NBER WORKING PAPER SERIES

UNCERTAINTY BUSINESS CYCLES - REALLY?

Rüdiger Bachmann
Christian Bayer

Working Paper 16862
http://www.nber.org/papers/w16862

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2011

We thank Nick Bloom for his discussion and Dirk Krüger, GiuseppeMoscarini, Gernot Müeller,Matthew Shapiro as well as Eric Sims for their comments. We are grateful to seminar/meeting participants at RWTH Aachen, the ASSA (San Francisco), Bundesbank, CESifoMacro Conference (2010), Cowles Summer Conference (2009), CSEF (Capri), Duke, ESEM (Barcelona), ESSIM 2009, Georgetown, Innsbruck, Konstanz Seminar on Monetary Theory and Policy, Mainz, Michigan-Ann Arbor, Minneapolis FED, NBER Summer Institute (2009), SED (Istanbul), Universitá Ca'Foscari Venezia, VfS (Magdeburg), Wisconsin-Madison and Zürich 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Uncertainty Business Cycles - Really?
Rüdiger Bachmann and Christian Bayer
NBER Working Paper No. 16862
March 2011
JEL No. E2,E22,E3,E32

ABSTRACT

Are fluctuations in firms' profitability risk a major cause of regular business cycles? We study this question within the framework of a heterogeneous-firm dynamic stochastic general equilibrium model with fixed capital adjustment costs. In such a model, surprise increases of risk lead to a wait-and-see policy for investment at the firm level and a decrease in aggregate economic activity. We calibrate the model using German firm-level data with a broader sectoral, size and ownership coverage than comparable U.S. data sets. The use of these data enables us to provide robust lower and upper bound estimates for the size of firm-level risk fluctuations. We find that time-varying firm-level risk on its own is unlikely to be a major quantitative source of regular business cycle fluctuations. When we augment a model with only aggregate productivity shocks by time-varying risk, the risk shocks dampen the high contemporaneous correlations of the productivity-shock-only model, but do not alter the other unconditional business cycle properties.

Rüdiger Bachmann
Department of Economics
University of Michigan
Lorch Hall 365B
Ann Arbor, MI 48109-1220
and NBER
rudib@umich.edu

Christian Bayer
Universität Bonn
Department of Economics
Adenauerallee 24-42
53113 Bonn
Germany
christian.bayer@uni-bonn.de

1 Introduction

Is time-varying firm-level profitability risk a major cause of regular business cycle fluctuations? Shocks to firm risk have the appealing theoretical property that they can generate naturally bust-boom cycles, as shown in a seminal paper by Bloom (2009). After a surprise increase in risk, firms, more uncertain about future profitability, will halt or slow down all activities that cannot be easily reversed, they wait and see. Investment in equipment and structures is an important example. After the heightened uncertainty is resolved, pent-up demand for capital goods leads to an investment boom. In this paper we evaluate this mechanism quantitatively.

We start from a heterogeneous-firm dynamic stochastic general equilibrium model with persistent idiosyncratic productivity shocks and fixed capital adjustment costs. In such an environment, time-varying firm-level risk is naturally modeled as fluctuations in the variance of future firm-level productivity shocks. We develop the numerical tools to solve such a model in general equilibrium. The model features 'wait-and-see' when firm-level risk rises, because investment decisions cannot be reversed easily. The *conditional* effect of increases in firms' risk is thus a bust-boom cycle in aggregate economic activity. While important, conditional moments paint an incomplete picture of the business cycle. We study the *unconditional* business cycle implications of time-varying firm-level risk and compare them to the data and the business cycle properties of a model with aggregate productivity shocks only.

We use the *Deutsche Bundesbank* balance sheet data base of German firms, USTAN, to calibrate the model – in particular the capital adjustment costs and the idiosyncratic risk process. USTAN is a private sector, annual, firm-level data set that allows us to use 26 years of data (1973-1998), with cross-sections that have, on average, over 30,000 firms per year. USTAN has a broader ownership, size and industry coverage than the available comparable U.S. data sets from Compustat and the Annual Survey of Manufacturers. The richness of USTAN lets us take into account measurement error and sample selection issues. It also allows us to formulate lower and upper bound scenarios for the size of firm-level risk fluctuations.

We find that risk shocks alone do not produce recognizable business cycles. They generate only 15 per cent of the volatility of aggregate output, with investment and employment being too volatile relative to output. They lead to negative correlations between aggregate consumption on the one hand and output, investment and employment on the other. We then introduce risk shocks as an independent process alongside standard aggregate productivity shocks. In such an environment, risk shocks help to dampen the notoriously too high contemporaneous correlations in the productivity-shocks-only model. Otherwise the business cycle properties are unaltered. Moreover, the conditional impulse responses to surprise increases in firm-level risk are inconsistent with at least the point estimates of their data counterparts. This can be amended by allowing for correlation between aggregate productivity and firm-level risk and then feeding their *joint* dynamics into the model. In this case, firm-level risk shocks contribute substantially to aggregate fluctuations. Yet, when we isolate the contribution of the 'wait-and-see' effect to these fluctuations, we find that it is again small.

We also show that including time-varying *aggregate* risk has negligible effects since the average level of idiosyncratic risk is estimated to be an order of magnitude larger than aggregate risk. Relative to the large average idiosyncratic risk that firms face, even the sizeable fluctuations of aggregate risk in the data, with a percentage volatility between 30 and 40 per cent, have a negligible impact on the total risk in firms' future profitability and hence also negligible effects on firms' optimal policies.

There is now a growing literature arguing that various measures of firm-level risk both across countries and across data sources, e.g. balance sheet and survey data, are unconditionally countercyclical.[1] While interesting and pervasive, these facts do not, however, directly speak to the question whether risk fluctuations generate regular business cycle fluctuations. Some authors have tackled this question using structural VARs and (linearized) DSGE models. In Christiano et al. (2009), a DSGE estimation exercise, risk shocks have a low frequency and a rather small business cycle impact. This is similar to the SVAR findings in Bachmann, Elstner and Sims (2010), who use business survey data to measure firms' risk. They also argue that observed risk increases might be systematic reactions to first-moment shocks, rather than autonomous drivers of the business cycle.

Our approach, by contrast, is to quantitatively evaluate the 'wait-and-see' effect caused by capital adjustment frictions. We thus build on the literature that highlights *physical* frictions as a propagation mechanism for risk shocks: Bernanke (1983), Dixit and Pindyck (1994), Hassler (1996 and 2001), Bloom (2009), Bloom et al. (2010) and Schaal (2010). Bloom (2009) structurally estimates a rich heterogeneous firm model that features the 'wait-and-see' effect in partial equilibrium. Bloom et al. (2010) show that this conditional effect survives general equilibrium price movements. Schaal (2010) uses a directed search model with uncertainty shocks to understand the labor market in the so-called Great Recession.[2]

The remainder of this paper is organized as follows: Section 2 explains the model. Section 3 describes its calibration and Sections 4 and 5 discuss the results. Appendices provide details on the data as well as the robustness of the calibration and the simulation results.

[1] Bachmann and Bayer (2011), Bachmann, Elstner and Sims (2010), Berger and Vavra (2010), Bloom et al. (2010), Doepke et al. (2005), Doepke and Weber (2006), Gilchrist, Yankow and Zakrajsek (2009), Gourio (2008), Higson et al. (2002, 2004) and Kehrig (2010).

[2] The literature has considered other channels, for example financial frictions in Arellano et al. (2010), Chugh (2009) and Gilchrist, Sim and Zakrajsek (2009); or agency problems in Narita (2010). Fernandez-Villaverde et al. (2009) argue that positive shocks to the interest rate volatility depress economic activity in several Latin American economies. Another literature stresses the importance of rare, but drastic changes in the economic environment, disaster risk: Barro (2007), Barro et al. (2010), Gourio (2010). There is also a literature that studies low frequency movements in both idiosyncratic and aggregate risk, see Davis et al. (2006) as well as Carvalho and Gabaix (2010).

2 The Model

Our model follows closely Khan and Thomas (2008) as well as Bachmann, Caballero and Engel (2010). The main departure from either paper is the introduction of time-varying idiosyncratic and aggregate productivity risk. Specifically, we assume that firms *today* observe the standard deviations of aggregate and idiosyncratic productivity shocks *tomorrow*, respectively, $\sigma(z')$ and $\sigma(\epsilon')$. Notice the timing assumption: if firms learn their productivity levels at the beginning of a period, an increase in today's standard deviation of idiosyncratic shocks does not constitute higher *risk* for firms. It merely leads to a higher cross-sectional dispersion of idiosyncratic productivity today. In contrast, higher standard deviations tomorrow are true risk today. We make this stark timing assumption to give risk shocks the best chance to have the most direct effect possible. None of our main results depend on it.[3]

2.1 Firms

The economy consists of a unit mass of small firms. There is one commodity in the economy that can be consumed or invested. Each firm produces this commodity, employing its predetermined capital stock (k) and labor (n), according to the following Cobb-Douglas decreasing-returns-to-scale production function ($\theta > 0$, $\nu > 0$, $\theta + \nu < 1$):

$$y = z\epsilon k^{\theta} n^{\nu}, \tag{1}$$

where z and ϵ denote aggregate and idiosyncratic revenue productivity, respectively.

The idiosyncratic log productivity process is first-order Markov with autocorrelation ρ_ϵ and time-varying conditional standard deviation, $\sigma(\epsilon')$. Idiosyncratic productivity shocks are otherwise independent from aggregate shocks. The aggregate log productivity process is an AR(1) with autocorrelation ρ_z and time-varying conditional standard deviation, $\sigma(z')$. Idiosyncratic productivity shocks are independent across productive units. The processes for $\sigma(\epsilon') - \bar{\sigma}(\epsilon)$ and $\sigma(z') - \bar{\sigma}(z)$ are also modeled as AR(1) processes, where $\bar{\sigma}(\epsilon)$ denotes the time-average of idiosyncratic risk and $\bar{\sigma}(z)$ the same for aggregate risk.

We denote the trend growth rate of aggregate productivity by $(1-\theta)(\gamma-1)$, so that aggregate output and capital grow at rate $\gamma - 1$ along the balanced growth path. From now on we work with k and y (and later aggregate consumption, C) in efficiency units.

[3] In Table 10 in Appendix B we explore a timing assumption, where firms today know only today's standard deviations, but predict tomorrow's using persistence in the process for the standard deviation of idiosyncratic productivity shocks.

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

Upon investment, i, the firm incurs a fixed cost of $\omega\xi$, where ω is the current real wage. 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, $\big(z, \sigma(z'), \sigma(\epsilon'), \mu\big)$ constitutes the current aggregate state and μ evolves according to the law of motion $\mu' = \Gamma\big(z, \sigma(z'), \sigma(\epsilon'), \mu\big)$, which firms take as given.

To summarize: 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.

Next we describe the dynamic programming problem of a firm. We will take two shortcuts (details can be found in Khan and Thomas, 2008). We state the problem in terms of utils of the representative household (rather than physical units), and denote the marginal utility of consumption by $p = p\big(z, \sigma(z'), \sigma(\epsilon'), \mu\big)$. Also, given the i.i.d. nature of the adjustment costs, continuation values can be expressed without future adjustment costs.

