

Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model

Ruediger Bachmann Ricardo J. Caballero Eduardo M.R.A. Engel*

April 12, 2010

Abstract

The sensitivity of U.S. aggregate investment to shocks is procyclical: the response upon impact increases by approximately 50% from the trough to the peak of the business cycle. This feature of the data follows naturally from a DSGE model with lumpy microeconomic capital adjustment. Beyond explaining this specific time variation, our model and evidence provide a counterexample to the claim that microeconomic investment lumpiness is inconsequential for macroeconomic analysis.

JEL Codes: E10, E22, E30, E32, E62.


Keywords: Ss model, RBC model, time varying impulse response function, history dependence, conditional heteroscedasticity, aggregate shocks, sectoral shocks, idiosyncratic shocks, adjustment costs.

*Respectively: University of Michigan; MIT and NBER; Yale University and NBER. We are grateful to Fabián Duarte, David Rappoport and José Tessada for excellent research assistance, and to Olivier Blanchard, William Brainard, Jordi Galí, Pete Klenow, John Leahy, Giuseppe Moscarini, Anthony Smith, Julia Thomas, three anonymous referees, and workshop participants at the AEA, Bonn, Cornell, Econometric Society, Karlsruhe, Mainz, NBER-EFG, NYU, SITE, Stanford, Toulouse, U. de Chile-CEA and Yale for helpful comments. Financial support from the NSF is gratefully acknowledged.

1 Introduction

U.S. nonresidential private fixed investment exhibits conditional heteroscedasticity. Figure 1 depicts a smooth, nonparametric, normalized estimate of the heteroscedasticity of the residual from fitting an AR process to quarterly aggregate investment rates from 1960 to 2005, as a function of the average recent investment rate (see Appendix B for details). This figure shows that investment is significantly more responsive to shocks in times of high investment.¹

Figure 1: Conditional Heteroscedasticity




In this paper we show that this nonlinear feature of the data follows naturally from a DSGE model with lumpy microeconomic investment. The reason for conditional heteroscedasticity in the model, is that the impulse response function is history dependent, with an initial response that increases by approximately 50% from the bottom to the peak of the business cycle. In particular, the longer an expansion, the larger the response of investment to further shocks. Conversely, recovering from investment slumps is hard.

More broadly, our calibrated model suggests that over the 1960-2005 period the average initial response of investment to a productivity shock in the top quartile is 32% higher than the average response in the bottom quartile. The left and center panels in Figure 2 depict the response over five quarters to a one standard deviation shock taking place at selected points of the U.S. investment cycle, for an ARCH-type time series model and our calibrated lumpy investment model, respectively.² The periods considered are the trough in 1961, a period of average investment activity in 1989 and the peak in 2000. The variability of these impulse responses is large and similar in the left and center panels. For example, the immediate response to a shock in the trough in 1961 and the peak in 2000 differ by roughly 50%. The contrast with the

¹The dotted lines depict one-standard error bands.

²See Appendix B for details on the time series model.

Figure 2: Impulse Response in Different Years - Time Series, Lumpy and Frictionless Models



right panel of this figure, which depicts the impulse responses for a model with no microeconomic frictions in investment (essentially, the standard RBC model), is evident: For the latter, the impulse responses vary little over time.³

Beyond explaining the rich nonlinear dynamics of aggregate investment rates, our model provides a counterexample to the claim that microeconomic investment lumpiness is inconsequential for macroeconomic analysis. This is relevant, since even though Caballero and Engel (1999) found substantial aggregate nonlinearities in a partial equilibrium model with lumpy capital adjustment, recent and important methodological contributions by Veracierto (2002), Thomas (2002) and Khan and Thomas (2003, 2008) have provided examples where general equilibrium undoes the partial equilibrium features.

Why do we reach a different conclusion? Because, implicitly, earlier calibrations imposed that the bulk of investment dynamics was determined by general equilibrium constraints rather than by adjustment costs. Instead, we focus our calibration effort on gauging the relative importance of these forces, and conclude that *both* adjustment costs and general equilibrium forces play a relevant role.

Our calibration begins by noting that the objective in any dynamic macroeconomic model is to trace the impact of aggregate shocks on aggregate endogenous variables (investment in our context). The typical response is less than one-for-one upon impact, as a variety of microeconomic frictions and general equilibrium constraints attenuate and spread over time the response of the endogenous variable. We refer to this process as *smoothing*, and decompose it into its partial equilibrium (PE) and general equilibrium (GE) components. In the context

³The figures in the three panels are normalized so that the impulse response in 1989:I is one upon impact.

of nonlinear lumpy-adjustment models, PE-smoothing does *not* refer to the existence of microeconomic inaction and lumpiness per se, but to their impact in smoothing the response of aggregates. This is a key distinction in this class of models, as in many instances microeconomic inaction translates into limited aggregate inertia (recall the classic Caplin and Spulber (1987) result, where price-setters follow *Ss* rules but the aggregate price level behaves as if there were no microeconomic frictions). In a nutshell, our key difference with the previous literature is that the latter explored combinations of parameter values that implied microeconomic lumpiness but left almost no role for PE-smoothing, thereby precluding the possibility of fitting facts such as the conditional heteroscedasticity of aggregate investment rates depicted in Figures 1 and 2.

Table 1: CONTRIBUTION OF PE AND GE FORCES TO SMOOTHING OF I/K

No frictions (0.0425) 0%		
	↓	
Only PE smoothing (0.0040) 81.0%		Only GE smoothing (0.0036) 84.6%
	↓	
PE and GE smoothing (0.0023) 100%		

Table 1 illustrates our model's decomposition into PE- and GE-smoothing. The lower entry shows quarterly volatility of aggregate investment rates in our model with adjustment costs and price responses. The upper entry reports this statistic when neither smoothing mechanism is present, that is, when adjustment costs are set to zero and prices to their average value in our model with both sources of smoothing. The intermediate entries consider only one source of smoothing at a time, for example, "only PE-smoothing" retains adjustment costs but sets prices at their average values in the economy that leads to the lower entry. The reduction of the quarterly standard deviation of the aggregate investment rate achieved by PE-smoothing alone amounts to 81.0% of the reduction achieved by the combination of both smoothing mechanisms. Alternatively, the additional smoothing achieved by PE-forces, compared with what GE-smoothing achieves by itself, is 15.4% of the total.

It is clear from Table 1 that both sources of smoothing do not enter additively, so some care is needed when quantifying their relative importance. Nonetheless, averaging the upper and lower bounds mentioned above suggests roughly similar roles for both. By contrast, as discussed in detail in Section 3, the contribution of PE-smoothing is very small in the recent literature—typically the upper bound is under 20% while the lower bound is zero.

Our calibration strategy is designed to capture the role of PE-smoothing as directly as possible. To this effect, we use sectoral data to calibrate the parameters that control the impact of micro-frictions on aggregates, *before* general equilibrium forces have a chance to play a significant smoothing role. Specifically, we argue that the response of semi-aggregated (e.g., 3-digit) investment to corresponding sectoral shocks is less subject to general equilibrium forces, and hence serves to identify the relative importance of PE-smoothing.

Table 2: VOLATILITY AND AGGREGATION

Model	3-digit	Aggregate	3-dig. Agg. Ratio
<i>Data</i>	<i>0.0163</i>	<i>0.0098</i>	<i>1.66</i>
This paper:	0.0163	0.0098	1.66
Frictionless:	0.1839	0.0098	18.77
Khan-Thomas (2008):	0.4401	0.0100	44.01

The first row in Table 2 shows the observed volatility of annual sectoral and aggregate investment rates, and their ratio.⁴ The second and third rows show these values for our baseline lumpy model and the model with no microeconomic frictions in investment, respectively. The fourth row reports these statistics for the model in Khan and Thomas (2008).⁵ It is apparent from this table that the frictionless model fails to match the sectoral data (it was never designed to do so). In contrast, by reallocating smoothing from GE- to PE-forces, the lumpy investment model is able to match both aggregate and sectoral volatility. This pins down our decomposition and is, together with the conditional-heteroscedasticity feature, the essence of our calibration strategy.

The remainder of the paper is organized as follows. In the next section we present our dynamic general equilibrium model. Section 3 discusses the calibration method in detail. Section 4 presents the main macroeconomic implications of the model. Section 5 concludes and is followed by several appendices.

⁴Sectoral investment data are only available at an annual frequency. The numbers in rows two and three come from the annual analogues of our quarterly baseline models. For details, see Appendices A.2 and A.3.

⁵The lumpy model in Kahn and Thomas (2008) exhibits larger sectoral volatility than the frictionless counterpart of our lumpy model because of differences in the curvature of the revenue function (see Section 3.1 for details). The volatility of aggregate investment rates in the Kahn-Thomas (2008) entry of Table 2 is taken from table III in their paper. The volatility of sectoral investment rates is based on our calculations.

2 The Model

In this section we describe our model economy. We start with the problem of the production units, followed by a brief description of the households and the definition of equilibrium. We conclude with a sketch of the equilibrium computation. We follow closely Kahn and Thomas (2008), henceforth KT, both in terms of substance and notation. Aside from parameter differences, we have three main departures from KT. First, production units face persistent sector-specific productivity shocks, in addition to aggregate and idiosyncratic shocks. Second, production units undertake some within-period maintenance investment which is necessary to continue operation (some parts and machines that break down need to be replaced, see, e.g., McGrattan and Schmitz (1999) and Letterie, Pfann and Verick (2004) for evidence on the importance of maintenance and replacement investment). Third, the distribution of aggregate productivity shocks is continuous rather than a Markov discretization, which allows us to back out the aggregate shocks that are fed into the model to produce Figures 2 and 3.

2.1 Production Units

The economy consists of a large number of sectors, which are each populated by a continuum of production units. Since we do not model entry and exit decisions, the mass of these continua is fixed and normalized to one. There is one commodity in the economy that can be consumed or invested. Each production unit produces this commodity, employing its pre-determined capital stock (k) and labor (n), according to the following Cobb-Douglas decreasing-returns-to-scale production function ($\theta > 0$, $\nu > 0$, $\theta + \nu < 1$):

$$y_t = z_t \epsilon_{S,t} \epsilon_{I,t} k_t^\theta n_t^\nu, \quad (1)$$

where z , ϵ_S and ϵ_I denote aggregate, sectoral and unit-specific (idiosyncratic) productivity shocks.

We denote the trend growth rate of aggregate productivity by $(1 - \theta)(\gamma - 1)$, so that y and k grow at rate $\gamma - 1$ along the balanced growth path. From now on we work with k and y (and later C) in efficiency units. The detrended aggregate productivity level, which we also denote by z , evolves according to an AR(1) process in logs, with normal innovations ν with zero mean and variance σ_A^2 :

$$\log z_t = \rho_A \log z_{t-1} + \nu_t. \quad (2)$$

The sectoral and idiosyncratic technology processes follow Markov chains, that are approximations to continuous AR(1) processes with Gaussian innovations.⁶ The latter have standard

⁶We use the discretization in Tauchen (1986), see Appendix C for details.

deviations σ_S and σ_I , and autocorrelations ρ_S and ρ_I , respectively. Productivity innovations at different aggregation levels are independent. Also, sectoral productivity shocks are independent across sectors and idiosyncratic productivity shocks are independent across productive units.

Each period a production unit draws from a time invariant distribution, G , its current cost of capital adjustment, $\xi \geq 0$, which is denominated in units of labor. G is a uniform distribution on $[0, \bar{\xi}]$, common to all units. Draws are independent across units and over time, and employment is freely adjustable.

At the beginning of a period, a production unit is characterized by its pre-determined capital stock, the sector it belongs to and the corresponding sectoral productivity level, its idiosyncratic productivity, and its capital adjustment cost. Given the aggregate state, it decides its employment level, n , production occurs, workers are paid, and investment decisions are made. Upon investment the unit incurs a fixed cost of $\omega\xi$, where ω is the current real wage rate. Capital depreciates at a rate δ and a fraction of depreciated capital is replaced to continue operation. Then the period ends.

We also introduce replacement and maintenance investment as an essential feature of actual production units. This is justified when each productive unit can be viewed as a composite of core and peripheral components, where core components need to be replaced immediately for the unit to continue production. Alternatively, maintaining certain components of a productive unit on a regular basis so that they do not depreciate at all, can be considerably more cost effective than using a stop-go approach to maintenance.⁷

We define $\bar{\psi} \equiv \frac{\gamma}{1-\delta} > 1$ as the investment rate needed to fully compensate depreciation and trend growth. The degree of necessary maintenance or replacement, χ , can then be conveniently defined as a fraction of $\bar{\psi}$. If $\chi = 0$, no maintenance investment is needed; if $\chi = 1$, all depreciation and trend growth must be replaced for a production unit to continue operation. We can now summarize the evolution of the unit's capital stock (in efficiency units) between two consecutive periods, from k to k' , after non-maintenance investment i and maintenance investment $i^M = \chi(\gamma - 1 + \delta)k$ take place, as follows:

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

If $\chi = 0$, then $k' = (1 - \delta)k/\gamma$, while $k' = k$ if $\chi = 100\%$. We treat χ as a primitive parameter.⁸

⁷For instance, maintaining the roof of a structure on a regular basis is likely to dominate over the alternative of repairing it only when it begins to leak.

