Investment Dispersion and the Business Cycle

Rüdiger Bachmann[a,*] Christian Bayer[b,†‡]

[a]RWTH Aachen University, Templergraben 64, Rm. 513, 52062 Aachen, Germany.

[b]University of Bonn, Adenauerallee 24-42, 53113 Bonn, Germany.

June 3, 2013

Abstract

The cross-sectional dispersion of firm-level investment rates is procyclical. This makes investment rates different from productivity, output and employment growth, which have countercyclical dispersions. A calibrated heterogeneous-firm business cycle model with nonconvex capital adjustment costs and countercyclical dispersion of firm-level productivity shocks replicates these facts and produces a correlation between investment dispersion and aggregate output of 0.53, close to 0.45 in the data. We find that small shocks to the dispersion of productivity, which in the model constitutes firm risk, suffice to generate the mildly procyclical investment dispersion in the data but do not produce serious business cycles.

JEL Codes: E20, E22, E30, E32.
Keywords: Ss model, RBC model, cross-sectional firm dynamics, lumpy investment, aggregate shocks, idiosyncratic shocks, risk shocks, heterogeneous firms.

*Corresponding author. Phone: +49-241-8096203. Fax: +49-241-8092649. E-mail: ruediger.bachmann@rwth-aachen.de. Also: NBER (United States), CESifo (Germany), and ifo (Germany).

†E-mail: christian.bayer@uni-bonn.de.

‡We thank Dirk Krüger, Giuseppe Moscarini, Matthew Shapiro as well as four anonymous referees for their comments. We are grateful to seminar/meeting participants at Amsterdam, the ASSA (San Francisco), Bonn, Duke, ECARES (Brussels), ESSIM 2009, Georgetown, Groningen, Johns Hopkins, Innsbruck, Michigan-Ann Arbor, Michigan State, the Minneapolis Fed, Notre Dame, Regensburg, the SED (Istanbul), 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. This paper formerly circulated under the current title as NBER-WP 17063 and under: "The Cross-section of Firms over the Business Cycle: New Facts and a DSGE Exploration" as CESifo-WP 2810. Any remaining errors are our own.

This paper establishes a novel business cycle fact: the cross-sectional dispersion of firm-level investment rates is robustly and significantly positively correlated with the business cycle.[1] The procyclicality of the firm-level investment rate dispersion is noteworthy for at least two reasons: first, it relates to a growing empirical literature on the time-series dynamics of the distribution of micro-level variables; second, as we will argue, it has important implications for the recent macroeconomic literature on the role of uncertainty or risk shocks as drivers of the business cycle. The procyclicality of investment dispersion places a robust and tight upper bound on the aggregate importance of firm-level risk shocks.

Researchers have documented that, across different countries and data sets, the dispersion of changes in firm- (or plant-) level variables, such as output, productivity, prices and business forecasts, is robustly countercyclical.[2] We find similar comovement patterns for firm-level output and productivity growth and add firm-level employment growth to the list of variables the dispersion of which is countercyclical. Since a simple frictionless environment predicts the dispersions of all decision variables of a firm to comove over the cycle, finding procyclical investment rate dispersion is interesting and suggestive of an important friction.

We argue that a friction in the capital adjustment technology, nonconvex costs of capital adjustment, is likely behind this new empirical regularity. First, we show that the strength of the comovement between investment rate dispersion and the cycle varies with empirical proxies for the importance of lumpiness in investment. Second, a calibrated heterogeneous-firm, lumpy investment general equilibrium model with shocks to the idiosyncratic productivity dispersion, which in the model constitute firm-level risk shocks, can closely match the observed cyclical comovement of the investment rate dispersion as well as the cyclical comovement of the dispersion of output and employment growth.

For the calibration of this model, we put front and center another well-established empirical fact about investment at the micro level: the long-run distribution of investment rates is positively skewed and has excess kurtosis (Caballero et al., 1995). While this previous research has highlighted the role of nonconvex adjustment costs in shaping the *long-run* distribution of investment rates, we add in this paper a fact about the *time-series* dynamics of the cross-sectional distribution of micro level investment and again relate it to nonconvex capital adjustment costs. More generally, we argue that a microfounded business cycle theory

[1]The literature has documented several related facts: Doms and Dunne (1998) show that the Herfindahl index of U.S. plant-level manufacturing investment is positively correlated with aggregate investment. Beaudry et al. (2001) show that cross-sectional investment dispersion in an unbalanced panel of roughly 1,000 U.K. manufacturing plants is negatively correlated with conditional inflation volatility. Eisfeldt and Rampini (2006) document that capital reallocation in U.S. Compustat data is procyclical.

[2]See Bachmann and Bayer (2013), Bloom et al. (2012), Doepke et al. (2005), Doepke and Weber (2006), Gourio (2008), Higson et al. (2002, 2004) and Kehrig (2011) for output and/or productivity, Berger and Vavra (2011) for prices, and Bachmann et al. (2013b) for business forecasts.

can - as we show - successfully speak to the dynamics of more than just the cross-sectional means of the distributions that underlie macroeconomic aggregates. We view this paper as a step toward such a research program for the firm sector.[3]

The procyclicality of the firm-level investment rate dispersion also has implications for the aggregate business cycle more generally. Arellano et al. (2012), Bloom et al. (2012), Christiano et al. (2010), Chugh (2012), Gilchrist et al. (2010), Narita (2011), Panousi and Papanikolaou (2012), Schaal (2011), and Vavra (2012) are examples of recent papers that have studied the business cycle implications of a time-varying dispersion of firm-specific variables, often interpreted as and used to calibrate shocks to firm risk, propagated through various frictions: wait-and-see effects from capital adjustment frictions, financial frictions, search frictions in the labor market, nominal rigidities and agency problems. We use the recent literature on risk shocks that are propagated through real options effects as an example to show that the procyclicality of investment dispersion helps researchers to gauge the power of risk shocks to generate or alter business cycle fluctuations. To be precise, our model in the spirit of Bloom (2009) and Bloom et al. (2012) matches the comovement of the investment rate dispersion with the business cycle only with fluctuations in the cross-sectional productivity growth dispersion that are in line with direct empirical evidence but lower than what the previous literature has advocated. More generally, our new fact imposes a natural overidentifying restriction on any model that uses shocks to firm-level risk that operate through an investment channel. In a world of noisy micro data, where especially the direct measurement of firm-level productivity is difficult and invariably assumption-laden, matching the cyclical behavior of not just the productivity growth dispersion but also the cyclical behavior of the dispersion of outcome variables, e.g. investment, is a challenge that models with idiosyncratic risk shocks should meet.

Our primary data source is the Deutsche Bundesbank balance-sheet database of German firms, USTAN. Hence, the data frequency we observe is annual. This database includes detailed accounting data that allow us to measure a firm's value-added, its stock of capital and its revenue productivity. Another strength of this database is that it covers virtually the entire nonfinancial private business sector of the German economy, unlike, e.g., the U.S. Annual Survey of Manufacturing (ASM), and it includes information from many non-traded medium size companies, unlike COMPUSTAT. This broad coverage permits sample splits that help us correlate the procyclicality of the investment rate dispersion with empirical proxies for the importance of lumpiness in investment, namely, industry or firm size. This makes USTAN uniquely suitable for our purposes. We show nevertheless that the investment rate dispersion is also procyclical in U.S. (COMPUSTAT) and UK (Cambridge DTI) data.

[3]Castaneda et al. (1998) have done this for the household side, documenting and explaining the business cycle dynamics of the U.S. income distribution.

Table 1: Cyclicality of Cross-Sectional Moments

Correlation with Cycle			
Cross-Sectional Standard Deviation of ...		Fraction of ...	
Investment rates	0.45**		
Output growth	-0.45*	Adjusters	0.73***
Employment growth	-0.50**	Spike adjusters	0.61***
Invest. rates cond. on spike adj.	-0.55***		
Productivity growth	-0.47**		

Notes: The left panel refers to the correlation with the cycle of the cross-sectional standard deviations, linearly detrended, of the investment rate, the log-change of real gross value-added, the net employment change rate, the investment rate conditional on its absolute value exceeding 20% (spike adjustment), and the log-change of Solow residuals, all at the firm level. Data are from the Bundesbank's USTAN database. We removed firm fixed and 2-digit industry-year effects from each variable. The right panel refers to the correlation with the cycle of the fraction, linearly detrended, of firms exhibiting an investment rate exceeding 1% (adjusters) and 20% (spike adjusters) in absolute value. The cyclical indicator is the HP(100)-filtered aggregate real gross value-added in the nonfinancial private business sector, computed from German VGR (*Volkswirtschaftliche Gesamtrechnungen*) data. ***,** ,* indicate significance at the 1%, 5%, and 10% level, resulting from an overlapping block bootstrap of four-year windows with 10,000 replications.

Table 1 summarizes our main empirical findings based on USTAN. The firm-level investment rate dispersion is procyclical. In contrast, the dispersions of output growth, employment growth, productivity growth and investment rates conditional on an investment spike, which we define as an investment rate larger than 20% in absolute value, are countercyclical. Finally, measures of the extensive margin of investment, namely, the fraction of firms with an investment rate larger than 1% in absolute value and the fraction of spike adjusters are significantly procyclical.

How do nonconvex capital adjustment costs and a countercyclical idiosyncratic productivity shock dispersion interact to generate what we find in Table 1? Even abstracting from risk shocks, nonconvex capital adjustment costs lead to two-step investment rules at the firm level. Firms first choose whether to adjust or not (extensive margin) and second, conditional on adjustment, they decide by how much to adjust (intensive margin). The cross-sectional investment dispersion will in general be a complicated nonlinear function of both steps, but, as we will show, it is the extensive margin choice that drives the procyclicality in investment rate dispersions. The difference in the sign of the cyclicalities of investment dispersion and investment dispersion conditional on an investment spike in Table 1 attests to this fact.

To fix ideas, approximate an investment distribution where most investment activity is concentrated in large lumps by assuming that firms can only decide whether to have an investment spike or not, i.e., to increase their capital stock by a given and large percentage or let it depreciate. Under this assumption, both the cross-sectional average and the dispersion of investment are solely determined by how many firms adjust. Aggregate investment is increasing in the fraction of firms exhibiting an investment spike, which is why the extensive margin measures in Table 1 are procyclical. And so is the dispersion of investment if less than half of the firms exhibit such a spike. Hence, as long as aggregate investment is procyclical and driven mainly by spike investments, the investment rate dispersion is also procyclical.

Adding a countercyclical productivity shock dispersion lets three additional effects come into play. With a decrease in the dispersion of productivity shocks, on the one hand firms will tend to adjust less often, as they simply move more slowly over their adjustment triggers. If low dispersion is concentrated in booms, then this *volatility* effect will tend to counteract the procyclicality of the extensive margin and the dispersion of investment. On the other hand, to the extent that a decrease in dispersion also constitutes a decrease in firm-level risk, firms see a decline in the option value of waiting, narrow their adjustment triggers, and thus adjust their capital stock more frequently. With a countercyclical productivity shock dispersion, this *real options* effect will tend to strengthen the procyclicality of the extensive margin and the dispersion of investment. Finally, there is the *intensive margin* effect that goes in the same direction as the volatility effect. Conditional on adjustment, costs are sunk and firms behave similarly to a frictionless setup. Hence conditional on adjustment, the distribution of investment rates follows the distribution of shocks, i.e., the conditional investment rate dispersion is negatively correlated with the cycle. The same argument holds for all firm-level decision variables that are not subject to large fixed costs of adjustment.

After expounding on the empirical findings in Section I, we ask in Sections II-IV whether the qualitative intuition described here holds up in a fully specified heterogeneous-firm real business cycle model that features both the extensive and the intensive margins of capital adjustment as well as shocks to aggregate productivity and countercyclical dispersion of firm-level productivity shocks. We find that indeed such a model, whose adjustment costs are calibrated to match the skewness and kurtosis of the long-run investment rate distribution, can quantitatively match the empirical procyclicality of the investment rate dispersion. We also show that small risk shocks to firm-level productivity are necessary for this result, in that the volatility and the intensive margin effect must be sufficiently yet not too strong to be quantitatively consistent with the procyclicality of the dispersion of investment rates.

I The Facts

We showed in the Introduction that the dispersion of firm-level investment rates is procyclical despite the dispersion of productivity, output and employment growth being countercyclical. In this section, we first fill in some information about our primary data source (USTAN), which is complemented by a more detailed description in Appendix A. We then link the procyclicality of investment dispersion to proxies for the lumpiness of investment, i.e., we identify firms where investment lumpiness should be more prevalent and show that the investment rate dispersion is more procyclical for those firms. We then show that investment rate dispersions are also procyclical in the U.S. and the UK and conclude with a summary of the robustness checks that we present in detail in Appendix B.

I.A Data and Sample Selection

Our primary data source is the Deutsche Bundesbank balance-sheet database of German firms, USTAN. USTAN is an annual private-sector, firm-level data set that allows us to make use of 26 years of data (1973-1998), with cross-sections that have, on average, over 30,000 firms per year. For the U.S. and UK evidence we use the COMPUSTAT sample from 1970-1994 and the Cambridge DTI database for 1978-1990, respectively. After the data undergo a similar data treatment to that for USTAN, the former covers roughly 2,150 firms per year and the latter only 850.[4] As usual, we compute economic capital stocks by a perpetual inventory method from balance-sheet data (see Appendix A.4 for details). We exploit the fact that all three data sets provide information for capital disaggregated into structures and equipment. This makes our measure of the economic capital stock robust to heterogeneity in capital portfolios, when, for instance, some firms have a larger fraction of their capital invested in structures than others.[5] The size of USTAN and its broad coverage in terms of ownership, firm size and industry allow us to study the cyclicality of investment rate dispersions in various sample splits that are meant to capture putative differences in the relevance of investment lumpiness.

