Firm-Specific Productivity Risk over the Business Cycle: Facts and Aggregate Implications

Ruediger Bachmann Christian Bayer*

February 15, 2009

Abstract

Is time-varying uncertainty a major cause or amplifier of the business cycle? This paper investigates this question in the context of a heterogeneous-firm RBC model with persistent firm-level productivity shocks and lumpy capital adjustment, where cyclical changes in uncertainty correspond naturally to cyclical changes in the cross-sectional dispersion of firm-specific Solow residual innovations. We use a unique German firm-level data set to investigate the extent to which firm-level uncertainty varies over the cycle. This allows us to put empirical discipline on the numerical simulations. We find that, while firm-level uncertainty is indeed countercyclical, it does not fluctuate enough to strongly alter the dynamics of a standard RBC model with only first moment shocks. The changes we do find are mainly caused by a bad news effect: higher uncertainty today predicts lower aggregate Solow residuals tomorrow. This effect dominates the real option value effect of time-varying uncertainty.

JEL Codes: E20, E22, E30, E32.

Keywords: Ss model, RBC model, cross-sectional firm dynamics, lumpy investment, countercyclical risk, aggregate shocks, idiosyncratic shocks, heterogeneous firms, news shocks, uncertainty shocks.

*Respectively: University of Michigan; Bonn University. We are grateful to seminar/meeting participants at the ASSA (San Francisco), Duke, Georgetown, Innsbruck, Wisconsin-Madison, Minneapolis FED, Michigan-Ann Arbor, Zuerich, for their comments. We thank the staff of the Research Department of Deutsche Bundesbank for their assistance. Special thanks go to Timm Koerting for excellent research assistance.

1 Introduction

> *"Crises feed uncertainty. And uncertainty affects behaviour, which feeds the crisis. [...] But there is more at work. If you think that another Depression might be around the corner, better to be careful and save more."*
>
> IMF Chief Economist Olivier Blanchard in: The Economist, Jan. 29, 2009

Is time-varying uncertainty a major cause or amplifier of the business cycle? This paper investigates this question in the context of a heterogeneous-firm RBC model with persistent firm-level productivity shocks and lumpy capital adjustment. We make three contributions to the existing literature: first, we discipline the calibration of time-varying uncertainty as a driving force through a detailed analysis of the behavior of the cross-section of firms over the cycle; secondly, we analyze the implications of time-varying uncertainty in the context of a general equilibrium framework; and, thirdly, we show a new mechanism through which uncertainty shocks have aggregate effects: a bad news effect, rather than the time-varying option value effect that has been highlighted thus far (Dixit and Pindyck, 1994, and, more recently, Bloom, 2009). We find the latter to be dominated by the former. In the aforementioned model context, cyclical changes in uncertainty correspond naturally to cyclical changes in the cross-section of firms – more specifically the dispersions of change rates of firm-level real sales, real value added and Solow residuals.

Thus, using the balance sheet data set of Deutsche Bundesbank (USTAN) – a unique private sector, annual, firm-level data set that allows us to investigate 26 years of data (1973-1998), in which the cross-sections of the panel have over 30,000 firms per year on average, and which has a broad ownership, size and sectoral coverage –, we first show that the cross-sectional standard deviations of the firm-level innovations in the Solow residual, value added and sales are indeed robustly and significantly countercyclical. This confirms US-related findings by Bloom (2009) and Bloom et al. (2009) for the implied volatility of the stock market, for Compustat data on publicly traded firms and for sectoral data. We interpret these changes as countercyclical fluctuations in firm-specific productivity uncertainty. Second, we show that for our broad German cross-section of firms the volatility in uncertainty is substantially lower than has been considered in the literature. These two results are robust to different choices for the cyclical indicator, to alternative calculations of the Solow residual, and to various changes in the sample selection criteria. Figure 1 illustrates the countercyclicality of firm-level risk. It displays the distribution of the log-difference in the firm specific Solow residual for the boom and the recession years 1979 and 1975, respectively.

Figure 1: Non-parametric estimates of the Solow residual-change distribution in booms and recessions

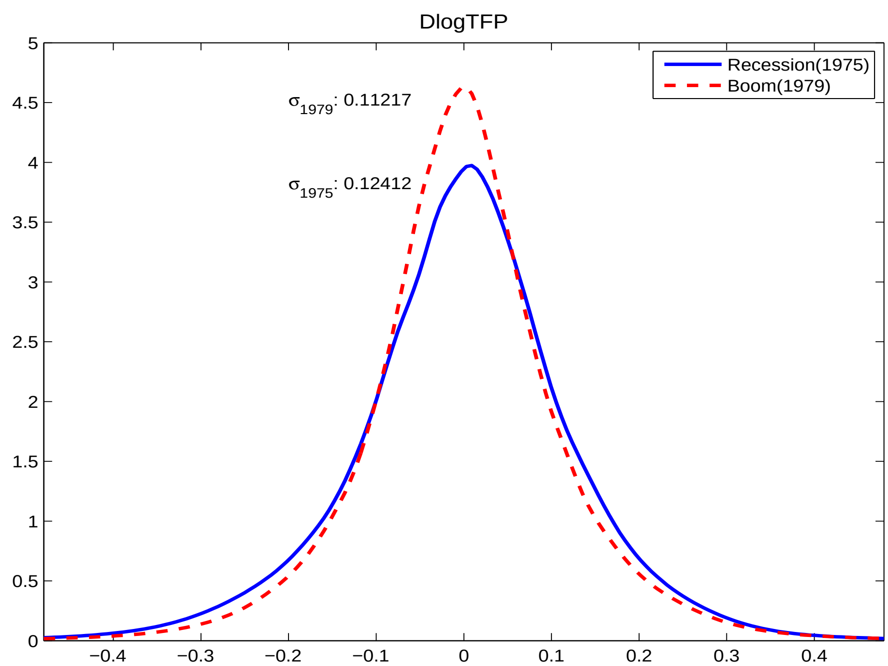

Data is centered around time mean. Densities are estimated using an adaptive variable bandwidth (see Pagan and Ullah, 1999, pp. 12) kernel density estimator, using a normal kernel. The pilot bandwidth is chosen by Silverman's (1986) rule of thumb and data dependent bandwidths are based on a parametric normal estimate of the density.

We then explore the quantitative importance of these shocks to uncertainty in a heterogeneous-firm RBC model similar to the one in Khan and Thomas (2008) and Bachmann et al. (2008). We have developed the computational techniques to study time-varying uncertainty in general equilibrium in parallel with Bloom et al. (2009). We find that uncertainty shocks alone - if of a size in line with our empirical evidence - do not introduce quantitatively significant business cycle fluctuations. Moreover, uncertainty driven business cycles lead to counterfactual output-consumption correlations.

Shocks to uncertainty nonetheless modestly alter the aggregate behavior of the model if they are introduced alongside standard first moment fluctuations in aggregate productivity and compared to a model with only these first moment shocks (essentially the standard RBC model). Yet, this change does not come through the real options effect often associated with time-varying uncertainty. Instead, we identify a new channel by which fluctuations in uncertainty influence the business cycle that results only from their correlation with aggregate productivity. Since changes in uncertainty empirically correlate with future developments in the aggregate Solow residual, the former constitute news about the latter.

Specifically, we find that an increase in productivity risk essentially acts as a bad news shock. In general equilibrium, this means that households decrease their consumption and increase their labor supply. In fact, we find a decrease in the real wage as a result of a shock to uncertainty not only in our model but also in the impulse response function of the actual data. As a result of this decrease in wages, we find – again both in the data as well as in our model – an increase in economic activity at the moment of the uncertainty shock, followed by a recession, when the predicted downturn in productivity occurs. The effects of uncertainty shocks that we find in our model almost entirely operate through this channel.

Comparing a partial equilibrium version to our general equilibrium model in terms of the impulse response functions of investment with respect to shocks to uncertainty illustrates this mechanism further. While in partial equilibrium the bad news on impact causes aggregate investment activity to collapse, in general equilibrium, this effect is reversed by the decrease in the real wage, such that investment rates even expand on impact and only turn negative as lower productivity actually realizes. This shows that general equilibrium analysis is paramount in understanding the aggregate effects of firm-level uncertainty shocks.

The remainder of this paper is organized as follows: Section 2 describes our data set, the USTAN data, discusses briefly the selection of the final sample and details our empirical results. Section 3 explains the model. Section 4 describes its calibration and Section 5 discusses the numerical results. Appendices provide more details on the data set, give further robustness checks for the empirical findings as well as well as for the simulation results.

Related Literature

Over the last two decades, the economic literature has used models with fixed costs of capital adjustment or irreversibility as a natural starting point to study the effects of uncertainty on economic activity (see e.g. Dixit and Pindyck, 1994). More recently, the literature has highlighted the effects of differences in uncertainty both cross-sectionally as well as across time. For example, Bloom (2009) and Bloom et al. (2009) document that increases in stock market volatility are correlated with a reduction in aggregate economic activity. The former in partial equilibrium, the latter in general equilibrium then provide a formal model qualitatively similar to the one used here, where an increase in uncertainty leads to a higher real option value of investment, letting more firms 'wait and see' after such an increase, which in turn leads to a fall in aggregate investment and employment. In a similar vein and also in general equilibrium, but in a different model context, Sim (2008) puts forth a model that explains the cyclical patterns in firms' entry and exit with cyclical variations in uncertainty. None of the above draws its empirical evidence on broad-based cross-sectional firm data. Instead, stock market data, analyst data available only for a subset of large publicly traded companies or sectoral data is used. Using stock market data to identify fluctuations in productivity, however, implies strong market efficiency assumptions (see for instance Shiller, 1981). The empirical part of this paper is most closely related to a series of papers by Higson and Holly et al. (2002, 2004), Doepke and Holly et al. (2005, 2008), Doepke and Weber (2006), as well as Holly and Santoro (2008). Higson and Holly et al. (2002), using Compustat data, study empirically the cyclicality of the standard deviation, skewness and kurtosis of the sales growth rate distribution and find them to be countercyclical, countercyclical and procyclical, respectively. Higson and Holly et al. (2004) repeat this analysis for UK data on quoted firms, and Doepke and Holly et al. (2005) for Germany, using the USTAN database, with similar findings. Doepke and Weber (2006) study, again using USTAN data, the cyclicality of transitions between sales growth regimes in firm-level data. In contrast to these papers, we focus on the cyclicality of cross-sectional second moments only, but include real value added and Solow residuals into the analysis. The quantitative part of this paper draws heavily on the recent literature on heterogenous-firm RBC models, developed in Khan and Thomas (2008) as well as Bachmann et al. (2008). Finally, this paper is related to the work by Beaudry and Portier (2006), Jaimovich and Rebelo (2008), Sims (2008) as well as Schmidt-Grohe and Uribe (2008) on the impact of news shocks on business cycle dynamics. In a companion paper (Bachmann and Bayer, 2009), we focus on the implications of countercyclical dispersion in the firm-level Solow residual innovations for cross-sectional dynamics as opposed to aggregate dynamics.

