Firm Dispersion and Business Cycles: Estimating Aggregate Shocks Using Panel Data*

Simon Mongey

New York University

Jerome Williams

New York University

January 5, 2017

Click here for most recent version

Abstract

Are fluctuations in firm-level dispersion a cause or an effect of the business cycle? To answer this question, we build a general equilibrium model rich enough to jointly explain characteristics of the firm distribution and the dynamics of macroeconomic aggregates. The model includes frictions that generate movements in dispersion following standard macroeconomic shocks such as aggregate productivity, as well as a direct shock to the dispersion of firm level productivity growth. This type of general equilibrium model with heterogeneous agents and aggregate shocks is computationally difficult to solve, which typically keeps likelihood-based estimation out of reach. We exploit recent advances in solution techniques to obtain a characterization for which estimation is feasible. To answer our question, we estimate the model using time series of macroeconomic aggregates and newly constructed cross-sectional time series, which reflect movements in the firm distribution over time. Now able to account for firm dispersion and the business cycle, we find that (i) standard macroeconomic aggregate shocks explain almost all variation in macroeconomic aggregates, (ii) an uncertainty shock explains almost all variation in firm-level dispersion.

*Corresponding author: jmw612@nyu.edu, 19 W 4th St, New York City, NY 10012. We thank Tim Cogley, Gianluca Violante, Mark Gertler, Jaroslav Borovička, Miguel Faria-e-Castro and Daniel Greenwald for helpful comments and suggestions. We also thank seminar participants at the NYU Macro Student Lunch Seminar and NYU Stern Macroeconomics Lunch Seminar.

1 Introduction

Recently it has been well-documented that measures of firm-level dispersion are cyclical: the cross-sectional dispersion of firm-level output growth, profit growth, employment growth, stock returns and price changes are countercyclical and the dispersion of investment rates is procyclical.[1] It is less clear why this is the case. Increased dispersion of firm-level outcomes may be the result of shocks to the dispersion of productivities, as argued by Bloom (2009), Bloom et al. (2012), Gilchrist et al. (2014) and Christiano et al. (2014), or may be the result of heterogeneous responses across firms to first-moment shocks, as argued by Bachmann and Moscarini (2012). In particular, Bloom et al. (2012) document that firm-level output growth and a revenue-based measure of firm-level total factor productivity (TFP) are countercyclical, and argue that this is consistent with an "uncertainty" shock—a shock to the volatility of idiosyncratic productivities—which simultaneously increases the dispersion of firm-level sales growth growth, and, by inducing wait-and-see effects in firm-level investment policies, lowers investment and output.[2] On the other hand, Bachmann and Bayer (2011) document that the dispersion of investment rates is *procyclical*, seemingly at odds with an uncertainty-driven recession, and argue that the data is best explained by a joint process for TFP and uncertainty which is negatively correlated.[3]

In this paper, we contribute to the question by bringing new information and methods to bear. We study a general equilibrium model rich enough to jointly explain characteristics of firm investment distribution and the dynamics of macroeconomic aggregates. We subject the model to uncertainty shocks as well as to standard macroeconomic shocks used widely in the literature: aggregate productivity, discount rate, financial and labor disutility shocks.[4] This is an environment in which we can ask how important firm-level dispersion shocks are for macroeconomic measures, and how important macroeconomic shocks are for firm-level dispersion measures.

The model includes investment frictions in that firms must pay a fixed cost to invest. This lumpy investment model has been at the core of the recent investment dynamics literature and studied in the papers listed above. Importantly this model has the capacity to deliver

[1] See Bloom (2009), Bloom et al. (2012), Vavra (2014), Basu and Bundick (2012) and Bachmann and Bayer (2011).

[2] Gilchrist et al. (2014) and Christiano et al. (2014) also argue for the presence of an uncertainty shock—they show that financial frictions can amplify shocks to the volatility of firm-level productivities and generate large drops in output.

[3] Bachmann and Moscarini (2012) show that increased dispersion can be the result of an endogenous response of firms to a drop in aggregate TFP.

[4] See Chari et al. (2007) for a description of how aggregate productivity, discount rate and labor disutility shocks map into the three key 'wedges' in an RBC model that are necessary for the explanation of aggregate output, consumption and hours worked

countercyclical movements in the dispersion of firm sales growth—the outcome most often linked to an uncertainty shock—in response to TFP shocks. If firm investment is lumpy enough, then a positive TFP shock can cause firms with lower optimal investment rates to invest, which compresses the distribution of firm growth rates. It also has the capacity to deliver decreases in output following a positive uncertainty shock: the *wait-and-see* effect of Bloom (2009) can lead some firms to hold off on investment as uncertainty increases.[5] This expands the distribution of firm growth rates. Can these theoretical predictions be consistent with the data in a model which allows for both direct shocks to uncertainty and TFP, while matching the time series properties of firm dispersion and output?

To answer this question we estimate the model using Bayesian methods, using as observables times series of both aggregates and moments of the distribution of firms in Compustat data. This allows us to answer two questions in a precise likelihood sense. Which shocks account for fluctuations in moments that reflect the dispersion of firm-level outcomes? Do direct shocks to the dispersion of firm level productivity growth (uncertainty shocks) account for significant movements in output?

We find that almost all of the variation in macroeconomic aggregates is explained by standard macroeconomic shocks rather than the uncertainty shock. The reason for this is relatively simple. The two forces discussed above, which can qualitatively generate first (second) order fluctuations from second (first) order shocks, are only quantitatively large under parameterizations of the model which are inconsistent with the frequency of firm level capital adjustment and dispersion in sales growth found in the data. In particular, although capital adjustment is lumpy, it is not *lumpy enough* for this friction to generate large countercyclical sales growth dispersion from standard macroeconomic shocks. Additionally, firm level productivity must have a very low persistence in order to generate large declines in output following a rise in uncertainty. When firm level productivity is persistent, increases in uncertainty fan out the distribution of realized productivities, leading output to increase as firms grow larger in the right tail of the distribution. That is, the economy benefits from the upside of the spread in the distribution of productivities. A low persistence of firm level productivity reduces these dispersion effects and leaves the negative wait-and-see effect. But such low levels of productivity persistence are rejected by the empirical standard deviation of investment rates and standard deviation of sales growth.

A key advantage of our approach is that it allows us to clearly ask how much of movements in the dispersion in sales growth or investment rates is left to be explained by an uncertainty shock once we have accounted for the role of standard macroeconomic shocks. Rather than studying shocks in isolation, we match the relative volatility and comovement of output,

[5]For a similar effect in a labor markets model see Schaal (2015).

consumption and investment using three macroeconomic shocks—a shock to households' preference for labor supply, a shock to the rate of time preferences, and a shock to aggregate TFP—and treating these series as observables in the likelihood based estimation of the model.[6]

Since the distribution of firms across idiosyncratic capital and productivity is itself a state of the economy, solving for the equilibrium of the model is challenging. The principal method used in the literature for models of this type is the method of Krusell and Smith (1998).[7] This method is computationally intensive (see Algan et al. (2014) for a discussion) and quantitative work that uses this method typically relies on calibration and simulated method of moments. Likelihood-based estimation has so far remained out of reach due to this computational infeasibility. Likelihood-based estimation is also a key tool in accounting for business cycles, which is the exercise at the core of this paper.

We present a novel approach to estimating heterogeneous agents models with aggregate shocks. We exploit the method of Reiter (2009), which combines elements of the perturbation methods and projection methods to produce a first-order approximate solution.[8] Given a first-order approximate solution, we can then apply the standard toolbox of estimation and analysis familiar from the study of linearized DSGE models. A similar approach is followed by Winberry (2016a).

The Reiter (2009) method works as follows. First, we use projection methods to solve for the model's recursive stationary equilibrium in the absence of aggregate shocks. Second, we construct a finite representation of the equilibrium. We construct a finite approximation of the firm distribution and policy functions and a corresponding discretization of the law of motion of the distribution. Third, we compute the solution in the presence of aggregate shocks by perturbing elements of the finite representation of the equilibrium around the steady state of the model which obtains in the absence of aggregate shocks.

An advantage of the Reiter (2009) method is that the model solution is in linear state space form, which lends itself conveniently to estimation. The perturbation techniques that the Reiter method draws on have long been used in the literature to solve representation-agent models, and there is an extensive literature on estimating models in linear state-space form.[9] In particular, we use Bayesian estimation techniques, described in detail in An and Schorfheide (2007). To the best of our knowledge, this is the first full-information estimation

[6]These shock are commonly used in the literature to capture movements in aggregates. See Smets and Wouters (2007), Christiano et al. (2001) and Justiniano et al. (2011).

[7]See Khan and Thomas (2008), Bachmann and Bayer (2011), Bloom et al. (2012) for applications of the Krusell and Smith (1998) method.

[8]See Fernández-Villaverde et al. (2016). For recent applications of this method, see Winberry (2016a), Winberry (2016b), McKay and Reis (2013).

[9]See Fernández-Villaverde et al. (2016) and An and Schorfheide (2007) for summaries.

of a firm dynamics model.

As mentioned above we can use aggregate time series as observables in the estimation, but since our model has predictions for movements in the firm distribution we also include newly-constructed cross-sectional time series. In particular, we include in the set of observable time series the cross-sectional standard deviation of sales growth. This allows us to ask, say, whether an uncertainty shock can jointly explain movements in output and sales growth dispersion *in the presence of* other standard shocks that would move output. We also consider estimating the model with the cross-sectional standard deviation of investment rates as an observable time series.

Before we describe the model, we make a brief detour regarding terminology. In our baseline treatment, we adopt the approach of Bloom (2009) and Gilchrist et al. (2014), in using the expression "uncertainty shock" to describe a shock to the volatility of idiosyncratic shocks. In particular, firms in our model are subject to an autoregressive productivity process with Gaussian innovations. All firms are subject to the same process, although of course, individual firms will have different realizations. An "uncertainty shock" refers to a change in the volatility of innovations to that process. The timing of the shock is also important. We assume that agents observe their current realization of idiosyncratic productivity z and the volatility of the next draw. That is, an uncertainty shock today reflects news about next period's productivity distribution.

"Uncertainty" shock is somewhat of a misnomer, since the shock reflects changes in the actual volatility process affecting firms. This turns out to be important for our results. Some elements of the literature have referred to these as "risk shocks" (see Christiano et al. (2014)). The shock certainly has an uncertainty interpretation, since in the face of a more volatile productivity process, agents are indeed more uncertain about future outcomes, but it also has a second effect, that realized productivities will be more dispersed in future. We use the the terminology "uncertainty shock" for the sake of consistency with previous literature.

The remainder of this paper is organized as follows. Section 2 presents our model of heterogeneous firms facing investment frictions and a number of macroeconomic shocks. Section 3 discusses the solution method, which makes estimation feasible. In Section 4 we discuss the Compustat data used to estimate the model and, in stages, estimated the unknown parameters of the model. Section 5 presents our main results, which consist of forecast error variance decompositions and discussion of the impulse response properties of the model. Section 6 concludes.

2 Model

Our model is similar to that of Bachmann and Bayer (2014) and Bloom et al. (2012): an equilibrium model where firms are heterogeneous in their capital holdings and productivity, and face a fixed cost of investment. Households supply labor and own the firms. In what follows we present the economic environment in detail, then describe the firm and household problems. Finally we define an equilibrium of the aggregate economy. The contribution of this paper is to study and estimate this economy with multiple aggregate shocks. We focus on these in the description of the model.

2.1 Environment

Time is discrete and the horizon is infinite. Two types of agents populate the economy: firms and households. All firms produce the same final good, which can be used either for consumption or investment.

Firms

Production technology There is a fixed unit mass of competitive firms, indexed by $i \in [0, 1]$. Firms are heterogeneous in their productivity z_{it}, and operate a decreasing returns to scale production technology $y\left(X_t^Z, z_{it}, k_{it}, n_{it}\right)$, which uses inputs of capital k_{it} and labor n_{it}. The output of production is a homogeneous final good, whose competitive price is the numeraire of the economy. The productivity of this technology is shifted by idiosyncratic productivity z_{it} and aggregate productivity (TFP) X_t^Z, which both follow AR(1) processes in logs. Aggregate productivity X_t^Z is common to all firms and evolves according to

$$\log X_{t+1}^Z = \rho^Z \log X_t^Z + \varepsilon_{t+1}^Z, \quad \varepsilon_{t+1}^Z \sim N\left(0, \sigma^Z\right).$$

Idiosyncratic productivity z_{it} is independent across firms. Each firm's productivity follows an AR(1) process in logs:

$$\log z_{it} = \rho^z \log z_{it} + \varepsilon_{it+1}^z, \quad \varepsilon_{it+1}^z \sim N\left(0, X_t^\sigma \bar{\sigma}_z\right). \tag{1}$$

The volatility of this process $X_t^\sigma \bar{\sigma}_z$ has two components, a permanent component $\bar{\sigma}_z$, which reflects average volatility over time, and a time-varying aggregate component X_t^σ, which is common to all firms, and follows an AR(1) process in levels:

$$X_{t+1}^\sigma = \rho^\sigma X_t^\sigma + \varepsilon_{t+1}^\sigma, \quad \varepsilon_{t+1}^\sigma \sim N\left(0, \sigma^\sigma\right).$$

We refer to fluctuations in X_t^σ as uncertainty shocks. The timing of the shock is such that the standard deviation of the innovations to $t+1$ idiosyncratic productivity, ε_{it+1}^z, is known to firms at the beginning of period t.

Investment The firm owns its capital stock and the timing is such that k_{it} is predetermined. The capital stock which the firm will operate in the period $t+1$ is determined by depreciation and investment in period t. Capital depreciates at a rate δ, so to enter period $t+1$ with capital stock k_{it+1}, a firm must invest $i_{it} = k_{it+1} - (1-\delta)\, k_{it}$. Investment is costly. In particular, if investment i_{it} lies outside the interval $[-bk_{it}, bk_{it}]$, the firm must pay a fixed cost ξ_{it} in units of labor. Each firm draws a new fixed cost at the beginning of period t from the distribution $F(\xi)$. This draw is *iid* across firms and time.

Labor The firm hires labor n_{it} from the date t frictionless labor market at the prevailing wage W_t.

Financing The firm is owned by the household and proceeds from production net of depreciation, wages and investment are paid out to the household as dividends d_{it}. In this version of the model we place no constraints on the values taken by d_{it}. In the appendix, we extend the model to include financial frictions, and introduce a cost of negative dividends and an aggregate shock to this cost.

