I document that innovation activity and firm-level uncertainty are related as predicted by the stylized model of Section 3. Second, this section also shows that the model predicted reallocation dynamics are consistent with aggregated firm-level data. Finally, having provided empirical support for the model mechanism, this section ends by quantifying the extent to which uncertainty fluctuations are growth-driven in the data.

5.1 Innovation and uncertainty at the industry level

The stylized example of Section 3 derived an analytical expression for the components of the cross-sectional variance of firm-specific productivity shocks, i.e. for firm-level uncertainty. For convenience, I repeat equation (7) here

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

The above expression determines firm-level uncertainty as depending on the distribution of productivity gaps (with mean μ_γ and variance σ_γ^2), the variance of frontier technology growth ($\overline{Z}_t^2 = (\overline{Z} + \epsilon_{Z,t})^2$) and the probability of adopting the leading technology (p). If there is no or full technology adoption, i.e. a pure vintage or homogeneous technology model, then firm-level uncertainty is constant. In fact, both the average level and the variation in firm-level uncertainty increases with the speed of technology adoption.[15] This constitutes a testable prediction that can be verified in the data.

In particular, as an uncertainty proxy I use the dispersion of plant-level productivity shocks at the 4-digit industry-level in manufacturing constructed by Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). As a measure of the speed of frontier technology adoption, I use average R&D expenditures and R&D stocks as a share of sales, also at the 4-digit industry level, computed by Bloom, Schankerman, and van Reenen (2013).[16] Figure 5 shows that the predictions of the model can also be found in the data: industries with larger R&D shares are also industries where uncertainty is on average larger (top panels) and more volatile (bottom panels).[17]

[15] This holds true for $p \leq 0.5$.

[16] Bloom, Schankerman, and van Reenen (2013) use data on R&D expenditures from Compustat and impute the stock of R&D using the perpetual inventory method.

[17] There is some evidence that R&D may be related to uncertainty also due to growth-options (see

Figure 5: Impulse responses to an aggregate TFP shock: aggregate variables

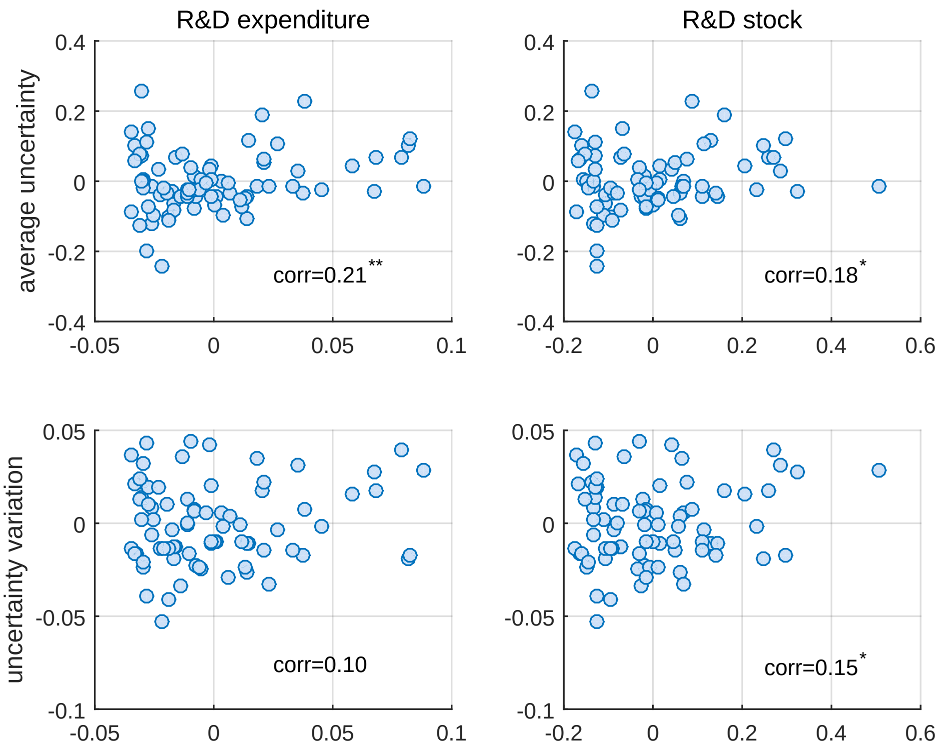

Notes: The scatter plots show the relation between the speed of frontier technology adoption, proxied by R&D expenditures (left panels) and R&D stocks (right panels) as a share of sales constructed by Bloom, Schankerman, and van Reenen (2013), and average firm-level uncertainty (top panels) and the standard deviation of firm-level uncertainty (bottom panels) at the 4-digit industry level for manufacturing. The uncertainty measures are taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). One and two stars indicate statistical significance of the one-sided test at the 10 and 5 percent level, respectively.

5.2 Reallocation dynamics in the aggregate

The full quantitative model provides testable predictions about the nature of the reallocation process, conditional on frontier technology shocks. This subsection investigates whether these predictions are present in the data.

e.g. Kraft, Schwartz, and Weiss, 2013). However, Bloom (2007) notes that it is unclear whether high uncertainty increases or decreases R&D expenditures.

5.2.1 Estimation strategy and data

In order to analyze the responses of firm dynamics to technology shocks I estimate a series of structural vector autoregressions (VARs) using long-run restrictions to identify technology shocks as in e.g. Blanchard and Quah (1989), Gali (1999), Fisher (2006) and Lopez-Salido and Michelacci (2007).

The identification is based on assuming that only technology shocks determine productivity in the long-run, which is consistent with the structural model. The vector of variables in the structural VAR is given by $Y_t = (\Delta q_t, \Delta a_t, e_t, x_t)'$, where Δ is the first-difference operator, q is the log of the price of investment goods, a is the log of labor productivity, e is the log of the employment rate and x represents different firm dynamics variables described below. In order to economize on the number of estimated parameters, I estimate a separate structural VAR for each different firm dynamics variable.[18]

By including the investment price, the estimation accounts for investment-specific technology shocks which are not considered in the model but which may also affect productivity in the long-run Fisher (2006).[19] While the inclusion of labor productivity growth facilitates the identification of the technology shocks, the employment rate is included in order to confirm that technology shocks indeed lead to a temporary aggregate downturn.

The price of investment goods is measured as ratio of the investment and the consumption deflator. Labor productivity is given by output per hour in the non-farm business sector and the employment rate is defined as 100 minus the unemployment rate. The main focus of the empirical analysis is aimed at the variables in x which include different measures of firm dynamics taken from the Business Dynamics Statistics (BDS) database. In particular, I use data on job creation in entrants, job destruction in exiting firms and job reallocation among continuing firms from the Business Dynamic Statistics database.