⁸We note that our version of maintenance investment differs from what KT call "constrained investment". Here,

As we will discuss in Section 4, replacement and maintenance investment play an important role in shaping aggregate investment dynamics, since it determines the effective (i.e., after maintenance) depreciation rate. This differs from what happens with linear investment models, where the depreciation rate plays a minor role. We have introduced these determinants of investment in an admittedly stylized manner with a single structural parameter, and leave for future research a more detailed study of these issues.

Given the i.i.d. nature of the adjustment costs, it is sufficient to describe differences across production units and their evolution by the distribution of units over $(\epsilon_S, \epsilon_I, k)$. We denote this distribution by μ . Thus, (z, μ) constitutes the current aggregate state and μ evolves according to the law of motion $\mu' = \Gamma(z, \mu)$, which production units take as given.

Next we describe the dynamic programming problem of each production unit. We take two shortcuts (details can be found in KT). First, we state the problem in terms of utils of the representative household (rather than physical units), and denote by $p = p(z, \mu)$ the marginal utility of consumption. This is the relative intertemporal price faced by a production unit. Second, given the i.i.d. nature of the adjustment costs, continuation values can be expressed without explicitly taking into account future adjustment costs.

It will simplify notation to define an additional parameter, $\psi \in [1, \bar{\psi}]$:

$$\psi = 1 + (\bar{\psi} - 1)\chi, \quad (3)$$

and write maintenance investment as:⁹

$$i^M = (\psi - 1)(1 - \delta)k. \quad (4)$$

Let $V^1(\epsilon_S, \epsilon_I, k, \xi; z, \mu)$ denote the expected discounted value—in utils—of a unit that is in idiosyncratic state (ϵ_I, k, ξ) , and is in a sector with sectoral productivity ϵ_S , given the aggregate state (z, μ) . Then the expected value prior to the realization of the adjustment cost draw is given by:

$$V^0(\epsilon_S, \epsilon_I, k; z, \mu) = \int_0^{\bar{\xi}} V^1(\epsilon_S, \epsilon_I, k, \xi; z, \mu) G(d\xi). \quad (5)$$

With this notation the dynamic programming problem is given by:

$$V^1(\epsilon_S, \epsilon_I, k, \xi; z, \mu) = \max_n \{CF + \max_{k'} (V_i, \max[-AC + V_a])\}, \quad (6)$$

maintenance refers to the replacement of parts and machines without which production cannot continue, while in KT it is an extra margin of adjustment for small investment projects.

⁹Note that if $\psi = 1$, then $i^M = 0$, and if $\psi = \bar{\psi}$, then $i^M = (\gamma - 1 + \delta)k$, undoing all trend devaluation of the capital stock.

where CF denotes the firm's flow value, V_i the firm's continuation value if it chooses inaction and does not adjust, and V_a the continuation value, net of adjustment costs AC , if the firm adjusts its capital stock. That is:

$$CF = [z\epsilon_S\epsilon_I k^\theta n^\nu - \omega(z, \mu)n - i^M]p(z, \mu), \quad (7a)$$

$$V_i = \beta E[V^0(\epsilon'_S, \epsilon'_I, \psi(1 - \delta)k/\gamma; z', \mu')], \quad (7b)$$

$$AC = \xi\omega(z, \mu)p(z, \mu), \quad (7c)$$

$$V_a = -ip(z, \mu) + \beta E[V^0(\epsilon'_S, \epsilon'_I, k'; z', \mu')], \quad (7d)$$

where both expectation operators average over next period's realizations of the aggregate, sectoral and idiosyncratic shocks, conditional on this period's values, and we recall that $i^M = (\psi - 1)(1 - \delta)k$ and $i = \gamma k' - (1 - \delta)k - i^M$. Also, β denotes the discount factor from the representative household.

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

2.2 Households

We assume a continuum of identical households that have access to a complete set of state-contingent claims. Hence, there is no heterogeneity across households. Moreover, they own shares in the production units and are paid dividends. We do not need to model the household side explicitly (see KT for details), and concentrate instead on the first-order conditions to determine the equilibrium wage and the intertemporal price.

Households have a standard felicity function in consumption and (indivisible) labor:

$$U(C, N^h) = \log C - AN^h, \quad (8)$$

where C denotes consumption and N^h the fraction of household members that work. Households maximize the expected present discounted value of the above felicity function. By definition we have:

$$p(z, \mu) \equiv U_C(C, N^h) = \frac{1}{C(z, \mu)}, \quad (9)$$

and from the intratemporal first-order condition:

$$\omega(z, \mu) = -\frac{U_N(C, N^h)}{p(z, \mu)} = \frac{A}{p(z, \mu)}. \quad (10)$$

2.3 Recursive Equilibrium

A *recursive competitive equilibrium* is a set of functions

$$\left(\omega, p, V^1, N, K, C, N^h, \Gamma\right),$$

that satisfy

1. *Production unit optimality*: Taking ω , p and Γ as given, $V^1(\epsilon_S, \epsilon_I, k; z, \mu)$ solves (5) and the corresponding policy functions are $N(\epsilon_S, \epsilon_I, k; z, \mu)$ and $K(\epsilon_S, \epsilon_I, k, \xi; z, \mu)$.
2. *Household optimality*: Taking ω and p as given, the household's consumption and labor supply satisfy (8) and (9).
3. *Commodity market clearing*:

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

4. *Labor market clearing*:

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

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

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

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

2.4 Solution

As is well-known, (6) is not computable, since μ is infinite dimensional. Hence, we follow Krusell and Smith (1997, 1998) and approximate the distribution μ by its first moment over

capital, and its evolution, Γ , by a simple log-linear rule. In the same vein, we approximate the equilibrium pricing function by a log-linear rule:

$$\log \bar{k}' = a_k + b_k \log \bar{k} + c_k \log z, \quad (11a)$$

$$\log p = a_p + b_p \log \bar{k} + c_p \log z, \quad (11b)$$

where \bar{k} denotes aggregate capital holdings. Given (10), we do not have to specify an equilibrium rule for the real wage. As usual with this procedure, we posit this form and verify that in equilibrium it yields a good fit to the actual law of motion (see Appendix C for details).

To implement the computation of sectoral investment rates, we simplify the problem further and impose two additional assumptions: 1) $\rho_S = \rho_I = \rho$ and 2) enough sectors, so that sectoral shocks have no aggregate effects. Combining both assumptions reduces the state space in the production unit's problem further to a combined technology level $\epsilon \equiv \epsilon_S \epsilon_I$. Now, $\log \epsilon$ follows an AR(1) with first-order autocorrelation ρ and Gaussian innovations $N(0, \sigma^2)$, with $\sigma^2 \equiv \sigma_S^2 + \sigma_I^2$. Since the sectoral technology level has no aggregate consequences by assumption, the production unit cannot use it to extract any more information about the future than it has already from the combined technology level. Finally, it is this combined productivity level that is discretized into a 19-state Markov chain. The second assumption allows us to compute the sectoral problem independently of the aggregate general equilibrium problem.¹⁰

Combining these assumptions and substituting \bar{k} for μ into (6) and using (11a)–(11b), we have that (7a)–(7d) become

$$CF = [z\epsilon k^\theta n^\nu - \omega(z, \bar{k})n - i^M]p(z, \bar{k}), \quad (12a)$$

$$V_i = \beta E[V^0(\epsilon', \psi(1 - \delta)k/\gamma; z', \bar{k}')], \quad (12b)$$

$$AC = \xi \omega(z, \bar{k})p(z, \bar{k}), \quad (12c)$$

$$V_a = -ip(z, \bar{k}) + \beta E[V^0(\epsilon', k'; z', \bar{k}')]. \quad (12d)$$

With the above expressions, (6) becomes a computable dynamic programming problem with policy functions $N = N(\epsilon, k; z, \bar{k})$ and $K = K(\epsilon, k, \xi; z, \bar{k})$. We solve this problem via value function iteration on V^0 and Gauss-Hermitian numerical integration over $\log(z)$ (see Appendix C for details).

Several features facilitate the solution of the model. First, as mentioned above, the employment decision is static. In particular it is independent of the investment decision at the end of the period. Hence we can use the production unit's first-order condition to maximize out the

¹⁰In Appendix C.3 we show that our results are robust to this simplifying assumption.

optimal employment level:

$$N(\epsilon, k; z, \bar{k}) = \left(\frac{\omega(z, \bar{k})}{\nu z \epsilon k^\theta} \right)^{1/(\nu-1)}. \quad (13)$$

Next we comment on the computation of the production unit's decision rules and value function, given the equilibrium pricing and movement rules (11a)–(11b). From (12d) it is obvious that neither V_a nor the optimal target capital level, conditional on adjustment, depend on current capital holdings. This reduces the number of optimization problems in the value function iteration considerably. Comparing (12d) with (12b) shows that $V_a(\epsilon; z, \bar{k}) \geq V_i(\epsilon, k; z, \bar{k})$.¹¹ It follows that there exists an adjustment cost factor that makes a production unit indifferent between adjusting and not adjusting:

$$\hat{\xi}(\epsilon, k; z, \bar{k}) = \frac{V_a(\epsilon; z, \bar{k}) - V_i(\epsilon, k; z, \bar{k})}{\omega(z, \bar{k}) p(z, \bar{k})} \geq 0. \quad (14)$$

We define $\xi^T(\epsilon, k; z, \bar{k}) \equiv \min(\bar{\xi}, \hat{\xi}(\epsilon, k; z, \bar{k}))$. Production units with $\xi \leq \xi^T(\epsilon, k; z, \bar{k})$ will adjust their capital stock. Thus,

$$k' = K(\epsilon, k, \xi; z, \bar{k}) = \begin{cases} k^*(\epsilon; z, \bar{k}) & \text{if } \xi \leq \xi^T(\epsilon, k; z, \bar{k}), \\ \psi(1 - \delta)k/\gamma & \text{otherwise.} \end{cases} \quad (15)$$

We define *mandated investment* for a unit with current state (ϵ, z, \bar{k}) and current capital k as:

$$\text{Mandated investment} \equiv \log \gamma k^*(\epsilon; z, \bar{k}) - \log \psi(1 - \delta)k.$$

That is, mandated investment is the investment rate the unit would undertake, after maintaining its capital, if its current adjustment cost draw were equal to zero.

Now we turn to the second step of the computational procedure that takes the value function $V^0(\epsilon, k; z, \bar{k})$ as given, and pre-specifies a randomly drawn sequence of aggregate technology levels: $\{z_t\}$. We start from an arbitrary distribution μ_0 , implying a value \bar{k}_0 . We then recompute (6), using (12a)–(12d), at every point along the sequence $\{z_t\}$, and the implied sequence of aggregate capital levels $\{\bar{k}_t\}$, *without* imposing the equilibrium pricing rule (11a):

$$\tilde{V}^1(\epsilon, k, \xi; z_t, \bar{k}_t; p) = \max_n \left\{ \left[z_t \epsilon k^\theta n^\nu - i^M \right] p - An + \max \left\{ \beta V_i, \max_{k'} \left(-\xi A - ip + \beta E[V^0(\epsilon', k'; z', \bar{k}'(k_t))] \right) \right\} \right\},$$

¹¹The production unit can always choose $i = 0$ and thus $k^* = \psi(1 - \delta)k/\gamma$.

with V_i defined in (7b) and evaluated at $\bar{k}' = \bar{k}'(k_t)$. This yields new “policy functions”

$$\begin{aligned}\tilde{N} &= \tilde{N}(\epsilon, k; z_t, \bar{k}_t, p) \\ \tilde{K} &= \tilde{K}(\epsilon, k, \xi; z_t, \bar{k}_t, p).\end{aligned}$$

We then search for a p such that, given these new decision rules and after aggregation, the goods market clears (labor market clearing is trivially satisfied). We then use this p to find the new aggregate capital level.