Households

Preferences We assume a unit measure of identical households, which value consumption and leisure, supply labor and own the firms. The household maximizes the expected present discounted value of utility, given by

$$\mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \left[\prod_{s=0}^{t} \left(X_{s-1}^{\beta} \beta \right) \right] \left[u\left(C_t - H_t\right) - X_t^{\psi} \nu\left(N_t\right) \right] \right\}.$$

Period t utility is given by three components: a discount factor, and two additively separable components that depend on consumption and labor supply.

The household inter-period discount factor $X_t^\beta \beta$ has two components, a permanent component β and a stochastic component X_t^β, which follows an AR(1) process in logs:

$$\log X_{t+1}^{\beta} = \rho^{\beta} \log X_t + \varepsilon_{t+1}^{\beta}, \quad \varepsilon_{t+1}^{\beta} \sim N\left(0, \sigma^{\beta}\right).$$

Period utility from consumption, depends on the deviation of consumption C_t from the accumulated habit stock H_t. The function u is increasing and concave. Habit formation is modelled in the form described in Campbell and Cochrane (1999), where the evolution of the habit stock can be expressed in terms of the surplus consumption ratio $S_t = (C_t - H_t)/C_t$, which has the following law of motion

$$\log S_{t+1} = \left(1-\rho^S\right)\log \bar{S} + \rho^S \log S_t + \lambda^S \log\left(\frac{C_{t+1}}{C_t}\right). \tag{2}$$

We assume that, when making its consumption decision, the household takes the evolution of S_t as given. Winberry (2016b) shows that habit preferences of this form ensure that the correlation of the real interest with output is not counterfactually large.

Period disutility of labor supply is given by the function $v(N_t)$, which is increasing and concave and is subject to a stochastic component X_t^{ψ}, which follow an AR(1) process in logs:

$$\log X_{t+1}^{\psi} = \rho^{\psi}\log X_t^{\psi} + \varepsilon_{t+1}^{\psi}.$$

An increase X_t^{ψ} increases the disutility that households incur from supplying labor.

Assets The household trades shares in a mutual fund portfolio of all firms in the economy. Shares in this portfolio are denoted A_t, and the household is paid out the sum of all dividends from firms, denoted Π_t. The price of a share in the mutual fund is given by Q_t, which the household takes as given.

Resources Each member of the household is endowed with one unit of time which can be allocated to labor or leisure. Labor is paid a wage W_t which is determined in the labor market. The household's total resources are therefore given by returns on the mutual fund and labor income. Its total expenditures are given by consumption and new shares in the mutual fund.

2.2 Household optimization

Let $\mathbf{S}_t$ denote the aggregate state, which consists of the distribution of firms over idiosyncratic states $\mu_t = \mu(k_{it}, z_{it}, \xi_{it})$, the aggregate shocks $\mathbf{X}_t = \left(X_t^Z, X_t^{\psi}, X_t^{\beta}, X_t^{\sigma}\right)$ and level of habit stock S_t. We now describe the household and firm problems recursively.

Problem Let $\mathbf{W}(\mathbf{S})$ be the household's expected present discounted utility when the aggregate state is $\mathbf{S}$. Then $\mathbf{W}(\mathbf{S})$ satisfied the following Bellman equation:

$$\begin{aligned}
\mathbf{W}(\mathbf{S}) &= \max_{C,N,A'} u(C-H) - X^{\psi}(\mathbf{S})\, v(N) + \beta X^{\beta}(\mathbf{S}')\, \mathbb{E}\left[\mathbf{W}(\mathbf{S}')\,|\mathbf{S}\right], && (3)\\
\text{subject to} & && (4)\\
C + Q(\mathbf{S})\, A' &= W(\mathbf{S})\, N + (Q(\mathbf{S}) + \Pi(\mathbf{S}))\, A.
\end{aligned}$$

Solution The solution to the household problem consists of a labor supply condition and an Euler equation for shares prices. Given a wage $W(\mathbf{S})$, the labor supply condition is

$$W(\mathbf{S}) = X^{\psi}(\mathbf{S})\, \frac{v'(N(\mathbf{S}))}{u'(C(\mathbf{S}) - H(\mathbf{S}))}. \tag{5}$$

The Euler equation is

$$\mathbb{E}\left[\beta X^{\beta}(\mathbf{S})\, \frac{u'(C(\mathbf{S}') - H(\mathbf{S}'))}{u'(C(\mathbf{S}) - H(\mathbf{S}))}\, \frac{Q(\mathbf{S}') + \Pi(\mathbf{S}')}{Q(\mathbf{S})}\right] = 0,$$

which delivers the household's discount factor

$$M(\mathbf{S}, \mathbf{S}') = \beta X^{\beta}(\mathbf{S})\, \frac{u'(C(\mathbf{S}') - H(\mathbf{S}'))}{u'(C(\mathbf{S}) - H(\mathbf{S}))}. \tag{6}$$

Since markets are complete, this stochastic discount factor will be used by firms to price future payoffs. Note that we can rewrite $C(\mathbf{S}) - H(\mathbf{S})$ more conveniently in terms of the surplus consumption ratio:

$$C(\mathbf{S}) - H(\mathbf{S}) = C(\mathbf{S})\, S(\mathbf{S}),$$

which allows us to write the household's labor supply condition and discount factor

$$W(\mathbf{S}) = X^{\psi}(\mathbf{S})\, \frac{v'(N(\mathbf{S}))}{u'(C(\mathbf{S})\, S(\mathbf{S}))}, \qquad M(\mathbf{S}, \mathbf{S}') = \beta X^{\beta}(\mathbf{S})\, \frac{u'(C(\mathbf{S}')\, S(\mathbf{S}'))}{u'(C(\mathbf{S})\, S(\mathbf{S}))}. \tag{7}$$

This makes clear how habit formation affects the real interest rate in the economy—which is given by the inverse of the discount factor. Even if consumption is expected to grow quickly, the habit stock moves slowly, leading to a muted response of interest rates.

2.3 Firm optimization

Problem Let $v(k, z, \xi; \mathbf{S})$ be the present discounted value of dividends of the firm, as valued by the household's discount factor, given realizations of the current aggregate state

$\mathbf{S}$ and idiosyncratic states k, z and ξ. Since the firm must decide whether to adjust its capital or not, it is convenient to consider separately the value of adjusting, the value of not adjusting, and the adjustment decision.

Value of adjusting Let $v^{\text{adj}}(k, z; \mathbf{S})$ be the value of adjusting, conditional on (k, z) and aggregate state $\mathbf{S}$:

$$\begin{aligned} v^{\text{adj}}(k, z; \mathbf{S}) &= \max_{k' \geq 0, n} d + \mathbb{E}\left[M(\mathbf{S}, \mathbf{S}') v(k', z', \xi'; \mathbf{S}')\right], \qquad (8) \\ \text{subject to} & \\ d &= \pi(k, zk; \mathbf{S}) - i, \qquad (9) \\ i &= k' - (1-\delta) k, \qquad (10) \\ \mathbf{S}' &= \Gamma(\mathbf{S}' | \mathbf{S}), \qquad (11) \end{aligned}$$

where $M(\mathbf{S}, \mathbf{S}')$ is the household's one period stochastic discount factor defined in (6), and Γ is the firm's perceived law of motion of the aggregate state. Recall that capital is predetermined, which means that operating profits $\pi(z, k; \mathbf{S})$ are the outcome of the firm's static labor demand choice

$$\pi(k, z; \mathbf{S}) = \max_n \, y\left(X^Z, z, k, n\right) - W(\mathbf{S}) n. \qquad (12)$$

The solution to these problems give the firm's optimal level of capital next period, which we denote $k^{\text{adj}}(k, z; \mathbf{S})$, and its labor demand $n(k, z; \mathbf{S})$.

Value of not adjusting The value of not adjusting, $v^{\text{stay}}(k, z; \mathbf{S})$ is the same as (8) to (11) above, subject to the additional constraint that the investment rate of the firm is constrained within a small interval

$$\frac{i}{k} \in [-b, b]. \qquad (13)$$

Let $k^{\text{stay}}(k, z; \mathbf{S})$ denote the firm's optimal choice of capital conditional on not adjusting. Note that the labor demand decision n and operating profits π are the same for adjusting and non-adjusting firms.

Adjustment After observing k, z, ξ and the aggregate state $\mathbf{S}$, the firm chooses whether to adjust or not. If the firm adjusts, it pays a fixed cost ξ in units of labor. The value of the firm at the start of the period can therefore be expressed as

$$v(k,z,\xi;\mathbf{S}) = \max\left\{-\xi W(\mathbf{S}) + v^{\text{adj}}(k,z;\mathbf{S}), v^{\text{stay}}(k,z;\mathbf{S})\right\}.$$

The firm chooses to adjust if and only if the value of adjusting is greater than its cost, that is, if and only if

$$v^{\text{adj}}(k,z;\mathbf{S}) - \xi W(\mathbf{S}) \geq v^{\text{stay}}(k,z;\mathbf{S}). \tag{14}$$

For every $(k,z;\mathbf{S})$ there is a threshold value of ξ, which we denote $\xi^*(k,z;\mathbf{S})$, at which the firm is indifferent between adjusting and not adjusting. The firm adjusts if $\xi \leq \xi^*(k,z;\mathbf{S})$ and does not adjust if $\xi > \xi^*(k,z;\mathbf{S})$. We can see from (14) that the threshold is given by

$$\xi^*(k,z;\mathbf{S}) = \frac{v^{\text{adj}}(k,z;\mathbf{S}) - v^{\text{stay}}(k,z;\mathbf{S})}{W(\mathbf{S})}. \tag{15}$$

Let $k'(k,z,\xi;\mathbf{S})$ denote the capital choice of the firm, conditional on k, z, ξ and $\mathbf{S}$:

$$k'(k,z,\xi;\mathbf{S}) = \begin{cases} k^{\text{adj}}(k,z;\mathbf{S}) & \text{if } \xi < \xi^*(k,z;\mathbf{S}), \\ k^{\text{stay}}(k,z;\mathbf{S}) & \text{if } \xi \geq \xi^*(k,z;\mathbf{S}). \end{cases} \tag{16}$$

Finally, let $d(k,z,\xi;\mathbf{S})$ denote the net payout of the firm, conditional on k, z, ξ and $\mathbf{S}$,

$$d(k,z,\xi;\mathbf{S}) = \pi(k,z;\mathbf{S}) - [k'(k,z,\xi;\mathbf{S}) - (1-\delta)k].$$

2.4 Equilibrium

Wage rigidity In a competitive equilibrium, the wage W_t adjusts so that total labor demand from the firms, $N_t^D = \int_0^1 n_{it} di$, is equal to the total labor supply from the household N_t^S. In order to capture the widely documented fact that wages appear to adjust slowly over the cycle, we include some wage rigidity in the model by distorting these equilibrium conditions. In the interests of tractability, we adopt the parsimonious specification used in Beraja et al. (2016).[10]

Let W_t^* be the wage that would determine labor supply from the household's intratemporal first order condition (5)

$$W_t^* = X_t^{\psi} \frac{v'(N_t)}{u'(C_t - H_t)}. \tag{17}$$

[10] As noted in Beraja et al. (2016), this specification can be thought of as a reduced form that stands in for the endogenous wage stickiness that arises from the wage bargaining model of Hall and Milgrom (2008) or the model of monopsonistic competition discussed in Gali (2011) under the assumption of myopia on the part of the agents.

We assume that the wage paid by firms W_t has an autoregressive component and adjusts slowly towards W_t^*:

$$\log W_t = \omega \log W_t^* + (1-\omega) \log W_{t-1}. \tag{18}$$

This means that in equilibrium we clear labor markets, but the household is off its labor supply condition. Given W_t the firm demands some amount of labor $N_t^D = \int_0^1 n_{it}(W_t)\, di$, and the household is forced to supply this amount of labor. In this sense the household will be taking the firm's labor demand as given, and so in its budget constraint will have $W_t N_t^D$ as its wage payment. Its labor supply condition—for a problem where the household could choose its labor supply—is then used to index how W_t evolves. Note that if $\omega = 1$, wages are fully flexible and the labor market clears.

Definition A *recursive competitive equilibrium* of the model is a firm value function $v(k, z, \xi; \mathbf{S})$ and associated policy functions $k'(k, z, \xi; \mathbf{S})$, $n(k, z; \mathbf{S})$ and cut-off rule $\xi^*(k, z; \mathbf{S})$, household policy function $C(\mathbf{S})$, and associated stochastic discount factor $M(\mathbf{S}, \mathbf{S}')$, wage $W(\mathbf{S})$ and flexible wage $W^*(\mathbf{S})$, surplus consumption $S(\mathbf{S})$, labor demand $N(\mathbf{S})$, lagged wage $W_{-1}(\mathbf{S})$, lagged consumption $C_{-1}(\mathbf{S})$ and lagged surplus consumption $S_{-1}(\mathbf{S})$, firm payouts $\Pi(\mathbf{S})$, a distribution of firms $\mu(k, z, \xi; \mathbf{S})$ and a law of motion for the distribution of firms $\Gamma(\mu, \mu'; \mathbf{S})$, such that

1. Taking $W(\mathbf{S})$, $\mu(\mathbf{S})$, $M(\mathbf{S}, \mathbf{S}')$ and $\Gamma(\mu, \mu'; \mathbf{S})$ as given, $k'(k, z, \xi; \mathbf{S})$, $n(k, z; \mathbf{S})$ and $\xi^*(k, z; \mathbf{S})$ solve the firm's problem (8)-(11), and $v(k, z, \xi; \mathbf{S})$ is the associated value function. Aggregate dividends $\Pi(\mathbf{S})$, and labor demand $N(\mathbf{S})$ are

$$\begin{aligned}
\Pi(\mathbf{S}) &= \int \pi(k, z; \mathbf{S})\, d\mu(k, z, \xi; \mathbf{S}) \\
&\quad - W(\mathbf{S}) \int \mathbf{1}\{\xi < \xi^*(k, z; \mathbf{S})\}\, \xi d\mu(k, z, \xi; \mathbf{S}) \\
&\quad - \int (k'(k, z, \xi; \mathbf{S}) - (1-\delta) k)\, d\mu(k, z; \mathbf{S}), \\
N(\mathbf{S}) &= \int n(k, z; \mathbf{S})\, d\mu(k, z, \xi; \mathbf{S}) + \int \xi \mathbf{1}\{\xi \leq \xi^*(k, z; \mathbf{S})\}\, d\mu(k, z; \mathbf{S}).
\end{aligned}$$

2. Taking labor demand $N(\mathbf{S})$, wage $W(\mathbf{S})$, dividends $\Pi(\mathbf{S})$ and habit stock $S(\mathbf{S})$ as given, $C(\mathbf{S})$ solves the household problem (3)-(4), and $M(\mathbf{S}, \mathbf{S}')$ is the corresponding stochastic discount factor.

