‘Wait-and-See’ Business Cycles?[☆]

Rüdiger Bachmann[a], Christian Bayer[b]

[a] RWTH Aachen University, NBER, CESifo, and ifo
[b] Bonn University

Received Date; Received in Revised Form Date; Accepted Date

Abstract

Are shocks to firms’ profitability risk, propagated by physical capital adjustment costs, a major source of business cycle fluctuations? We study this question using a heterogeneous-firm dynamic stochastic general equilibrium model, where firms face fixed capital adjustment costs. Surprise increases in idiosyncratic risk lead firms to adopt a ‘wait-and-see’ policy for their investment. The model is calibrated using a German firm-level data set with broader coverage than comparable U.S. data sets. The main result is that time-varying firm-level risk through ‘wait-and-see’ dynamics is unlikely a major source of business cycle fluctuations.

JEL Codes: E20, E22, E30, E32.

Keywords: Ss model, RBC model, lumpy investment, aggregate shocks, idiosyncratic shocks, heterogeneous firms, risk shocks.

[☆]We thank Nick Bloom, Bob King, Dirk Krüger, Francois Gourio, Giuseppe Moscarini, Gernot Müller, Ricardo Reis, Matthew Shapiro, Eric Sims, and an anonymous referee for their helpful comments. We are grateful to seminar/meeting participants at RWTH Aachen University, the ASSA (2009 meeting in San Francisco), Bundesbank, CESifo Macro 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), Universit 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. Any remaining errors are our own. This paper formerly circulated under: CESIFO-WP 2844 “Firm-Specific Productivity Risk over the Business Cycle: Facts and Aggregate Implications” and NBER-WP 16862 “Uncertainty Business Cycles - Really?”

[*]Corresponding Author: Rüdiger Bachmann, Templergraben 64, RWTH Aachen University, 52062 Aachen, Germany, Phone: +49 (0)241- 80-96203, Email: ruediger.bachmann@rwth-aachen.de.

1. Introduction

Is time-varying firm-level profitability risk, propagated by physical capital adjustment frictions, a major cause of business cycle fluctuations? Shocks to firm risk have the appealing theoretical property that they can naturally generate bust-boom cycles, as shown in a seminal paper by Bloom (2009). These bust-boom cycles feature sharp recessions and longer-lasting recoveries, an asymmetry that is typical of most observed business cycles. After a surprise increase in risk, firms, more uncertain about future profitability, will halt or slow down activities that cannot be easily reversed - they 'wait and see'.[1] 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.

The propagation of firm-level risk shocks in the 'wait-and-see' story has the additional appeal that it is based on well-established capital adjustment frictions at the micro level, such as non-convex adjustment costs and irreversibilities (see Doms and Dunne, 1998, as well as Cooper and Haltiwanger, 2006). There is also microeconometric evidence that links increased firm-level risk to investment declines mediated through physical adjustment frictions: Guiso and Parigi (1999), Fuss and Vermeulen (2004), Bloom et al. (2007), as well as Bontempi et al. (2010). Finally, says Bloom (2009): "More generally, the framework in this paper also provides one response to the 'where are the negative productivity shocks?' critique of real business cycle theories. ...Recessions could simply be periods of high uncertainty without negative productivity shocks." This paper provides a quantitative evaluation of this mechanism.

We use the *Deutsche Bundesbank's* firm balance sheet data base, USTAN, to measure firms' profitability risk and its cyclical fluctuations. USTAN is a private sector, annual 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

[1] The basic idea goes back to Bernanke (1983), Dixit and Pindyck (1994), Hassler (1996 and 2001).

coverage than the available comparable U.S. data sets from Compustat and the Annual Survey of Manufacturers (ASM).

Firm-level Solow residual growth and firm-level real output growth in USTAN display countercyclical dispersion. The richness of USTAN also allows us to formulate lower and upper bound scenarios for the size of firm-level risk fluctuations. The empirical analysis suggests that existing estimates of the size of firm-level risk fluctuations, based on U.S. data sets for large publicly traded or manufacturing firms, are likely biased upward.

A growing literature argues that various measures of firm-level risk are pervasively countercyclical in an unconditional sense. Bloom et al. (2010) document this for the sales growth of large publicly traded firms (Compustat) and manufacturing plants (ASM). Gilchrist et al. (2010) as well as Gourio (2008) use Compustat data, and Chugh (2011) uses ASM data to establish related facts for various productivity measures. Kehrig (2010) shows that the level of productivity in the ASM is countercyclically disperse. Berger and Vavra (2011), using the underlying micro data of the CPI, show that the dispersion of price changes is countercyclical. Davis et al. (1996) find that the dispersion of employment growth rates across manufacturing establishments was significantly larger in 1982 (recession) than in 1978 (expansion). Finally, Bachmann, Elstner and Sims (2011) show that both disagreement and forecast error dispersion indices from business surveys are countercyclically disperse.[2] While interesting and suggestive, these facts do not, however, directly speak to the question whether risk fluctuations generate business cycle fluctuations.

Our approach is to quantitatively evaluate the business cycle implications of the 'wait-and-see' effect caused by capital adjustment frictions. The USTAN data are used to calibrate 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. The necessary numerical tools are developed to solve such a

[2] Doepke et al. (2005), Doepke and Weber (2006), Higson et al. (2002, 2004) are additional examples.

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. We then study the *unconditional* business cycle implications of time-varying firm-level risk and compare them to the data as well as the business cycle properties of a model with aggregate productivity shocks only. While the conditional dynamics establish 'wait-and-see' as an interesting mechanism, the conditional dynamics are not sufficient to show that time-varying firm-level profitability risk, mediated by capital adjustment costs, is a major cause of business cycle fluctuations. The results from this comparison of unconditional business cycle moments are the reason why this paper reaches a different conclusion from Bloom (2009) and Bloom et al. (2010) regarding the promise of 'wait-and-see' uncertainty-driven business cycles.

Firm-level risk shocks in the model generate roughly 2 percent of the observed time series variance of output (equivalent to 15 percent of the standard deviation) when introduced alone or alongside independent aggregate productivity shocks. In other words, firm-level risk shocks propagated through 'wait-and-see' dynamics leave the unconditional business cycle statistics of the model basically unaltered. This holds true also without equilibrium real wage and real interest rate movements, i.e. in partial equilibrium.

While the baseline model, focusing exclusively on 'wait-and-see' as a transmission channel, suggests that risk shocks have negligible effects for the business cycle, a forecast error variance decomposition on the actual data based on simple Choleski-identified vector autoregressions reveals that risk shocks, at longer horizons, account for roughly one third of output, investment and hours fluctuations, because risk and TFP are correlated in the actual data and not orthogonal as the baseline model assumes. In fact, data and model nicely align once one feeds the observed correlation between aggregate productivity and firm-level risk and their joint dynamics into the model. Then, firm-level risk shocks contribute as substantially to aggregate fluctuations in the model as in the data. Moreover, the conditional impulse responses to surprise increases in firm-level risk in the model become consistent with their data counterparts. Nonetheless, the isolated contribution of the 'wait-and-see' mechanism remains small. Rather, risk shocks contribute to aggregate volatility as they help forecast future aggregate productivity over and above what firms can predict from today's state of aggregate productivity.

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

To be clear about what these findings mean for 'wait-and-see' uncertainty-driven business cycles in particular, and for uncertainty-driven business cycles more generally: these findings leave open the possibility that particular historical recessions were driven mainly by 'wait-and-see' bust-boom cycles. Also, since variance decompositions of the data show that firm-level risk shocks explain roughly one third of output, investment and employment fluctuations, at least at longer horizons,[3] the results in this paper open up room for other (propagation) mechanisms that are currently discussed in the literature. One angle is to explore more structurally the reason for the observed dynamic correlations between firm-level risk and aggregate productivity found in the data. Bachmann and Moscarini (2011), using as their starting point the SVAR results of Bachmann, Elstner and Sims (2011), which suggest that increases in uncertainty might be caused by aggregate first-moment shocks, provide a learning model where firms that are subject to negative first-moment shocks or news thereof react with increased risk-taking and experimentation with their prices. In this view, observed risk fluctuations are an epiphenomenon of aggregate first-moment shocks, not autonomous drivers of the business cycle.