This procedure generates a time series of $\{p_t\}$ and $\{\bar{k}_t\}$ endogenously, with which assumed rules (11a)–(11b) can be updated via a simple OLS regression. The procedure stops when the updated coefficients a_k, b_k, c_k and a_p, b_p, c_p are sufficiently close to the previous ones. We show in Appendix C that the implied R^2 of these regressions are high for all model specifications, generally well above 0.99, indicating that production units do not make large mistakes by using the rules (11a)–(11b). This is confirmed by the fact that adding higher moments of the capital distribution does not increase forecasting performance significantly.

3 Calibration

Our calibration strategy and parameters are standard with two additional features: We combine sectoral and aggregate data in order to infer the relative importance of PE- and GE-smoothing, and we calibrate the maintenance parameter by matching the conditional heteroscedasticity of investment in U.S. data.

3.1 Calibration Strategy

The model period is a quarter. The following parameters have standard values: $\beta = 0.9942$, $\gamma = 0.004$, $\nu = 0.64$, and $\rho_A = 0.95$. The depreciation rate δ matches the average quarterly investment rate in the data, 0.026, which leads to $\delta = 0.022$. The disutility of work parameter, A , is chosen to generate an employment rate of 0.6.

Next we explain our choices for θ and the parameters of the sectoral and idiosyncratic technology process (ρ_S, σ_S, ρ_I and σ_I). The output elasticity of capital, θ , is set to 0.18, in order to capture a revenue elasticity of capital, $\frac{\theta}{1-\nu}$, equal to 0.5, while keeping the labor share at its 0.64-value.¹² We determine σ_S and ρ_S by a standard Solow residual calculation on annual 3-

¹²In a world with constant returns to scale and imperfect competition this amounts to a markup of approximately 22%. The curvature of our production function lies between the values considered by KT and Gourio and Kashyap (2007). Cooper and Haltiwanger (2006), using LRD manufacturing data, estimate this parameter to be 0.592; Hennessy and Whited (2005), using Compustat data, find 0.551.

digit manufacturing data, taking into account sector-specific trends and time aggregation. This leads to values of 0.0273 for σ_s and 0.8612 for ρ_s .¹³ For computational convenience we set $\rho_I = \rho_s$, and σ_I to 0.0472, which leads to an annual standard deviation of the sum of sectoral and idiosyncratic shocks equal to 0.10.¹⁴

We turn now to the calibration of the two key parameters of the model, the adjustment cost parameter, $\bar{\xi}$, and the maintenance parameter, χ . We also describe how we calibrate the volatility of aggregate productivity shocks, σ_A .

With the availability of new and more detailed establishment level data, researchers have calibrated adjustment costs by matching establishment level moments (see, e.g., KT). This is a promising strategy in many instances, however, there are two sources of concern in the context of this paper’s objectives. First, one must take a stance regarding the number of productive units in the model that correspond to one productive unit in the available micro data. Some authors assume that this correspondence is one-to-one, while others match a large number of model-micro-units to one observed productive unit.¹⁵

Second, in state dependent models the frequency of microeconomic adjustment is not sufficient to pin down the object of primary concern, which is the aggregate impact of adjustment costs. Parameter changes in other parts of the model can have a substantial effect on this statistic, even in partial equilibrium. For example, anything that changes the drift of mandated investment (such as the maintenance investment parameter), changes the mapping from microeconomic adjustment costs to aggregate dynamics. Caplin and Spulber (1987) provide an extreme example of this phenomenon, where aggregate behavior is totally unrelated to microeconomic adjustment costs. In Appendix D we present a straightforward extension of this paper’s main model that provides a good fit of observed establishment level moments. This extension adds two micro parameters which, as in the Caplin and Spulber model, have no aggregate (or sectoral) consequences, yet can alter significantly establishment level moments.

Because of these concerns, we follow an alternative approach where we use 3-digit sectoral rather than plant level data to calibrate adjustment costs. More precisely, given a value of χ , we choose $\bar{\xi}$ to match the volatility of sectoral U.S. investment rates. Having done this, we choose σ_A to match the volatility of the aggregate U.S. investment rate. This leads to $\sigma_A = 0.0080$.¹⁶ The novelty in our calibration strategy is that it focuses on matching the relative importance of PE

¹³See Appendix A.3 for details and Appendix B.4 for robustness checks.

¹⁴In Table 16 in Appendix B.4 we consider values of 0.075 and 0.15 for the annual total standard deviation, with no significant changes to our baseline calibration.

¹⁵See Cooper and Haltiwanger (2006) and KT for an example of the former, and Abel and Eberly (2002) and Bloom (2009), who respectively assume that a continuum and 250 model micro units correspond to one observed plant or firm, for examples of the latter.

¹⁶For the frictionless model we also choose σ_A to match aggregate investment volatility, which leads to $\sigma_A = 0.0051$.

and GE smoothing directly. This approach assumes that the sectors we consider are sufficiently disaggregated so that general equilibrium effects can be ignored while, at the same time, there are enough micro units in them to justify the computational simplifications that can be made with a large number of units. Hence the choice of the 3-digit level.¹⁷

Given a set of parameters, the sequence of sectoral investment rates is generated as follows: First, the units' optimal policies are determined as described in Section 2.4, working in general equilibrium. Next, starting at the steady state, the economy is subjected to a sequence of sectoral shocks. Since sectoral shocks are assumed to have no aggregate effects and $\rho_I = \rho_S$, productive units perceive them as part of their idiosyncratic shock and use their optimal policies with a value of one for the aggregate shock and a value equal to the product of the sectoral and idiosyncratic shock—i.e. $\log(\epsilon) = \log(\epsilon_S) + \log(\epsilon_I)$ —for the idiosyncratic shock.¹⁸

The value of sectoral volatility of annual investment rates we match is 0.0163. To obtain this number we compute the volatilities of the linearly detrended 3-digit sectoral investment rates and take a weighted average. As noted in the introduction, this number is one order of magnitude smaller than the one predicted by the frictionless model. To match this annual sectoral volatility in the model simulations, we aggregate over time the quarterly investment rates generated by the model.

As shown in Figure 1, the residuals from estimating an autoregressive process for aggregate U.S. investment exhibit time-varying heteroscedasticity. Appendix B finds evidence of heteroscedasticity, and therefore of time-varying impulse responses, for overall U.S. investment and for equipment and structures separately, using two families of ARCH-type models.¹⁹ We calibrate the maintenance parameter χ by matching the logarithm of the ratio between the 95th and the 5th percentile of the estimated values for the conditional heteroscedasticity over the 1960-2005 period in the simple ARCH model described below; we refer to this statistic as the *heteroscedasticity range* in what follows.

Concretely, given a quarterly series of aggregate investment-to-capital ratios, x_t , the moment we match is obtained—both for actual and model-simulated data—by first regressing the series on its lagged value and then regressing the squared residual from this regression, $\hat{\epsilon}_t^2$, on x_{t-1} . Denoting by σ_{95} and σ_5 the 95th and 5th percentile of the fitted values from the latter

¹⁷Table 10 in Appendix A.3 provides information on the average number of establishments per 2-digit, 3-digit and 4-digit sector, both in absolute terms as well as in relation to the whole U.S. economy. Table 16 in Appendix B.4 shows that our calibration results do not change significantly if we work with 2 or 4-digit sectors.

¹⁸Appendix C.3 describes the details of the sectoral computation. There we also document a robustness exercise where we relax the assumption that sectoral shocks have no general equilibrium effects, and recompute the optimal policies when micro units consider the distribution of sectoral productivity shocks—summarized by its mean—as an additional state variable. Our main results are essentially unchanged.

¹⁹Applying these models to U.S. output shows no evidence of heteroscedasticity, suggesting that our heteroscedasticity finding is not determined by properties of the underlying shocks, which can be assumed to also drive output, but of the mechanism that leads from the shocks to investment.

regression, the heteroscedasticity range is equal to $\log(\sigma_{95}/\sigma_5)$. The target value for the heteroscedasticity range in the data is 0.3021.

3.2 Calibration Results

The upper bound of the adjustment cost distribution, $\bar{\xi}$, and the maintenance parameter, χ , that jointly match the sectoral investment volatility and the conditional heteroscedasticity statistic are $\bar{\xi} = 8.8$ and $\chi = 0.50$, respectively.²⁰ The average cost actually paid is much lower than the average adjustment cost, $\bar{\xi}/2$, as shown in Table 3, since productive units wait for good draws to adjust. The third row shows that, conditional on adjusting, in our calibrated model a production unit pays 3.6% of its annual output (column 1) or, equivalently, 5.6% of its regular wage bill (column 2).²¹ These costs are at the lower end of previous estimates, as shown by comparing them with rows 4 through 6.

Table 3: THE ECONOMIC MAGNITUDE OF ADJUSTMENT COSTS - ANNUAL

Model	Cond. Adj. Costs/ Unit's Output (1)	Cond. Adj. Costs/ Unit's Wage Bill (2)
This paper ($\chi = 0$):	38.9%	60.9%
This paper ($\chi = 25\%$):	12.7%	19.8%
This paper ($\chi = 50\%$):	3.6%	5.6%
Caballero-Engel (1999):	16.5%	—
Cooper-Haltiwanger (2006):	22.9%	—
Bloom (2009):	35.4%	—
Khan-Thomas (2008):	0.5%	0.8%

Notes: Based on Table IV in Bloom (2009). For Cooper-Haltiwanger (2006) and Bloom (2009) we report the sum of costs associated with two sources of lumpy adjustment: fixed adjustment costs and partial irreversibility. The remaining models only have fixed adjustment costs.

The first two rows of Table 2 in the introduction and Table 4 below show that our model fits both the sectoral and aggregate volatility of investment, as well as the range of conditional heteroscedasticity in aggregate data. This is not surprising, since our calibration strategy is designed to match these moments. In contrast, the bottom two rows in each of these tables show that neither the frictionless counterpart of our model nor the KT model match these features.²²

²⁰The values of χ we consider are multiples of 0.10, while the grid for $\bar{\xi}$ is finer.

²¹To compare our findings with the annual adjustment cost estimates in the literature, we report these numbers for an annual analogue of the quarterly model.

²²See footnote 5 for an explanation of why KT exhibits slightly lower nonlinearity than our calibration of a frictionless model.

Our calibration exercise yields a maintenance coefficient of 0.5 while the limited evidence available suggests values somewhere between 0.25 and 0.66.²³ This suggests our calibrated χ is at the upper end of the values in the literature and motivates considering variants of our lumpy adjustment model with a smaller role for replacement and maintenance.

Table 4: HETEROSCEDASTICITY RANGE

Model	$\log(\sigma_{95}/\sigma_5)$
<i>Data</i>	<i>0.3021</i>
This paper ($\chi = 0$):	0.1830
This paper ($\chi = 25\%$):	0.2178
This paper ($\chi = 50\%$):	0.2901
Frictionless:	0.0539
Khan-Thomas (2008):	0.0468

The first two rows in Table 3 report the magnitude of adjustment costs for $\chi = 0$ and $\chi = 0.25$. When calibrating these models, we no longer match the heteroscedasticity range in the data, but continue to match both sectoral and aggregate investment volatilities. For $\chi = 0.25$, the magnitude of adjustment costs lies slightly below the average of those estimated in the literature, for $\chi = 0$ they are slightly above the maximum value, but still within the ballpark.

The second and third rows in Table 4 show the range of heteroscedasticity values for versions of our model with values of χ smaller than in the benchmark case; the first row shows the values obtained directly from the data using our ARCH model. Even though these ranges now are smaller than those in the data, they continue being significantly larger than those implied by a frictionless model. The model with $\chi = 0$ has a heteroscedasticity range three times as large as in the frictionless model, for the model with $\chi = 25\%$ it is four times as large. The latter is much closer to the value in the data than in the frictionless model.

Ultimately, the main difference between our calibration and KT is the size of the adjustment cost. Table 5 makes this point, by reporting upper and lower bounds for the contribution of PE-smoothing to total smoothing, for several models, at different frequencies. The upper and

²³Cooper and Haltiwanger (2006) find the mode in the distribution of annual establishment level investment rates at 0.04. With an effective annual drift of 0.104, this suggests a maintenance parameter just below 40%. Alternatively, McGrattan and Schmitz (1999) show for Canadian data that maintenance and repair expenditures for equipment and structures amounts to roughly 30% of expenditures on new equipment and structures. This suggests just below 25% maintenance as a fraction of overall investment. And Letterie et al. (2004) report that replacement investment in Germany accounts for 66% of all investment, which suggests a value for χ of 0.66.

lower bounds for the contribution of PE-smoothing are calculated as follows:

$$UB = \log[\sigma(\text{NONE})/\sigma(\text{PE})]/\log[\sigma(\text{NONE})/\sigma(\text{BOTH})],$$

$$LB = 1 - \log[\sigma(\text{NONE})/\sigma(\text{GE})]/\log[\sigma(\text{NONE})/\sigma(\text{BOTH})]$$

where σ denotes the standard deviation of aggregate investment rates, NONE refers to the partial equilibrium model with no microeconomic frictions, PE to the model that only has microeconomic frictions so that prices are fixed at their average levels of the GE specification, GE to the model with endogenous price movements governed by the first-order conditions of the representative household (10), and BOTH to the model with both micro frictions and GE constraints.