3. Given $W_{-1}(\mathbf{S})$, the wage $W(\mathbf{S})$ evolves according to (18), where $W^*(\mathbf{S})$ is given by (17) under $C(\mathbf{S})$ and $N(\mathbf{S})$.

4. Given $S_{-1}(\mathbf{S})$ and $C_{-1}(\mathbf{S})$, the surplus consumption ratio $S(\mathbf{S})$ evolves according to (2) under $C(\mathbf{S})$.

5. The law of motion of the distribution Γ is consistent with the firm's policies. For all measurable sets $\mathcal{K} \times \mathcal{Z} \times \mathcal{X}$,

$$\mu'(\mathcal{K} \times \mathcal{Z} \times \mathcal{X}) = Q(\mathcal{K}, \mathcal{Z}, \mathcal{X}, k, z, \xi; \mathbf{S}) \times \mu(dk, dz, d\xi),$$

where

$$Q(\mathcal{K}, \mathcal{Z}, \mathcal{X}, k, z, \xi) = \int P(z' \in \mathcal{Z}|z)\, dz' \times \mathbf{1}\{k'(k, z, \xi; \mathbf{S}) \in \mathcal{K}\} \times G(\mathcal{X}).$$

6. The aggregate shocks in $\mathbf{X}$ each evolve according to the exogenous process:

$$\begin{aligned}
\log\left(X^j\right)' &= \rho^j \log X^j + \left(\varepsilon^j\right)', \qquad \varepsilon^j \sim N\left(0, \left(\sigma^j\right)^2\right), \qquad \forall j \in \{Z, \psi, \beta\} \\
(X^\sigma)' &= \rho^\sigma X^\sigma + \varepsilon^\sigma, \qquad \varepsilon^\sigma \sim N\left(0, (\sigma^\sigma)^2\right).
\end{aligned}$$

3 Solution method

In general, it is difficult to solve for the recursive competitive equilibrium for models of this type, since the firm's policies depend on firm's forecasts of the aggregate consumption and wage, and next period's aggregate consumption and wage depend on next period's distribution of firms, which is an infinite-dimensional object μ. The pre-eminent method in the literature is the method of Krusell and Smith (1998). We follow an alternative method for two reasons, both of which are motivated by the aim of the paper, which is to estimate this equilibrium heterogeneous firms model using moments of the distribution of firms as observable.

First, given the large number of aggregate shocks in our model, the approach of Krusell and Smith (1998) is impractical. Following Krusell-Smith as it iss applied in the literature would require stipulating a law of motion for the marginal utility as a function of one moment of the distribution—for example, past aggregate capital—and the exogenous states. With many exogenous states, a forecasting rule would include many level and cross-product terms, rendering the approach both impractical and slow. Since we will have to solve the model many times in order to estimate the processes for aggregate shocks, our procedure must be fast.

Second, for estimation, it is desirable to have a linearized state space representation of the model's equilibrium conditions. Krusell-Smith leads to a non-linear representation of the equilibrium.

Instead we turn to a new approach in the literature, pioneered by Reiter (2009), which involves solving firm's policies globally at the deterministic steady state, and then perturbing the solution with respect to aggregate shocks. The approach therefore maintains the full non-linearity of the firm's policies with respect to idiosyncratic productivity and capital, and perturbs these policies linearly with respect to aggregates.

This approach proceeds in three steps. First, we construct a finite representation of the equilibrium conditions which, as they stand are equations in infinite dimensional objects: the value function and the distribution of firms. Second, we solve for the steady-state of the model, which is given by the values that satisfy the discretized equilibrium conditions in the absence of aggregate shocks. Third, we linearize the equilibrium conditions around the steady state to obtain a state-space representation of the macroeconomic dynamics of the model.

Finite representation The equilibrium conditions contain two infinite dimensional objects that we must approximate: the value function and the distribution of firms. We construct a finite representation of the expected value function $\tilde{v}\left(k, z; \mathbf{S}\right)$ using cubic splines, where $\tilde{v}\left(k, z; \mathbf{S}\right) = \int v\left(k, z, \xi; \mathbf{S}\right) dG\left(\xi\right)$ (since the adjustment cost ξ is *iid*, it is convenient to integrate it out and work with the expected value function). Denote by θ_{ij}^{V} the finite vector of coefficients of the cubic spline representation for $\tilde{v}\left(k, z; \mathbf{S}\right)$. Here $i \in (1, \ldots, n_k)$ and $j \in (1, \ldots, n_z)$, index the nodes in the state-spaces for capital and productivity over which we form this approximation. We approximate the distribution $\mu\left(k, z; \mathbf{S}\right)$ with a histogram, parameterized by λ_{ij}. The law of motion of this histogram is then computed following the method of Young (2010), which preserves aggregation. This law of motion and the other discretized equilibrium conditions are given in Appendix A.

We note that the choice of a histogram rather than a functional approximation of the distribution, as considered by Algan et al. (2014), is motivated by our use of moments of the distribution of firms in estimation. If we only wanted to solve the model, then our approximation only needs to keep track of the features of the distribution that are relevant for approximating how aggregate prices evolve. Here we want to estimate the model using complicated moments of the distribution, such as the standard deviation of sales growth. This requires keeping track of the entire distribution of sales growth. In practice, we have found the histogram method to be more appropriate for achieving this.

In practice, the bins over which we construct the histogram representation of μ may be different from the points used in the approximation of $\tilde{v}$. For ease of exposition, here we assume they are the same.

Linearizing around steady state Using the finite representations of the value function and distribution, we are able to write the equilibrium conditions as a system of differences equations in the finite vector of variables. The equilibrium conditions are variously backward-looking (e.g. wage rigidity, evolution of distribution), and forward-looking (e.g. the Bellman equations). Following the standard approach of using expectational errors in forward-looking equations (denoted η_{t+1}), the equilibrium can be written as the finite non-linear system

$$\Gamma\left(\Theta_t, \Theta_{t+1}, \eta_{t+1}, \varepsilon_{t+1}\right) = \mathbf{0} \tag{19}$$

where the vector $\Theta_t = \left[\left(\theta_{ij}^V\right)_t, \left(\lambda_{ij}\right)_t, \mathbf{X}_t, C_t, S_t, W_t, W_t^*, \mathbf{g}_t\right]$ contains state and jump variables, and ε_{t+1} is a $n_\varepsilon \times 1$ vector of Gaussian disturbances to the exogenous processes. The vector Θ_t also contains an $M \times 1$ vector $\mathbf{g}_t$, which will contain observables that will be used in the estimation step, and are potentially non-linear functions of the other elements of Θ_t.

With this representation in hand, we first solve for the equilibrium value of Θ_t when the aggregate shocks $\mathbf{X}_t$ are zero, using nonlinear global methods. We call the value of Θ_t when the aggregate shocks are zero the deterministic steady state, and denote it $\bar{\Theta}$, which satisfies

$$\Gamma\left(\bar{\Theta}, \bar{\Theta}, \mathbf{0}, \mathbf{0}\right) = \mathbf{0}$$

We then express (19) in terms of log deviations from steady state, $\hat{\Theta}_t = \log \Theta_t - \log \bar{\Theta}$, and take a first-order Taylor expansion. This delivers a linear system of equations, which give a SVAR representation of the model,

$$\Gamma_0 \hat{\Theta}_{t+1} = \Gamma_1 \hat{\Theta}_t + \Psi \varepsilon_{t+1}. \tag{20}$$

The matrices Γ_0 and Γ_1 contain first-order partial derivatives of the equilibrium conditions with respect to the elements of Θ_t, which are computed numerically. To understand the composition of Γ_0 and Γ_1, consider the example of the condition which defines aggregate output Y_t, which we include in $\mathbf{g}_t$. In its discretized form, this condition reads as follows, where the subscript t picks up the dependence on the aggregate state:

$$Y_t = \sum_{i=1}^{n_k} \sum_{j=1}^{n_z} \underbrace{X_t^Z}_{(1)} \int \underbrace{y_t\left(k_i, z_j, n\left(k_i, z_j, X_t^z, W_t\right) \xi\right)}_{(2)} dF(\xi) d \underbrace{\lambda_{ijt}}_{(3)}.$$

Following a shock to X_t^Z, output responds for three reasons: (1) there is a direct effect on output in period t, holding firm decisions and the distribution of firms constant, (2) firm policies respond—which can be separated into direct responses to the shock and indirect responses due to the movement of prices—and (3) the response of policies will shift the distribution of firms in future periods, which will effect output in future periods. Numerically differentiating this condition to compute one element of Γ_0 requires perturbing X^Z and computing these responses. In addition, other elements of these matrices pick up how prices respond to shocks, how policies respond to prices, and how the distribution responds to changes in firm policies.

VAR representation Given the finite, linear formulation of the model's equilibrium conditions (20) we can use standard methods, such as the method described by Sims (2002), to obtain a linear Gaussian state-space representation

$$\hat{\Theta}_{t+1} = A\hat{\Theta}_t + B\varepsilon_{t+1} \tag{21}$$

where A is a $n_x \times n_x$ matrix and and B is a $n_x \times n_\varepsilon$ matrix.

With this representation in hand, we can compute the likelihood of any sequence of $\hat{\Theta}_t$. In our estimation, we assume that we observe only the elements of $\hat{\Theta}_t$ that correspond to $\hat{\mathbf{g}}_t$, where $\hat{\mathbf{g}}_t = D\hat{\Theta}_t$ and D is a selection matrix. Given that the system is linear and Gaussian, it is straightforward to compute the likelihood of a sequence of $\hat{\mathbf{g}}_t$, using the Kalman filter.

This solution method has therefore met both of our criteria. First, the model can be solved quickly with many aggregate shocks. Second, its amenable to estimation using Bayesian methods, where we can incorporate functions of the distribution of firms as observables in the estimation.

4 Estimation

4.1 Functional forms

We assume that the period household utility function is given by

$$u(C_t - H_t) - X_t^\psi \psi v(N_t) = \log\left(C_t - H_t\right) - X_t^\psi \psi \frac{N_t^{1+\eta}}{1+\eta}.$$

That is, households have a unit coefficient of relative risk aversion and a positive elasticity of labor supply, where η is the inverse of the Frisch elasticity.

The firm's production technology is given by

$$y\left(X_t^Z, z_{it}, k_{it}, n_{it}\right) = X_t^Z z_{it}^{\frac{1-\kappa}{\kappa}} \left(k_{it}^{1-\nu} n_{it}^{\nu}\right)^{\kappa}.$$

where κ controls the decreasing returns to scale in the production technology, and ν the output elasticity of labor. We assume that z_{it} enters with the exponent $(1-\kappa)/\kappa$ so that the firm's sales y_{it} are linear in productivity when we substitute in the first order condition for labor demand and the firm's optimal capital stock. This assumption is a normalization of productivity, and helps with the approximation of the distribution of capital.[11]

Finally, we assume that the fixed cost of investment is uniformly distributed, $\xi \sim U\left[0, \bar{\xi}\right]$, and we refer to $\bar{\xi}$ simply as the adjustment cost.

4.2 Estimation strategy

We estimate the model in three stages. First, we externally calibrate a number of parameters. These include preference parameters, depreciation rate δ, and labor output elasticity ν. Second, we estimate a subset of the model parameters by simulated method of moments. These are the remaining parameters needed to solve the steady state of the model: the adjustment cost $\bar{\xi}$, the parameters of the firm idiosyncratic productivity process ρ_z and σ_z, and the decreasing returns to scale parameter κ. The moments used in the estimation capture key properties of firm-level investment behavior. Third, we estimate the parameters of the aggregate shock processes using Bayesian methods with time series data for macroeconomic aggregates and moments of the firm distribution.[12]

4.3 Data

For the moments of the distribution of firms, we use Compustat data, which collects accounting data for the universe of publicly listed firms in the United States. For aggregate time series, we use NIPA data. In both cases, data are quarterly and for the time period 1985:I to 2014:IV.

4.4 Externally calibrated parameters

Externally calibrated parameters are reported in Table 1. The model is quarterly, so we set the discount factor β to 0.99. This results in an average real interest of 4% annually.

[11]If this exponent was equal to one, then the firm's desired capital stock would be highly convex in the firm's productivity, with a coefficient of $1/(1-\kappa) \approx 5$ on the productivity term. This makes approximating the movement of firms in the tail of the distribution challenging. Since we are interested in time series of moments such as the standard deviation of investment rates, then getting a good approximation of how firms with large levels of capital move about the distribution is important.

[12]In future versions of the paper we aim to estimate some of these parameters—especially $\bar{\xi}$—in the Bayesian estimation step, since, as we show later, the value of this parameter is important for the aggregate dynamics of the model. Therefore the information in the aggregate time series used in the Bayesian estimation should be informative for $\bar{\xi}$.

Table 1: Externally calibrated parameters.

Parameter		Value	Parameter		Value
Discount factor	β	0.99	Output elasticity of labor	ν	0.65
Curvature of utility function	σ	1	Wage flexibility	$1-\omega$	0.69
Inverse Frisch elasticity	η	2	Average surplus ratio	$\bar{S}$	0.65
Depreciation rate	δ	0.03	Surplus autocorrelation	ρ^S	0.95

We set the depreciation rate δ to 0.03, which gives an aggregate investment rate of 12 percent annually, consistent with aggregate US data (NIPA). The output elasticity of labor ν is chosen to match a labor share of 0.65. For wage rigidity, we take the estimate of the autoregressive coefficient of 0.31 from Beraja et al. (2016). The inverse Frisch elasticity is set to 2, consistent with the range of values surveyed in Chetty et al. (2011). For the habit process, we follow Winberry (2016a), who shows that a surplus consumption ratio $\bar{S} = 0.65$ and surplus autocorrelation $\rho^S = 0.95$ deliver a correlation of interest rates and output close to zero, as is the case in the data. We verify in Section 4 that this holds in our model. (b??)

