

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

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

June 11, 2008

Abstract

The sensitivity of U.S. aggregate investment to shocks is procyclical: the initial response 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 Olivier Blanchard, William Brainard, Jordi Galí, Pete Klenow, John Leahy, Giuseppe Moscarini, Anthony Smith, Julia Thomas and seminar/meeting participants at the AEA (Chicago), Bonn, Cornell, Econometric Society (Bogotá), Karlsruhe, Mainz, NBER-EFG, NYU, SITE, U. de Chile (CEA) and Yale for their comments on an earlier version (April, 2006) of this paper, entitled "Lumpy Investment in Dynamic General Equilibrium." Financial support from NSF is gratefully acknowledged.

1 Introduction

U.S. non-residential private fixed investment exhibits conditional heteroscedasticity. Figure 1 depicts a smooth, nonparametric estimate of the heteroscedasticity of the residual from fitting an AR(1) process to quarterly aggregate investment rate from 1960 to 2005, as a function of the average recent investment rate (see Appendix C 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 roughly 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, investment slumps are hard to recover from.

More broadly, our calibrated model suggests that over the 1960-2005 period the initial response of investment to a productivity shock in the top quartile is 32% higher than the average response in the bottom quartile. Differences go beyond the initial response. The left panel in Figure 2 depicts the response over time to a one standard deviation shock taking place at selected points of the U.S. investment cycle: the trough in 1961, a period of average investment activity in 1989 and the peak in 2000.² The variability of these impulse responses is apparent

¹The dotted lines depict \pm one standard deviation confidence bands.

²The figure depicts the impulse responses of the aggregate investment rate at each year, normalized by the average aggregate quarterly investment rate: 0.026.

Figure 2: Impulse Response in Different Years



and large. 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 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.

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 of situations where general equilibrium undoes the partial equilibrium features.

Why do we reach such a different conclusion? Because, implicitly, their particular calibrations impose that the bulk of investment dynamics is 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.

Concretely, 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 smooth and spread over time the response of the endogenous variable. We refer to this process as *smoothing*, and decompose it into its pre-general equilibrium (PE) and general equilibrium (GE) components. In the context of nonlinear lumpy-

adjustment models, PE-smoothing does *not* refer to the existence of microeconomic inaction and lumpiness per se, but to the impact these have on aggregate smoothing. 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 S_s 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 Figure 1.

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 upper entry shows the volatility of aggregate investment rates in our model when neither smoothing mechanism is present (in other words, when there are no adjustment costs at the microeconomic level and no price adjustments in the economy). The intermediate entries incorporate PE- and GE-smoothing, one at a time, while the lower entry considers both sources of smoothing simultaneously. 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 smoothing achieved by both sources.

It is clear that both sources of smoothing do not enter additively, so some care is needed

when quantifying the contribution of each source to overall smoothing. Nonetheless, averaging the upper and lower bounds mentioned above suggests roughly similar roles for both sources of smoothing in our model.³ 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.

Given its centrality in differentiating our answer from that of previous models, 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 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 row shows the same values for our baseline lumpy model and the third row does the same for a model with no microeconomic frictions in investment. The fourth row reports the same statistics for the model in Khan and Thomas (2008), which we discuss later in the paper. It is apparent from this table that the frictionless model

³The upper and lower bounds for the contribution of PE-smoothing are calculated as follows:

$$\begin{aligned}
 UB &= \log[\sigma(\text{NONE})/\sigma(\text{PE})]/\log[\sigma(\text{NONE})/\sigma(\text{BOTH})], \\
 LB &= 1 - \log[\sigma(\text{NONE})/\sigma(\text{GE})]/\log[\sigma(\text{NONE})/\sigma(\text{BOTH})]
 \end{aligned}$$

where NONE refers to the pre-general equilibrium model with no microeconomic frictions, PE to the model that only has microeconomic frictions so that prices are fixed, GE to the model with only GE constraints, and BOTH to the model with both micro frictions and GE constraints.

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 and B. The volatility of aggregate investment rates in Table 2 for Kahn and Thomas is taken from table III in Kahn and Thomas (2008). The volatility of sectoral investment rates is based on our calculations. We add their idiosyncratic shock and our sectoral shock to compute the total standard deviation for the PE-innovations. The lumpy model in Kahn and Thomas (2008) exhibits larger sectoral volatility than the frictionless counterpart of our lumpy model because of parameter differences between our model and theirs, such as the curvature of the revenue function (see details in section 3). What matters for our purposes is that either one fails to match sectoral volatilities by an order of magnitude.

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 and Section 5 shows the robustness of the main results. Section 6 concludes and is followed by several appendices.

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 (there is fixed proportions and some parts and machines that break down need to be replaced, see, e.g., McGrattan and Schmitz (1999) for evidence on the importance of maintenance investment). Third, the distribution of aggregate productivity shocks is continuous rather than a Markov discretization.⁵

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 ($\theta > 0$, $\nu > 0$, $\theta + \nu < 1$):

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

⁵This allows us to do computations that are not feasible with a Markov discretization. For example, backing out the aggregate shocks that are fed into the model to produce Figure 3.

where z_t , 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 v with zero mean and variance σ_A^2 :

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

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 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, maintenance is carried out, workers are paid, and investment decisions are made. Then the period ends.

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 , but units may find it necessary to replace certain items during the production process.

Define $\bar{\psi} \equiv \frac{\gamma}{1-\delta} > 1$ as the maintenance investment rate needed to fully compensate depreciation and trend growth. The degree of necessary maintenance, 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 undone 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:

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

	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.⁷

Notice that χ is obviously irrelevant for the units that actually adjust at the end of the period. This is not to say that these units do not have to spend on maintenance within the production period, but rather their net capital growth, conditional on incurring the fixed cost and optimal adjustment, is independent of this expenditure. This is essentially a feature of only having fixed adjustment costs, as opposed to more general adjustment technologies that also include a component that depends on the magnitude of capital adjustments.

