

Endogenous profitability dispersion

Florian Kuhn

May 19, 2015

Abstract

Recent research has shown that higher uncertainty — which increases profitability dispersion across firms — may cause recessions. This paper explores how recessions can cause an endogenous rise in uncertainty about profitability. If heterogeneous firms face real and financial frictions, then a first moment shock leads to countercyclical profitability dispersion through firms' heterogeneous responses in price setting. Additionally, the mechanism endogenously generates countercyclical credit spreads and credit spread dispersion. The model explains two-thirds (69%) of the observed increase in profitability dispersion by fluctuations in aggregate levels holding productivity risk constant. This suggests that small uncertainty shocks may still be needed to explain the remaining volatility in profitability risk.

1 Introduction

In recent years the empirical finding that recessions are associated with an increase in volatility of firm profitability has been widely documented. A number of papers has since shown that this correlation can be driven by the effects of exogenous movements in firm risk (fittingly called “uncertainty shocks”). A smaller strand of the literature has examined the other direction of causation, investigating how cross-sectional dispersion among firms can increase following a worsening of aggregate conditions. In both cases, most papers have focused on firms' physical productivity as a measure of firm profitability, and how this cross-sectional productivity dispersion interacts with the business cycle.

This paper explicitly models firms’ pricing decisions to account for differences in physical productivity (TFP) and revenue productivity (TFPR or “profitability”)¹. The aim is to assess how much of the measured profitability dispersion can potentially arise from conventional first-moment shocks to TFP even when the underlying TFP dispersion is constant. To do so, the model combines the standard simple pricing behavior derived from monopolistic competition with two frictions often used in uncertainty-shock models: non-convex adjustment costs and costly financial intermediation. I choose these two frictions in order to make the model comparable to models that rely on uncertainty shocks to generate dispersion in TFPR. Moreover, the inclusion of a financial channel allows me to assess consistency of the model with data on the financial state of firms that are frequently interpreted as a channel through which firm-risk operates.

The countercyclicity of a number of cross-sectional dispersion measures has emerged as a fairly robust finding. Eifeldt and Rampini (2006) document that capital productivity is more dispersed in recessions. In seminal papers, Bloom (2009) and Bloom et al. (2012) show that the distribution of stock returns, firm sales growth, shocks to plant TFPR, and sectoral output all become wider when aggregate output is low. Kehrig (2011) establishes that dispersion in the level of profitability is countercyclical, especially for unproductive firms.

If one takes increases in firm risk as exogenous, there are two prominent ways in which these uncertainty shocks can generate recessions. The first channel is a real-options effect: If it is costly to adjust production inputs it maybe worth holding off doing so until the economic environment is less uncertain (Bachmann and Bayer (2013b) have dubbed this the “wait-and-see effect”). In this way high firm risk can reduce investment and hiring, thereby lowering output in subsequent periods. This is the mechanism used in Bloom (2009), Bloom et al. (2012), and Bachmann and Bayer (2013a). A second channel is through financial intermediation. If firms need to borrow funds in order to invest, then high firm risk will make it more likely that the firm ends up defaulting on its debt. This drives up the risk premium which in turn dampens investment. Models utilizing this effect include Christiano et al. (2013), Arellano et al. (2012) and Gilchrist et al. (2013).

There are also a variety of papers that endogenize countercyclical pro-

¹Foster et al. (2005) have shown that this distinction can be important for firm dynamics for the case of entry and exit decisions.

ductivity dispersion. Kehrig (2011) shows that in a model of entry and exit under overhead costs the marginal entrant's productivity can be procyclical, implying a tighter productivity distribution in a boom. Cui (2013) builds a vintage model of capital in which capital reallocation is procyclical and consequently recessions are times when relatively unproductive machinery is still being utilized. In ?'s, intangible capital investment is needed to access markets. Firms find it optimal to invest more heavily in intangibles during a boom, giving them access to a larger number of markets such that market-specific risk smooths out. Finally, Bachmann and Moscarini (2012) provide a model of demand uncertainty in which firms, when unsure if their drop in demand is due to weak aggregate demand, which is transient, or due to weak private demand, which is permanent. In order to learn about their permanent demand elasticity, firms find it profitable to experiment by setting a higher price in recessions when the opportunity cost of foregoing profits is low. This in turn increases dispersion in sales among firms.

This paper is similar to Bachmann and Moscarini (2012) in the sense that to my knowledge they are the only two papers explicitly modeling prices in a framework of endogenous firm risk. Here however, firms face idiosyncratic TFP shocks, allowing me to investigate the relationship between physical productivity and profitability, whereas in Bachmann and Moscarini firms differ in their (constant) demand elasticity and an iid idiosyncratic demand shock. This paper is also close to Gilchrist et al. (2013) in the sense that their paper includes both real adjustment costs and costly financial intermediation (in fact the mechanism of the financial friction here is taken from their paper). Their focus however is on the effect of second-moment shocks on aggregate outcomes. While they do consider aggregate TFP shocks, these shocks can not generate revenue productivity dispersion among firms since they all produce the same output good and thus no pricing mechanism exists.

The outline of the paper is as follows: The next section 2 describes the model setup. Section 3 aims to give the intuition how adjustment costs generate endogenous TFPR dispersion. In section 4 I discuss calibration and results. I find that the model can generate substantial countercyclical dispersion in revenue productivity. It also matches the cyclicity, but not the size of movements in credit spreads and credit spread dispersion among firms. I interpret this as evidence that the distinction between productivity and profitability risk is important in the context of models featuring uncertainty shocks.

2 Model

2.1 Households

The model's household and final goods sectors are kept as simple as possible. The representative household works for a fixed amount of hours and owns the firms in the other sectors. Hence, in any period her only decision is between consumption C_t and saving S_t . Savings are deposited in the representative bank discussed below. Normalizing her labor supply to one, the household's labor income is equal to the wage w_t . She also earns capital income from past savings $R_{t-1}S_{t-1}$, where R_{t-1} is the risk-free real interest rate determined last period. Finally, she receives all profits from the final goods, banking, and intermediate goods sectors, respectively $\Pi_t \equiv \pi_{ft} + \pi_{bt} + \pi_{it}$. Consequently her flow budget constraint is

$$C_t + S_t = w_t + R_{t-1}S_{t-1} + \Pi_t$$

Subject to the sequence of budget constraints, the household maximizes

$$U(\{C_t\}_{t=0}^{\infty}) = E_t \left[\sum_{t=0}^{\infty} \beta^t \frac{1}{1-\psi} C_t^{1-\psi} \right]. \quad (1)$$