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

$$V^0\big(\epsilon, k; z, \sigma(z'), \sigma(\epsilon'), \mu\big) = \int_{\underline{\xi}}^{\bar{\xi}} V^1\big(\epsilon, k, \xi; z, \sigma(z'), \sigma(\epsilon'), \mu\big) G(d\xi). \tag{2}$$

With this notation the dynamic programming problem becomes:

$$V^1\big(\epsilon, k, \xi; z, \sigma(z'), \sigma(\epsilon'), \mu\big) = \max_n \{\mathrm{CF} + \max(V_{\text{no adj}}, \max_{k'}[-\mathrm{AC} + V_{\text{adj}}])\}, \tag{3}$$

where CF denotes the firm's flow value, $V_{\text{no adj}}$ the firm's continuation value if it chooses inaction and does not adjust, and V_{adj} the continuation value, net of adjustment costs AC, if the firm adjusts its capital stock. That is:

$$\mathrm{CF} = \big[z\epsilon k^{\theta} n^{\nu} - \omega\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big)n\big]\, p\big((z,\sigma(z'),\sigma(\epsilon'),\mu\big), \tag{4a}$$

$$V_{\text{no adj}} = \beta \mathrm{E}\big[V^0\big(\epsilon',(1-\delta)k/\gamma; z',\sigma(z''),\sigma(\epsilon''),\mu'\big)\big], \tag{4b}$$

$$AC = \xi\omega\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big)\, p\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big), \tag{4c}$$

$$V_{\text{adj}} = -i\, p\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big) + \beta \mathrm{E}\big[V^0\big(\epsilon',k'; z',\sigma(z''),\sigma(\epsilon''),\mu'\big)\big], \tag{4d}$$

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 $i = \gamma k' - (1-\delta)k$. The discount factor, β, reflects the time preferences of the representative household.

Taking as given $\omega\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big)$ and $p\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big)$, and the law of motion $\mu' = \Gamma\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big)$, 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: $N = N\big(\epsilon,k;z,\sigma(z'),\sigma(\epsilon'),\mu\big)$ and $K = K\big(\epsilon,k,\xi;z,\sigma(z'),\sigma(\epsilon'),\mu\big)$. Since capital is pre-determined, the optimal employment decision is independent of the current adjustment cost draw.

2.2 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. They own shares in the firms and are paid dividends. We do not need to model the household side in detail (see Khan and Thomas (2008) for that), we just use the first-order conditions that determine the equilibrium wage and the marginal utility of consumption.

Households have a standard felicity function in consumption and labor:[4]

$$U(C,N^h) = \log C - AN^h, \tag{5}$$

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:

$$p\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big) \equiv U_C(C,N^h) = \frac{1}{C\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big)}, \tag{6}$$

[4]We have experimented with a CRRA of 3 without much impact on our results.

and from the intratemporal first-order condition:

$$\omega\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big) = -\frac{U_N(C,N^h)}{p\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big)} = \frac{A}{p\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big)}. \tag{7}$$

2.3 Recursive Equilibrium

A *recursive competitive equilibrium* for this economy 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, $V^1\big(\epsilon, k; z, \sigma(z'), \sigma(\epsilon'), \mu\big)$ solves (3) and the corresponding policy functions are $N\big(\epsilon, k; z, \sigma(z'), \sigma(\epsilon'), \mu\big)$ and $K\big(\epsilon, k, \xi; z, \sigma(z'), \sigma(\epsilon'), \mu\big)$.

2. *Household optimality*: Taking ω and p as given, the household's consumption and labor supply satisfy (6) and (7).

3. *Commodity market clearing*:

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

4. *Labor market clearing*:

$$N^h\big(z,\sigma(z'),\sigma(\epsilon'),\mu\big) = \int N\big(\epsilon, k; z, \sigma(z'), \sigma(\epsilon'), \mu\big) d\mu + \int\int_{\underline{\xi}}^{\bar{\xi}} \xi \mathscr{J}\Big(\gamma K\big(\epsilon, k, \xi; z, \sigma(z'), \sigma(\epsilon'), \mu\big) - (1-\delta)k\Big) dG d\mu,$$

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

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

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

2.4 Solution

It is well-known that (3) is not computable, because μ is infinite dimensional. We follow Krusell and Smith (1997, 1998) and approximate the distribution, μ, by a finite set of its moments, and its evolution, Γ, by a simple log-linear rule. As usual, we include aggregate capital holdings, $\bar{k}$. We also find that it improves the fit of the Krusell-Smith-rules to add the standard deviation of the natural logarithm of idiosyncratic productivity, $std(\log(\epsilon))$. This is of course owing to the now time-varying nature of the distribution of idiosyncratic productivity. In the same vein, we approximate the equilibrium pricing function by a log-linear rule, discrete aggregate state by discrete aggregate state:

$$\log \bar{k}' = a_k\big(z, \sigma(z'), \sigma(\epsilon')\big) + b_k\big(z, \sigma(z'), \sigma(\epsilon')\big) \log \bar{k} + c_k\big(z, \sigma(z'), \sigma(\epsilon')\big) \log std(\log(\epsilon)), \tag{8a}$$

$$\log p = a_p\big(z, \sigma(z'), \sigma(\epsilon')\big) + b_p\big(z, \sigma(z'), \sigma(\epsilon')\big) \log \bar{k} + c_p\big(z, \sigma(z'), \sigma(\epsilon')\big) \log std(\log(\epsilon)). \tag{8b}$$

Given (7), we do not have to specify an equilibrium rule for the real wage. We posit the log-linear forms (8a)–(8b) and check that in equilibrium they yield a good fit to the actual law of motion. The R^2 for capital in our baseline calibration are all above 0.999. For the marginal utility of consumption they exceed 0.995.[5]

Substituting $\bar{k}$ and $std(\log(\epsilon))$ for μ into (3) and using (8a)–(8b), (3) becomes a computable dynamic programming problem with corresponding policy functions $N = N\big(\epsilon, k; z, \sigma(z'), \sigma(\epsilon'), \bar{k}, std(\log(\epsilon))\big)$ and $K = K\big(\epsilon, k, \xi; z, \sigma(z'), \sigma(\epsilon'), \bar{k}, std(\log(\epsilon))\big)$. We solve this problem by value function iteration on V^0. We do so by applying multivariate spline techniques that allow for a continuous choice of capital when the firm adjusts.

With these policy functions, we can then simulate a model economy *without* imposing the equilibrium pricing rule (8b). Rather, we impose market-clearing conditions and solve for the pricing kernel at every point in time of the simulation. We simulate the model economy for a large number of time periods. This generates a time series of $\{p_t\}$ and $\{\bar{k}_t\}$ endogenously, on which the assumed rules (8a)–(8b) can be updated with a simple OLS regression. The procedure stops when the updated coefficients $a_k\big(z, \sigma(z'), \sigma(\epsilon')\big)$ to $c_p\big(z, \sigma(z'), \sigma(\epsilon')\big)$ are sufficiently close to the previous ones.

[5] Of course, $std(\log(\epsilon))$ has an analytically known law of motion, given the AR(1) specification for $\sigma(\epsilon')$. The lowest R^2 for the capital rule without $std(\log(\epsilon))$ is just above 0.94 and for the marginal utility of consumption just above 0.99.

3 Calibration

In this Section we discuss the calibration of those model parameters that remain the same across all specifications and for the baseline model specification presented in Section 4. Our firm-level data source is the USTAN database from *Deutsche Bundesbank*. USTAN is a large annual firm-level balance sheet data base (*Unternehmensbilanzstatistik*). It has broader coverage in terms of firm size, industry and ownership structure than comparable U.S. data sets.[6] From USTAN we compute a time series of the cross-sectional dispersion of firm-level Solow residual growth for 26 years, spanning 1973-1998.

Standard Parameters

The model period is a year. This corresponds to the data frequency in USTAN. Most firm-level data sets that are based on balance sheet data are of that frequency. The following parameters then have standard values: $\beta = 0.98$ and $\delta = 0.094$, which we compute from German national accounting data (*VGR*) for the nonfinancial private business sector. Given this depreciation rate, we pick $\gamma = 1.014$, in order to match the time-average aggregate investment rate in the nonfinancial private business sector: 0.108. $\gamma = 1.014$ is also consistent with German long-run growth rates. The disutility of work parameter, A, is chosen to generate an average time spent at work of 0.33: $A = 2$. We set the output elasticities of labor and capital to $\nu = 0.5565$ and $\theta = 0.2075$, respectively, which correspond to the measured median labor and capital shares in manufacturing in the USTAN data base.[7]

We measure the steady state standard deviation of idiosyncratic productivity shocks as $\bar{\sigma}(\epsilon) = 0.0905$. In the calculation of this number we take measurement error and 2-digit industry-year effects as well as firm-level fixed effects in Solow residual growth rates into account.[8] Since idiosyncratic productivity shocks in the data also exhibit above-Gaussian kurtosis - 4.4480 on

[6] Davis et al. (2006) show that studying only publicly traded firms (Compustat) can lead to wrong conclusions, when cross-sectional dispersion is concerned. Also, just under half of our firms are from manufacturing. We focus instead on the nonfinancial private business sector. Specifically, we include firms that are in one of the following six 1-digit industries: agriculture, mining and energy, manufacturing, construction, trade, transportation and communication. For details on the data set and the calculation of $\sigma(\epsilon)$ in the data, see Appendix A as well as Bachmann and Bayer (2011). An additional advantage of these data is easy access: while on-site, it is otherwise practically unrestricted for researchers, so that results derived from this data base can be easily checked.

[7] 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 $\frac{1}{\theta+\nu} = 1.31$. The implied capital elasticity of the revenue function, $\frac{\theta}{1-\nu}$ is 0.47. Cooper and Haltiwanger (2006), using LRD manufacturing data, estimate this parameter to be 0.592; Henessy and Whited (2005), using Compustat data, find 0.551. We have experimented with both elasticities within conventional ranges, but have not found any of our main results to depend on them. Simulation results are available on request.

[8] See Appendix A for details. Removing fixed effects here serves two purposes. First, it removes differences in idiosyncratic productivity growth that are predictable for the firm. Second, it homogenizes the sample in the sense that we can read these numbers as if the sample composition was fixed. Appendix A also deals with sample selection issues.

average -, and since the fixed adjustment costs parameters will be identified by the kurtosis of the firm-level investment rate (together with its skewness), we want to avoid attributing excess kurtosis in the firm-level investment rate to lumpy investment, when the idiosyncratic driving force itself has excess kurtosis. We incorporate the measured excess kurtosis into the discretization process for the idiosyncratic productivity state by using a mixture of two Gaussian distributions: $N(0, 0.0586)$ and $N(0, 0.1224)$ - the standard deviations are 0.0905 ± 0.0319, with a weight of 0.4118 on the first distribution. Finally, we set $\rho_\epsilon = 0.95$. This process is discretized on a 19−state-grid, using Tauchen's (1986) procedure with mixed Gaussian normals. Heteroskedasticity in the idiosyncratic productivity process is modeled with time-varying transition matrices between idiosyncratic productivity states, where the matrices correspond to different values of $\sigma(\epsilon')$.