[3]In Christiano et al. (2010), a DSGE estimation exercise, risk shocks have also a strong low frequency impact. This is similar to the SVAR findings in Bachmann, Elstner and Sims (2011), who use business survey data to measure firms' risk, as well as Bond and Cummins (2004), who use surveys of financial analysts.

There is also a growing literature that stresses the interaction of risk and economic activity propagated through financial, rather than physical frictions. Using a model with financial frictions, Gilchrist et al. (2010) argue that increases in firm risk lead to an increase in bond premia and the cost of capital which, in turn, triggers a decline in investment activity and measured aggregate productivity. Arellano et al. (2011) show that firms downsize investment projects to avoid default when faced with higher risk. Chugh (2011), who explains the dynamics of leverage with shocks to micro-level risk, finds in accordance with the results in this paper only a small business cycle impact of risk shocks. Dorofeenko et al. (2008) is another example in this literature.

More recently, the literature has also started investigating risk shocks in environments with nominal rigidities and their interaction with precautionary saving. For example, Basu and Bundick (2011) study the impact of shocks to aggregate risk and argue that in a model without nominal rigidities (and without capital adjustment frictions) increases in aggregate risk lead to an investment boom, induced by the interest rate decline from increased precautionary saving. The result is a lack of comovement between consumption on the one hand and output, investment and employment on the other hand. In a model with sticky prices where output is essentially demand-determined, however, a decline in consumption leads to a decline in output, employment and investment.[4]

The remainder of this paper is organized as follows: Section 2 presents the empirical analysis from the USTAN data. Section 3 explains the model. Section 4 describes its calibration and Section 5 discusses the results.

2. Some Facts about Firm-Level Risk

Our firm-level data source is the USTAN data base from *Deutsche Bundesbank*. USTAN is a large annual balance sheet data base, which originated as a by-product of the Bundesbank's

[4]Other channels for the propagation of risk shocks have been considered in the literature: for example, search frictions in Schaal (2011), investment opportunities in Lee (2011), as well as agency problems in Narita (2011). Another literature stresses the importance of rare, but drastic surprise changes in the economic environment: see Gourio (2010).

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 data were then archived and collected into a database.[5]

USTAN has broader coverage in terms of firm size, industry and ownership structure than comparable U.S. data sets. 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 the firms in USTAN are from manufacturing. USTAN allows us to study instead virtually the entire nonfinancial private business sector. Specifically, firms that are in one of the following six 1-digit industries are included into the sample: agriculture, mining and energy, manufacturing, construction, trade, transportation and communication.

We model fluctuations in idiosyncratic risk as fluctuations in the cross-sectional standard deviation of firm-specific Solow residual growth, and use a standard Cobb-Douglas revenue production function at the firm level as a measurement device (as in the model used to evaluate the quantitative importance of risk fluctuations):

$$\log(y_{i,t}) = \log(z_t) + \log(\epsilon_{i,t}) + \theta \log(k_{i,t}) + \nu \log(n_{i,t}), \tag{1}$$

where $\epsilon_{i,t}$ is firm-specific and z_t aggregate productivity.[6] Labor input $n_{i,t}$ is assumed to be immediately productive, whereas capital $k_{i,t}$ is pre-determined and inherited from last period. The output elasticities of the production factors, ν and θ, are estimated as median shares of factor expenditures over gross value added within each 1-digit industry.[7] The

[5] Details about the USTAN data base, data treatment, including measurement error correction and sample selection issues, as well as the actual time series of the baseline firm-level risk measure are available in an online appendix.

[6] Disentangling firm-specific demand and supply shocks is not possible with USTAN, because firm-level prices are not observed. The notion of productivity here is revenue productivity. Firms are indifferent in their investment decisions as to whether higher revenues come from an increased idiosyncratic demand for their products or higher productivity of their input factors.

[7] These are, respectively: $\nu = 0.2182, \theta = 0.7310$ (agriculture); $\nu = 0.3557, \theta = 0.5491$ (mining and energy); $\nu = 0.5565, \theta = 0.2075$ (manufacturing); $\nu = 0.6552, \theta = 0.1771$ (construction); $\nu = 0.4536, \theta = 0.2204$ (trade); $\nu = 0.4205, \theta = 0.2896$ (transportation and communication).

resulting time series of the cross-sectional dispersion of firm-level log Solow residual growth spans 26 years from 1973 to 1998, and is based on a sample of 854,105 firm-year observations, which means an average cross-section size of 32,850 firms per year.

To focus on idiosyncratic changes that do not capture differences in industry-specific responses to aggregate shocks or permanent ex-ante firm heterogeneity, firm fixed and industry-year effects are removed from the observations on firm-level Solow residual growth. Measured Solow residuals will likely reflect true firm productivity with some error. We take this into account and perform a measurement error correction, estimating the size of the measurement error by comparing the variances of one- and two-year Solow residual growth rates.

Intuitively, the importance of fluctuations in idiosyncratic risk depends on its volatility and its cyclicality. Since any data treatment necessarily involves many decisions, this paper reports a range of results. This will allow us in the model section to compute lower and upper bound scenarios for the aggregate importance of idiosyncratic risk fluctuations. The results from the whole sample after removing fixed effects and the measurement error correction are used in the 'Baseline' calibration. The first row of Table 1 shows that firm-level risk fluctuates on average 4.72 percent as a fraction of average firm risk, 0.09. It is also countercyclical, as measured by the contemporaneous correlation of firm-level risk with the cyclical component of the real gross value added of the nonfinancial private business sector. This confirms the aforementioned results in the literature that have found various dispersion measures of firm-level realization or expectation variables to be countercyclical.

–Table 1 about here–

Table 1 also displays the cyclical properties of the cross-sectional standard deviation of Solow residual growth as well as average firm-level risk for various ways of cutting the sample and treating the data. For instance, the small differences between row (1) and (4) indicate that the observed dispersion in the raw data mostly comes from idiosyncratic shocks. Regarding firm size (rows 2, 3 and 6), larger firms tend to have larger risk fluctuations. Row 7 checks to what extent the cyclicality results are driven by cyclical changes in sample composition (e.g. small, high-risk firms dropping out in recessions) by restricting the analysis to firms that are almost always in the sample, i.e. have at least 20 observations of Solow residual changes.[8] Finally, focusing on specific industries (manufacturing, row 8) and ownership structures (row 5), tends to increase the strength of measured risk fluctuations.

We base our lower bound ('LB') calibration scenario (where the coefficient of variation of firm-level risk is halved) loosely on the second row, which displays the cyclical properties of firm-level risk for small firms. Small firms are still underrepresented in USTAN. The upper bound ('UB') calibration scenario is loosely based on the third row, which delivers the strongest risk fluctuations. To be conservative this value is roughly doubled 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', does not substantially increase the volatility of risk over and above what each of these features alone does (see row (9) of Table 1). Any other combination of characteristics would not have left sufficient data to yield reliable results. These results show that one should be cautious when inferring the importance of risk fluctuations from data sources that are heavily geared towards manufacturing, publicly traded firms or large firms. One might overstate risk fluctuations.[9]

– Table 2 about here –

Are the micro-level risk processes in Germany and the U.S. comparable? Focussing on output-based growth measures (which is what Bloom et al., 2010, make publicly available),

[8]Since the sample design of USTAN does not lead to a strictly representative sample, Heckman (1976) style regressions are run to check whether sample selection is important for the results. Correcting for sample selection leaves the series of productivity dispersion virtually unchanged.

[9]One of the strengths of the USTAN data set is that it allows for a comparison of the extent of firm-level risk fluctuations across industries. The combined retail and wholesale trade sector, for example, features a similar volatility and cyclicality of risk as the overall USTAN data set. 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 a smaller share in aggregate production and employment than it has in Germany.