USTAN is a byproduct of the Bundesbank's rediscounting activities. The Bundesbank had to assess the creditworthiness of all parties backing promissory notes or bills of exchange put up for rediscounting (i.e., as collateral for overnight lending). It implemented this regulation by requiring balance-sheet data of all parties involved, and the data were then

[4] When cross-sectional dispersions are concerned, Davis et al. (2006) show that studying only publicly traded firms (COMPUSTAT) can lead to misleading conclusions, which is why we view USTAN as very suitable for our purposes.

[5] The heterogeneity of capital portfolios is also the reason why we restrict the COMPUSTAT sample to 1970-1994. After 1994 no separate information for structures and equipment is available there.

collected and archived (see Appendix A.1, Stoess (2001) and von Kalckreuth (2003) for details).

From the original USTAN data, we select firms that report information on payroll, gross value-added and capital stocks, and for which we have at least five observations in first differences. Moreover, we drop outliers and observations from East German firms to avoid a break in the series in 1990 (see Appendix A.2 for details). This leaves us with a sample of 854,105 firm-year observations from 72,853 different firms, i.e., on average a firm is observed in the sample for 11.7 years. The average number of firms in the cross-section of any given year is 32,850. The resulting sample covers roughly 70% of the West German real gross value-added in the nonfinancial private business sector and 50% of its employment.[6]

Throughout the paper, we follow Bloom (2009) and define the investment rates of a firm j at time t as $i_{j,t} = \frac{I_{j,t}}{0.5(k_{j,t}+k_{j,t+1})}$.[7] Investment rates exhibit cross-sectional skewness (2.19) and kurtosis (20.04), which, according to Caballero et al. (1995), is a strong indication of investment lumpiness at the firm level. Table 20 in Appendix A.6 shows that the investment histograms for USTAN and the manufacturing sector in USTAN look very similar to the investment histogram for the U.S. in Cooper and Haltiwanger (2006).

For firm-level employment growth rates we use the symmetric adjustment rate definition proposed in Davis et al. (1996), $\frac{\Delta n_{j,t}}{0.5*(n_{j,t-1}+n_{j,t})}$. Firm-level productivity and output growth rates are simple log-differences of, respectively, Solow residuals and real gross value-added, for which we deflate the balance-sheet item nominal gross value-added by the price index for gross value-added from German national accounting data (*VGR*).[8] To focus on idiosyncratic changes that do not capture differences in industry-specific responses to aggregate shocks or ex-ante firm heterogeneity, firm fixed and industry-year effects are removed from investment rates, as well as from the employment, output and productivity growth rates.

[6] Throughout we will refer to Agriculture, Mining and Energy, Manufacturing, Construction, Trade and Transportation and Communication collectively as the *nonfinancial private business sector (NFPBS)*.

[7] Spike adjusters are defined relative to this investment rate definition, i.e., $i_{j,t} < -20\%$ or $i_{j,t} > 20\%$. Strictly speaking, the literature, e.g., Cooper and Haltiwanger (2006) and Gourio and Kashyap (2007), has used the 20% threshold with respect to $\frac{I_{j,t}}{k_{j,t}}$, but we show in Appendix A.6 that the investment rate histograms in USTAN look similar for either definition of the investment rate.

[8] To compute firm-level Solow residuals, we start, in accordance with the model in Section II, from a firm-level Cobb-Douglas production function: $y_{j,t} = \exp(z_t + \epsilon_{j,t})k_{j,t}^{\theta}n_{j,t}^{\nu}$, where ϵ is firm-specific and z aggregate log productivity. We assume that labor input n is immediately productive, whereas capital k is predetermined and inherited from the last period. This difference is reflected in the different timing convention in the definitions of the investment and employment adjustment rates. We estimate the output elasticities of the production factors, ν and θ, as median factor expenditure shares over gross value-added within each industry. 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. See for details Appendices A.7 and A.8.

I.B Procyclicality of Investment Rate Dispersions and Proxies for Lumpy Investment

The literature has typically focused on the manufacturing sector to find evidence for nonconvex adjustment technologies (see Doms and Dunne (1998), Caballero et al. (1995), Caballero and Engel (1999), Cooper and Haltiwanger (2006), and Gourio and Kashyap (2007)), and for good reason: manufacturing is where heavy-duty machinery needs to be installed and large production halls need to be built, which may lead to disruptions in the production process. Indeed, Table 2 shows that in manufacturing and also in construction, the correlation of the investment rate dispersion with the industry cycle is particularly strong and statistically significant.

Table 2: Cyclicality of Cross-Sectional Investment Rate Dispersion - By One-Digit Industries

Correlation of Real Gross Value-Added with $std(i_{j,t})$					
		All Industries	0.45**		
Primary Sector		Secondary Sector		Tertiary Sector	
Agriculture	−0.19	Manufacturing	0.48***	Trade	0.21
Mining & Energy	0.04	Construction	0.44**	Transport & Communication	0.40**

Notes: See notes to Table 1. The table displays correlation coefficients with the cyclical component of aggregate real gross value-added of the nonfinancial private business sector in the first row, thereafter with the real gross value-added of the corresponding one-digit industry.

Another dimension that is likely correlated with the relevance of adjustment frictions is firm size. Larger firms may partially outgrow fixed adjustment costs or can smooth the effects of nonconvex capital adjustment costs and the extensive margin over several production units. In Table 3 we see that the procyclicality of the investment rate dispersion is falling in firm size. The very large firms, in contrast to the small ones, have an almost acyclical investment dispersion. This distinction is statistically significant in the sense that if size is measured in terms of employment or value-added, neither the point estimate for the smallest size class lies in the $[5\%, 95\%]-$band of the largest size class nor vice versa.

Another way to see how the extensive margin of investment and the dispersion of firm-level productivity growth interact to generate a procyclical investment rate dispersion is to

Table 3: Cyclicality of Cross-Sectional Investment Dispersion - By Firm Size

Size Class / Criterion	Employment	Value-Added	Capital
Smallest 25%	0.58***	0.60***	0.39**
25% to 50%	0.46**	0.47**	0.42**
50% to 75%	0.37*	0.33	0.39*
Largest 25%	0.19	0.22	0.40**
Largest 5%	0.05	0.05	0.18

Notes: See notes to Tables 1 and 2. Just as for the aggregate numbers in Table 1, we use the cyclical component of the aggregate output of the private nonfinancial business sector as the cyclical indicator.

exploit the sectoral information in our data. We disaggregate the data by years and 14 two-digit industries and then regress in a pooled OLS regression the dispersion of investment rates on the dispersion of firm-level productivity growth and the fraction of (spike) adjusters in a given industry and year. Table 4 shows that the larger the dispersion of shocks is and the more frequent investment activities in an industry-year are, the larger is the dispersion of investment rates. Importantly, conditional on the fraction of (spike) adjusters, i.e., the extensive margin of investment, investment dispersion and the dispersion of firm-level productivity growth comove positively. Since the productivity growth dispersion is countercyclical, the procyclicality of the dispersion of investment rates has to be driven by the procyclicality of the cross-sectional investment frequency. Put differently, the fluctuations in firm-level risk cannot be too large so as to undo the extensive margin effect.

Table 4: Evidence from Disaggregation by Two-Digit Industry and Year

Regression of $std(i_{j,t})$ on ...			
(a)		(b)	
Fraction of Adjusters	.23***	Fraction of Spike Adjusters	.28***
$std(\Delta\epsilon_{j,t})$	.37***	$std(\Delta\epsilon_{j,t})$	.20***

Notes: The table displays the estimated coefficients of a pooled OLS regression of the cross-sectional investment rate dispersion for each two-digit industry and year on the fraction of (spike) adjusters in that industry and year and the dispersion of idiosyncratic productivity shocks. All data have been linearly detrended at the industry level. See Table 15 in Appendix A.3 for more details on the two-digit industries in USTAN.

Table 5: Evidence from U.S. and UK data

Correlation of $std(i_{j,t})$ with HP(100)-Y			
U.S. (COMPUSTAT)	0.60^{***}	UK (Cambridge DTI)	0.45^{**}

Notes: See notes to Table 1. The cyclical indicator HP(100)-Y refers to the cyclical component of aggregate real gross value-added of the nonfinancial private business sector from NIPA data for the U.S. and, because the DTI has a high fraction of manufacturing firms, to the manufacturing production index for the UK.

Collectively, the evidence in this section at least suggests that nonconvex capital adjustment costs play a role in explaining procyclical investment dispersion.

To conclude, we show that the procyclicality of the investment rate dispersion is robust across different data sets. We use the U.S. COMPUSTAT data and the Cambridge DTI data to document procyclical investment dispersions for the U.S. and the UK. Wherever possible, we treat the data analogously to the way we treated the USTAN data. As cross-sections are much smaller in the U.S. and UK data, the sample splits that we do for USTAN are not possible. Again we select firms from the nonfinancial private business sector and correlate their investment rate dispersions with the corresponding aggregate gross real value-added. Table 5 displays the results.

I.C Robustness

How robust is the procyclicality of the investment rate dispersion? Potential issues for robustness are: our measure of the cycle, our measure of dispersion, and the representativeness of the sample both cross-sectionally and in the time-series dimension. We establish the robustness of our findings in all these dimensions in Appendix B in Tables 21, 22, and 23.

We check robustness to alternative measures of the cycle (HP-filter parameters, aggregate variables indicating the cycle, detrending of the dispersion series, etc.), to looking at dynamic correlations, to excluding the two most extreme investment dispersion years, and to excluding the post-reunification period. Moreover, we check robustness with respect to sample composition. We study a sample where we include firms only three years after they entered the sample the first time, to make sure our result is not driven by firm entry. We also look at firms that are stable in the sample so that cyclicality is not driven by the systematic exit of firms. Finally, we try different criteria for excluding outliers, using a non-centralized definition of investment rates, alternative ways of treating movements in the price of capital goods, and replacing the standard deviation as our measure of dispersion with the interquartile range.

II The Model

We follow closely Khan and Thomas (2008) and Bachmann et al. (2013a). The main departure from both papers is the introduction of a second aggregate shock, namely, time-varying idiosyncratic productivity risk.

II.A 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 labor (n) and its pre-determined capital stock (k), according to the following Cobb-Douglas decreasing-returns-to-scale production function:

$$y = \exp(z + \epsilon)k^{\theta}n^{\nu};\ \text{with } \theta, \nu > 0 \text{ and } \theta + \nu < 1, \tag{1}$$

where z and ϵ are log aggregate and log idiosyncratic revenue productivity, respectively.

The idiosyncratic log productivity process is first-order Markov with autocorrelation ρ_ϵ and time-varying conditional standard deviation, $\sigma(\epsilon)$. We assume two exogenous aggregate states (z, s), which evolve jointly according to an unrestricted VAR(1) process, with normal innovations u that have zero mean and covariance Ω:[9]

$$\begin{pmatrix} z' \\ s' \end{pmatrix} = \varrho_A \begin{pmatrix} z \\ s \end{pmatrix} + u,\ cov(u) = \Omega. \tag{2}$$

In line with the production function (1) z is the trend deviation of the natural logarithm of aggregate productivity, while s drives the dispersion of idiosyncratic productivity shocks, which is given by $\sigma(\epsilon) = \sigma_\sigma s + \bar{\sigma}(\epsilon)$, where $\bar{\sigma}(\epsilon)$ denotes the steady-state standard deviation of the innovations to idiosyncratic productivity, and σ_σ scales the size of the fluctuations in $\sigma(\epsilon)$. The shocks to the exogenous aggregate states, u, and idiosyncratic productivity shocks are independent. Idiosyncratic productivity shocks are independent across productive units. We do not impose any restrictions on Ω or $\varrho_A \in \mathbb{R}^{2\times 2}$.

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

We model employment as freely adjustable but assume that capital adjustment is costly. 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.

[9] Curdia and Reis (2011) recently pointed to using correlated shocks for understanding business cycles.

Upon investment the firm incurs fixed costs $\omega\xi$, where ω is the current real wage. Capital depreciates at rate δ. We denote the firms' distribution over (ϵ, k) by μ. Thus, $\big(z, s, \mu\big)$ constitutes the current aggregate state and μ evolves according to the law of motion $\mu' = \Gamma\big(z, s, \mu\big)$, which firms take as given.

Next we describe the dynamic programming problem of a firm. Following Khan and Thomas (2008), we state this problem in terms of utils of the representative household (rather than physical units) and denote the marginal utility of consumption by $p = p\big(z, s, \mu\big)$. This is the kernel that firms use to price output streams. Also, given the i.i.d. nature of adjustment costs, continuation values can be expressed without future adjustment costs.

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

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

With this notation the dynamic programming problem becomes:

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

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:

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

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

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

$$V_{\text{adj}} = -\big(\gamma k' - (1 - \delta)k\big)p(z, s, \mu) + \beta \text{E}[V^0(\epsilon', k'; z', s'), \mu')], \tag{5d}$$

where both expectation operators average over the next period's realizations of the aggregate and idiosyncratic shocks, conditional on this period's values. The discount factor, β, reflects the time preferences of the representative household.

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

II.B Households

We assume a continuum of identical households. They have a standard felicity function in consumption and labor:

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

where C denotes consumption and N^h the households' labor supply. Households maximize the expected present discounted value of the above felicity function, yielding:

$$p\big(z, s, \mu\big) \equiv U_C(C, N^h) = \frac{1}{C\big(z, s, \mu\big)}, \quad \text{and} \qquad \omega\big(z, s, \mu\big) = -\frac{U_N(C,N^h)}{p\big(z,s,\mu\big)} = \frac{A}{p\big(z,s,\mu\big)}. \tag{7}$$

II.C 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, $V^1(\epsilon, k, \xi; z, s, \mu)$ solves (4) and the corresponding policy functions are $N(\epsilon, k; z, s, \mu)$ and $K(\epsilon, k, \xi; z, s, \mu)$.