In what follows, we describe our baseline choices for the parameters that characterize the aggregate shock processes and adjustment costs. In Section 5 as well as Appendix B we discuss how our model behaves under various alternative choices for these parameters.

Aggregate Shocks

In the baseline case we abstract from time-varying aggregate risk and correlation between aggregate productivity and idiosyncratic risk. Both themes will be taken up in Section 5. Thus, to compute ρ_z and $\bar{\sigma}(z)$, we estimate an AR(1)-process for the linearly detrended cross-sectional average of the natural logarithm of firm-level Solow residuals, again taking industry as well as firm-level fixed effects in Solow residuals into account. The estimation of the AR(1)-process leads to $\rho_z = 0.7530$ and $\bar{\sigma}(z) = 0.0133$.[9] This process is discretized on a 5−state grid, using Tauchen's (1986) procedure.

We also estimate an AR(1)-process for the linearly detrended cross-sectional standard deviation of the first differences of the natural logarithm of firm-level Solow residuals. This leads to $\rho_{\sigma(\epsilon)} = 0.5800$ and $\sigma_{\sigma(\epsilon)} = 0.0037$.[10] Again, this process is discretized on a 5−state grid, using Tauchen's (1986) procedure. This finer discretization compared to a two-state one has the advantage that we do not need to define the high-risk state as a certain multiple of the size of the low-risk state, in order to match the overall volatility of firm-level risk. We do not want to take a stand on how 'catastrophic', i.e. strong but rare, a risk shock is. Instead, we opt for assuming normality of risk shocks, which is supported by the data. Both a Shapiro-Wilk-test and a Jarque-Bera-test do not reject at conventional levels. In fact, Bloom et al. (2010) show that catastrophic risk events such as a doubling of firm-level risk has not occurred in U.S. post war data, and we do not find it in German data, either.[11]

[9] Without taking out the fixed effects in the cross-section these numbers would be, respectively, $\rho_z = 0.7209$ and $\bar{\sigma}(z) = 0.0147$. In Table 11 in Appendix B we report results, where we use an AR(1) based on aggregate Solow residuals calculated from national accounting data. They are basically the same as our baseline results.

[10] Without the fixed effects these numbers would be, respectively, $\rho_{\sigma(\epsilon)} = 0.5710$ and $\sigma_{\sigma(\epsilon)} = 0.0037$.

[11] Figure 5 in Appendix A.2 shows the time path of firm-level risk and average productivity.

To gauge the importance of shocks to firm-level risk for aggregate fluctuations we use its time series coefficient of variation, which for our baseline case equals: $CV_{risk} = 4.72\%$. We will show below that the business cycle relevance of firm-level risk shocks is essentially an increasing function of this statistic.

Pinning down the value of CV_{risk} from firm-level data is invariably laden with assumptions and decisions during the data treatment process. We view our baseline number for CV_{risk} as a middle case. In order to assess how our results depend on CV_{risk}, we consider two additional scenarios: a 'Lower Bound' scenario, where we halve CV_{risk}, and an 'Upper Bound' scenario, where CV_{risk} is quadrupled. The 'Lower Bound' scenario corresponds roughly to a case where we do not eliminate fixed effects nor measurement error and focus only on the smallest 25 percent of firms ($CV_{risk} = 1.97\%$). The idea behind this scenario is to stay as close as possible to the raw data, using minimal assumptions, and to compensate, albeit somewhat crudely, for the unavoidable overrepresentation of large firms even in USTAN. To compute the 'Upper Bound' scenario we take again measurement error and a full set of fixed effects in Solow residual growth rates into account and capital-weight the cross-sectional standard deviation of firm-level Solow residual shocks. This is to give more importance to large firms, which roughly doubles the baseline CV_{risk} to 8.38%. To be conservative, we double this again and use four times the baseline CV_{risk} as the 'Upper Bound' scenario. We show in Section 5.3 that these bounds also cover the available U.S. numbers.

Adjustment Costs

In our baseline specification, we set the lower bound of the adjustment cost distribution, $\underline{\xi}$, to zero. Given the aforementioned set of parameters $\left(\beta, \delta, \gamma, A, \nu, \theta, \bar{\sigma}(\epsilon), \rho_\epsilon, \bar{\sigma}(z), \rho_z, \sigma_{\sigma(\epsilon)}, \rho_{\sigma(\epsilon)}\right)$, we calibrate the remaining adjustment costs parameter, $\bar{\xi}$, to minimize a quadratic form in the normalized 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:[12]

$$\min_{\bar{\xi}} \Psi(\bar{\xi}) \equiv 0.5 \cdot \Bigg[\bigg(\Big(\frac{1}{T}\sum_t skewness(\frac{i_{i,t}}{0.5*(k_{i,t}+k_{i,t+1})})(\bar{\xi}) - 2.1920\Big)/0.6956\bigg)^2 + \\ \bigg(\Big(\frac{1}{T}\sum_t kurtosis(\frac{i_{i,t}}{0.5*(k_{i,t}+k_{i,t+1})})(\bar{\xi}) - 20.0355\Big/5.5064)\bigg)^2\Bigg]. \quad (9)$$

As can be seen from (9), the distribution of firm-level investment rates exhibits both substantial positive skewness – 2.1920 – as well as excess kurtosis – 20.0355. Caballero et al. (1995) doc-

[12]The normalization constants in (9) are, respectively, the time series standard deviation of the cross-sectional investment rate skewness and the time series standard deviation of the cross-sectional investment rate kurtosis in the data.

ument 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 firm-level shock process. With fixed adjustment costs, firms have an incentive to lump their investment activity together over time in order to economize on these adjustment costs. Therefore, typical capital adjustments are large, which creates excess kurtosis. Making use of depreciation, firms can adjust their capital stock downward without paying adjustment costs. This makes negative investments less likely and hence leads to positive skewness in firm-level investment rates. We therefore use the skewness and kurtosis of firm-level investment rates to identify $\bar{\xi}$.

The following Table 1 shows that $\bar{\xi}$ is indeed identified in this calibration strategy, as cross-sectional skewness and kurtosis of the firm-level investment rates are both monotonically increasing in $\bar{\xi}$. The minimum of Ψ is achieved for $\bar{\xi} = 0.25$, which constitutes our baseline case.[13] This implies average costs conditional on adjustment equivalent to roughly 7% of annual firm-level value added, which is well in line with estimates from the U.S. (see Bloom (2009), Table IV, for an overview).

Table 1: CALIBRATION OF ADJUSTMENT COSTS - $\bar{\xi}$

$\bar{\xi}$	Skewness	Kurtosis	$\Psi(\bar{\xi})$	Adj. costs/ Unit of Output
0	-0.0100	3.5696	18.9640	0%
0.01	0.8961	5.1365	10.7922	0.74%
0.1	2.2612	9.6531	3.5651	3.53%
0.25 (BL)	2.8847	12.3966	2.9162	6.97%
0.5	3.3398	14.75196	3.6431	12.09%
0.75	3.5958	16.2382	4.5482	16.97%
1	3.7735	17.3476	5.4069	21.90%
5	4.7616	24.8953	14.4252	110.31%

Notes: 'BL' denotes the baseline calibration. Skewness and kurtosis refer to the time-average of the corresponding cross-sectional moments of firm-level investment rates. The fourth column displays the value of Ψ in (9). The last column shows the average adjustment costs conditional on adjustment as a fraction of the firm's annual output.

[13] Table 12 in Appendix B shows results where we quadruple the adjustment costs, $\bar{\xi} = 1$. This is to give firms more of a wait-and-see motive. Table 13 in Appendix B shows results for the case $\underline{\xi} = \bar{\xi}$. Our baseline specification has stochastic adjustment costs, but their uncertainty does not change over time. This may reduce the time-varying 'wait-and-see' effect. We check this, by making adjustment costs deterministic in this alternative specification.

4 Baseline Results

With this set-up we can now answer our initial question concerning the importance of risk shocks as drivers of the business cycles. We do so in two steps. First, we study a model with only risk shocks ('Risk Model'). Then we add risk shocks as an independent process alongside standard aggregate productivity shocks ('Full Model').

4.1 Risk Model

Partial equilibrium models feature 'wait-and-see' dynamics as their conditional response to a risk shock: a collapse of economic activity on impact, then a strong rebound and overshooting (Bloom, 2009). We confirm in Figure 1 that this characteristic impulse response survives general equilibrium real interest rate and wage adjustments. In fact, the initial investment collapse is somewhat stronger in general equilibrium due to the usual wealth effect. Households perceive the prolonged rebound and overshooting of economic activity in the future, are wealthier and increase consumption of goods and leisure today. Less output is produced, more of it consumed and investment decreases. The rebound is weaker in general equilibrium due to consumption smoothing.

Figure 1: Response of Aggregate Investment to a Shock in Idiosyncratic Risk

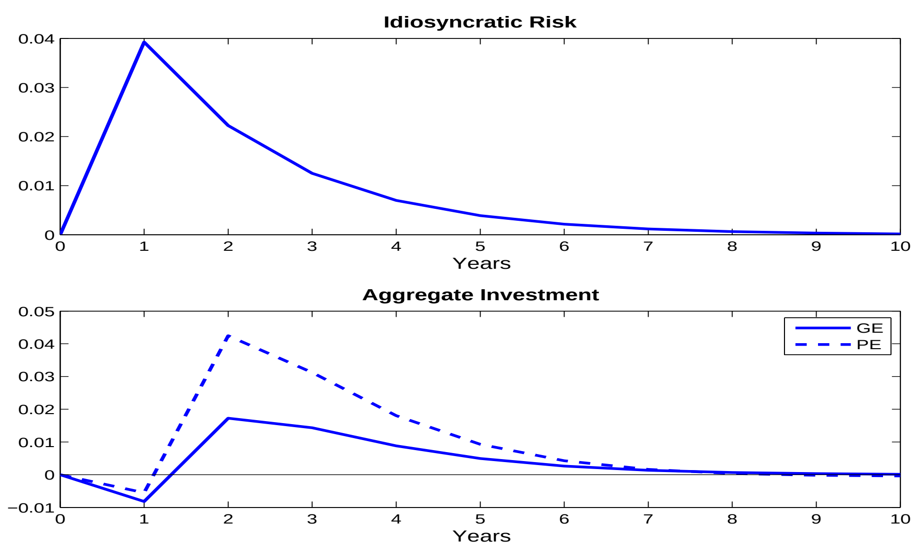

Notes: impulse responses are computed by increasing $\sigma(\epsilon')$ by one standard deviation and letting it return to its steady state value, according to the AR(1) process estimated in Section 3. 'GE' stands for general equilibrium and takes real wage and interest rate movements into account. 'PE' stands for partial equilibrium and fixes the real wage and the interest rate at its steady state level.