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 will 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:⁸

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

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

⁷We note that our version of maintenance investment differs from what KT call “constrained investment”. Here, 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.

⁸Note 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.

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:

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

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 (6) 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 D 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. Both assumptions combined allow us to reduce 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.⁹

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

⁹In Appendix D.3 we show that our results are robust to this simplifying assumption.

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 D 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 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 $V_a(\epsilon; z, \bar{k}) \geq V_i(\epsilon, k; z, \bar{k})$.¹⁰ 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,

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

¹⁰The production unit can always choose $i = 0$ and thus $k^* = \psi(1 - \delta)k/\gamma$.

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 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):

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

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 D 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 decomposition of PE- and GE-smoothing, and we calibrate the conditional heteroscedasticity of investment in U.S. data.

3.1 Calibration Strategy

The model period for the baseline model is a quarter. The following parameters have standard values: $\beta = 0.9942$, $\gamma = 0.004$, $\nu = 0.64$, and $\rho_A = 0.95$. The log-felicity function features an elasticity of intertemporal substitution (EIS) of one. The depreciation rate δ is picked to match 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 , σ_A and the parameters of the sectoral and idiosyncratic technology process (ρ_S , σ_S , ρ_I and σ_I). This is followed by a detailed discussion of how we calibrate the adjustment cost parameter, $\bar{\xi}$, and the maintenance parameter, χ , which are at the heart of our calibration strategy.

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.¹¹ For comparability in the second moments, σ_A is picked to make both the lumpy and the frictionless models match the volatility of the quarterly aggregate investment rate (0.0023) perfectly.¹²

We determine σ_S and ρ_S by a standard Solow residual calculation on annual 3-digit manufacturing data, taking into account sector-specific trends and time aggregation (see Appendices A and B for details). σ_S equals 0.0273 and ρ_S 0.8612. For computational feasibility we set $\rho_I = \rho_S$, and σ_I to 0.0472, which makes the annual total standard deviation of sectoral and idiosyncratic shocks 0.10. We turn now to the calibration of the two key parameters of the model, $\bar{\xi}$ and χ .

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 (e.g., KT) assume that this correspondence is one-to-one, while other authors (e.g., Abel and Eberly (2002) and Bloom (2007)) match a large number—a continuum and 250, respectively—of model-micro-units to one observed productive unit (firm or establishment).

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

¹¹In 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).

¹²See Appendix A for the values. For annual calibrations, we target 0.0098 as the volatility of the aggregate investment rate.

statistic, even in partial equilibrium. For example, anything that changes the drift of mandated investment (such as maintenance investment), 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.¹³

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. In this approach we assume 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.

Given a set of parameters, the sequence of sectoral investment rates is generated as follows: 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.¹⁴

The value of sectoral volatility of annual investment rates we match is 0.0163.¹⁵ As noted in the introduction, this number is one order of magnitude smaller than the one predicted by the frictionless model.

Finally, we calibrate the maintenance parameter χ by matching the logarithm of the ratio between the maximum and minimum of the estimated values for the conditional heteroscedasticity; we refer to this statistic as the *heteroscedasticity range* in what follows. That is, given a quarterly series of aggregate investment rates, x_t , the moment we match is obtained by first regressing the series on its lagged value and then regressing the absolute residual from this re-

¹³In Appendix E we present a simple extension of the paper's main model, to show how by adding two micro parameters with no macroeconomic or sectoral consequences one can obtain a very good fit of observed micro moments. The problems of matching micro moments and matching aggregate moments are orthogonal in this extension.

¹⁴Appendix D.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 by this extension.

¹⁵We time-aggregate the quarterly investment rates generated by the model to obtain this number. For details, on how we compute this number on the data, see Appendix B.2.

gression, $|\hat{e}_t|$, on x_{t-1} (both regressions are estimated via OLS):

$$|\hat{e}_t| = \hat{\alpha}_0 + \hat{\alpha}_1 x_{t-1} + \text{error}. \quad (16)$$

Denoting by σ_{\max} and σ_{\min} the largest and smallest fitted values from the regression in (16), the heteroscedasticity range is equal to $\pm \log(\sigma_{\max}/\sigma_{\min})$, with a positive sign when the maximum lies to the right of the minimum and a negative sign otherwise. The target value for the heteroscedasticity range in the data is 0.3971, which implies a variation in the initial response to shocks that increases by approximately 50% from the trough to the peak of the business cycle ($e^{0.3971} \simeq 1.49$). Of course, when simulating our model to calculate the average heteroscedasticity range for given parameter values, the length of the simulated series is equal to the length of the actual data (184 quarterly observations).

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. The average cost actually paid is much lower, as shown in Table 3, since productive units wait for good draws to adjust. Conditional on adjusting, a production unit pays 9.53% of its quarterly output (column 3) or, equivalently, 14.88% of its regular wage bill (column 4). To be able to compare these findings with the annual adjustment cost estimates in the literature, we also report these numbers for an annual analogue of the quarterly model. With 3.60% and 5.62%, respectively, they appear to be at the lower end of the literature (see Caballero and Engel (1999), Cooper and Haltiwanger (2006) as well as Bloom (2007)). The first two columns report the aggregate resources spent on adjustment, as a fraction of aggregate output and aggregate investment, respectively.

Table 3: THE ECONOMIC MAGNITUDE OF ADJUSTMENT COSTS

Model	Tot. adj. costs/ Aggr. Output (1)	Tot. adj. costs/ Aggr. Investment (2)	Adj. costs/ Unit Output (3)	Adj. costs/ Unit Wage Bill (4)
Lumpy quarterly:	0.35%	2.41%	9.53%	14.88%
Lumpy annual:	0.41%	2.84%	3.60%	5.62%

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 degree of conditional het-

eroskedasticity in aggregate data. 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 of the data.