2 The Facts

In Section 2.1 we briefly describe the USTAN data set and the main sample selection criteria we use. Details are relegated to Appendix A. In Section 2.2.1 we present the baseline fact: the contemporaneous correlations of cyclical aggregate output and the cross-sectional standard deviations of firm-level Solow residual and real value added innovations as well as sales change rates are negative, but the cyclical variations are small in size. In Section 2.2.2 we perform extensive robustness checks and also show, that these facts do not depend on observable firm characteristics.

2.1 A Brief Data Description

2.1.1 USTAN Data

USTAN is a large annual firm-level balance sheet data base (*Unternehmensbilanzstatistik*) collected by *Deutsche Bundesbank*. It is unique in its size and coverage. It provides annual firm level data from 1971 to 1998 from the balance sheets and the profit and loss accounts of over 60,000 firms per year (see Stoess (2001) and von Kalckreuth (2003) for further details). In the days when the discounting of commercial bills were one of the principal instruments of German monetary policy, Bundesbank law required the Bundesbank to assess the credit-worthiness of all parties backing a commercial bill put up for discounting. The Bundesbank implemented this regulation by requiring balance sheet data of all parties involved. These balance sheet data were then archived and collected into a database.

Although the sampling design – one's commercial bill being put up for discounting – does not lead to a perfectly representative selection of firms in a statistical sense, the coverage of the sample is very broad. USTAN covers incorporated firms as well as privately-owned companies, which sets it apart from Compustat data.[1] Its sectoral coverage – while still somewhat biased to manufacturing firms – includes the construction, the service as well as the primary sectors. This makes it different from, for instance, the Annual Survey of Manufacturing (ASM) in the U.S.[2] The following table 1 displays the sectoral coverage of our final baseline sample.

Moreover, while there remains a bias to somewhat larger and financially healthier firms, the size coverage is still fairly broad: 31% of all firms in our final baseline sample have less than 20 employees and 57% have less than 50 employees (see Appendix A for details). Finally,

[1] Davis et al. (2006) show that studying only publicly traded firms can lead to wrong conclusions, in particular when higher cross-sectional moments are investigated.

[2] An additional advantage of these data is easy access: while access is on-site, it is practically free for researchers, so that results derived from this data base can be easily tested and checked.

Table 1: Sectoral Coverage

1-digit Sector	Firm-year observations	Percentage
Agriculture	12,291	1.44
Mining & Energy	4,165	0.49
Manufacturing	405,787	47.50
Construction	54,569	6.39
Trade (Retail & Wholesale)	355,208	41.59
Transportation & Communication	22,085	2.59

the Bundesbank itself frequently uses the USTAN data for its macroeconomic analyses and for cross-checking national accounting data. We take this as an indication that the bank considers the data sufficiently representative and of sufficiently high quality.[3] This makes the USTAN data a uniquely suitable data source for the study of cross-sectional business cycle dynamics.

2.1.2 Selection of the Baseline Sample

From the original USTAN data, we select only firms that report complete information on payroll, gross value added and capital stocks. Moreover, we drop observations from East German firms to avoid a break of the series in 1990.[4] In addition, we remove observations that stem from irregular accounting statements, e.g. when filing for bankruptcy or when closing operations. We deflate all but the capital data by the implicit deflator for gross value added from the German national accounts.

Capital is deflated with one-digit sector- and capital-good specific investment good price deflators within a perpetual inventory method. Even though USTAN data can be considered as particularly high quality data, we cannot directly use capital stocks as reported. Tax motivated depreciation and price developments of capital goods lead to a general understatement of the stock of capital a firm holds. Thus, capital stocks have to be recalculated using a perpetual inventory method. Similarly, we recover the amount of labor inputs from wage bills, as information on the number of employees (as opposed to payroll data) is only updated infrequently for some companies (see Appendix A for details). Finally, the firm-level Solow residual has to be calculated from data on gross value added and factor inputs.

We remove outliers according to the following procedure: we calculate log changes in real

[3] For example, data entry of the balance sheet data is double-checked by a second Bundesbank staff member.

[4] We identify a West German firm as a firm that has a West German address or has no address information but enters the sample before 1990.

gross value added, the Solow residual, real capital and employment and drop all observations where a change falls outside a 3 standard deviations interval around the year-specific mean.[5] We also drop those firms for which we do not have at least five observations in first differences. This leaves us with a sample of 854,105 firm-year observations, which corresponds to observations on 72,853 firms, i.e. the average observation length of a firm in the sample is 11.7 years. The average number of firms in the cross-section of any given year is 32,850. We perform numerous robustness checks with respect to each of these choices: we use sectoral deflators for value added, an aggregate investment good price deflator, change the cut-off rule to 2.5 and 5 standard deviations and leave all firms in the sample with two and twenty observations in first differences, respectively. None of these choices change our baseline results (see Appendix B for details).

2.1.3 Solow Residual and Productivity Innovations

We compute the firm-level Solow residual based on the following Cobb-Douglas production function in accordance with our model:

$$y_t = z_t \epsilon_t k_t^{\theta} n_t^{\nu},$$

where ϵ_t is the firm-specific component and z_t is aggregate productivity. We assume that labor input n_t is immediately productive, whereas capital k_t is pre-determined and inherited from last period. In our main specification, we estimate the output elasticities of the production factor, ν and θ, as median shares of factor expenditures over gross value added within each industry.[6] We use log-differences in the Solow residual to capture Solow residual innovations, as the persistence of firm-level Solow residuals exhibits behavior close to a unit root. Alternatively we measure productivity innovations (and hence productivity risk) by looking at log-differences in value added and sales. We remove firm fixed and sectoral-year effects from these first-difference variables to focus on idiosyncratic fluctuations that do not capture differences in sectoral responses to aggregate shocks or permanent ex-ante heterogeneity between firms.[7]

[5] This outlier removal is done after removing firm and sectoral fixed effects. Centering the outlier removal around the year mean is important to avoid artificial and countercyclical skewness of the respective distributions.

[6] To check the robustness of our results, we try alternative specifications with predefined elasticities common across sectors. We also change the timing assumption to include a predetermined employment stock, as well as immediate adjustment in both factors. All results are very robust to the alternative ways to generate the firm-specific Solow residual (for a detailed discussion, see the appendix).

[7] We decomposed the outcome of firm i in sector j at time t into firm fixed effects α_i, sectoral time effects

2.1.4 Macro data

When combining this micro data with aggregate data, we have to take a stance on what sectoral aggregate we view as the empirical counterpart to our model. We chose to include firms from the following six sectors in our analysis: agriculture, mining and energy, manufacturing, construction, trade (both retail and wholesale) as well as the transportation and communication sector. This aggregate can be roughly characterized as the non-financial private business sector in Germany. Whenever we use the term aggregate in the following, we mean this sector.

German national accounting data per one-digit sector allow us to compute real value added, investment, capital and employment data for this sectoral aggregate, and therefore also an aggregate Solow residual. Real private consumption data are taken from table 3.2 in the *Volkswirtschaftliche Gesamtrechnungen*, source: *Statistisches Bundesamt*. Our USTAN sample captures on average 70% of sectoral value added, 44% of sectoral investment, 71% of its capital stock and 49% of sectoral employment.

In addition to representing a large part of the non-financial private business sector in Germany, USTAN also represents its cyclical behavior very well, see the right hand panel in Table 2.

2.2 Main Facts

2.2.1 Idiosyncratic Productivity Risk: Large and Countercyclical

Since firm-level productivity has an autocorrelation close to one - even after controlling for fixed effects -, we use *first differences* in the Solow residual, real value added and real sales to measure stochastic innovations that firms face. We focus on the cross-sectional dispersion of these variables and measure the dispersion in terms of standard deviations.[8] In a neoclassical model, these standard deviations characterize the uncertainty (or 'risk' which we use interchangeably) of a firm with regard to its idiosyncratic productivity growth. In line with this, the literature (e.g. Bloom et al., 2007, Sim, 2008, Bloom et al. , 2009, and Bloom, 2009) has modeled fluctuations in uncertainty as fluctuations in the standard deviation of firm-specific innovations in the Solow residual. We follow this view and investigate the quan-

γ_{jt} and an idiosyncratic unpredictable error term ε_{ijt} which are orthogonal to each other

$$x_{ijt} = \alpha_i^x + \gamma_{jt}^x + \varepsilon_{ijt}^x.$$

We then focus on the error term ε_{ijt}.

[8]Measuring the dispersion in terms of interquartile ranges gives no different results. See Appendix B for details.

Table 2: Business Cycle Statistics of Cross Sectional Distributions

	Solow residual $\Delta\epsilon_{it}$	value added Δy_{it}	sales $\Delta sales_{it}$		Solow residual $\Delta\epsilon_{it}$	value added Δy_{it}	sales $\Delta sales_{it}$
$E\left(s_t^x\right)$	12.0	14.2	18.7				
$\sigma\left(s_t^x\right)$	0.321	0.528	0.715	$\sigma\left(\mu_t^x\right)$	2.374	3.136	1.427
$\sigma\left(s_t^x\right)/E\left(s_t^x\right)$	2.671	3.725	3.823				
$\rho\left(s_t^x, Y_t\right)$	-0.481	-0.450	-0.405	$\rho\left(\mu_t^x, Y_t\right)$	0.592	0.663	0.410
$\rho\left(s_t^x, z_t\right)$	-0.624	-0.563	-0.313	$\rho\left(\mu_t^x, z_t\right)$	0.657	0.731	0.655
$ac\left(s_t^x\right)$	0.539	0.597	0.191	$ac\left(\mu_t^x\right)$	0.225	0.335	0.181

s_t^x : cross-sectional standard deviation of variable x;
$E\left(s_t^x\right)$: time-series average of s_t^x, multiplied by 100;
μ_t^x : cross-sectional mean of variable x;
$\sigma\left(x\right)$: time series standard deviation (after linear detrending), multiplied by 100.
Y_t HP-100 filtered gross aggregate output, z_t HP-100 filtered aggregate Solow residual.
$\rho\left(x_1, x_2\right)$: correlation of x_1 and x_2, $ac\left(x\right)$ autocorrelation after linear detrending.

titative aggregate importance of fluctuations in the cross-sectional dispersion of productivity innovations.