To answer our initial question and to understand the importance of time-varying risk for the business cycle, however, conditional responses are not sufficient. Table 2 displays the uncondi-

Table 2: AGGREGATE BUSINESS CYCLE STATISTICS FOR THE 'RISK MODEL'

	Risk Model Baseline	Risk Model Lower Bound	Risk Model Upper Bound	Data
Volatility of Output	0.34%	0.17%	1.20%	2.30%
Volatility of aggregate variables relative to output volatility				
Consumption	0.79	0.79	0.81	0.78
Investment	7.26	7.26	7.27	1.90
Employment	1.48	1.47	1.51	0.78
Persistence				
Output	0.47	0.47	0.47	0.48
Consumption	0.42	0.42	0.40	0.67
Investment	0.18	0.18	0.19	0.42
Employment	0.16	0.16	0.17	0.61
Contemporaneous Correlation with Aggregate Output				
Consumption	-0.12	-0.12	-0.14	0.66
Investment	0.86	0.86	0.85	0.83
Employment	0.82	0.82	0.81	0.68
Contemporaneous Correlation with Aggregate Consumption				
Investment	-0.62	-0.61	-0.63	0.60
Employment	-0.67	-0.67	-0.69	0.36

Notes: 'Risk Model-Baseline' refers to a simulation, where the only aggregate shock is to $\sigma(\epsilon')$, whose time series coefficient of variation is 4.72%. 'Risk Model-Lower Bound' halves this coefficient of variation and 'Risk Model-Upper Bound' quadruples it. 'Data' refers to the nonfinancial private business sector's aggregates. All series, both from data and model simulations, have been logged and HP-filtered with smoothing parameter 100.

tional business cycle properties of models that feature the conditional 'wait-and-see'-response shown in Figure 1. Risk fluctuations in the 'Upper Bound' scenario explain somewhat over half of the output volatility in the data. In the baseline calibration, however, risk shocks produce only 15% of the output volatility in the data. Interestingly, output volatility is essentially a linear function of the size of risk fluctuations. The relative volatilities of investment and employment are too high. Their persistence is too low. Consumption is negatively correlated with the other macroeconomic aggregates in this model.

This constitutes a negative result. The literature has argued that risk shocks might generate cycles through the concentration of economic activity in periods of relatively stable economic environments. However, our quantitative results show that risk fluctuations do not keep this promise when introduced in a relatively standard general equilibrium environment.

We note that going beyond a partial equilibrium analysis and taking into account general equilibrium price movements is important to understanding the relatively mild fluctuations from risk shocks. With fixed real interest rates and real wages the output fluctuations in each scenario roughly double: 0.67%, 0.34% and 2.42% for the 'Baseline'-, 'Lower Bound'- and 'Upper Bound'-scenarios, respectively.[14]

Table 2, in its last column, also shows that the business cycle properties in Germany are roughly the same as in the U.S., so that our results are not due to idiosyncracies in the German business cycle. The only exception is the (relative) volatility of investment, which is indeed lower than in the U.S. However, in a very open economy such as Germany it is unclear what the best data counterpart of model investment is; indeed, the relative volatility of national saving in Germany is 4.62, much closer to the U.S. number for investment.

We conclude with our first result: firm-level risk fluctuations alone, mediated through capital adjustment frictions, are unlikely to be major drivers of the business cycle.

4.2 Full Model

We next ask whether and how exogenous fluctuations in firm-level risk alter the business cycle dynamics of a standard RBC model with fixed capital adjustment costs, when they are added as a second independent aggregate shock process.

Table 3 shows that for an intermediate estimate of the CV_{risk} the business cycle is essentially identical to the one from the RBC model. The ability of risk fluctuations to proportionally rescale output fluctuations has vanished, when first moment fluctuations are present. Only in the extreme case of a $CV_{risk} = 18.88\%$ can risk fluctuations contribute to dampening the notoriously too high comovement of aggregate quantities in the one-shock RBC model, albeit not enough to match the data.

This is our second result: firm-level risk fluctuations added to first moment productivity shocks do not alter significantly RBC business cycle dynamics, with the exception of comovement in the case of highly volatile risk.

[14] Detailed simulation results are available on request.

Table 3: AGGREGATE BUSINESS CYCLE STATISTICS FOR THE 'FULL MODEL'

	Full Model Baseline	Full Model Lower Bound	Full Model Upper Bound	RBC Model	Data
Volatility of Output	2.26%	2.26%	2.39%	2.26%	2.30%
Volatility of aggregate variables relative to output volatility					
Consumption	0.50	0.50	0.51	0.50	0.78
Investment	3.74	3.71	4.14	3.70	1.90
Employment	0.60	0.59	0.71	0.59	0.78
Persistence					
Output	0.41	0.42	0.42	0.42	0.48
Consumption	0.59	0.59	0.56	0.59	0.67
Investment	0.34	0.34	0.31	0.35	0.42
Employment	0.33	0.34	0.28	0.35	0.61
Contemporaneous Correlation with Aggregate Output					
Consumption	0.91	0.92	0.79	0.92	0.66
Investment	0.96	0.97	0.93	0.97	0.83
Employment	0.93	0.94	0.87	0.94	0.68
Contemporaneous Correlation with Aggregate Consumption					
Investment	0.77	0.79	0.51	0.80	0.60
Employment	0.70	0.74	0.38	0.75	0.36

Notes: see notes to Table 2. 'Full Model' refers to a simulation, where there are two orthogonal aggregate shocks, to z and $\sigma(\epsilon')$. The fluctuations of z in 'Full Model-Baseline' have been rescaled to roughly match the volatility of output. All other models use the same rescaling factor. 'RBC Model' refers to a simulation, where the only aggregate shock is to z.

5 Extensions and Robustness

5.1 A Model With Time-Varying Aggregate Risk

In this section, we add time-varying *aggregate* risk to the 'Full Model' with time-varying firm-level risk and productivity shocks. Formally, we allow $\sigma(z')$ to deviate from $\bar{\sigma}(z)$. For computational simplicity, to save on one state variable, we introduce this additional shock as perfectly correlated with the state of firm-level risk. We expect to maximize the impact of time-varying aggregate risk this way. The impact of time-varying risk – wait-and-see – can only be diluted, when both types of risk can move in opposite directions. Thus, in the implementation, whenever $\sigma(\epsilon')$ moves around on its 5-state grid, centered around $\bar{\sigma}(\epsilon) = 0.0905$, we have $\sigma(z')$ move around in the same way on a 5-state grid, centered around $\bar{\sigma}(z) = 0.0133$. We use the grid width of the latter to calibrate the time series coefficient of variation of aggregate risk to roughly 35%.[15] Relative to its average, aggregate risk is thus more than seven times as variable as idiosyncratic risk. One might expect large aggregate effects from these risk fluctuations. The following Table 4 shows that this is not the case. The business cycle statistics of the 'Full Model' with time-varying aggregate and idiosyncratic risk are very similar to those from the 'Full Model' with time-varying idiosyncratic risk only, which are similar to those from the 'RBC Model'.

Table 4: AGGREGATE BUSINESS CYCLE STATISTICS FOR THE 'FULL MODEL' WITH TIME-VARYING AGGREGATE RISK

	Full Model AGGR-RISK	Full Model Baseline	RBC Model	Data
Volatility of Output	2.35%	2.26%	2.26%	2.30%
Volatility of aggregate variables relative to output volatility				
Consumption	0.50	0.50	0.50	0.78
Investment	3.73	3.74	3.70	1.90
Employment	0.59	0.60	0.59	0.78
Persistence				
Output	0.43	0.41	0.42	0.48
Consumption	0.60	0.59	0.60	0.67
Investment	0.36	0.34	0.35	0.42
Employment	0.34	0.33	0.35	0.61
Contemporaneous Correlation with Aggregate Output				
Consumption	0.91	0.91	0.92	0.66
Investment	0.96	0.96	0.97	0.83
Employment	0.93	0.93	0.94	0.68
Contemporaneous Correlation with Aggregate Consumption				
Investment	0.77	0.77	0.80	0.60
Employment	0.70	0.70	0.75	0.36

Notes: see notes to Tables 2 and 3. 'Full Model-AGGR-RISK' refers to a variant of the 'Full Model', where also $\sigma(z')$ varies over time. It is perfectly correlated with $\sigma(\epsilon')$ and its time series coefficient of variation is 34.72%.

[15] We use rolling window standard deviation estimates for the growth rates of aggregate output and employment in Germany and the U.S. The precise number is somewhat sensitive to the data frequency and window size used - higher frequencies and larger window sizes tend to give lower coefficients of variation for aggregate volatility. But most results lie between 30 and 40 per cent.

To understand this result note that the average idiosyncratic risk, $\bar{\sigma}(\epsilon)$, is almost an order of magnitude larger than the average aggregate risk, $\bar{\sigma}(z)$. Since standard deviations are not additive, the combined small aggregate and large idiosyncratic conditional risk, i.e. the standard deviation of the combined productivity shock, is close to the one of idiosyncratic risk. For example, starting from a situation of average aggregate and idiosyncratic risk, the combined conditional risk the firm faces is 0.0915. Jumping from here to a situation with highest aggregate risk (and average idiosyncratic risk) would lead to a combined conditional risk of 0.0940, a 2.7% increase. Moving from the average situation to a situation with highest idiosyncratic risk (and average aggregate risk), leads to an increase in the combined risk to 0.1049 or almost 15%.

We conclude with our third result: aggregate risk fluctuations added to first moment productivity shocks and idiosyncratic risk fluctuations do not alter significantly RBC business cycle dynamics.

5.2 A Model with Correlated Risk and Productivity Shocks

In the previous sections we have investigated the *unconditional* business cycle properties of models with risk shocks. In this Section we study the *conditional* responses of the model and the data to an innovation in firm-level risk.

We estimate three-variable VARs with the cross-sectional average of the natural logarithm of firm-level Solow residuals, idiosyncratic risk and various aggregate activity variables. This ordering is then used in a simple Choleski-"identification", which is, obviously, not meant to have a structural interpretation. It is rather a different, but convenient and instructive way to summarize the data, albeit, given the annual frequency of the data and thus relatively few data points, invariably with some imprecision.

Figure 2 shows this exercise for aggregate output and total hours (using aggregate employment leads to essentially the same picture). Figure 3 does the same for aggregate investment and consumption. The responses in the data of output, hours, investment and consumption to a risk innovation are positive, positive, positive and negative, respectively. The model responses for the 'Full Model - BL', i.e. independent first and second moment shocks, are just the opposite; they feature wait-and-see dynamics. Moreover, the risk responses of the 'Full Model - BL' are not nearly as pronounced as in the data and have overall the wrong shape.

The impulse responses estimated on simulated model data are much closer to those in the data, however, when we allow for correlated risk and productivity processes and feed into the model the joint dynamics we estimate from the data for these two time series ('CORR-BL'). The impulse responses from simulated data now qualitatively match the shape of the impulse responses from actual data for all four macroeconomic quantities.