An additional challenge arises from the fact that the uncertainty measures are available only at an annual frequency. Therefore, in order to exploit the quarterly information

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

[19] The identifying assumption is that investment-specific technology shocks impact not only productivity in the long-run, but also the relative investment price.

in the other variables, the structural VAR is estimated using the Kalman filter and Maximum likelihood with the within-year observations of the uncertainty proxies treated as missing values. The presented results are thus for the Kalman filtered time-series of the underlying (unobserved) quarterly uncertainty measures. The sample period is dictated by the availability of the BDS data and it therefore runs from 1976Q2-2013Q1. In all cases, the structural VARs are estimated with four lags.

Finally, it has been argued that low-frequency movements in productivity may impair the identification of technology shocks (see e.g. Fernald, 2007; Canova, Lopez-Salido, and Michelacci, 2013). Following the latter two studies, the estimation accounts for this possibility by allowing for breaks in the intercepts of the structural VAR.

The Appendix provides further details on the estimation procedure as well as additional results such as the responses of other variables to neutral technology shocks and the impact of investment-specific technology shocks. Moreover, it also shows that the baseline results are robust to a set of alternatives: using local projections as in Jorda (2005) together with technology shocks constructed by Basu, Fernald, and Kimball (2006), estimating the structural VAR using only annual information and using alternative uncertainty measures constructed by Jurado, Ludvigson, and Ng (2015).

5.2.2 Technology shocks and firm dynamics in the data

Figure 6 displays that as in the model, also in the data, a technological improvement leads to a statistically significant increase in the number of jobs created among startups, the number of jobs lost due to exiting firms and the number of jobs reallocated between continuing firms. In light of the structural model, this last piece of evidence provides support for the mechanism behind growth-driven uncertainty shocks.

In particular, to the extent that part of the evolution of firm-specific productivity is driven by the adoption of a stochastically evolving leading technology, aggregate technology shocks result in firm-level uncertainty fluctuations. The reason why such fluctuations are counter-cyclical stems from the fact that such growth shocks spur a process of creative destruction. While firms which did not adopt the leading technology shrink or shut down,

Figure 6: Impulse responses to a technology shock: firm dynamics

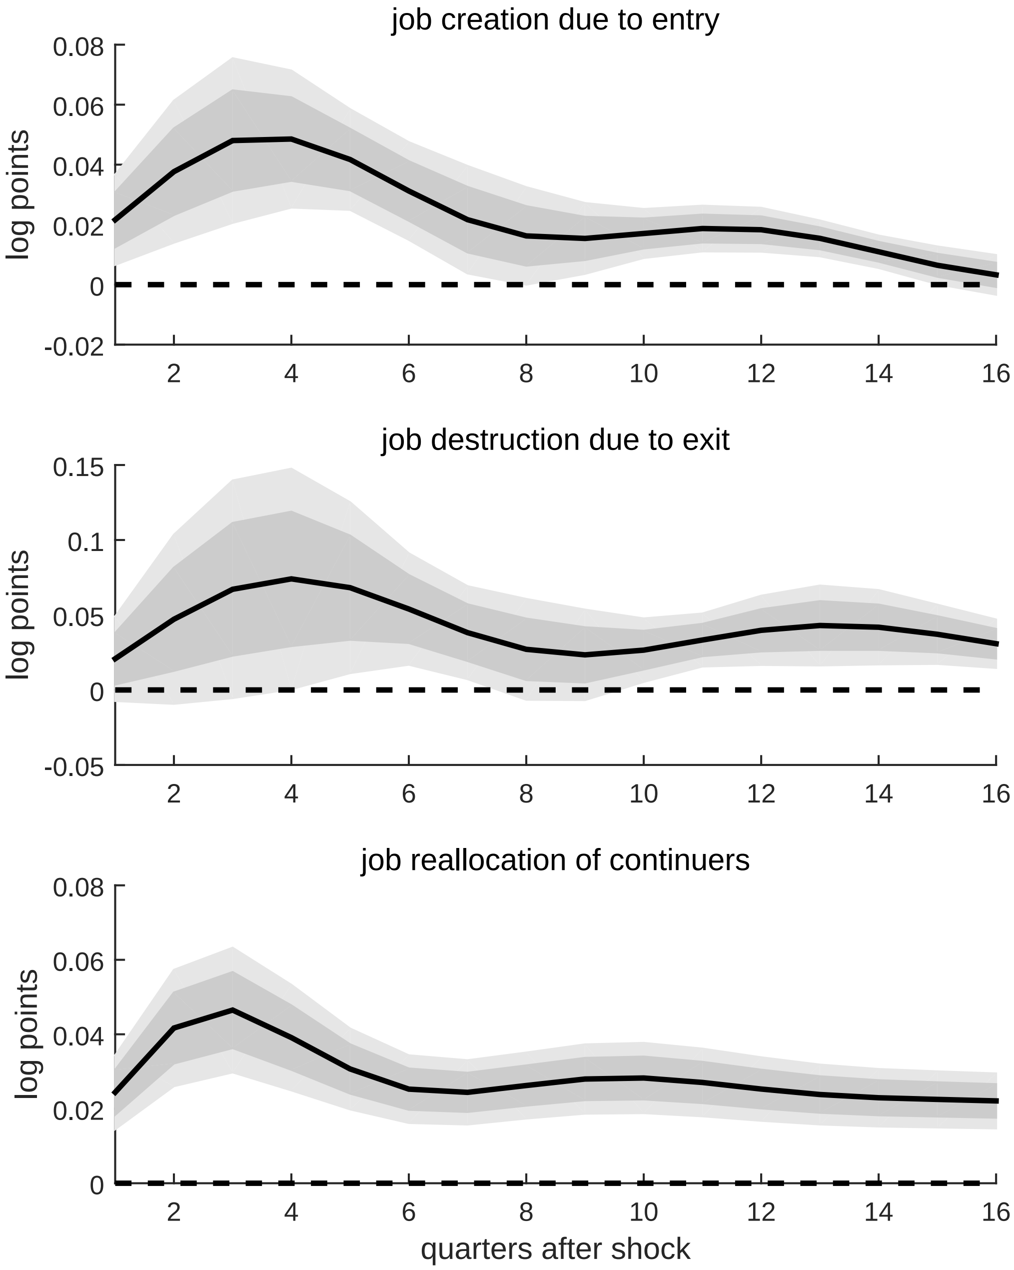

Notes: Impulse responses to a positive one-standard-deviation shock to the frontier technology. The firm dynamics data are taken from the Business Dynamics Statistics. The dark and light shaded areas indicate one-standard-deviation and 90% confidence intervals, respectively.