Table 4: HETEROSCEDASTICITY RANGE

Model	$\log(\sigma_{\max}/\sigma_{\min})$
<i>Data</i>	0.3971
This paper:	0.4008
Frictionless:	0.0767
Khan-Thomas (2008):	0.0998

Ultimately, the main difference between our calibration and KT is the size of the adjustment cost. Tables 5 and 6 make this point. The former reports upper and lower bounds for the contribution of PE-smoothing to total smoothing, for several models, at different frequencies. The main message can be gathered from the first two rows of these tables. In Table 5 we see that by changing the adjustment cost distribution in KT's model for ours,¹⁶ 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 seven 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 parameters. Table 6 shows the economic magnitudes of the different assumptions on adjustment costs.

Table 5: SMOOTHING DECOMPOSITION: KT

<u>Model</u>	<u>PE/total smoothing</u>		
	Lower bd.	Upper bd.	Avg.
KT-Lumpy annual (ECMA 2008):	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 (50% maint.):	20.0%	76.7%	48.3%
Our model quarterly (0% maint.):	14.5%	80.9%	47.7%
Our model quarterly (50% maint.):	15.4%	81.0%	48.2%

¹⁶Since 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.

Table 6: THE ECONOMIC MAGNITUDE OF ADJUSTMENT COSTS: KT

Model	Tot. adj. costs/ Aggr. Output	Tot. adj. costs/ Aggr. Investment	Adj. costs/ Unit Output	Adj. costs/ Unit Wage Bill
KT-Lumpy annual:	0.22%	1.13%	0.50%	0.77%
Our model annual (0% maint.):	1.80%	12.86%	38.95%	60.86%
Our model annual (50% maint.):	0.41%	2.84%	3.60%	5.62%
Our model quarterly (0% maint.):	1.49%	10.50%	97.08%	151.69%
Our model quarterly (50% maint.):	0.35%	2.41%	9.53%	14.88%

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 7 and 8 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 so as to match the volatility of the aggregate investment rate.¹⁷

Table 7: 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

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.¹⁸ It also increases slightly the persistence of most aggregate variables, bringing these statistics closer to their values in the data.

¹⁷The 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.

¹⁸Consistent 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.

Table 8: 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

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 defined in Caballero and Engel (1993b) for the 1960-2005 period (in percentage deviations from its “steady state” value). The solid and dashed lines represent the index for the lumpy and frictionless models, respectively, while the vertical lines denote NBER business cycle dates.¹⁹ This index captures the response upon impact of the aggregate investment rate to an innovation. At each point in time, this index is calculated conditional on the history of shocks, summarized by the current dis-

¹⁹We use the term “steady state” to refer to 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 zero aggregate innovations for 300 periods, but using the policy functions under the assumption of an economy subject to aggregate shocks.

tribution of capital across units (see Appendix F for the formal definition). That is, the index corresponds to the first element of the impulse response conditional on the cross-section of capital in the given year. The shocks fed into the model are those that allow us 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.

The figure confirms the statement in the introduction according to which the initial response to an aggregate shock varies significantly over time, as does the responsiveness index which takes values between 0.0161 and 0.0243; this means the responsiveness of the economy differs by 51% between trough and peak. By contrast, the frictionless model's responsiveness index and impulse responses exhibit very little variation: they vary by only 12% between trough and peak.

To explain how lumpy adjustment models generate time-varying impulse responses, we consider a particular sample path that is roughly designed to mimic the boom-bust investment episode in the U.S. during the last decade. For this, we simulate the paths of the frictionless and lumpy economies that result from a sequence of twenty consecutive positive aggregate productivity innovations half the size of the respective model's standard deviation, followed by a long period where the innovations are equal to zero. The peak investment rate in the path of the lumpy model is 2.96%, compared to 3.09% in the data. Both economies start from their respective steady states.

Figure 4: The Aggregate Investment Rate in a Boom-Bust Episode




Figure 4 shows the evolution of the aggregate investment rates (as log-deviation from their

steady state values) for both economies. There are important difference between them: While at the outset of the boom phase their values are similar, eventually the investment rate in the lumpy economy reacts by more than the frictionless economy to further positive shocks. The flip side of the lumpy economy's larger boom is a more protracted decline in investment during the bust phase. Let us discuss these two phases in turn.

Figure 5: The Responsiveness Index in a Boom-Bust Episode





Figure 5 plots the evolution of the responsiveness index (its log-deviation from steady state), both for the lumpy model and for the frictionless model. 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 index of responsiveness in the lumpy economy exceeds that of the frictionless economy, which explains the larger investment boom observed in the lumpy economy after a history of positive shocks.

Figure 6 illustrates why the PE-smoothing mechanism weakens as the boom progresses. It shows the cross-section of mandated investment (and the probability of adjusting, conditional on mandated investment) at three points in time: the beginning of the episode with the economy at its steady state (solid line), the peak of the boom (dashed line) and the trough of the

Figure 6: Investment Boom-Bust Episode: Cross-section and Hazard



cycle (dotted line).²⁰ It is apparent from this figure that during the boom the cross-section of 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 6 depict the state-dependent adjustment hazard; that is, the probability of adjusting conditional on the corresponding value of 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

²⁰See 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.

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 7, which decomposes the responsiveness index 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 variations in the fraction of adjusters, that is, by the extensive margin.

Figure 7: Decomposition of Responsiveness Index: 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 8. Both series are in log-deviations from their steady state values. This is consistent with what Doms and Dunne (1998) documented for establishment level investment in the U.S., 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.

Figure 8: Decomposition of I/K into Intensive and Extensive Margins






Figure 9: Aggregate Capital



Let us now turn to the bust phase. Figure 9 illustrates the “overaccumulation” of capital resulting from the large investment boom in the lumpy economy. As a result of this boom, once the positive shocks subside, the economy experiences an “overhang” that leads to the protracted investment slump shown in Figure 4 and the sharp drop in the responsiveness index shown in Figure 5.