Figure 2: Impulse Responses to an Innovation in Idiosyncratic Risk - Data and Models

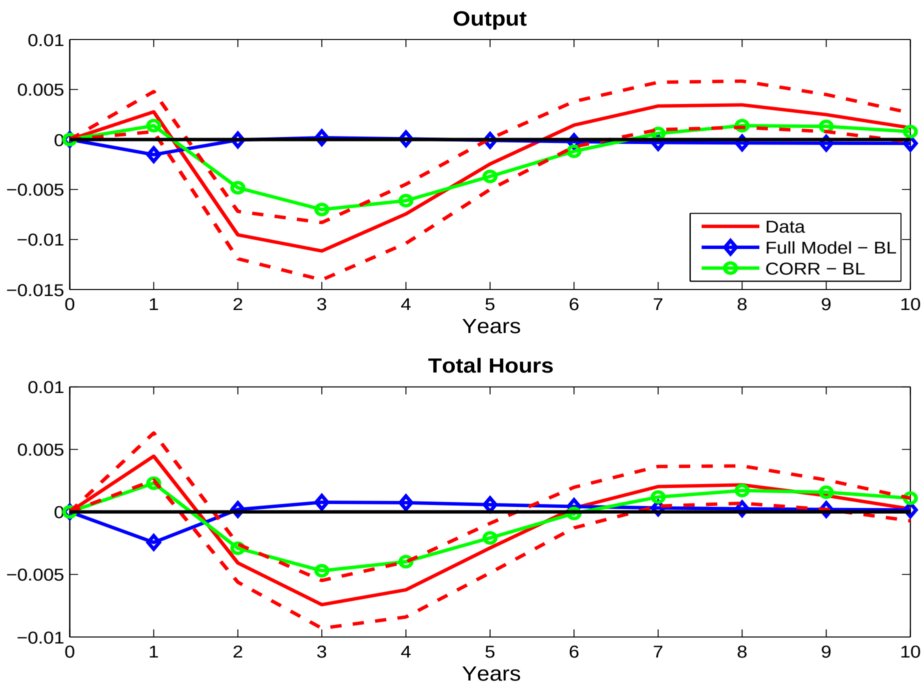

Notes: impulse response functions from SVARs with the linearly detrended cross-sectional average of the natural logarithm of firm-level Solow residuals (ordered first), the linearly detrended idiosyncratic risk (ordered second) and HP(100)-filtered aggregate output/total hours (ordered third). The dotted lines reflect one standard deviation confidence bounds for the estimates on the data from 10,000 bootstrap replications. We employ a bias correction a la Kilian (1998). Estimates from data are in red, estimates from simulated model data in blue ('Full Model-BL') and green ('CORR-BL'), respectively. 'CORR-BL' refers to a simulation, where there are two correlated aggregate shocks, to z and $\sigma(\epsilon')$. 'CORR-BL' is based on a time series coefficient of variation for $\sigma(\epsilon')$ of 4.72%. The joint process is given by: $\begin{pmatrix} 0.8749 & -1.4708 \\ 0.1382 & 0.5101 \end{pmatrix}$, for the VAR-coefficients, where the first row is for the z-equation, and $\begin{pmatrix} 0.0088 & 0.2010 \\ 0.2010 & 0.0029 \end{pmatrix}$ for the matrix of standard deviations and the correlation coefficient. The joint process for z and $\sigma(\epsilon')$ is discretized by a two-dimensional analog of Tauchen's (1986) procedure.

Figure 3: Impulse Responses to an Innovation in Idiosyncratic Risk - Data and Models

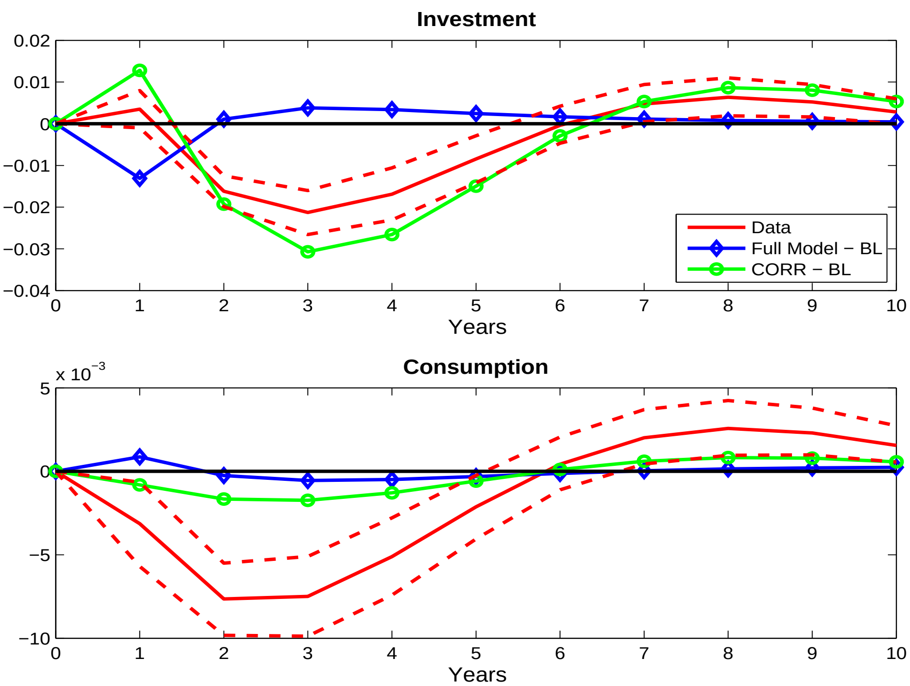

Notes: see notes to Figure 2.

Unlike in the 'Full Model', the introduction of risk shocks in 'CORR-BL' also changes the stochastic properties of aggregate productivity. This effect is very strong, as can be seen in Figure 4, where we compute the impulse response of a risk shock on aggregate investment in a model, where actual firm-level risk is fixed at $\bar{\sigma}(\epsilon)$ and $\sigma(\epsilon)$ is re-interpreted as a latent state variable, which jointly evolves with z just as in 'CORR-BL'. This specification is denoted 'Forecast Model', because "risk" today merely predicts productivity tomorrow, but does not change the *idiosyncratic* stochastic environment of the firms. In other words, "risk" is just a signal of future productivity in this specification. The impulse responses for 'CORR-BL' and 'Forecast Model' are almost identical, which suggests that the conditional effects of risk on aggregate activity are mainly driven by this signalling effect.

This signalling effect – the coefficient of risk today on productivity tomorrow is negative (-1.4708) – has important general equilibrium implications. Figure 4 shows that without market-clearing real interest rates and wages, the investment response to a risk shock would be strongly negative. Since higher risk today forecasts lower productivity tomorrow, a general equilibrium wealth effect makes agents consume less and work more (the real wage declines both in the data and the model), which drives up output and – through a decrease in the real interest rate – investment on impact.

Figure 4: Impulse Response of Aggregate Investment to an Innovation in Idiosyncratic Risk

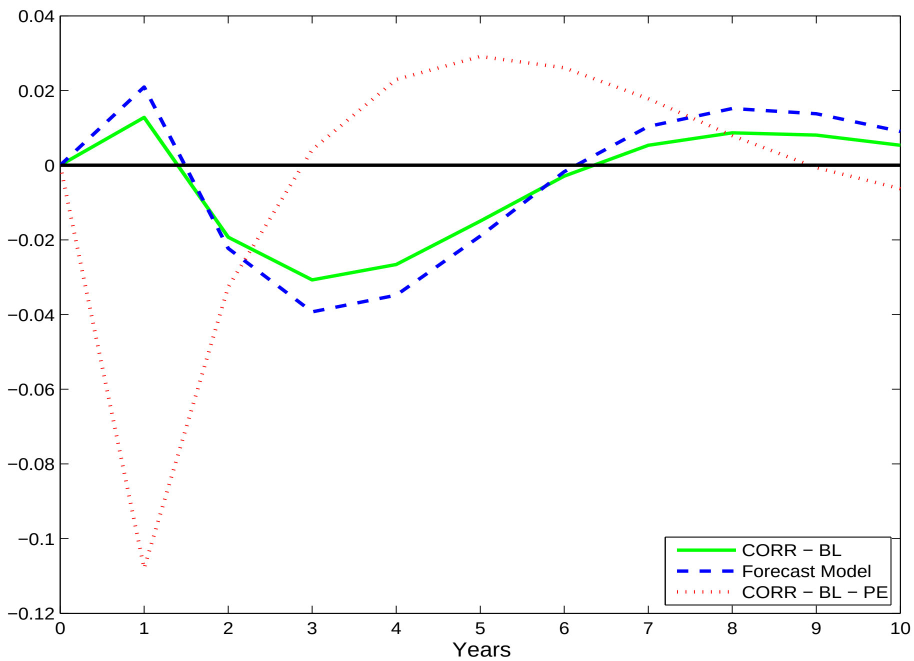

Notes: see notes to Figure 2. 'Forecast Model' uses the same aggregate driving process as 'CORR-BL', but sets the actual value of $\sigma(\epsilon)$ constant at $\bar{\sigma}(\epsilon)$. $\sigma(\epsilon)$ is simply a second random variable that is correlated with z. 'Full Model - BL - PE' is 'Full Model - BL' with a fixed real interest rate and real wage.

Table 5: AGGREGATE BUSINESS CYCLE STATISTICS FOR THE MODEL WITH CORRELATED RISK AND PRODUCTIVITY SHOCKS

	CORR BL	CORR LB	CORR UB	Forecast Model	Naive Model	RBC Model	Data
Volatility of Output	2.34%	2.52%	1.67%	2.71%	2.42%	1.75%	2.30%
Volatility of aggregate variables relative to output volatility							
Consumption	0.42	0.40	0.59	0.39	0.51	0.50	0.78
Investment	3.97	4.12	3.32	4.24	3.67	3.69	1.90
Employment	0.62	0.66	0.53	0.69	0.58	0.58	0.78
Persistence							
Output	0.61	0.62	0.47	0.64	0.64	0.42	0.48
Consumption	0.68	0.70	0.56	0.73	0.72	0.59	0.67
Investment	0.58	0.60	0.42	0.61	0.60	0.35	0.42
Employment	0.58	0.59	0.45	0.60	0.60	0.35	0.61
Contemporaneous Correlation with Aggregate Output							
Consumption	0.94	0.92	0.93	0.90	0.92	0.92	0.66
Investment	0.98	0.98	0.95	0.98	0.96	0.97	0.83
Employment	0.98	0.97	0.82	0.97	0.94	0.94	0.68
Contemporaneous Correlation with Aggregate Consumption							
Investment	0.86	0.83	0.76	0.79	0.78	0.80	0.60
Employment	0.84	0.80	0.55	0.75	0.72	0.75	0.36

Notes: see notes to Tables 2 and 3, as well as Figure 2. 'CORR-BL' is based on CV_{risk} = 4.72%. 'Risk Model-Lower Bound' halves this coefficient of variation and 'Risk Model-Upper Bound' quadruples it. 'Naive Model' is the same as 'Forecast Model', except that agents do not take into account that there is a second random variable that shocks the economy. The fluctuations of z in 'CORR-BL' have been rescaled to roughly match the volatility of output. All the models use the same rescaling factor.