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

3. *Commodity market clearing*:

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

4. *Labor market clearing*:

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

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

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

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

II.D Solution

It is well-known that (4) 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 find that adding the unconditional cross-sectional standard deviation of the natural logarithm of the level of idiosyncratic productivity, $std(\epsilon)$, not only improves the fit of the Krusell-Smith rules but it also matters for our economic results (see Appendix C for details).[10] This is owing to the now time-varying nature of the distribution of idiosyncratic productivity. We surmise that this is a general insight and that simple Krusell-Smith rules are likely inappropriate in models with firm-level risk shocks. In the same vein, we approximate the equilibrium pricing function by a log-linear rule:

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

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

Given (7), we do not have to specify a rule for the real wage. We posit the rules (8a)–(8b) and check that in equilibrium they yield a good fit to the actual law of motion.[11]

Substituting $\bar{k}$ and $std(\epsilon)$ for μ and using (8a)–(8b), (4) becomes a computable dynamic programming problem with corresponding policy functions $N = N\big(\epsilon, k; z, s, \bar{k}, std(\epsilon)\big)$ and $K = K\big(\epsilon, k, \xi; z, s, \bar{k}, std(\epsilon)\big)$. We solve this problem by value function iteration on V^0 and apply multivariate spline techniques that allow for a continuous choice of capital when the firm adjusts.

With these policy functions, we can 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. 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 $a_k\big(z,s'\big)$ to $c_p\big(z,s'\big)$ are sufficiently close to the previous ones.

III Calibration

The model frequency is annual, which corresponds to the data frequency in USTAN. Some model parameters are directly calculated or estimated from VGR and/or USTAN data (such

[10] For a similar insight, see Zhang (2005).

[11] $std(\epsilon)$ is a function of $std(\epsilon_{-1})$ and $\sigma(\epsilon)$. Further details on the numerical solution method and on the quality of the approximation are available in Appendix C.

Table 6: Model Parameters: Baseline Calibration

Parameter		Value	Calibrated from / to
discount factor	β	0.97	real interest rate on corporate bonds: 4.6%
disutility of labor	A	2	average time spent at work: 1/3
depreciation rate	δ	0.094	VGR Data: depreciation rate
long-run growth factor	γ	1.014	VGR Data: aggregate investment rate
output elasticity of labor	ν	0.5565	USTAN: value-added share
output elasticity of capital	θ	0.2075	USTAN: value-added share
time-average idiosyncratic risk	$\bar{\sigma}(\epsilon)$	0.0905	USTAN: Solow residual growth dispersion
autocorrelation of idiosyncratic produc'ty	ρ_ϵ	0.9675	USTAN: Solow residual growth
Joint process of			
aggregate productivity and	ϱ_A, Ω	see below	VGR Data: Solow residuals
volatility of idiosyncratic risk			USTAN: Solow residual growth dispersion
Scaling of risk fluctuations	σ_σ	1	Normalization
adjustment cost parameter	$\bar{\xi}$	0.2	USTAN: investment rate skewness and kurtosis

as the depreciation rate, the output elasticities and the parameters of the aggregate and idiosyncratic driving processes). The remaining parameters are jointly calibrated to match the real interest rate, the average time spent at work and the aggregate investment rate in the German nonfinancial private business sector as well as the skewness and kurtosis of the firm-level investment rate in the USTAN data. Table 6 gives an overview of the model parameters, their values and the data sources. Importantly, the cyclicality of the investment rate dispersion is not targeted, nor are the cyclicalities of the fraction of spike adjusters, the output growth dispersion and the employment growth dispersion.

III.A Technology and Preference Parameters

We take the depreciation rate, $\delta = 0.094$, directly from German national accounting (VGR) data for the nonfinancial private business sector. Given this depreciation rate, $\gamma = 1.014$ matches the time-average aggregate investment rate in the nonfinancial private business sector: 0.108.[12] Since the average real interest rate in Germany over the period 1973-1998

[12] $\gamma = 1.014$ is also consistent with German long-run growth rates.

was 4.6%, we obtain for the discount factor $\beta = 0.97$.[13] 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 database. Our model simulations show that value-added shares are good estimators of the output elasticities even in the presence of fixed adjustment costs and outperform other estimators of the production function; see Appendix A.7 for details.

III.B Idiosyncratic Shocks

We calibrate the standard deviation of idiosyncratic productivity shocks to $\bar{\sigma}(\epsilon) = 0.0905$, which we obtain from measured firm-level Solow residual growth in USTAN cleansed of measurement error. We set $\rho_\epsilon = 0.9675$, which we estimate again from measured productivity in USTAN (details are available in Appendix A.8). This process is discretized on a 19−state-grid, using Tauchen's (1986) procedure with a mixture of two Gaussian normals to capture above-Gaussian kurtosis - 4.4480 on average - in idiosyncratic productivity shocks (details are available in Appendix C). 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 $\sigma(\epsilon)$.

III.C Aggregate Shocks

To calibrate the parameters of the two-state aggregate shock process, we estimate a bivariate, unrestricted VAR with the linearly detrended natural logarithm of the aggregate Solow residual[14] and the linearly detrended $\sigma(\epsilon)$-process, i.e., the process for the standard deviation of the innovations to idiosyncratic productivity, from the USTAN data, where we normalize $\sigma_\sigma = 1$ for the baseline model. The estimated parameters of this VAR are:[15]

$$\varrho_A = \begin{pmatrix} 0.2791 & -1.3439^{**} \\ 0.1059^{**} & 0.8072^{***} \end{pmatrix} \quad \Omega = \begin{pmatrix} 0.0115^{***} & -0.5459^{***} \\ -0.5459^{***} & 0.0036^{***} \end{pmatrix} \tag{9}$$

Importantly, both the negative contemporaneous correlation and the negative coefficient of firm risk on future TFP are significant. This process is discretized on a $[5 \times 5]$−state grid, using a bivariate analog of Tauchen's procedure.

[13]We calculate the real interest rate as the return on corporate bonds minus the ex-post inflation rate.

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

[15]With a slight abuse of notation, but for the sake of readability, Ω has standard deviations on the main diagonal and correlations on the off diagonal. ** and *** denote the usual significance levels. Notice the high persistence in the $\sigma(\epsilon)$-process. Today's idiosyncratic productivity shock dispersion has strong predictive power about tomorrow's idiosyncratic productivity shock dispersion and thus reflects risk shock.

Table 7: CALIBRATION OF ADJUSTMENT COSTS - $\bar{\xi}$

$\bar{\xi}$	Skewness	Kurtosis	$\Psi(\bar{\xi})$
0.00	-0.01	3.61	18.93
0.01	1.13	5.83	8.98
0.10	2.45	10.88	2.90
0.20 (BL)	2.90	13.05	2.66
0.30	3.17	14.46	3.01
0.50	3.51	16.40	4.01
1.00	3.96	19.37	6.51

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.

III.D Adjustment Costs

The distribution of firm-level investment rates exhibits both substantial positive skewness – 2.1920 – as well as kurtosis – 20.0355. Caballero et al. (1995) document a similar fact for U.S. manufacturing plants. They also argue that nonconvex capital adjustment costs are an important ingredient for explaining 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 $\bar{\xi}$.

Since, as a practical matter, the adjustment cost parameter, $\bar{\xi}$, hardly impacts long-run variables, such as the average real interest rate, the average time spent at work or the average aggregate investment rate in the model, it is convenient to proceed as follows: given the following set of parameters $\{\beta, \delta, \gamma, A, \nu, \theta, \bar{\sigma}(\epsilon), \rho_\epsilon, \varrho_A, \Omega, \sigma_\sigma\}$, we find $\bar{\xi}$ by minimizing 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 we estimate skewness and kurtosis, we weigh both with the inverse of their time-series standard deviation. Table 7 shows that $\bar{\xi}$ is indeed identified in this calibration strategy, as cross-sectional skewness and kurtosis of the firm-level investment rates are monotonically increasing in $\bar{\xi}$. The minimum of Ψ is achieved at $\bar{\xi} = 0.2$, our baseline.

Before investigating what this calibration implies for the correlation of the dispersion of various firm-level activity variables with aggregate economic activity, it is also instructive to see what it entails for other statistics related to nonconvex capital adjustment costs, which have not been targeted here. The average adjustment costs conditional on adjustment amount to roughly 11% as a fraction of annual firm-level output, which is at the lower end of estimates from the U.S. (see Bloom (2009), Table IV, for an overview). Moreover, the baseline model implies a fraction of spike adjusters of 11.3%, i.e., firms with an investment rate that is larger than 20% in absolute value, which is well in line with the 13.4% in the USTAN data. Finally, our model produces basically zero autocorrelation of firm-level investment rates (-0.05), compared to -0.03 in the USTAN data, a typical feature of lumpy investment at the micro-level.[16]

IV Results

IV.A Baseline Results

Can a general equilibrium model with standard aggregate productivity shocks, persistent idiosyncratic productivity shocks, countercyclical aggregate shocks to their dispersion and fixed capital adjustment costs, calibrated to the long-run non-Gaussianity of the investment rate distribution, reproduce the non-targeted cyclicality of the cross-sectional dispersion of firm-level investment rates, output growth and employment growth and the cyclicality of the extensive and intensive margins of spike investment? Table 8 says yes.

That the model can closely match the cyclicality of the investment rate dispersion is an example of the larger premise of this paper: that cross-sectional dynamics are an important aspect of the data that heterogeneous firm models should address. With quantitatively realistic shocks to the dispersion of firm-level Solow residuals, it is one level of adjustment costs that makes the model jointly consistent with the (targeted) time-average skewness and kurtosis of the investment rate distribution – two statistics closely related to the relevance of nonconvexities at the micro-level – and the time-series correlation between the cross-sectional standard deviation of investment rates and output (not targeted).

Let us reiterate the intuition for the results in Table 8. In a world with aggregate productivity shocks only, and firms merely having to make a decision about whether to realize an investment spike or not, both the fraction of spike adjusters and the dispersion of investment rates are procyclical, as long as spike adjustment is sufficiently infrequent:

[16]Cooper and Haltiwanger (2006) found an autocorrelation of plant-level investment rates of 0.06 in the U.S. LRD data and use this number as one of the characteristic moments in their GMM procedure to identify (non-)convex capital adjustment costs.

Table 8: Cyclicality of Cross-Sectional Dispersions and the Margins of Investment - Baseline Model

Correlation with the cycle		
Cross-sectional Moment	Model	Data
$std(i_{j,t})$	0.53	0.45
Fraction of spike adjusters	0.63	0.61
$std(\Delta \log y_{j,t})$	-0.36	-0.45
$std(\frac{\Delta n_{j,t}}{0.5*(n_{j,t-1}+n_{j,t})})$	-0.38	-0.50
$std(i_{j,t})$ conditional on spike adjustment	-0.74	-0.55

Notes: See notes to Table 1. The table displays correlation coefficients with HP(100)-filtered aggregate output. The column 'Model' refers to the correlation coefficients from a simulation of the baseline model.

aggregate gross investment is given by $\lambda\kappa$, where λ is the fraction of spike adjusters and κ the size of the investment spike. Investment dispersion in such an economy would be $\lambda(1-\lambda)\kappa^2$, which is increasing in the fraction of spike adjusters as long as $\lambda \leq 0.5$, which is the case in the data. This intuition carries over into a more realistic economy where there is an intensive margin of investment that is nevertheless of secondary (to the extensive margin) importance for aggregate investment, again as in the data.

With a countercyclical productivity shock dispersion, the real options effect of higher firm-level risk will strengthen the procyclicality of the investment rate dispersion, as lower risk increases the fraction of spike adjusters. In contrast, both the volatility effect – lower risk decreases the probability that firms hit their adjustment triggers – and the intensive margin effect counteract the procyclicality of the investment rate dispersion. The volatility and the intensive margin effects appear to dominate the real-options effect, which is suggested by the fact that for both the investment rate dispersion and for the fraction of spike adjusters we see positive correlations with the cycle that are substantially smaller than those same correlations in a model without risk shocks (see Table 9 below).

Going back to Table 8, the countercyclicality of the output and employment growth dispersion follows directly from the countercyclical productivity shock dispersion, as both output and employment growth are simple functions of productivity growth and the growth of the capital stock of a firm, essentially the investment rate one period ago. Ignoring the cross-sectional covariances and their time-series behavior, the impact of which is small with idiosyncratic productivity being close to a random walk, the cyclicality of the output and employment growth dispersion is then simply a function of the cyclicality of the dispersions

Table 9: Adjustment Costs and the Cyclicality of Cross-Sectional Variables

	with 2nd moment shocks			w/o 2nd moment shocks		
$\bar{\xi}$	$std(i_{j,t})$	$std(\Delta \log y_{j,t})$	Fraction of spike adjusters	$std(i_{j,t})$	$std(\Delta \log y_{j,t})$	Fraction of spike adjusters
0	-0.41	-0.46	-0.41	-	-	-0.01
0.01	-0.23	-0.44	0.17	0.85	0.11	0.69
0.1	0.23	-0.39	0.34	0.88	0.20	0.88
0.2	0.53	-0.36	0.63	0.88	0.21	0.91
0.3	0.68	-0.35	0.76	0.89	0.23	0.91
0.5	0.82	-0.32	0.86	0.89	0.18	0.92
1	0.91	-0.29	0.92	0.90	0.17	0.92

Notes: See notes to Table 8. 'with 2nd moment shocks' refers to a simulation with aggregate productivity shocks and shocks to the dispersion of firm-level Solow residuals, as specified in equation (9). 'w/o 2nd moment shocks' refers to a simulation with only aggregate productivity shocks, where $\varrho_A = 0.5223$ and $\Omega = 0.0121$. Note that in this case, with $\bar{\xi} = 0$, $std(i_{j,t})$ and $std(\Delta \log y_{j,t})$ are constant, which means that their correlation coefficients with output are not defined.

of productivity shocks and investment rates. Since as usual with Cobb-Douglas production functions the coefficient on factor growth is an order of magnitude smaller than that on productivity growth, the cyclicality of the productivity shock dispersion will dominate.