Table 5: Forecast error variance decomposition: contribution of neutral technology shocks

	forecast horizon in quarters				
	1	4	8	12	16
establishment-level TFP shocks	31.9	38.9	37.2	35.3	34.9
establishment-level employment growth rates	32.6	30.1	25.6	24.5	24.1
firm-level stock returns	29.3	25.8	21.6	20.2	19.6

Notes: The table reports forecast error variance decompositions at different horizons (in columns) as a fraction (in percent) of total variation of a given uncertainty measure (in rows).

new firms emerge and new jobs are created. However, because firm growth takes time and because firm entry is costly, the initial destruction phase dominates, generating a temporary downturn.

5.3 Growth-driven uncertainty fluctuations in the data

Having provided empirical support for the proposed model mechanism, it remains to be investigated whether also uncertainty measures can be found to increase following technology improvements in the data. Towards this end, I repeat the above empirical exercise, but this time instead of firm dynamics variables, the structural VARs include various measures of micro-uncertainty. In particular, the benchmark specifications use the interquartile ranges of establishment-level TFP shocks, establishment-level employment growth rates, and firm-level stock returns all taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). The sample periods used in the estimations are again dictated by the availability of the uncertainty measures.[20]

Figure 7 shows impulse response functions of the three benchmark uncertainty proxies to a positive one-standard-deviation frontier technology shock. In all three cases

[20]Specifically, the estimation is based on quarterly data over the period 1972Q1-2009Q4 when using the dispersion of establishment-level TFP shocks and employment growth rates as uncertainty measures. When using the dispersion of firm-level stock returns as the uncertainty proxy, the sample period is 1960Q1-2010Q3.

Figure 7: Impulse response functions of uncertainty measures

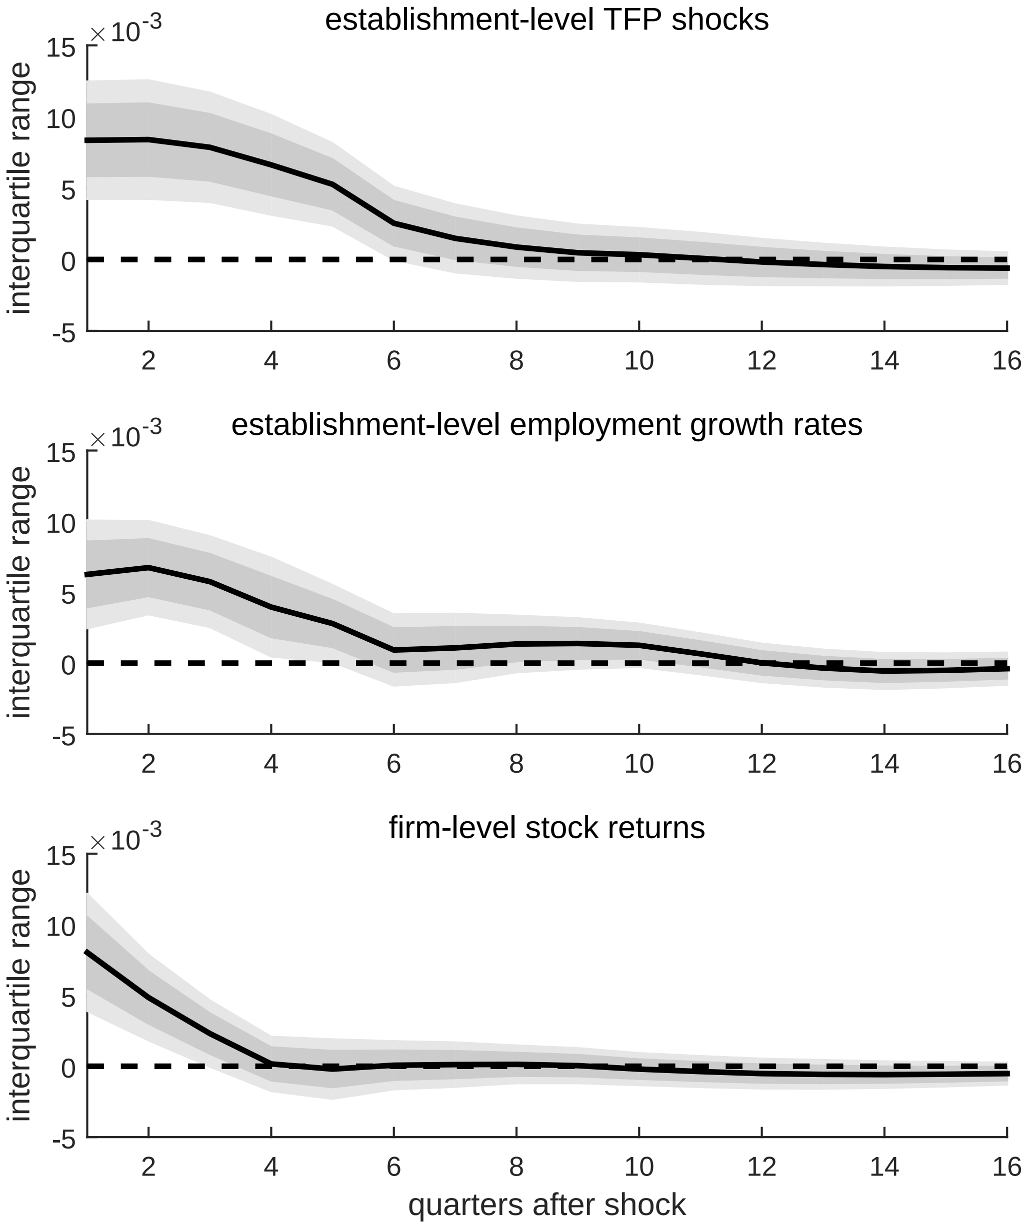

Notes: Impulse response functions to a positive one-standard-deviation frontier technology shock. All uncertainty proxies are taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). The dark and light shaded areas indicate one-standard-deviation and 90% confidence intervals, respectively.

establishment- and firm-level uncertainty rises significantly and persistently in response to the technology shock.

Moreover, looking at the forecast error variance decomposition in Table 5 shows that frontier technology shocks are important drivers of uncertainty fluctuations also quantitatively. In particular, they alone explain between 20 and 40 percent of the overall uncertainty variability depending on the uncertainty measure and the forecast horizon. These magnitudes are in line with previous studies analyzing the quantitative importance of technology shocks for the business cycle variation of aggregate variables (see e.g. Fisher, 2006; Lopez-Salido and Michelacci, 2007) and they corroborate the evidence in Ludvigson, Ma, and Ng (2016) who also find that a large part of uncertainty fluctuations are endogenous responses to other shocks.