The left-hand panel of Table 2 presents the average cross-sectional standard deviations and their cyclical behavior for the three variables of interest. The right-hand panel of Table 2 displays the time-series volatility (in terms of standard deviations) and the cyclical correlations of the sample means for comparison. We find that idiosyncratic productivity risk (in terms of standard deviations) is 12% and thus large compared to the fluctuations in the average Solow residual.[9] It is significantly countercyclical (both in terms the correlation with the aggregate Solow residual and aggregate value added (-62% and -48% respectively) and its fluctuations are relatively persistent. However, the size of annual fluctuations in risk are relatively small, with a coefficient of variation of roughly 2.7%.[10]

One could be concerned that the Solow residual does not measure productivity sufficiently

[9]The volatility of the linearly detrended aggregate Solow residual is 2.06%. This order-of-magnitude difference also suggests that the fluctuations-in-uncertainty-story has to be likely one about idiosyncratic risk, rather than aggregate risk. Firms' policy functions are mainly determined by the idiosyncratic state given this difference, and fluctuations in anyway small aggregate risk do not seem to be a plausible candidate for business cycle fluctuations.

[10]For comparison, Bloom et al. (2009) find a correlation of their uncertainty index with GDP of -49% and Bloom (2009) considers a productivity process for monthly data, which exhibits a time-series coefficient of variation of the dispersion in annual productivity growth of roughly 17.4%.

well. For this reason, we provide evidence also from log-differences in real value added and real sales, which as balance sheet items are less subject to measurement error. Measuring productivity risk in terms of the cross-sectional dispersion in these variables provides no different results: idiosyncratic risk is large, it is significantly countercyclical, but the business cycle fluctuations of risk are small. In the next subsection we show that all these findings in Table 2 are robust to subsamples stratified by observable firm characteristics, which means they are not a result of the sectoral or size composition of our sample.

2.2.2 Robustness

Sample Splits USTAN is no random sample. Consequently, we need to check whether our findings are specific to certain sectors or types of firms overrepresented in this sample. Table 3 shows that this is not the case. The left panel of the table displays the behavior of productivity risk for the three main single-digit sectors (which together represent 96% of all firms in the data). Both the relatively large idiosyncratic risk and the countercyclicality of risk is not specific to any sector. More importantly, also the result that risk fluctuations are not large is robust across sectors. Manufacturing and Construction exhibit somewhat larger risk fluctuation, but the difference to the trade sector is small. The range across industries of the time series coefficient of variation of between 2.6% and 3.6% with the average standard deviation in idiosyncratic productivity growth being between 11.2% and 12.4%.

The non random design of the USTAN sample manifests itself not only in the sectoral but also the size composition of the sample. To understand the effect of this, we split the sample according to firm size. Again this provides no different results. The right panel of Table 3 gives results for the smallest 33% and the largest 5% of firms in each year. Large firms exhibit a somewhat smaller average risk that fluctuates more than the risk of small firms, though.[11]

Finally, we focus on those firms who did not change the intensity of material usage between two periods. This sample split eliminates potential effects of changes in capacity utilization that might have been falsely attributed to productivity fluctuations in our specification of the production function. Also for this group of firms our results carry over in general. These firms exhibit a somewhat smaller productivity risk that fluctuates slightly more than the productivity risk of the average firm in the sample, but the difference is again small.[12]

[11] The large firm sample is closer in terms of typical firm size to data bases on publicly traded firms like COMPUSTAT. We tried alternative splits according to the stock of capital and value added and taking different cut-off values to define the group of small firms. Results are very similar, details can be found in Appendix B.

[12] The same results holds for restricting the sample to those firms that did not change the absolute usage

Table 3: Business Cycle Statistics of Cross Sectional Standard Deviations: Results by Sector and Firm Size

		$\Delta\epsilon_{it}$	Δy_{it}	$\Delta sales_{it}$			$\Delta\epsilon_{it}$	Δy_{it}	$\Delta sales_{it}$
Manu-	$E\left(s_t^x\right)$	11.5	14.0	17.0	Small	$E\left(s_t^x\right)$	13.7	15.8	20.4
facturing	$\frac{\sigma(s_t^x)}{E(s_t^x)}$	3.543	5.028	4.664	Firms[1]	$\frac{\sigma(s_t^x)}{E(s_t^x)}$	2.006	2.663	4.391
	$\rho\left(s_t^x, z_t\right)$	-0.659	-0.617	-0.440		$\rho\left(s_t^x, z_t\right)$	-0.525	-0.443	-0.243
	$ac\left(s_t^x\right)$	0.568	0.570	0.324		$ac\left(s_t^x\right)$	0.495	0.454	0.126
Con-	$E\left(s_t^x\right)$	11.2	15.0	32.8	Large	$E\left(s_t^x\right)$	9.8	12.0	14.6
struction	$\frac{\sigma(s_t^x)}{E(s_t^x)}$	3.565	4.774	6.994	Firms[2]	$\frac{\sigma(s_t^x)}{E(s_t^x)}$	5.550	6.542	5.508
	$\rho\left(s_t^x, z_t\right)$	-0.379	-0.423	0.117		$\rho\left(s_t^x, z_t\right)$	-0.579	-0.519	-0.191
	$ac\left(s_t^x\right)$	0.584	0.591	0.452		$ac\left(s_t^x\right)$	0.347	0.484	-0.166
Trade	$E\left(s_t^x\right)$	12.4	14.1	17.6	constant	$E\left(s_t^x\right)$	10.3	11.7	12.4
	$\frac{\sigma(s_t^x)}{E(s_t^x)}$	2.676	3.477	5.934	material	$\frac{\sigma(s_t^x)}{E(s_t^x)}$	4.192	5.940	6.642
	$\rho\left(s_t^x, z_t\right)$	-0.450	-0.325	-0.308	intensity[3]	$\rho\left(s_t^x, z_t\right)$	-0.444	-0.432	-0.419
	$ac\left(s_t^x\right)$	0.417	0.514	0.411		$ac\left(s_t^x\right)$	0.602	0.669	0.638

See notes to Table 2.
[1] Lowest 33-percentile in employment by year.
[2] Top 5-percentile in employment by year.
[3] Firms whose material expenditures over sales $\frac{M_{it}}{Sales_{it}}$ have changed by less then one percentage point.

Cyclical Sample Selection The bias towards manufacturing is more of a permanent phenomenon in the USTAN data and can be healed by looking at the different sectors separately. The possibility of systematic cyclical variations in sample selection poses a more severe problem to our analysis.

This sample can be thought of as the result of a two-stage selection. First, we can think of firms being selected into the set of active firms. Second from this set of firms there is a selection into the USTAN sample.

The USTAN sample is biased towards financially healthy firms and financial health can be expected to have cyclical fluctuations. However, this sample selection should be important rather for the cross-sectional distribution of *levels* of productivity, employment, and capital. Since we find unit-root behavior of these variables, however, levels have no predictive power for the first-differences on which our analysis focuses.

Sample selection will depend on changes in productivity only insofar as these changes influence the level of productivity. For example, firms that have a lower average productivity-growth rate are more likely drop out of the sample.

We check the importance of these effects in two ways, based on the dispersion of firm specific averages - the estimates of firm fixed effects. First, we check, how large the dispersion of these fixed effects is. The total cross-sectional variance of a variable x, $s_t^2(x)$, is composed of ex-ante heterogeneity due to differences in firm fixed effects, $s_{\alpha,t}^2(x)$,[13] and ex-post heterogeneity due to sectoral differences, $s_{\gamma,t}^2(x)$, and idiosyncratic risk, $s_{\varepsilon,t}^2(x)$. Since the decomposition is orthogonal,

$$s_t^2(x) = s_{\alpha,t}^2(x) + s_{\gamma,t}^2(x) + s_{\varepsilon,t}^2(x)$$

holds for all t asymptotically in i.

Second, we decompose the variation of raw data heterogeneity in variations of risk, variations in fixed effects etc. If the set of firms in the sample was constant or the sample was a random collection of firms from the universe of all German firms, this dispersion would be constant over time.[14]

The results of these experiments are displayed in Table 4 and 5. Compared to within firm changes in productivity ('idiosyncratic risk'), fixed differences in productivity growth

of material, see Appendix B.

[13]Note that we remove fixed effects in growth rates. Therefore we allow for ex-ante heterogeneous time profiles of firms. Ex-ante heterogeneous firm sizes are controlled for by focussing on first differences.

[14]We make use of the decomposition

$$\sigma^2\left(s_{x,t}^2\right) = \sigma^2\left(s_{\alpha,t}^2\right) + \sigma^2\left(s_{\gamma,t}^2\right) + \sigma^2\left(s_{\varepsilon,t}^2\right) + 2\left(cov\left(s_{\alpha,t}^2, s_{\gamma,t}^2\right) + cov\left(s_{\varepsilon,t}^2, s_{\gamma,t}^2\right) + cov\left(s_{\alpha,t}^2, s_{\varepsilon,t}^2\right)\right).$$

As we work with an unbalanced panel, this decomposition holds only approximately.

Table 4: Decomposition of observed heterogeneity

	Solow residual $\Delta\epsilon_{it}$	value added Δy_{it}	sales $\Delta sales_{it}$	Solow residual $\Delta\epsilon_{it}$	value added Δy_{it}	sales $\Delta sales_{it}$
absolute average dispersion				*% of raw dispersion ...*		
raw data, s_t^2	1.582	2.354	3.960	100%	100%	100%
fixed effects, $s_{\hat{\alpha},t}^2$	0.115	0.301	0.360	7.2%	12.8%	9.1%
sectoral effects, $s_{\hat{\gamma},t}^2$	0.024	0.037	0.103	1.5%	1.6%	2.6%
idiosyncratic risk $s_{\hat{\varepsilon},t}^2$	1.444	2.013	3.497	91.3%	85.5%	88.3%

Variances $s_{x,t}^2$ multiplied by 100.

have little dispersion and also fluctuate little both in absolute terms as well as in terms of the coefficient of variation. Therefore, it is unlikely that selection strongly affects our results. Correspondingly, a decomposition the time-series variation in heterogeneity in the raw data shows that the main source of these fluctuations stems from variations in idiosyncratic, within firm uncertainty.