To further investigate the mechanism behind the procyclical investment rate dispersion, Table 9 displays the cyclicality of the investment rate and output growth dispersions as well as the cyclicality of the fraction of spike adjusters that the model generates for various levels of adjustment costs, both with and without second moment shocks.[17]

Two findings are important:

1. The right panel of Table 9 shows that, without second moment shocks, neither the procyclicality of the investment dispersion, the procyclicality of the fraction of lumpy adjusters, nor the countercyclicality of the output growth dispersion can be quantitatively replicated. Already a very small nonconvex capital adjustment cost factor generates procyclical investment dispersion. The model overshoots the number in the data considerably. Also, without countercyclical second moment shocks, the dispersion of value-added growth is slightly procylical. This follows immediately from the considerations above: without time-varying dispersion of productivity growth, the cyclicality

[17] The cyclicality of the employment growth dispersion, for space reasons not shown in Table 9, behaves similarly to that of the output growth dispersion.

of the dispersion of output growth is determined by the cyclicality of (lagged) investment rates. Hence, countercyclical second moment shocks play an important role in understanding cross-sectional firm dynamics.

2. In the presence of countercyclical second moment shocks (left panel), in the frictionless case, $\bar{\xi} = 0$, the dispersions of investment and output growth merely mirror the countercyclicality of the dispersion of the idiosyncratic driving force. Because it is more likely to observe an investment spike when the dispersion of shocks goes up (at $\bar{\xi} = 0$ the volatility and the intensive margin effect are the same), the fraction of spike adjusters is also countercyclical. Increasing the fixed adjustment costs leads to a gradual increase in the procyclicality of the investment rate dispersion and the procyclicality of the fraction of spike adjusters. The real options effect of countercyclical risk shocks becomes stronger; more firms invest in low-risk times, i.e., in booms. At $\bar{\xi} = 1$ this effect eventually starts to dominate the volatility effect, and the model with second moment shocks has a more procyclical investment dispersion than the one without.

Table 10 shows how our baseline results change with more volatile firm-level risk. To this end, we double and quadruple, respectively, the scaling parameter, σ_σ, in the definition of $\sigma(\epsilon)$, and re-estimate (2). This means that we double/quadruple the time-series coefficient of variation of firm-level risk. Given that the data suggest a correlated shock process between aggregate productivity and firm-level risk (see equation 9), it is important to keep the information structure on aggregate productivity the same and isolate the pure effect of time-varying firm-level risk. This is conveniently done by setting $\sigma_\sigma = 1, 2$, or 4.

Table 10: Volatility of Risk and Cyclicality of Investment Dispersion/Extensive Margin

	Volatility of $std(\Delta\epsilon_{j,t})$		
Cross-sectional Moment	Baseline $\sigma_\sigma = 1$	Double $\sigma_\sigma = 2$	Quadruple $\sigma_\sigma = 4$
$std(i_{j,t})$	0.53	-0.09	-0.37
$std(\Delta \log y_{j,t})$	-0.36	-0.42	-0.45
Fraction of spike adjusters	0.63	-0.10	-0.41

Notes: See notes to Table 8.

Increasing the time-series volatility of countercyclical firm-level risk leads, counterfactually, to acyclical or even countercyclical behavior of the investment rate dispersion. The table also shows that the dispersion of productivity growth affects the investment dispersion not only through the intensive margin of investment – conditional on adjustment, investment becomes more disperse when shocks are more disperse – but also through the extensive margin: with more volatile firm-level risk shocks, the volatility effect makes the fraction of spike adjusters countercyclical – more – not less – firms adjust when $\sigma(\epsilon)$ is high.

In summary, Table 10 shows that only our baseline specification with relatively small fluctuations in idiosyncratic productivity growth dispersions matches the cyclicality of the investment dispersion, the cyclicality of the output growth dispersion, and the cyclicality of the fraction of spike adjusters jointly. This means effectively that the procyclicality of the investment rate dispersion places a tight bound on the volatility and thus the importance of firm-level risk shocks. If we increase this volatility and make second moment shocks "more important," the procyclicality of the investment rate dispersion disappears.

Why is this important? In the Introduction we mentioned a growing literature that studies the aggregate consequences of shocks to firm-level productivity risk. To get quantitatively realistic predictions out of models with such shocks, researchers have to measure the size of these risk fluctuations. This measurement is invariably assumption-laden. Bachmann and Bayer (2013) demonstrate that, for instance, depending on whether only continuing firms are taken into account in the sample, the volatility of firm-level risk can vary substantially. In contrast, as we show in this paper, the mild procyclicality of the investment rate dispersion is robust to this issue. Moreover, in order to measure firm-level productivity risk, researchers have to go through all the steps that need to be taken to calculate investment rates and, in addition, estimate a production function.

In this paper we use a mid-range estimate for the volatility of firm-level risk: the time-series coefficient of variation of firm-level risk in our baseline calibration is 4.72%.[18] One approach to dealing with this parameter uncertainty is to report lower and upper bound scenarios. Our alternative approach is to use more cross-sectional information as an overidentifying restriction and investigate whether the models with risk shocks not only match the cyclical behavior of the dispersion of the driving force (firm-level revenue productivity growth) or merely one outcome variable (firm-level sales growth) but jointly the cyclical behavior of the dispersion of all major firm-level outcome variables. The next section will illustrate the implications of this alternative approach for the literature on risk shocks in environments with physical capital adjustment frictions.

[18]Bloom (2009) uses roughly a four times more volatile shock process for firm-level risk.

IV.B Aggregate Fluctuations

How does our baseline model perform in terms of standard aggregate time-series moments? Table 11 compares unconditional second moments of aggregate output, consumption, investment and employment between the model in our baseline calibration, a version of the model without second moment shocks and with frictionless capital adjustment (essentially the RBC model), and the data. As in the standard RBC model, output volatility is roughly matched, consumption and employment are somewhat too smooth, and aggregate investment is too volatile. The model falls short in terms of persistence and overpredicts the comovement of aggregates with output. Thus we have similar success or a lack thereof as in the standard RBC model in matching unconditional second moments. Our new focus is on the cross-sectional dynamics of the model and its aggregate implications, which is why the simple RBC model, despite its shortcomings, is the right point of departure for our analysis.

Table 11: Aggregate Fluctuations

		Y	C	I	N
volatility	BL-Model	2.19	0.82	10.22	1.52
	RBC-Model	2.02	0.76	9.64	1.45
	Data	2.30	1.79	4.37	1.80
persistence	BL-Model	0.29	0.55	0.22	0.20
	RBC-Model	0.25	0.56	0.18	0.18
	Data	0.48	0.67	0.42	0.61
correlation with Y	BL-Model	1.00	0.87	0.98	0.96
	RBC-Model	1.00	0.83	0.97	0.96
	Data	1.00	0.66	0.83	0.68

Notes: The table displays the percent standard deviations (volatility), autocorrelation (persistence), and correlation with aggregate output of HP(100)-filtered log aggregate output (Y), consumption (C), investment (I), and employment (N) of the model under the baseline calibration ('BL-Model'), a version of the baseline model without second moment shocks and with frictionless capital adjustment ('RBC-Model'), as well as German aggregate data from *VGR*.

More interesting are the implications of firm-level risk shocks in our model. To understand these implications, we perform two separate exercises. First, we investigate in our baseline model the effect of an increase in the volatility of firm-level risk – making it "more important" – on the volatility of aggregate output. We proceed exactly as described for the exercises in

Table 10, i.e., we vary σ_σ over $0, 1, 2, 4$. In the extreme case of $\sigma_\sigma = 0$, firm risk is constant over time, yet aggregate productivity is still the result of a bivariate stochastic process.

In the second exercise we, counterfactually, eliminate shocks to aggregate productivity in order to understand the impact on aggregate volatility of risk shocks in isolation. Of course, this yields not literally a variance decomposition, given the correlated nature of the aggregate shock process and the nonlinearity of the model, but we believe that we can thus nevertheless gauge the relevance of idiosyncratic risk shocks for aggregate fluctuations. Again, we vary σ_σ over $\sigma_\sigma = 0, 1, 2, 4$.

Table 12: Aggregate Output Volatilities and the Volatility of Firm-Level Risk Shocks

	Full Model		Risk-Only Model	
	Standard Dev. of Output	Decline in Variance of Output	Standard Dev. of Output	Variance Explained
$\sigma_\sigma = 0$	2.29	0%	-	-
$\sigma_\sigma = 1$ (BL)	2.19	9%	0.24	1%
$\sigma_\sigma = 2$	2.12	15%	0.46	4%
$\sigma_\sigma = 4$	2.08	17%	0.82	13%

Notes: The table displays in the left panel ('Full Model') the percent standard deviations of HP(100)-filtered log aggregate output (Y) in our model under the baseline calibration and for $\sigma_\sigma = 0, 2, 4$. It also displays the percentage decline in the variance of aggregate output for $\sigma_\sigma = 1, 2, 4$ relative to $\sigma_\sigma = 0$: $|\frac{var(Y)[\sigma_\sigma=1,2,4]-var(Y)[\sigma_\sigma=0]}{var(Y)[\sigma_\sigma=0]}|$. The right panel ('Risk-Only') displays the output volatilities for a version of the baseline model without any aggregate productivity shocks. It also displays the percentage of the variance of aggregate output in the data that is explained by the 'Risk-Only Model'-model: $\frac{var(Y)[\sigma_\sigma=1,2,4]}{var(Y)[Data]}$.

Table 12 shows the results from these exercises. We focus on the fluctuations in aggregate output.[19] Introducing countercyclical dispersions in idiosyncratic productivity shocks decreases the variance of output by roughly 9% relative to the case where there are no risk shocks.[20] This change identifies the total business cycle effect of the risk shocks. This effect includes the real options effect as well as the volatility and the intensive margin effect of risk. Since risk shocks are countercyclical, a decline in aggregate volatility means that the volatility effect also dominates in aggregate fluctuations.[21] Looking at the right panel of

[19]The results for the volatility of other aggregates are similar. The comovement patterns and persistence remain basically unchanged when we vary σ_σ.

[20]We focus on the comparison of variances instead of standard deviations, as only the former are meaningfully additive.

[21]There are also so-called Hartman-Abel effects at work here: higher risk concentrates economic activity in highly productive firms and increases aggregate output through Jensen's inequality.

Table 12 we see that the contribution to the output fluctuations in the data for the baseline calibration of risk shocks is only just above 1%.

If we were to entertain more volatile risk shocks, the effects would get larger. For $\sigma_\sigma = 4$ the variance of aggregate output would be dampened by over 17%. Risk shocks would explain 13% of output fluctuations, which in terms of importance would put them in the neighborhood of conventional monetary policy shocks (see, for instance, Smets and Wouters (2007)). However, such relatively strong risk fluctuations would be difficult to reconcile with the procyclicality of the investment rate dispersion, as we have seen in the previous section. We believe this is a more general insight: the procyclicality of the investment rate dispersion constrains any model with risk shocks and, thus, a fortiori, any model where risk shocks have important aggregate implications. This makes the procyclicality of the investment rate dispersion an important statistic for researchers studying the aggregate effects of risk shocks to firm-level driving forces.

IV.C Robustness

We have checked the robustness of our quantitative results with respect to a range of modeling assumptions and calibration choices. The results of these robustness checks are summarized in Table 13.[22] The bottom line is: both the (not targeted) procyclicality of the investment rate dispersion and the (not targeted) small relevance of risk shocks for aggregate fluctuations, given the procyclicality of the investment rate dispersion, are robust features.

First, we check whether the timing assumption – firms observe today's productivity dispersion – is important. When we alternatively assume that firms observe the dispersion of productivity shocks one period ahead, as some of the literature has done (Bloom (2009), for instance), we find that the aggregate "importance" of risk shocks increases slightly. But in terms of the broader economic picture, the results are unchanged.

Second, we check whether the assumption that adjustment costs are in units of labor is driving the results. In one experiment, we increase the disutilty of labor so that average time spent at work equals 0.25; without much impact for the baseline results. As a second experiment, we denote adjustment costs in units of the numeraire. In the baseline specification, with high wages adjustment is more costly in booms, which somewhat dampens investment activity: adjustment costs are procyclical in the baseline. In the alternative specification with acyclical adjustment costs, the procyclicality of the investment rate dispersion (through the increased procyclicality of the fraction of spike adjusters) increases slightly.

Third, we check whether our results change for deviations from the estimated production function parameters. Starting from Bloom's (2009) choice, $\theta = 0.25$, $\nu = 0.50$, as a refer-

[22]Unless stated otherwise, we recalibrate the adjustment cost parameter to minimize $\Psi(\xi)$ for each robustness check.

Table 13: Robustness of Quantitative Implications

	Full Model	Full Model	Risk Only Model
	Cyclicality of Investment Dispersion	Decline in Variance of Output	Variance Explained
Baseline	0.53	9%	1%
$\sigma(\epsilon_{j,t})$ observed in $t-1$	0.43	11%	2%
High disutility of labor	0.35	9%	1%
Adj. Costs in numeraire units	0.58	9%	1%
Elasticities $\theta = 0.25, \nu = 0.50$	0.56	11%	1%
Elasticities $\theta = 0.27, \nu = 0.53$	0.46	13%	2%
Elasticities $\theta = 0.28, \nu = 0.56$	0.33	15%	3%
Frictionless small investment	0.46	9%	1%
High Adj. Costs: $\bar{\xi} = 0.5$, $\sigma_\sigma = 2$	0.30	16%	4%
High Adj. Costs: $\bar{\xi} = 0.5$, $\sigma_\sigma = 4$	-0.24	20%	11%

Notes: See notes to Tables 8 and 12. We display $|\frac{var(Y)[\sigma_\sigma=\ldots]-var(Y)[\sigma_\sigma=0]}{var(Y)[\sigma_\sigma=0]}|$ (second column) and $\frac{var(Y)[\sigma_\sigma=\ldots]}{var(Y)[Data]}$ (third column). Unless stated otherwise, $\sigma_\sigma = 1$.

ence point[23] we increase the returns-to-scale from 0.75 to 0.83 (equivalently, we decrease the markup from 1.33 to 1.20) and thereby also increase the capital elasticity of static revenue from 0.50 to 0.63.[24] We then recompute the properties of idiosyncratic productivity in the USTAN data set, imposing the corresponding production function parameters, and recalibrate the adjustment costs parameter, $\bar{\xi}$.[25] It is clear that while the cyclicality of investment

[23]If one views the DRTS assumption in our model as a stand-in for a CRTS production function with monopolistic competition, then these choices correspond to an employment elasticity of the underlying production function of two-thirds and a markup of $\frac{1}{\theta+\nu} = 1.33$; returns-to-scale, $\theta + \nu$, are 0.75, and the capital elasticity of static revenue, $\frac{\theta}{1-\nu}$, is 0.5.