Table 5 summarizes and compares the unconditional business cycle moments for the 'RBC Model' and 'CORR-BL'. It does so in several steps, as the introduction of a second correlated shock changes several features at once relative to the one-shock 'RBC Model'. The intermediate steps help identify these different effects. The 'Naive Model' uses the jointly estimated data generating process for risk and productivity in the model simulations, under two assumptions: first, the agents in the economy – naively – continue to use the univariate process for productivity from the 'RBC Model' when they compute their optimal policies; and, secondly, $\sigma(\epsilon)$ is constant at $\bar{\sigma}(\epsilon)$. The 'Forecast Model' lifts the first assumption, while keeping the second. It corresponds to a model where productivity is driven by two latent random processes instead of one, and agents know that. In addition, 'CORR-BL' lifts the second assumption. 'CORR - LB' and 'CORR - UB' halve and quadruple, respectively, the time series coefficient of variation of firm-level risk. The changes from the 'Forecast Model' to 'CORR-BL' identify the specific effects of time-varying firm risk on aggregate fluctuations.

It is mostly volatilities and relative volatilities that are changed by introducing the second shock. Output fluctuates more, but these output fluctuations are dampened, when actual risk shocks hit the economy, the more so the more volatile risk is. The responsiveness of the economy to productivity shocks decreases in the overall volatility of risk shocks. In terms of relative volatilities, the more volatile actual risk, the less fluctuates aggregate investment and the more aggregate consumption. The correlation structure of aggregate quantities is the same across models, and the increase in persistence from the 'RBC Model' to a model with risk shocks is largely mechanical, as it is manifest already in the 'Naive Model'.

We summarize this section with our fourth result: the conditional impulse responses of aggregate quantities to a risk innovation in a model where risk and productivity shocks are uncorrelated are inconsistent with their data counterparts. A model with correlated risk and productivity shocks matches the data better in terms of conditional impulse responses and leads to a reduction of the ability of productivity shocks to generate aggregate fluctuations.

5.3 Discussion

Are the firm-level risk processes in Germany and the U.S. comparable?[16]

Table 6: COMPARISON GERMANY - U.S.

	CV_{risk}		$Cyclicality_{risk}$	
	STD	IQR	STD	IQR
Baseline Calibration	4.72%		-0.47	
Manufacturing USTAN	6.08%		-0.61	
Manufacturing USTAN Output-based	5.01%	8.00%	-0.54	-0.50
Manufacturing ASM Ouptut-based		9.80%		-0.36

Notes: CV_{risk} is the time series coefficient of variation for the corresponding firm-level risk measure, which can be a cross-sectional standard deviation ('STD') or the interquartile range ('IQR'). $Cyclicality_{risk}$ is the correlation coefficient of the corresponding firm-level risk measure with HP(100)-filtered GDP. 'Manufacturing USTAN' is similar to 'Baseline Calibration' in that it take fixed effects and measurement error into account, as described in Appendix A, but restricts the sample to manufacturing. 'Manufacturing USTAN Output-based' uses the raw firm-level real gross value added data, for better comparison with the available U.S. evidence. 'Manufacturing ASM Ouptut-based' is the 1973-2005 IQR series for firm-level output growth rates in the Annual Survey of Manufacturers, available from $http://www.stanford.edu/\ nbloom/index\_files/Page315.htm$.

[16] We note that German business cycle statistics look rather similar to those in the U.S., which is at least prima facie inconsistent with risk fluctuations being important and different in the two countries.

Table 6 compares our results with the available limited U.S. evidence and shows that both economies have similar firm-level risk processes. The first important fact to note is that all measures of firm-level risk are countercyclical. Second, volatility of the cross-sectional interquartile range of output growth from the Annual Survey of Manufacturers, 9.80%, is close to the corresponding number in the USTAN data, 8.00%. Third, this table also shows in rows one and two that focusing on manufacturing is likely to lead to an overestimation of firm-level risk fluctuations. The USTAN data set allows for a comparison of the extent of firm-level risk fluctuations across industries and our analysis demonstrates that manufacturing is different from services. The combined retail and wholesale trade sector, for example, features a similar volatility and cyclicality of risk as the overall USTAN data set (see Table 8 in Appendix A.1). The combined transportation and communication sector has somewhat higher risk volatility (albeit lower than manufacturing), but firm-level risk is essentially acyclical there. Restricting the analysis to manufacturing data is thus problematic and even more so in the U.S., where this industry has an even smaller share in aggregate production and employment than it has in Germany. Finally, Table 6 shows that the lower and upper bound scenarios we use – half and quadruple the CV_{risk} of the baseline scenario – comfortably cover the available U.S. evidence.

6 Conclusion

This paper argues that shocks to firm-level risk, mediated through physical capital adjustment frictions, are unlikely to be major drivers of the business cycle. We arrive at this conclusion by studying a heterogeneous-firm dynamic stochastic general equilibrium model with persistent idiosyncratic shocks, fixed capital adjustment costs and time-varying firm-level risk. We discipline the model using a rich German firm-level data set. The model features the 'wait-and-see' property for investment after surprise increases in firm-level risk that the recent literature has highlighted. We focus on the *unconditional* business cycle dynamics generated by firm-level risk fluctuations. On its own, time-varying firm-level risk does not produce quantitatively realistic year-to-year business cycle fluctuations, and when juxtaposed to standard aggregate productivity shocks it does little to alter these fluctuations. We leave open the possibility that in different model environments and/or for specific historical episodes risk shocks are important for understanding aggregate developments (see Arellano et. al., 2010, Gilchrist, Sim and Zakrajsek, 2010, as well as Schaal, 2010).

References

[1] Arellano, C., Y. Bai, and P. Kehoe (2010). "Financial Markets and Fluctuations in Uncertainty", Federal Reserve Bank of Minneapolis Research Department Staff Report.

[2] Bachmann, R., S. Elstner and E. Sims (2010). "Uncertainty and Economic Activity: Evidence from Business Survey Data", NBER-WP 16143.

[3] Bachmann, R. and C. Bayer (2011). "The Cross-section of Firms", mimeo University of Michigan.

[4] Bachmann, R., R. Caballero and E. Engel (2010). "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model", mimeo Yale University.

[5] Barro, R. (2007). "On the Welfare Costs of Consumption Uncertainty", mimeo Harvard University.

[6] Barro, R., E. Nakamura, J. Steinsson and J. Ursua (2010). "Crises and Recoveries in an Empirical Model of Consumption Disaster", mimeo Harvard University.

[7] Basu, S. (1996). "Procyclical Productivity: Increasing Returns or Cyclical Utilization?", *The Quarterly Journal of Economics*, **111/3**, 719–751.

[8] Berger, D., and J. Vavra (2010). "The Dynamics of the U.S. Price Distribution", mimeo Yale University.

[9] Bernanke, B. (1983)."Irreversibility, Uncertainty, and Cyclical Investment", *The Quarterly Journal of Economics*, 98/1, 85–106.

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

[11] Bloom, N., M. Floetotto and N. Jaimovich (2010). "Really Uncertain Business Cycles", mimeo Stanford University.

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

[13] Carvalho, V., and X. Gabaix (2010). "The Great Diversification and Its Undoing", mimeo NYU Stern.

[14] Clementi, G., and D. Palazzo (2010). "Entry, Exit, Firm Dynamics, and Aggregate Fluctuations", mimeo NYU Stern.

[15] Christiano, L., R. Motto and M. Rostagno (2009). "Financial Factors in Economic Fluctuations", mimeo Northwestern University.

[16] Chugh, S. (2009). "Firm Risk and Leverage-Based Business Cycles", mimeo University of Maryland.

[17] 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*.

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

[19] 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.

[20] 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.

[21] Fernandez-Villaverde, J., P. Guerron-Quintana, J. Rubio-Ramirez and M. Uribe (2009). "Risk Matters: The Real Effects of Volatility Shocks", mimeo University of Pennsylvania.

[22] Gilchrist, S., J. Sim and E. Zakrajsek (2010). "Uncertainty, Financial Frictions, and Investment Dynamics", mimeo Boston University.

[23] Gilchrist, S., V. Yankov and E. Zakrajsek (2010). "Credit Market Shocks and Economic Fluctuations: Evidence from Corporate Bond and Stock Markets", *Journal of Monetary Economics*, forthcoming.

[24] Gourio, F. (2008). "Estimating Firm-Level Risk", mimeo Boston University.

[25] Gourio, F. (2010). "Disaster Risk and Business Cycles", mimeo Boston University.

[26] Hassler, J. (1996). "Variations in Risk and Fluctuations in Demand - a Theoretical Model", *Journal of Economic Dynamics and Control*, **20**, 1155–1143.

[27] Hassler, J. (2001). "Uncertainty and the Timing of Automobile Purchases", *The Scandinavian Journal of Economics*, **103**, 351–366.

[28] Heckman, J. (1976). "The common structure of statistical models of truncation, sample selection, and limited dependent variables and a simple estimator for such models", *Annals of Economic and Social Measurement*, 5, 475–492.

[29] Henessy, C. and T. Whited, (2005). “Debt Dynamics”, *The Journal of Finance*, **LX/3**, 1129–1165.

[30] 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.

[31] 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.

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

[33] Kehrig, M. (2010). “The Cyclicality of Productivity Dispersion”, mimeo Northwestern University.

[34] 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.

[35] Kilian, L. (1998). “Small Sample Confidence Intervals for Impulse Response Functions”, *Review of Economics and Statistics*, 80, 218–230.

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

[37] Krusell, P. and A. Smith (1998). “Income and Wealth Heterogeneity in the Macroeconomy”, *Journal of Political Economy*, **106** (5), 867–896.

[38] Narita, F. (2010). “Hidden Actions, Risk-Taking and Uncertainty Shocks”, mimeo University of Minnesota.

[39] Schaal, E. (2010). “Uncertainty, Productivity and Unemployment in the Great Recession”, mimeo Princeton University.

[40] 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.

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

A Appendix - Data

A.1 Description of the Sample

Our firm-level data source is USTAN (*Unternehmensbilanzstatistik*) of *Deutsche Bundesbank*, which is a large annual firm-level balance sheet data base. 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. It originated as a by-product of the Bundesbank's rediscounting and lending activities. Bundesbank law required the Bundesbank to assess the creditworthiness of all parties backing a commercial bill put up for discounting. It implemented this regulation by requiring balance sheet data of all parties involved. These balance sheet data were then archived and collected into a database (see Stoess (2001) and von Kalckreuth (2003) for details).

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. Its industry coverage – while still somewhat biased towards manufacturing firms – includes the construction, the service as well as the primary sectors. The following Table 7 displays the industry coverage of our final baseline sample.

Table 7: INDUSTRY COVERAGE

One-digit Industry	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

While there remains a bias towards larger and financially healthier firms, the size coverage is still fairly broad: 31% of all firm-year observations in our final baseline sample have less than 20 employees and 57% have less than 50 employees. In terms of ownership structure, only 2% of firm-year observations are from publicly traded firms, just under 60% from limited liability companies and just under 40% from private firms with fully liable partners. Finally, 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 as sufficiently representative and of high quality. This makes the USTAN data a suitable data source for the study of cross-sectional business cycle dynamics.