2.2 Final goods sector

The final goods sector is represented by a single competitive firm functioning as a standard CES aggregator. It uses as input a continuum of intermediate goods $\{y_{jt}\}_{j \in [0,1]}$, where j indexes the variety of the good. Production according to a function F_f yields output $Y_t = F_f(\{y_{jt}\}_j)$ which can be used for consumption and investment. All prices in the economy are expressed in units of the final good, i.e its price is normalized to 1. As usual for this type of model the production function F_f is given by

$$F_f(\{y_{jt}\}) = \left(\int_{j=0}^1 y_{jt}^{\frac{\sigma-1}{\sigma}} dj \right)^{\frac{\sigma}{\sigma-1}}$$

where the parameter σ measures the elasticity of substitution between intermediate goods. Taking prices $\{p_{jt}\}$ of the intermediates goods as given, the standard inverse demand curve for input j follows from the final goods firm's

maximization problem as

$$p_{jt}(q_{jt}) = \left(\frac{Y_t}{q_{jt}} \right)^{\frac{1}{\sigma}}.$$

2.3 Banks

The banking sector constitutes the only source of financing for the interediate goods firms. By assumption, firms can not raise equity and instead have to rely on external finance. These loans are risky, because there is a chance that a firm will default and trigger a state verification process which generates additional costs to the lender. Because of this friction, which will be discussed in detail in the next section, each individual loan may be repaid to the bank either fully, partially, or not at all.

There is a large number of atomistic, perfectly competitive banks that each hold a fully diversified portfolio of loans and deposits — this is equivalent to modeling the sector by a representative bank operating under a zero-expected-profit condition. The bank receives deposits S_t from the households on which it pays the risk-free interest rate R_t determined this period. On the other side of the budget constraint stand loans to intermediate goods firms, and their aggregate repayments. The latter include the sum of all full and partial repayments minus the state-verification costs that are due in the case of default. Finally, while the bank absorbs all idiosynratic firm risk through the law of large numbers, it cannot hedge against the risk of aggregate shocks which influence the default rate. This risk is passed on to the bank's owners, the households, via positive or negative profits. These accounting profits are given by

$$\pi_t^b = S_t + \text{Rep}_t - R_{t-1}S_{t-1} - \text{Loans}_t$$

where Repayments and Loans will be defined as aggregates below.

By the assumption of perfect competition embedded in the representative bank, the risk-neutral bank takes loan rates that it can charge firms as given. In particular, if a firm approaches a bank requesting a loan of size b , the bank's optimization problem is simply to decide whether to grant the loan or to walk away. As discussed below, this implies that in equilibrium the expected return on a loan to a firm is just equal to the risk-free interest rate.

2.4 Intermediate goods producers

The intermediate goods sector is where the two main frictions, nonconvex adjustment costs and borrowing under costly state verification, are built into the model. Firms accumulate capital and are subject to idiosyncratic and aggregate productivity shocks. They have the option of borrowing from a bank. The interest rate of the loan is firm-specific; that is the firm borrows against its future profits and capital stock, and the loan rate reflects the size of that collateral and the risk of default.

Production and adjustment cost Production is assumed to follow the Cobb-Douglas form and is given by

$$f(z, A, k, l) = zAk^\alpha l^{1-\alpha}.$$

Capital k is quasi-fixed and can be changed only with a lag of one period via investment whereas labor input l is fully flexible and can be hired as needed every period.² Physical productivity evolves exogenously according to an idiosyncratic and an aggregate shock process $\{z_{jt}\}_t$ and $\{A_t\}_t$, respectively. Both will be specified as log-normal AR-1 processes. The firm takes into account the demand function for its goods from the final goods sector.

Next period's capital k' is determined today by the standard accumulation rule

$$k' = (1 - \delta)k + I,$$

where I represents investment in the current period. Finally, as far as production technology goes, there is a fixed cost ϕ to capital adjustment. I assume the adjustment cost has to be paid if $k' \neq (1 - g)k$ where $g \ll \delta$ is a small positive number (and matters only to avoid an indeterminate non-stochastic steady state). Intuitively, if the firm wants to change its capital stock at all (besides the small change through g), the fixed cost is due. This assumption aims to make capital adjustment upward and downward approximately symmetric.

²In order to keep the model's state space small, I do not consider the case of additional adjustment frictions in the labor input of the firm. This is despite the fact that, for example in Bloom et al. (2012) it is actually the labor friction that has the largest impact on aggregate dynamics. Allowing for stronger frictions on input factors would most likely strengthen the results here, as firms will find it even harder to achieve their desired revenue product generating more dispersion.

Borrowing problem In order to introduce a financial dimension to the firm problem I assume that for undertaking investment the firm has to rely on external finance. This loan market is subject to costly state verification in case of default, a commonly used type of financial friction³. The particular setup is taken from Gilchrist et al. (2013). Their version of the borrowing problem allows for substantial firm heterogeneity (in particular admitting persistent differences in physical productivity), and yields endogenous credit spreads while remaining computationally tractable.

Firms can sell bonds b at a price of q that entitle the buyer to 1 unit of tomorrow's consumption good. The firm-specific price of the bond determines the implicit interest rate the firm faces, and will be determined by the marginal risk of default. Default risk exists because, by assumption, there is a minimum amount of net worth denominated \bar{n} , that is not enforceable for repayment. In other words, lenders can enforce repayment only up to the point where the borrower's net worth is just \bar{n} , so that if $b > n - \bar{n}$ default occurs (where n denotes the firm's net worth just before repayment is due). If the firm is unable to repay the lender in full, default is partial if $n > \bar{n}$ and total if $n \leq \bar{n}$.

Besides the loss of the principal, there is an additional state-verification cost that the lender has to bear in case of default which represents his cost of determining the borrower's remaining net worth. As a simplifying assumption, the default cost is assumed to be proportional to the size of the original loan. A further assumption is that the cost is always due in case of default — in particular, there is no decision to be made by the lender about whether it might be worth to walk away from the loan entirely and avoid paying the state-verification cost.