As a final robustness check for cyclical sample attrition, we constrain our sample to those firms which we observe at least 20 times in the sample (see Table 6). This set of firms is in general more selected, e.g. due to a survivor bias, such that large firms are even more overrepresented. The advantage of this sample is that its selection is by construction less subject to cyclical fluctuations. This sample displays the same countercyclical pattern in productivity risk with a somewhat higher volatility of uncertainty. Yet, this increase partly reflects the stronger representation of large firms. On the basis of these three pieces of evidence - little importance of fixed differences in growth rates in explaining variations in heterogeneity and low volatility in these differences, as well as robustness of our results to a sample with surviving firms - we conclude that cyclical sample attrition is not central to our findings.

In Appendix B we provide further robustness checks for our results. In all these robustness checks we find a strongly countercyclical productivity risk with small cyclical fluctuations in this risk. Next, we explore the quantitative implications of this finding in a heterogeneous firm dynamic general equilibrium model.

Table 5: Decomposition of fluctuations in observed heterogeneity

	Solow residual $\Delta\epsilon_{it}$	value added Δy_{it}	sales $\Delta sales_{it}$
Time series fluctuations of fixed effects			
$\sigma\left(s_t^x\right)$	0.052	0.089	0.037
$\sigma\left(s_t^x\right)/E\left(s_t^x\right)$	1.537	1.625	1.495
$\rho\left(s_t^x, Y_t\right)$	0.267	0.191	0.181
$\rho\left(s_t^x, z_t\right)$	0.433	0.284	0.209
$ac\left(s_t^x\right)$	0.640	0.774	0.708

	Solow residual $\Delta\epsilon_{it}$	value added Δy_{it}	sales $\Delta sales_{it}$
Percentage of the time series fluctuations of observed heterogeneity explained by ...			
idiosyncratic risk	73.7%	81.8%	72.4%
fixed effects	2.8%	2.0%	4.7%
sectoral effects	0.3%	0.2%	0.1%
covariance terms	23.2%	16.0%	22.8%

See notes to Table 2.

Table 6: Business Cycle Statistics of Cross Sectional Distributions for Firms with At Least 20 Observations

	Solow residual $\Delta\epsilon_{it}$	value added Δy_{it}	sales $\Delta sales_{it}$
$E\left(s_t^x\right)$	10.6	12.0	14.9
$\sigma\left(s_t^x\right)/E\left(s_t^x\right)$	4.613	7.003	7.900
$\rho\left(s_t^x, Y_t\right)$	-0.341	-0.383	-0.466
$\rho\left(s_t^x, z_t\right)$	-0.491	-0.459	-0.448
$ac\left(s_t^x\right)$	0.761	0.760	0.660

See notes to Table 2.

3 A model with countercyclical risk and fixed cost of capital adjustment.

In this section we describe our model economy. We start with the firm's problem, followed by a brief description of the households and the definition of equilibrium. We conclude with a sketch of the equilibrium computation. Our model follows closely Khan and Thomas (2008) and Bachmann et al. (2008). Since there the model set up is discussed in detail, we will be rather brief here.

The main departure from either models is the introduction of a second exogenous aggregate state. Following Bloom (2009) we assume that firms observe the standard deviation of the distribution of idiosyncratic productivity shocks in period $t+1$, $s_{t+1} = s(\epsilon_{t+1})$, one period ahead, i.e. in period t. Following Khan and Thomas (2008), we approximate this now bivariate aggregate state process with a discrete Markov chain.

3.1 Firms

The economy consists of a unit mass of small firms. We do not model entry and exit decisions. There is one commodity in the economy that can be consumed or invested. Each firm produces this commodity, employing its pre-determined capital stock (k) and labor (n), according to the following Cobb-Douglas decreasing-returns-to-scale production function ($\theta > 0$, $\nu > 0$, $\theta + \nu < 1$):

$$y_t = z_t \epsilon_t k_t^{\theta} n_t^{\nu}, \tag{1}$$

where z_t and ϵ_t denote aggregate and firm-specific (idiosyncratic) technology, respectively.

The idiosyncratic technology process has autocorrelation ρ_I. It follows a Markov chain, whose transition matrix depends on the aggregate state of its time-varying standard deviation, $s_{t+1} = s(\epsilon_{t+1})$. In contrast, its support is independent of the aggregate state. To also capture observed excess kurtosis in the idiosyncratic productivity shocks, we use a mixture of two Gaussian distributions in the Tauchen-approximation algorithm instead of the usual normal distribution.[15]

We denote the trend growth rate of aggregate productivity by $(1-\theta)(\gamma-1)$, so that aggregate y and k grow at rate $\gamma - 1$ along the balanced growth path. From now on we work with k and y (and later C) in efficiency units. The linearly detrended logarithm of aggregate productivity levels, which we also denote by z, as well as linearly detrended $s(\epsilon)$ evolve according to a VAR(1) process, with normal innovations υ that have zero mean and covariance Ω:

$$\begin{pmatrix} \log z_t \\ s_{t+1} - \bar{s} \end{pmatrix} = \varrho_A \begin{pmatrix} \log z_{t-1} \\ s_t - \bar{s} \end{pmatrix} + \upsilon_t, \tag{2}$$

where $\bar{s}$ denotes the steady state standard deviation of idiosyncratic productivity innovations.[16]

[15] Tauchen (1986). For details, see Section 4.

[16] Specifying this process in terms of $\log(s_t)$, in order to avoid negativity of the standard deviation of idiosyncratic productivity shocks is – given its relatively low variability – an unnecessary precaution that

Productivity innovations at different aggregation levels are independent. Also, idiosyncratic productivity shocks are independent across productive units. In contrast, we do not impose any restrictions on Ω and ϱ_A.

Each period a firms draws from a time-invariant distribution, G, its current cost of capital adjustment, $\xi \geq 0$, which is denominated in units of labor. G is a uniform distribution on $[0, \bar{\xi}]$, common to all firms. Draws are independent across firms and over time, and employment is freely adjustable.

At the beginning of a period, a firm is characterized by its pre-determined capital stock, its idiosyncratic productivity, and its capital adjustment cost. Given the aggregate state, it decides its employment level, n, production and depreciation occurs, workers are paid, and investment decisions are made. Then the period ends.

Upon investment, i, the firm incurs a fixed cost of $\omega\xi$, where ω is the current real wage rate. Capital depreciates at a rate δ. We can then summarize the evolution of the firm's capital stock (in efficiency units) between two consecutive periods, from k to k', as follows:

	Fixed cost paid	$\gamma k'$
$i \neq 0$:	$\omega\xi$	$(1-\delta)k + i$
$i = 0$:	0	$(1-\delta)k$

Given the i.i.d. nature of the adjustment costs, it is sufficient to describe differences across firms and their evolution by the distribution of firms over (ϵ, k). We denote this distribution by μ. Thus, $\big(z, s, \mu\big)$ constitutes the current aggregate state and μ evolves according to the law of motion $\mu' = \Gamma(z, s, \mu)$, which firms take as given.

Next we describe the dynamic programming problem of each firm. We will take two shortcuts (details can be found in Khan and Thomas, 2008). First, we state the problem in terms of utils of the representative household (rather than physical units), and denote by $p = p(z, s, \mu)$ the marginal utility of consumption. This is the relative intertemporal price faced by a firm. Second, given the i.i.d. nature of the adjustment costs, continuation values can be expressed without explicitly taking into account future adjustment costs.

Let $V^1(\epsilon, k, \xi; z, s, \mu)$ denote the expected discounted value—in utils—of a firm that is in idiosyncratic state (ϵ, k, ξ), given the aggregate state (z, s, μ). Then the expected value prior to the realization of the adjustment cost draw is given by:

$$V^0(\epsilon, k; z, s, \mu) = \int_0^{\bar{\xi}} V^1(\epsilon, k, \xi; z, s, \mu) G(d\xi). \tag{3}$$

does not change the results. Simulation results are available upon request.

With this notation the dynamic programming problem is given by:

$$V^1(\epsilon, k, \xi; z, s, \mu) = \max_n \{\text{CF} + \max(V_i, \max_{k'}[-\text{AC} + V_a])\}, \tag{4}$$

where CF denotes the firm's flow value, V_i the firm's continuation value if it chooses inaction and does not adjust, and V_a the continuation value, net of adjustment costs AC, if the firm adjusts its capital stock. That is:

$$\text{CF} = [z\epsilon k^\theta n^\nu - \omega(z, s, \mu)n]p(z, s, \mu), \tag{5a}$$

$$V_i = \beta \text{E}[V^0(\epsilon', (1-\delta)k/\gamma; z', s', \mu')], \tag{5b}$$

$$AC = \xi\omega(z, s, \mu)p(z, s, \mu), \tag{5c}$$

$$V_a = -ip(z, s, \mu) + \beta \text{E}[V^0(\epsilon', k'; z', s', \mu')], \tag{5d}$$

where both expectation operators average over next period's realizations of the aggregate and idiosyncratic shocks, conditional on this period's values, and we recall that $i = \gamma k' - (1-\delta)k$. Also, β denotes the discount factor of the representative household.

Taking as given intra- and intertemporal prices $\omega(z, s, \mu)$ and $p(z, s, \mu)$, and the law of motion $\mu' = \Gamma(z, s, \mu)$, the firm chooses optimally labor demand, whether to adjust its capital stock at the end of the period, and the optimal capital stock, conditional on adjustment. This leads to policy functions: $N = N(\epsilon, k; z, s, \mu)$ and $K = K(\epsilon, k, \xi; z, s, \mu)$. Since capital is pre-determined, the optimal employment decision is independent of the current adjustment cost draw.