Table 5: SMOOTHING DECOMPOSITION: KT

<u>Model</u>	<u>PE/total smoothing</u>		
	LB	UB	Avg.
KT-Lumpy annual:	0.0%	16.1%	8.0%
KT-Lumpy annual, our $\bar{\xi}$:	8.1%	59.2%	33.7%
Our model annual (0% maint.), KT's $\bar{\xi}$:	0.8%	16.0%	8.4%
Our model annual (0% maint.):	18.9%	75.3%	47.0%
Our model annual (25% maint.):	19.1%	75.7%	47.4%
Our model annual (50% maint.):	19.9%	76.6%	48.3%
Our model quarterly (0% maint.):	14.5%	80.9%	47.7%
Our model quarterly (25% maint.):	15.4%	80.9%	48.2%
Our model quarterly (50% maint.):	15.4%	81.0%	48.2%

The main message can be gathered from the first two rows of these tables: By changing the adjustment cost distribution in KT's model for ours,²⁴ its ability to generate substantial PE-smoothing rises significantly. Conversely, introducing KT adjustment costs into an annual version of our lumpy model with zero maintenance (third row) leads to a similarly small role of PE-smoothing as in their model. Rows four to nine show the much larger role for PE-smoothing under our calibration strategy, robustly for annual and quarterly calibrations and low and high values of the maintenance parameter.

²⁴Since KT measure labor in time units (and therefore calibrate to a steady state value of 0.3), and we measure labor in employment units, the steady state value of which is 0.6, and adjustment costs in both cases are measured in labor units, we actually use half of our calibrated adjustment cost parameter. Conversely, when we insert KT adjustment costs into our model, we double it.

3.3 Conventional RBC Moments

Before turning to the specific aggregate implications and mechanisms of microeconomic lumpiness that are behind the empirical success of our model, we show that these gains do not come at the cost of sacrificing conventional RBC-moment-matching. Tables 6 and 7 report standard longitudinal second moments for both the lumpy model and its frictionless counterpart. We also include a model with no idiosyncratic shocks and the higher revenue elasticity of KT (we label it RBC). As with all models, the volatility of aggregate productivity shocks is chosen to match the volatility of the aggregate investment rate.²⁵

Table 6: VOLATILITY OF AGGREGATES IN PER CENT

Model	Y	C	I	N
Lumpy:	1.34	0.83	4.34	0.56
Frictionless:	1.11	0.44	5.39	0.73
RBC:	1.35	0.45	5.03	0.97
Data:	1.36	0.94	4.87	1.27

Table 7: PERSISTENCE OF AGGREGATES

Model	Y	C	I	N	I/K
Lumpy:	0.70	0.71	0.70	0.70	0.92
Frictionless:	0.69	0.79	0.67	0.67	0.86
RBC:	0.70	0.80	0.68	0.68	0.92
Data:	0.91	0.87	0.91	0.90	0.96

Overall, the second moments of the lumpy model are reasonable and comparable to those of the frictionless models. While the former exacerbates the inability of RBC models to match the volatility of employment (we use data from the establishment survey on total employment from the BLS), the lumpy model improves significantly when matching the volatility of consumption.²⁶ The lumpy model also increases slightly the persistence of most aggregate variables, bringing these statistics closer to their values in the data.

²⁵The value of σ_A required for the RBC model is 0.0058. For the lumpy model, the employment statistics are computed from total employment, that is including those workers who work on adjusting the capital stock. We work with all variables in logs and detrend then with an HP-filter using a bandwidth of 1600.

²⁶Consistent with our model, we define aggregate consumption as consumption of nondurables and service minus housing services. Also, we define output as the sum of this consumption aggregate and aggregate investment.

4 Aggregate Investment Dynamics

In this section we describe the mechanism behind our model's ability to match the conditional heteroscedasticity of aggregate investment rates. In particular, we show that lumpy adjustment models generate history dependent aggregate impulse responses.

Figure 3: Time Paths of the Responsiveness Index




Figure 3 plots the evolution of the quarterly responsiveness index for the 1960-2005 period (in log deviation from its average value). The solid and dashed lines represent the index for the lumpy and frictionless models, while the dotted line represents the index for the ARCH-type time series model discussed in Section 3.1.

Following Caballero and Engel (1993b), the responsiveness index at time t is defined as follows: Given an economy characterized by a joint distribution of capital and productivity μ_t , and an aggregate productivity level z_t , we denote the resulting aggregate investment rate by $\frac{I}{K}(\mu_t, \log z_t)$ and define the normalized response of this economy to a positive and negative one standard deviation aggregate productivity shock, respectively, as

$$\begin{aligned}\mathcal{I}^+(\mu_t, \log z_t) &\equiv \left(\frac{I}{K}(\mu_t, \log z_t + \sigma_A) - \frac{I}{K}(\mu_t, \log z_t) \right), \\ \mathcal{I}^-(\mu_t, \log z_t) &\equiv \left(\frac{I}{K}(\mu_t, \log z_t - \sigma_A) - \frac{I}{K}(\mu_t, \log z_t) \right),\end{aligned}$$

where σ_A is the standard deviation of the aggregate innovation. The Responsiveness Index at

time t then is defined as:

$$F_t \equiv 0.5(\mathcal{I}^+(\mu_t, \log z_t) - \mathcal{I}^-(\mu_t, \log z_t)). \quad (16)$$

That is, this index captures the response upon impact of the aggregate investment rate to an aggregate productivity innovation, conditional on the current state of the economy.

The shocks fed into the model are the ones backed out to match actual aggregate quarterly investment rates over the sample period. We initialize the process with the economy at its steady state in the fourth quarter of 1959.²⁷

The figure confirms the statement in the introduction, according to which in the lumpy capital adjustment model the initial response to an aggregate shock varies significantly more over time than in a frictionless model: The responsiveness index grows by 50.9% between trough and peak, which is similar to the 46.8% variation obtained from the simple ARCH-type model discussed in Section 3.1 and considerably larger than the 11.6% variation implied by the frictionless model.

To understand how lumpy adjustment models generate time varying impulse responses, two features of the time paths of the responsiveness index are important. Note first that the index fluctuates much less in the frictionless economy than in the lumpy economy. Recall also that the frictionless economy only has general equilibrium forces to move this index around. From these two observations we can conjecture that the contribution of the general equilibrium forces to the volatility of the index in the lumpy economy is minor.


It follows from this figure that it is the decline in the strength of the PE-smoothing mechanism that is responsible for the rise in the index during the boom phase. When this mechanism is weakened, the responsiveness index in the lumpy economy grows by more than that of the frictionless economy in a boom.

Figure 4 illustrates why the PE-smoothing mechanism weakens as the boom progresses. The figure shows the cross-section of mandated investment (and the probability of adjusting, conditional on mandated investment) at three points in time: a period of average investment in the first quarter of 1989 (solid line), a period of booming aggregate investment, the second quarter of 2000 (dashed line), and a period of depressed aggregate investment in the first quarter of 1961 (dotted line).²⁸ It is apparent from this figure that during the boom the cross-section of

²⁷By “steady state” we mean the ergodic (time-average) distribution, which we calculate as follows: starting from an arbitrary capital distribution and the ergodic distribution of the idiosyncratic shocks, we simulate the development of an economy with no aggregate innovations for 300 periods, but using the policy functions under the assumption of an economy subject to aggregate shocks.

²⁸See Section 2.4 for the formal definition of mandated investment. Also note that the scale on the left of the figure is for the mandated investment densities, while the scale on the right is for the adjustment hazards. See Appendix A.2, Figure 10 for a time path of the quarterly aggregate investment rate in the U.S.

Figure 4: Investment Boom-Bust Episode: Cross-section and Hazard




mandated investment moves toward regions where the probability of adjustment is higher and steeper. The fraction of micro units with mandated investment close to zero decreases considerably during the boom, while the fraction of units with mandated investment rates above 40% increases significantly. Also note that the fraction of units in the region where mandated investment is negative decreases during the boom, since the sequence of positive shocks moves units away from this region.

The convex curves in Figure 4 depict the state-dependent adjustment hazard; that is, the probability of adjusting conditional on mandated investment. It is clear that the probability of adjusting increases with the (absolute) value of mandated investment. This is the ‘increasing hazard property’ described in Caballero and Engel (1993a). The convexity of the estimated state-dependent adjustment hazards implies that the probability that a shock induces a micro unit to adjust is larger for units with larger values of mandated investment. Since units move into the region with a higher slope of the adjustment hazard during the boom, aggregate investment becomes more responsive. This effect is further compounded by the fact that the adjustment hazard shifts upward as the boom proceeds, although this mechanism is small.


In summary, the decline in the strength of PE-smoothing during the boom (and hence the larger response to shocks) results mainly from the rise in the share of agents that adjust to further shocks. This is in contrast with the frictionless (and Calvo style) models where the only margin of adjustment is the average size of these adjustments. This is shown in Figure 5, which decomposes the time path of the responsiveness index of the lumpy model into two components: one that reflects the response of the fraction of adjusters (the extensive margin) and another that captures the response of average adjustments of those who adjust (the intensive margin). It is apparent that most of the change in the responsiveness index is accounted for by

Figure 5: Decomposition of Responsiveness Index: Intensive and Extensive Margins



variations in the fraction of adjusters, that is, by the extensive margin.

Figure 6: Decomposition of I/K into Intensive and Extensive Margins




The importance of fluctuations in the fraction of adjusters is also apparent in the decomposition of the path of the aggregate investment rate into the contributions from the fluctuation of the fraction of adjusters and the fluctuation of the average size of adjustments, as shown in Figure 6. Both series are in log-deviations from their average values. This is consistent with what Doms and Dunne (1998) documented for establishment level investment in the U.S. and Gourio and Kashyap (2007) for the U.S. and Chile, where the fraction of units undergoing major investment episodes accounts for a much higher share of aggregate (manufacturing in their case) investment than the average size of their investment.²⁹

²⁹Doms and Dunne (1998) show that the number of plants that have their highest investment in a given year has a correlation with aggregate investment of roughly 60%. Gourio and Kashyap (2007) show in their Figure 2 that

Next we illustrate the time variation of the investment response during the turn-of-the-millennium boom-bust cycle. Figure 7 depicts the responses over five quarters of the baseline lumpy model to a one standard deviation shock taking place during the peak of this cycle in the second quarter of 2000 and the trough in the first quarter of 2003, normalized by the average impulse response upon impact over the entire sample. The response of investment to a stimulus (e.g., an investment credit) varies systematically over the cycle, being least responsive during a slowdown. Using a linear model to gauge the effect of a stimulus is likely to overestimate the investment response during a downturn, by approximately 20%. This is because the response to a sequence of average shocks, which corresponds to the standard impulse response function calculated for a linear model, is in between both cases and fails to capture the significant time variation of the impulse responses in a world with lumpy investment.

Figure 7: Impulse Responses of the Aggregate Investment Rate in the 2000 Boom-Bust Cycle




We end this section by analyzing the role of the maintenance parameter in determining aggregate investment dynamics. As discussed in Section 3.2, the magnitude of adjustment costs decreases with χ , while the extent to which the investment response varies over the cycle increases with χ . The insights we have gained earlier in this section provide an explanation for these findings.

The negative correlation between adjustment costs and maintenance follows from the fact that a higher maintenance parameter lowers the effective drift of mandated investment, defined as depreciation that is not necessarily undone in a given period. Without maintenance, the drift dominates over microeconomic uncertainty shocks and the cross section of mandated investment is closer to the Caplin and Spulber extreme where there is no PE-smoothing (see Caballero and Engel, 2007). Figure 8 shows the average cross-section distribution of mandated investment for our baseline model and for the model with $\chi = 0$. The former is clearly farther

aggregate investment is mainly driven by investment spikes and those to a large degree are accounted for by the fraction of units undergoing major investment episodes.


away from the Caplin-Spulber uniform limit, leaving more space for PE-smoothing. Thus, in order to keep PE-smoothing constant as we reduce the value of the maintenance parameters, we need to compensate it by increasing the adjustment cost.

Figure 8: Ergodic cross-section: Zero and Baseline Maintenance



Note, however, that the compensation via an increase in adjustment costs is not enough to preserve the volatility of the impulse response as we drop maintenance (see Table 4), since this feature is more sensitive to the shape of the cross section distribution. However, Figure 9 shows that even with zero maintenance, the responsiveness index of the lumpy economy varies considerably more than in the frictionless economy.