Timing The period starts with draws of the aggregate productivity A_t and the set of idiosyncratic productivity states $\{z_{it}\}$. With the resolution of aggregate uncertainty all aggregate variables in period t are determined, and the economy's aggregate state is given by $\Sigma \equiv (A, \mu)$. The function μ is the density of the distribution of firms over (z, k, b) . Next, intermediate goods firms hire the optimal amount of labor l_{it} given their capital stock k_{it} and productivity z_{it} , regardless of their level of debt. After production occurs at the intermediate and final goods level, intermediate goods firms consider their net worth $n_{it} = (1 - \delta)k_{it} + \pi(z_{it}, k_{it}, \Sigma_t)$ composed of undepreciated

³See Carlstrom and Fuerst (1997) and Bernanke et al. (1999) for seminal papers.

capital and revenue net of wages. As described in the previous section a firm then defaults if $n_{it} - \bar{n} < b_{it}$ or otherwise repays its debt b_{it} in full. In either case the firm is left with an end-of-period net worth \tilde{n}_{it} , and since there is no ‘punishment’ for default in terms of an exclusion from credit markets, the firm’s state after production is captured by the end-of-period net worth together with the undepreciated capital stock and physical productivity. The firm now has to pick its desired levels of investment and borrowing which, through the budget constraint, also determines dividend payments to its owners.

Optimal choices and firm value With labor hired on the spot it is straightforward to write down a function for revenue net of wages as

$$\begin{aligned}\pi(z, A, k) &\equiv \max_l f(z, A, k, l)^{-\frac{\sigma-1}{\sigma}} Y^{\frac{1}{\sigma}} - wl \\ &= C(w) [(Azk^\alpha)^{\sigma-1} Y]^{\frac{1}{1-\alpha+\alpha\sigma}}.\end{aligned}$$

where $C(w) \equiv \frac{1-\alpha+\alpha\sigma}{\sigma} \left[\frac{(\sigma-1)(1-\alpha)}{\sigma w} \right]^{\frac{(\sigma-1)(1-\alpha)}{1-\alpha+\alpha\sigma}}$. The maximizing labor input follows as

$$l(z, A, k) = \left\{ \left[\frac{(\sigma-1)(1-\alpha)}{\sigma w} \right]^\sigma (Azk^\alpha)^{\sigma-1} Y \right\}^{\frac{1}{1-\alpha+\alpha\sigma}}. \quad (2)$$

Of course with capital quasi-fixed, with its choice of labor the firm simultaneously picks production y , price p and its revenue product pz .

The firm-specific bond price follows directly as a no-arbitrage constraint from the bank’s zero-profit condition. In expectation, a loan to the firm will yield the risk-neutral bank the same return as an investment at the safe interest rate R_t . The bond price as a function of the loan size, next period’s capital, and this period’s firm productivity is then given by

$$\begin{aligned}q_b(z, k', b', \Sigma) &= \frac{1}{R} E_A \left\{ 1 - F(\bar{\epsilon}) (1 + \chi) + [F(\bar{\epsilon}) - F(\underline{\epsilon})] \left[\frac{(1-\delta)k - \bar{n}}{b'} \right] + \right. \\ &\quad \left. + \frac{C(w') [(A'z^{\rho_z} k'^\alpha)^{\sigma-1} Y']^{\frac{1}{1-\alpha+\alpha\sigma}} e^{\sigma_\nu^2/2}}{b'} \times \right. \\ &\quad \left. \times \left[\operatorname{erf} \left(\frac{\sigma_\nu^2 - \log(\underline{\nu})}{\sqrt{2}\sigma_\nu} \right) - \operatorname{erf} \left(\frac{\sigma_\nu^2 - \log(\bar{\nu})}{\sqrt{2}\sigma_\nu} \right) \right] \right\}. \quad (3)\end{aligned}$$

where $\bar{\epsilon}$, $\underline{\epsilon}$, $\bar{\nu}$ and $\underline{\nu}$ are cutoff values for productivity shocks that lead to partial and total default, respectively, and erf denotes the error function. For details see the appendix and Gilchrist et al. (2013).

The firm's optimal borrowing and investment choices are then represented by the firm's value function. The firm's value depends on the idiosyncratic state variables (z, k, b) and on the aggregate state Σ , and it chooses next period's capital k' and debt b' . The choice of capital can be thought of as two sequential decisions: A discrete one whether to adjust at all, and, if the answer is yes, an unconstrained choice about the level of investment. Representing the investment decision as this two-step process makes it easy to write down the value function as the maximum of the value of adjusting and the value of not adjusting capital, denoted by V_a and V_n respectively. Denote the binary decision whether to adjust $\mathbb{1}_{\text{adj}}(z, k, b, \Sigma)$.

$$V(z, k, b, \Sigma) = \max_{\mathbb{1}_{\text{adj}} \in \{0,1\}} \mathbb{1}_{\text{adj}} V_a(z, \tilde{n}(z, k, b, \Sigma), \Sigma) + [1 - \mathbb{1}_{\text{adj}}] V_n(z, k, b, \Sigma) \quad (4)$$

If the firm chooses adjustment and the fixed cost has been paid its current capital stock relevant only in so far as it contributes to the firm's net worth. Its value is therefore a function of just productivity and net worth

$$V_a(z, \tilde{n}, \Sigma) = \max_{k', b'} \tilde{n} + q(z, k', b', \Sigma) b' - k' + E_{z, \Sigma} [d(\Sigma', \Sigma) V(z', k', b', \Sigma')]. \quad (5)$$

The policy functions for next period's capital and debt conditional on adjustment are denoted $k'_a(z, k, b, \Sigma)$ and $b'_a(z, k, b, \Sigma)$, respectively.

On the other hand if the firm has made the decision to save the adjustment cost and stay at its current level of capital, then both k and b remain relevant state variables. while the firm's choice is now only over the tradeoff debt/dividend payments.

$$V_n(z, k, b, \Sigma) = \max_{b'} \tilde{n} + (z, k, b, \Sigma) q(z, k', b', \Sigma) b' - (1 - g)k + E_{z, \Sigma} [d(\Sigma', \Sigma) V(z', (1 - g)k, b', \Sigma')] \quad (6)$$

Denote the firm's bond supply choice conditional on not adjusting the capital stock $b'_n(z, k, b, \Sigma)$. The corresponding policy function for capital is trivially given as $k'_n(z, k, b, \Sigma) = (1 - g)k$. The function $d(S', \Sigma)$ represents the household's stochastic discount factor.

Finally, from (4)-(6) the unconditional policy functions follow as $k'(z, k, b, \Sigma) = \mathbb{1}_{\text{adj}}(z, k, b, \Sigma) k'_a(z, k, b, \Sigma) + [1 - \mathbb{1}_{\text{adj}}(z, k, b, \Sigma)] k'_n(z, k, b, \Sigma)$ and $b'(z, k, b, \Sigma) = \mathbb{1}_{\text{adj}}(z, k, b, \Sigma) b'_a(z, k, b, \Sigma) + [1 - \mathbb{1}_{\text{adj}}(z, k, b, \Sigma)] b'_n(z, k, b, \Sigma)$.