Finally, let us gage the extent to which the proposed mechanism impacted the specific time-path of uncertainty fluctuations over time. Figure 8 shows the three uncertainty measures together with associated counterfactuals. The counterfactuals are simply the predicted values of the uncertainty measures from the estimated structural VARs based on variation in frontier technology shocks alone.

The figure shows that these "growth-driven" uncertainty fluctuations track the respective patterns of the unconditional uncertainty measures quite closely with correlation coefficients between 0.42 and 0.47. Moreover, growth-driven uncertainty fluctuations are characterized by recessionary spikes similar to those of the unconditional time-series. However, other factors influencing firm-level uncertainty are also clearly important. This is apparent in the 1980's recessions and especially in the Great Recession during which the strong increases in uncertainty were, by and large, not growth-driven.[21]

The provided 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, but the latter are important drivers of uncertainty fluctuations also quantitatively.

[21] This is in line with Ludvigson, Ma, and Ng (2016), who argue that the Great Recession increase in uncertainty was primarily related to financial uncertainty.

Figure 8: Actual and "growth-driven" uncertainty measures

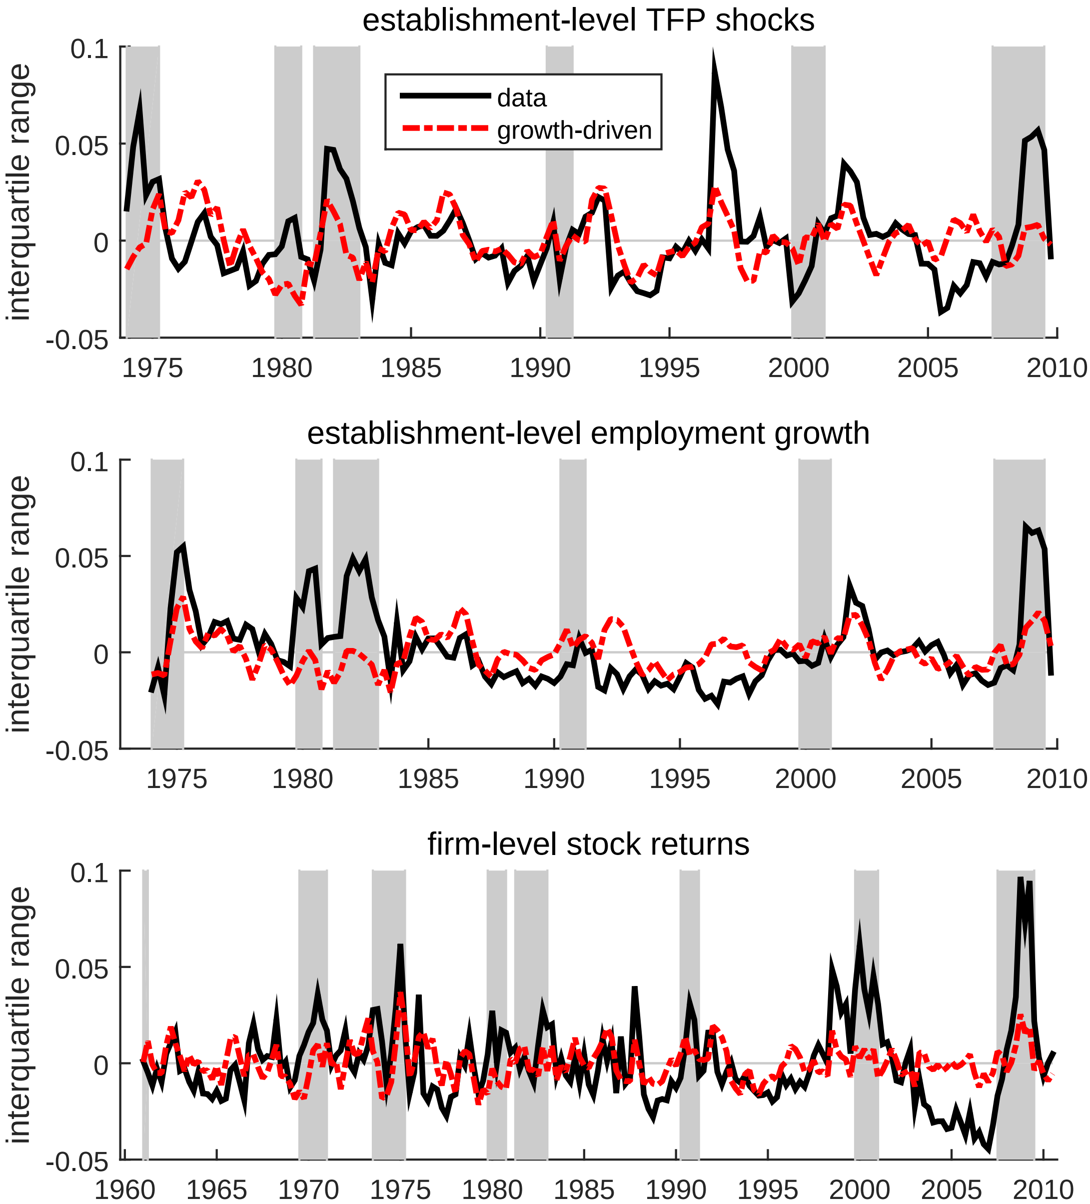

Notes: "Actual" refers to data taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014), while "growth-driven" refers to counterfactuals constructed as predicted values from the estimated structural VARs using variation in the identified neutral technology shocks alone. All variables are expressed in deviations from the respective estimated (time-varying) means.

6 Conclusion

This paper documents that uncertainty co-moves positively with technology growth. Consistent with this observation, counter-cyclical fluctuations in uncertainty can be understood as an epiphenomenon of creative destruction, providing an alternative view to the "wait-and-see" hypothesis. In particular, such growth-driven uncertainty shocks are associated with *more* worker and firm reallocation ultimately leading to productivity growth. These predictions are confirmed in the data and technology shocks are shown to explain up to 40 percent of uncertainty fluctuations, limiting the extent to which uncertainty shocks can be exogenous drivers of the business cycle.

This does not mean that other factors do not play an important role in shaping uncertainty fluctuations. Especially 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 variation in uncertainty.

Appendix to
Creative Destruction and Uncertainty

A Robustness of uncertainty co-movement

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

A.1 Data used for baseline results

The benchmark measure of micro-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.

The data used to proxy frontier technology growth in the main text comes from three distinct approaches: patent application growth, estimation using a structural VAR and utilization-adjusted TFP growth. Patent applications are taken directly from the United States Patent and Trademark Office (USPTO).[22] The benchmark results correlate HP-filtered patent applications (in logs) with the above-mentioned uncertainty measure (also HP-filtered for comparability).