From the original USTAN data, we select only firms that report information on payroll, gross value added (before depreciation) and capital stocks. Moreover, we drop observations from East German firms to avoid a break of the series in 1990. We deflate all but the capital and investment data by the implicit deflator for gross value added from the German national accounts.

Capital is deflated with one-digit industry- and capital-good specific investment good price deflators within a perpetual inventory method. Similarly, we recover the amount of labor inputs from wage bills (we calculate an average wage for cells of firms described by industry, year, firm-size, and region and then divide the payroll by this average), as information on the number of employees is only updated infrequently for some companies. Finally, the firm-level Solow residual is calculated from data on real gross value added and factor inputs.

We remove outliers according to the following procedure: we calculate log changes in real gross value added, the Solow residual, real capital and employment, as well as the firm-level investment rate and drop all observations where a change falls outside a three standard deviations interval around the year-specific mean. 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. Details on the implementation as well as the representativeness of the resulting sample can be found in Bachmann and Bayer (2011).

A.2 The Measurement of Firm-Level Risk

A.2.1 Dispersion of Innovations in Measured Firm-Level Solow Residuals

We model fluctuations in idiosyncratic risk as fluctuations in the cross-sectional standard deviation of firm-specific innovations to Solow residuals. Our first step is thus to calculate firm-specific Solow residuals. In accordance with our model, we use the Cobb-Douglas production function from Section 2:

$$y_{i,t} = z_t \epsilon_{i,t} k_{i,t}^{\theta} n_{i,t}^{\nu},$$

where $\epsilon_{i,t}$ is firm-specific and z_t aggregate productivity. We assume that labor input $n_{i,t}$ is immediately productive, whereas capital $k_{i,t}$ is pre-determined and inherited from last period. We estimate the output elasticities of the production factors, ν and θ, as median shares of factor expenditures over gross value added within each industry. We use log-differences in the Solow residual to capture Solow residual innovations, as the persistence of firm-level Solow residuals is high, close to a unit root.

Table 8 displays the cyclical properties of the cross-sectional standard deviation of measured Solow residual innovations for various ways of cutting the sample and treating the data.

Table 8: THE CYCLICAL PROPERTIES OF $std(\Delta \log \epsilon_{i,t})$

Specification	$CV(std(\Delta \log \epsilon_{i,t}))$	$Correl(std(\Delta \log \epsilon_{i,t}), HP(100) - Y)$
Raw data	2.89%	-0.45
Industry and firm fixed effects removed	2.67%	-0.48
Two observations of first differences	2.33%	-0.43
Twenty observations of first differences	5.42%	-0.29
Smallest 25% firms (capital)	1.97%	-0.40
Largest 5% firms (capital)	5.41%	-0.41
Size weighted (capital)	5.92%	-0.68
Publicly traded	5.03%	-0.27
Limited liability	3.44%	-0.45
Privately owned	2.72%	-0.40
Manufacturing	3.83%	-0.57
Trade	2.86%	-0.22
Transportation and Communication	3.20%	0.20
Constant material intensity	4.57%	-0.26

Notes: the first column displays the time series coefficient of variation of the cross-sectional standard deviation of firm-specific Solow residual innovations. The second column displays the time series correlation of this cross-sectional standard deviation with HP(100)-filtered aggregate real gross value added for the nonfinancial private business sector. The first row, 'Raw data', is the baseline *relative* to which the other rows of the table change.

The first row of Table 8, 'Raw data', is the baseline *relative* to which the other rows change.[17] For the second row we remove firm fixed and industry-year effects from these first-difference variables to focus on idiosyncratic fluctuations that do not capture differences in industry-specific responses to aggregate shocks or permanent ex-ante firm heterogeneity. The small differences between the first and second row indicate that the raw data are indeed mostly driven by truly idiosyncratic shocks.

For a firm to be in our baseline sample, we required it to have at least five observations in first differences of payroll, gross value added and capital stocks. When we focus on firms that are in the data base for almost the entire time horizon, i.e. for twenty observations in first differences, the volatility of firm-level risk almost doubles. The reason for this is explained in the next two panels of Table 8, which, respectively, analyze the time series properties of firm risk by firm-size and ownership. It is indeed large and publicly traded firms that display stronger risk fluctuations than smaller and privately owned ones. Since smaller firms seem to face weaker

[17] We also explore different subsamples, for example only the pre-reunification period, industry-specific deflators for firm-level gross value added and various ways to remove outliers – 2.5 and 5 standard deviations or, alternatively, the largest 1% and 5% of observations. None have any significant effect on the results.

fluctuations in risk,[18] our baseline sample in which small firms are underrepresented (and indeed all data sets with an overrepresentation of large and or publicly traded firms) is likely to yield an overestimation of the true cyclical risk fluctuations.

The next to last panel tells a cautionary tale about relying exclusively on manufacturing data when measuring firm-level risk fluctuations. In services, firm-level risk is either less volatile than in manufacturing ('Trade') or not countercyclical ('Transportation and Communication').

Finally, in any Solow residual calculation that is based on a simple Cobb-Douglas production function with only labor and capital, there is the potential problem of attributing optimal changes in utilization, hours per worker or effort to random productivity changes and therefore of overstating (average) firm-level risk. Using ideas from Basu (1996), we calculate the dispersion of the Solow residual innovations for those firms that keep the intensity of material usage constant between two periods, i.e. firms for which the fraction of material usage to sales does not change. One can see that the coefficient of variation goes up but it is still below the estimate for large firms.

In the reminder of this subsection, we check whether the particular sample selection of the USTAN data has any impact on our findings. Clementi and Palazzo (2010) show in a structural model with firm entry and exit that selection on productivity would produce *procyclical* measured risk fluctuations in an actually homoskedastic world. This result provides some indication that our findings are not driven by selection on productivity. The econometric evidence confirms this.

The sample consists of those firms whose bills of exchange were put up for discounting at the Bundesbank. These are likely to be financially healthier and more productive than the average firm. However, this implies a bias for our results only if the level of productivity or financial health predicts productivity changes. To asses whether such a bias is present, we estimate the following simple Heckman (1976)-selection model with a maximum likelihood estimator for each year $t = 1973, \ldots, 1998$. We take all firms present in $t-1$ and observe the estimated level of the Solow residual $\hat{\epsilon}_{i,t-1}$. We assume that the selection to remain in the sample at time t is based on a latent variable $\vartheta_{i,t}$ composed of current productivity $\alpha\epsilon_{i,t}$ and some normally distributed noise term $u_{i,t}$.

$$\vartheta_{it} = \alpha\epsilon_{i,t} + u_{i,t} = \alpha\epsilon_{i,t-1} + \underbrace{u_{i,t} + \alpha\Delta\epsilon_{i,t}}_{=\eta_{i,t}}$$

Moreover, we assume that productivity follows a random walk whose increments $\Delta\epsilon_{i,t} = v_{i,t}$ will be correlated with the composed error term of the selection equation, $\eta_{i,t} = u_{i,t} + \alpha v_{i,t}$. We

[18] Bachmann and Bayer (2011) show that they have higher *average* risk.

assume that $v_{i,t}$ is orthogonal to $\epsilon_{i,t-1}$. Let m_t be the expected growth rate of $\epsilon_{i,t}$, then:

$$\Delta\epsilon_{i,t} = m_t + v_{it}.$$

A (latent) random walk fulfils the exclusion restriction necessary to use the Heckman estimator with $\epsilon_{i,t-1}$ being independent of the innovation $v_{i,t}$ and thus a valid instrument for selection. It is indeed more likely for the firm to remain in the sample with higher levels of $\epsilon_{i,t-1}$. However, this does not influence the estimated variance of $v_{it}, \hat{\sigma}_v^{ML}(t)$. Its correlation with the sample variance $std(\Delta \log \epsilon_{i,t})$ is almost perfect and the cyclical properties of $\hat{\sigma}_v^{ML}(t)$ are almost identical to those of $std(\Delta \log \epsilon_{i,t})$, namely: $CV = 2.63\%$ vs. $CV = 2.67\%$ and a cyclicality of -0.44 versus -0.48.[19] We conclude that, although the sample is clearly no random sample with respect to productivity levels, it is sufficiently random with respect to productivity changes.

A.2.2 Dealing with Measurement Error

Measured Solow residuals will reflect true firm productivity only with error. We take this into account and assume that measured Solow residuals $\hat{\epsilon}_{i,t}$ are composed of the true productivity $\epsilon_{i,t}$ that follows a random walk and a white-noise error term $u_{i,t}$. We assume that $u_{i,t}$ has a time-constant variance σ_u^2, while the innovations to $\epsilon_{i,t}$, $v_{i,t}$, have a time-varying variance $\sigma_v^2(t)$. Under these assumptions we can recover the variance of measurement error, σ_u^2, from

$$\begin{aligned} s_1(t) \quad &: \quad = E\left(\hat{\epsilon}_{i,t} - \hat{\epsilon}_{i,t-1}\right)^2 = E\left(\epsilon_{i,t} - \epsilon_{i,t-1}\right)^2 + E\left(u_{i,t} - u_{i,t-1}\right)^2 \\ &= \quad \sigma_v^2(t) + 2\sigma_u^2 \end{aligned}$$

and

$$\begin{aligned} s_2(t) \quad &: \quad = E\left(\hat{\epsilon}_{i,t+1} - \hat{\epsilon}_{i,t-1}\right)^2 = E\left(\epsilon_{i,t+1} - \epsilon_{i,t-1}\right)^2 + E\left(u_{i,t+1} - u_{i,t-1}\right)^2 \\ &= \quad \sigma_v^2(t+1) + \sigma_v^2(t) + 2\sigma_u^2 \end{aligned}$$

by estimating σ_u^2 from the sample analogues $\hat{s}_1$ and $\hat{s}_2$ of s_1 and s_2, averaging over time:

$$2\hat{\sigma}_u^2 = \frac{1}{T-2}\sum_{t=2}^{T-1}\left\{\left[\hat{s}_1(t) + \hat{s}_1(t+1)\right] - \hat{s}_2(t)\right\}.$$

[19] We used the case with industry and firm fixed effects removed for this analysis.

This estimated measurement error is then subtracted from the variance of measured Solow residuals $\hat{s}_1(t)$ in order to obtain an estimate of the variance of productivity innovations:

$$\hat{\sigma}_{v}^{2}(t) = \hat{s}_1(t) - \frac{1}{T-2}\sum_{s=2}^{T-1}\{[\hat{s}_1(t) + \hat{s}_1(t+1)] - \hat{s}_2(t)\}. \tag{10}$$

We take (10) as our measure of firm-level risk, which is depicted in Figure 5 below, together with average firm-level productivity and aggregate output.