[24]We do so in a way that remains consistent with an underlying CRTS production function with an employment elasticity of two-thirds. Interestingly, with the exception of the two primary sector industries Agriculture and Mining and Energy, which exhibit a very low procyclicality in their investment rate dispersion, all other two-digit industries have $\frac{\theta}{1-\nu}$ lower than 0.63 and $\theta + \nu$ lower than or equal to 0.83.

[25]Parameters for $(\theta = 0.25, \nu = 0.5)$: $\bar{\xi} = 0.2$, $\bar{\sigma}(\epsilon) = 0.0919$, $\rho_\epsilon = 0.9525$. Parameters for $(\theta =$

dispersion decreases with a production function with less curvature,[26] it remains procyclical and in the ballpark of the baseline calibration. Importantly, a higher volatility of firm-level risk remains inconsistent with a procyclical investment dispersion under these production function specifications.

Fourth, our results are robust to an extension of the model that allows for costless small adjustments of a firm's capital stock, as proposed by Khan and Thomas (2008), to match the fraction of small investment rates in the data.

Finally, could larger idiosyncratic risk fluctuations be reconciled with the observed procyclicality of investment dispersion in the data, if we assume higher adjustment costs, as Table 9 suggests? The answer is yes, in principle, but quantitatively the evidence still points toward small risk fluctuations. The last two rows of Table 13 show that raising the adjustment cost parameter from $\bar{\xi} = 0.2$ to $\bar{\xi} = 0.5$ and at the same time doubling the volatility of firm risk leads to a correlation of investment dispersion with output of 0.3, compared to 0.45 in the data. Quadrupling the volatility of firm-level risk would lead to countercyclical investment dispersion (-0.24), as in the baseline calibration. We note that $\bar{\xi} = 0.5$ means that firms pay average adjustment costs per unit of output, conditional on adjustment, in the upper end of the estimated numbers in the literature: 26% versus 11% in the baseline (see Bloom (2009), Table IV).

V Final Remarks

The cross-sectional standard deviation of firm-level investment is robustly and significantly procyclical. This is likely the result of lumpy investment at the micro level and a modest amount of countercyclical fluctuations in firm-level risk. In an example, namely, the recent literature on risk shocks in environments with physical adjustment frictions, we show that the cyclicalities of cross-sectional dispersion measures are jointly informative about the strengths of these risk shocks. This is a more general insight: the procyclicality of the investment rate dispersion constrains any model that features shocks to firms' productivity dispersions. To the extent that such risk shocks matter for aggregate fluctuations, yet are difficult to measure directly, the procyclicality of the investment rate dispersion is an informative statistic for macroeconomics. More generally, our findings suggest that the time-series behavior of the entire cross-section of firm-level investment is informative on how investment and investment frictions at the micro level should be modeled. We leave an exploration of this hypothesis for future research.

$0.2667, \nu = 0.5333$): $\bar{\xi} = 0.3$, $\bar{\sigma}(\epsilon) = 0.0914$, ρ_ϵ =0.9375. Parameters for ($\theta = 0.2778, \nu = 0.5556$): $\bar{\xi} = 0.4$, $\bar{\sigma}(\epsilon) = 0.0912$, $\rho_\epsilon = 0.9200$.

[26]Interestingly, across the two-digit industries in USTAN we find a similar pattern between the curvature of the production function and the procyclicality of investment dispersion.

References

Ackerberg, D., Caves, K., and Frazer, G. (2006). Structural identification of production functions. MPRA Paper 38349, University Library of Munich, Germany.

Arellano, C., Bai, Y., and Kehoe, P. (2012). Financial markets and fluctuations in uncertainty. *Federal Reserve Bank of Minneapolis Research Department Staff Report 466.*

Bachmann, R. and Bayer, C. (2013). Wait-and-see business cycles? *Journal of Monetary Economics, forthcoming.*

Bachmann, R., Caballero, R., and Engel, E. (2013a). Aggregate implications of lumpy investment: New evidence and a DSGE model. *American Economic Journal: Macroeconomics, forthcoming in October.*

Bachmann, R., Elstner, S., and Sims, E. (2013b). Uncertainty and economic activity: Evidence from business survey data. *American Economic Journal: Macroeconomics*, 5(2):217–249.

Bayer, C. (2006). Investment dynamics with fixed capital adjustment cost and capital market imperfections. *Journal of Monetary Economics*, 53(8):1909 – 1947.

Beaudry, P., Caglayan, M., and Schiantarelli, F. (2001). Monetary instability, the predictability of prices, and the allocation of investment: An empirical investigation using U.K. panel data. *American Economic Review*, 91(3):648–662.

Berger, D. and Vavra, J. (2011). Dynamics of the US price distribution. *mimeo, Yale University.*

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Bloom, N., Floetotto, M., Jaimovich, N., S.-E. I., and Terry, S. (2012). Really uncertain business cycles. *NBER WP 18245.*

Caballero, R. J. and Engel, E. M. R. A. (1999). Explaining investment dynamics in U.S. manufacturing: A generalized (S,s) approach. *Econometrica*, 67(4):783–826.

Caballero, R. J., Engel, E. M. R. A., Haltiwanger, J. C., Woodford, M., and Hall, R. E. (1995). Plant-level adjustment and aggregate investment dynamics. *Brookings Papers on Economic Activity*, 1995(2):1–54.

Castaneda, A., Diaz-Gimenez, J., and Rios-Rull, J.-V. (1998). Exploring the income distribution business cycle dynamics. *Journal of Monetary Economics*, 42(1):93 – 130.

Christiano, L., Motto, R., and Rostagno, M. (2010). Financial factors in economic fluctuations. *ECB Working Paper 1192*.

Chugh, S. (2012). Firm risk and leverage-based business cycles. *mimeo, University of Maryland*.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *Review of Economic Studies*, 73(3):611–633.

Curdia, V. and Reis, R. (2011). Correlated disturbances and U.S. business cycles. *mimeo, Columbia University*.

Davis, S., Haltiwanger, J., and Schuh, S. (1996). *Job creation and destruction*. The MIT Press.

Davis, S. J., Haltiwanger, J., Jarmin, R., and Miranda, J. (2006). Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. Working Paper 12354, National Bureau of Economic Research.

Den Haan, W. J. (2010). Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents. *Journal of Economic Dynamics and Control*, 34(1):79–99.

Doepke, J., Funke, M., Holly, S., and Weber, S. (2005). The cross-sectional dynamics of German business cycles: a bird's eye view. *Discussion Paper Series 1: Economic Studies*, 23.

Doepke, J. and Weber, S. (2006). The within-distribution business cycle dynamics of German firms. *Discussion Paper Series 1: Economic Studies*, 29.

Doms, M. and Dunne, T. (1998). Capital adjustment patterns in manufacturing plants. *Review of Economic Dynamics*, 1(2):409–429.

Eisfeldt, A. and Rampini, A. (2006). Capital reallocation and liquidity. *Journal of Monetary Economics*, 53(3):369–399.

Gandhi, A., Navarro, S., and Rivers, D. (2011). On the identification of production functions: How heterogeneous is productivity? University of Western Ontario, CIBC Centre for Human Capital and Productivity Working Papers 20119, University of Western Ontario, CIBC Centre for Human Capital and Productivity.

Gilchrist, S., Sim, J., and Zakrajšek, E. (2010). Uncertainty, financial frictions, and investment dynamics. *mimeo, Boston University.*

Gourio, F. (2008). Estimating firm-level risk. *mimeo, Boston University.*

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

Heckman, J. J. (1976). The common structure of statistical models of truncation, sample selection and limited dependent variables and a simple estimator for such models. In *Annals of Economic and Social Measurement, Volume 5, number 4*, NBER Chapters, pages 475–492. National Bureau of Economic Research, Inc.

Higson, C., Holly, S., and Kattuman, P. (2002). The cross-sectional dynamics of the US business cycle: 1950-1999. *Journal of Economic Dynamics and Control*, 26(9-10):1539–1555.

Higson, C., Holly, S., Kattuman, P., and Platis, S. (2004). The business cycle, macroeconomic shocks and the cross-section: The growth of UK quoted companies. *Economica*, 71(282):299–318.

Kehrig, M. (2011). The cyclicality of productivity dispersion. *mimeo, Northwestern University.*

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

Krusell, P. and Smith, A. (1997). Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns. *Macroeconomic dynamics*, 1(02):387–422.

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

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, 70(2):317–341.

Narita, F. (2011). Hidden actions, risk-taking and uncertainty shocks. *mimeo, University of Minnesota.*

Olley, G. S. and Pakes, A. (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica*, 64(6):1263–1297.

Panousi, V. and Papanikolaou, D. (2012). Investment, idiosyncratic risk, and ownership. *Journal of Finance*, 67(3):1113–1148.

Schaal, E. (2011). Uncertainty, productivity and unemployment in the great recession. *mimeo, Princeton University*.

Smets, F. and Wouters, R. (2007). Shocks and frictions in U.S. business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606.

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.

Tauchen, G. (1986). Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181.

Vavra, J. (2012). Inflation dynamics and time-varying uncertainty: New evidence and an Ss interpretation.

von Kalckreuth, U. (2003). Exploring the role of uncertainty for corporate investment decisions in Germany. *Swiss Journal of Economics and Statistics*, 139(II):173–206.

Zhang, L. (2005). The value premium. *Journal of Finance*, LX(1):67–103.

A Data

A.1 Description of the Database

Our German firm-level data source is USTAN (*Unternehmensbilanzstatistik*) of Deutsche Bundesbank. 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. USTAN captures all major balance-sheet items, the major items of the profit and loss statements, and employment. Importantly, USTAN provides separate investment data for structures and equipment. As we will show below, the USTAN sample covers a large fraction of the nonfinancial private business sector (NFPBS).

It originated as a by-product of the Bundesbank's rediscounting, i.e. (overnight-)lending activities. By law, the Bundesbank was required to assess the creditworthiness of all parties backing a *Wechsel*, a promissory note or commercial bill of exchange, put up for discounting. It implemented this regulation by requiring balance-sheet data of all parties involved. These balance-sheet data were then collected and archived into a database.

Promissory notes were a form of trade credit with widespread use throughout the sample period. From the volume of a 0.15% stamp tax on promissory notes and bills of exchange, one can infer that a volume of these titles of roughly 10% of German GDP was issued each year. Moreover, rediscounting promissory notes was a commonly used instrument of monetary policy in Germany. Thus, unlike the Federal Reserve, the Bundesbank did not use T-bills as the major form of collateral but rather private debt. As far as potential cyclical sample selection is concerned, it is important to note that it had to happen only once in a given year that a promissory note from a given firm was used as collateral by someone in order for that firm to appear in USTAN, i.e. it is irrelevant how often that firm issued trade credit and in what volumes.

The quality of the data is particularly high. All mandatory data collected for USTAN have been double-checked by Bundesbank staff. 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.

One drawback of USTAN is that with the introduction of the euro, the Bundesbank stopped buying commercial bills and collected firm balance-sheet data only irregularly and only from publicly available sources. For this reason, the data set stops being useful in 1999. Therefore, we only use data from 1971 to 1998, which leaves us, after lagging and first-differencing, with 26 years of observations from 1973 to 1998.

The coverage of the sample is broad, although it is technically not a representative sample due to the nonrandom sample design. It was also more common to use promissory notes as trade credit in manufacturing and for incorporated companies, which biases our data somewhat toward these kinds of firms. And, of course, the Bundesbank would only rediscount notes to which it gave a good rating, so that the set of firms in USTAN is also somewhat biased toward financially healthy and larger firms. Nevertheless, USTAN covers a wide range of firms, in fact a wider range in some dimensions (size, ownership, industry) than comparable U.S. data sets (ASM, COMPUSTAT), since short-term financing through promissory notes was a common practice for many German companies across most business sectors. Due to the Bundesbank's rediscounting policy, bills of exchange were very liquid for the creditor.

A.2 Sample Selection

We start with the universe of observations in the USTAN data, merging the files for 1971-1986 and 1987-1998. In a first pass, we then drop all balance sheets that are irregular, e.g., bankruptcy or closing balance sheets, or that stem from a group/holding (*Konzernbilanz*). This leaves us with only regular balance sheets (*Handelsbilanz* or *Steuerbilanz*). We also drop all firms with missing payroll data or missing or negative sales data, which are basically nonoperating firms. A small amount of duplicate balance sheets is removed as well. Finally, we drop the following sectors: hospitality (hotels and restaurants), which only has a small number of firms in the database, financial and insurance institutions, the mostly public health and education sectors, as well as other public companies like museums, etc., and some other small service industries, such as hair cutters, dry cleaners and funeral homes;[27] or when sectoral information was missing. The sectoral aggregate we are studying can be roughly characterized as the nonfinancial private business sector in Germany. This sample selection leaves us with an initial data set of 1,764,846 firm-year observations and 259,614 different firms. The average number of firms per year is 63,030.

From this initial sample we remove step-by-step observations, in order to get an economically meaningful data set. We first drop observations from likely East German firms to avoid a break in the series in 1990. We identify a West German firm as a firm that has a West German address or has no address information but enters the sample before 1990. Then we recompute capital stocks with a perpetual inventory method (PIM). In the PIM we drop a small amount of outliers. We remove observations that do not have a log value-added and a log capital stock after PIM.