4.5 Simulated method of moments

The vector of parameters to be estimated is $\theta_{SMM} = \left(\bar{\xi}, \rho_z, \sigma_z, \kappa\right)$. To estimate these parameters we proceed by simulated method of moments. We specify a vector of moments $\mathbf{h}_t$ which we compute for each quarter in our sample. Our estimate $\hat{\theta}_{SMM}$ minimizes the minimum distance criterion function

$$\mathcal{Q}(\theta) = \left(\mathbf{h}\left(\bar{\mathbf{S}};\theta\right) - T^{-1}\sum_{t=1}^{T}\mathbf{h}_t\right)' \mathbf{W} \left(\mathbf{h}\left(\bar{\mathbf{S}};\theta\right) - T^{-1}\sum_{t=1}^{T}\mathbf{h}_t\right),$$

where $\mathbf{h}\left(\bar{\mathbf{S}};\theta\right)$ is the corresponding vector of moments computed from the model when aggregate shocks are all set to zero. Note that given our linear solution of the aggregate dynamics of the model, these also correspond to the time-series average of simulations of the model with aggregate shocks turned on, that is $h\left(\bar{\mathbf{S}}\right) = \lim_{T\to\infty} T^{-1}\sum_{t=1}^{T}\mathbf{h}\left(\mathbf{S}_t\right)$. The weighting matrix is diagonal with entries equal to $1/h_{mt}^2$ for each moment m.

Four moments are chosen to provide power in identifying θ_{SMM} and capture key properties of both firm level behavior. Investment moments include the fraction of non-adjusting firms and standard deviation of investment rates at adjusting firms. We also include the ratio of dividends to output which is informative for decreasing returns to scale. Finally we include the standard deviation of sales growth. The standard deviation of sales growth and the standard deviation of investment rates will be used as observables in the Bayesian estimation step. Since the representation of the model used in the estimation is linearized around the steady state, it is important that they are included and matched in the steady state of the model. We now detail how we construct these moments.

Investment The investment rate of firm j in period t, ir_t is defined as the ratio of gross investment of the firm between periods t and $t+1$ to the average capital stock in the two periods:

$$\mathrm{ir}_{jt} = \frac{(k_{jt+1} - k_{jt})}{\frac{1}{2}(k_{jt+1} + k_{jt})}.$$

In the model, the equivalent variable is:

$$\mathrm{ir}(k, z, \xi; \mathbf{S}) = \frac{k'(k, z, \xi; \mathbf{S}) - k}{\frac{1}{2}(k'(k, z, \xi; \mathbf{S}) + k)}.$$

The investment inaction rate in period t, $\mathrm{inaction}_t$, is the fraction of firms with investment rate ir_{jt} less than 1 percent in absolute terms:

$$\mathrm{inaction}_t = \frac{1}{N_t} \sum_{j \in J_t} \mathbf{1}\{|\mathrm{ir}_{jt}| < 0.01\}.$$

where J_t is the set of firms of firms in the economy at time t and N_t is the number of firms at time t. In the model, the corresponding moment is

$$\mathrm{inaction}(\mathbf{S}) = \int \mathbf{1}\{|\mathrm{ir}(k, z, \xi; \mathbf{S})| < 0.01\} \, d\mu(k, z, \xi; \mathbf{S}).$$

The standard deviation of investment rates, conditional on adjustment, at time t, is the standard deviation of ir_{jt} across adjusting firms:

$$\sigma_t^{\mathrm{ir}} = \frac{\sum_{j \in J_t} \left(\mathrm{ir}_{jt} - \bar{\mathrm{ir}}_t\right)^2 \mathbf{1}\{|\mathrm{ir}_{jt}| > 0.01\}}{\sum_{j \in J_t} \mathbf{1}\{|\mathrm{ir}_{jt}| > 0.01\}}$$

where $\bar{\mathbf{ir}}_t$ is the mean investment rate. In the model, the corresponding moment is

$$\sigma^{\mathrm{ir}}(\mathbf{S}) = \frac{\int \left[\mathrm{ir}(k, z, \xi; \mathbf{S}) - \bar{\mathrm{ir}}(\mathbf{S})\right]^2 \mathbf{1}\{|\mathrm{ir}(k, z, \xi; \mathbf{S})| < 0.01\} \, d\mu(k, z, \xi; \mathbf{S})}{\mathbf{1}\{|\mathrm{ir}(k, z, \xi; \mathbf{S})| < 0.01\} \, d\mu(k, z, \xi; \mathbf{S})}.$$

where $\bar{\mathrm{ir}}(\mathbf{S})$ is the mean investment rate, conditional on adjustment.

Sales and output The sales growth of a firm j at time t is the change in the firm's sales between t and $t-1$, as a fraction of average sales in the two periods:

$$\mathrm{sg}_{jt} = \frac{\mathrm{s}_{jt} - \mathrm{s}_{jt-1}}{\frac{1}{2}\left(\mathrm{s}_{jt} + \mathrm{s}_{jt-1}\right)},$$

where s_{jt} is sales for firm j in period t (see Section B.1 for details). The corresponding quantity in the model is the sales growth of a firm with state (k, z, ξ) and previous state $(k_{-1}, z_{-1}, \xi_{-1})$, when the aggregate state is $\mathbf{S}$ and last period's state was $\mathbf{S}_{-1}$:

$$\mathbf{sg}\left(k, z, \xi, k_{-1}, z_{-1}, \xi; \mathbf{S}, \mathbf{S}_{-1}\right) = \frac{y\left(k, z, \xi; \mathbf{S}\right) - y\left(k_{-1}, z_{-1}, \xi_{-1}; \mathbf{S}_{-1}\right)}{\frac{1}{2}\left(y\left(k', z', \xi'; \mathbf{S}\right) + y\left(k, z, \xi; \mathbf{S}_{-1}\right)\right)},$$

where $y\left(k, z, \xi; \mathbf{S}\right) = X^{Z}\left(\mathbf{S}\right) z \left(n\left(k, z; \mathbf{S}\right)^{\nu} k^{1-\nu}\right)^{\kappa}$. At each period t, the cross-sectional standard deviation of sales growth, σ_t^{sg}, is the standard deviation of sg_{jt} across all firms:

$$\sigma_t^{\mathrm{sg}} = \frac{1}{N_t} \sum_{j \in J_t} \left(\mathrm{sg}_{jt} - \frac{1}{N} \sum \mathrm{sg}_{jt}\right).$$

The corresponding moment in the model is

$$\sigma^{\mathrm{sg}}\left(\mathbf{S}, \mathbf{S}_{-1}\right) = \int \left\{ \int \left[\mathrm{sg}\left(k, z, \xi, k_{-1}, z_{-1}, \xi_{-1}; \mathbf{S}, \mathbf{S}_{-1}\right) - \bar{\mathrm{sg}}\left(\mathbf{S}, \mathbf{S}_{-1}\right)\right]^2 d\mu\left(k, z, \xi; \mathbf{S}\right) \right\} d\mu\left(k_{-1}, z_{-1}, \xi_{-1}; \mathbf{S}_{-1}\right).$$

Finally, the ratio of dividends to sales in the data and model are computed

$$\mathrm{dy}_t = \frac{\sum_{j \in J_t} d_{jt}}{\sum_{j \in J_t} s_{jt}}, \quad \mathrm{dy}\left(\mathbf{S}\right) = \frac{\int d\left(k, z, \xi\right) d\mu\left(k, z, \xi; \mathbf{S}\right)}{\int y\left(k, z, \xi\right) d\mu\left(k, z, \xi; \mathbf{S}\right)}.$$

Cleaning In practice we compute these moments within sectors (as designated by 2-digit SIC codes) then remove seasonal and sectoral effects. Again, more details can be found in B.1. The vectors of moments are

$$\begin{aligned} \mathbf{h}_t &= \left(\mathrm{inaction}_t, \sigma_t^{\mathrm{ir}}, \sigma_t^{\mathrm{sg}}, \mathrm{dy}_t\right), \\ \mathbf{h}\left(\bar{\mathbf{S}}\right) &= \left(\mathrm{inaction}\left(\bar{\mathbf{S}}\right), \sigma^{\mathrm{ir}}\left(\bar{\mathbf{S}}\right), \sigma^{\mathrm{sg}}\left(\bar{\mathbf{S}}\right), \mathrm{dy}\left(\bar{\mathbf{S}}\right)\right). \end{aligned}$$

Parameter estimates Table 2 summarizes the estimated parameters and vector of moments in the data and model. At these parameters the model very closely matches the data. The adjustment cost parameter $\bar{\xi}$ implies that adjustment costs are 0.4% of output, which is close to the 1% reported by Khan and Thomas (2008). The degree of decreasing returns to scale is 0.744, which lies within the range of values in the literature: Bachmann and Bayer

Table 2: Targeted moments.

Parameter	θ_{SMM}	**Value**	**Target**	**Data**	**Model**
Upper bound of adjustment costs	$\bar{\xi}$	0.00037	Investment inaction rate	0.263	0.270
Decreasing returns to scale	κ	0.744	Dividends / Output	0.070	0.073
Persistence of firm-level shocks	ρ_z	0.31	SD sales growth	0.285	0.243
Average volatility of firm-level shocks	$\bar{\sigma}_z$	0.32	SD investment rates	0.154	0.181

(2014) estimate decreasing returns to scale of 0.75.[13]

The estimated parameters for the idiosyncratic productivity process imply a cross-sectional dispersion of productivity in steady state of 0.33. This is close to the average dispersion of productivity of 0.375 estimated by Imrohoroglu and Tuzel (2014) for the US (also in Compustat data), but substantially larger than that considered in other papers in the firm level investment literature. This number must be large in order to match the large observed standard deviation of sales growth found in the data. An order of magnitude smaller dispersion in productivity emerge from other papers that do not match this moment.[14]

The model struggles to match both the high dispersion in sales growth and the low dispersion in investment rates. Keeping ρ_z fixed, increasing $\bar{\sigma}_z$ increases both. Keeping $\bar{\sigma}_z$ fixed, increasing ρ_z strongly reduces the standard deviation of sales *growth* - since productivity is highly correlated from one period to the next. But increasing ρ_z increases the standard deviation of investment rates. As ρ_z increases, firms wait for large productivity shocks then make large lumpy investments, increasing the standard deviation of the investment rate. The model cannot get ρ_z low enough to get the model up to the empirical standard deviation of sales growth, without leading to large values of the standard deviation of investment rates. Notable in this exercise is how matching the standard deviation of sales growth constrains the parameter estimates.[15]

4.6 Bayesian estimation

There are four exogenous shocks in the model: TFP $\left(X_t^Z\right)$, labor disutility $\left(X_t^\psi\right)$, time prefrence $\left(X_t^\beta\right)$, and the uncertainty shock (X_t^σ). We estimate the parameters of the shock

[13] By comparison, Khan and Thomas (2008) calibrated to 0.9.

[14] Although unreported, the dispersion in productivity due to idiosyncratic shocks in these papers is as follows: Bloom et al. (2012) - 0.12, Winberry Winberry (2016a) - 0.07, Khan and Thomas (2008) - 0.04. These appear counterfactually small.

[15] Adding the standard deviation of sales growth also gets around a problem detailed in Clementi and Palazzo (2010) regarding the inability to identify $\bar{\sigma}_z$, ρ_z and $\bar{\xi}$ from the inaction rate, dispersion of investment rates and kurtosis of the investment rate distribution alone.

processes $\theta_{MLE} = \left(\rho^Z, \sigma^Z, \rho^\psi, \sigma^\psi, \rho^\beta, \sigma^\beta, \rho^\sigma, \sigma^\sigma\right)$ using Bayesian methods, as surveyed in An and Schorfheide (2007).

Computing the likelihood As described in Section 3, for a given set of parameters θ, we obtain a model solution of the form

$$\hat{\Theta}_t = A(\theta)\,\hat{\Theta}_{t-1} + B(\theta)\,\mathcal{E}_t \tag{22}$$

where $\hat{\Theta}_t$ is a $n_\xi \times 1$ vector of latent states, and $\mathcal{E}_t$ is a $n_\varepsilon \times 1$ vector of standard normal Gaussian innovations $\mathcal{E}_t \sim N\left(\mathbf{0}, \mathbb{I}_{n_\varepsilon \times n_\varepsilon}\right)$. We define an observation equation

$$\Upsilon_t = D\hat{\Theta}_t + C\eta_t \tag{23}$$

where Υ_t is an $n_y \times 1$ vector of observables, D is an $n_y \times n_x$ matrix which selects elements of $\hat{\Theta}_t$. We include measurement error through the vectorη_t which is an $n_y \times 1$ vector of *iid* Gaussian innovations with covariance matrix C.

Given a set of time series data corresponding to the elements of Υ_t, the likelihood of a model in the form (22)-(23) is easily computed using the Kalman filter.[16]

Observable time series In order to estimate the parameters associated with four aggregate shocks we must specify four observable time-series, that is $n_y = 4$. We include three variables that are standard in the literature—output, consumption, and hours worked, which are constructed from NIPA—and one variable which we construct from the Compustat microdata. In our baseline estimation this *cross-sectional time-series* is the standard deviation of sales growth. We also contrast our baseline results to an alternative estimation using the standard deviation of investment rates in place of the standard deviation of sales growth.

For all variables take logs, and detrend to remove fluctuations that are not at a business cycle frequency. For the aggregate variables we apply a one-sided HP-filter with a smoothing parameter of 1600 which is the conventional parameter for quarterly data. For the cross-sectional time series we apply a quadratic trend. We do this because it turns out that the cross-sectional time-series have much lower frequency trends. This means it is unclear what the appropriate smoothing parameter should be for these series. For the cross-sectional time-series we also remove seasonal effects at a quarterly frequency and industry fixed effects (see Appendix B.1 for details). The resulting time series are plotted in Figure 9-7.

[16] Initializing the filter requires that we specify a prior mean and variance for the state in the initial period. We pick the prior distribution to have zero mean (since the state is expressed as deviations from steady state), and variance Σ, where Σ is the unconditional variance-covariance matrix of $\hat{\Theta}_t$, which solves $\Sigma = A\Sigma A' + BB'$, which we obtain by simulation.

Table 3: Shock process parameters.