2.5 Aggregation and Recursive Equilibrium

There are three aggregate markets in the economy: the market for final goods, the market for labor, and the market for loans. Aggregate supply and demand on each of these markets are:

$$L^s = 1$$

$$L^d = \int_{(z,k,b)} l(z, k, \Sigma) d\mu(z, k, b) \quad (7)$$

$$Y^s = \left[\int_{(z,k,b)} (y^s(z, k, \Sigma))^{\frac{\sigma-1}{\sigma}} d\mu(z, k, b) \right]^{\frac{\sigma}{\sigma-1}} \quad (8)$$

$$Y^d = C + I + \text{Mon} + \Phi$$

$$\text{Loan}^s = S$$

$$\text{Loan}^d = \int_{(z,k,b)} q_b(z, k'(z, k, b, \Sigma), b'(z, k, b, \Sigma), \Sigma) \times \quad (9)$$

$$b'(z, k, b, \Sigma) d\mu(z, k, b).$$

In equation (8), aggregate monitoring cost and aggregate adjustment cost paid are defined as

$$\text{Mon} = \chi \int_{(z,k,b)} b \mathbb{1}_{z < \bar{z}(k,b,\Sigma)} d\mu(z, k, b)$$

and

$$\Phi = \phi \int_{(z,k,b)} \mathbb{1}_{\text{adj}}(z, k, b, \Sigma) d\mu(z, k, b),$$

respectively.

A recursive equilibrium is composed of a set of value functions, policy functions, pricing functions, as well as a law of motion that are consistent with agent optimization, market clearing, and rational expectations.

Specifically,

- $S(\Sigma)$, $C(\Sigma)$ are the policy functions derived from the household's problem (1) and $U^*(\Sigma)$ is maximized lifetime utility; $V(z, k, b, \Sigma)$, $V_a(z, \tilde{n}, \Sigma)$, $V_n(z, k, b, \Sigma)$ are an intermediate goods firm's value functions (4)-(6) and $l(z, k, b, \Sigma)$, $k'(z, k, b, \Sigma)$, $b'(z, k, b, \Sigma)$ are the corresponding policy functions;

- Wages are given by $w(\Sigma)$, interest rates by $R(\Sigma)$, bond prices by $q_b(z, k', b', \Sigma)$
- Markets in equations (7) - (9) clear, i.e. $L^s = L^d$, $Y^s = Y^d$, and $\text{Loan}^s = \text{Loan}^d$.
- The law of motion $\mu'(\mu, A)$ for the evolution of distribution over (z, k, b) follows from the firms' policy functions and the exogenous process for z .

3 Determinants of the profitability distribution

The aim of this section is to give an intuition for how cross sectional dispersion in revenue productivity can arise when one includes firm pricing. I will first discuss the general TFPR distribution and how it interacts with the real friction, and then turn to the financial friction.

Non-convex adjustment costs The general mechanism is that, if there are frictions, firms' investment needs not react symmetrically in response to shocks. Considering the example of fixed adjustment costs, physically productive firms are more likely to adjust, since for these firms the relative benefit of adjustment is larger. This in turn will reset their revenue product towards the level chosen if there were no adjustment costs. The non-adjusters on the other hand will tend to be physically unproductive firms. Not adjusting fully to a, say, negative shock in physical productivity levels will leave these firms with too much factor inputs and causing them to overproduce relative to the frictionless case. This drives down the price and therefore revenue productivity. Hence the left tail of the TFPR distribution becomes wider in a recession, and conversely it becomes shorter in a boom.

Turning to the specifics of the model, we can use the fact that labor is a fully flexible input factor and is chosen optimally as given by (2). One can then write down an analytic expression for revenue productivity pAz as being proportional the productivity states and capital:

$$pAz \sim (Az)^{\frac{\alpha(\sigma-1)}{1+\alpha(\sigma-1)}} k^{-\frac{\alpha}{1+\alpha(\sigma-1)}}. \quad (10)$$

Given capital, TFPR increases in physical productivity; and given productivity, TFPR decreases in the amount of capital held by the firm.

It can be informative to consider this relationship under some extreme cases. First, if A and z are fixed (the variance of the respective shocks is set to 0) and there are no adjustment costs, then firms will choose the same k equalizing their revenue product. This also corresponds to the setup often used in models where production is linear in the fully flexible factor labor and the TFPR distribution is degenerate (case a in figure 1). Second, if there is variation in A and z while still holding adjustment costs at 0, then using the firm's profit function it can be shown that firms pick k in a way that equalizes their expected revenue product. This means that any dispersion in TFPR will be caused by contemporaneous shocks to idiosyncratic productivity z , implying a symmetric profitability distribution (case b in figure 1). It is important to note that the firms' ranks in the TFP distribution are preserved in the TFPR distribution: a higher physical productivity implies a higher revenue productivity. Additionally, for a given variance of shocks to z , dispersion is increasing in α (the more the quasi-fixed factor is used in production, the more dispersion) and increasing in σ (the higher the price elasticity of demand, the more dispersion).

Third, say the level of idiosyncratic TFP z displays some persistence and that there are positive adjustment costs. As discussed above, assume they are such that firms with high z find it profitable to adjust their capital stock in response to a given change in A , while firms with low z do not. Then next period's capital is going to be correlated with productivity: For example, following a positive shock to A , physically unproductive firms will have relatively little capital, and therefore a high revenue product. Underlying this is of course the fact that these firms do not produce as much as they would have if capital adjustment were free, and therefore drive up the price of their good.

Financial friction As just outlined, the effect of the aggregate state on the profitability of firms depends on the response of investment and the associated marginal costs of production in the following period. The question is therefore whether a recession more strongly affects productive firms or unproductive firms. Under the borrowing friction investment is influenced by the firm-specific loan price; so what matters is whether unproductive firms' loan price reacts more to a change in A than a productive firm's price.




Figure 1: Profitability distribution in special cases

Consider a case without adjustment costs: A firm with low expected z will invest little, therefore requiring only a small loan. Of course its expected revenues are also small. Together these effects determine the firm’s default risk and its loan price.

If a recession drives up productive firms’ credit spreads more than that of unproductive firms, productive firms will have less capital and higher prices next period, increasing dispersion in profitability. On the other hand it is possible that a recession mainly affects unproductive firms default risk. In this case by the same argument the TFPR distribution would become narrower making its dispersion procyclical. Analytically it is not clear which of the cases applies, and in general this will depend on the specific model environment. The calibration used below shows countercyclical TFPR dispersion when only the financial friction is active, indicating that the effect on productive firms dominates.