Figure 9: Time Paths of the Responsiveness Index - Lower Maintenance



5 Final Remarks

This paper begins by presenting time series evidence showing that the impulse response function for U.S. investment is history dependent: investment responds more to a given shock during persistent booms than during slumps.

Next we argue that it is important to identify the relative contribution of partial and general equilibrium forces in smoothing the impact of shocks on aggregate variables. In particular, in the case of investment models with lumpy capital adjustment we find that only models that allow a non-trivial role for partial equilibrium smoothing can match the time series evidence on history dependent impulse responses.

Finally, we show that the reason why models that add realistic lumpy capital adjustment to an otherwise standard RBC model generate procyclical impulse responses is that, relative to the standard RBC model, in the lumpy models investment booms feed into themselves and lead to significantly larger capital accumulation following a string of positive shocks. During busts, on the other hand, the economy is largely unresponsive to positive shocks. These are exactly the patterns we observe in U.S. aggregate data.

References

- [1] Abel, A. and J. Eberly (2002). "Investment and q with Fixed Costs: An Empirical Analysis," mimeo.
- [2] Bloom, N. (2009). "The Impact of Uncertainty Shocks," *Econometrica*, **77** (3), 623–685.
- [3] Caballero, R. and E. Engel (1993a). "Microeconomic Adjustment Hazards and Aggregate Dynamics", *Quarterly Journal of Economics*, **108** (2), 359–383.
- [4] Caballero, R. and E. Engel (1993b). "Microeconomic Rigidities and Aggregate Price Dynamics", *European Economic Review*, **37** (4), 697–711.
- [5] Caballero, R. and E. Engel (1999). "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach," *Econometrica*, July 1999, **67**(4), 741–782.
- [6] Caballero, R. and E. Engel (2007). "Price Stickiness in Ss Models: New Interpretations of Old Results," *J. of Monetary Economics*, **54**, 100–121.
- [7] Caplin, A. and D. Spulber (1987). "Menu Costs and the Neutrality of Money," *Quarterly Journal of Economics*, **102**, 703–26.
- [8] Cooper, R. and J. Haltiwanger (2006). "On The Nature of Capital Adjustment Costs," *Review of Economic Studies*, 2006, **73**, 611–633.
- [9] Doms, M. and T. Dunne (1998). "Capital Adjustment Patterns in Manufacturing Plants", *Review of Economic Dynamics*, **1**, 1998, 409–429.
- [10] Gourio, F. and A.K. Kashyap, (2007). "Investment spikes: New facts and a general equilibrium exploration", *Journal of Monetary Economics*, **54**, 2007, 1–22.
- [11] Hennessy, C. and T. Whited, (2005). "Debt Dynamics", *The Journal of Finance*, **LX/3**, 2005, 1129–1165.
- [12] Judd, K. (1998). *Numerical Methods in Economics*. MIT Press: Cambridge, Mass.
- [13] Khan, A. and J. Thomas, (2003). "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?", *Journal of Monetary Economics*, March 2003, **50**(2), 331–360.
- [14] Khan, A. and J. Thomas, (2008). "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics", *Econometrica*, **76**(2), March 2008, 395–436.

- [15] Krusell, P. and A. Smith (1997). "Income and Wealth Heterogeneity, Portfolio Choice and Equilibrium Asset Returns". *Macroeconomic Dynamics* **1**, 387–422.
- [16] Krusell, P. and A. Smith (1998). "Income and Wealth Heterogeneity in the Macroeconomy", *J. of Political Economy*, **106** (5), 867-896.
- [17] Letterie, W., G. Pfann and S. Verick (2004). "Non-Linearities in the Expansion of Capital Stock," IZA Discussion Paper No. 1132, May 2004.
- [18] McGrattan, E. R. and J. A. Schmitz (1999). "Maintenance and Repair: Too Big to Ignore," *Federal Reserve Bank of Minneapolis Quarterly Review*, **23**(4), 2–13.
- [19] Smith, A. (1993). "Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions," *J. of Applied Econometrics*, **8**, S63–S84.
- [20] Tauchen, G. (1986). "Finite State Markov-Chain Approximations To Univariate and Vector Autoregressions". *Economics Letters* **20**, 177–181.
- [21] Thomas, J.K. (2002). "Is Lumpy Investment Relevant for the Business Cycle?", *Journal of Political Economy*, **110**(3), 2002, 508–534.
- [22] Veracierto, M., (2002). "Plant Level Irreversible Investment and Equilibrium Business Cycles," *American Economic Review*, **92** (1), March 2002, 181–197.

A Parameter and Data Appendix

A.1 Parameters

Table 8 summarizes the *common* parameters of the models explored in the paper:

Table 8: COMMON PARAMETERS

Calibration	ρ_A	$\rho_S = \rho_I$	σ_S	σ_I	δ	γ	β	θ	ν
Quarterly	0.9500	0.8612	0.0273	0.0472	0.0220	1.0040	0.9942	0.1800	0.6400
Yearly	0.8145	0.5500	0.0501	0.0865	0.0880	1.0160	0.9770	0.1800	0.6400

Persistence parameters have the following relation between quarterly and annually: $\rho_q = \rho_y^{0.25}$ (the same holds true for β). For standard deviations the following relationship holds: $\sigma_q = \frac{\sigma_y}{\sqrt{1 + \rho_q + \rho_q^2 + \rho_q^3}}$. For ρ_S and σ_S the yearly parameters are primitive because of the merely annual availability of sectoral data. Notice that for the yearly specification $\sqrt{\sigma_S^2 + \sigma_I^2} = 0.1$. Finally, the production function for quarterly output is one fourth of the one for yearly output.

The calibration of the other parameters, $\sigma_A, \chi, \bar{\xi}$ and A is explained in Section 3. When we refer in the main text to a *quarterly calibration* (our benchmark models), then we use – given the quarterly parameters in the table above – σ_A and $\bar{\xi}$ to match jointly the standard deviation of the quarterly aggregate investment rate and the standard deviation of the yearly sectoral investment rate, which is aggregated up over four quarters in the sectoral simulations (we do not have quarterly sectoral data). This amounts to $\sigma_A = 0.0080$ for the baseline lumpy model and $\sigma_A = 0.0051$ for its frictionless counterpart. When we refer to a *yearly calibration*, then we use – given the yearly parameters in the table above – σ_A and $\bar{\xi}$ to match jointly the standard deviation of the yearly aggregate investment rate and the standard deviation of the yearly sectoral investment rate. This amounts to $\sigma_A = 0.0186$ for the baseline lumpy model and $\sigma_A = 0.0120$ for its frictionless counterpart. The parameter that governs conditional heteroscedasticity, χ , is calibrated only for the quarterly specifications, because we estimate conditional heteroscedasticity on quarterly aggregate data to have enough data points to detect possible nonlinearities.

A.2 Aggregate Data

Since they are not readily available from standard sources, we construct quarterly series of the aggregate investment rate using investment and capital data from the national account and fixed asset tables, available from the Bureau of Economic Analysis (BEA). The time horizon is

1960:I–2005:IV. The quarterly aggregate investment rate in period t is defined as $I_t^{Q,\text{real}}/K_{t-1}^{Q,\text{real}}$, where the denominator is the real capital stock at the end of period $t-1$ and the numerator is real investment in period t .

The information we used is (a) nominal annual private fixed nonresidential investment, I^Y , from table 1.1.5 Gross Domestic Product line 9; (b) the annual private nonresidential capital stock at year-end prices, \tilde{K}^Y , from table 1.1 Fixed Assets and Consumer Durable Goods line 4; (c) nominal annual private nonresidential depreciation, D^Y , from table 1.3 Fixed Assets and Consumer Durable Goods line 4; (d) quarterly nominal fixed nonresidential investment seasonally adjusted at annual rates, \tilde{I}^Q , from table 1.1.5 Gross Domestic Product line 9; and (e) the quarterly implicit price deflator of nonresidential investment, P^Q , from table 1.1.9 Gross Domestic Product line 9.

Quarterly figures for investment are obtained as follows. Since seasonally adjusted quarterly nominal investment does not add up to annual nominal investment, we impose this adding up constraint by calculating nominal investment in quarter t of year y as $I_t^Q = (I_y^Y / \sum_{t \in y} \tilde{I}_t^Q) \tilde{I}_t^Q$, where y denotes both the year and all quarters in that year. Real investment is then calculated as, $I_t^{Q,\text{real}} = I_t^Q / P_t^Q$.

To calculate the quarterly real capital stock we proceed as follows. Let π_t denote quarterly investment price inflation between period $t-1$ and t , which is obtained from the implicit price deflator data by $1 + \pi_t = P_t^Q / P_{t-1}^Q$. We assume that annual depreciation figures reported by the BEA are at average prices of the year. Quarterly depreciation series are constructed using nominal annual depreciation and quarterly investment inflation, under the assumptions that quarterly nominal depreciation numbers add up to annual figures and that real depreciation is the same for every quarter of a given year. That is, nominal depreciation in the four quarters of a year, denoted D_1, D_2, D_3, D_4 , are given by,

$$\begin{aligned} D_4 &= D_3(1 + \pi_4) = D_2(1 + \pi_3)(1 + \pi_4) = D_1(1 + \pi_2)(1 + \pi_3)(1 + \pi_4), \\ D^Y &= D_1 + D_2 + D_3 + D_4, \end{aligned}$$

where D^Y denotes total depreciation during that year. To compute quarterly nominal capital stocks, K_t^Q , during the first three quarters we use the following identity:

$$K_t^Q = K_{t-1}^Q(1 + \pi_t) + I_t^Q - D_t^Q,$$

where all variables are nominal. For fourth quarter capital stocks we use the annual end-of-year data. Year-end prices reported by the BEA are the average of fourth-quarter prices in the current year and first-quarter prices in the following year, thus nominal end-of-year capital,

K_4^Q , for any given year is obtained from $K_4^Q = 2P_4^Q \tilde{K}^Y / (P_4^Q + P_1^Q)$, where P_1^Q corresponds to the nominal price of investment in the first-quarter of next year. Real capital is then calculated as, $K_t^{Q,real} = K_t^Q / P_t^Q$.

As Figure 10 shows (the vertical lines denote NBER business cycle dates), the aggregate investment rate does not appear to exhibit any trend, which is why we do not filter any statistics related to it (both for real and simulated data).

Figure 10: The Quarterly Aggregate Investment Rate




Table 9 summarizes statistics of the aggregate investment rate.³⁰

Table 9: AGGREGATE INVESTMENT RATE

	Mean	STD	Persistence	Max	Min
Quarterly	0.026	0.0023	0.96	0.031	0.022
Yearly	0.104	0.0098	0.73	0.125	0.086

A.3 Sectoral Data

For sectoral data the best available source is the NBER manufacturing data set, publicly available on the NBER website. It contains yearly 4-digit industry data for the manufacturing sector,

³⁰The maximum is achieved in 2000:II, the minimum in 1961:I.

according to the SIC-87 classification. We look at the years 1960-1996, later years are not available. We take out industry 3292, the asbestos products, because this sector essentially dies out in the nineties. This leaves us with 458 4-digit industries altogether.

Since the sectoral model analysis has to (a) be isolated from general equilibrium effects, and (b) contain a large number of production units, we take the 3-digit level as the best compromise aggregation level. This leaves us with 140 industries.³¹ Hence, we sum employment levels, real capital, nominal investment and nominal value added onto the 3-digit level. The deflator for investment is aggregated by a weighted sum (weighted by investment). Value added is deflated by the GDP deflator instead of the sectoral deflators for shipments (the data do not contain separate deflators for value added). We do this, because our model does not allow for relative price movements between sectors, so by deflating sectoral value added with the GDP deflator the resulting Solow residual is essentially a composite of true changes in sectoral technology and relative price movements. Since value added and deflators are negatively correlated, we would otherwise overestimate the volatility of sectoral innovations and thus overcalibrate adjustment costs.³²

TFP-Calculation: Since our model is about value added production as opposed to output production—we do not model utilization of materials and energy—we do not use the TFP-series in the data set, which are based on a production function for output. Rather, we use a production function for real value added in employment and real capital with payroll as a fraction of value added as the employment share, and the residual as capital share, and perform a standard Solow residual calculation for each industry separately.