3.2 Households

We assume a continuum of identical households that have access to a complete set of state-contingent claims. Hence, there is no heterogeneity across households. Moreover, they own shares in the firms and are paid dividends. We do not need to model the household side explicitly (see Khan and Thomas (2008) for details), and concentrate instead on the first-order conditions to determine the equilibrium wage and the intertemporal price.

Households have a standard felicity function in consumption and (indivisible) labor:

$$U(C, N^h) = \log C - AN^h, \tag{6}$$

where C denotes consumption and N^h the household's labor supply. Households maximize the expected present discounted value of the above felicity function. By definition we have:

$$p(z, s, \mu) \equiv U_C(C, N^h) = \frac{1}{C(z, s, \mu)}, \tag{7}$$

and from the intratemporal first-order condition:

$$\omega(z, s, \mu) = -\frac{U_N(C, N^h)}{p(z, s, \mu)} = \frac{A}{p(z, s, \mu)}. \tag{8}$$

3.3 Recursive Equilibrium

A *recursive competitive equilibrium* is a set of functions

$$\Big(\omega, p, V^1, N, K, C, N^h, \Gamma\Big),$$

that satisfy

1. *Firm optimality*: Taking ω, p and Γ as given, $V^1(\epsilon, k; z, s, \mu)$ solves (4) and the corresponding policy functions are $N(\epsilon, k; z, s, \mu)$ and $K(\epsilon, k, \xi; z, s, \mu)$.

2. *Household optimality*: Taking ω and p as given, the household's consumption and labor supply satisfy (7) and (8).

3. *Commodity market clearing*:

$$C(z, s, \mu) = \int z\epsilon k^\theta N(\epsilon, k; z, s, \mu)^\nu d\mu \; - \; \int\int_0^{\bar{\xi}} [\gamma K(\epsilon, k, \xi; z, s, \mu) - (1-\delta)k] dG d\mu.$$

4. *Labor market clearing*:

$$N^h(z, s, \mu) \; = \; \int N(\epsilon, k; z, s, \mu) d\mu \; + \; \int\int_0^{\bar{\xi}} \xi \mathcal{J}\Big(\gamma K(\epsilon, k, \xi; z, s, \mu) - (1-\delta)k\Big) dG d\mu,$$

   where $\mathcal{J}(x) = 0$, if $x = 0$ and 1, otherwise.

5. *Model consistent dynamics*: The evolution of the cross-section that characterizes the economy, $\mu' = \Gamma(z, s, \mu)$, is induced by $K(\epsilon, k, \xi; z, s, \mu)$ and the exogenous processes for z, s as well as ϵ.

Conditions 1, 2, 3 and 4 define an equilibrium given Γ, while step 5 specifies the equilibrium condition for Γ.

3.4 Solution

As is well-known, (4) is not computable, since μ is infinite dimensional. Hence, we follow Krusell and Smith (1997, 1998) and approximate the distribution μ by its first moment over capital, and its evolution, Γ, by a simple log-linear rule. In the same vein, we approximate the equilibrium pricing function by a log-linear rule discrete – aggregate state by discrete aggregate state:

$$\log \bar{k}' = a_k\big(z, s\big) + b_k\big(z, s\big) \log \bar{k}, \tag{9a}$$

$$\log p = a_p\big(z, s\big) + b_p\big(z, s\big) \log \bar{k}, \tag{9b}$$

where $\bar{k}$ denotes aggregate capital holdings. Given (8), we do not have to specify an equilibrium rule for the real wage. As usual with this procedure, we posit this form and check that in equilibrium it yields a good fit to the actual law of motion. We use this simple forecasting rule because it is much less computationally involved than a rule that includes higher moments of the capital distribution. In models without second moment shocks, it has been extensively shown that the first moment suffices. Unfortunately, we show here that the pure R^2 goodness-of-fit metric does not favor the simple rule ($9a$): R^2 below 0.9 are possible, as we shall see in Section 5.2. However, we show for our baseline calibration that the aggregate dynamics of the economy are hardly affected, when higher moments of the capital distribution are included and the R^2 are pushed closer to unity (see Bachmann et al. (2008) for a similar observation). Therefore, we prefer the increase in computational speed and report our results, unless otherwise noted, with the first moment only as a state variable.

Combining these assumptions and substituting $\bar{k}$ for μ into (4) and using (9a)–(9b), we have that (4) becomes a computable dynamic programming problem with policy functions $N = N(\epsilon, k; z, s, \bar{k})$ and $K = K(\epsilon, k, \xi; z, s, \bar{k})$. We solve this problem via value function iteration on V^0.

With these policy functions, we can then simulate a model economy *without* imposing the equilibrium pricing rule (9b), but rather solve for the equilibrium prices along the way. We simulate the model economy for 1,600 time periods and discard the first 100 observations, when computing any statistics. This procedure generates a time series of $\{p_t\}$ and $\{\bar{k}_t\}$ endogenously, with which assumed rules (9a)–(9b) can be updated via a simple OLS regression. The procedure stops when the updated coefficients $a_k\big(z, s\big)$ and $b_k\big(z, s\big)$, as well as $a_p\big(z, s\big)$ and $b_p\big(z, s\big)$ are sufficiently close to the previous ones. We skip the details of this procedure, as this has been outlined elsewhere – see Khan and Thomas (2008) and Bachmann et al. (2008).

4 Calibration

4.1 Baseline

The model period for the baseline model is a year – in congruence with our data frequency. The following parameters have standard values: $\beta = 0.98$ and $\delta = 0.094$, which we compute from NIPA data for the sectoral aggregate that the USTAN sample corresponds to. Given this depreciation rate, we pick $\gamma = 1.014$, in order to match the time-average aggregate investment rate of 0.108. This number is also consistent with German long-run growth rates. The log-felicity function features an elasticity of intertemporal substitution (EIS) of one. The disutility of work parameter, $A = 2$, is chosen to generate an average time spent at work of 1/3, i.e. 8 hours a day, for the baseline calibration.

We set the output elasticities of labor and capital to $\nu = 0.5565$ and $\theta = 0.2075$, respectively, which correspond to the measured average labor and capital shares in manufacturing in the USTAN data base. While our data also include a considerable amount of firms from other sectors, any weighted average or median of these shares would still be close to the manufacturing values, which is why we decided to use them in our baseline calibration. We discuss robustness to this parameter choice in Section 5.2.[17]

Next, we discuss the parameters of the two-state aggregate shock process. Here we simply estimate a bivariate, unrestricted VAR with the linearly detrended natural logarithm of the aggregate Solow residual[18] and the linearly detrended s-process from the USTAN data.[19] The parameters of this VAR are as follows:[20]

$$\varrho_A = \begin{pmatrix} 0.4474 & -3.7808 \\ 0.0574 & 0.7794 \end{pmatrix} \quad \Omega = \begin{pmatrix} 0.0146 & 0.1617 \\ 0.1617 & 0.0023 \end{pmatrix} \tag{10}$$

This process is discretized on a $[5 \times 5]-$grid, using the bivariate analog of Tauchen's procedure.

We measure the steady state standard deviation of idiosyncratic technology innovations, corresponding to the baseline choices for ν and θ, as $\bar{s} = 0.1201$. Since these innovations

[17]If one views the DRTS assumption as a mere stand-in for a CRTS production function with monopolistic competition, than these choices would correspond to an employment elasticity of the underlying production function of 0.7284 and a markup of 31%. Given the regulated product markets in Germany, this is a reasonable value. The implied capital elasticity of the revenue function, $\frac{\theta}{1-\nu}$ is 0.47. Finally, model simulations show that using the capital share as an estimate for the output elasticity of capital under the null hypothesis of the model leads only to a very small overestimation of the latter (roughly 0.011).

[18]We use again $\nu = 0.5565$ and $\theta = 0.2075$ in these calculations.

[19]After firm-level and sectoral fixed effects have been removed, as described in Section 2.1.

[20]With a slight abuse of notation, but for the sake of readability, Ω displays standard deviations on the main diagonal and correlations on the off diagonal.

also exhibit mild excess kurtosis – 4.4480 on average over our time horizon –,[21] and since the adjustment cost parameter $\bar{\xi}$ will be identified by the kurtosis of the firm-level investment rate (next to its skewness), we want to avoid attributing excess kurtosis in the firm-level investment rate to nonlinearities in the adjustment technology, when the driving force itself has kurtosis. Hence, we incorporate the measured excess kurtosis into the discretization process for the idiosyncratic technology state.[22] Finally, we set $\rho_I = 0.95$, in accordance with the high persistence of Solow residual innovations we find in the data. This process is discretized on a 19−state-grid, using the Tauchen's procedure with mixed Gaussian normals. The general equilibrium setup precludes modeling strict unit-root behavior in productivity. Choosing higher levels of autocorrelation below one, is known to be problematic in terms of the quality of the approximation. [23]

Given the aforementioned set of parameters $\Big(\beta, \delta, \gamma, A, \nu, \theta, \rho_A, \Omega, \bar{s}, \rho_I\Big)$[24], we then calibrate the adjustment costs parameter $\bar{\xi}$ to minimize a quadratic form in the logarithmic differences between the time-average firm-level investment rate skewness produced by the model and the data, as well as the time-average firm-level investment rate kurtosis:

$$\min_{\bar{\xi}} \Psi(\bar{\xi}) \equiv 0.5 \cdot \left[\left(\log\left(\frac{1}{26} \sum_t skewness(\frac{I_{i,t}}{0.5 * (K_{i,t} + K_{i,t+1})})(\bar{\xi}) - 1.6645 \right) \right)^2 + \left(\log\left(\frac{1}{26} \sum_t kurtosis(\frac{I_{i,t}}{0.5 * (K_{i,t} + K_{i,t+1})})(\bar{\xi}) - 19.1046 \right) \right)^2 \right]. \quad (11)$$

As can be seen from (11), the histogram of firm-level investment rates exhibits both substantial positive skewness – 1.6645 – as well as excess kurtosis – 19.1046. Caballero et al. (1995) document a similar fact for U.S. manufacturing plants. They also argue that non-convex capital adjustment costs are an important ingredient to explain such a strongly non-Gaussian distribution, given a close-to-Gaussian shock process. We therefore, use the deviation from Gaussianity in firm-level investment rates to identify $\bar{\xi}$.