Table 2 compares USTAN results with the readily available U.S. evidence and shows that both economies have similar idiosyncratic risk processes, with the U.S. exhibiting slightly higher risk fluctuations than Germany: the volatility of the cross-sectional interquartile range of output growth from the USTAN data (row 3), 8.00%, is close to the corresponding number in the ASM (row 4), 9.80%.[10]

Table 2 demonstrates that the lower and upper bound scenarios – half and quadruple the coefficient of variation ('CV') of the baseline scenario – comfortably cover the available U.S. evidence. This means that to the extent that our model simulations reveal little aggregate effects of risk fluctuations, these results are not driven by the use of German data.

3. The Model

The empirical results from the previous section are used to calibrate a simple heterogeneous-firm model that features 'wait-and-see' effects of risk. The models in Khan and Thomas (2008) as well as Bachmann, Caballero and Engel (2011) serve as starting points. 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. This stark timing assumption is made to give risk shocks the best chance to have the most direct effect possible.[11]

[10] Measured *average* micro-level risk is higher in the ASM as it is plant-level data, while USTAN is firm-level data.

[11] The alternative 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, has little impact on the main results.

3.1. *Firms - The Physical Environment*

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 pre-determined capital stock (k) and labor (n), according to the following Cobb-Douglas decreasing-returns-to-scale production function ($\theta > 0$, $\nu > 0$, $\theta + \nu < 1$):

$$y = z\epsilon k^{\theta} n^{\nu}, \tag{2}$$

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 of 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.

The trend growth rate of aggregate productivity is denoted by $(1 - \theta)(\gamma - 1)$, so that aggregate output and capital grow at rate $\gamma - 1$ along the balanced growth path. From now on k and y (and later aggregate consumption, C) are understood to be denoted in efficiency units.

Each period a firm draws its current cost of capital adjustment, $0 \leq \xi \leq \bar{\xi}$, which is denominated in units of labor, from a time-invariant distribution, G. 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.[12]

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

[12] An experiment with a specification, where adjustment costs are deterministic, shows little impact on the results of the paper.

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

The firms' distribution over (ϵ, k) is denoted 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.

3.2. Firms - The Dynamic Programming Problem

Following Khan and Thomas (2008), we state the dynamic programming problem of a firm in terms of utils of the representative household (rather than physical units), and denote the marginal utility of consumption by $p = p\big(z, \sigma(z'), \sigma(\epsilon'), \mu\big)$. This is the kernel that firms use to price output streams.

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:

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

With this notation the dynamic programming problem becomes:

$$V^1\big(\epsilon, k, \xi; z, \sigma(z'), \sigma(\epsilon'), \mu\big) = \max_n \{\mathrm{CF} + \max(V_{\text{no adj}}, \max_{k'}[-\mathrm{AC} + V_{\text{adj}}])\}, \tag{4}$$

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:

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

$$V_{\text{no adj}} = \beta \mathrm{E}\big[V^0\big(\epsilon', (1-\delta)k/\gamma; z', \sigma(z''), \sigma(\epsilon''), \mu'\big)\big], \tag{5b}$$

$$AC = \xi\omega\big(z, \sigma(z'), \sigma(\epsilon'), \mu\big) p\big(z, \sigma(z'), \sigma(\epsilon'), \mu\big), \tag{5c}$$

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

where both expectation operators average over next period's realizations of the aggregate and idiosyncratic shocks, conditional on this period's values. 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 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.

3.3. Households

There is a continuum of identical households. They have a standard felicity function in consumption and labor:[13]

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

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:

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

and from the intratemporal first-order condition:

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

3.4. Solution

The recursive competitive equilibrium of this economy requires the usual optimality and market clearing conditions, which is omitted here for the sake of brevity. It is well-known that (4) is not computable, because μ is infinite dimensional. Following Krusell and Smith (1997, 1998) the distribution, μ, is approximated by a finite set of its moments, and its evolution,

[13] Experiments with a CRRA of 3 showed little impact on the results of the paper.

Γ, by a simple log-linear rule. As usual, aggregate capital holdings, $\bar{k}$, are included into this rule. 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, the equilibrium pricing function is approximated by a log-linear rule:

$$\log \bar{k}' = a_k\big(z, \sigma(z'), \sigma(\epsilon')\big) + b_k\big(z, \sigma(z'), \sigma(\epsilon')\big) \log \bar{k} + c_k\big(z, \sigma(z'), \sigma(\epsilon')\big) \log std(\log(\epsilon)), \quad \text{(9a)}$$

$$\log p = a_p\big(z, \sigma(z'), \sigma(\epsilon')\big) + b_p\big(z, \sigma(z'), \sigma(\epsilon')\big) \log \bar{k} + c_p\big(z, \sigma(z'), \sigma(\epsilon')\big) \log std(\log(\epsilon)). \quad \text{(9b)}$$

Given (8), it is unnecessary to specify an equilibrium rule for the real wage. The log-linear forms (9a)– (9b) are posited, and then it is checked that in equilibrium they yield a good fit to the actual law of motion. The R^2 for capital in the baseline calibration are all above 0.999. For the marginal utility of consumption they exceed 0.993.[14]

Substituting $\bar{k}$ and $std(\log(\epsilon))$ for μ into (4) and using (9a)–(9b), (4) 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 $K = K\big(\epsilon, k, \xi; z, \sigma(z'), \sigma(\epsilon'), \bar{k}, std(\log(\epsilon))\big)$. This problem is solved by value function iteration on V^0, applying multivariate spline techniques that allow for a continuous choice of capital when the firm adjusts.

With these policy functions, a model economy can be simulated *without* imposing the equilibrium pricing rule (9b). Rather, the market-clearing conditions are imposed and the pricing kernel solved for at every point in time of the simulation. This generates a time series of $\{p_t\}$ and $\{\bar{k}_t\}$ endogenously, on which the assumed rules (9a)–(9b) 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.

[14] Of course, $std(\log(\epsilon))$ has an analytically known law of motion. The lowest R^2 for the capital rule without $std(\log(\epsilon))$ is just above 0.88, and for the marginal utility of consumption it is just above 0.97.

4. Calibration

This Section discusses the calibration of the model from Section 3: standard technology and preference parameters, followed by the idiosyncratic and aggregate shock processes and then the fixed capital adjustment cost parameter.

4.1. Technology and Preference Parameters

The model period is a year, which 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: The discount factor is $\beta = 0.98$, and the depreciation rate is $\delta = 0.094$. δ is computed from German national accounting data (VGR) for the nonfinancial private business sector. Given this depreciation rate, the long-run growth factor $\gamma = 1.014$ is picked, in order to match the time-average aggregate investment rate in the nonfinancial private business sector of 10.8%. This number for γ is also consistent with German long-run growth rates. The disutility of work parameter $A = 2$ generates an average time spent at work of 1/3. The output elasticities of labor and capital are set 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.[15] Table 3 summarizes these and the following parameter choices.

– Table 3 about here –

4.2. Idiosyncratic Shocks

The standard deviation of idiosyncratic productivity shocks is calibrated to $\bar{\sigma}(\epsilon) = 0.0905$ (see the first row of Table 1 in Section 2) and set $\rho_\epsilon = 0.95$. This process is discretized on

[15] 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.

a 19−state-grid, using Tauchen's (1986) procedure with mixed Gaussian normals.[16] 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 $\sigma(\epsilon')$.

4.3. Aggregate Shocks

The baseline case (Section 5.1) abstracts from time-varying aggregate risk and correlation between aggregate productivity and idiosyncratic risk. Both themes will be taken up in Sections 5.3 and 5.4, respectively. Thus, to compute ρ_z and $\bar{\sigma}(z)$, an AR(1)-process is estimated for the linearly detrended natural logarithm of aggregate Solow residuals computed from German national accounting data. This estimation leads to $\rho_z = 0.5223$ and $\bar{\sigma}(z) = 0.0121$. This process is discretized on a 5−state grid, using Tauchen's (1986) procedure.