Next, in order to extract the residual industry-specific and uncorrelated-with-the-aggregate component for each industry, we regress each industry time series of logged Solow residuals on the time series of the value added-weighted cross-sectional average of logged Solow residuals and a constant. Since the residuals of this regression still contain sector-specific effects, but our model features ex-ante homogenous sectors, we take out a deterministic quadratic trend on these residuals for each sector. We use a deterministic quadratic trend because it makes persistence and volatility of the estimated residuals smaller than with a linear trend or no detrending. This is a conservative approach for our purposes, as this will make, ceteris paribus, the calibrated adjustment costs and therefore aggregate nonlinearities smaller. Not detrending the sectoral Solow residuals would increase both annual persistence and the annual standard deviation of the sectoral shock innovation from 0.55 to 0.65, and from 0.0501 to 0.0518, respec-

³¹Aggregating to the 2-digit levels leaves us with 20 industries.

³²Indeed, using a weighted sum of 3-digit level value added deflators instead of the GDP deflator would increase the standard deviation of the sectoral shock innovation from the 0.0501 we are using to 0.0564 and the persistence of sectoral technology from 0.55 to 0.61, other things being equal. We thank Julia Thomas for this suggestion.

tively. The residuals of this trend regression are then taken as the pure sectoral Solow residual series. By construction, they are uncorrelated with the cross-sectional average series. We then estimate an AR(1)-specification for each of these series, and, to come up with a single value for σ_S and ρ_S , set σ_S equal to the value-added-weighted average of the estimated standard deviations of the corresponding innovations, which results in $\sigma_S = 0.0501$ (annual), and ρ_S equal to the value-added-weighted average of the estimated first-order autocorrelation, which leads to $\rho_S = 0.55$ (annual).

Since this computation is subject to substantial measurement error and somewhat arbitrary choices, we perform a number of robustness checks: 1) We fix the employment share and capital share to $\nu = 0.64$ and $\theta = 0.18$, as in our model parametrization for all industries. 2) Instead of using an OLS projection onto the cross-sectional mean, we simply subtract the latter. 3) We look at unweighted means. 4) We look at medians instead of means, again weighted and unweighted. The resulting numbers remain in the ballpark of the parameters we use (see Table 16 in Appendix B.4 for a robustness analysis with some of these alternative choices).

Calculation of I/K-Moments: To extract a pure sectoral component of the time series of the industry investment rate, which like the aggregate data includes equipment and structures, we perform the same regressions that were used for TFP-calculation, except that we use a deterministic linear trend to extract sector specific effects. A quadratic detrending of the driving force and a linear detrending for the outcome variable is a conservative approach, as it will make calibrated adjustment costs and aggregate nonlinearities smaller. We do not log or filter the investment rate series. The common component we regress the sectoral investment rate series on is now a capital-weighted average of the industry investment rates. Again, we perform robustness checks with fairly stable results. The resulting standard deviation of sectoral investment rates – our target of calibration – is 0.0163.³³

Data for Different Digit Levels: Finally, Table 10 provides information on the number of establishments per sector and the size of each sector within the U.S. economy for the 2-digit, 3-digit and 4-digit levels.³⁴ It justifies our choice to use 3-digit data in the baseline calibration.

The mean and median (across industries) number of establishments in the 3-digit industries are 2,671 and 1,147, respectively. While at the 4-digit level industries still contain a fairly large number of establishments on average, the continuum assumption is certainly more justifiable for the 3-digit level. Conversely, the across industries average fraction of industry establishments over the number of total establishments in the U.S. is 0.04% at the 3-digit level, the

³³Their persistence is 0.55.

³⁴We use the County Business Pattern data from 1996 to generate the numbers in this table.

Table 10: SUMMARY STATISTICS FOR MANUFACTURING ESTABLISHMENTS

	Mean# Est.	Median# Est.	Mean Frac. Est.	Median Frac. Est.	Max. Frac. Est.
2-digit	19041	14455	0.28%	0.21%	0.94%
3-digit	2671	1147	0.04%	0.02%	0.51%
4-digit	780	333	0.01%	0.00%	0.38%

median 0.02% and the maximum 0.51%. In other words, the manufacturing industry with the largest number of establishments has a share of half a per cent in the total number of U.S. establishments.³⁵ The table thus shows that the choice of the 3-digit level is a good compromise between our two assumptions: small enough to not have general equilibrium impacts and large enough to justify the assumption of a large number of units. Nevertheless, in Appendix B.4 we report calibration results also for the 2-digit and 4-digit levels and show that our results do not hinge on this choice.

B Conditional Heteroscedasticity and Aggregate Investment

In this appendix we first present time series evidence for conditional heteroscedasticity in aggregate U.S. investment to capital ratios. Then we explain how we calibrated the maintenance parameter using this feature of the data.

B.1 Time series models

We consider two stationary time series models within the ARCH family to explore whether aggregate investment exhibits the kind of heteroscedasticity predicted by S_s -type models, namely that investment responds more to a shock during a boom than during a slump. Both models assume that

$$x_t = \sum_{j=1}^p \phi_j x_{t-j} + \sigma_t e_t, \quad (17)$$

where $x_t \equiv I_t/K_t$ denotes the investment to capital ratio, the e_t are i.i.d. with zero mean and unit variance, and σ_t is a simple function of recent values of x_t as summarized by the index

$$\bar{x}_{t-1}^k \equiv \frac{1}{k} \sum_{j=1}^k x_{t-j}. \quad (18)$$

³⁵These numbers would be, respectively, 0.07%, 0.03% and 9%, had we used the average fraction of industry establishments over the number of total establishments in the manufacturing sector. Had we used share in total U.S. employment as our metric, the numbers for the 3-digit sector would have been: 0.12%, 0.07% and 0.76%.

For model 1 we stipulate

$$\sigma_t = \alpha_1 + \eta_1 \bar{x}_{t-1}^k, \quad (19)$$

while for model 2 we posit

$$\sigma_t^2 = \alpha_2 + \eta_2 \bar{x}_{t-1}^k. \quad (20)$$

When $\eta_1 = \eta_2 = 0$, the above models simplify to a standard autoregressive time series, with an impulse response that does not vary over time.

It follows from (17) that the impulse response of x to e upon impact at time t , denoted by $\text{IRF}_{0,t}$, is equal to σ_t . Hence:

$$\text{IRF}_{0,t} = \begin{cases} \alpha_1 + \eta_1 \bar{x}_{t-1}^k, & \text{for model 1;} \\ \sqrt{\alpha_2 + \eta_2 \bar{x}_{t-1}^k}, & \text{for model 2.} \end{cases}$$

The models with lumpy adjustment developed in this paper (and earlier models such as Caballero and Engel, 1999) predict positive values for η_1 and η_2 . The reason is that, as shown in Figure 4 in the main text, in these models the cross-section of mandated investment concentrates in a region with a steeper likelihood of adjusting when recent investment has been high, which implies that investment becomes more responsive to shocks during these times.

B.2 Estimation and Results

Assume observations for x_t are available for $t = 1, \dots, T$, and denote by p_{\max} and k_{\max} the largest values considered for p and k in (17) and (18), respectively. For all pairs (p, k) with $p \leq p_{\max}$ and $k \leq k_{\max}$ we estimate an $\text{AR}(p)$ using OLS, and then use the residuals from this regression, denoted ε_t , to estimate α and η via OLS from:³⁶

$$\begin{aligned} \text{Model 1:} \quad & |\varepsilon_t| = \sqrt{\frac{2}{\pi}} (\alpha_1 + \eta_1 \bar{x}_{t-1}^k) + \text{error}, \\ \text{Model 2:} \quad & \varepsilon_t^2 = \alpha_2 + \eta_2 \bar{x}_{t-1}^k + \text{error}. \end{aligned} \quad (21)$$

We choose the optimal values for p and k , denoted by p^* and k^* , using the Akaike Information Criterion (AIC).

³⁶The first equation is based on

$$\text{E}[|\varepsilon_t| | \bar{x}_{t-1}^k] = \sqrt{\frac{2}{\pi}} (\alpha_1 + \eta_1 \bar{x}_{t-1}^k),$$

while the second equation comes from

$$\text{E}[\varepsilon_t^2 | \bar{x}_{t-1}^k] = \alpha_2 + \eta_2 \bar{x}_{t-1}^k.$$

Also note that we use the same number of observations when estimating all regressions: $T - \max(p_{\max}, k_{\max})$.

Table 11 presents the estimates obtained for both models, for U.S. private, fixed, nonresidential investment, and for equipment and structures separately. The frequency is quarterly, from 1960:I to 2005:IV. We use $p_{\max} = k_{\max} = 12$.

The first and second rows report the optimal values for p and k . The following seven rows report statistics related to the magnitude and significance of the parameter that captures heteroscedasticity and time-variation in impulse responses, η . The third row has the point estimate for η and the fourth row the corresponding t -statistic, obtained from OLS estimates for (21). The latter may overstate the significance of η , since it ignores variations in the first stage regressions that determine the autoregressive order, p^* . For this reason we use 10,000 bootstrap simulations for the investment rate series, starting from our estimates for the e_t s in (17) and (18), to provide an alternative measure of the precision of our estimates for η .³⁷ The fifth row presents the p -values we obtain for $\eta > 0$ via bootstrap simulations, we report one-sided p -values since Ss -type models predict $\eta > 0$. The last 4 rows present measures for the range of values taken by the estimated impulse response upon impact: σ_{\max} , and σ_{\min} denote the largest and smallest heteroscedasticity estimates over the sample considered (172 observations), σ_p the p -th percentile. We sign the range estimates by the estimated sign of η .

Table 11: Evidence of heteroscedasticity - U.S. Investment

Series:	All	All	Equip	Equip	Str	Str
Model:	1	2	1	2	1	2
p^* :	6	6	7	7	6	6
k^* :	6	6	8	8	2	2
$\eta \times 10^3$:	45.93	0.03731	30.62	0.05380	39.95	0.02581
$t\text{-}\eta$:	3.121	2.496	2.089	1.724	4.097	3.245
$p\text{-value}(\eta > 0)\text{-bootstrap}$:	0.0088	0.0236	0.0375	0.0742	0.0043	0.0094
$\pm \log(\sigma_{\max}/\sigma_{\min})$:	0.7367	0.5933	0.5521	0.4395	1.1167	1.1169
$\pm \log(\sigma_{95}/\sigma_5)$:	0.6118	0.4816	0.4520	0.3547	0.9194	0.8894
$\pm \log(\sigma_{90}/\sigma_{10})$:	0.5203	0.4082	0.3355	0.2254	0.8003	0.7403
no. obs. lhs. 1st and 2nd reg.:	172	172	172	172	172	172

Table 11 shows that nonresidential investment exhibits significant (both statistically and economically) heteroscedasticity for both models. This is also the case for structures, and for equipment under Model 1. The range of heteroscedasticity values implied by the estimated models is large. For example, the estimates for model 2 imply that the 95th percentile is 61.9% larger ($e^{0.4816} \simeq 1.619$) than the 5th percentile.

³⁷For each series generated via bootstrap we estimate the $p_{\max} \times k_{\max}$ models and determine the optimal values for p , k and, most important, η .

Table 12 shows the estimates we obtain when applying the methodology described above to the cyclical component of log-GDP. We consider the three most commonly used filters to detrend GDP: the HP-1600 filter, the Baxter-King's bandpass filter and first-differences. By contrast with investment rates, there is no evidence of the heteroscedastic behavior predicted by Ss -type models for GDP, in fact, in 3 out of 6 cases considered the estimated value for η has the wrong sign. This suggests that it is not the shocks, which presumably affect both investment and GDP, that drive our heteroscedasticity findings for investment, but the mechanism that transmits these shocks into aggregate investment.

Table 12: Evidence of heteroscedasticity: U.S. GDP

Detrending:	HP	HP	BK	BK	Diff	Diff
Model:	1	2	1	2	1	2
p^* :	12	5	2	7	2	2
k^* :	10	10	12	2	2	7
$\eta \times 10^3$:	60.07	1.135	32.25	-0.113	-117.5	-4.012
t- η :	1.3937	1.4729	1.4054	-0.9021	-1.5291	-1.8326
p -value($\eta > 0$)-bootstrap:	0.173	0.140	0.224	0.740	0.867	0.938
no. obs. lhs. 1st and 2nd reg.:	172	172	172	172	172	172

The time series models (19) and (20) provide simple and robust approaches to test for the presence of conditional heteroscedasticity. More sophisticated options can be used as well. For example, instead of assuming the parametric relationships in (19) and (20), we could allow for a more general expression of the form $\sigma_t = h(\bar{x}_{t-1}^k)$, where h is estimated non-parametrically. Figure 1 in the introduction plots the estimate we obtain for h (normalized by the average fitted value for σ_t) using $p = k = 6$, a Gaussian kernel, and cross-validation to determine the appropriate bandwidth.