		Prior			Posterior	
Parameter name	Parameter	**Type**	**Mean**	**SD**	**Mode**	**SD**
Autocorrelation, TFP shock	ρ^Z	Beta	0.600	0.260	0.983	0.015
Autocorrelation, labor supply shock	ρ^ψ	Beta	0.600	0.260	0.986	0.006
Autocorrelation, time preference shock	ρ^β	Beta	0.600	0.260	0.739	0.054
Autocorrelation, uncertainty shock	ρ^σ	Beta	0.600	0.260	0.775	0.257
Standard deviation, TFP shock	σ^Z	Exp	0.100	0.100	0.012	0.001
Standard deviation, labor supply shock	σ^ψ	Exp	0.100	0.100	0.012	0.001
Standard deviation, time preference shock	σ^β	Exp	0.100	0.100	0.002	0.001
Standard deviation, uncertainty shock	σ^σ	Exp	0.100	0.100	0.002	0.005

We compute the mode and standard deviation of the posterior by drawing from the posterior distribution using the Metropolis-Hastings algorithm. We use 10 chains of 100,000 draws each.

Since we have four time series and four shocks, we could proceed by setting all elements of the measurement equation matrix C to zero. However, we find in preliminary calculations that, for some sets of parameters, this gives likelihoods which are indistinguishable from zero to machine precision. In order to avoid this artifact of the numerical estimation, we allow for some measurement error in the estimation: in particular, we allow for error in the measurement of the cross-sectional time-series since we conjecture that this is measured with more error than the aggregate variables. We set all elements of C to zero except for the diagonal element corresponding to σ_t^{sg}, which we denote η_σ, and include this as a parameter in our estimation.

Priors and posteriors

We estimate the parameters using Bayesian methods following An and Schorfheide (2007). The independent prior distributions for each of the estimated parameters are given in the first three columns of Table 3. For all the autocorrelation parameters $\left(\rho^Z, \rho^\psi, \rho^\beta, \rho^\sigma\right)$ we set a Beta prior with a mean of 0.6 and standard deviation of 0.26. This is a dispersed prior which restricts the parameters to be between 0 and 1, as autocorrelation parameters must be for the process to be stationary, and is the close to the prior chosen by Smets and Wouters (2007). For the standard deviation parameters $\left(\sigma^Z, \sigma^\psi, \sigma^\beta, \sigma^\sigma\right)$, we set an exponential prior with mean 0.1. The conventional choice in the literature is to set priors that follow the inverse Gamma distribution for standard deviation parameters. However, this is not an appropriate choice for us, since the inverse Gamma function tends to 0 as the argument approaches zero and therefore tends to move the standard deviation away from zero.[17] Given prior densities

[17]See Gelman (2006) for a discussion of the inverse Gamma distribution as a prior.

Figure 1: Countercyclicality of cross-sectional sales growth dispersion.

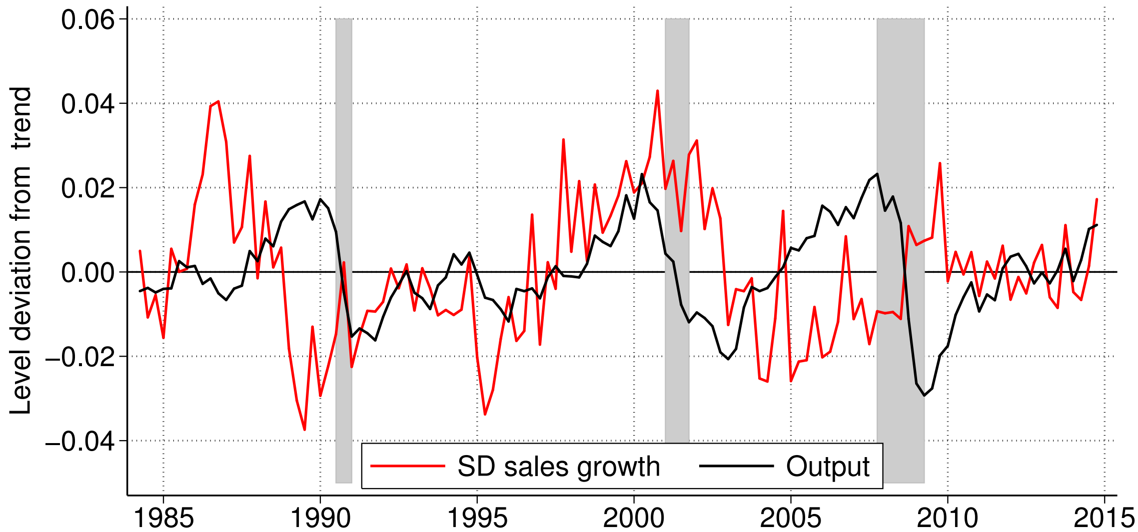

we compute the posterior using the Metropolis-Hastings algorithm. This is implemented by using 10 parallel chains of 100,000 draws, of which we discard the first 10,000 draws of each chain and splice together the remaining draws. We calculate convergence diagnostics as described in Gelman and Rubin (1992) and verify that the $\sqrt{\hat{R}}$ statistic is less than 1.05 for each parameter separately.

The prior and posterior densities for each parameter are plotted in Figure 8. The estimated persistence for the TFP shock, ρ^Z, and for the labor supply shock, ρ^ψ, are very high and reasonably well-identified, at 0.983 and 0.986, respectively. The posterior modes of the persistence parameters are similar to estimates in related RBC and New Keynesian models: Smets and Wouters (2007), for instance, we find high persistence in the TFP shock and wage markup shocks (0.95 and 0.97, respectively). The estimated persistence of the time preference shock is 0.739, in line with the persistence of the risk premium shock in Smets and Wouters (2007), which, like our time preference factor shock, appears as the a wedge in the household's Euler equation.

Finally, consider the persistence and volatility of the uncertainty shocks. We find that the persistence and volatility of the uncertainty shocks are $\rho^\sigma = 0.774$ and $\sigma^\sigma = 0.002$.

5 Results

With the model estimated we can quantitatively approach our main question: Are movements in firm-level dispersion a source of or a response to business cycle fluctuations? This

Figure 2: Forecast error variance decomposition of aggregate variables

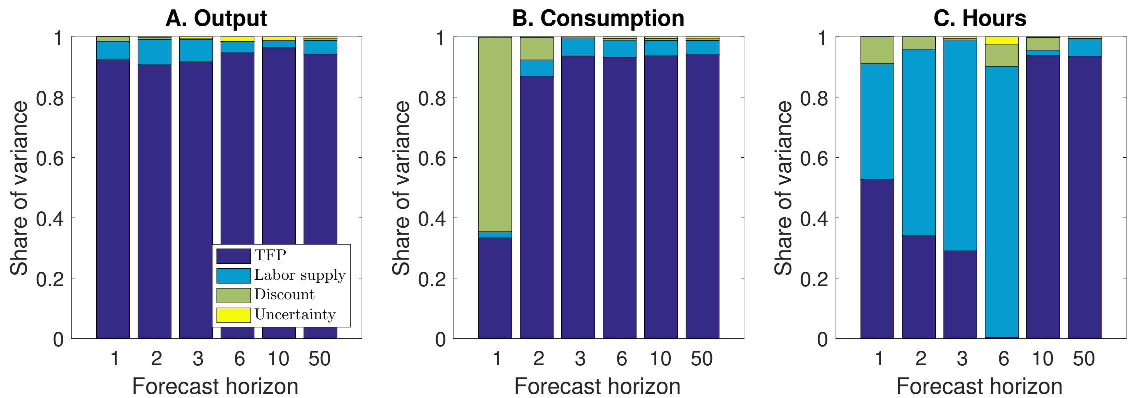

can be answered by studying variance decompositions of the model. These attribute the percentage of the variance of a given model-series due to each shock. We compute these for a number of different horizons. To be clear, we would conclude that firm-level dispersion is a source of business cycle fluctuations if the variance decomposition attributed a large fraction of fluctuations in output to the uncertainty shock. While firm-level dispersion is a response to business cycle fluctuations if the variance decomposition attributing a large fraction of fluctuations in sales growth dispersion to the productivity, labor disutility and preference shocks. We first describe our results then return to explaining them. A relevant figure to return to is Figure 1 which shows the time-series for the standard deviation of sales growth and output.

5.1 Forecast error variance decompositions

First, we consider what drives fluctuations in macroeconomic aggregates. Figure 2 shows the Forecast Error Variance Decompositions (FEVDs) for output, consumption and hours (Tables 4 and 5 in the appendix detail the exact one period and two period ahead decompositions). Since output is equal to hours plus consumption in our model, then the FEVD for investment is essentially the same as output.[18]

Our first main result it that uncertainty shocks play almost no role in explaining fluctuations in these macroeconomic aggregates. The uncertainty shock explains 1%, 1% and 2% of the movements in output, consumption and labor, respectively, at a 1-period horizon. The decomposition of the aggregates is similar to what we would find in the model without

[18]We could have included investment as an additional observable in the estimation of the model and, following the literature, included a shock to generate a wedge in the resouce identity $Y_t = C_t + I_t + wedge_t$. However since the goal of the paper is not explaining business cycles per se, this would not change our results in a way that they answer our main questions.

Figure 3: Forecast error variance decomposition of cross-sectional variables

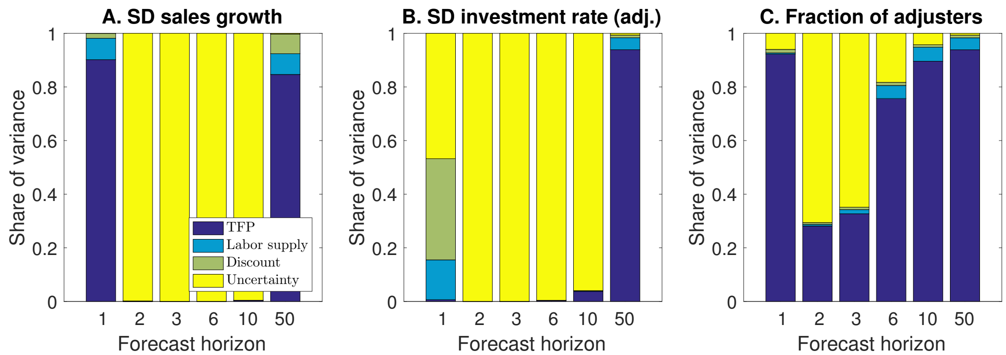

uncertainty shocks: TFP drives output and consumption, and the labor supply shock is important for short run fluctuations in hours, allowing the model to generate larger likelihoods when output falls a little and hours fall a lot (i.e., the labor wedge). The discount rate shock is important for short-term consumption fluctuations.

Second, we consider what is driving fluctuations in movements in the cross-section of firms. Figure 3 plots the FEVDs for the standard deviation of sales growth, standard deviation of investment rates and the fraction of firms adjusting their capital. Recall that we use the standard deviation of sales growth in the data, but we can compute FEVDs for any series in the state-space representation of the model. Including the standard deviation of sales growth and fraction of adjusting firms as elements of $\mathbf{g}_t$ allows us to compute their FEVDs.

Almost all of the fluctuations in the standard deviation of sales growth are explained by the uncertainty shock. The timing of the uncertainty shock is such that an uncertainty shock at time t means that productivities will be more dispersed at time $t+1$. Therefore it is not surprising that the uncertainty shock has a very small effect on impact. In period t, the only margin available to firms which might influence σ_t^{sg} is the hiring decision. At all medium horizons ($t = 2, \ldots, 10$) the uncertainty shock is the dominant source of fluctuations in the standard deviation of sales growth, and very little is explained by the 'RBC shocks'. The uncertainty shock, does, however effect investment in period t, and contributes to immediate movements in the standard deviation, which are also effected by the discount rate shock which shifts the timing of investment at some firms.

Despite the fact that the uncertainty shock does not have large effects on output, we do find that—in the model—it explains a large amount of the fluctuations in the fraction of adjusting firms. The result that the uncertainty shock does not lead to movements in output, therefore will be related to the fact that movements in the fraction of adjusters do not generate large movements in output.

In Figure **??**, we plot what the model implies for σ_t^{sg} when the uncertainty shock is shut down (i.e. X_t^σ is set to 0). This confirms the result of the forecast error variance decomposition. The RBC shocks explain a negligible portion of the movement in σ_t^{sg} across the entire time period.

5.2 Impulse response functions

The fact that uncertainty shocks to not generate large fluctuations in output is not due to an uncertainty shock generating counterfactual predictions for the direction of responses of other variables used in the estimation. Figure 4A shows that a positive uncertainty shock leads to a steep decline in investment, a decline in output and a modest decline in consumption. These responses are consistent with Bloom (2009). Figure 4B shows that a negative TFP shock leads to movements of these variables in the same direction and in approximately the same proportion. The issue, instead, is in the magnitudes. The positive uncertainty shock leads to a large increase in the standard deviation of sales growth, but the movements in aggregate are small. Therefore, if it were to be the sole source of business cycle fluctuations we would see counterfactually larger fluctuations in the standard deviation of sales growth than we see in the data. Similarly, the negative TFP shock leads to a large movement in aggregates but the movement in the cross-section is small. Therefore we need uncertainty shocks to match the observed fluctuations in the dispersion of sales growth. Our result is due to an estimation procedure which forces both RBC and uncertainty shocks to account—in a likelihood sense—for the data.

For completeness, Figure 11 in Appendix C gives the full set of impulse responses of output, consumption, hours, investment and the dispersion of sales growth in response to each of the four shocks in the model. These help in making clear the above point regarding the relative magnitudes of the responses.