Figure 10: Impulse Responses of the Aggregate Investment Rate at the Trough



The observed slump in the responsiveness index has important implications for the economy’s ability to return to its steady state investment rate, as the latter becomes unresponsive to positive stimuli, such as a positive aggregate shock or policy intervention (e.g., an investment tax credit). Figure 10 illustrates this mechanism by plotting the impulse responses of the frictionless and lumpy economies following a positive aggregate shock that takes place at their respective troughs.²¹ The more sluggish response of investment in the lumpy economy is apparent.

As shown in Figure 3 at the beginning of this section, the insights we obtained from our study of the boom-bust episode apply more generally. On average, investment responds more to aggregate shocks after a sequence of above-average shocks, than after a sequence of below-average shocks. 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 our model.

²¹These impulse responses are plotted in deviations from the paths without the new shock and normalized by the data’s average aggregate quarterly investment rate.

5 Robustness

Our calibration exercise yields a relatively large maintenance coefficient, $\chi = 0.5$, while the limited evidence available suggests values somewhere between 0.25 and 0.40.²² In this section we show that allowing for larger values of the EIS yields smaller values for the maintenance parameter. We also show that even when the maintenance parameter is set to zero, the model exhibits substantially more conditional heteroscedasticity than a frictionless model.

5.1 High Elasticity of Substitution

Here we consider an alternative calibration to the one we used for the baseline model. We fix the standard deviation of aggregate shocks at the level that is required for the frictionless model to match the observed quarterly aggregate investment volatility in U.S. data ($\sigma_A = 0.0051$) and instead use the EIS to calibrate quarterly aggregate investment volatility. Then the maintenance parameter is again chosen to match the heteroscedasticity range described in Section 3.1. This procedure yields a high value of 20 for the EIS, which is better interpreted as a reduced form parameter for the elasticity of supply of capital faced by the country rather than the elasticity of intertemporal substitution. That is, the EIS in this model is best viewed as a reduced form parameter that measures the strength of GE-smoothing relative to PE-smoothing, where the latter is calibrated directly by $\bar{\xi}$.²³ We also report results we obtain when we apply the calibration strategy described in Section 3.1 with a value of 2 for the EIS; Gruber (2006) obtained this estimate using careful identification based on individual tax data. Table 9 reports the main parameters and moments associated to these calibrations.

The main point to highlight is that our alternative calibration strategy cuts the estimated maintenance parameter in half, from 0.50 to 0.25; using the same calibration strategy with an EIS of 2 also lowers the maintenance parameter, to a value of 0.4. The nonlinearities associated with conditional heteroscedasticity are now reached at a lower level of maintenance. This, in turn, raises the magnitude of both calibrated ($\bar{\xi}$) and realized adjustment costs (rows three to

²²Cooper 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.

²³To implement this calibration strategy, we consider the following family of felicity functions: $\frac{1}{1-\sigma_c} C^{1-\sigma_c} - X^{1-\sigma_c} AN^h$, where σ_c is the inverse of the EIS and X the trend level of aggregate technology. The latter feature guarantees the existence of balanced growth. In a static setting this is easy to see: employment depends on the wage in efficiency units only and the latter is constant under balanced growth. Obviously, if $\sigma_c = 1$, we are back to a standard log-felicity specification. Finally, the discount factor, $\beta \equiv \beta^* \gamma^{(1-\sigma_c)}$ should be interpreted as an effective discount factor, whose value is kept constant across calibrations, so that β^* is a preference parameter that changes with σ_c .

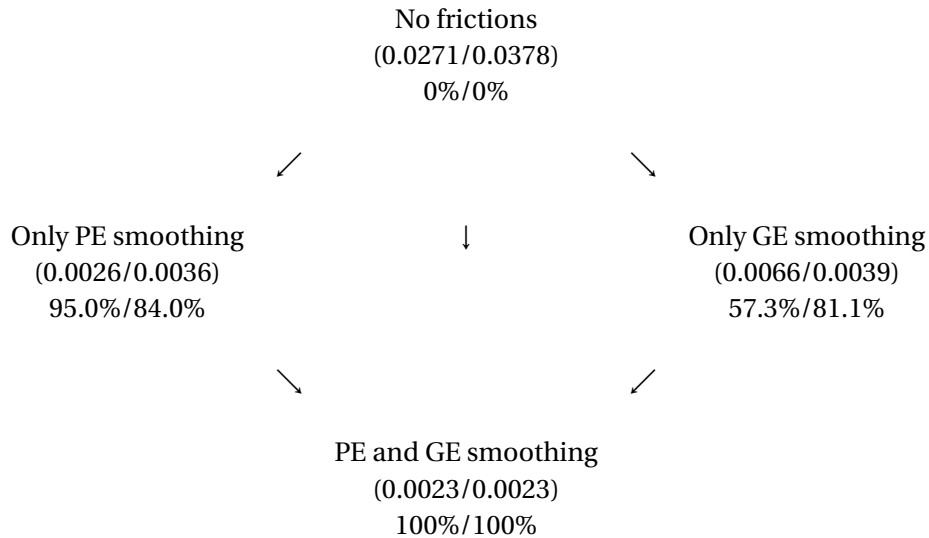
Table 9: HIGH EIS

	χ	$\bar{\xi}$	Total Adj. Costs/ Aggr. Output	Total Adj. Costs/ Aggr. Investment	Adj. Costs/ Unit Ourput	Adj. Costs/ Unit Wage Bill
EIS = 20:	0.25	27.5	0.80	5.51	33.1%	18.2%
EIS = 2:	0.40	14.3	0.50	3.46	51.7%	25.4%

	$\sigma(Y)$	$\sigma(C)$	$\sigma(I)$	$\sigma(N)$	$\rho(Y)$	$\rho(C)$	$\rho(I)$	$\rho(N)$	$\rho(I/K)$
EIS = 20:	1.69	1.25	4.36	1.68	0.71	0.72	0.70	0.71	0.93
EIS = 2:	1.55	1.08	4.35	1.06	0.71	0.71	0.70	0.70	0.92

six).²⁴ Notice, however, that this rise is not due to the increase in the EIS per se — in fact, adjustment cost statistics are largely unresponsive to changes in the EIS —, it is the change of the maintenance parameter that causes the rise. The corresponding values for a model with high EIS but a maintenance parameters of 0.5 are: 0.35%, 2.41%, 9.56% and 14.93%, which are very close to the baseline scenario. For the longitudinal second moments, persistence is almost unchanged with respect to the baseline model, but the volatilities of output, consumption and employment have increased substantially.