The following Table 7 demonstrates identification of $\bar{\xi}$, as cross-sectional skewness and kurtosis of the firm-level investment rates are both monotonically increasing in $\bar{\xi}$. We pick $\bar{\xi} = 0.25$ as our baseline case.[25]

[21] We find no skewness.

[22] We achieve this by using a mixture of two Gaussian distributions: $N(0, 0.0777)$ and $N(0, 0.1625)$ – the standard deviations are 0.1201 ± 0.0424 – with a weight of 0.4118 on the first distribution.

[23] The results do not change significantly with either an increase in the fineness of the aggregate grid to $[9 \times 9]$, nor with one in the idiosyncratic grid to a 35−state-grid.

[24] Plus the mixing parameter.

[25] We searched over a much finer grid of $\bar{\xi}$ than displayed in the table, in order to find the optimal $\bar{\xi}$.

Table 7: Calibration of Adjustment Costs - $\bar{\xi}$

$\bar{\xi}$	Skewness	Kurtosis	$\Psi(\bar{\xi})$	Adj. costs/ Unit of Output
0.01	0.7840	5.0383	1.0824	1.5%
0.05	1.5155	7.6443	0.6511	4.2%
0.10	1.9329	9.3329	0.5175	6.8%
0.25 (BL)	2.5590	12.1591	0.4411	13.3%
0.5	3.0683	14.7695	0.4692	23.3%
1	3.5927	17.8153	0.5463	43.3%

4.2 Variants of the Driving Processes

The empirical analysis has shown a co-movement between risk and productivity. The dynamics of this co-movement - as in (2) and (10) - imply that one variable conveys information about the future development of the other. In other words, a shock to productivity risk will affect aggregate activity not only through the option value channels highlighted by Bloom (2009) but also as a news shock in the spirit of Beaudry and Portier (2006), Jaimovich and Rebelo (2008), or Schmidt-Grohe and Uribe (2008). Upon observing an increase in risk, households and firms rationally expect a future decrease in productivity ceteris paribus.

To disentangle the effects from the news about productivity contained in movements in risk, we solve a variant of our model where the actual risk that firms face remains constant, but firms still observe s as a variable without further economic content that nonetheless contains information about future productivity according to (2). The parameters that describe the law of motion (2) remain as given in (10) , i.e. as in the model with risk. This can be achieved by re-parameterizing the model for the standard deviation of idiosyncratic productivity shocks σ_ϵ, and making use of the following definition:

$$\sigma_\epsilon = \bar{s} + \alpha\left(s_t - \bar{s}\right).$$

For $\alpha = 0$ uncertainty remains constant and in the following we use the term 'News Model' to describe this parameterization. Using an $\alpha > 1$ we can scale up the risk fluctuations without changing the underlying dynamics of the exogenous state variables, i.e. keeping the news content of s constant.

For the 'News Model' ($\alpha = 0$), the first element of the two-dimensional state describes actual aggregate productivity as before. The second state no longer influences actual risk, but only the conditional expectation of future productivity z_{t+1}. The baseline specification in which risk fluctuates as well is termed ‘Full Model’ ($\alpha = 1$). Moreover, for comparison we consider a model with a standard univariate AR(1) process for the aggregate Solow residual (‘RBC Model’), and a specification where the only driving force of business cycles are shocks to uncertainty (‘Risk Model’).[26]

5 Results

5.1 Baseline Results

The main set of results is displayed in Tables 8, 10, and 11. These tables compare the simulation results for output, consumption, investment and employment between the four model variants (‘Full’, ‘News’, ‘RBC’, and ‘Risk’) and German aggregate data in terms of a standard second moment analysis.

Risk Model Our first results concern the model with only uncertainty fluctuations. The pure Risk Model, in which the only driving force of business cycles is fluctuations in uncertainty yields little fluctuations. Output fluctuations are an order of magnitude smaller compared to the data. Table 9 shows that doubling the risk fluctuations ($\alpha = 2$) leads to an almost linear increase in all volatilities of the pure Risk Model. Consequently, we would need to observe uncertainty fluctuations that are roughly 8 times as large as the fluctuations we observe for risk, in order for fluctuations in uncertainty to produce as much output volatility as in the data. This number would then be close to the volatility of uncertainty suggested by Bloom (2009) for the US. However, it is unlikely that we underestimate the risk fluctuations in such an order of magnitude. Even focusing on different subsamples we never find an increase in the volatility of uncertainty by more than a factor of two relative to the baseline calibration.

In terms of relative volatility the Risk Model performs better than the standard, first moment shock RBC Model for consumption, but it performs worse for employment and investment. Moreover, in terms of contemporaneous correlations, it leads to a counterfactual de-coupling of consumption and output. In terms of persistence, the ‘Risk Model’ improves upon the ‘RBC Model’. Finally, increasing the adjustment cost parameter $\bar{\xi}$ by factor four to $\bar{\xi} = 1$, leads to no different performance of the ‘Risk Model’, it only increases the volatility of investment somewhat.

[26]The estimated autocorrelation of the aggregate Solow residual is 0.5259 and the standard deviation of its innovations 0.0182. The corresponding moments for the univariate risk process are: 0.5685 and 0.0028.

Table 8: Simulation Results I: relative standard deviations

	Full Model	News Model	RBC-Model	Risk Model	German Economy All Years	German Economy Pre 1990
Volatility of Output	4.0534	4.253	3.1778	0.2985	2.3032	2.467
Volatility of aggregate variables relative to output volatility						
Consumption	0.311	0.297	0.380	0.768	0.775	0.866
Investment	4.737	4.793	4.379	6.718	1.899	1.834
Employment	0.750	0.763	0.699	1.367	0.781	0.804
Investment-Rate	0.518	0.524	0.494	0.815	0.371	0.238

All aggregate data (except for the investment rate) is HP-filtered with a smoothing parameter of $\lambda = 100$.

'Full Model' and 'News Model' Relative to the data from the actual German economy, Full, News, and RBC Model yield somewhat too strong cyclical fluctuations, see Table 8. This is a result of two effects. First, we generate the aggregate Solow residual from the private sector data only. Second the coarse 5×5 grid for the aggregate state. This is somewhat mitigated when using a 9×9 grid.[27] We run a robustness check, where we decrease the volatility of the aggregate Solow residual in order to match the volatility of aggregate output to see whether this too volatile first moment shock has an influence on the other statistics (relative volatilities, correlations with output, autocorrelations). Another common shortcoming of all models (including the 'Risk Model') is that the relative volatilities for consumption and aggregate investment show that there is not enough smoothing in the baseline calibration, which is a known issue of the standard RBC model. The introduction of uncertainty shocks does not improve this as the level of non-convexities essentially puts the model well in a parameter range, where the Khan and Thomas neutrality result still holds (see Khan and Thomas, 2008).

Overall, there is little difference between the 'Full Model' and the 'News Model' in terms of volatility of the key macro aggregates, their correlation with output as well as their persistence. There is some difference, though, to the standard RBC model. Compared to the 'RBC Model' there is some increase in aggregate volatility and in persistence. The correlations with output remain unchanged by and large, see Tables 10 and 11.

[27]Since the relative standard deviations and correlations are hardly affected by choosing a larger grid, we opted for the increase in computational speed.

Table 9: Robustness Checks Risk Model

	RBC-Model	Risk Model	Risk-Model	Risk Model	German Economy All Years	German Economy Pre 1990
Variant	baseline	baseline	(1) $\alpha = 2$	(2) $\bar{\xi} = 1$		
Volatility of Output	3.1778	0.2985	0.597	0.295	2.3032	2.467
Volatility of aggregate variables relative to output volatility						
Consumption	0.380	0.768	0.762	0.781	0.775	0.866
Investment	4.379	6.718	6.712	7.003	1.899	1.834
Employment	0.699	1.367	1.363	1.468	0.781	0.804
Investment-Rate	0.494	0.815	0.815	0.851	0.371	0.238

All aggregate data (except for the investment rate) is HP-filtered with a smoothing parameter of $\lambda = 100$.
(1) double uncertainty fluctuations $\alpha = 2$.
(2) quadruple adjustment costs $\bar{\xi} = 1$.

Table 10: Simulation Results II: Correlations with Aggregate Output

	Full Model	News Model	RBC-Model	Risk Model	German Economy All Years	German Economy Pre 1990
Consumption	0.834	0.845	0.861	0.032	0.670	0.788
Investment	0.979	0.981	0.974	0.839	0.834	0.844
Employment	0.974	0.978	0.962	0.784	0.681	0.639
Investment-Rate	0.944	0.944	0.913	0.720	0.452	0.661

All aggregate data (except for the investment rate) is HP-filtered with a smoothing parameter of $\lambda = 100$.

Table 11: Simulation Results III: Autocorrelations

	Full Model	News Model	RBC-Model	Risk Model	German Economy All Years	German Economy Pre 1990
Output	0.416	0.458	0.306	0.551	0.477	0.525
Consumption	0.665	0.657	0.542	0.514	0.677	0.752
Investment	0.343	0.406	0.255	0.291	0.420	0.498
Employment	0.331	0.401	0.252	0.262	0.606	0.604
Investment-Rate	0.406	0.466	0.349	0.436	0.688	0.590

All aggregate data (except for the investment rate) is HP-filtered with a smoothing parameter of $\lambda = 100$.

The fact that there is little difference in 'News Model' and 'Full Model' demonstrates that the option value effect of uncertainty fluctuations that has been emphasized in the literature is small. Most of the difference to the 'RBC Model' is driven by the fact that changes in uncertainty s alter the conditional expectations of productivity. Hence, an increase in risk acts as a bad news shock. Figure 2 illustrates this by plotting a trajectory of aggregate output simulations from both models for identical shock series.

If we increase the actual fluctuations in uncertainty, setting $\alpha = 2$, the difference between Full and News Model becomes somewhat more pronounced, see Table 12. Interestingly, the overall volatility of the model deceases. This is a result of the mechanism highlighted in Bloom et al. (2007): When uncertainty increases firms react less to first moment shocks. Moreover, the contemporaneous correlation of consumption and output drops to 77.5%.[28]

To further understand the mechanisms behind the effect of shocks to uncertainty, we run a series of trivariate SVAR(1)s where we augment aggregate Solow residual and uncertainty by a third variable, which is the actual endogenous variable of interest.[29] We estimate these SVARs both from the actual aggregate and uncertainty data for 1973 - 1998 and for 26-year samples from simulations of our model.