Another part of the data is removed when firms did not have changes in log firm-level

[27] The number of firms from the public sector and these small industries is tiny to begin with, as they did not regularly use bills of exchange as a financing instrument.

employment (N), capital (K) and real value-added (VA), which obviously requires us to observe firms for two consecutive years. Then we remove outliers in factor changes and real value-added changes. Specifically, we identify as outliers in our sample a firm-year in which the firm-level investment rate or log changes in firm-level real value-added, employment and capital stock fall outside a three-standard-deviations band around the firm and sectoral-year mean. Then we compute firm-level Solow residuals and similarly remove observations with missing log changes in Solow residuals as well as outliers therein. We finally remove – before and after each step of the outlier removal – firms that have less than five observations in firm-level Solow residual changes. We conduct extensive robustness checks of our results to the choices for the outlier and observation thresholds. Table 14 summarizes how many observations are dropped in each step.

Table 14: Sample Creation

Criterion	Firm-Year Observations
Initial Sample	1,764,846
East Germany	-104,299
Outliers in PIM	-7,539
Missing log value-added	-1,349
Missing log capital	-31,819
Missing log-changes in N, K, VA	-161,668
Outliers in N, K and VA log-changes	-41,453
Missing log-changes in Solow residual	-126,086
Outliers in Solow residual log-changes	-18,978
Not enough observations	-417,550
Final Sample	854,105

A.3 Sample Composition

The final sample then consists of 854,105 firm-year observations, which amounts to observations on 72,853 different firms. The average observation length of a firm in the sample is 11.7 years. The average number of firms per year is 32,850. The following Tables 15, 16 and 17 show the average industry,[28] the legal form and the size distributions in our final sample.

USTAN's industry coverage, while somewhat biased toward manufacturing firms, includes the construction, service and the primary sectors. While a bias toward larger firms remains, the size coverage is still fairly broad: 31% of all firm-year observations in our final baseline

[28]WZ 2003 is the industry classification from 2003 that the German national accounting system (*Volkswirtschaftliche Gesamtrechnung, VGR*) uses.

sample have fewer than 20 employees and 57% have fewer 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.

Table 15: Two-Digit Industry Distribution

ID	Sector	Observations	Frequency	WZ 2003
10	Agriculture	12,291	1.44%	A, B
20	Energy & Mining	4,165	0.49%	C, E
31	Chemical Industry, Oil	14,721	1.72%	DF, DG
32	Plastics, Rubber	23,892	2.80%	DH
33	Glass, Ceramics	28,623	3.35%	DI
34	Metals	30,591	3.58%	DJ
35	Machinery	162,407	19.01%	DK, DL, DM, DN
36	Wood, Paper, Printing	61,672	7.22%	DD, DE
37	Textiles, Leather	46,173	5.41%	DB, DC
38	Food, Tobacco	37,708	4.41%	DA
40	Construction	54,569	6.39%	F
61	Wholesale Trade	213,071	24.95%	G51
62	Retail Trade & Cars	142,137	16.64%	G50, G51
70	Transportation & Communication	22,085	2.59%	I
	Total	854,105	100%	

Table 16: Legal Form Distribution

Legal Form	Observations	Frequency
Publicly Traded (AG, KGaA, etc.)	18,582	2.18%
Limited Liability Companies (GmbH, GmbH&Co., etc.)	506,184	59.26%
Fully Liable Partnerships (OHG, KG, etc.)	327,526	38.35%
Other: unincorporated associations (e.V.) municipal agencies (Körperschaften öR) etc.	1,813	0.21%
Total	854,105	100%

Table 17: Size Distributions of Firms

Number of Employees	1-4	5-9	10-14	15-19	20-49	50-99	100-249	250-499	500+
Fraction	6.14%	9.46%	8.24%	7.30%	26.28%	17.04%	14.37%	5.68%	5.49%
Capital Stock (in 1000 1991-Euro)	0-299	300-599	600-999	1,000-1,499	1,500-2,499	2,500-4,999	5,000-9,999	10,000-24,999	25,000+
Fraction	8.23%	9.01%	9.67%	9.36%	13.08%	17.71%	13.87%	11.08%	7.99%
Real Value Added (in 1000 1991-Euro)	0-299	300-499	500-999	1,000-1,499	1,500-2,499	2,500-4,999	5,000-9,999	10,000-24,999	25,000+
Fraction	8.17%	7.93%	16.38%	11.56%	14.45%	16.28%	11.20%	8.25%	5.79%

A.4 Perpetual Inventory Method

In order to obtain economically meaningful stocks of capital series for each firm, we have to re-calculate capital stocks in a Perpetual Inventory Method (PIM); see Bayer (2006), for instance. The first step is to compute firm-level investment series, $I_{j,t}$, from the corporate balance sheets, which contain data only on accounting capital stocks, $k^a_{j,t}$, and accounting total depreciation, $d^a_{j,t}$. The following accumulation identity for the book value of capital allows us to back out nominal firm-level investment, $p^I_t I_{j,t}$:[29]

$$k^a_{j,t+1} = k^a_{j,t} - d^a_{j,t} + p^I_t I_{j,t}. \tag{10}$$

The next step is to recognize that capital stocks from corporate balance sheets are not directly usable for economic analysis for two reasons: 1) accounting depreciation, $d^a_{j,t}$, in corporate balance sheets is often motivated by tax reasons and is typically higher than economic depreciation, $\delta^e_{j,t}$, expressed as a rate; 2) accounting capital stocks are reported at historical prices. Both effects would lead to an underestimation of the real firm-level capital stock, if one were to simply deflate the current accounting capital stock, $k^a_{j,t}$, with a current investment price deflator, p^I_t (assuming that p^I_t increases over time). We therefore apply a Perpetual Inventory Method (PIM) to compute economic real capital stocks:

$$k^{(1)}_{j,1} = k^a_{j,1}. \tag{11}$$

$$k^{(1)}_{j,t+1} = (1 - \delta^e_t)\, k^{(1)}_{j,t} + \frac{p^I_t}{p^I_{1991,t}} I_{j,t}. \tag{12}$$

$k^a_{j,1}$ is the accounting capital stock in 1991 prices at the beginning of an uninterrupted sequence of firm observations – if for a firm-year we have a missing investment observation, the PIM is started anew when the firm appears again in the data set. The investment-good-price deflator is $p^I_{1991,t}$, with 1991 as the base year. We estimate the economic depreciation rate δ^e_t for each year from national accounting data, *VGR*, separately for equipment and nonresidential structures (Table 3.1.3, *VGR*, *Nettoanlagevermögen nach Vermögensarten in jeweiligen Preisen, Ausrüstungen und Nichtwohnbauten*; Table 3.1.4, *VGR*, *Abschreibungen nach Vermögensarten in jeweiligen Preisen, Ausrüstungen und Nichtwohnbauten*). *VGR* contains sectoral and capital-good-specific depreciation data only after 1991, which is why

[29] Specifically, $k^a_{j,t}$ is the sum of balance-sheet items ap65, *Technische Anlagen und Maschinen*, and ap66, *Andere Anlagen, Betriebs-und Geschäftsausstattung*, for equipment; and balance-sheet item ap64, *Grundstücke, Bauten*, for structures. Since balance-sheet data are typically end-of-year stock data, notice that $k^a_{j,t}$ is the end-of-period capital stock in year $t-1$. $d^a_{j,t}$ is profit and loss account item ap156, *Abschreibungen auf Sachanlagen und immaterielle Vermögensgegenstände des Anlagevermögens*. In contrast to $k^a_{j,t}$, $d^a_{j,t}$ is not given for each capital good separately. For the solution of this complication, see below.

we decided to use only capital-good-specific depreciation rates for the entire time horizon. For the data sources for investment price deflators, see footnote 32 below. The drawback to this procedure is that we do not directly observe capital-good specific $d^a_{j,t}$ in the balance sheets, so that (10) is not directly applicable to the two types of capital good separately. We therefore split up $d^a_{j,t}$ according to the fraction that each type of capital good accounts for in the book value of total capital, weighting each type of capital good by its VGR depreciation rate. We finally aggregate both types of capital into a single capital good at the firm level.

There is a final complication, which arises through relying on $k^a_{j,1}$ as the starting value of the PIM. The recalculation of capital stocks is motivated by the bias that historical cost accounting and tax depreciation induce, i.e., that the book value of capital is typically not a good estimate of the productive real capital stock of the firm at that time. To take this issue into account also for the first observation of a firm, we calculate the time-average factor ϕ (for each sector), by which $k^{(1)}_{j,t}$ is larger than $k^a_{j,t}$, and replace $k^a_{j,1}$ by $\phi k^a_{j,1}$ in the perpetual inventory method. We do this iteratively until ϕ converges, i.e., we calculate (using $k^{(0)}_{j,t} = k^a_{j,t}$ and $\phi^{(0)} = 1$):

$$
\begin{aligned}
k^{(n)}_{j,t+1} &= (1-\delta^e_t)\, k^{(n)}_{j,t} + \frac{p^I_t}{p^I_{1991,t}} I_{j,t} && (13)\\
k^{(n)}_{j,1} &= \phi^{(n-1)} k^{(n-1)}_{j,1} && (14)\\
\phi^{(n)} &= (NT)^{-1} \sum_{j,t} \frac{k^{(n)}_{j,t}}{k^{(n-1)}_{j,t}} && (15)
\end{aligned}
$$

We stop when for each sector and each capital good category $\phi < 1.1$.[30]

Since we want to compute economic, i.e. productive, capital stocks, we then – as a final step – add to the capital stock series from this iterative PIM the net present value of the real expenditures for renting and leasing equipment and structures.[31]

[30] Extreme ϕ's indicate that constant depreciation is not a good approximation for this particular firm. Such a firm will have had an episode of extraordinary depreciation (e.g., fire, accident, etc.) and the capital stocks by PIM will be a bad measure of the actual capital stock after the accident. That is why we drop a small number of observations from the top and bottom of the ϕ-distribution (see Table 14).

[31] We take item ap161, *Miet- und Pachtaufwendungen*, from the profit and loss accounts, deflate it by the implicit investment good price deflator, which we compute from Tables 3.2.8.1 and 3.2.9.1 from *VGR*, and then divide it by a measure of the user cost of capital. The latter is simply the sum of real interest rates for a given year, which we compute from nominal interest rates on corporate bonds and ex-post CPI inflation data, and the time-average, accounting capital-good-weighted depreciation rate per firm.

A.5 Representativeness

How well does the USTAN aggregate represent the nonfinancial private business sector (NFPBS) in Germany? Table 18 shows that USTAN represents on average 70% of the value-added of the NFPBS, 44% of its investment, etc.[32]

Table 18: USTAN and the NFPBS

	USTAN/NFPBS
Value-Added	70%
Investment	44%
Capital	71%
Employment	49%
Payroll	54%

Table 19 shows that the cross-sectional *averages* of investment as well as output, employment and productivity growth, computed from USTAN, are strongly positively correlated with the cyclical component of the real gross value-added of the nonfinancial private business sector. This means that USTAN represents well the cyclical behavior of the sectoral aggregate it is meant to represent.

[32] NFPBS value-added is taken from *Bruttowertschöpfung in jeweiligen Preisen*, Table 3.2.1 of *VGR*, deflated by the implicit deflator for aggregate value-added, Table 3.1.1 of *VGR* (we apply the same deflator to the USTAN data). The base year is always 1991. NFPBS investment is *Bruttoanlageinvestitionen in jeweiligen Preisen* from Table 3.2.8.1, deflated with the implicit sector-specific investment price deflators given by *Bruttoanlageinvestitionen - preisbereinigt*, a chain index, from Table 3.2.9.1, *VGR*. NFPBS capital is *Nettoanlagevermögen in Preisen von 2000* from Table 3.2.19.1, *VGR*, re-chained to 1991 prices. In computing both the investment and the capital data for USTAN in the PIM, we use the implicit sector and capital-good-specific (equipment and nonresidential structures) deflators for investment: Tables 3.2.8.2, 3.2.9.2., 3.2.8.3 and 3.2.9.3., *VGR*. We also experiment with deflating USTAN data with a uniform investment price deflator, the *Preisindex der Investitionsgüterproduzenten*, source: GP-X002, *Statistisches Bundesamt*. NFPBS employment is the number of employed, *Arbeitnehmer*, from Table 3.2.13, *VGR*. Finally, payroll is taken from *Arbeitnehmerentgelt*, Table 3.2.10., *VGR*, deflated by the same general implicit deflator for aggregate value-added that we use to deflate value-added numbers.

Table 19: Cyclicality of Cross-Sectional Averages

Cross-sectional Moment	Correlation with Cycle
$mean(i_{j,t})$	0.756***
$mean(\Delta \log y_{j,t})$	0.663***
$mean(\frac{\Delta n_{j,t}}{0.5*(n_{j,t-1}+n_{j,t})})$	0.602***
$mean(\Delta \log \epsilon_{j,t})$	0.592***

Notes: The table shows the correlation with the cycle of the cross-sectional averages, linearly detrended, of, respectively, the investment rate, the log-change of real gross value-added (we deflate the profit and loss account item ap153, Rohergebnis, with the aggregate value-added deflator from VGR data), the net employment change rate, and the log-change of Solow residuals, all at the firm level. We have removed firm fixed and 2-digit industry-year effects from each variable. As a cyclical indicator we use the HP(100)-filtered aggregate real gross value-added in the German nonfinancial private business sector, computed from German VGR (*Volkswirtschaftliche Gesamtrechnungen*) data. *** indicates significance at the 1%level, resulting from an overlapping block bootstrap of four-year windows with 10,000 replications.

A.6 The Cross-sectional Investment Rate Distribution

As Table 20 shows, the distribution of firm-level investment rates from our USTAN sample is comparable to the one calculated for the U.S. from the LRD, reported in Cooper and Haltiwanger (2006). For comparability with Cooper and Haltiwanger (2006), we also show the distribution of investment rates for $\frac{I_{j,t}}{k_{j,t}}$. The USTAN sample exhibits less disinvestment activity, slightly more inactivity, and a little less frequent spikes. There are various reasons for these differences. First, LRD is plant-level data, whereas our data is firm level, so some difference may come from unit aggregation, in particular, observing fewer spikes in USTAN. Second, the U.S. and German manufacturing sector were facing fairly different compositional trends over the respective sample periods, which may explain why more negative investment due to stronger reallocation is observed in the LRD.