Two more subtle points arise from studying these IRFs. First, we see that at medium horizons of $5-10$ years, the uncertainty shock has a small positive effect on output, accounting for around 5% of its variance. The IRF in Figure 4A confirms that the uncertainty shocks can lead to more protracted movements in output than the TFP shock in Figure 4B. The uncertainty shock generates a fall in the fraction of adjusting firms, leading to a misallocation of capital that takes time to wash out of the economy. Second, this persistence in output is not shared by the standard deviation of sales growth. The half-life of the output response is about twice that of sales growth dispersion. However when we simply eyeball the data in Figure 1, it does not seem like the standard deviation of sales growth has a particularly different persistence to that of output. This suggests that even if the output responses to an

Figure 4: Responses of aggregates and dispersion to Uncertainty and TFP shocks

(a) Responses to a positive Uncertainty shock (X_t^{σ})

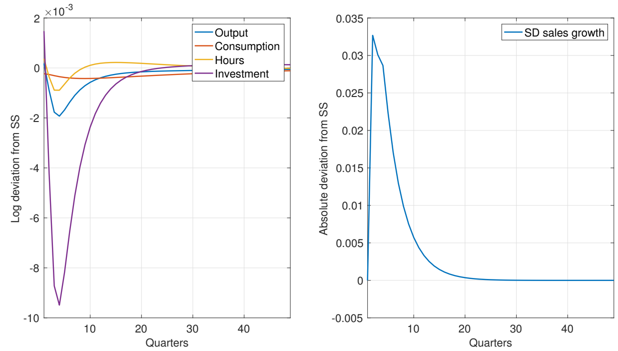

(b) Responses to a positive TFP shock (X_t^{Z})

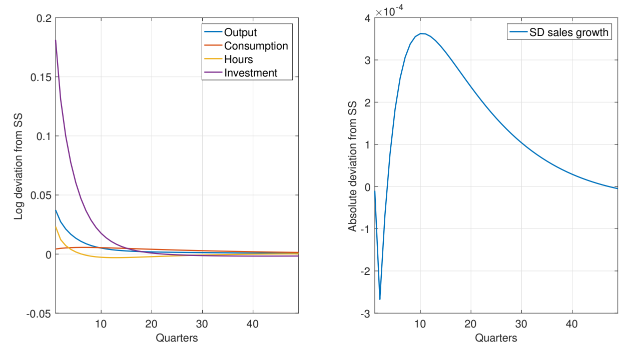

uncertainty shock were larger, the fact that the persistence of these two observable series is different in the model, but similar in the data, would make an uncertainty shock an unlikely explanation for the business cycle.

Finally, note that even though the magnitude of the response of sales growth dispersion to a TFP shock is small—which leads the estimation to reject TFP shocks as a meaningful source of fluctuations in the dispersion of sales growth—the response is also *pro-cyclical.* This is also counterfactual and runs against the intuitive explanation for how TFP shocks can generate counter-cyclical dispersion which we provided in the introduction. Below we discuss this feature of the model, and how it is sensitive to the model's calibration.

5.3 How can these results change?

Our results are that the model essentially separates into a model where aggregate shocks move aggregate variables and the uncertainty shock moves dispersion in the distribution of firms. As described above, the results hinge on the magnitude and direction of the response of sales growth dispersion to RBC shocks, and the magnitude of the response of aggregates in response to an uncertainty shock. Are there other parameterizations of the model that can undo either of these and so undo this 'separation' result? And if so, what are the features of the data that move the estimation of the model away from these parameterizations?

Two parameters that we consider here are the size of the adjustment cost $\bar{\xi}$ and the persistence of firm level shocks ρ_z. In Figure 5 we plot the impulse responses of the standard deviation of sales growth (panels A and B) and output (panels C and D) in response to a positive uncertainty shock and positive TFP shock, where in each figure we consider four different pairs of values for these parameters. We consider cases where we set $\bar{\xi} = 0.05$ and $\rho_z = 0.9$, once at a time, and then together. Panels A and D show that these parameters

Figure 5: Forecast error variance decomposition of cross-sectional variables

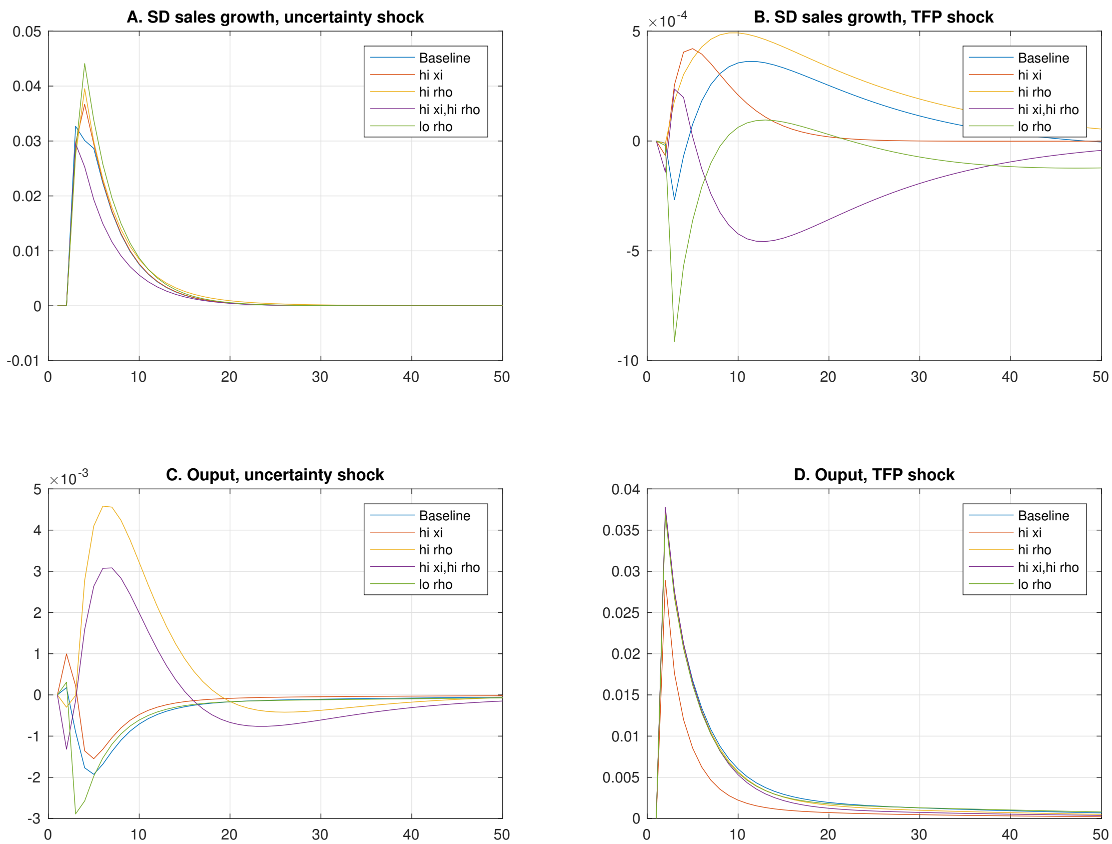

do not matter much for the aggregate dynamics of output (dispersion) in response to first (second) order shocks.[19]

Panel B shows that under a calibration with a high fixed cost of investment *and* a high persistence in productivity, we can get a *counter-cyclical* standard deviation of sales growth from a TFP shock. When both of these parameters are high, firm investment is very lumpy (high $\bar{\xi}$), but also responds more to idiosyncratic shocks (high ρ_z). Therefore, in steady state, only firms with a very high z and low k will invest. Only some firms invest, and these firms grow a lot. In this environment an increase in aggregate productivity—which is also persistent—encourages firms with lower z and higher k to invest. That is, some additional firms invest that are growing less than the firms that were already investing in steady state, while the growth rates of firms that were already investing in steady state do not change

[19] With respect to panel D this is a small result in itself, since it suggests that the response of the economy following TFP shocks is to some extent independent of the parametric details of the firm's problem.

substantially. This lowers the *dispersion* in growth rates.

This is a step in the right direction, relative to the data, but the magnitudes are still small. A one standard deviation TFP shock causes an decline in σ_t^{sg} around a tenth of the magnitude of a one standard deviation shock to X_t^σ. The reason that such high fixed cost, high persistence parameters were rejected when estimating the steady state of the model was that they would imply counterfactually high inaction rates. In other words, investment in the data is *not lumpy enough*—the investment frictions are not large enough—to generate conter-cyclical dispersion from first moment shocks.

Panel C shows that changing these parameters—at least by these magnitudes—cannot increase the response of output to uncertainty shocks. In fact, the figure shows that a necessary condition for a negative response of output to an uncertainty shock is a low persistence of firm level shocks. In the two cases where ρ_z is high, the response of output to an uncertainty shock is positive. We now discuss why an increase in uncertainty generates a recession only when firm-level productivity persistence is low.

5.4 Uncertainty or dispersion?

As we noted in the introduction to the paper, the term *uncertainty shock* is a slight misnomer with respect to the shock we model in this paper, and which has been studied in the firm investment literature. The uncertainty shock plays a dual role. Yes, it increases the uncertainty of the firm with regards to the variance of the next period productivity innovation. But, since this increase in the variance of innovations is realized, it also leads to a mean preserving spread in the distribution of firm level productivities. If firms face adjustment frictions, and productivity is persistent, then an increase in dispersion can lead investment to jump up at firms with even higher productivities. Meanwhile firms with low productivities no longer invest, reducing the pressure on wages—a general equilibrium effect which compounds the positive effects of the shock at high productivity firms. This could potentially lead to an increase in output if the dispersion component of the shock is the most important, and firm level productivity is very persistent: the economy benefits from the up-side of the increase in dispersion.

We get a small decline in output due to the uncertainty shock in this paper, which we want to understand, so find it useful to decompose the effect of the shock into these two parts. In Figure 6, we plot the response of output and investment to "news-only" uncertainty along with the original uncertainty shock. The "news-only" uncertainty is a shock to the beliefs of agents about future volatility *without* the corresponding realization of increased volatility. We repeat this under the baseline, low, and high values of firm level persistence

Figure 6: Response of output to an increase in "news-only" uncertainty

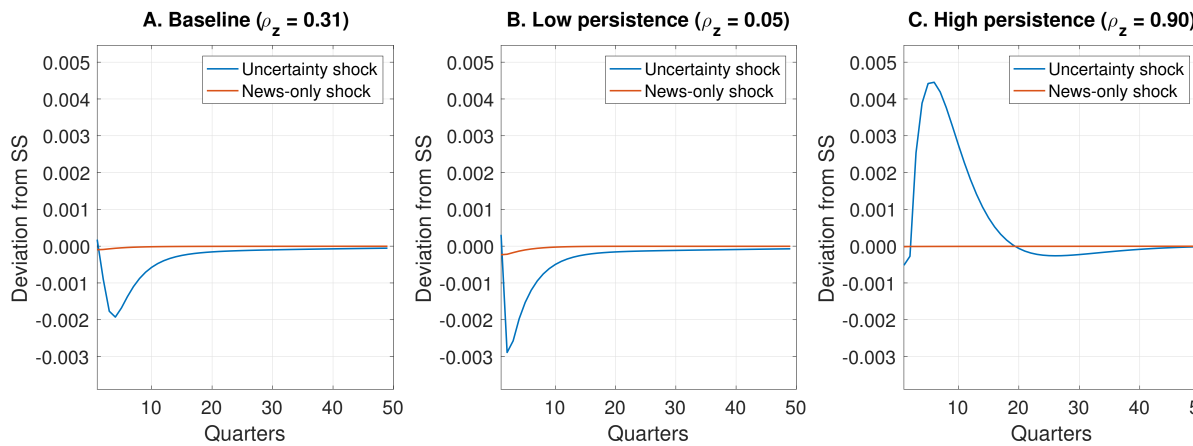

ρ_z.

Two clear points emerge. First, the response to the news shock is small. That is, the *wait-and-see effect* accounts for less than a tenth of the decline in output. Most of the effect is driven purely by the increased dispersion in realized productivities. This means that the misnomer of uncertainty shock is serious, since really the effects are coming from the dispersion component of the shock. Second, if firm level productivity is very persistent (panel C), then the positive dispersion shock generates a boom. To generate larger declines in output following an increase in dispersion, we require a counter-factual level of ρ_z which is below our baseline estimate (panel B). These findings push against the importance of the wait-and-see effect, and show the conditions under which a dispersion shock generates large negative co-movements between output and the standard deviation of sales growth, involve counter-factually low persistence of firm level shocks.

5.5 Estimation using dispersion in investment rates

In the above results we used the standard deviation of sales growth as an observable time series in the estimation of the model. In this section we consider using an alternative measure: the standard deviation of investment rates. As opposed to the standard deviation of sales growth, this measure is strongly pro-cyclical. This can be seen in Figure 7. Figure 10 shows that the fraction of firms adjusting their capital stock is also strongly pro-cyclical, a fact documented by Bachmann and Bayer (2014).

- *Results from this exercise are forthcoming.*

Figure 7: Procyclicality of dispersion of investment rate (conditional on adjustment)

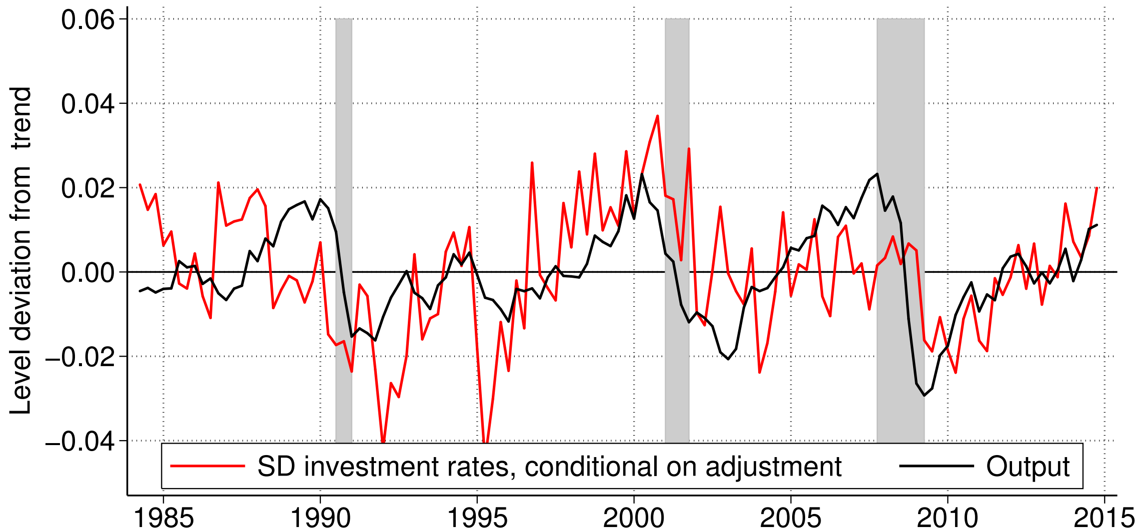

6 Conclusion

The cyclicality of firm dispersion suggests that firm dispersion may be a driving factor of business cycles or an endogenous response to business cycles or some combination of both. We build a general equilibrium model of firm investment and financing rich enough to quantify the relative strengths of these forces. We include investment and financing frictions so shock standard business cycle shocks (a TFP shock, a labor supply shock and a time preference shock) may have effects on the dispersion of firm-level sales growth and investment rates. We allow the model to respond to an uncertainty shock, i.e. a shock to the distribution of firm-level productivities.