B.3 Calibrating the Maintenance Parameter

To choose parameter values that match the heteroscedasticity present in aggregate U.S. investment, it is useful to summarize the estimated conditional heteroscedasticity schedules (19) and (20) by one statistic. We do this via the signed log-ratio of the 95th and 5th percentile of the fitted values for σ . Our calibration strategy is akin to the indirect inference approach proposed by Smith (1993), since we match a time-series moment informed by our DSGE model. We work with model 2 and consider $p = 1$, because the shocks in the DSGE model are AR(1), and $k = 1$. This is the time series model in Figures 2 and 3.

Table 13 reports estimates of the heteroscedasticity statistic for the frictionless model and the models we considered when calibrating the maintenance parameter χ . The first column reports the value for the range statistic in the actual U.S. investment series. The second column reports the value for a model where capital can be adjusted at no cost ('frictionless model'). Columns 3 through 9 consider various values for the maintenance parameter. In each case the simulated model matches the volatility of sectoral and aggregate investment. For each value of χ we generated a large number of time series of aggregate investment to capital ratios of the same length as the U.S. investment series in our data. We then estimated the range statistic for these series — Table 13 reports the average values.

It follows from the first row of Table 13 that our models with lumpy adjustment match the conditional heteroscedasticity in the actual data much better than a frictionless model. It also follows from the first row that a maintenance parameter of 0.50 generates a first moment of 0.3021 for the range statistic, which is close to the estimated value of 0.2901. We therefore choose $\chi = 0.50$ for our DSGE model with lumpy capital adjustment.

Table 13: HETEROSCEDASTICITY STATISTIC AND MAINTENANCE PARAMETER

Statistic	U.S. I/K	frictionless	χ						
			0	0.10	0.20	0.30	0.40	0.50	0.60
$\pm \log(\sigma_{95}/\sigma_5)$	0.3021	0.0539	0.1830	0.1955	0.2095	0.2261	0.2539	0.2901	0.3207

B.4 Robustness

This section addresses various issues related to the robustness of our calibration. First we show that using alternative moments to summarize the range of heteroscedasticity values does not affect our findings. Second, we show that had we used model 1 instead of model 2 would have biased our results slightly against finding large conditional heteroscedasticity in the data. Third, we show that our sectoral calibration choices do not drive the results. Specifically, we study robustness to using 2-digit and 4-digit data to compute the statistics for the sectoral Solow residuals and the investment rates we use. Fourth, we also discuss robustness to how we detrend sectoral data. We show that using a weighted median as opposed to a weighted average to aggregate sectoral moments leaves the results unaltered. And finally, we experiment with the choice for the total annual standard deviation production units face, studying cases with 0.075 and 0.15, compared to our baseline choice of 0.10.

Table 14 shows the implied values for χ for alternative definitions for the heteroscedastic-

ity range (log-ratio for 90th and 10th percentile, and log-ratio for largest and smallest values over the 184 quarters considered). It shows that $\chi = 0.50$ slightly overestimates the volatility of the responsiveness index, if we use the ratio of the maximum to minimum index – the 50% mentioned in the introduction and slightly underestimates it, had we use the ratio of the 90 percentile to the 10 percentile. But clearly, all calibration targets are close, whereas the frictionless and Khan-Thomas models cannot match conditional heteroscedasticity independently of how it is defined.

Table 14: HETEROSCEDASTICITY RANGE

Model	$\log(\sigma_{95}/\sigma_5)$	$\log(\sigma_{90}/\sigma_{10})$	$\log(\sigma_{\max}/\sigma_{\min})$
<i>Data</i>	0.3021	0.2558	0.3841
This paper:	0.2901	0.2183	0.4063
Frictionless:	0.0539	0.0405	0.0844
Khan-Thomas (2008):	0.0468	0.0391	0.0675

Table 15: HETEROSCEDASTICITY RANGE - TIME SERIES MODEL 1

Model	$\log(\sigma_{95}/\sigma_5)$	$\log(\sigma_{90}/\sigma_{10})$	$\log(\sigma_{\max}/\sigma_{\min})$
<i>Data</i>	0.3175	0.2689	0.3964
This paper:	0.2880	0.2184	0.4008
Frictionless:	0.0649	0.0493	0.0909
Implied χ	0.60	>0.60	0.50

Table 15 shows that had we used Model 1 in the conditional heteroscedasticity regression, we would have found a slightly stronger variation of the responsiveness index in the data, by 50% ($e^{0.3964} \simeq 1.49$), and, consequently, the calibrated maintenance parameter would have been higher.

We now go back to the calibration based on Model 2 and check robustness to the choices we made in obtaining the sectoral statistics we use in the calibration. Table 16 shows that our baseline calibrated value for the maintenance parameter, $\chi = 0.5$, is robust to these choices. Also, in any of these cases the frictionless model features a heteroscedasticity range around 0.05, well below the data's 0.3021.

Table 16: ROBUSTNESS - CALIBRATION

Specification	Cond. Adj. Costs/ Unit's Output	Calibrated χ	σ_S	ρ_S	Target Sectoral I/K Volatility	$\log(\sigma_{95}/\sigma_5)$ FL
<i>Baseline</i>	3.6%	0.50	0.0273	0.8612	0.0163	0.0539
Lin. detr. Solow res.	20.2%	0.40	0.0276	0.8979	0.0163	0.0537
Quadr. detr. I/K	4.4%	0.50	0.0273	0.8612	0.0152	0.0539
2-digit Data	7.4%	0.50	0.0166	0.8764	0.0098	0.0536
4-digit Data	1.4%	>0.60	0.0369	0.8492	0.0226	0.0539
Weighted Median	5.7%	0.50	0.0244	0.8727	0.0145	0.0537
Total $\sigma = 0.075$	3.3%	0.50	0.0273	0.8612	0.0163	0.0538
Total $\sigma = 0.15$	2.9%	>0.60	0.0273	0.8612	0.0163	0.0547

C Numerical Appendix

In this appendix, we describe in detail the numerical implementation of the model computation. Unless otherwise stated, the numerical specifications refer to the baseline calibration in the main text, although most of them are common across all models.

C.1 Decision Problem

Given the assumptions we made in the main paper: 1) $\rho_S = \rho_I = \rho$, and 2) approximating the distribution μ by the aggregate capital stock, \bar{k} , the dynamic programming problem has a 4-dimensional state space: $(k, \bar{k}, z, \epsilon)$. Since the employment problem has an analytical solution, there is essentially just one continuous control, k' .

We note that for all partial equilibrium computations the dimension of the state space collapses to three, \bar{k} is no longer needed to compute prices and aggregate movements. Instead, we follow KT in fixing the intertemporal price and the real wage at their average levels from the general equilibrium simulations.

Since we allow for a continuous control, k , and \bar{k} and z can take on any value continuously, we can only compute the value function exactly at the grid points above and interpolate for in-between values. This is done by using a multidimensional cubic splines procedure, with a so-called “not-a-knot”-condition to address the large number of degrees of freedom problem, when using splines (see Judd, 1998). We compute the solution by value function iteration, using 20 steps of policy improvement after each actual optimization procedure. The optimum is found by using a golden section search. Upon convergence, we check single-peakedness of the objective function, to guarantee that the golden section search is reasonable.

Table 17: Assessing agents' forecasting rules for capital

	FL	Baseline	Baseline-SKEW
$a_{\bar{k}}$	0.0065	0.0021	0.0112
$b_{\bar{k}}$	0.9061	0.9473	0.9388
$c_{\bar{k}}$	0.2199	0.1184	0.1106
$d_{\bar{k}}$	NaN	NaN	0.0075
$e_{\bar{k}}$	NaN	NaN	-0.0008
R^2	1.0000	0.9999	1.0000
SE	0.0000	0.003	0.0001
MAD(%)	0.09	0.61	0.34
MSE(%)	0.04	0.30	0.14
Correl.	1.0000	0.9956	0.9992

C.2 Equilibrium Simulation

For the calibration of the general equilibrium models we draw one random series for the aggregate technology level and fix it across models. We use $T = 600$ and discard the first 100 observations. For computing the conditional heteroscedasticity in the model simulations we use a much longer simulation horizon of $T = 10000$. We find that, generally, the statistics are robust to T . We start from an arbitrary individual capital distribution and the stationary distribution for the combined productivity level. The model economies typically settle fast into their stochastic steady state after roughly 50 observations. Since with idiosyncratic shocks, adjustment costs and necessary maintenance some production unit may not adjust for a very long time, we take out any individual capital stock in the distribution that has a marginal weight below 10^{-10} , in order to save on memory. We re-scale the remaining distribution proportionally.

As in the production unit's decision problem, we use a golden section search to find the optimal target capital level, given p . We find the market clearing intertemporal price, using a combination of bisection, secant and inverse quadratic interpolation methods. Precision of the market-clearing outcome is better than 10^{-7} .

To further assess the quality of the assumed log-linear equilibrium rules, we perform the following simulation: for each point in the $T = 500$ (we discard the first 100 observations) time series, we iterate for a time series of $\tilde{T} = 100$ aggregate capital and the intertemporal price forward, using only the equilibrium rules and assuming the actual time path for aggregate technology. We then compare the aggregate capital and p after \tilde{T} steps with the actually simulated ones, when the equilibrium price was updated at each step. We then compute maximum absolute percentage deviations, mean squared percentage deviations, and the correlation between the simulated values and the out-of-sample forecasts. Tables 17 and 18 summarize the numer-

Table 18: Assessing agents' forecasting rules for p

	FL	Baseline	Baseline-SKEW
a_p	1.8438	1.8489	1.8748
b_p	-0.3357	-0.2442	-0.2701
c_p	-0.5836	-0.8020	-0.8215
d_p	NaN	NaN	0.0213
e_p	NaN	NaN	-0.0031
R^2	1.000	0.9992	0.9999
SE	0.0000	0.0006	0.0002
MAD(%)	0.02	0.19	0.11
MSE(%)	0.01	0.06	0.03
Correl.	1.000	0.9997	0.9999

ical results for each model. The rows contain: the coefficients of the log-linear regression, its R^2 and standard error and the three above measures that assess the out-of-sample quality of the equilibrium rules. They assess the log-linear approximation for future capital and current p , respectively. Baseline-SKEW refers to our baseline calibration, where agents use additionally the log standard deviation and skewness of the capital distribution for forecasting.

Table 17 shows that there exists a good log-linear approximation for aggregate capital as a function of last period's capital and the current aggregate shock. This may seem surprising in light of the time-varying impulse response functions we described in the main text. However, the numbers also show that in particular out-of-sample forecasts improve, when higher moments of the capital distribution are introduced. Furthermore, as we argue next, the goodness-of-fit for an equation analogous to (11a) and (11b), but with the aggregate investment rate as the dependent variable, is worse, even though the poorer fit has no bearing on aggregate investment dynamics.

Table 19: Assessing agents' forecasting rules for I/K

Highest moment	R^2						Autocorrelation	
	all	1st quart.	2nd quart.	3rd quart	4th quart.	average	1st	2nd
Baseline								
Mean:	0.9896	0.9535	0.7859	0.7259	0.9501	0.8538	0.906	0.816
St. deviation:	0.9992	0.9947	0.9869	0.9822	0.9961	0.9900	0.922	0.846
Skewness:	0.9998	0.9986	0.9975	0.9978	0.9988	0.9982	0.919	0.841

We simulated a series of 500 observations for our baseline model, assuming that agents use

the first, the first two and the first three moments of capital in their forecasting rules.³⁸ We divided the simulated series into quartiles based on the magnitude of the actual investment rate, and calculated, for each quartile, the R^2 -goodness-of-fit statistic between the aggregate investment rate series implied by the forecasting rule and the “true” aggregate investment rate series, which we assume to be the one generated, when agents use three moments of the capital distribution for forecasting.

Table 19 shows our results. The average (across quartiles) R^2 between the log-linear approximation and the true investment rate is only 0.85 for the baseline model. This average increases to 0.99 (0.97) when the log-standard-deviation of capital is added as a regressor, and to well above 0.99 when the skewness statistic is included as well.³⁹ The last two columns of Table 19 show that the estimated first and second order autocorrelations of the investment rate also improve significantly when using higher moments in the forecasting rules: the corresponding values for the actual investment rate series are 0.919 and 0.842, respectively, for the baseline calibration.

However, we also recomputed the evolution of the aggregate investment rate, when agents use the rules that include higher moments of capital, and found no discernible differences with what we obtained with the log-linear forecasting rules: the correlation coefficient between the sample paths of I/K generated with forecasting rules with and without higher moments is above 0.9999.