4 Simulation and Results

4.1 Simulation

Calibration For most parameter values I take existing estimates from the literature as a baseline case. In particular, I draw on Khan and Thomas (2008)’s model of fixed adjustment costs for a range of technology parameters, as do Bloom et al. (2012). Table 1 summarizes the baseline parameter choices.

The model is calibrated to an annual frequency in order to facilitate computation. The household’s rate of time preference β is set to 0.96 generating an average annual interest rate of around 4%. The household is given an elasticity of intertemporal substitution of 1 corresponding to a log utility function. As in Khan and Thomas (2008). I set the persistence of the aggregate productivity process $\rho_A = 0.86$. For the standard deviation of shocks to A I choose a value of $\sigma_A = 0.027$ which targets the detrended time-series volatility of annual output in the United States of 2.2%. The elasticity of substitution in the final goods sector is chosen to be $\sigma = 4$ implying a price elasticity of demand for intermediate goods of -4 . While this is on the lower end of estimates in the literature it corresponds to the value implied by Bloom et al. (2012)’s choice of decreasing returns to scale on the firm level.⁴

⁴Bloom et al. point out that the source of decreasing returns of their firms’ production

Parameter		Value	Target / Source
Household and final goods sector			
Rate of time preference	β	0.96	Real interest rate of 4%
EIS	ψ	1	Log utility
Price elasticity of demand	σ	4	Bloom et al. (2012)
Technology			
Persistence of aggregate TFP	ρ_A	0.86	Khan and Thomas (2008)
SD of innovations to A	σ_A	0.027	Volatility of output 2.2%
Persistence of idiosyncratic TFP	ρ_z	0.86	Khan and Thomas (2008)
SD of innovations to z	σ_z	0.022	Khan and Thomas (2008)
Capital share	α	0.2	Labor income 60%
Depreciation rate	δ	0.1	Standard value
Frictions			
Adjustment cost	ϕ	0.04	2.6% spike adjusters
Depreciation when not adjusting	g	0.01	Small value
Verification cost in default	χ	0.10	Gilchrist et al. (2013)
Firms' dividend preference	γ	0.05	Gilchrist et al. (2013)
Borrower's protected net worth	\bar{n}	0.0	Gilchrist et al. (2013)

Table 1: Parameter values

Turning to the intermediate goods firms, I again follow Khan and Thomas (2008) in choosing a persistence for the idiosyncratic productivity process that is equal to the one of the aggregate process and set $\rho_z = 0.86$. There is a range of estimates for the variance of innovations to firm TFP in the literature. For example, Khan and Thomas (2008) use a σ_z of 2.2%, and Bloom et al. (2012) use 4%. As will be seen below, the relative TFPR dispersion over the cycle in this model is somewhat sensitive to the choice of σ_z . For now I stick to the calibration by Khan and Thomas (2008) and will discuss higher values of σ_z below. As is standard in the literature, the annual depreciation rate is set to $\delta = 0.10$. For the parameter g that describes capital shrinkage in the case of non-adjustment I ad-hoc pick 0.01 as a ‘small value’. The elasticity of intermediate firm output with respect to capital is set to $\alpha = 0.2$. Given demand elasticity σ above, this value matches a 60% labor share of output as α and σ determine the monopolistic firms’ profits jointly.

What remains is to set the parameters governing the model’s frictions. For the financial friction I follow Gilchrist et al. (2013) in setting the default cost to $\chi = 0.10$, the preference rate for dividend payments to $\gamma = 0.05$ and the level of protected net worth to $\bar{n} = 0$. The final parameter that needs to be chosen is the fixed adjustment cost ϕ . These costs are usually calibrated to match a moment of the distribution of investment rates; oftentimes this is the share of ‘spike adjusters’ whose investment rates exceed 20%. For the US this value is around 15% for equipment capital and around 10% for all types of capital (including structures). These numbers however include replacement investment. Since in the model depreciation is paid by all firms every period (unless they adjust downward) I target a fraction of adjusters significantly lower between 2% and 3%. For this parameter, too, I consider alternative specifications below.

Numerical approximation I solve the model approximately using standard techniques for the computation of heterogeneous agent DSGE models. I briefly outline the procedure here and put additional description into the appendix.

The state space is discretized using a uniformly spaced grid for the endogenous variables k and b . The exogenous state variables z and A and their evolution over time are approximated as discrete Markov chains using Tau-

function could be derived from monopolistic competition.

chen’s method. I use a particularly dense grid for z since the realization of the idiosyncratic shock is essential for the default decision. Aggregate TFP A is modeled as a process over three discrete states standing for recessions, normal times, and booms. The distribution μ of capital and debt among firms as endogenous aggregate variable is approximated using a grid for the first moment of the marginal distribution of capital, K .

Agents in the economy are assumed to use the aggregate state (A, K) to forecast other aggregate variables. Making household’s marginal utility the numéraire as in Khan and Thomas (2008), the variables that need to be forecast are next period’s capital stock K' , the wage w , the price of the final good (in utils), as well as the output of the final good Y . The latter is relevant because the intermediate goods firms’ output decisions and aggregate output are interdependent through final goods firm’s demand function. With agents thus taking all aggregates as given functions of the state variables, I use value function iteration to derive the monopolists’ policy functions for k' and b' . The economy is then simulated over a long time horizon. This procedure is then repeated, updating the forecast rules iteratively until forecasts match simulated prices and quantities as closely as possible.

4.2 Results

This section discusses the model economy’s response to shocks in three ways. First, I compare average booms and recessions to get an idea of the cyclicity of central variables. Then I graph impulse response functions as a way to capture the response of the economy to an isolated shock. Finally I look at unconditional volatilities and the correlations of aggregates over the cycle to assess the magnitude and cyclicity of the measures of interest quantitatively.

Average booms and recessions I now consider fluctuations in the level of aggregate TFP. With TFP A modeled as a three-point process, table 2 presents model statistics for the recessionary and expansionary state of the economy relative to the ‘neutral’ state in which A is normalized to 1, respectively.

TFPR dispersion is measured as the coefficient of variation of cross-sectional revenue productivity.⁵ The mean credit spread is the average of

⁵That is, $sd(p_i A z_i) / mean(p_i A z_i)$. The coefficient of variation is useful in this context

	Recession	Normal times	Boom
Mean of . . .			
TFP A_t	-3.92%	1	3.92%
Output Y_t	-2.44%	0.99	2.61%
TFPR dispersion	3.7%	0.019	-4.4%
Mean credit spread	22 bps	88 bps	-24 bps
Credit spread dispersion	21.5%	0.016	-17.9%

Table 2: Recession vs boom

the differences of firm-specific interest rates as implied by their bond price and the risk-free interest rate in the same period. Finally, dispersion of credit spreads is again measured as the cross-sectional coefficient of variation.