An AR(1)-process is also estimated for the linearly detrended cross-sectional standard deviation of the first differences of the natural logarithm of firm-level Solow residuals, as computed in the baseline case, i.e. eliminating sectoral and firm-level fixed effects from Solow residual growth and correcting the data for measurement error (see the first row of Table 1 in Section 2). This leads to $\rho_{\sigma(\epsilon)} = 0.5800$ and $\sigma_{\sigma(\epsilon)} = 0.0037$. 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 there is no 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 (Bloom et al., 2010). We do not want to take a stand on how 'catastrophic', i.e. strong but rare, a risk shock is. Instead, normality of risk shocks, which is supported by the data, is assumed. 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

[16] Since idiosyncratic productivity shocks in the data also exhibit above-Gaussian kurtosis - 4.4480 on average -, and since the fixed adjustment costs parameters will be identified by the kurtosis of the firm-level investment rate (together with its skewness), it is important to avoid attributing excess kurtosis in the firm-level investment rate to lumpy investment, when the idiosyncratic driving force itself has excess kurtosis. The measured excess kurtosis is incorporated into the discretization process for the idiosyncratic productivity state 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.

events such as a doubling of firm-level risk has not occurred in U.S. post war data, and the German data do not exhibit extreme risk shocks, either.

The importance of shocks to firm-level risk for aggregate fluctuations is gauged by its time series coefficient of variation, which for the 'Baseline' case equals: $CV_{risk} = 4.72\%$. It will be shown 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 the results depend on CV_{risk}, two additional scenarios are considered: a 'Lower Bound' scenario with half the CV_{risk} (roughly based on the second row of Table 1 in Section 2), and an 'Upper Bound' scenario, where CV_{risk} is quadrupled relative to the baseline case. To be conservative, the highest CV_{risk} found in the data, namely the size-weighted cross-sectional standard deviation of firm-level Solow residual growth (see the third row of Table 1 in Section 2), is doubled for this scenario.

4.4. Adjustment Costs

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) 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 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. Therefore the skewness and kurtosis of firm-level investment rates are used to identify $\bar{\xi}$. Given the following set of parameters $\left\{\beta, \delta, \gamma, A, \nu, \theta, \bar{\sigma}(\epsilon), \rho_\epsilon, \bar{\sigma}(z), \rho_z, \sigma_{\sigma(\epsilon)}, \rho_{\sigma(\epsilon)}\right\}$, the adjustment costs, $\bar{\xi}$, are calibrated to minimize the Euclidean distance, $\Psi(\bar{\xi})$, between the time-average firm-level investment rate skewness and kurtosis produced by the model and the data. To take into account the different precision at which skewness and kurtosis are estimated, both are weighted with the inverse of their time-series standard deviation.

The following Table 4 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 $\bar{\xi} = 0.25$, our baseline case.[17]

– Table 4 about here –

This implies average costs conditional on adjustment equivalent to roughly 7% of annual firm-level output, which is well in line with estimates from the U.S. (see Bloom, 2009, Table IV, for an overview). Moreover, the last column of Table 4 shows that the baseline model matches fairly well the fraction of firms with lumpy investment in any given year, measured as the investment rate that is larger than 20 percent in absolute value. This statistic is 13.80% in the USTAN data and has been commonly used in the literature as a measure of the fraction of investment episodes that can reasonably be considered lumpy (see Cooper and Haltiwanger, 2006, as well as Gourio and Kashyap, 2007).

5. Results

This set-up can be used to evaluate the quantitative importance of capital adjustment frictions in propagating risk shocks through 'wait-and-see' effects. This is done in two steps. First, a model is analyzed where firm-level risk shocks are introduced as an independent process alongside standard aggregate productivity shocks, see Table 5. In extensions, Table 7, aggregate productivity shocks and risk shocks have the correlation structure observed in the data ('Correlated Processes'). Finally a case with time-varying aggregate risk added to the 'Baseline' model is analyzed ('Aggregate Risk').

[17] Quadrupling $\bar{\xi}$ to experiment with a stronger 'wait-and-see' motive has little impact on results.

5.1. *Baseline Results - Independent Shocks*

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). Figure 1 confirms that this characteristic impulse response survives both the introduction of independent standard first-moment aggregate productivity shocks as well as general equilibrium real interest rate and wage adjustments. The 'Baseline Model' thus features the expected conditional response to risk shocks. In fact, the initial investment collapse is somewhat stronger in general equilibrium due to a wealth effect, whereas overall fluctuations are dampened. 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 about here –

To answer our initial question and to understand the importance of time-varying risk for the business cycle, however, studying the sign of the conditional responses might not be sufficient.

– Table 5 about here –

Table 5 displays the unconditional business cycle properties of various variants of the baseline model, i.e. firm-level risk shocks introduced as an independent process alongside standard aggregate productivity shocks. A comparison between column (1), the baseline calibration featuring a model with an intermediate estimate of the $CV_{risk} = 4.72\%$, and the constant-risk, RBC-style model in column (4) shows that their business cycle statistics are essentially identical. Only in the extreme case of a $CV_{risk} = 18.88\%$, column (3), the upper bound calibration, can risk fluctuations contribute somewhat to the volatility of output.

Result 1. *Firm-level risk fluctuations added to first moment productivity shocks lead to similar business cycle dynamics as RBC models, unless firm-level risk is (counterfactually) highly volatile.*

Table 5, in its last column, also shows that the business cycle properties in Germany are roughly the same as in the U.S. (for instance, aggregate technology shocks alone explain a large fraction of the business cycle volatility of output in the 'RBC Model'), so that our results are not likely due to idiosyncrasies 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, and what the model predicts.

Could the mild effects of firm-level risk shocks be driven by general equilibrium price adjustments? After all, 'wait-and-see' is a partial equilibrium mechanism. The partial equilibrium counterparts of the general equilibrium output volatilities in columns (1)-(4) in Table 5 (the baseline, the lower and the upper bound calibration, as well as the RBC style model) are, respectively, 3.18%, 3.16%, 3.79%, and 3.17%. These numbers show that general equilibrium effects as usual dampen aggregate fluctuations, but that they are not causing firm-level risk shocks to be essentially neutral, when introduced into a standard RBC model – 'Baseline' and 'RBC model' behave also in partial equilibrium almost identical.

5.2. Risk Shocks Alone

Column (5) of Table 5, 'Risk Model', shows business cycle statistics for the same set of risk shocks as in the baseline model, but with aggregate productivity shocks shut down. Column (6), 'Psych. Risk', studies the same case, but here risk shocks only change the risk *perceptions* of firms, yet never materialize. This model specification features purely subjective uncertainty, whose fluctuations have almost no effects on aggregate volatility. The 'Risk model', featuring actual shocks to firm-level risk but no aggregate productivity shocks, produces in terms of standard deviations 15% of the output volatility in the data (i.e. 2% of its variance).

Result 2. *The literature has argued that firm-level risk fluctuations, propagated through capital adjustment frictions, might generate cycles through the concentration of economic activity in periods of relatively stable economic environments. We show that this mechanism is unlikely to be a major driver of the business cycle.*

Risk shocks alone also lead to a lack of comovement between consumption and the other macroeconomic aggregates. When a firm-level risk shock hits the economy, aggregate investment demand declines through the 'wait-and-see' mechanism, which leads to declining interest rates and higher consumption on impact. The opposite lack of comovement for an aggregate risk shock – decreased consumption due to increased precautionary saving and increased investment – has recently been discussed in a model without capital adjustment frictions by Basu and Bundick (2011). They argue that nominal rigidities can fix this lack of comovement. We surmise that the same line of argument applies to our model setup. In the short run, risk shocks decrease aggregate investment demand at all real rates. With nominal rigidities, it depends on monetary policy reactions whether the resulting multiplier on output is large enough to also cause a consumption decline.

5.3. Extensions: Correlated Shocks

Consider now the more general correlation structure between risk shocks and aggregate productivity shocks in the data. It will be shown that this not only makes the model fit better the *conditional* responses of various macroeconomic aggregates to an innovation in firm-level risk,[18] but it also reveals a new potential mechanism how risk shocks can generate aggregate fluctuations: they negatively predict aggregate productivity in the future.