The second proxy comes from estimating a bi-variate structural VAR as in Gali (1999). The data includes labor productivity growth, measured as real output per hour in the non-farm business sector, and the employment rate (defined as 1 minus the unemployment rate). The identifying assumption is that only technology shocks affect labor productivity in the long-run. Therefore, this methodology is the same as that used in the main text to identify the responses of firm dynamics and uncertainty measures to technology shocks without, however, accounting for investment-specific technology shocks.

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

Table 6: Cyclicality frontier technology proxies

	patents	Gali	BFFK
ΔY	0.08	0.20*	0.19*
Δ N	0.04	−0.13	0.13
Δ Y/N	0.10	0.44***	0.16

Notes: Δ indicates growth rates, Y is real GDP, N is aggregate employment, "patents" refers to the total number of patent applications taken from the USPTO, "Gali" refers to technology shocks identified following the bivariate specification in Gali (1999) and "BFFK" refers to the utilization-adjusted TFP measure constructed by Basu, Fernald, Fisher, and Kimball (2013). One, two and three stars indicate that the correlation is significant at the 10, 5 and 1% level, respectively.

Finally, the third proxy is taken directly from Basu, Fernald, Fisher, and Kimball (2013), who construct a utilization-adjusted measure of TFP growth.

Before we move on to the robustness exercises, it is worth noting that the correlation of the three proxies for frontier technology growth are mildly pro-cyclical (Table 6). 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.

A.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 7 shows that the results are robust to such alternative definitions.

Table 7: Co-movement of micro-level uncertainty with frontier technology: robustness

	patents	Gali	BFFK
	alternative uncertainty definitions		
s.d. 38Y	0.34**	0.29**	0.25*
s.d. 2Y	0.37**	0.17	0.09
IQR 25Y	0.08	0.24*	0.27*
IQR 38Y	0.08	0.24*	0.28**
IQR 2Y	0.29**	0.10	0.08
	correlations after 1984		
corr(σ_t,x)	0.30**	0.32**	0.31**
	patent grants and application growth		
corr(σ_t,grants)		0.21	
corr(σ_t,Δ patents)		0.15	
	R&D expenditure growth		
corr(σ_t,ΔR&D)		0.36**	

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, "Gali" refers to technology shocks identified following the bivariate specification in Gali (1999) and "BFFK" refers to the utilization-adjusted TFP measure constructed by Basu, Fernald, Fisher, and Kimball (2013). "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, "R&D" refers to real expenditure on research and development taken from the BEA. One, two and three stars indicate that the correlation is significant at the 10, 5 and 1% level, respectively.

A.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 7 shows that the positive co-movement of frontier technology growth and micro-uncertainty does not vanish when considering a sample from 1984-2009.

A.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 7).

A.5 R&D expenditures

Yet another possible proxy for frontier technology growth are real expenditures on research and development. Such a measure is somewhat looser than the other since an increase in R&D does not necessitate a successful improvement in technology. Nevertheless, bottom panel of Table 7 shows that R&D expenditure growth is also positively correlated with micro-uncertainty.

B Model results: solution method and details

This section of the Appendix provides details of the solution method used in the main text as well as further model details.

B.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.

B.2 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 $f\left(\overline{y}, \overline{y}, \overline{x}, \overline{x}; \Upsilon, 0\right) = 0$. 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.

B.3 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.

B.4 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

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

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

C 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.

C.1 Estimation procedure

The empirical results presented in the main text are based on a structural VAR with long-run restrictions. Let us first discuss the identification of technology shocks and then explain the details of the mixed-frequency nature of the VAR.

C.1.1 Identification of technology shocks

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

$$\upsilon_t = A_0^{-1}\epsilon_t, \tag{16}$$

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

$$A_0 A_0' = \Sigma \tag{17}$$

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.

C.1.2 Estimation methodology

As mentioned in the main text, the uncertainty measures are available only at an annual frequency. In order not to loose information by aggregating other variables (available at higher frequencies), such as productivity growth and employment, the above-discussed structural VAR is estimated with Maximum Likelihood utilizing the Kalman filter to handle the true mixed-frequency nature of the data.

In particular, the used frequency of the VAR is quarterly. In case of the uncertainty measures, which are dispersions of annually observed firm- or 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.[23] The resulting estimates are then used as starting values for the mixed frequency VAR. 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.[24]

C.2 Further impulse response functions

The main text focused primarily on the responses of firm dynamic measures and uncertainty proxies to (neutral) technology shocks. This subsections shows the remaining IRFs plotted together for all three estimated structural VARs in Section 5. First, Figure 9 shows that neutral technology shocks permanently increase (the level of) labor productivity (middle left panel). This is the identifying assumption. Second, they have a temporary positive impact on the investment price (top left panel) and they are recessionary in the sense that the employment rate falls persistently (bottom left panel). The right panels show the same responses, but for the case of a positive investment-specific shock. The investment price falls and labor productivity increases permanently, again due to the identifying assumptions. The employment rate somewhat increases initially (except for the specification with the dispersion of TFP shocks) and then falls in response to investment-specific technology shocks. All these response are roughly in line with previous studies (see e.g. Lopez-Salido and Michelacci, 2007).

Second, Figure 10 shows the IRFs of the three benchmark uncertainty proxies to a positive investment-specific technology shock. In contrast to the responses to neutral technology shocks, there is no clear pattern in the response of uncertainty proxies to investment-specific innovations. This highlights the importance of accounting for investment specific shocks in the estimation.

[23]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.

[24]The break points are in 1973Q1, 1997Q1 and 2005Q1. Average productivity growth is statistically different in all three sub-periods. Shifting the break points forward or backwards by 1-2 years does little to the results.