Table 10: CONTRIBUTION OF PE AND GE FORCES TO SMOOTHING OF (I/K) : HIGH EIS




In this variation of the model the relative importance of PE-smoothing raises over that of

²⁴The values for rows three to six from an equivalent annual calibration are: 0.96%, 6.70%, 12.67% and 19.80%, which are still very much in line with estimates from the literature.

GE-smoothing. This is apparent in Table 10 (which reports the EIS=20 numbers followed by the EIS=2 numbers). Compared to the baseline scenario, the upper and lower bounds for PE-smoothing increase from 81.0% to 95.0% and 15.4 to 42.7%, respectively, for the high EIS case. The average of both bounds is now close to 70%, compared to approximately 50% for the original model. Finally, Figure 11 shows that the paths of the first element of the impulse response function are nearly identical for the main calibration and the model obtained with the alternative calibration strategy of this section.

Figure 11: Time Paths of the Responsiveness Index - High EIS



5.2 Nonlinearities and Zero Maintenance

Let us now return to the EIS=1 model but lower the maintenance parameter. Thus we no longer match the heteroscedasticity range in the data, but continue to match both sectoral and aggregate investment volatilities. This configuration implies that in order to match PE-smoothing we need to increase the size of adjustment costs. Table 11 illustrates this trade-off: to keep roughly the same magnitude of PE-smoothing (last three columns), adjustment costs decrease as maintenance becomes more important. This mechanism holds for any level of the EIS.

The negative correlation between adjustment costs and maintenance follows from the fact that a higher maintenance parameter lowers the effective drift of mandated investment (since part of depreciation is undone in each period). This is important in these models, as it implies that the cross-section distributions of mandated investment are far from the Caplin-Spulber uniform limit, and hence there is plenty of space for them to vary over time in response to shocks. By contrast, 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.²⁵

Table 11: THE IMPACT OF MAINTENANCE

Model	Adj. costs	Adj. costs	PE/total smoothing		
	Unit Output	Unit Wage Bill	Lower bd.	Upper bd.	Avg.
EIS=1, 0% maint.:	97.08%	151.69%	14.5%	80.9%	47.7%
EIS=1, 25% maint.:	33.02%	51.59%	15.4%	80.9%	48.2%
EIS=1 50% maint.:	9.53%	14.88%	15.4%	81.0%	48.2%

The same reasoning explains why the volatility of the impulse response function increases with the maintenance parameter. An increase in the latter allows for larger countercyclical fluctuations in the degree of PE-smoothing, which exacerbates the magnitude of the response of aggregate investment to shocks in the face of an unusually long string of positive aggregate shocks.

Figure 12: Time Paths of the Responsiveness Index - No Maintenance



The main conclusion for this subsection, however, is shown in figure 12. Even with zero maintenance the responsiveness index of the lumpy economy varies by 30% between trough and peak (from 0.0185 to 0.0241).

²⁵The Caplin-Spulber model is the only S_s model where the impulse response functions do not vary over time, see Caballero and Engel (2007) for details.

6 Final Remarks

We have shown that adding realistic lumpy capital adjustment at the microeconomic level to an otherwise standard RBC model has important macroeconomic implications. In particular, the extended model is able to match the fact that the impulse response functions of aggregate investment, conditional on the history of shocks, varies considerably over time in US data.

The reason for this success is that, relative to the standard RBC model, in the lumpy model investment booms feed into themselves and lead to significantly larger capital accumulation following a string of positive shocks. Busts, on the other hand, can lead to protracted periods of depressed investment, which are largely unresponsive to positive shocks.

References

- [1] Abel, A. and J. Eberly (2002). "Investment and q with Fixed Costs: An Empirical Analysis," mimeo.
- [2] Bloom, N. (2007). "The Impact of Uncertainty Shocks," NBER-WP 13385.
- [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] Gruber, J., (2006). "A Tax-Based Estimate of the Intertemporal Elasticity of Substitution", NBER-WP 11945.
- [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] 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.
- [18] Tauchen, G. (1986). "Finite State Markov-Chain Approximations To Univariate and Vector Autoregressions". *Economics Letters* **20**, 177–181.
- [19] Thomas, J.K. (2002). "Is Lumpy Investment Relevant for the Business Cycle?", *Journal of Political Economy*, **110**(3), 2002, 508–534.
- [20] Veracierto, M., (2002). "Plant Level Irreversible Investment and Equilibrium Business Cycles," *American Economic Review*, **92** (1), March 2002, 181–197.

A Parameter Appendix

The following table 12 summarizes the *common* parameters of the models explored in the paper:

Table 12: 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 (see Appendix B below) 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. For the EIS=2 case we use $\sigma_A = 0.0071$. 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 (see Appendix B below) 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. The use of the EIS as a calibration object is explored in Section 5.1, but the basic principle is the same as with the σ_A -calibration.