To obtain conditional responses, three-variable VARs are estimated with the natural logarithm of aggregate Solow residuals, idiosyncratic risk and various aggregate activity variables. This ordering is then used in a simple Choleski-"identification". This is similar to the VAR in Bloom (2009), who orders stock market returns before stock market volatility,

[18]See Curdia and Reis (2011), who argue that allowing for correlated shocks in a standard medium-scale DSGE model leads to a better fit to U.S. macroeconomic data.

in order to identify pure second-moment shocks after controlling for first-moment shocks. While not deeply structural, this is as a different, but convenient and instructive way to summarize the data.

Figure 2 shows this exercise for aggregate output, aggregate investment, total hours, and consumption for the 'Baseline' specification with orthogonal productivity and risk shocks. The impact 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 'Baseline' calibration, i.e. independent first and second moment shocks, are just the opposite; they feature relatively short-run 'wait-and-see' dynamics, similarly to Bloom (2009). Given the short time series available, some of these impact estimates, notably for investment and consumption, are obtained with some imprecision. Eventually uncertainty shocks lead to protracted declines of economic activity in all four major macroeconomic aggregates. This finding is consistent with the results in Bond and Cummins (2004), Gilchrist et al. (2010) as well as Bachmann, Elstner and Sims (2011). Nevertheless, the risk responses of the 'Baseline' calibration appear broadly inconsistent with the data: they are not as pronounced and protracted.

However, the 'Baseline Model' abstracts from one important feature of the data: aggregate productivity and idiosyncratic risk are not orthogonal processes. We therefore solve and simulate an alternative model specification with correlated firm risk and aggregate productivity processes and feed into the model the joint dynamics estimated from the data for these two time series ('Correlated Processes'), see Figure 3. In this alternative specification, the impulse responses estimated on simulated model data are much closer to those in the data and qualitatively match the shape of the impulse responses from actual data for all four macroeconomic quantities.

–Figure 2 about here –

–Figure 3 about here –

A forecast error variance decomposition of the same VAR (see Table 6) both on the actual data and on the model simulated data supports the conclusion from Section 5.1 and from the impulse responses. In the data, risk fluctuations contribute a significant amount to the fluctuations of output, investment and total hours, especially at longer horizons. Yet, the 'Baseline' with only 'wait-and-see' dynamics does not generate the rising forecast error variance contribution profile found in the data; the 'Correlated Processes' calibration does.

– Table 6 about here –

Why does the 'Correlated Processes' calibration generate this better fit to the data? The introduction of risk shocks in 'Correlated Processes' also changes the stochastic properties of aggregate productivity – a feature absent in the calibrations with orthogonal shock processes. Shocks to risk change the conditional expectation of future productivity – higher risk signals a productivity decline as the coefficient of risk today on aggregate productivity tomorrow is negative (-1.9735) – which, in turn, has important general equilibrium implications: a general equilibrium wealth effect makes agents consume less and work more,[19] which on impact drives up output and – through a decrease in the real interest rate – investment.

To better understand the consequences of assuming a richer contemporaneous and dynamic correlation structure between aggregate productivity and firm-level risk, we also compute a specification, where actual firm-level risk is fixed at $\bar{\sigma}(\epsilon)$ and where $\sigma(\epsilon')$ is reinterpreted as a latent state variable that jointly evolves with z as in 'Correlated Processes'. This specification is denoted 'Forecast Model', because "risk" today then merely predicts productivity tomorrow, but does not change the *idiosyncratic* stochastic environment of the firms. In other words, in this case "risk" is exclusively a signal of future productivity. The impulse responses of 'Correlated Processes' and 'Forecast Model' in Figure 3 are similar,

[19]The real wage decreases after a risk shock both in the data and in the model simulated data from 'Correlated Processes'.

which suggests that the conditional effects of risk on aggregate activity are mainly driven by this signalling effect.

In a second specification, 'Naive Model', the same joint process for aggregate productivity and firm risk is used as in the 'Forecast Model', but the agents in the economy – naively – continue to use the univariate process for productivity from the 'RBC Model' when computing their optimal policies. Table 7 shows how the unconditional business cycle moments evolve from the 'RBC Model' to 'Correlated Processes'. The changes from the 'Forecast Model' to 'Correlated Processes' identify the specific effects of time-varying firm risk on aggregate fluctuations.

– Table 7 about here–

It is mostly volatilities that are affected by introducing the second shock. Output fluctuates more, but these output fluctuations are dampened, when actual risk shocks hit the economy. The responsiveness of the economy to productivity shocks decreases in the volatility of risk shocks. The correlations of aggregate quantities are 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'.

Result 3. *The conditional impulse responses of aggregate quantities to a risk innovation in a model where risk and productivity shocks are uncorrelated appear to be inconsistent with their data counterparts. A model with correlated risk and productivity shocks matches the data better in terms of conditional impulse responses.*

5.4. Extensions: Aggregate Risk

Finally, in column (5) of Table 7, we add time-varying *aggregate* risk ('Aggregate-Risk Model') to the 'Baseline' calibration with time-varying firm-level risk and aggregate productivity shocks, but maintain the independence assumption between the latter. To save on one state variable in the computation, this additional shock is introduced as perfectly correlated with the state of firm-level risk. This way, the impact of time-varying aggregate risk can be expected to be maximized. 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 $\bar{\sigma}(\epsilon) = 0.0905$, $\sigma(z')$ moves around in the same way on a 5-state grid, centered around $\bar{\sigma}(z) = 0.0121$. The grid width of the latter is used to calibrate the time series coefficient of variation of aggregate risk to roughly 36%.[20] 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. Table 7 shows that again the effects are mild. The business cycle statistics of the 'Baseline' with time-varying aggregate and idiosyncratic risk are very similar to those from the 'RBC Model', with some increase in aggregate volatility.

To understand this result note that the average idiosyncratic risk, $\bar{\sigma}(\epsilon) = 0.0905$, is almost an order of magnitude larger than the average aggregate risk, $\bar{\sigma}(z) = 0.0121$. Since standard deviations are not additive, the combined small aggregate and large idiosyncratic risk, i.e. the standard deviation of the combined productivity shock, is close to the one of idiosyncratic productivity. For example, starting from a situation of average aggregate and idiosyncratic risk, the combined risk the firm faces is 0.0913. Jumping from here to a situation with highest aggregate risk (and average idiosyncratic risk) would lead to a combined risk of 0.0935, a 2.4% 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 0.1048 or almost 15%.

Result 4. *Aggregate risk fluctuations added to aggregate productivity shocks and idiosyncratic risk fluctuations lead to similar business cycle dynamics as RBC models.*

[20] 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. Yet, most results lie between 30 and 40 percent.

6. Conclusion

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. The model features a 'wait-and-see' effect for investment after surprise increases in firm-level risk. The model is disciplined using a rich German firm-level data set. Relative to the literature this data set allows us to uncover upward biases, when the volatility of firm-level risk, and thus its importance for aggregate fluctuations, is measured from data that focus on manufacturing or publicly traded firms. However, the main reason why this paper arrives at a somewhat different conclusion from the literature is our 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 business cycle volatility, and when juxtaposed to standard aggregate productivity shocks it does little to alter business cycle fluctuations. Correlated firm-level risk and aggregate productivity shocks improve the model fit, which suggests as a direction for future research to understand better the precise structure underlying the dynamic correlations and the direction of causality between first- and second-moment shocks.

References

Arellano, C., Bai, Y., Kehoe, P., 2011. Financial markets and fluctuations in uncertainty. Federal Reserve Bank of Minneapolis Research Department Staff Report.

Bachmann, R., Caballero, R., Engel, E., 2011. Aggregate implications of lumpy investment: New evidence and a DSGE model. Mimeo Yale University.

Bachmann, R., Elstner, S., Sims, E., 2011. Uncertainty and economic activity: Evidence from business survey data. NBER-WP 16143.

Bachmann, R., Moscarini, G., 2011. Business cycles and endogenous uncertainty. Mimeo Yale University.

Basu, S., Bundick, B., 2011. Uncertainty shocks in a model of effective demand. Mimeo Boston College.

Berger, D., Vavra, J. 2011. The dynamics of the U.S. price distribution. Mimeo Yale University.