These statistics were generated by simulating the full economy, and then simply averaging over all periods in which A was, say, low. The table qualitatively confirms that the model displays countercyclical TFPR dispersion and worsening credit conditions: Revenue productivity is on average 9.7% more dispersed comparing recessions to booms, and credit spreads are 46 basis points higher as well as 39% percent more dispersed. This table does not tell us much about the timing of the responses to shocks, nor do we readily observe A in the real world. Next I therefore consider impulse response functions, and will afterwards compare the unconditional volatilities of the aggregates to the data.

Impulse response functions I obtain simulated impulse response functions by manually holding the aggregate shock A at its long-run mean of 1 and simulating the economy for enough periods that all aggregate variables become approximately constant. The firms are still being hit by idiosyncratic shocks, and in contrast to the non-stochastic steady state, they expect movements in A according to its regular distribution (which simply don't materialize). A large number of economies is then seeded to this 'neutral' aggregate state as period 0. In period 1 all economies are hit by a negative shock, i.e. A is set to its low value. The economies are then simulated forward, differing in how long they remain in the recessionary state. Specifically, for each economy, the chance of remaining in the low state for one more period is given by the transition probabilities of A , and once an economy gets a draw

because it is a scale invariant dispersion measure.




Figure 2: Impulse response functions to negative shock in A

that would put it into the ‘neutral’ or ‘high’ TFP state, its A is set back to 1 for the remainder of the simulation. This part of the procedure ensures that all aggregate variables eventually return to their pre-impact levels and there is no sampling randomness. The impulse response functions for several variables are shown in figure 2.

The dispersion measures react as implied by theory. Revenue productivity dispersion increases by around 6%, credit spreads by around 50 bps and credit spread dispersion by 50%. Because the profitability distribution is only affected through the capital stock, TFPR dispersion does not react on impact, but increases as firms carry out their different investment policies in response to the shock. Since loan prices are determined in the period in which investment takes place, credit spreads and credit spread dispersion react immediately, shooting up and flattening back out quickly.

Time series volatility I now compare the time-series variance of the model-simulated aggregates to the ones measured in the data. Table 3 displays second moments of the dispersion measures as well as the main business cycle aggregates.⁶

⁶In case of the dispersion measures (TFPR dispersion and credit spread dispersion), the time-series variance is the longitudinal second moment of a cross-sectional second moment.

	Model	Data
Standard deviation of		
TFPR dispersion	3.24%	4.7%
Credit spreads	36 bps	92 bps
Credit spread dispersion	24.6%	?
Y	2.2%	2.2%
C	2.0%	1.8%
I	6.0%	8.3%
Cyclicality		
$corr(Y, \text{TFPR disp})$	-0.32	-0.40
$corr(Y, \text{credit spreads})$	-0.53	-0.46
$corr(Y, \text{CS disp})$	-0.60	-0.25

Table 3: Time-series volatility and cyclicality

The empirical moments in table 3 are compiled from several sources. The data on revenue productivity dispersion is calculated using Kehrig (2011)'s annual time series on the median sectoral coefficient of variation of TFPR from 1972 to 2010.⁷ Similarly, the information on the level of the average credit spread uses the GZ-credit spread index released by Gilchrist and Zakrajšek (2011).⁸ This time series ranges from 1973 to 2012. While information on the volatility of the GZ-credit spread dispersion is not readily available, Gilchrist and Zakrajšek (2011) report the correlation of output with both mean credit spreads and the cross-sectional standard deviation of credit spreads. Finally, the numbers for output, consumption and investment are derived from FRED⁹ using HP(100)-filtered data from 1950 to 2013. I apply the same HP(100)-filter to the model-generated data for output, consumption, investment, and TFPR dispersion.

Consumption is a bit too volatile in the model, at the expense of an undershoot in the variance of investment (with the volatility of Y being targeted in the calibration). Looking at profitability dispersion, I find that the model can explain a little more than two thirds of the empirically observed volatility (3.24% versus 4.7%), and exhibits a similar correlation coefficient with output as in the data (-0.32 versus the observed -0.40). While the

⁷Kehrig (2011)'s data is HP(100)-filtered. Data available at <https://sites.google.com/site/matthiaskehrig/research>

⁸Data available at <http://people.bu.edu/sgilchri/Data/data.htm>

⁹<http://research.stlouisfed.org/fred2/>

model does not generate strong movements in the level of credit spreads, it matches the empirical cyclicalities fairly closely. Moreover the model delivers significant swings in credit spread dispersion. I interpret this as indication that the model’s first-moment shocks generate the observed negative relationship between recessions and firms’ financial health. Moreover, in order to generate a larger amplitude in credit spreads additional financial frictions may be needed. For example, Gilchrist et al. (2013) in their uncertainty-shock model supplement the borrowing friction with frictions on raising and lowering equity (issuing stocks and paying dividends, respectively).

Contribution of individual frictions Figure 3 displays the impact of frictions separately. For this exercise, I compare four simulations: First, the model in the baseline calibration as before; second, a model with only adjustment costs where firms can borrow on a frictionless credit market (there is no default, i.e. $\bar{n} = -\infty$); third, a calibration with only financial frictions (adjustment costs $\phi = 0$); and fourth, a simple model of monopolistic competition without real nor financial friction.

The most notable feature is that the paths of real variables between the case of both frictions and only adjustment costs are very similar. In particular, as long as the adjustment cost is present the responses generate a significant increase in profitability dispersion. The financial friction on its own, however, only leads to a small increase. Moreover, including adjustment costs raises the baseline level of credit spreads and increases their response to a negative shock compared to the case of the borrowing friction alone.

Table 4 confirms this finding: Adding the borrowing friction to a model of non-convex adjustment costs does not change the implications for the profitability distribution much. Conversely the adjustment cost improves the fit compared to only the financial friction.¹⁰ Notably, the borrowing friction itself does not generate the observed countercyclicality of credit spreads, a point also found by Gilchrist et al. (2013).

Sensitivity to parameters I consider alternative specifications for two central parameters; namely the adjustment cost ϕ as well as the standard

¹⁰The excessive TFPR dispersion appears to be result of a nonlinearity: In the simulation, a more than proportionate jump in credit spreads occurs when the aggregate state quickly switches from boom to recession without spending much time in the intermediate state. Without these episodes volatility in TFPR dispersion is small.