B Data Appendix

B.1 Aggregate Data

Since they are not readily available from standard sources, we construct quarterly and yearly series of the aggregate investment rate, based on investment data from the national account and fixed asset tables, available from the Bureau of Economic Analysis (BEA). Time horizon for the aggregate data is 1960-2005. The aggregate investment rate at period t is defined as:

$$\frac{I_t^{real}}{K_{t-1}^{real}},$$

where the denominator is the capital stock at the end of period $t - 1$, deflated by the investment price index of $t - 1$, and the numerator is investment in period t , deflated by period t prices.

For annual investment rates we use the investment and capital data that are readily available from the BEA. The investment data is private fixed nonresidential investment at historical costs from fixed asset table 1.5, line 4 (this includes equipment and structures). The capital data (in current costs, end-of-year estimates) are from table 1.1., line 4 of the fixed asset table. The investment price index is from NIPA table 1.1.9, line 8.

For quarterly investment rates, we have to compute quarterly capital stocks, which are not readily available. We use the following identity:


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

where all variables are nominal, π_t^I denotes the investment price inflation rate from period $t - 1$ to period t , and D_t depreciation in period t . Again, depreciation data are only available on an annual basis. We make the following present value assumption:

$$\frac{D_{Q1}}{(1 + \pi_{Q2})(1 + \pi_{Q3})(1 + \pi_{Q4})} = \frac{D_{Q2}}{(1 + \pi_{Q3})(1 + \pi_{Q4})} = \frac{D_{Q3}}{1 + \pi_{Q4}} = D_{Q4} = \frac{D_Y}{4}.$$

With these two equations quarterly aggregate investment rates can be computed, using NIPA tables 1.1.5 and 1.1.9, line 8, for investment and investment price data, respectively, and fixed asset price tables 1.1. and 1.3., line 4, for capital stocks and depreciations, respectively. For fourth quarter capital stocks, we actually use the annual end-of-year data. The quarterly investment series is available on the authors' websites. As Figure 13 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 13: The Quarterly Aggregate Investment Rate



The following table 13 summarizes statistics of the aggregate investment rate:²⁶

Table 13: 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

B.2 Sectoral Data

For lack of sectoral data outside of manufacturing, the data source here is the NBER manufacturing data set, publicly available on the NBER website. It contains yearly 4-digit industry data for the manufacturing sector, 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 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. Hence, we sum employment levels, real capital, nominal investment and nominal value added onto the 3-digit level. The deflator for

²⁶The maxima are achieved in II/00 (2000), respectively, the minima in I/61 (1975).

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.²⁷

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.²⁸ 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 set σ_S equal to the value-added-weighted average of the estimated standard deviations of the corresponding innovations, which results in $\sigma_S = 0.0501$, and ρ_S equal to the value-added-weighted average of the estimated first-order autocorrelation, which leads to $\rho_S = 0.55$.

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.

²⁷Indeed, using a weighted sum of 4-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.

²⁸Linear trends are flat, as the regression on the manufacturing average takes essentially care of the linear component. Not de-trending the sectors would increase both persistence and the standard deviation of the sectoral shock innovation from 0.55 to 0.65, and from 0.0501 to 0.0518, respectively. We thank Pete Klenow for this suggestion.

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. 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 0.0163.²⁹

C Conditional Heteroscedasticity in the Aggregate Investment Rate

In this appendix we first present 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.

C.1 Conditional Heteroscedasticity in Aggregate U.S. Investment

Denote the (demeaned) private fixed nonresidential investment to capital ratio series by x_t . We work with the 1960.I–2005.IV period and estimate the following time-series model to quantify the extent to which this series exhibit conditional heteroscedasticity:³⁰

$$x_t = \phi x_{t-1} + \sigma_t e_t, \quad (17a)$$

$$\sigma_t = \alpha_0 + \alpha_1 \bar{x}_{t-1}^k, \quad (17b)$$

where \bar{x}_{t-1}^k denotes the following estimate for “recent investment”:

$$\bar{x}_{t-1}^k = \frac{1}{k} \sum_{j=0}^{k-1} x_{t-1-j},$$

and the e_t are i.i.d. with zero mean and variance equal to one. Of course, the first element of the impulse response of x_t to e_t -shocks is equal to σ_t , thereby illustrating the close connection between conditional heteroscedasticity and time-varying impulse response functions.

²⁹Their persistence is 0.55.

³⁰ α_0 and α_1 must satisfy conditions that ensure that σ_t defined below is positive for all t . For example, a well known result in the ARCH literature is that for $k = 1$ this condition is $\phi^2 + \alpha_1^2 < 1$.

The models with lumpy adjustment developed in this paper (and earlier models such as Caballero and Engel (1999)) predict that fitting the above time-series model to aggregate investment data will lead to a positive and economically significant value for the parameter α_1 . The reason is that 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.

To estimate the parameters of the time-series model (17a)-(17b), we first regress x_t on its lagged value and then regress the absolute residual from this regression, denoted by $|\hat{e}_t|$ in what follows, on \bar{x}_{t-1}^k (both regressions are estimated via OLS).³¹ The second regression then looks like

$$|\hat{e}_t| = \hat{\alpha}_0 + \hat{\alpha}_1 \bar{x}_{t-1}^k + \text{error},$$

and provides estimates for α_0 and α_1 .

Table 14 presents the estimates we obtain for values of k ranging from 1 to 7. The standard deviations are calculated based on 10,000 bootstrap simulations and reported in parentheses. Since the bootstrap simulations generate a significant fraction of large estimates for the parameter that captures conditional heteroscedasticity, α_1 , leading to a standard deviation that overstates the case against α_1 being positive, we also report the fraction of bootstrap simulations that generate a negative estimate for α_1 . This fraction is the p -value for rejecting the one-sided hypothesis $\alpha_1 > 0$. For six out of the seven values of k we consider, this p -value is below the classic 5% threshold, in three cases it is below 2%. More important, the estimated conditional heteroscedasticity schedule (17b) implies economically significant variations in σ_t over time. Thus, for example, the ratio of the largest and smallest values for σ_t when $k = 1$ is 1.49.