Bernanke, B., 1983. Irreversibility, uncertainty, and cyclical investment. The Quarterly Journal of Economics 98/1, 85–106.

Bloom, N., 2009. The impact of uncertainty shocks. Econometrica 77/3, 623–685.

Bloom, N., Bond, S., Van Reenen, J., 2007. Uncertainty and investment dynamics. Review of Economic Studies 74, 391–415.

Bloom, N., Floetotto, M., Jaimovich, N., 2010. Really uncertain business cycles. Mimeo Stanford University.

Bond, S., Cummins, J., 2004. Uncertainty and investment: an empirical investigation using data on analysts' profits forecasts. Mimeo Federal Reserve Board.

Bontempi, M., Golinelli, R., Parigi, G., 2010. Why demand uncertainty curbs investment: Evidence from a panel of Italian manufacturing firms. Journal of Macroeconomics 32(1), 218–238.

Caballero, R., Engel, E., Haltiwanger, J., 1995. Plant-level adjustment and aggregate investment dynamics. Brookings Paper on Economic Activity 2, 1–54.

Cooper, R., Haltiwanger, J., 2006. On the nature of capital adjustment costs. Review of Economic Studies 73, 611–633.

Christiano, L., Motto, R., Rostagno, M., 2010. Financial factors in economic fluctuations. ECB WP 1192.

Chugh, S., 2011. Firm risk and leverage-based business cycles. Mimeo Boston College.

Curdia, V., Reis, R., 2011. Correlated disturbances and U.S. business cycles. Mimeo Columbia University.

Davis, S., Haltiwanger, J., Schuh, S., 1996. Job Creation and Destruction. MIT Press, Cambridge, Massachusetts.

Davis, S., Haltiwanger, J., Jarmin, R., Miranda, J., 2006. Volatility and dispersion in business growth rates: Publicly traded and privately held firms. NBER Macroeconomics Annual, 107–180.

Dixit, A., Pindyck, R., 1994. Investment under Uncertainty. Princeton University Press, Princeton, New Jersey.

Doepke, J., Weber, S., 2006. The within-distribution business cycle dynamics of German firms. Discussion Paper Series 1: Economic Studies, No 29/2006, Deutsche Bundesbank.

Doepke, J., Funke, M., Holly, S., Weber, S., 2005. The cross-sectional dynamics of German business bycles: A bird's eye view. Discussion Paper Series 1: Economic Studies, No 23/2005, Deutsche Bundesbank.

Doms, M., Dunne, T., 1998. Capital adjustment patterns in manufacturing plants. Review of Economic Dynamics 1, 409–429.

Dorofeenko, V., Lee, G., Salyer, K., 2008. Time-varying uncertainty and the credit channel. Bulletin of Economic Research 60:4, 375–403.

Fuss, C., Vermeulen, P., 2004. Firm's investment decisions in response to demand and price uncertainty. Applied Economics 40(18), 2337–2351.

Gilchrist, S., Sim, J., Zakrajsek, E., 2010. Uncertainty, credit spreads and investment dynamics. Mimeo Boston University.

Gourio, F., 2008. Estimating firm-level risk. Mimeo Boston University.

Gourio, F., 2010. Disaster risk and business cycles. The American Economic Review, forthcoming.

Gourio, F., Kashyap, A., 2007. Investment spikes: New facts and a general equilibrium exploration. Journal of Monetary Economics 54, 1–22.

Guiso, L., Parigi, G., 1999. Investment and demand uncertainty. The Quarterly Journal of Economics 114(1), 185–227.

Hassler, J., 1996. Variations in risk and fluctuations in demand - a theoretical model. Journal of Economic Dynamics and Control 20, 1155–1143.

Hassler, J., 2001. Uncertainty and the timing of automobile purchases. The Scandinavian Journal of Economics 103, 351–366.

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.

Henessy, C., Whited, T., 2005. Debt dynamics. The Journal of Finance LX/3, 1129–1165.

Higson, C., Holly, S., Kattuman, P., 2002. The cross-sectional dynamics of the US business cycle: 1950–1999. Journal of Economic Dynamics and Control 26, 1539–1555.

Higson, C., Holly, S., Kattuman, P., Platis, S., 2004: The business cycle, macroeconomic shocks and the cross section: The growth of UK quoted companies. Economica 71/281, 299–318.

Kehrig, M., 2010. The cyclicality of productivity dispersion. Mimeo Northwestern University.

Khan, A., Thomas, J., 2008. Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. Econometrica 76(2), 395–436.

Kilian, L., 1998. Small sample confidence intervals for impulse response functions. Review of Economics and Statistics 80, 218–230.

Krusell, P., Smith, A., 1997. Income and wealth Heterogeneity, portfolio choice and equilibrium asset returns. Macroeconomic Dynamics 1, 387–422.

Krusell, P., Smith, A., 1998. Income and wealth heterogeneity in the macroeconomy. Journal of Political Economy 106 (5), 867–896.

Lee, J., 2011. Does an aggregate increase in idiosyncratic volatility cause a recession? Mimeo University of Chicago.

Narita, F., 2011. Hidden actions, risk-taking and uncertainty shocks. Mimeo University of Minnesota.

Schaal, E., 2011. Uncertainty, productivity and unemployment in the great recession. Mimeo Princeton University.

Tauchen, G., 1986. Finite state Markov-chain approximations to univariate and vector autoregressions. Economics Letters 20, 177–181.

Table 1: The (Cyclical) Properties of Firm-Level Risk

		CV	Cyclicality	Mean
(1)	Baseline - FE and ME	4.72%	-0.47	0.09
(2)	LB: Smallest 25% firms (capital) - FE and ME	2.73%	-0.48	0.11
(3)	UB: Size weighted (capital) - FE and ME	8.38%	-0.62	0.08
(4)	Raw - ME	4.10%	-0.44	0.11
(5)	Publicly traded - FE and ME	7.34%	-0.29	0.08
(6)	Largest 5% firms (capital) - FE and ME	7.28%	-0.46	0.08
(7)	20 obs. - FE and ME	7.26%	-0.38	0.08
(8)	Manufacturing - FE and ME	6.08%	-0.61	0.08
(9)	20 obs., manufacturing - FE and ME	7.52%	-0.50	0.08

Notes: Entries refer to the linearly detrended 1973-1998 series of cross-sectional standard deviation of firm-specific log Solow residual growth purged of measurement error ('ME') and firm-specific as well as industry-year fixed effects ('FE'), in short: firm-level risk. Columns displays the time-series coefficient of variation ("CV") of firm-level risk, its correlation with HP(100)-filtered series of real gross value added of the non-financial private business sector ("Cyclicality"), and the time average firm-level risk ("Mean"). Row (1) is the baseline data treatment; (2) restricts the sample to the 25% smallest firms in terms of capital stock; (3) considers a capital-weighted cross-sectional standard deviation of Solow residual changes; (4) has no 'FE'-treatment. Rows (5)-(8) restrict the sample to publicly traded firms, to large firms, to firms with at least 20 observations of Solow residual changes, and to manufacturing firms, respectively. The last row (9) combines the last two criteria.

Table 2: MICRO-LEVEL OUTPUT GROWTH DISPERSIONS: GERMANY - U.S. COMPARISON

		CV		Cyclicality		Mean	
		STD	IQR	STD	IQR	STD	IQR
(1)	USTAN Basesample - FE and ME	5.66%		-0.45		0.11	
(2)	USTAN Manufacturing - FE and ME	7.09%		-0.59		0.11	
(3)	USTAN Manufacturing	5.01%	8.00%	-0.54	-0.50	0.15	0.16
(4)	ASM Manufacturing		9.80%		-0.22		0.37

Notes: see notes to Table 1; however, here entries refer to output growth instead of Solow residual growth. In addition to risk being measured as the cross-sectional standard deviation ('STD'), rows (3) and (4) also report the interquartile range ('IQR'), for which our measurement error correction is not possible. Row (3) does not correct for fixed effects or measurement error, because the available U.S. evidence in Row (4) does not either. Row (4) refers to the 1973-2005 IQR series for plant-level output growth rates in the ASM, available from $http://www.stanford.edu/\ nbloom/index\_files/Page315.htm$. The U.S. output measure is the cyclical component of industrial production.