Figure 9: Impulse responses to a technology shocks

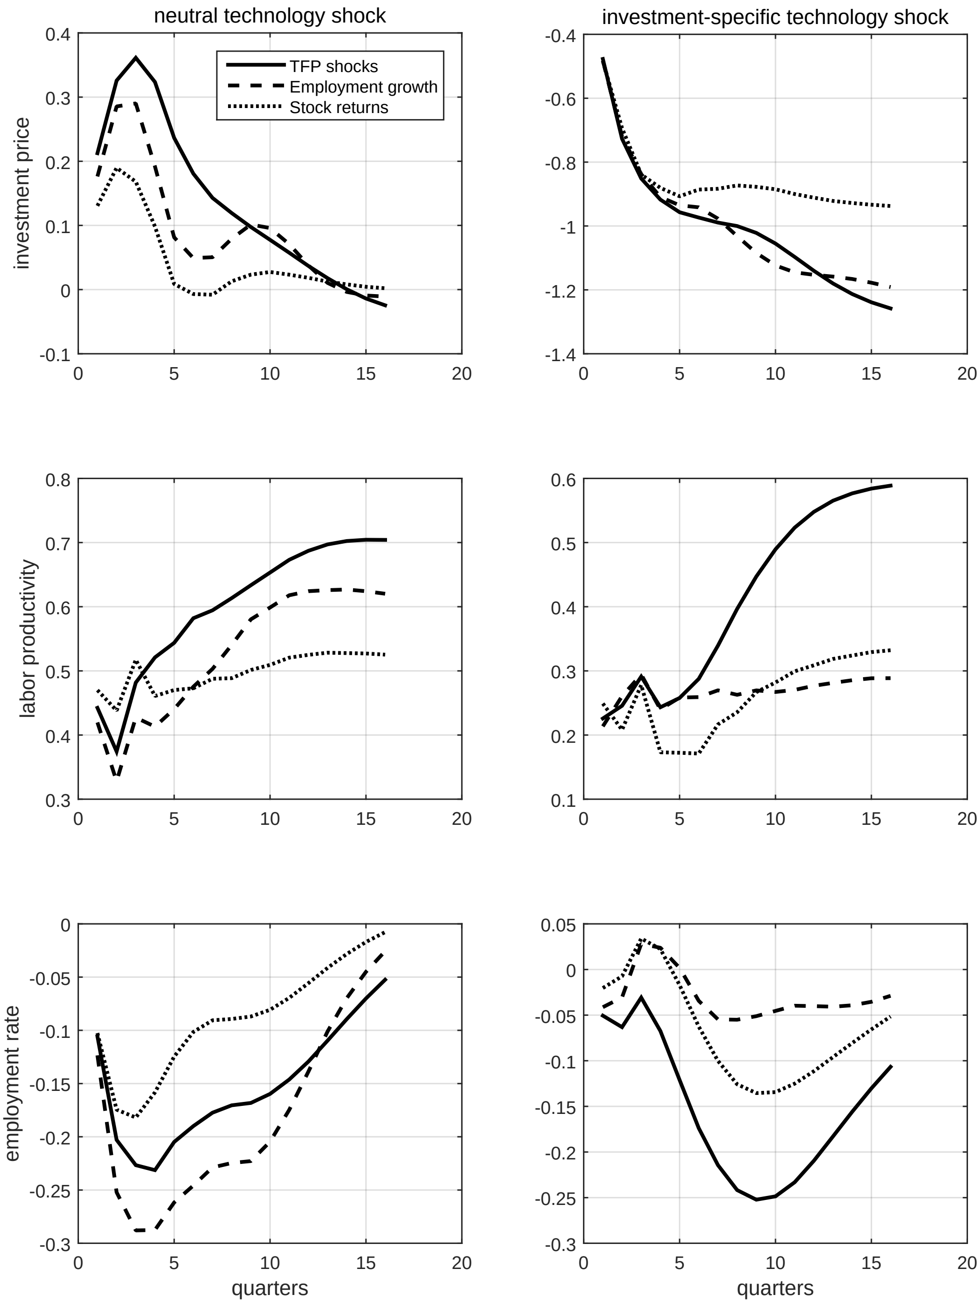

Notes: impulse response functions (in percent) of the investment price (top row), labor productivity (middle row) and the employment rate (bottom row) to positive one-standard-deviation neutral (left column) and investment-specific (right column) shocks in the three estimated structural VARs in Section 5. The legend refers to the 4th variable, i.e. the specific uncertainty proxy, in the respective VARs, which is the only one different across estimations.

Figure 10: Impulse responses to an investment-specific technology shock: uncertainty

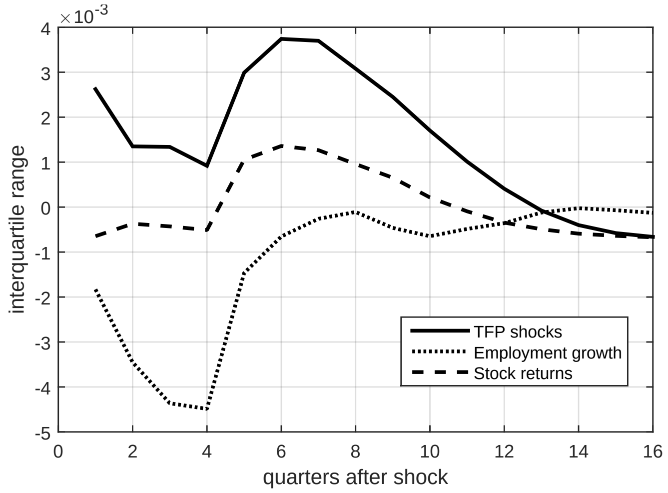

Notes: impulse response functions of the employment rate in the six estimated structural VARs to a positive one-standard-deviation investment specific shock. The legend refers to the 4th variable in the respective VAR, which is the only one different across estimations.

C.3 Alternative uncertainty measures

The main text provided results for uncertainty measures based on the dispersion of firm- or establishment-level variables. Not only are such measures commonly used in the literature, but the structural model of the main text has direct implications for such proxies. However, there are other measures of firm-level uncertainty available and one could investigate whether also such measures react to technology shocks, despite the model not being able to provide a direct theoretical explanation.

Figure 11 shows IRFs of firm-level uncertainty measures constructed by Jurado, Ludvigson, and Ng (2015) which are based on the common variation in the unforecastable component of a rich set of firm-level variables. The difference in the uncertainty proxies is merely in the forecasting horizon they refer to. In all cases, also these uncertainty measures increase significantly and persistently following a positive technology shock when considering one-standard-deviation confidence bands.

Figure 11: Impulse responses to a technology shock: alternative uncertainty proxies