Table 14: CONDITIONAL HETEROSCEDASTICITY IN U.S. (1960.I–2005.IV) I/K SERIES

	k						
	1	2	3	4	5	6	7
ϕ :	0.9599 (0.0299)	0.9599 (0.0300)	0.9599 (0.0298)	0.9599 (0.0302)	0.9599 (0.0300)	0.9599 (0.0295)	0.9599 (0.0296)
$10^3 \times \alpha_0$:	0.4828 (0.0351)	0.4849 (0.0353)	0.4823 (0.0355)	0.4822 (0.0363)	0.4805 (0.0386)	0.4820 (0.0390)	0.4837 (0.0397)
α_1 :	0.0209 (0.0172)	0.0239 (0.0173)	0.0282 (0.0172)	0.0300 (0.0174)	0.0330 (0.0179)	0.0336 (0.0179)	0.0331 (0.0182)
p -value ($\alpha_1 > 0$):	0.0507	0.0369	0.0222	0.0179	0.0170	0.0189	0.0210
No. obs. 1st regr.:	183	183	183	183	183	183	183
No. obs. 2nd regr.:	183	182	181	180	179	178	177

³¹To limit the influence of outliers on our estimates, we work with the absolute residual instead of the squared residual. The estimates we obtain are essentially the same if we work with the squared residuals.

The time series model (17a)-(17b) is a particularly simple and robust approach to test for the presence of conditional heteroscedasticity. More sophisticated options can be used as well. For example, instead of assuming the linear relation (17b) we could allow for a more general expression of the form $\sigma_t = h(\bar{x}_{t-1}^k)$, where h is a smooth increasing function. Figure 1 in the introduction plots the estimate we obtain for h (normalized by the average fitted value for σ_t) using $k = 5$,³² a Gaussian kernel and cross-validation to determine the appropriate bandwidth.

C.2 Calibrating the Maintenance Parameter

To choose parameter values that match the heteroscedasticity present in aggregate U.S. investment series, it is useful to summarize the estimated conditional heteroscedasticity schedule (17b) by one statistic. We do this via the signed log-ratio statistic, which we refer to as the *heteroscedasticity range*. We calculate this statistic as follows: The absolute value of this statistic is the log-ratio of the largest and smallest fitted values for the conditional heteroscedasticity regression (17b), when $k = 1$. The sign of the statistic is positive if the estimated slope is positive and negative otherwise. In this way we obtain one statistic that captures, simultaneously, the magnitude and sign of heteroscedastic behavior in lumpy adjustment models.

Table 15 reports estimates for the first moment of the range statistic for the models we consider throughout the paper. For each parameter configuration we generate a time-series of aggregate investment to capital ratios of the same length as the U.S. investment series we work with in the preceding subsection. We then estimate the range statistic for this series (we simulated a sequence of 10,000 observations for each parameter configuration, which leads to 54 series of length 184).

The first row in Table 15 reports the range statistics for the main model discussed in sections 3.2 and 4, that is, the model where we impose an EIS of 1. 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. Columns 3 through 9 consider various values for the maintenance parameter, in each case the simulated model matches perfectly the volatility of sectoral and aggregate investment. It follows from the first row 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 0.4008 for the range statistic, which is close to the estimated value of 0.3971. We therefore choose $\chi = 0.50$ in this case.

The second row reports a similar exercise when the value of the EIS is set to 2 instead of 1. This time $\chi = 0.40$ provides the best fit among the values of the maintenance parameter we

³²Which corresponds to the lowest p -value in Table 14.

consider. Finally, the third row considers the calibration strategy followed in Section 5.1, where the elasticity of intertemporal substitution is among the parameters used to match the selected data moments. In this case the best match of the range statistic is provided by the model with $\chi = 0.25$.

Table 15: HETEROSCEDASTICITY RANGE. CASE: EIS=1

	U.S. I/K	frictionless	χ						
			0	0.10	0.20	0.25	0.30	0.40	0.50
EIS = 1:	0.3971	0.0767	0.2527	0.2722	0.2945	0.3072	0.3215	0.3573	0.4008
EIS = 2:	0.3971	0.0767	0.2745	0.2940	0.3183	0.3330	0.3493	0.3896	0.4409
EIS-estimated:	0.3971	0.0767	0.3203	0.3494	0.3812	0.3998	0.4257	0.4803	0.5532

D 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.

D.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 discretize the state space as follows:

1. k : $n_k = 35$ grid points from $[0, 7.5]$, with a lower grid width at low capital levels, where the curvature of the value function is highest.
2. \bar{k} : $n_{\bar{k}} = 13$ grid points in $[0.8, 1.4]$, equi-spaced.
3. z : $n_z = 11$ grid points in $[0.92, 1.08]$ with closer grid points around unity. For the Gauss-Hermitian integration (see Judd, 1998) we use 7 integration nodes.
4. ϵ : $n_\epsilon = 19$. The grid points are equi-spaced (in logs) and the total grid width is given by $3 \times \sqrt{\frac{\sigma^2}{1-\rho^2}}$, the unconditional variance of the combined technology process. For the transition matrix we use the procedure proposed in Tauchen (1986).

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.