Table 9: THE CYCLICAL PROPERTIES OF FIRM-LEVEL RISK

Specification	$CV(risk)$	$Correl(risk, HP(100)-Y)$
Baseline - FE and ME	4.72%	-0.47
Smallest 25% firms (capital) - LB	1.97%	-0.40
Size weighted (capital) - FE and ME - UB	8.38%	-0.62
Raw - ME	4.10%	-0.44
20 obs. of first differences - FE and ME	7.26%	-0.38
Largest 5% firms (capital) - FE and ME	7.28%	-0.46
Manufacturing - FE and ME	6.08%	-0.61
Publicly traded - FE and ME	7.34%	-0.29
20 obs. of first differences manufacturing - FE and ME	7.52%	-0.50

Notes: the first column displays the time series coefficient of variation of the cross-sectional standard deviation of firm-specific Solow residual growth purged of measurement error ('ME') and firm-specific as well as industry-year fixed effects ('FE'): firm-level risk. The second column displays the time series correlation of firm-level risk with HP(100)-filtered GDP.

Table 9 shows the cyclical properties of firm-level risk, i.e. the innovations of the Solow residual purged of measurement error. The first row represents our baseline calibration. We base our lower bound calibration scenario loosely on the second row, which displays the cyclical properties of firm-level risk for small firms, which are underrepresented in USTAN, and using the raw data, which are based on a minimum amount of assumptions. We base our upper bound calibration scenario loosely on the third row, which delivers the strongest risk fluctuations. To be conservative we roughly double this value when computing the upper bound models. Interestingly, combining features that increase risk fluctuations, such as 'being almost always in the sample' and 'being in manufacturing' (see Table 8 in the previous subsection), does not substantially increase the volatility of risk over and above what one of these features alone does (see the last row of Table 9). Any other combination would not have left sufficient data to yield reliable results.

Figure 5 depicts the time series of firm-level productivity risk, average productivity and cyclical aggregate output.

Figure 5: Time Series of Firm-Level Risk, Average Productivity and Cyclical Aggregate Output

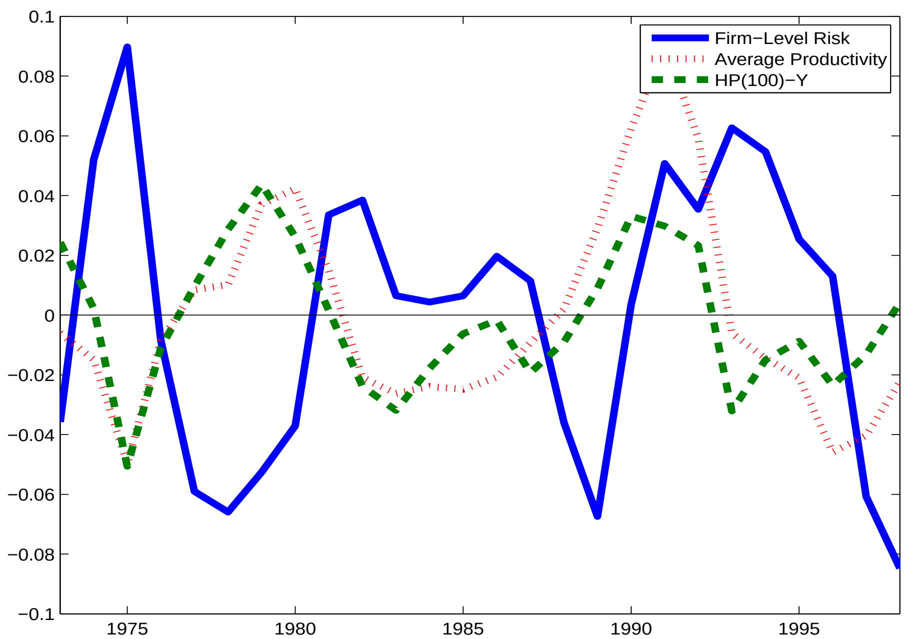

Notes: 'Firm-Level Risk', the solid line, is the time series of our baseline measure of firm-level risk, linearly detrended and normalized by time-average risk. 'Average Productivity', the dotted line, is the time series of firm-level average productivity, linearly detrended. 'HP(100)-Y' is HP(100)-filtered aggregate real gross value added for the nonfinancial private business sector.

B Appendix - Robustness

The following tables display results from simulations of our model for (i) a different timing assumption for when firms learn the dispersion of idiosyncratic shocks, (ii) aggregate productivity being calibrated from aggregate Solow residuals, (iii) a calibration with higher fixed costs of capital adjustment and (iv) a calibration with deterministic fixed costs of adjustment.

Table 10: AGGREGATE BUSINESS CYCLE STATISTICS FOR THE 'FULL MODEL' - DIFFERENT TIMING FOR RISK

	Full Model Diff. Timing	Full Model Baseline	RBC Model	Data
Volatility of Output	2.25%	2.26%	2.26%	2.30%
Volatility of aggregate variables relative to output volatility				
Consumption	0.50	0.50	0.50	0.78
Investment	3.68	3.74	3.70	1.90
Employment	0.58	0.60	0.59	0.78
Persistence				
Output	0.42	0.41	0.42	0.48
Consumption	0.59	0.59	0.59	0.67
Investment	0.34	0.34	0.35	0.42
Employment	0.34	0.33	0.35	0.61
Contemporaneous Correlation with Aggregate Output				
Consumption	0.92	0.91	0.92	0.66
Investment	0.97	0.96	0.97	0.83
Employment	0.94	0.93	0.94	0.68
Contemporaneous Correlation with Aggregate Consumption				
Investment	0.80	0.77	0.80	0.60
Employment	0.75	0.70	0.75	0.36

Notes: 'Full Model-Baseline' refers to a simulation, where there are two orthogonal aggregate shocks, to z and $\sigma(\epsilon')$. This is the baseline model discussed in Section 4.2. 'Full Model-Baseline' and 'Full Model-Diff. Timing' differ in that the latter allows agents to know only $\sigma(\epsilon)$ (and not $\sigma(\epsilon')$). 'RBC Model' refers to a simulation, where the only aggregate shock is to z; obviously, here there is no timing issue. 'Data' refers to the nonfinancial private business sector's aggregates.

Table 11: AGGREGATE BUSINESS CYCLE STATISTICS FOR THE 'FULL MODEL' - AGGREGATE SOLOW RESIDUALS

	Full Model Aggr. SR	RBC Model Aggr. SR	Full Model Baseline	RBC Model Baseline	Data
Volatility of Output	2.34%	2.34%	2.26%	2.26%	2.30%
Volatility of aggregate variables relative to output volatility					
Consumption	0.41	0.41	0.50	0.50	0.78
Investment	4.19	4.18	3.74	3.70	1.90
Employment	0.69	0.68	0.60	0.59	0.78
Persistence					
Output	0.30	0.31	0.41	0.42	0.48
Consumption	0.55	0.55	0.59	0.59	0.67
Investment	0.23	0.25	0.34	0.35	0.42
Employment	0.22	0.24	0.33	0.35	0.61
Contemporaneous Correlation with Aggregate Output					
Consumption	0.87	0.88	0.91	0.92	0.66
Investment	0.97	0.97	0.96	0.97	0.83
Employment	0.95	0.96	0.93	0.94	0.68
Contemporaneous Correlation with Aggregate Consumption					
Investment	0.73	0.75	0.77	0.80	0.60
Employment	0.68	0.71	0.70	0.75	0.36

Notes: see notes to Table 10. 'Full Model-Aggr. SR' refers to a simulation, where there are two orthogonal aggregate shocks, to z and $\sigma(\epsilon')$, but in contrast to the baseline case we use Solow residuals calculated from German industry national accounting data that correspond to the nonfinancial private business sector to calibrate the exogenous aggregate process. We use $\nu = 0.5565$ and $\theta = 0.2075$. 'RBC Model-Aggr. SR' is the analog of 'RBC Model-Baseline', again with Solow residuals from national account data. The fluctuations of z in 'Full Model-Aggr. SR' have been rescaled to roughly match the volatility of output. 'RBC Model-Aggr. SR' uses the same rescaling factor.

Table 12: AGGREGATE BUSINESS CYCLE STATISTICS FOR THE 'FULL MODEL' - HIGHER ADJUSTMENT COSTS

	Full Model High AC	RBC Model High AC	Full Model Baseline	RBC Model Baseline	Data
Volatility of Output	2.16%	2.16%	2.26%	2.26%	2.30%
Volatility of aggregate variables relative to output volatility					
Consumption	0.54	0.54	0.50	0.50	0.78
Investment	3.48	3.41	3.74	3.70	1.90
Employment	0.56	0.54	0.60	0.59	0.78
Persistence					
Output	0.42	0.42	0.41	0.42	0.48
Consumption	0.54	0.55	0.59	0.59	0.67
Investment	0.34	0.36	0.34	0.35	0.42
Employment	0.33	0.36	0.33	0.35	0.61
Contemporaneous Correlation with Aggregate Output					
Consumption	0.94	0.95	0.91	0.92	0.66
Investment	0.96	0.97	0.96	0.97	0.83
Employment	0.93	0.95	0.93	0.94	0.68
Contemporaneous Correlation with Aggregate Consumption					
Investment	0.81	0.86	0.77	0.80	0.60
Employment	0.75	0.81	0.70	0.75	0.36

Notes: see notes to Table 10. 'Full Model-High AC' refers to a simulation, which is similar to 'Full Model-Baseline', but the upper adjustment cost factor, $\bar{\xi}$, is quadrupled. 'RBC Model-High AC' is the analog of 'RBC Model-Baseline', again with quadrupled adjustment costs.

Table 13: AGGREGATE BUSINESS CYCLE STATISTICS FOR THE 'FULL MODEL' - DETERMINISTIC ADJUSTMENT COSTS

	Full Model Det. AC	RBC Model Det. AC	Full Model Baseline	RBC Model Baseline	Data
Volatility of Output	2.30%	2.31%	2.26%	2.26%	2.30%
Volatility of aggregate variables relative to output volatility					
Consumption	0.49	0.49	0.50	0.50	0.78
Investment	3.84	3.81	3.74	3.70	1.90
Employment	0.63	0.61	0.60	0.59	0.78
Persistence					
Output	0.41	0.42	0.41	0.42	0.48
Consumption	0.61	0.62	0.59	0.59	0.67
Investment	0.33	0.34	0.34	0.35	0.42
Employment	0.31	0.34	0.33	0.35	0.61
Contemporaneous Correlation with Aggregate Output					
Consumption	0.89	0.90	0.91	0.92	0.66
Investment	0.96	0.96	0.96	0.97	0.83
Employment	0.93	0.94	0.93	0.94	0.68
Contemporaneous Correlation with Aggregate Consumption					
Investment	0.73	0.76	0.77	0.80	0.60
Employment	0.66	0.70	0.70	0.75	0.36

Notes: see notes to Table 10. 'Full Model-Det. AC' refers to a simulation, which is similar to 'Full Model-Baseline', but adjustment costs are deterministic. 'RBC Model-Det. AC' is the analog of 'RBC Model-Baseline', again with deterministic adjustment costs. Adjustment costs are again calibrated to match a weighted quadratic form in the skewness and kurtosis of the average investment rate distribution (see Section 3). The fluctuations of z in 'Full Model-Det. AC' have been rescaled to roughly match the volatility of output. 'RBC Model-Det. AC' uses the same rescaling factor.