Figure 3: Impulse response functions of frictions separately

Friction	Both	Financial	Real	None
Standard deviation of				
TFPR dispersion	3.24%	8.04%	3.16%	0
Credit spreads	36 bps	17 bps	-	-
Credit spread dispersion	24.6%	45%	-	-
Cyclicality				
$corr(Y, \text{TFPR disp})$	-0.32	-0.17	-0.31	0
$corr(Y, \text{credit spreads})$	-0.53	-0.06	-	-
$corr(Y, \text{CS disp})$	-0.60	-0.41	-	-

Table 4: Impact of individual frictions

deviation of idiosyncratic productivity shocks σ_z . Table 5 contains results for a few different parameter values.

Overall, the model economy does not respond too strongly to changes in the adjustment cost. Lower ϕ mainly causes a larger share of firms to adjust each period, generating higher volatility in aggregate investment. Volatility of the profitability distribution reacts only mildly in a non-monotonic way except for very low values of the parameter.

Increasing σ_z reduces the magnitude of swings in TFPR dispersion. Intuitively, a high variance of firm-specific shocks makes the aggregate state less important to the firm in making its investment decisions. These decisions are now mainly driven by the firm's idiosyncratic state. For example, when using the parameter value from Bloom et al. (2012) with $\sigma_z = 0.04$, the standard deviation of TFPR dispersion is 2.7% or 55% of the empirical estimate, and decreases further for higher σ_z .

5 Conclusion

This paper shows that profitability dispersion among firms can arise endogenously in a response to a change in aggregate production levels. Therefore recessions can look like times of increased firm risk even when underlying productivity risk is constant over the cycle. In general this result comes from heterogeneity in how firms' pricing responds to an aggregate shock. The particular structure chosen in this paper demonstrates that this differential response can result under the same setup used in models of uncertainty shocks, employing non-convex adjustment costs and costly financial inter-

	<i>Time-series volatility (sd)</i>					<i>(mean)</i>
	TFPR disp	Cred. spread	Y	C	I	Frac. adjusters
Baseline	3.2%	36bps	2.2%	2.0%	6.0%	2.5%
ϕ						
0.03	3.8%	35bps	2.2%	2.0%	5.7%	3.0%
0.02	3.6%	32bps	2.2%	1.9%	6.3%	4.3%
0.015	2.6%	35bps	2.2%	1.8%	8.1%	8.4%
0.01	7.8%	26bps	2.2%	2.0%	14.7%	75%
$\sigma_{\mathbf{z}}$						
0.04	2.7%	22bps	2.2%	1.8%	6.8%	2.8%
0.08	0.8%	22bps	2.2%	1.5%	9.0%	7.3%

Table 5: Alternative parameter values

mediation. The baseline calibration generates two-thirds of the empirical observed cyclicity in levels of revenue productivity as well as procyclicality in the financial health of firms. These results suggest that accounting for the difference between productivity and profitability is relevant when assessing firm risk over the business cycle and calibrating uncertainty shocks to cross sectional moments.

References

- Arellano, C., Y. Bai, and P. Kehoe (2012). Financial frictions and fluctuations in volatility. *Federal Reserve Bank of Minneapolis Research Department Staff Report* (466).
- Bachmann, R. and C. Bayer (2013a). Investment dispersion and the business cycle. *American Economic Review* 104(4), 1392–1416.
- Bachmann, R. and C. Bayer (2013b, June). 'Wait-and-See' business cycles? *Journal of Monetary Economics* 60(6), 704–719.
- Bachmann, R. and G. Moscarini (2012). Business cycles and endogenous uncertainty. *manuscript, Yale University*.
- Bernanke, B., M. Gertler, and S. Gilchrist (1999). The financial accelerator in a quantitative business cycle framework. *Handbook of macroeconomics* 1.
- Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica* 77(3), 623–685.
- Bloom, N., M. Floetotto, N. Jaimovich, I. Saporta-Eksten, and S. J. Terry (2012). Really uncertain business cycles. *NBER* (w18245).
- Carlstrom, C. T. and T. S. Fuerst (1997). Agency costs, net worth, and business fluctuations: A computable general equilibrium analysis. *American Economic Review*, 893–910.
- Christiano, L., R. Motto, and M. Rostagno (2013). Risk shocks. *NBER* (w18682).
- Cui, W. (2013). Delayed Capital Reallocation. *2013 Meeting Papers Society fo*(No. 500).
- Eisfeldt, A. L. and A. a. Rampini (2006, April). Capital reallocation and liquidity. *Journal of Monetary Economics* 53(3), 369–399.
- Foster, L., J. Haltiwanger, and C. Syverson (2005). Reallocation, firm turnover, and efficiency: Selection on productivity or profitability? *NBER* (w11555).

- Gilchrist, S., J. Sim, and E. Zakrajšek (2013). Uncertainty, financial frictions, and irreversible investment. *Boston University and Federal Reserve Board, mimeo*.
- Gilchrist, S. and E. Zakrajšek (2011, May). Credit spreads and business cycle fluctuations. *NBER* (No. w17021).
- Kehrig, M. (2011). The cyclicalities of productivity dispersion. *US Census Bureau Center for Economic Studies Paper No. CES-WP-11-15*.
- Khan, A. and J. Thomas (2008). Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics. *Econometrica* 76(2), 395–436.

A Derivation of the firm-specific bond price

As mentioned in section 2.4 firm revenue net of wages is given by

$$\pi(z, k, \Sigma) = \frac{1 - \alpha + \alpha\sigma}{\sigma} \left\{ \left[\frac{(\sigma - 1)(1 - \alpha)}{\sigma w} \right]^{(\sigma - 1)(1 - \alpha)} (Azk^\alpha)^{\sigma - 1} Y \right\}^{\frac{1}{1 - \alpha + \alpha\sigma}}$$

$$C(w) [(Azk^\alpha)^{\sigma - 1} Y]^{\frac{1}{1 - \alpha + \alpha\sigma}}.$$

Then the firm's assets after production in any period are $\pi(z, k, A) + (1 - \delta)k$ whereas its liabilities consist of debt b carried over from last period (although b could theoretically be negative if the firm decided to save). Consequently net worth is $n = \pi(z, k, A) + (1 - \delta)k - b$. The assumption behind the friction is that for exogenous institutional reasons there is a lower bound of net worth \bar{n} which is not enforceable for repayment — so if the firm's net worth would fall below \bar{n} it partially or fully defaults instead.