We use a novel solution method that allows us to characterize the model solution in linear form, for which estimation is computationally feasible. We find that uncertainty shocks account for almost all of the variation in firm-level dispersion, and that standard business cycle shocks account for all of the movement in aggregates. The reason is that, at parameter values of the model which are consistent with the dispersion in investment rates, sales growth, and investment inaction in the data, the model cannot generate (i) large movements in sales growth dispersion from 'RBC' shocks, or (ii) generate large movements in output from uncertainty shocks. This leads the estimation to prefer parameter values at which the model decomposes into one in which fluctuations in first (second) order moments of the distribution of firms are generated by first (second) moment shocks.

References

Yann Algan, Olivier Allais, Wouter J. Den Haan, and Pontus Rendahl. Chapter 6 - solving and simulating models with heterogeneous agents and aggregate uncertainty. In Karl Schmedders and Kenneth L. Judd, editors, *Handbook of Computational Economics Vol. 3*, volume 3 of *Handbook of Computational Economics*, pages 277 – 324. Elsevier, 2014. doi: http://dx.doi.org/10.1016/B978-0-444-52980-0.00006-2. URL http://www.sciencedirect.com/science/article/pii/B9780444529800000062.

Sungbae An and Frank Schorfheide. Bayesian analysis of dsge models. *Econometric Reviews*, 26(2-4):113–172, 2007. doi: 10.1080/07474930701220071. URL http://dx.doi.org/10.1080/07474930701220071.

Rüdiger Bachmann and Christian Bayer. Uncertainty business cycles - really? *WP*, 2011.

Rüdiger Bachmann and Christian Bayer. Investment dispersion and the business cycle. *American Economic Review*, 104(4):1392–1416, April 2014. doi: 10.1257/aer.104.4.1392. URL http://www.aeaweb.org/articles?id=10.1257/aer.104.4.1392.

Rüdiger Bachmann and Giuseppe Moscarini. Business cycles and endogenous uncertainty. *WP*, 2012.

Susanto Basu and Brent Bundick. Uncertainty shocks in a model of effective demand. Working Paper 18420, National Bureau of Economic Research, September 2012. URL http://www.nber.org/papers/w18420.

Martin Beraja, Erik Hurst, and Juan Ospina. The aggregate implications of regional business cycles. Working Paper 21956, National Bureau of Economic Research, February 2016. URL http://www.nber.org/papers/w21956.

Nicholas Bloom. The Impact of Uncertainty Shocks. *Econometrica*, 77(3):623–685, 05 2009. URL https://ideas.repec.org/a/ecm/emetrp/v77y2009i3p623-685.html.

Nicholas Bloom, Max Floetotto, Nir Jaimovich, Itay Saporta-Eksten, and Stephen J. Terry. Really Uncertain Business Cycles. NBER Working Papers 18245, National Bureau of Economic Research, Inc, July 2012. URL https://ideas.repec.org/p/nbr/nberwo/18245.html.

John Y. Campbell and John H. Cochrane. By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, 107(2):205–251, 1999. doi: 10.1086/250059. URL http://dx.doi.org/10.1086/250059.

V. V. Chari, Patrick J. Kehoe, and Ellen R. McGrattan. Business cycle accounting. *Econometrica*, 75(3):781–836, 2007. ISSN 1468-0262. doi: 10.1111/j.1468-0262.2007.00768.x. URL http://dx.doi.org/10.1111/j.1468-0262.2007.00768.x.

Raj Chetty, Adam Guren, Dayanand S. Manoli, and Andrea Weber. Does indivisible labor explain the difference between micro and macro elasticities? a meta-analysis of extensive margin elasticities. Working Paper 16729, National Bureau of Economic Research, January 2011. URL http://www.nber.org/papers/w16729.

Lawrence J. Christiano, Martin Eichenbaum, and Charles Evans. Nominal rigidities and the dynamic effects of a shock to monetary policy. Working Paper 8403, National Bureau of Economic Research, July 2001. URL http://www.nber.org/papers/w8403.

Lawrence J. Christiano, Roberto Motto, and Massimo Rostagno. Risk shocks. *American Economic Review*, 104(1):27–65, January 2014. doi: 10.1257/aer.104.1.27. URL http://www.aeaweb.org/articles?id=10.1257/aer.104.1.27.

Gianluca Clementi and Dino Palazzo. Entry, exit, firm dynamics, and aggregate fluctuations, September 2010. Mimeo, New York University.

Jesús Fernández-Villaverde, Juan F. Rubio Ramírez, and Frank Schorfheide. Solution and estimation methods for dsge models. Working Paper 21862, National Bureau of Economic Research, January 2016. URL http://www.nber.org/papers/w21862.

Jordi Gali. Monetary policy and unemployment. In *Handbook of Monetary Economics*. Elsevier, 2011.

Andrew Gelman. Prior distributions for variance parameters in hierarchical models (comment on article by browne and draper). *Bayesian Anal.*, 1(3):515–534, 09 2006. doi: 10.1214/06-BA117A. URL http://dx.doi.org/10.1214/06-BA117A.

Andrew Gelman and Donald B. Rubin. Inference from iterative simulation using multiple sequences. *Statist. Sci.*, 7(4):457–472, 11 1992. doi: 10.1214/ss/1177011136. URL http://dx.doi.org/10.1214/ss/1177011136.

Simon Gilchrist, Jae W. Sim, and Egon Zakrajšek. Uncertainty, financial frictions, and investment dynamics. Working Paper 20038, National Bureau of Economic Research, April 2014. URL http://www.nber.org/papers/w20038.

Robert E. Hall and Paul R. Milgrom. The limited influence of unemployment on the wage bargain. *American Economic Review*, 2008.

Ayse Imrohoroglu and Selale Tuzel. Firm-level productivity, risk, and return. *Management Science*, 60(8):2073–2090, 2014. doi: 10.1287/mnsc.2013.1852. URL http://dx.doi.org/10.1287/mnsc.2013.1852.

Alejandro Justiniano, Giorgio Primiceri, and Andrea Tambalotti. Investment Shocks and the Relative Price of Investment. *Review of Economic Dynamics*, 14(1):101–121, January 2011. URL https://ideas.repec.org/a/red/issued/09-248.html.

Aubhik Khan and Julia K. Thomas. Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics. *Econometrica*, 76(2):395–436, 03 2008. URL https://ideas.repec.org/a/ecm/emetrp/v76y2008i2p395-436.html.

Per Krusell and Anthony A. Smith. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, 1998. doi: 10.1086/250034. URL http://dx.doi.org/10.1086/250034.

Alisdair McKay and Ricardo Reis. The Role of Automatic Stabilizers in the U.S. Business Cycle. NBER Working Papers 19000, National Bureau of Economic Research, Inc, April 2013. URL https://ideas.repec.org/p/nbr/nberwo/19000.html.

Michael Reiter. Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665, March 2009. URL https://ideas.repec.org/a/eee/dyncon/v33y2009i3p649-665.html.

Christopher A. Sims. Solving linear rational expectations models. *Computational Economics*, 20(1):1–20, 2002. ISSN 1572-9974. doi: 10.1023/A:1020517101123. URL http://dx.doi.org/10.1023/A:1020517101123.

Frank Smets and Rafael Wouters. Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach. CEPR Discussion Papers 6112, C.E.P.R. Discussion Papers, February 2007. URL https://ideas.repec.org/p/cpr/ceprdp/6112.html.

Joseph Vavra. Inflation dynamics and time-varying volatility: New evidence and an ss interpretation. *The Quarterly Journal of Economics*, 129(1):215–258, 2014. doi: 10.1093/qje/qjt027. URL http://qje.oxfordjournals.org/content/129/1/215.abstract.

Thomas Winberry. A toolbox for solving and estimating heterogeneous agent macro models. *WP*, 2016a.

Thomas Winberry. Lumpy investment, business cycles, and stimulus policy. Technical report, University of Chicago - Booth Business School, 2016b.

Eric R. Young. Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm and non-stochastic simulations. *Journal of Economic Dynamics and Control*, 34(1):36 – 41, 2010. ISSN 0165-1889. doi: http://dx.doi.org/10.1016/j.jedc.2008.11.010. URL http://www.sciencedirect.com/science/article/pii/S0165188909001316. Computational Suite of Models with Heterogeneous Agents: Incomplete Markets and Aggregate Uncertainty.

Appendix

This appendix is organized as follows Section A provides more details on the solution method. Section B provides more details on our construction of the moments used in the estimation of the model, the data used, how it is cleaned, and various adjustments made. Section C provides all figures referenced, but not included, in the main text.

A Method details

A.1 Discretized equilibrium conditions

We reproduce in this section the full set of equilibrium conditions in terms of the finite representation of the model state. We follow the Khan and Thomas (2008) procedure of writing the Bellman equation in terms of an adjusted value function which is scaled by marginal utility $\hat{v}_t \equiv u'(C_t)\,\tilde{v}_t$. Let Φ^v be the (time-invariant) basis function for the cubic spline that approximates the adjusted value function, and let θ_t^V be the corresponding vector of approximating coefficients. Denote by K the set of approximating points for capital k, and by Z the set of approximating points for productivity z.

(Bellman equation.) For all $k_i \in K$ and $z_j \in Z$,

$$
\begin{aligned}
\Phi^V(k_i, z_j)\,\theta_t^V &= \frac{\xi_t^*(k_i, z_j)}{\bar{\xi}}\left[V_t^{\text{adj}}(k_i, z_j) - \frac{1}{2}W_t\xi_t^*(k_i, z_j)\right] + \left(1 - \frac{\xi_t^*(k_i, z_j)}{\bar{\xi}}\right)V_t^{\text{stay}}(k_i, z_j) \\
\xi_t^*(k_i, z_j) &= \frac{V_t^{\text{adj}}(k_i, z_j) - V_t^{\text{stay}}(k_i, z_j)}{W_t} \\
V_t^{\text{adj}}(k_i, z_j) &= p_t\left(d^{\text{adj}}(k_i, z_j) - \mathbf{1}\left\{d^{\text{adj}}(k_i, z_j) < 0\right\}X_t^{\phi}\phi\left(d^{\text{adj}}(k_i, z_j)\right)^2\right) \\
&\quad + X_t^{\beta}\beta\sum_{z' \in Z}P(z, z'; X_t^{\sigma})\,\Phi^V\left(k^{\text{adj}}(k_i, z_j), z'\right)\theta_{t+1}^V \\
V_t^{\text{stay}}(k_i, z_j) &= p_t\left(d^{\text{stay}}(k_i, z_j) - \mathbf{1}\left\{d^{\text{stay}}(k_i, z_j) < 0\right\}X_t^{\phi}\phi\left(d^{\text{stay}}(k_i, z_j)\right)^2\right) \\
&\quad + X_t^{\beta}\beta\sum_{z' \in Z}P(z, z'; X_t^{\sigma})\,\Phi^V\left(k^{\text{stay}}(k_i, z_j), z'\right)\theta_{t+1}^V \\
d_t^{\text{adj}}(k_i, z_j) &= \pi_t(k_i, z_j) - \left(k_t^{\text{adj}}(k_i, z_j) - (1-\delta)k_i\right) \\
d_t^{\text{stay}}(k_i, z_j) &= \pi_t(k_i, z_j) - \left(k_t^{\text{stay}}(k_i, z_j) - (1-\delta)k_i\right) \\
\pi_t(k_i, z_j) &= X_t^Z z_j\left(n_t(k_i, x_j)^{\nu}k_i^{1-\nu}\right)^{\kappa} - W_t n_t(k_i, x_j) - \chi \\
n_t(k_i, z_j) &= \left(X_t^Z z_j \nu\kappa k_i^{\kappa(1-\nu)}W_t^{-1}\right)^{\frac{1}{1-\nu\kappa}}
\end{aligned}
$$

(Firm optimality.) For all $k_i \in K$ and $z_j \in Z$,

$$
p_t\left(1 - 2\phi d_t^{adj}(k_i, z_j)\,\mathbf{1}\left\{d_t^{adj}(k_i, z_j) < 0\right\}\right) = \beta\sum_{z' \in Z}P(z, z')\,\frac{\partial\Phi^V}{\partial k}\left(k^{adj}(k_i, z), z'\right)\theta_t^V
$$

(Law of motion.) For all $k_{i'} \in K$ and $z_{j'} \in Z$,

$$\lambda_t\left(k_{i'}, z_{j'}\right) = \sum_{k_i \in K} \sum_{z_j \in Z} Q\left(k_{i'}, z_{j'}, k_i, z_j\right) \lambda_{t-1}\left(k_i, z_j\right)$$

$$Q\left(k_{i'}, z_{j'}, k_i, z_j\right) = P\left(z_{j'}, z_j\right) \times \begin{cases} \xi^*\left(k_i, z_j\right) \times \frac{k_{i'+1} - k^{adj}\left(k_i, z_j\right)}{k_{i'+1} - k_{i'}} & \text{if } k^{adj}\left(k_i, z_j\right) \in \left[k_{i'}, k_{i'+1}\right] \\ \xi^*\left(k_i, z_j\right) \times \frac{k^{adj}\left(k_i, z_j\right) - k_{i'-1}}{k_{i'} - k_{i'-1}} & \text{if } k^{adj}\left(k_i, z_j\right) \in \left[k_{i'-1}, k_{i'}\right] \\ \left(1 - \xi^*\left(k_i, z_j\right)\right) \times \frac{k_{i'+1} - k^{stay}\left(k_i, z_j\right)}{k_{i'+1} - k_{i'}} & \text{if } k^{stay}\left(k_i, z_j\right) \in \left[k_{i'}, k_{i'+1}\right] \\ \left(1 - \xi^*\left(k_i, z_j\right)\right) \times \frac{k^{stay}\left(k_i, z_j\right) - k_{i'-1}}{k_{i'} - k_{i'-1}} & \text{if } k^{stay}\left(k_i, z_j\right) \in \left[k_{i'-1}, k_{i'}\right] \end{cases}$$

(Labor market clearing.)

$$N_t = \sum_{k_i \in K} \sum_{z_j \in Z} n_t\left(k_i, z_j\right) \lambda_{t-1}\left(k_i, z_j\right) + \sum_{k_i \in K} \sum_{z_j \in Z} \xi \mathbf{1}\left\{\xi \le \xi^*\left(k_i, z_j\right)\right\} \lambda_{t-1}\left(k_i, z_j\right)$$

(Wage law of motion.)

$$W_t = \omega W_t^* + (1 - \omega) W_{t-1}$$

$$W_t^* = \frac{\psi X_t^{\psi} N_t^{\eta}}{p_t}$$

(Stochastic discount factor.)

$$p_t = \left(C_t S_t\right)^{-\sigma}$$

(Habit stock law of motion.)

$$\log S_t = \left(1 - \rho^S\right) \log \bar{S} + \rho^S \log S_{t-1} + \lambda^S \log\left(C_t / C_{t-1}\right)$$

(Resource constraint.)

$$C_t = Y_t - I_t - A_t$$

$$Y_t = \sum_{k_i \in K} \sum_{z_j \in Z} X_t^Z z_i \left(n\left(k_i, z_j\right)^{\nu} k_i^{1-\nu}\right)^{\kappa} \lambda_{t-1}\left(k_i, z_j\right)$$

$$I_t = \sum_{k_i \in K} \sum_{z_j \in Z} \frac{\xi^*\left(k_i, z_j\right)}{\bar{\xi}} \left(k^{adj}\left(k_i, z_j\right) - (1 - \delta) k_i\right) \lambda_{t-1}\left(k_i, z_j\right) + \sum_{k_i \in K} \sum_{z_j \in Z} \left(1 - \frac{\xi^*\left(k_i, z_j\right)}{\bar{\xi}}\right) \left(k^{stay}\left(k_i, z_j\right) - (1 - \delta) k_i\right) \lambda_{t-1}\left(k_i, z_j\right)$$

$$A_t = \sum_{k_i \in K} \sum_{z_j \in Z} \frac{\xi^*\left(k_i, z_j\right)}{\bar{\xi}} \frac{1}{2} W_t \xi^*\left(k_i, z_j\right) \lambda_{t-1}\left(k_i, z_j\right)$$

B Data

B.1 Data

The investment and issuance data targets are computed for the sample of firms in the Compustat North American quarterly database for the period 1985:I to 2014:IV.