Table 20: Investment Rate Distributions in USTAN

	Negative		Inactivity	Positive	
	spike	intermediate		intermediate	spike
$i_{j,t}$					
All firms (USTAN)	0.4%	2.6%	14.8%	68.9%	13.4%
Manufacturing (USTAN)	0.4%	2.0%	11.1%	74.7%	11.9%
$\frac{I_{j,t}}{k_{j,t}}$					
All firms (USTAN)	0.3%	2.6%	15.1%	67.7%	14.2%
Manufacturing (USTAN)	0.3%	2.0%	11.4%	73.6%	12.7%
LRD	1.8%	8.6%	8.1%	62.9%	18.6 %

Notes: $i_{j,t} = \frac{I_{j,t}}{0.5(k_{j,t}+k_{j,t+1})}$ denotes our baseline investment rate definition. $\frac{I_{j,t}}{k_{j,t}}$ denotes the definition of the investment rate used by Cooper and Haltiwanger (2006). An investment spike is defined, for either investment rate, as being larger than 20% in absolute value. Inactivity is defined as an investment rate that is smaller than 1% in absolute value. The LRD data are taken from Cooper and Haltiwanger (2006).

A.7 Estimating the Production Function

We estimate the coefficients θ, ν of the production function by the median of the firm average share of factor expenditure in total value-added, as defined by:[33]

$$\hat{\nu}_j = T_j^{-1} \sum_t \frac{w_{j,t} n_{j,t}}{y_{j,t}}$$

$$\hat{\theta}_j = T_j^{-1} \sum_t \frac{(r_t + \delta_j) k_{j,t}}{y_{j,t}}$$

In a frictionless setup, this is estimating the production function coefficients from the first-order conditions. Importantly, this estimator is robust to classical measurement error in capital and labor. We take as the real interest rate, r_t, the average return on corporate bonds minus the ex post inflation rate and calculate firm-specific depreciation rates, δ_j, from capital-good-specific *VGR* depreciation rates, weighted by the firm-specific capital good portfolio.

[33] We use profit and loss account item ap153, Rohergebnis, for firm-level value-added and profit and loss account item ap154, Personalaufwand, for the firm-level wage bill.

Under the null hypothesis of our model, i.e., nonconvex capital adjustment frictions, these first-order conditions do not hold exactly for capital. However, when we estimate the production function coefficients from data simulated from our model in data sets of comparable size to USTAN, we find that this simple estimation procedure from factor expenditure shares works remarkably well. In fact, the labor share is estimated exactly, as labor can be adjusted without frictions in the model. The capital share estimated from the model-simulated data is $\hat{\theta} = 0.2238$, with the true θ being 0.2075.[34]

Alternative estimation approaches, such as those advocated by Olley and Pakes (1996) or Levinsohn and Petrin (2003), suffer from the collinearity issues discussed in Ackerberg et al. (2006) and Gandhi et al. (2011). In fact, when we estimate the production function parameters from model-simulated data using Olley and Pakes' (1996) estimator, we obtain a greatly upwardly biased estimate for ν: $\hat{\nu}_{OP} = 0.9962$. The estimated coefficient for capital is virtually zero. Also for the USTAN data we obtain fairly high coefficients on labor and low coefficients on capital. For manufacturing, for instance, we obtained Olley and Pakes estimates of $\nu = .744$ and $\theta = .069$, which is another indication that our model is a good description of firm-level behavior in the USTAN data set. The Olley and Pakes estimates are essentially unaffected by the inclusion (or lack thereof) of a selection term (cf. Appendix B) or the order of approximation used in the third stage of the estimator. In the first stage, we use only observations where the absolute value of the investment rate exceeds 20%.[35]

In summary, under the null hypothesis of our model, the factor shares are good estimators of the production function parameters despite the capital adjustment frictions.

A.8 The Idiosyncratic Productivity Process

We estimate the log-productivity residuals implied by the estimates for the production function coefficients as:

$$\hat{\epsilon}_{j,t} = \log(y_{j,t}) - \hat{z}_t - \hat{\nu} \log n_{j,t} - \hat{\theta} \log k_{j,t}.$$

We then specify for $\hat{\epsilon}_{j,t}$

$$\hat{\epsilon}_{j,t} = \epsilon_{j,t} + x_{j,t} + \mu_j$$

$$\epsilon_{j,t} = \rho \epsilon_{j,t-1} + u_{j,t},$$

[34] We thus find a small upward bias for the capital coefficient θ. Yet, an exact calibration alongside the adjustment cost parameter using our model would be prohibitively time-consuming and, given the negligible size of the bias, would not change our results substantively.

[35] Using a 10% threshold yields basically the same results.

where $x_{j,t}$ denotes measurement error. In order to estimate ρ we allow the measurement error to be a second- or third-order moving average process. This yields an estimation equation

$$\hat{\epsilon}_{j,t} = \rho\hat{\epsilon}_{j,t-1} + \sum_{s=1\ldots J} \xi_s \Delta\hat{\epsilon}_{j,t-s} + (1-\rho)\mu_j + \zeta_{j,t},$$

where we experiment with $J = 3$ and $J = 4$. An unbiased estimate of ρ can be obtained by an instrumental variable regression, where one uses lagged differences of $\Delta\hat{\epsilon}_{j,t-J-1}$ as instruments for $\epsilon_{j,t-1}$. The estimated ρ is 0.960 and 0.974 for $J = 3$ and $J = 4$, respectively. The coefficient used in our calibration is 0.9675 - the midpoint of these two estimates.

When measuring the size and the cyclicality of the dispersions of productivity growth, we apply a slightly simplified approach, given the relatively high values of ρ (and low values of ξ_s) that we estimate. Since a specification with a long moving average term for the measurement error would force us to discard many aggregate data points for generating the necessary lags, we specify $\epsilon_{j,t}$ as a random walk cum classical measurement error, for estimating the variance of the shocks as[36]

$$E\Delta\hat{\epsilon}^2_{j,t} - \sigma^2_{me},$$

where the variance of the measurement error, $\Delta x_{j,t}$, is estimated by the sample analogue to

$$\begin{aligned}
\sigma^2_{me} &= -E(\hat{\epsilon}_{j,t} - \hat{\epsilon}_{j,t-2})^2 + 2E\Delta\hat{\epsilon}^2_{j,t} \\
&= -E(\epsilon_{j,t} - \epsilon_{j,t-2} + x_{j,t} - x_{j,t-2})^2 + 2E(\Delta\epsilon_{j,t} + \Delta x_{j,t})^2 \\
&= -E(\Delta\epsilon_{j,t} + \Delta\epsilon_{j,t-1} + x_{j,t} - x_{j,t-2})^2 + 2E(\Delta\epsilon_{j,t} + \Delta x_{j,t})^2 \\
&= -E(-\sigma_{\Delta\epsilon,t} - \sigma_{\Delta\epsilon,t-1} - 2\sigma_x + 2\sigma_{\Delta\epsilon,t} + 4\sigma_x) = 2\sigma_x.
\end{aligned}$$

B Robustness of the Empirical Findings

This appendix provides the details for the robustness checks discussed in Section I.C. Table 21 shows robustness with respect to the cyclical indicator. We experiment with different HP-filter smoothing parameters, with dynamic correlations, with excluding the post-reunification period and with excluding the years with the most extreme investment dispersion observations.[37]

Table 22 provides robustness checks with respect to sample composition. The first two

[36] We have conducted a Monte Carlo analysis that shows that the mistake one makes when our measurement error estimation is applied to not-quite-unit-root data is small. We therefore prefer the simple estimation procedure for the idiosyncratic shock variance.

[37] For the sake of brevity, we do not show results for firm-level Solow residual growth and fractions of adjusters. Results for both are robust: for the former, similar to the values for output growth, for the latter, similar to the values for the fraction of spike adjusters.

rows (dropping the first three observations per firm or looking only at firms that are virtually always in the sample) show that our results are neither driven by firm entry (into the sample) nor by firm exit, nor do our results depend on how we remove outliers from the sample.

Further robustness checks for the procyclicality of the investment rate dispersion are available in Table 23. We check for the robustness with respect to replacing the standard deviation with the interquartile range as the measure of dispersion, and to alternative ways of detrending the dispersion series. Perhaps most important, the last row of Table 23 shows that the procyclicality of the investment rate dispersion is not due to cyclical variations in the sample composition. In the scenario 'Selection correction' we control for sample selection in the following way: we estimate a simple selection model, where lagged firm-level Solow residuals determine selection and the firm-level investment rate is modeled as a mean regression. We use the maximum likelihood estimator by Heckman (1976) to infer the selection-corrected variance of the residual in the firm-level investment rate equation. The latter is very close to the sample variance of firm-level investment rates, indicating that our results are not influenced by systematic sample drop outs. While the first stage of the regression shows that there is a positive selection in terms of levels (more productive firms being more likely to be in the sample), there is no strong selection with respect to changes or investment rates and their dispersions.

Finally, Table 24 shows that the ownership structure matters for cross-sectional results (focusing on publicly traded firms in Germany would eliminate the procyclicality of investment dispersion), making it important to use broader data sets for the study of cross-sectional facts (see Davis et al. (2006), for a similar point). Figure 1 plots the time-series of the cross-sectional standard deviation of firm-level investment rates, linearly detrended and detrended with an HP(100)-filter, the time-series of the cross-sectional interquartile range of firm-level investment rates, linearly detrended, and the fraction of spike adjusters, linearly detrended, against the cyclical component of aggregate real gross value-added.

Table 21: Robustness Checks I – Cyclical Indicator

Cyclical Indicator / Sample	Correlation of ... with			
	$std(i_{j,t})$	$std(\Delta y_{j,t})$	$std(\frac{\Delta n_{j,t}}{0.5*(n_{j,t-1}+n_{j,t})})$	Fraction of Spike Adjusters
Baseline: HP(100)-filtered Real Gross Value-Added, Y	0.45^{**}	-0.45^{*}	-0.50^{**}	0.61^{***}
HP(6.25)-filtered Y	0.37^{**}	-0.47^{**}	-0.53^{**}	0.44^{***}
HP(100)-filtered I	0.72^{***}	-0.30	-0.31	0.78^{***}
Lag of HP(100)-filtered Y	0.26	0.12	0.04	0.36^{*}
Lead of HP(100)-filtered Y	0.35	-0.56^{**}	-0.54^{**}	0.48^{**}
Pre-reunification: 1973-1990	0.30	-0.65^{***}	-0.63^{***}	0.51^{**}
Exclude min and max $std(i_{j,t})$ years	0.54^{***}	-0.49^{**}	-0.53^{**}	0.66^{***}
Exclude two max $\|std(i_{j,t})\|$ years	0.50^{***}	-0.58^{***}	-0.57^{**}	0.63^{***}

Notes: This table shows the correlation with the cycle of the cross-sectional standard deviations, linearly detrended, of, respectively, the investment rate, the log-change of real gross value-added and the net employment change rate, all at the firm level. It also shows, in the last column, the correlation with the cycle of the average fraction, linearly detrended, of firms exhibiting an investment spike. All data are from the Bundesbank's USTAN database. We have removed firm fixed and 2-digit industry-year effects from each variable. As a cyclical indicator we use aggregate real gross value-added and aggregate investment in the German nonfinancial private business sector, computed from German VGR (*Volkswirtschaftliche Gesamtrechnungen*) data. ***,**,* indicate significance at the 1%, 5%, and the 10% level, respectively, resulting from an overlapping block bootstrap of four-year windows with 10,000 replications.

Table 22: Robustness Checks II – Sample Composition

Sample	Correlation of ... with HP(100)-filtered Y			
	$std(i_{j,t})$	$std(\Delta y_{j,t})$	$std(\frac{\Delta n_{j,t}}{0.5*(n_{j,t-1}+n_{j,t})})$	Fraction of Spike Adjusters
W/o entry: drop first 3 obs. per firm	0.39^{*}	-0.48^{*}	-0.52^{*}	0.61^{***}
Only firms with 20+ obs.	0.39^{***}	-0.38^{**}	-0.37^{**}	0.77^{***}
Stricter outlier removal (2.5 std.)	0.45^{**}	-0.45^{*}	-0.55^{**}	0.62^{***}
Looser outlier removal (5 std.)	0.42^{**}	-0.44^{*}	-0.22	0.62^{***}
Percentile outlier removal (5%)	0.59^{***}	-0.45^{**}	-0.60^{***}	0.64^{***}
Percentile outlier removal (1%)	0.48^{**}	-0.47^{**}	-0.42^{*}	0.62^{***}

Notes: See notes to Table 21.