Given capital, debt and aggregate state this implies a cutoff \bar{z} for the level of idiosyncratic productivity that triggers default, implied by $\bar{n} = (1 - \delta)k' + C(w) [(A\bar{z}k^\alpha)^{\sigma - 1} Y]^{\frac{1}{1 - \alpha + \alpha\sigma}} - b$. Solving for \bar{z} one has

$$\bar{z} = \left(\frac{\bar{n} + b - (1 - \delta)k'}{C(w)} \right)^{\frac{1 - \alpha + \alpha\sigma}{\sigma - 1}} (Ak^\alpha)^{-1} Y^{-\frac{1}{\sigma - 1}}.$$

The autoregressive nature of z in turn determines a cutoff value $\bar{\epsilon}$ for the lognormally distributed shock ϵ given as

$$\rho_z \log z_{-1} + \log \bar{\epsilon} \equiv \log \bar{z}$$

$$\bar{\epsilon} = e^{\log \bar{z} - \rho_z \log z_{-1}},$$

determining the likelihood of default.

The second piece of information needed to determine the risk premium is the fraction of the loan that is recoverable in case of default (i.e. the recovery rate). In default, the lender can claim all but the minimum level \bar{n} of the borrower's assets so that the actual repayment \bar{b} is defined as

$$\bar{b}(z, k, \Sigma) \equiv \max \left\{ (1 - \delta)k + C(w) [(Azk^\alpha)^{\sigma - 1} Y]^{\frac{1}{1 - \alpha + \alpha\sigma}} - \bar{n}, 0 \right\}.$$

Repayments are bounded below by zero. This matters only for the rare case of total default, i.e. a realization of z which is so low that $(1 - \delta)k + \pi(z, k, A) <$

\bar{n} . This inequality implies a second cutoff value defining the threshold for total default

$$\underline{\epsilon} = \left(\frac{\bar{n} - (1 - \delta)k}{C(w)} \right)^{\frac{1-\alpha+\alpha\sigma}{\sigma-1}} [Az^{\rho z}k^{\alpha}]^{-1} Y^{-\frac{1}{\sigma-1}}.$$

There is a default cost χb amounting to a fraction χ of the original loan. So the recovery rate $\tilde{R}(z, k, b, \Sigma)$ can be defined as

$$\tilde{R}(z, k, b, \Sigma) \equiv \frac{\bar{b}(z, k)}{b} - \chi.$$

The price of the bond, then, makes the lender indifferent between lending to a firm that chooses k' and b' and is in state z today, and lending at the risk-free interest rate R .

$$\begin{aligned} q(z, k', b', \Sigma) &= \frac{1}{R} E_A \left[1 + \int_{\epsilon' < \bar{\epsilon}} \tilde{R}(z'(\epsilon'), (1 - \delta)k', b', \Sigma) - 1 dF(\epsilon') \right] \\ &= \frac{1}{R} E_A \left[1 - F(\bar{\epsilon})(1 + \chi) \right. \\ &\quad \left. + \int_{\underline{\epsilon} < \epsilon' < \bar{\epsilon}} \frac{(1 - \delta)k' + C(w') [(A'z'k'^{\alpha})^{\sigma-1} Y']^{\frac{1}{1-\alpha+\alpha\sigma}} - \bar{n}}{b'} dF(\epsilon') \right] \\ &= \frac{1}{R} E_A \left[1 - F(\bar{\epsilon})(1 + \chi) + [F(\bar{\epsilon}) - F(\underline{\epsilon})] \left[\frac{(1 - \delta)k - \bar{n}}{b'} \right] \right. \\ &\quad \left. + z^{\frac{\rho z(\sigma-1)}{1-\alpha+\alpha\sigma}} \frac{C(w') [(A'k'^{\alpha})^{\sigma-1} Y']^{\frac{1}{1-\alpha+\alpha\sigma}}}{b'} \int_{\underline{\epsilon} < \epsilon' < \bar{\epsilon}} \epsilon'^{\frac{\sigma-1}{1-\alpha+\alpha\sigma}} dF(\epsilon') \right], \end{aligned}$$

from which equation (3) follows with $\nu \equiv \epsilon^{\frac{\sigma-1}{1-\alpha+\alpha\sigma}}$ and $\bar{\nu}$ and $\underline{\nu}$ defined correspondingly.

B Outline of numerical model solution

Using Khan and Thomas (2008)'s approach of normalizing the price of output with the household's marginal utility of consumption define $P \equiv u'(C)$. With

the household's discount factor given as $d(\Sigma', \Sigma) \equiv \beta u'(C') / u'(C)$ equations (4)-(6) can be rewritten as

$$v(z, k, b, \Sigma) = \max_{\mathbb{1}_{\text{adj}}} \mathbb{1}_{\text{adj}} [v_a(z, -\phi + \tilde{n}, \Sigma)] + (1 - \mathbb{1}_{\text{adj}}) v_n(z, k, b, \Sigma), \quad (11)$$

$$v_a(z, \tilde{n}, \Sigma) = \max_{k', b'} P[\tilde{n} + q(z, k', b', \Sigma) b' - k' - \phi] + \beta E_{\Sigma} [v(z', k', b', \Sigma')], \quad (12)$$

$$v_n(z, k, b, \Sigma) = \max_{b'} P[\tilde{n} + q(z, k, b', \Sigma) b' - k] + \beta E_{\Sigma} [v(z', k, b', \Sigma')]. \quad (13)$$

As outlined in section 4.1, the distribution μ is approximated using a grid over values for aggregate capital K . The algorithm then proceeds as follows:

1. Guess an initial set of log-linear functions $K'(A, K)$, $C(A, K)$, $Y(A, K)$ and $w(A, K)$ which can be represented by their coefficients. Agents use these functions to forecast aggregate variables given the aggregate state.
2. Given the approximating functions the remaining aggregate variables $P(A, K)$ and $R(A, K)$ can be computed as functions of the aggregate state.
3. Derive the firms' value functions (11)-(13) and associated policy functions by value function iteration.
4. Simulate the economy for a large number of periods using the firms' policy functions. To this purpose a sequence of aggregate TFP A is drawn at random for the first simulation and held constant throughout all following iterations. The discretized steady state distribution is simulated forward using the policy functions for next period's endogenous state variables k' and b' and the stochastic transition rule for z' . The first few hundred periods are being discarded, and the aggregate variables of the remaining simulated periods are stored.
5. Regress the stored values for K' , C , Y and w from the simulation onto A and K to obtain new estimates for the coefficients of the log-linear relationships from 1. If the new estimates and the previously used coefficients are close, stop. Otherwise update the coefficients by using a convex combination of the previous ones and the new estimates, and return to 2.

Code for the numerical solution can be downloaded from the author's webpage.