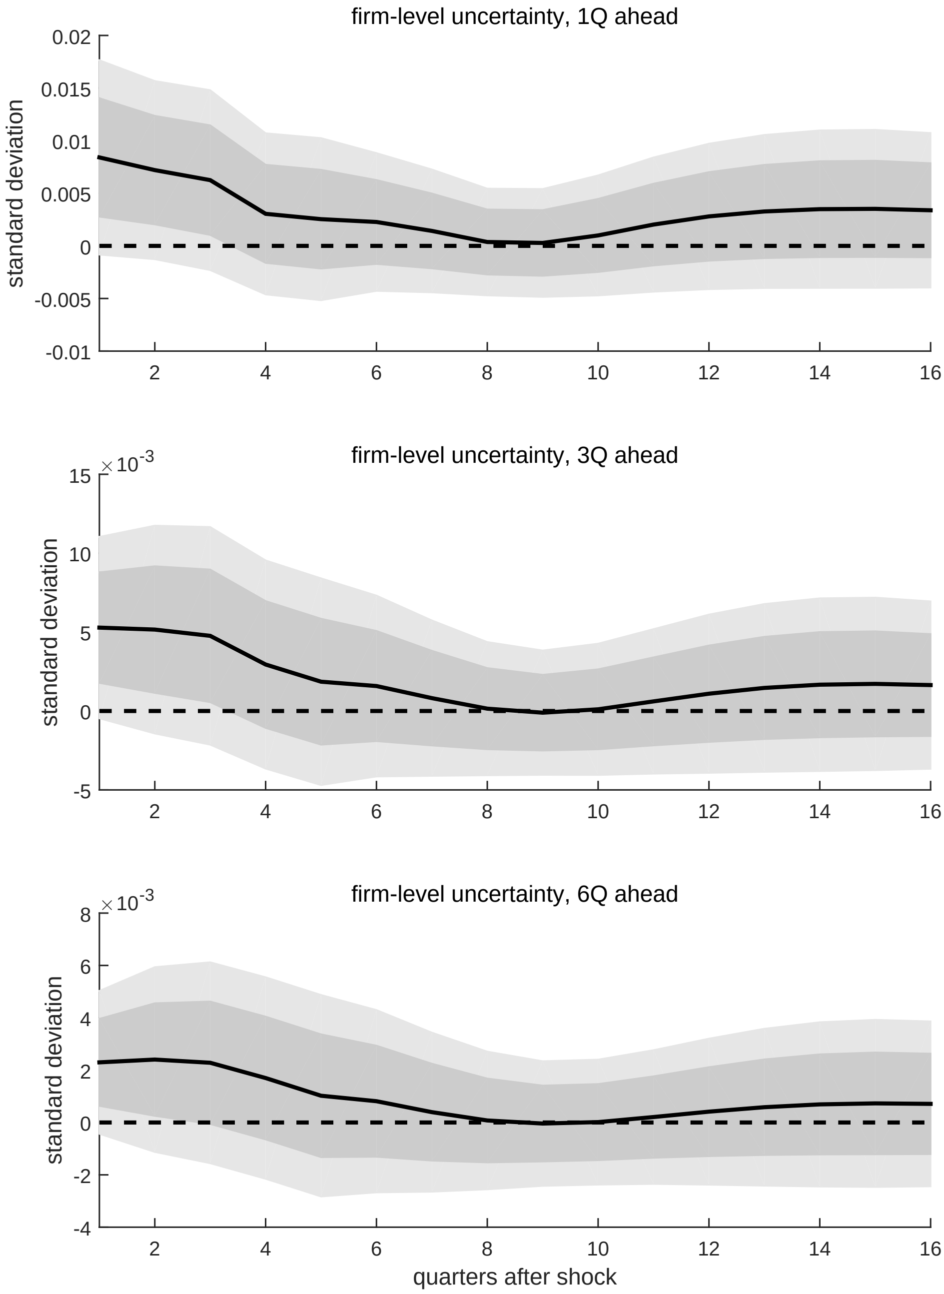

Notes: impulse response functions of firm-level uncertainty measures by Jurado, Ludvigson, and Ng (2015) to a positive one-standard-deviation neutral technology shock. Shaded areas represent one-standard deviation and 90% confidence bands, respectively.

Figure 12: Impulse responses to a technology shock: uncertainty

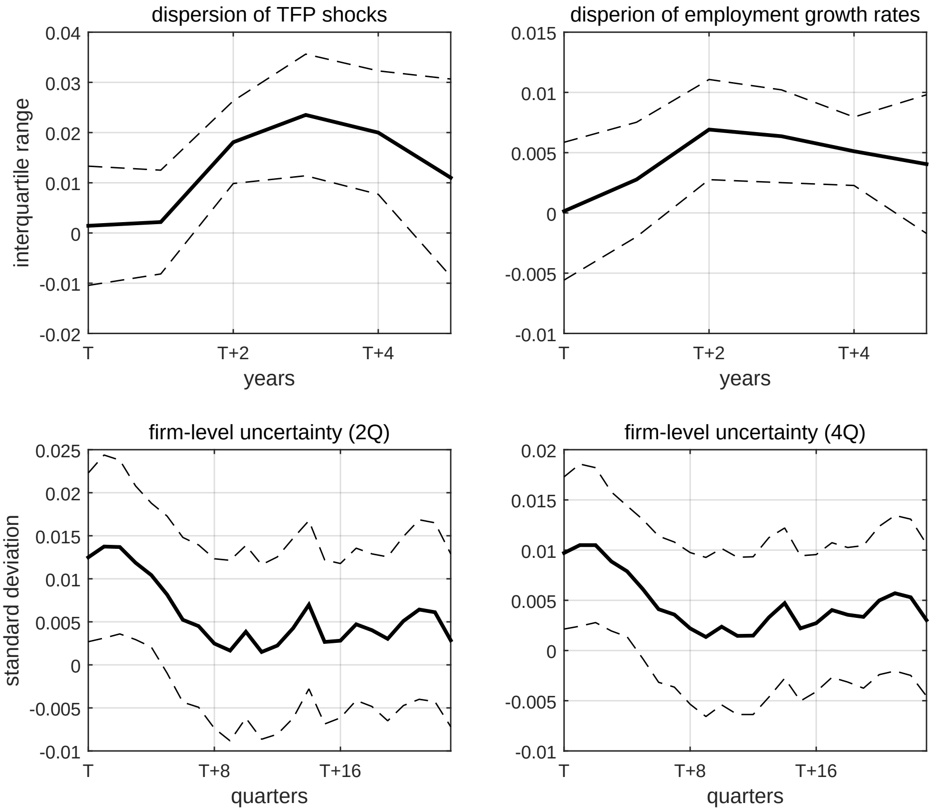

Notes: impulse response functions estimated using local projections and one-standard-deviation confidence bands (dashed lines). The top row displays, respectively, the responses of the interquartile range of establishment-level TFP shocks and employment growth rates taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). The bottom row shows the responses of firm-level uncertainty at a 2 and 4 quarter horizon constructed by Jurado, Ludvigson, and Ng (2015). The dashed lines represent one-standard-deviation bands.

C.4 Alternative empirical strategy

As an alternative empirical strategy it is possible to 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). This subsection shows that such an empirical strategy delivers very similar results.