Table 23: Robustness Checks III – Other

Correlation of investment dispersion with HP(100)-filtered Y	
Baseline	0.45 **
$iqr(i_{j,t})$	0.57 ***
Raw data - no fixed effects	0.45 ***
Uniform price index for investment	0.43 **
$std(\frac{I_{j,t}}{k_{j,t}})$ (dropping 5% outliers)	0.60 ***
$std(\frac{I_{j,t}}{k_{j,t}})$ (dropping 1% outliers)	0.48 **
$std(i_{j,t})_{\text{quadratic detrending}}$	0.56 ***
$std(i_{j,t})_{\text{cubic detrending}}$	0.60 ***
$std(i_{j,t})_{\text{HP(100)-detrending}}$	0.62 ***
Outlier $\geq$ 3 std means merger	0.42 **
Shorter in sample (2 obs.)	0.44 **
Selection correction	0.38 **

Notes: See notes to Table 21. $iqr(i_{j,t})$ refers to the cross-sectional interquartile range of firm-level investment rates. 'Raw data - no fixed effects' uses the standard deviation of the raw firm-level investment rates, no fixed effects removed. '$std(\frac{I_{j,t}}{k_{j,t}})$ (dropping 5% outliers)' uses $\frac{I_{j,t}}{K_{j,t}}$ the definition of the investment rate used by Cooper and Haltiwanger (2006). To take care of the higher sensitivity to outliers, we use a 5% outlier criterion here. '$std(\frac{I_{j,t}}{k_{j,t}})$ (dropping 1% outliers)' does the same with a 1% outlier criterion. 'Uniform price index for investment' refers to a scenario, in which we deflate firm-level investment and capital with an aggregate price deflator for investment goods, not with one-digit industry- and capital-good-specific ones. The next three rows show the results, when we detrend $std(i_{j,t})$ not with a linear trend, but, respectively, with a quadratic, cubic trend and an HP(100)-filter. 'Outlier $\geq$ 3 std means merger' refers to a scenario, in which we treat an observation of 3 standard deviations above or below the year-specific mean as indicating a merger and mark the firm henceforth as a new one. 'Shorter in sample (2 obs.)' refers to a scenario, in which we require firms to have two observations in first differences (instead of five) to be in the sample. 'Selection correction' refers to a scenario where we estimate a simple selection model, where lagged firm-level Solow residuals determine selection and the firm-level investment rate is modeled as a mean regression. We use the maximum likelihood estimator by Heckman (1976) to infer the selection-corrected variance of the residual in the firm-level investment rate equation.

Table 24: Cyclicality of Cross-Sectional Investment Dispersion – Legal Form

Aggregate	Publicly Traded	Limited Liability Companies	Fully Liable Partnerships
0.45**	0.10	0.32*	0.64***

Notes: See notes to Table 21. 'Publicly Traded' means the German legal forms of *AG* and *KGaA*. 'Limited Liability Companies' means the German legal forms of *GmbH* and *GmbH & Co KG*. 'Fully Liable Partnerships' means the German legal forms of *GBR*, *OHG* and *KG*.

Figure 1: Cyclicality of Cross-sectional Moments - Data

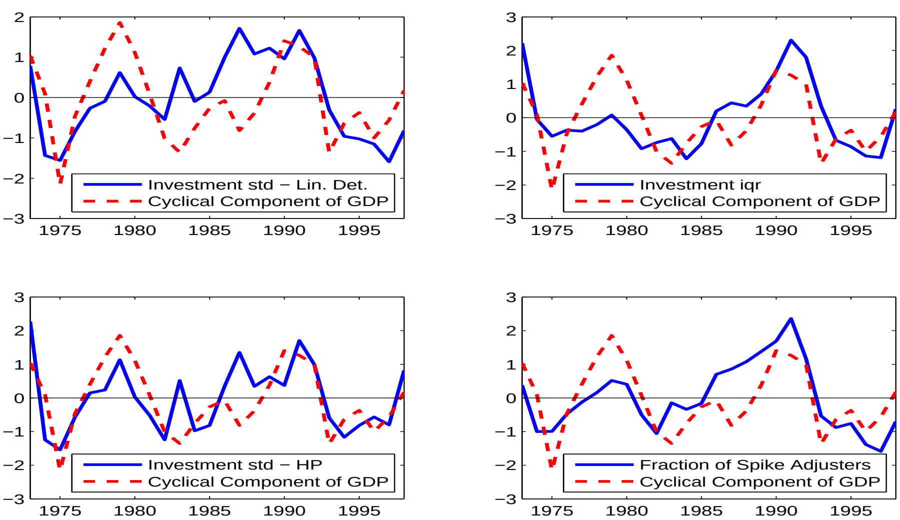

Notes: See notes to Table 21. For better readability, in each panel we normalize the cyclical components of the cross-sectional moments and of aggregate real gross value-added by their respective standard deviations.

C Numerical Implementation

C.1 Details

C.1.1 Decision Problem

The computable dynamic programming problem of a firm in our model has a 6-dimensional state space: $(k, \bar{k})$ are the endogenous states, while $(\epsilon, z, s, std(\epsilon))$ are the exogenous states. Since the employment problem has an analytical solution, there is essentially just one continuous control, k'. We discretize the state space as follows:

1. k: $n_k = 41$ grid points from $[0, 40]$, with a smaller grid width at low capital levels, where the curvature of the value function is higher.

2. $\bar{k}$: $n_{\bar{k}} = 4$ grid points: $[0.8, 1.0, 1.2, 1.4]$.

3. ϵ: $n_\epsilon = 19$. The grid points are equi-spaced (in logs) and the width between the midpoint, which is normalized to unity, and the extreme points is given by $3 \times \sqrt{\frac{\bar{\sigma}(\epsilon)^2}{1-\rho_\epsilon^2}}$, i.e., three times the unconditional variance of idiosyncratic productivity.

4. z : $n_z = 5$ grid points: $[0.9561, 0.9778, 1.0000, 1.0227, 1.0459]$.

5. s : $n_s = 5$ grid points: $[0.0764, 0.0834, 0.0905, 0.0976, 0.1047]$.

6. $std(\epsilon)$: $n_{std(\epsilon)} = 3$ grid points: $[0.30, 0.35, 0.40]$.

The various transition matrices for the stochastic processes are calculated as follows: first, using a bivariate version of the procedure in Tauchen (1986), we compute a discrete bivariate Markov chain on $z \times s$, using the results in equation (9). Second, we then compute (in the baseline case) for each s a transition matrix on the fixed (across s) ϵ-grid. The transition matrix takes two features into account: time-varying $\sigma(\epsilon)$ and the (small) excess kurtosis of the idiosyncratic productivity process in the data: 4.4480 on average. Since in our calibration strategy the fixed adjustment costs parameter is 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.

Since we allow for a continuous control, k, and $\bar{k}$ can take on any value continuously, we can only compute the value function exactly at the grid points above and interpolate for in-between values. This is done by using a multidimensional cubic splines procedure, with a so-called "not-a-knot" condition, to address the degrees-of-freedom problem that arises when splines are used. We compute the solution by value function iteration, using 10 steps of policy improvement after each actual optimization step. The optimum is found by using a golden section search. Upon convergence, we check single-peakedness of the objective function, to guarantee that the golden section search is reasonable. We have also experimented with finer grids for the baseline case and found our results to be robust.

C.1.2 Equilibrium Simulation

For the equilibrium simulations we draw one random series for the aggregate states and fix it across models. We use $T = 1600$ and discard the first 500 observations, when we compute statistics from these simulations.

As in the firm's decision problem, we use a golden section search to find the optimal target capital level, given p, at every point in time during the simulation. We find the market-clearing intertemporal price, using a combination of bisection, secant and inverse quadratic interpolation methods. The precision of the market-clearing algorithm is better than 10^{-7} at every point in time during the simulation.

There is a final complication due to the nature of the bivariate aggregate shock process: given the correlated processes for aggregate productivity and idiosyncratic firm-level risk, not all of the 25 (5×5) distinct aggregate state combinations are reached with sufficient frequency during our $T = 1600$-simulations to compute the regressions (8a) and (8b) state by state. Since going much beyond $T = 1600$ would be prohibitively burdensome in terms of computational resources and time, we proceed as follows: if we have at least five observations on an aggregate state combination, we run the state-by-state regressions. Otherwise, we use a version of (8a) and (8b) where we treat z and $\sigma(\epsilon)$ as if they were continuous variables, include several, but not all possible interaction terms, and run OLS regressions on these modified rules to extrapolate the coefficients for the remaining aggregate state combinations. As we will show below, these somewhat restricted regressions nevertheless provide a good fit for the time-series of aggregate capital and the marginal utility of consumption.

C.2 Quality of Numerical Approximations

What role does $\log std(\epsilon)$ play in the Krusell and Smith (1998) or KS rules, (8a) and (8b)? Table 25 shows that both our cross-sectional results and, to some extent, our aggregate results depend crucially on the addition of $\log std(\epsilon)$ in the KS rules. For instance, the procyclicality of the investment rate dispersion is substantially lower without it.

Tables 26, 27 and 28 illustrate why. First, the $R2$ is substantially higher and the standard error of estimation is substantially lower when $\log std(\epsilon)$ is added in the KS rules, especially the KS rules for the aggregate capital stock. Second, as Den Haan (2010) argues, it is important to check the quality not only of the one-step-ahead forecasts of the KS rules but also of forecasts at longer horizons. The performance of the KS rules – measured by the mean squared percentage deviation of applying the forecasting rules t times (using the actual realizations of the driving processes) from the actual value on the equilibrium simulation path (assuming households use the converged one-step-ahead forecasting rules), measured by the maximum absolute percentage deviation and measured by the correlation coefficient between the t-forecast and the actual value on the equilibrium simulation path – deteriorates substantially, especially at longer horizons, when $\log std(\epsilon)$ is not used.

Table 25: Economic Implications of $\log std(\epsilon)$ Being in the Krusell-Smith Rules

correlation between Y and $std(i_{j,t})$	BL-Model w $\log std(\epsilon)$	0.53			
	BL-Model w/o $\log std(\epsilon)$	0.28			
	Data	0.45			
		Y	C	I	N
volatility	BL-Model w $\log std(\epsilon)$	2.19	0.82	10.22	1.52
	BL-Model w/o $\log std(\epsilon)$	2.16	1.03	10.23	1.56
	Data	2.30	1.79	4.37	1.80
persistence	BL-Model w $\log std(\epsilon)$	0.29	0.55	0.22	0.20
	BL-Model w/o $\log std(\epsilon)$	0.22	0.65	0.08	0.06
	Data	0.48	0.67	0.42	0.61
correlation with Y	BL-Model w $\log std(\epsilon)$	1.00	0.87	0.98	0.96
	BL-Model w/o $\log std(\epsilon)$	1.00	0.74	0.93	0.89
	Data	1.00	0.66	0.83	0.68

Notes: The table displays – in the upper panel – the correlation of the dispersion of the investment rates with the cycle for the model under the baseline calibration ('BL-Model'), where $\log std(\epsilon)$ is included in the Krusell-Smith rules, a version of the baseline model, where $\log std(\epsilon)$ is not included in the Krusell-Smith rules, as well as from the USTAN data. In the lower panel the table does the same comparison for the percent standard deviations (volatility), autocorrelation (persistence), and correlation with aggregate output of HP(100)-filtered log aggregate output (Y), consumption (C), investment (I), and employment (N).

Table 26: Quality of KS Rules - R2 and the Standard Error

	Rule for $\bar{k}$					Rule for p				
	KS rules include $\log std(\epsilon)$									
	s_1	s_2	s_3	s_4	s_5	s_1	s_2	s_3	s_4	s_5
z_1	-	-	-	0.9999	1.0000	-	-	-	0.9996	1.0000
z_2	-	0.9997	0.9996	0.9997	1.0000	-	0.9999	0.9991	0.9994	0.9998
z_3	0.9997	0.9993	0.9997	0.9997	1.0000	0.9997	0.9989	0.9994	0.9995	1.0000
z_4	0.9995	0.9996	0.9997	0.9997	-	0.9994	0.9994	0.9996	0.9999	-
z_5	0.9997	0.9997	-	-	-	0.9995	0.9997	-	-	-
	Results for the regression w/o all interaction effects									
	$R2 = 0.9998$			$S.E. = 3.14 * 10^{-4}$		$R2 = 0.9997$				$S.E. = 2.20 * 10^{-4}$
	KS rules do not include $\log std(\epsilon)$									
	s_1	s_2	s_3	s_4	s_5	s_1	s_2	s_3	s_4	s_5
z_1	-	-	-	0.9586	0.9769	-	-	-	0.9984	0.9989
z_2	-	0.9420	0.8786	0.8782	0.8890	-	0.9992	0.9959	0.9960	0.9958
z_3	0.8256	0.8241	0.8729	0.8763	0.9029	0.9952	0.9935	0.9956	0.9956	0.9970
z_4	0.8398	0.8745	0.8878	0.9592	-	0.9940	0.9952	0.9959	0.9985	-
z_5	0.9376	0.7452	-	-	-	0.9972	0.9886	-	-	-
	Results for the regression w/o all interaction effects									
	$R2 = 0.9062$			$S.E. = 8.00 * 10^{-3}$		$R2 = 0.9976$				$S.E. = 8.00 * 10^{-4}$

Notes: All simulation results reported in the paper refer to the setup in the upper panel.

Table 27: Quality of Forecasting Rules at Various Horizons

Rule for ...	Evaluation Criterion	Forecast Horizon t $t=1$	$t=5$	$t=10$	$t=100$
$\bar{k}$	Mean Squared % Dev. forecast - actual	0.0003	0.0007	0.008	0.0008
$\bar{k}$	Max. Abs. % Dev. forecast - actual	0.0014	0.0034	0.0036	0.0037
$\bar{k}$	Correlation forecast - actual	0.9999	0.9995	0.9994	0.9994
p	Mean Squared % Dev. forecast - actual	0.0002	0.0003	0.0003	0.0003
p	Max. Abs. % Dev. forecast - actual	0.0007	0.00010	0.0011	0.0011
p	Correlation forecast - actual	0.9999	0.9998	0.9997	0.9997

Table 28: Quality of Forecasting Rules at Various Horizons - KS rules do not include $\log std(\epsilon)$

Rule for ...	Evaluation Criterion	Forecast Horizon t $t=1$	$t=5$	$t=10$	$t=100$
$\bar{k}$	Mean Squared % Dev. forecast - actual	0.0080	0.0220	0.0270	0.0296
$\bar{k}$	Max. Abs. % Dev. forecast - actual	0.0273	0.0832	0.0966	0.0943
$\bar{k}$	Correlation forecast - actual	0.9525	0.5995	0.3689	0.2404
p	Mean Squared % Dev. forecast - actual	0.0008	0.0087	0.0116	0.0130
p	Max. Abs. % Dev. forecast - actual	0.0027	0.0309	0.0402	0.0394
p	Correlation forecast - actual	0.9989	0.8468	0.7201	0.6471

Notes: All simulation results reported in the paper refer to the setup in Table 27.