We construct a measure of firm capital using a perpetual the inventory method, following Clementi and Palazzo (2010). For firm j, we set the initial value of the capital stock $k_{j,t}$ to the first available entry for $\text{PPEGT}_{j,t}$ (gross value of property, plant and equipment), and recursively construct the value for the capital stock using $\text{PPENTQ}_{j,t}$ (net value of property, plant and equipment)

$$k_{j,t+1} = k_{j,t} + \text{PPENT}_{j,t+1} - \text{PPENT}_{j,t},$$

where we interpolate linearly for $\text{PPENT}_{j,t}$ wherever $\text{PPENT}_{j,t}$ is missing and $\text{PPENT}_{j,t+1}$ and $\text{PPENT}_{j,t-1}$ are available.

We define the investment rate for firm j in period t, $\text{ir}_{j,t}$, as:

$$\text{ir}_{j,t} = \frac{k_{j,t} - k_{j,t-1}}{\frac{1}{2}\left(k_{j,t} + k_{j,t-1}\right)}$$

and sales growth $g_{j,t}$

$$\text{sg}_{j,t} = \frac{\text{s}_{j,t} - \text{s}_{j,t-1}}{\frac{1}{2}\left(\text{s}_{j,t} + \text{s}_{j,t-1}\right)},$$

where for sales $s_{j,t}$ we use the variable $\text{SALES}_{j,t}$.

For the external financing variables, we first construct an unadjusted measure of net cashflow from external finance. For unadjusted net cashflow from external financing, $\hat{\text{ef}}_{j,t}$, we use $\text{FINCFY}_{j,t}$ (net cash flow from financing activities) if it is available. If it is not available, we use the sum of its constituent parts,

$$\hat{\text{ef}}_{j,t} = \text{SSTK}_{j,t} - \text{PRSTKC}_{j,t} - \text{DV}_{j,t} + \text{DLTIS}_{j,t} - \text{DLTR}_{j,t} + \text{DLCCH}_{j,t} + \text{FIAO}_{j,t},$$

where $\text{SSTKY}_{j,t}$ is net cash flow from sale of common and preferred stock, $\text{PRSTKCY}_{j,t}$ is net cash flow from purchase of common and preferred stock, $\text{DVY}_{j,t}$ is net cash flow from cash dividends, $\text{DLTIS}_{j,t}$ is net cash flow from long-term debt issuance, $\text{DLTR}_{j,t}$ is net cash flow from long-term debt reduction, $\text{DLCCH}_{j,t}$ is net cash flow from changes in current debt and $\text{FIAO}_{j,t}$ is net cash flow from other financing activities. If any of the constituent parts is missing, we replace that part with zero.

In the model, firms raise external finance in order to invest, to hire labor or to pay costs.

In the data, however, a significant portion of net financing activity involves funding financial investment (that is, buying financial assets). We adjust the measure of external financing to account for this: we define adjusted external financing, $\mathrm{ef}_{j,t}$, as unadjusted external financing less net cashflow due to increases in financial assets:

$$\mathrm{ef}_{j,t} = \hat{\mathrm{ef}}_{j,t} - \Big(-\mathrm{IVCH}_{j,t} + \mathrm{SIV}_{j,t} + \mathrm{IVSTCH}_{j,t} - \mathrm{AQC}_{j,t} + \mathrm{IVACO}_{j,t} + \mathrm{CHECH}_{j,t}\Big)$$

where $\mathrm{IVCH}_{j,t}$ is net cashflow from increase in investments, $\mathrm{SIV}_{j,t}$ is net cashflow from sale of investments, $\mathrm{IVSTCH}_{j,t}$ is net cashflow from change in short-term investments, $\mathrm{AQC}_{j,t}$ is net cashflow from acquisitions, $\mathrm{IVACO}_{j,t}$ is net cashflow from other investing activities, and $\mathrm{CHECH}_{j,t}$ is increase in cash and cash equivalents.

B.1.1 Quarterly effects

Before computing moments and constructing the time series, we remove quarter effects at the firm level. (Note that we remove the effect of a particular quarter of the year $q \in \{1, 2, 3, 4\}$ not the effect of a particular time period.)

For each of $\mathrm{ir}_{j,t}$, $\mathrm{sg}_{j,t}$ and $\mathrm{ef}_{j,t}$, let x_{jyq} be the value for firm j in year y and quarter q. We first drop incomplete firm-years. We then define the firm-year average,

$$\bar{x}_{iy} \equiv \frac{1}{4}\sum_{q=1}^{4} x_{iqy}, \qquad \forall i \in I, y \in Y,$$

where, I is the set of all firms, and Y is is the set of all years in the sample. We then compute each firm-quarter's deviation from the firm-year average

$$h_{iqy} \equiv x_{iqy} - x_{iy}, \qquad \forall i \in I, q \in \{1, 2, 3, 4\}, y \in Y,$$

We compute firm i's average quarter-q deviation across years,

$$\gamma_{iq} \equiv \frac{1}{T_i}\sum_{y \in Y_i} h_{iqy}, \qquad \forall i \in I, q \in \{1, 2, 3, 4\},$$

where Y_i is the set of years for which firm i has observations. We define each firm i's excess deviation for a period qy as its actual quarterly deviation in qy less its average quarterly deviation in quarter q,

$$\delta_{iqy} = h_{iqy} - \gamma_{iq}, \qquad \forall i \in I, q \in \{1, 2, 3, 4\}, y \in Y$$

and define the quarterly-adjusted value as the firm-year mean plus the firm-quarter's excess deviation:

$$\tilde{x}_{iqy} = \bar{x}_{iy} + \delta_{iqy}.$$

B.1.2 Industry effects

We remove industry fixed effects at the level of 2-digit SIC industries.

C Figures

Figure 8: Prior and posterior densities

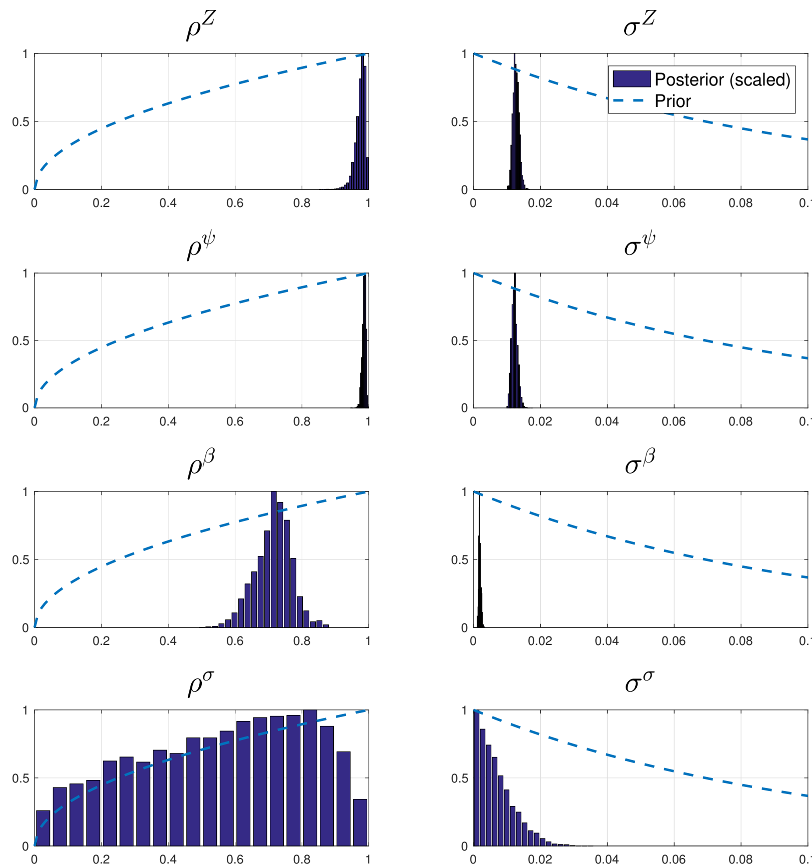

Prior density: dashed line. Posterior density: histogram.

Table 4: Forecast error variance decomposition, 1 quarter ahead

Shock	**Output**	**Consumption**	**Labor**	**SD sales growth**
TFP	0.36	0.23	0.03	0.05
Labor supply	0.56	0.36	0.85	0.08
Time preference	0.08	0.40	0.12	0.86
Uncertainty	0.00	0.00	0.00	0.00

Table 5: Forecast error variance decomposition, 2 quarters ahead

Shock	**Output**	**Consumption**	**Labor**	**SD sales growth**
TFP	0.37	0.31	0.03	0.00
Labor supply	0.57	0.47	0.89	0.00
Time preference	0.06	0.22	0.08	0.00
Uncertainty	0.00	0.00	0.00	1.00

Figure 9: Aggregate time series.

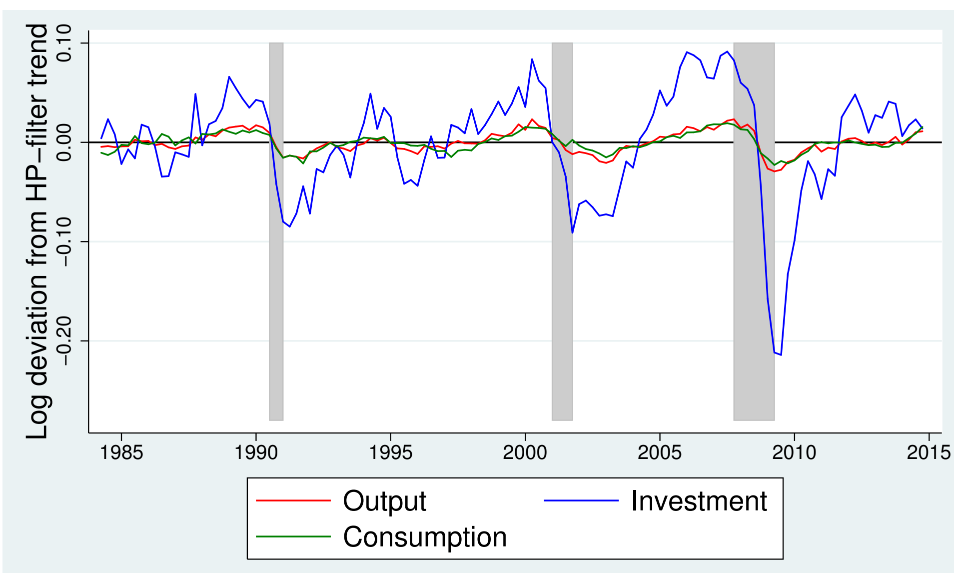

Figure 10: Procyclicality of fraction of adjusting firms.

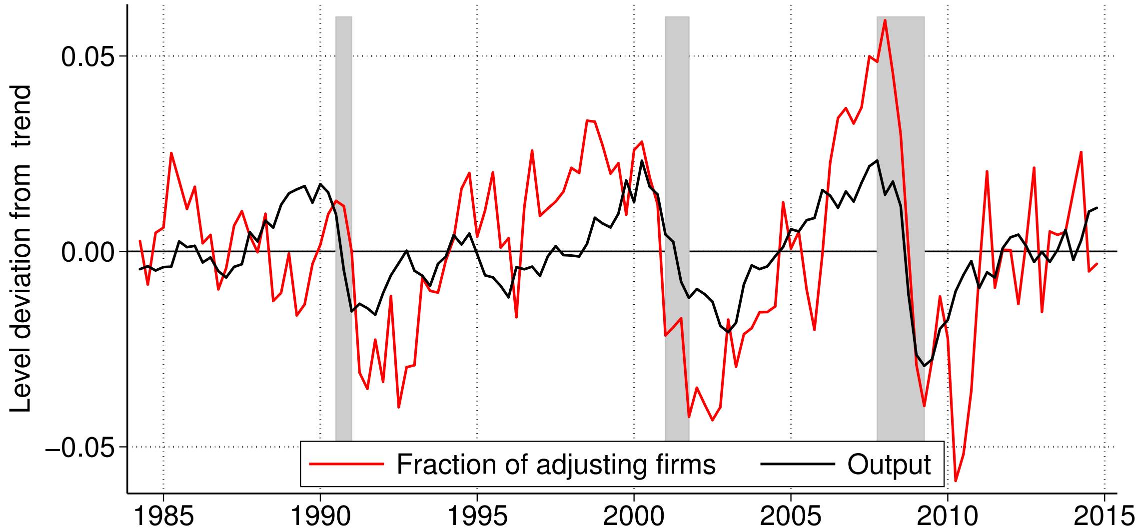

Figure 11: Impulse reponse functions for all shocks