In particular, the top row of Figure 12 shows the impulse responses of two popular annual measures of firm-level uncertainty taken from Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014), namely the interquartile range of establishment-level TFP shocks and of establishment-level employment growth rates. The figure shows that in both cases measured uncertainty increases persistently following a technological improvement. The bottom row of Figure 12 shows that also the uncertainty measures constructed by Jurado, Ludvigson, and Ng (2015) increase significantly upon a positive technology shock, albeit somewhat less persistently than the dispersion measures constructed by Bloom, Floetotto, Jaimovich, Saporta-Eksten, and Terry (2014). In all these cases the error bands are considerably wider than in the benchmark specification. This is because the results presented here are based on the annual data, rather than the mixed-frequency estimation.

C.5 Evidence on establishment dynamics

The empirical results in the main text concern firms, rather than establishments. Figure 13 shows that results for establishment dynamics are very similar to those of firms. The only difference being that the responses of job creation in new establishments and job destruction in exiting establishments increases somewhat less significantly than that of firms. The discrepancy comes from the fact that a part of the new and exiting establishments account for job creation and destruction among continuing firms.

Figure 13: Impulse responses to a technology shock: establishment dynamics

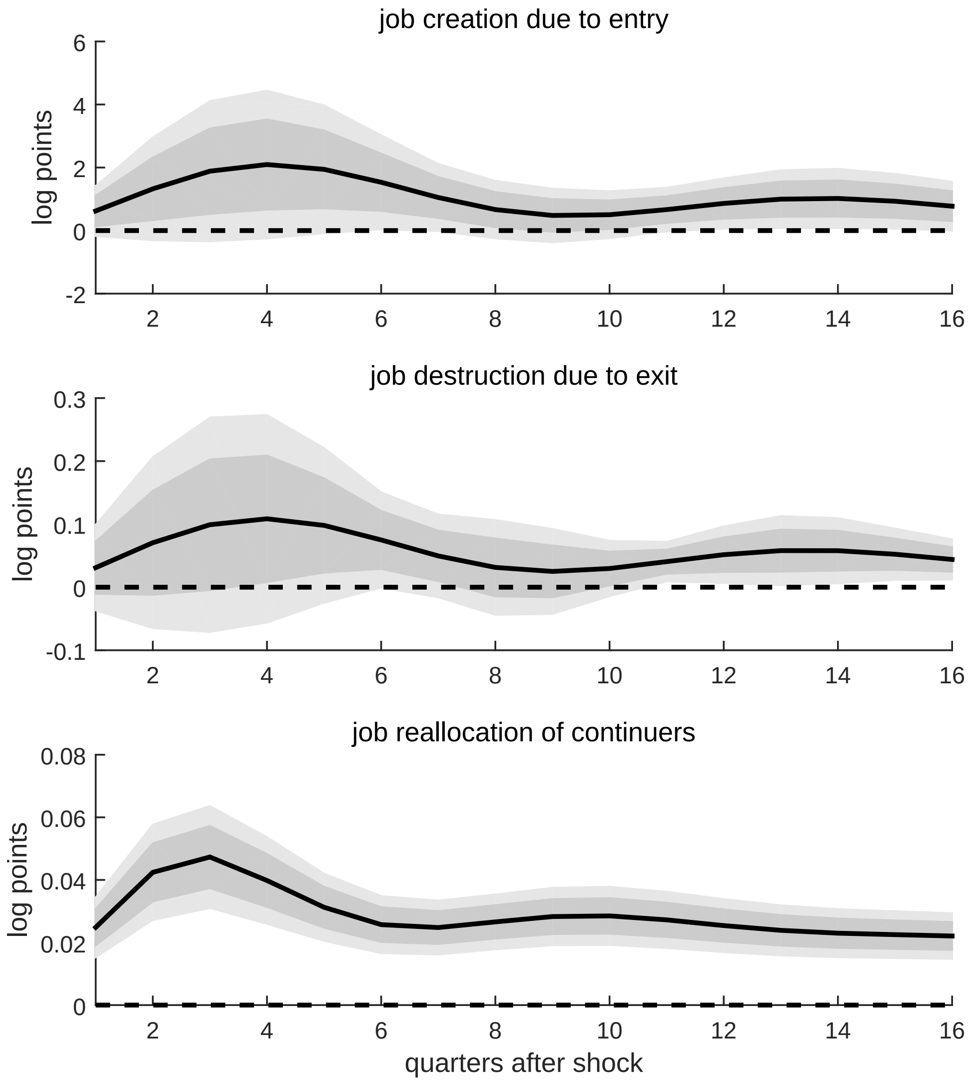

Notes: impulse response functions to a positive neutral technology shock. "Job reallocation" is defined as the sum of job creation and destruction. The light and dark shaded areas indicate one-standard-deviation and 90% confidence intervals, respectively.

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 (2015): "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., S. Elstner, and A. Hristov (2014): "Surprise, Surprise - Measuring Firm-level Investment Innvoations," mimeo.

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

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

Barlevy, G. (2004): "The Cost of Business Cycles under Endogenous Growth," *American Economic Review*, 94(4), 964–990.

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.

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

Bloom, N. (2007): "Uncertainty and the Dynamics of R&D," *American Economic Review: Papers and Proceedings*, 97(2), 250–255.

——— (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.

Bloom, N., M. Schankerman, and J. van Reenen (2013): "Identifying Technology Spillovers and Product Market Rivalry," *Econometrica*, 81(4), 1347–1393.

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.

Christiano, L., M. Eichenbaum, and R. Vigfusson (2003): "What Happens After a Technology Shock?," NBER Working Paper No. 9819.

Christiano, L., R. Motto, and M. Rostagno (2010): "Financial Factors in Economic Fluctuations," ECB Working Paper No. 1192.

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.

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.

Gilchrist, S., J. Sim, and E. Zakrajšek (2014): "Uncertainty, Financial Frictions, and Investment Dynamics," NBER Working Paper No. 20038.

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.

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.

Kehrig, M. (2015): "The Cyclical Nature of the Productivity Dispersion," mimeo.

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

Kraft, H., E. Schwartz, and F. Weiss (2013): "Growth Options and Firm Valuation," NBER Working Paper No. 18836.

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 (2016): "Uncertainty and Business Cycles: Exogenous Impulse or Endogenous Response?," mimeo.

Mecikovsky, A., and M. Meier (2015): "Do Plants Freeze Upon Uncertainty Shocks?," 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.

Riegler, M. (2014): "The Impact of Uncertainty Shocks on the Job-Finding Rate and Separation Rate," mimeo.

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

Schaal, E. (2012): "Uncertainty, Productivity and Unemployment in the Great Recession," mimeo.

Sedláček, P. (2015): "The Lost Generation of Firms and Aggregate Labor Market Dynamics," mimeo.

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.

Vavra, J. (2014): "Infation Dynamics and Time-Varying Volatility: New Evidence and an Ss Interpretation," *Quarterly Journal of Economics*, 129(1), 215–258.