Figure 3 displays the results for output. Three aspects are noteworthy from this figure: 1) Model and empirical impulse response look very much alike, i.e. our model replicates the aggregate dynamics resulting from a shock to uncertainty. 2) 'Full Model' and 'News Model'

[28]Tables are available upon request.

[29]We use a simple Cholesky decomposition, assuming the aggregate Solow residual has an instantaneous effect on uncertainty but not vice versa. This is to identify a pure innovation in contemporaneous uncertainty, when studying uncertainty shocks. Notice that studying SVARs in this context is really studying and comparing an interesting summary statistic between data and models. There is nor claim on causal empirical identification here.

Table 12: Robustness Checks Simulation I: Risk Fluctuation and Adjustment Costs

Alternative Specifications		(1)	Full Model $\alpha = 2$	News Model $\alpha = 2$	(2)	Full Model $\bar{\xi} = 1$	News Model $\bar{\xi} = 1$
	Standard deviation of Output		3.857	4.258		3.859	4.069
	Volatility of aggregate variables relative to output volatility						
	Consumption		0.355	0.296		0.360	0.337
	Investment		4.703	4.798		4.399	4.502
	Employment		0.743	0.764		0.696	0.721
	Investment-Rate		0.513	0.525		0.486	0.498

All aggregate data (except for the investment rate) is HP-filtered with a smoothing parameter of $\lambda = 100$.
(1) double uncertainty fluctuations $\alpha = 2$.
(2) quadruple adjustment costs $\bar{\xi} = 1$.

Figure 2: Comparison of realizations of the 'Full Model' and the'News Model'

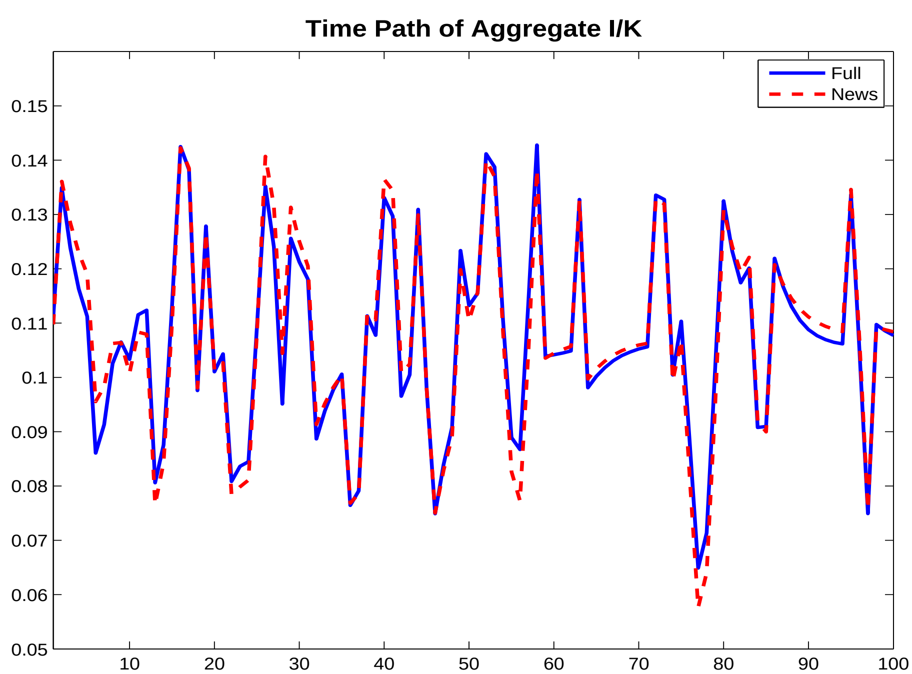

Simulation of 'Full Model' (solid line) and 'News Model' (dotted line) are indistinguishable in their responses. 3) On impact, output goes up after an increase in uncertainty (larger dispersion of productivity innovations in the following period), both in the model as well as empirically.

Figure 3: Impulse Response of Output to a Shock in Uncertainty

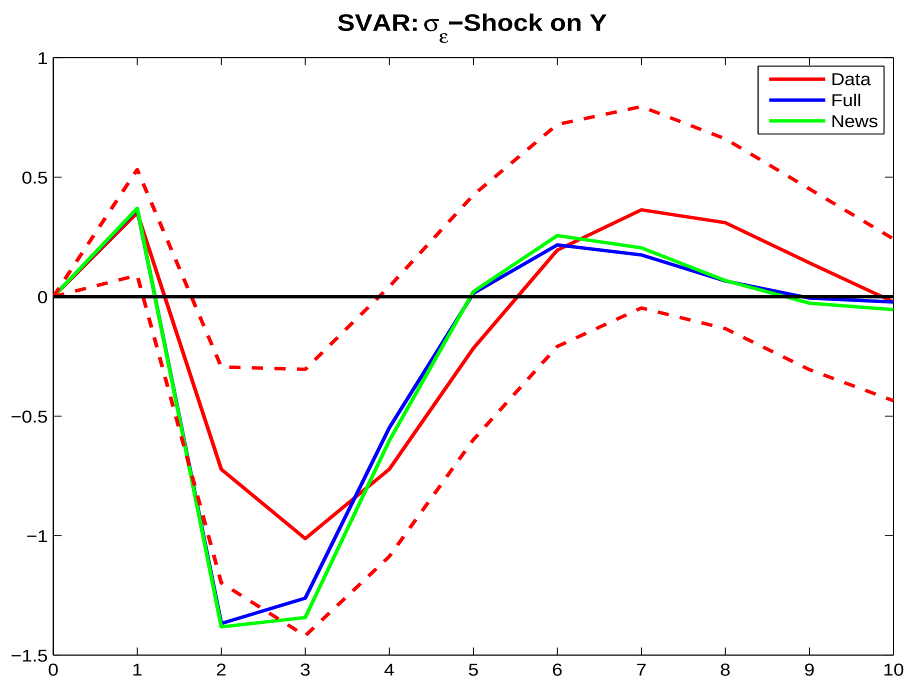

Impulse response functions from SVAR estimations of Value Added, the aggregate Solow residual, and risk. The dotted lines reflect 95% confidence bounds from 10,000 bootstrap replications. Estimates from real data in red, estimates from simulated data in blue and green, respectively.

Figure 4 shows a similar expansion in the aggregate investment rate as a response to an uncertainty shock, which is at least qualitatively consistent with the data, given the wide confidence bands around the point estimate.[30]

In contrast, Figure 5 demonstrates that the expansion in the aggregate investment rate is a general equilibrium effect. The increase in investment is absent in partial equilibrium. There the expectation of lower future productivity drives down the expected return on capital and hence investment, which can be seen from the fact that 'Full Model' and 'News Model' are almost indistinguishable.

Where does the expansionary effect in general equilibrium come from? The SVAR with the real wage provides the answer, as can be seen in Figure 6. Both in the data and in the

[30]An SVAR with aggregate investment instead of the aggregate investment rate shows a positive point estimate with similarly wide confidence bands.

Figure 4: Impulse Response of the Aggregate Investment Rate to a Shock in Uncertainty

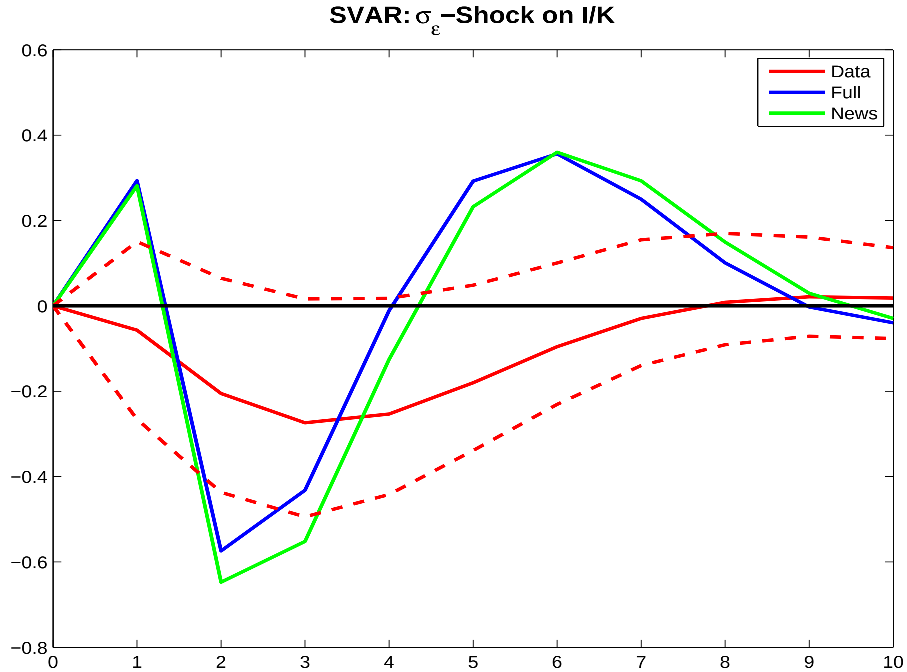

Figure 5: Impulse Response of the Aggregate Investment Rate to a Shock in Uncertainty - GE vs. PE