Table 3: Model Parameters

Parameter		Value	Data Source
discount factor	β	0.98	standard
disutility of labor	A	2	standard
depreciation rate	δ	0.094	VGR Data
long-run growth factor	γ	1.014	VGR Data
time-average aggregate risk	$\bar{\sigma}(z)$	0.0121	VGR Data
autocorrelation of aggregate productivity	ρ_z	0.5223	VGR Data
output elasticity of labor	ν	0.5565	USTAN
output elasticity of capital	θ	0.2075	USTAN
time-average idiosyncratic risk	$\bar{\sigma}(\epsilon)$	0.095	USTAN
autocorrelation of idiosyncratic productivity	ρ_ϵ	0.95	USTAN
volatility of idiosyncratic risk	$\sigma_{\sigma(\epsilon)}$	0.0037	USTAN
persistence of idiosyncratic risk	$\rho_{\sigma(\epsilon)}$	0.58	USTAN
adjustment cost parameter	$\bar{\xi}$	0.2	USTAN

Notes: this table summarizes the values of the model parameters and the data sources used to calibrate them.

Table 4: Calibration of Adjustment Costs - $\bar{\xi}$

$\bar{\xi}$	Skewness	Kurtosis	$\Psi(\bar{\xi})$	Adj. costs/ Unit of Output	Fraction Lumpy Adj.
0	-0.0100	3.5696	18.9638	0%	53.45%
0.01	0.8968	5.1370	10.7881	0.74%	31.80%
0.1	2.2625	9.6529	3.5656	3.53%	15.39%
0.25 (BL)	2.8859	12.3950	2.9202	6.97%	11.44%
0.5	3.3406	14.7488	3.6480	12.09%	9.20%
0.75	3.5964	16.2341	4.5525	16.97%	8.12%
1	3.7739	17.3431	5.4105	21.80%	7.45%
5	4.7614	24.8881	14.4205	110.32%	4.63%
Data	2.1920	20.0355			13.80%

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, the precision-weighted Euclidean distance of the model's cross-sectional skewness and kurtosis of investment-rates to their data counterparts. The fifth column shows the average adjustment costs conditional on adjustment as a fraction of the firm's annual output. The last column displays the fraction of firms with lumpy capital adjustments in any given year, i.e. firms with investment rates that are larger than 20 percent in absolute value.

Table 5: BUSINESS CYCLE STATISTICS FOR THE BASELINE RESULTS

	(1) Base-line (BL)	(2) Lower Bound (LB)	(3) Upper Bound (UB)	(4) RBC Model	(5) Risk Model	(6) Psych. Risk Model	(7) Data
Volatility of Output	2.07%	2.07%	2.20%	2.07%	0.34%	0.08%	2.30%
Volatility of aggregate variables relative to output volatility							
Consumption	0.41	0.41	0.44	0.41	0.79	0.99	0.78
Investment	4.20	4.18	4.50	4.18	7.26	10.29	1.90
Employment	0.69	0.69	0.79	0.68	1.48	2.01	0.78
First-order Autocorrelation							
Output	0.30	0.30	0.33	0.31	0.47	-0.01	0.48
Consumption	0.54	0.55	0.51	0.55	0.42	-0.06	0.67
Investment	0.23	0.24	0.23	0.25	0.18	-0.06	0.42
Employment	0.22	0.23	0.21	0.24	0.16	-0.06	0.61
Contemporaneous Correlation with Aggregate Output							
Consumption	0.86	0.88	0.72	0.88	-0.12	-0.92	0.66
Investment	0.97	0.97	0.94	0.97	0.86	0.98	0.83
Employment	0.95	0.96	0.89	0.96	0.82	0.98	0.68
Contemporaneous Correlation with Aggregate Consumption							
Investment	0.72	0.74	0.46	0.75	-0.62	-0.97	0.60
Employment	0.67	0.70	0.34	0.71	-0.67	-0.98	0.36

Notes: Columns (1),(2),(3) refer to simulations with two orthogonal aggregate shocks, to z and $\sigma(\epsilon')$. They differ in the time series coefficient of variation of $\sigma(\epsilon')$. In (1), it is 4.72%, which is halved in (2) and quadrupled in (3). Column (4) refers to a simulation, where the only aggregate shock is to z, and (5) refers to a simulation, where the only aggregate shock is to $\sigma(\epsilon')$, whose time series coefficient of variation is 4.72%. '(6) Psych. Risk Model' uses the same simulated firm-level risk series as (5), but the risk series enters only into the firms' policy functions, whereas actual $\sigma(\epsilon')$ is fixed at $\bar{\sigma}(\epsilon)$. 'Data' (except for consumption) refers to the nonfinancial private business sector's aggregates. Consumption is aggregate consumption. All series, from data and model simulations, have been logged and HP-filtered with a smoothing parameter 100.

Table 6: Forecast Error Variance Decompositions

		Forecast Horizon		
		1Y	2Y	10Y
Output	Data	4.26%	15.93%	35.40%
	Correlated Processes	0.88%	16.98%	32.31%
	Baseline	0.56%	0.53%	0.85%
Investment	Data	2.56%	8.24%	27.74%
	Correlated Processes	2.97%	18.22%	31.70%
	Baseline	2.28%	2.22%	2.80%
Total Hours	Data	12.84%	10.65%	24.51%
	Correlated Processes	3.79%	18.60%	31.90%
	Baseline	3.13%	3.03%	4.11%

Notes: see notes to Figures 2 and 3.

Table 7: Business Cycle Statistics for the Extensions and Robustness Results

	(1) Correlated Processes	(2) Forecast Model	(3) Naive Model	(4) RBC Model	(5) Aggregate Risk	(6) Data
Volatility of Output	2.41%	2.59%	2.25%	2.07%	2.14%	2.30%
Volatility of aggregate variables relative to output volatility						
Consumption	0.36	0.33	0.42	0.41	0.41	0.78
Investment	4.36	4.53	4.17	4.18	4.21	1.90
Employment	0.70	0.75	0.68	0.68	0.70	0.78
First-order Autocorrelation						
Output	0.38	0.40	0.43	0.31	0.29	0.48
Consumption	0.60	0.65	0.61	0.55	0.53	0.67
Investment	0.31	0.33	0.38	0.25	0.23	0.42
Employment	0.30	0.31	0.38	0.24	0.22	0.61
Contemporaneous Correlation with Aggregate Output						
Consumption	0.89	0.87	0.88	0.88	0.87	0.66
Investment	0.98	0.98	0.97	0.97	0.97	0.83
Employment	0.97	0.97	0.96	0.96	0.95	0.68
Contemporaneous Correlation with Aggregate Consumption						
Investment	0.80	0.77	0.74	0.75	0.73	0.60
Employment	0.77	0.73	0.70	0.71	0.67	0.36

Notes: see notes to Table 5. Columns (1)-(3) refer to simulations, where there are two correlated exogenous aggregate states, z and s, which follow $\left(\begin{smallmatrix} z' \\ s' \end{smallmatrix}\right) = \left(\begin{smallmatrix} 0.4497 & -1.9735 \\ 0.0693 & 0.6753 \end{smallmatrix}\right)\left(\begin{smallmatrix} z \\ s \end{smallmatrix}\right) + \zeta$ with the matrix of standard deviations and the correlation coefficients for ζ being $\left(\begin{smallmatrix} 0.0095 & 0.2372 \\ 0.2372 & 0.0034 \end{smallmatrix}\right)$. In the 'Correlated Processes' specification $\sigma(\epsilon') = \bar{\sigma}(\epsilon) + s$ such that $\sigma(\epsilon')$'s time series coefficient of variation is 4.72% as in 'Baseline' in Table 5. In the 'Forecast Model' (Column (2)) specification, actual firm-level risk is fixed at $\bar{\sigma}(\epsilon)$ and s is simply a latent state variable, which jointly evolves with z. The joint process for z and s is discretized by a two-dimensional analog of Tauchen's (1986) procedure. Column (3) is the same as (2), except that agents do not take into account that latent random variable. Column (4) refers to a variant of the 'Baseline' (in Table 5), where also $\sigma(z')$ varies over time. It is perfectly correlated with $\sigma(\epsilon')$ and its time series coefficient of variation is 36.05%.