C.3 Sectoral Simulation

Underlying the sectoral simulation are four assumptions: First, we assume a large enough number of sectors. Second, we assume a large σ_S/σ_A ratio, so that we can compute the sectoral implications of our model independently of the aggregate general equilibrium calculations. This is also reflected in our treatment of the sectoral data as residual values, which are uncorrelated with aggregate components. Third, we make use of the assumption that a sector is large enough to use a law of large numbers for the true idiosyncratic productivity shocks. Fourth, we assume that $\rho_S = \rho_I$, and that sectoral and idiosyncratic productivity shocks are independent, so that we can treat sectoral and idiosyncratic uncertainty as one state variable in the general equilibrium problem.

We start by fixing the aggregate technology level at its average level: $z^{SS} = 1$. The converged equilibrium law of motion for aggregate capital can then be used to compute the steady state

³⁸More precisely, the first case has the log-mean of capital holdings as a regressor, the second case adds the log-standard deviation and the third case also incorporates the skewness of capital holdings. Of course, $\log z_t$ is a regressor in all cases.

³⁹For the frictionless model the first part of the first row would read: 0.9981, 0.9887, 0.9774, 0.9796, 0.9841, 0.9825. And the autocorrelations for forecasted investment rates are almost identical to the ones for the actual series.

aggregate capital level that belongs to this aggregate productivity. It is the fixed point of the aggregate law of motion, evaluated at z^{SS} :

$$\bar{k}^{SS} \equiv \exp \frac{a_{\bar{k}}}{1 - b_{\bar{k}}}.$$

This, in turn, leads to the steady state $p^{SS} \equiv \exp(a_p + b_p \log(\bar{k}^{SS}))$.

Then we specify a separate grid for idiosyncratic and sectoral productivity in such a way that all new grid points and any product of them will lie on the original 19-state grid for the combined productivity, used in the general equilibrium problem. Given the equi-spaced (in logs) nature of the combined grid this is obviously possible. Thus, the idiosyncratic grid comprises 11 grid points, and the sectoral grid 9 grid points, both equi-spaced and centered around unity.

Next, we recompute optimal target capital levels as well as gross values of investment (see equation 12d) at z^{SS}, \bar{k}^{SS} , at the 19 values for ϵ . By construction, these are then also the values for any (ϵ_S, ϵ_I) -combination. Note that we use the value functions computed from the general equilibrium case. We draw a random series of $T = 2600$ for ϵ_S , which remains fixed across all models, start from an arbitrary capital distribution and the stationary distribution for the idiosyncratic technology level, and follow the behavior of this representative sector, using the sectoral policy rules. The details are similar to those of the equilibrium simulation.

Finally, we test the two main assumptions on which we base our sectoral computations: a continuum of sectors and fixing the aggregate environment at its steady state level. To this end, we compute the equilibrium with a finite number of sectors, N_S . Also, we introduce an additional state-variable, given by: $\bar{\epsilon}_{S,t} \equiv \sum_{i=1, \dots, N_S} \log(\epsilon_{S,t}(i))$, which captures changes in the aggregate environment, beyond the common aggregate shock. Obviously, $\bar{\epsilon}_{S,t} = 0, \forall t$, as $N_S \rightarrow \infty$, by the law of large numbers and assuming sectoral independence. This additional aggregate state is then integrated over by Gauss-Hermitian integration, which is facilitated by the fact that the $\bar{\epsilon}_{S,t}$ -process is independent of the aggregate technology process (by assumption). For computational reasons - following a large number of sectors with a large number of production units each is considerably more onerous in a quarterly calibration than in a yearly calibration -, we run these robustness checks for the annual equivalents of our baseline models.

We choose two different values for N_S . First, 400, which roughly equals the number of 3-digit SIC-87 sectors in the U.S. (395). Since, however, sectors are of very different size and overall importance, and also often correlated, we decrease, secondly, N_S to 100 for robustness reasons. The resulting residual $\sigma_{\bar{\epsilon}_S}$ is 0.0026 and 0.0052, respectively. Notice that in both cases $\sigma_{\bar{\epsilon}_S}$ is considerably smaller than $\sigma_A = 0.0120$, the σ_A for the annual frictionless calibration, so that we should not expect too large an effect from this additional source of aggregate uncertainty.

The following table shows the aggregate and sectoral standard deviations for annual invest-

ment rates for the frictionless model and our baseline lumpy model ($\chi = 0.5$). The raw sectoral standard deviations are shown as a capital-weighted average (the unweighted averages are only insignificantly different). The residual sectoral standard deviations are shown with the same filtering operations as discussed in Appendix A.3.

Table 20: ROBUSTNESS OF THE SECTORAL COMPUTATION

Model:	FL	FL	Lumpy	Lumpy
Number of sectors:	100	400	100	400
Aggr. St.dev.	0.0113	0.0102	0.0103	0.0099
Sect. St.dev. - raw	0.1824	0.1838	0.0190	0.0188
Sect. St.dev. - res.	0.1819	0.1834	0.0159	0.0160

The first important observation is that the numbers obtained here are not much different from what we have obtained in the simplified computation, which is in particular true for the lumpy model. Specifically, the frictionless model continues to fail to match observed sectoral volatility by an order of magnitude. And, secondly, the numbers deviate in the expected direction: the aggregate standard deviation increases (from 0.0098), because there is an additional aggregate shock, but only slightly so; the sectoral standard deviations decrease a little bit (from 0.0163), because now general equilibrium forces contribute also to sectoral smoothing. Overall, our simplified sectoral simulations seem justified.

D Matching Establishment Statistics

One argument we gave for using sectoral rather than plant level data to calibrate micro frictions is that matching micro moments may not be a robust way of pinning down microeconomic parameters when the goal is to use these parameters to identify aggregate effects of the mechanism. In this appendix we provide support to this claim by showing that a straightforward modification of the micro underpinnings of our baseline model leads to a satisfactory match of establishment level moments. More important, the match of sectoral and aggregate moments we obtained in the main text is unaffected by this extension. Our objective here is not to add realism to our original model, but to illustrate the potential lack of power of using (only) plant level data for our purpose.

D.1 A Simple Extension

A first choice we need to make when matching the model to micro data is how many micro units in the model correspond to one establishment. Choices by other authors have covered a wide range, going from one to a continuum (see footnote 15).

Two additional issues arise if we choose to model an establishment as the aggregation of many micro units. First, we must address the extent to which shocks—both to productivity and to adjustment costs—are correlated across units within an establishment.⁴⁰ Second, we must take a stance on the fact that establishments sell off and buy what in our model corresponds to one or more micro units.

Next we present a simple model that incorporates both elements mentioned above. The economy is composed of sectors (indexed by s), which are composed of establishments (indexed by e), which are composed of units (indexed by u). Data are available at the establishment level but not at the unit level. The log-productivity shock faced by unit u in establishment e in sector s at time t is decomposed into aggregate, sectoral, establishment and unit level shocks as follows:

$$\log z_{uest} = \log \varepsilon_t^A + \log \varepsilon_{st}^S + \log \varepsilon_{est}^E + \log \varepsilon_{uest}^U,$$

where $\log \varepsilon_t^A \sim \text{AR}(1; \rho_A, \sigma_A)$, $\log \varepsilon_{st}^S \sim \text{AR}(1; \rho_S, \sigma_S)$, $\log \varepsilon_{est}^E \sim \text{AR}(1; \rho_E, \sigma_E)$ and $\log \varepsilon_{uest}^U \sim \text{AR}(1; \rho_U, \sigma_U)$, and the usual orthogonality assumptions hold.⁴¹ Consistent with the assumptions we made in the paper, we set $\rho_S = \rho_E = \rho_U$ and denote the common value by ρ .

An establishment is composed of a large number (continuum) of units. The extent to which the behavior of units within an establishment is correlated varies with the relative importance of σ_U and σ_E . The larger σ_E , the higher the correlation of productivity shocks across units and the more coordinated their investment decisions will be. The sectoral and aggregate investment series generated by the extended model are the same as those generated by the model developed in the main text as long as $\sigma_E^2 + \sigma_U^2 = \sigma_I^2$, since all we do in this extension is group micro units into “establishments” in a way that has no implication for sectoral aggregates. We consider the polar cases with uncorrelated productivity shocks ($\sigma_U^2 = \sigma_I^2$, $\sigma_E^2 = 0$) and perfectly correlated shocks ($\sigma_U^2 = 0$, $\sigma_E^2 = \sigma_I^2$). The degree of coordination also depends on how correlated adjustment costs are across units within an establishment; again we consider the polar cases where adjustment costs are perfectly correlated and independent.⁴²

Regarding the sale and purchase of micro units, we assume that in every period the capi-

⁴⁰For tractability, we assume that decisions are made at the micro-unit level, not the establishment level.

⁴¹ $x_t \sim \text{AR}(1; \rho, \sigma)$ means that the process x_t follows an AR(1) with first order autocorrelation ρ and standard deviation of innovations equal to σ .

⁴²There is a one-to-one match between micro units in the model and establishment level data when adjustment costs and productivity shocks are perfectly correlated. Otherwise a continuum of model micro units correspond to one establishment in the data.

tal stock that is recorded at the establishment level is related to the capital stock for that unit determined by our baseline model via

$$K_{est}^r \equiv (1 + \tau_{est})K_{est}^m, \quad (22)$$

where the superscripts r and m stand for “recorded” and “model”. The τ ’s are i.i.d. draws from a normal distribution with zero mean and standard deviation σ_τ .⁴³

The capital accumulation identity for recorded investment and (22) lead to:

$$I_{est}^r = K_{es,t+1}^r - (1 - \delta)K_{est}^r = (1 + \tau_{es,t+1})K_{es,t+1}^m - (1 + \tau_{est})(1 - \delta)K_{est}^m. \quad (23)$$

It follows that if the establishment sells off a fraction of the units it holds in $t + 1$, these sales will show up as lower (or even negative) investment in the recorded investment data.

Dividing both sides of (23) by K_{est}^r , using (22) and denoting investment rates I_t/K_t by i_t yields

$$i_{est}^r = \frac{(1 + \tau_{es,t+1})K_{es,t+1}^m}{(1 + \tau_{est})K_{est}^m} - (1 - \delta).$$

Using the capital accumulation identity to express $K_{es,t+1}^m/K_{est}^m$ in terms of the model’s investment rate and ignoring second order terms in τ then leads to

$$i_{est}^r \approx (1 - \Delta\tau_{es,t+1})i_{est}^m + \Delta\tau_{es,t+1}(1 - \delta), \quad (24)$$

with $\Delta\tau_{es,t+1} \equiv \tau_{es,t+1} - \tau_{est}$.

Summing up, our (admittedly simple) extension introduces three parameters—the degree of correlation of productivity and adjustment costs across units within establishments, and the volatility parameter for unit purchases and sales—that can be used to fit establishment level moments without affecting the match of sectoral and aggregate statistics.

D.2 Matching Establishment Level Statistics

For the four combinations of correlation across both sources of shocks we generate a histogram with 2,500 realizations of establishment level I/K using our model.⁴⁴

Denote by f_i , $i = 1, \dots, 5$ the fraction of LRD establishments that adjusted less than -20% , between -20 and -1% , between -1% and 1% , between 1 and 20% and above 20% , respec-

⁴³We choose a symmetric distribution so that asymmetries in the histogram of investment rates cannot be attributed to this choice.

⁴⁴We compute these investment rates using the approximation described in Appendix C.3 with $\sigma_S^2 + \sigma_E^2$ in the role of σ_S^2 , and $\sigma_I^2 - \sigma_E^2$ in the role of σ_I^2 .

tively. And denote by $\pi_i(\sigma_\tau)$ the fraction of units with adjustment in the previous bins after applying the transformation described in (24). We choose the value of σ_τ that minimizes $\sum_i |f_i - \pi_i(\sigma_\tau)|/f_i$, that is, that minimizes the average absolute deviation.

Table 21 presents our results. We consider four combinations of correlation among adjustment and productivity shocks. Comparing the first four rows with the last row shows that the match we obtain for statistics of the plant level distribution is reasonable.⁴⁵ More important, this exercise illustrates that establishment level moments may not be useful to calibrate model parameters that play an important role determining aggregate dynamics.

Table 21: MATCHING LRD MOMENTS

Model		f_1	f_2	f_3	f_4	f_5	σ_τ	avge. abs. deviation
correl. adj. costs	correl. prod. shocks							
0	0	.022	.040	.015	.833	.090	.074	.380
0	1	.000	.090	.046	.729	.135	.056	.377
1	0	.031	.081	.043	.696	.149	.037	.331
1	1	.000	.090	.079	.726	.105	.035	.386
<i>Data</i>		.019	.090	.082	.622	.187	—	—

⁴⁵The goodness of fit is similar to that obtained by KT, which is 0.303.