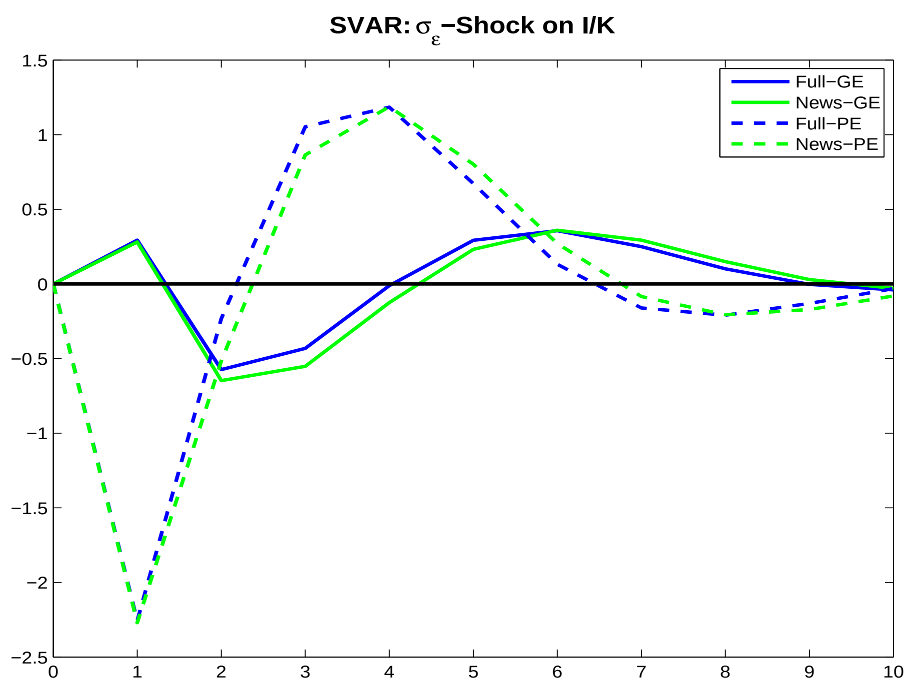

model, we see a decrease in the real wage, which is due to an intertemporal wealth effect that increases labor supply for any given real wage, as a reaction to the bad news heralded by an increase in uncertainty. The labor market, thus, clears at a lower real wage. This is the standard bad news effect on labor supply present in a standard RBC model with King-Rebelo-Plosser preferences (see Beaudry and Portier (2006), Jaimovich and Rebelo (2008) and Sims (2008) for a discussion of this issue). To summarize: quantitatively realistic uncertainty shocks have bad news effects in general equilibrium and it is this news effect that dominates the time-varying real-option value effect. Therefore, studying uncertainty shocks in general equilibrium is paramount to understanding their aggregate implications and the causal mechanism behind these implications.

Figure 6: Impulse Response of the Real Wage to a Shock in Uncertainty - GE vs. PE

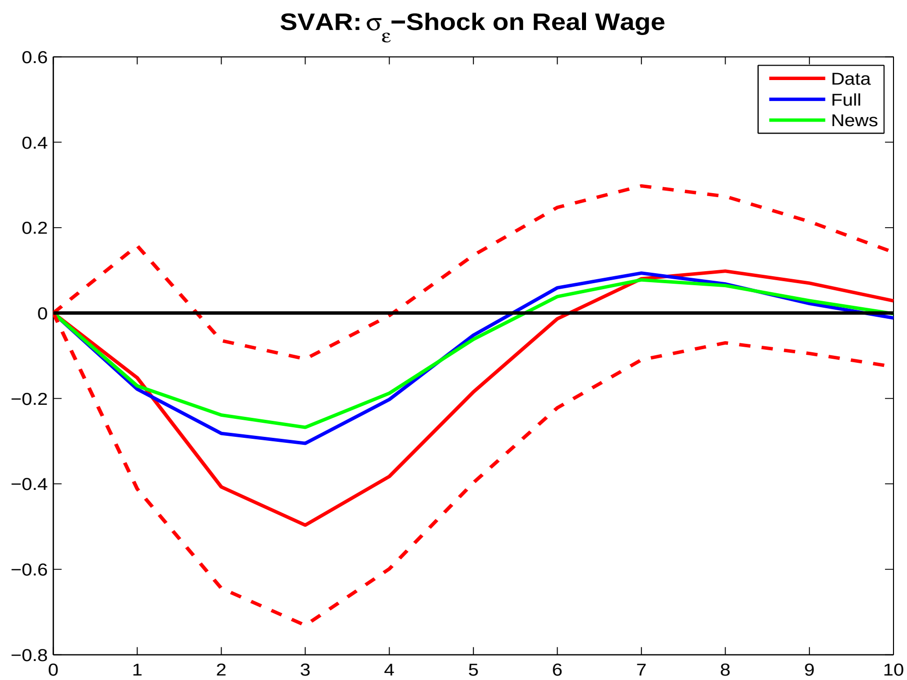

5.2 Further Robustness Checks

TO BE ADDED: Robustness to numerics.

6 Final Remarks

This paper, to the best of our knowledge, is the first to jointly study the cyclical behavior of the second moments of the cross-sections of firm-level innovations to real value added, Solow residuals, and real sales. We show that firm-specific Solow residual innovations (likewise innovations in value added and sales) are more disperse in recessions than in booms. In this sense, uncertainty is significantly and robustly countercyclical in the way Bloom (2009) finds this for US stock market and Compustat data. We also show that the volatility of uncertainty is much lower than has been previously found in this US data.

We then explore the quantitative importance of these risk fluctuations in an otherwise standard heterogeneous-firm dynamic stochastic equilibrium model in the spirit of Khan and Thomas (2008) and Bachmann et al. (2008). We find that empirically realistic fluctuations in uncertainty about productivity growth on their own are not sufficient to generate realistic business cycles. Adding the second moment shocks to a model that has standard first moment shocks, we find that some aspects of the cyclical behavior of the model are altered. In theory this difference can be due to two factors. First, the fact that with higher risk, firms delay investments, the real option value effect highlighted in Bloom (2009); secondly, due to the information risk shocks carry about future fluctuations in productivity, a news effect. For the size of risk fluctuations that we find in German firm-level data, the news effect is largely dominant and explains almost the entire difference between the model with second moment shocks and a standard RBC model with first moment shocks only.

In a companion paper (Bachmann and Bayer, 2009), we provide further indirect evidence that the magnitude of uncertainty shocks over the business cycle cannot be too large. We document there that the firm-level investment rates display significantly and robustly procyclical dispersion. We explain this fact trough a procyclical extensive margin effect caused by lumpy capital adjustment. However, we also demonstrate there that with volatilities of dispersion in the order of magnitude necessary to generate sizeable aggregate effects, this procyclical extensive margin is mitigated, driving down the procyclicality of investment dispersion, that the heterogeneous-firm RBC model here can produce, well below the level found in the data.

Since we base our results on German data, we leave open the possibility that shocks to uncertainty are a major driving force for the business cycle in the US and encourage future research in this cross-country dimension. We also leave open the possibility that the dispersion shocks we measure in the data can be a source of aggregate fluctuations in a significantly altered model environment. We thus view future research into the precise interaction between cross-sectional dynamics and aggregate dynamics as desirable and this paper as merely the beginning of a new research program.

References

[1] Bachmann, R. and C. Bayer (2009). "The Cross-section of Firms over the Business Cycle: New Facts and a DSGE Exploration", mimeo.

[2] Bachmann, R., Caballero, R. and E. Engel (2008). "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model", mimeo.

[3] Bloom, N. (2009). "The Impact of Uncertainty Shocks", *Econometrica*, forthcoming.

[4] Bloom, N., Bond, S. and Van Reenen, J.(2007),"Uncertainty and Investment Dynamics"", *Review of Economic Studies*, 74, 391-415.

[5] Bloom, N., M. Floetotto and N. Jaimovich (2009). "Really Uncertain Business Cycles", incomplete draft.

[6] Caballero, R., E. Engel and J. Haltiwanger (1995). "Plant-Level Adjustment and Aggregate Investment Dynamics", *Brookings Paper on Economic Activity*, 1995, (2), 1–54.

[7] Davis, S., J. Haltiwanger and S. Schuh (1996)."Job Creation and Destruction", Cambridge, MA: MIT Press.

[8] Davis, S., J. Haltiwanger, R. Jarmin and J. Miranda (2006). "Volatility and Dispersion in Business Growth Rates: Publicly Traded and Privately Held Firms", *NBER Macroeconomics Annual*.

[9] Dixit, Avinash K. and Robert S. Pindyck (1994)."Investment under Uncertainty", Princeton University Press, Princeton, New Jersey.

[10] Doepke, J. and S. Weber (2006). "The Within-Distribution Business Cycle Dynamics of German Firms", *Discussion Paper Series 1: Economic Studies*, No 29/2006. Deutsche Bundesbank.

[11] Doepke, J., M. Funke, S. Holly and S. Weber (2005). "The Cross-Sectional Dynamics of German Business Cycles: a Bird's Eye View", *Discussion Paper Series 1: Economic Studies*, No 23/2005. Deutsche Bundesbank.

[12] Doepke, J., M. Funke, S. Holly and S. Weber (2008). "The Cross-Section of Output and Inflation in a Dynamic Stochastic General Equilibrium Model with Sticky Prices", CWPE 0853.

[13] Higson, C., S. Holly and P. Kattuman (2002). "The Cross-Sectional Dynamics of the US Business Cycle: 1950–1999", *Journal of Economic Dynamics and Control*, **26**, 1539–1555.

[14] Higson, C., S. Holly, P. Kattuman and S. Platis (2004): "The Business Cycle, Macroeconomic Shocks and the Cross Section: The Growth of UK Quoted Companies", *Economica*, **71/281**, May 2004, 299–318.

[15] Holly, S. and E. Santoro (2008). "Financial Fragility, Heterogeneous Firms and the Cross-Section of the Business Cycle", CWPE 0846.

[16] von Kalckreuth, U. (2003). "Exploring the Role of Uncertainty for Corporate Investment Decisions in Germany", *Swiss Journal of Economics*, Vol. 139(2), 173–206.

[17] Khan, A. and J. Thomas, (2008). "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics", *Econometrica*, **76**(2), March 2008, 395–436.

[18] Krusell, P. and A. Smith (1997). "Income and Wealth Heterogeneity, Portfolio Choice and Equilibrium Asset Returns", *Macroeconomic Dynamics* **1**, 387–422.

[19] Krusell, P. and A. Smith (1998). "Income and Wealth Heterogeneity in the Macroeconomy", *Journal of Political Economy*, **106** (5), 867-896.

[20] Shiller, R. (1981)."Do Stock Prices Move Too Much to be Justified by Subsequent Changes in Dividends?", American Economic Review, 71, 421–436

[21] Sims, E. (2008)."Expectations Driven Business Cycles: An Empirical Evaluation", mimeo

[22] Stoess, E. (2001). "Deutsche Bundesbank's Corporate Balance Sheet Statistics and Areas of Application", *Schmollers Jahrbuch: Zeitschrift fuer Wirtschafts- und Sozialwissenschaften (Journal of Applied Social Science Studies)*, **121**, 131—137

[23] Tauchen, G. (1986). "Finite State Markov-Chain Approximations To Univariate and Vector Autoregressions", *Economics Letters* **20**, 177–181.

A Appendix A - Data Appendix

B Appendix B - Robustness of Cross-sectional Cyclicality