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

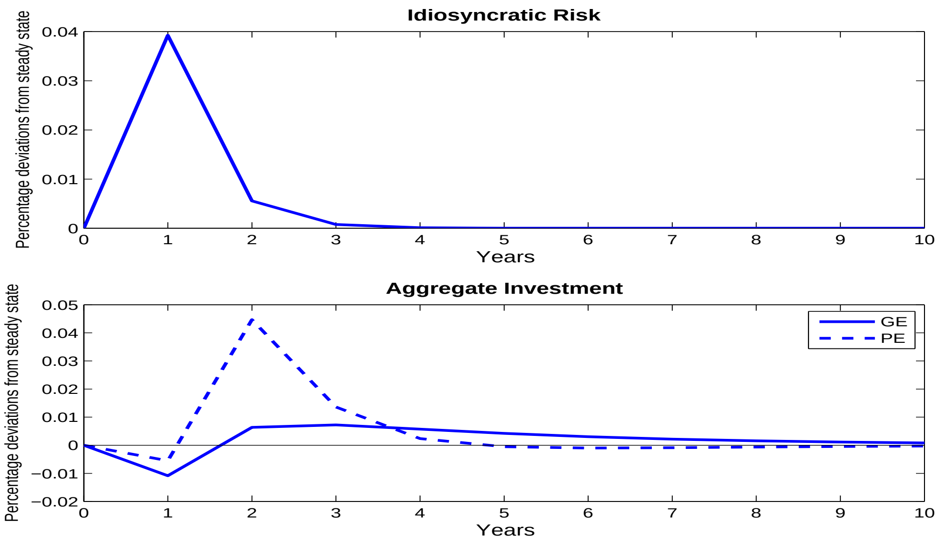

Notes: impulse responses are computed by increasing $\sigma(\epsilon')$ one standard deviation and letting it return to its steady state value, according to the AR(1) process estimated in Section 4. 'GE' stands for general equilibrium and takes real wage and interest rate movements into account. 'PE' stands for partial equilibrium and fixes the real wage and the interest rate at their time series averages from the 'GE' model.

Figure 2: Impulse Responses to an Innovation in Idiosyncratic Risk - Data and 'Baseline' Model

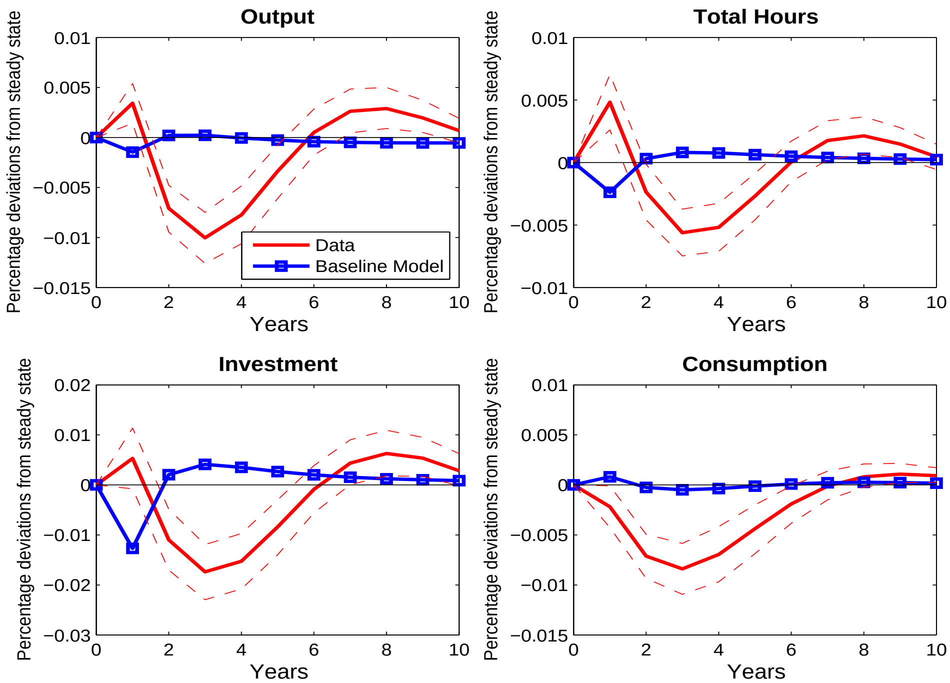

Notes: impulse response functions from SVARs with the linearly detrended natural logarithm of aggregate Solow residuals (ordered first), the linearly detrended idiosyncratic risk (ordered second) and HP(100)-filtered aggregate output/total hours/consumption/investment (ordered third). The dotted red 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 with squared markers. The model refers to the 'Baseline' calibration as in Table 5.

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

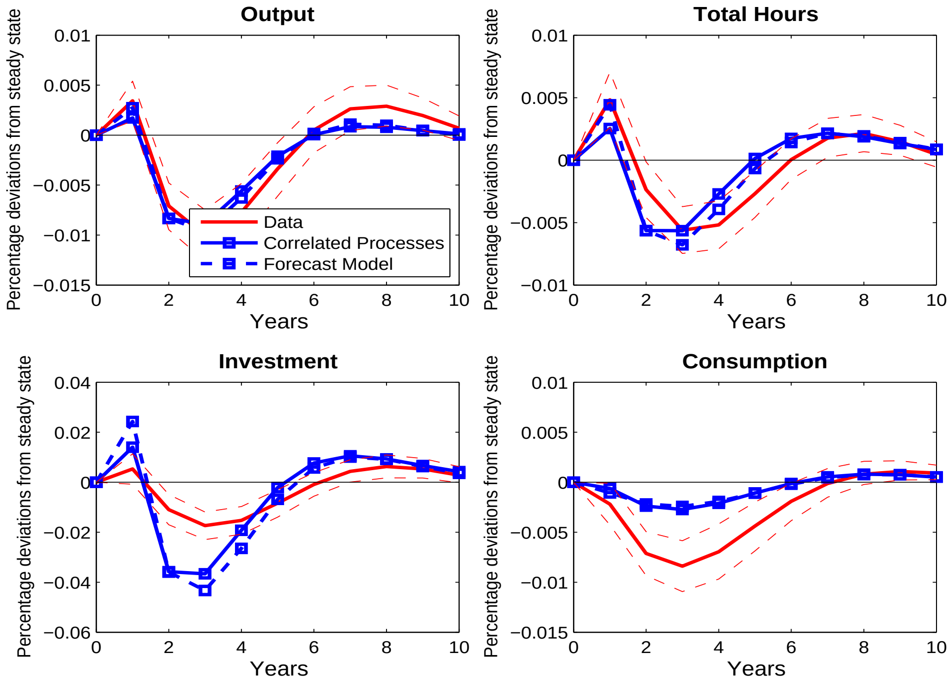

Notes: see notes to Figure 2. 'Correlated Processes' and 'Forecast Model' refer to simulations, where there are two correlated exogenous aggregate states, z and s, which follow $\binom{z'}{s'} = \begin{pmatrix} 0.4497 & -1.9735 \\ 0.0693 & 0.6753 \end{pmatrix} \binom{z}{s} + \zeta$ with the matrix of standard deviations and the correlation coefficients for ζ being $\begin{pmatrix} 0.0095 & 0.2372 \\ 0.2372 & 0.0034 \end{pmatrix}$. In the 'Correlated Processes' specification $\sigma(\epsilon') = \bar{\sigma}(\epsilon) + s$ such that $\sigma(\epsilon')$'s time series coefficient of variation is 4.72% as in 'Baseline' in Table 5. In the 'Forecast Model' specification, actual firm-level risk is fixed at $\bar{\sigma}(\epsilon)$ and s is simply a latent state variable, which jointly evolves with z. The joint process for z and s is discretized by a two-dimensional analog of Tauchen's (1986) procedure.