D.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

Table 16: Assessing agents' forecasting rules for capital

	FL	Baseline	Baseline-SKEW	High-EIS	High-EIS-SKEW
$a_{\bar{k}}$	0.0065	0.0021	0.0112	0.0013	0.0150
$b_{\bar{k}}$	0.9061	0.9473	0.9388	0.9457	0.9299
$c_{\bar{k}}$	0.2199	0.1184	0.1106	0.1925	0.1744
$d_{\bar{k}}$	NaN	NaN	0.0075	NaN	0.0164
$e_{\bar{k}}$	NaN	NaN	-0.0008	NaN	-0.0020
R^2	1.0000	0.9999	1.0000	0.9999	1.0000
SE	0.0000	0.003	0.0001	0.0004	0.0002
MAD(%)	0.09	0.61	0.34	0.92	0.37
MSE(%)	0.04	0.30	0.14	0.47	0.14
Correl.	1.0000	0.9956	0.9992	0.9897	0.9992

Table 17: Assessing agents' forecasting rules for p

	FL	Baseline	Baseline-SKEW	High-EIS	High-EIS-SKEW
a_p	1.8438	1.8489	1.8748	0.0930	0.0992
b_p	-0.3357	-0.2442	-0.2701	-0.0350	-0.0423
c_p	-0.5836	-0.8020	-0.8215	-0.0880	-0.0956
d_p	NaN	NaN	0.0213	NaN	0.0074
e_p	NaN	NaN	-0.0031	NaN	-0.0010
R^2	1.000	0.9992	0.9999	0.9944	0.9998
SE	0.0000	0.0006	0.0002	0.0002	0.0000
MAD(%)	0.02	0.19	0.11	0.04	0.01
MSE(%)	0.01	0.06	0.03	0.01	0.00
Correl.	1.000	0.9997	0.9999	0.9980	0.9998

the simulated values and the out-of-sample forecasts. Tables 16 and 17 summarize the numerical 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, and High-EIS-SKEW does the same for our high EIS=20 calibration with a maintenance parameter of $\chi = 0.25$.

Table 16 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, in particular for the high EIS calibration. Fur-

thermore, 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 less good, even though the poorer fit has no bearing on aggregate investment dynamics.

Table 18: 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
High EIS								
Mean:	0.9759	0.9007	0.5374	0.2470	0.8831	0.6421	0.904	0.813
St. deviation:	0.9979	0.9882	0.9705	0.9436	0.9901	0.9731	0.928	0.857
Skewness:	0.9998	0.9988	0.9976	0.9963	0.9991	0.9979	0.923	0.849

We simulated a series of 500 observations for our baseline model and the high EIS calibration, assuming that agents use the first, the first two and the first three moments of capital in their forecasting rules.³³ 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 18 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 and 0.64 for the high EIS 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.³⁴ The last two columns of Table 18 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, and 0.923 and 0.850, respectively, for the high EIS calibration.

We also recomputed the aggregate 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

³³More 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.

³⁴For 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.

the sample paths of I/K generated with forecasting rules with and without higher moments is above 0.9999.³⁵

D.3 Sectoral Simulation

Underlying the sectoral simulation are four assumptions: first, a large enough number of sectors and, secondly, that σ_S is large enough relative to σ_A , 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. Thirdly, we make use of the assumption that a sector is large enough to comprise a large number of production units by invoking a law of large numbers now for the true idiosyncratic productivity. Finally, $\rho_S = \rho_I$, and the independence of sectoral and the idiosyncratic productivity, so that we can treat sectoral and truly 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 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.

We then 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

³⁵Based on this result, when computing the R^2 mentioned in the preceding paragraph, we used the actual series that results when agents use the first three moments of capital as the “true” series. There are no significant differences in Table 18, if we use the actual series that results when agents only use the first, or the first two moments.

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 US (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.

Of course, in order to make the computation viable, we have to scale down the numerical specification of the computation, in particular the grid lengths. The grid length for the additional aggregate shock is 7, equi-spaced, between $[-0.03, 0.03]$ for $N_S = 100$, and $[-0.015, 0.015]$ for $N_S = 400$. We use 5 nodes for both continuous aggregate shocks in the Gauss Hermitian integration.

The following table shows the aggregate and sectoral standard deviations for annual investment 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 B.2.

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

Table 19: 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

0.0163), because now general equilibrium forces contribute also to sectoral smoothing. Overall, our simplified sectoral simulations seem justified.

E Matching Establishment Statistics

One argument we used to justify the use of sectoral rather than plant level data to calibrate micro frictions, is that there are many determinants of plant level moments which are less relevant for the macro dimensions we are concerned with. In this appendix we provide support to this claim by showing that minor modifications of the micro underpinnings of the model we presented in the main text can lead to a satisfactory match of establishment level moments as well. More importantly, the match of sectoral and aggregate moments we obtained in the main text are unaffected by the extension we develop next.

E.1 A Simple Extension

A first choice we need to make when matching the model to micro data is to decide 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 number large enough—sometimes a continuum—so that adding additional units makes no difference.³⁶

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.³⁷ 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.

³⁶Cooper and Haltiwanger (2005) and KT are examples of the former; Abel and Eberly (2002) and Bloom (2007) of the latter.

³⁷For tractability, all models assume that decisions are made at the micro-unit level, not the establishment level.

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). 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 $\log \varepsilon_{uest}^U \sim \text{AR}(1; \rho_U, \sigma_U)$.^{38,39} Consistent with the assumptions we made in the paper, we assume $\rho_S = \rho_E = \rho_U$ 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 will depend on the relative importance of σ_U and σ_E . The larger σ_E , the larger the correlation of productivity shocks across units within an establishment and the more coordinated their investment decisions will be. We consider two extreme scenarios for the degree of correlation of adjustment cost shocks across units within an establishment: perfect correlation and independence.

The sectoral and aggregate investment series generated by this model will be the same as those generated by the model we developed in the main text as long as $\sigma_E^2 + \sigma_U^2 = \sigma_I^2$, since all we are doing in this extension is grouping micro units into groups we call “establishments”, which has no implication for sectoral aggregates. We therefore can decompose σ_I^2 into the sum of σ_U^2 and σ_E^2 as we please, without affecting sectoral and aggregate statistics. We define $\zeta \in [0, 1]$ via $\sigma_U^2 = \zeta \sigma_I^2$, so that $\sigma_E^2 = (1 - \zeta) \sigma_I^2$. Productivity shocks are the same across units within an establishment when $\zeta = 0$, their correlation decreases as ζ increases.

Regarding the sale and purchase of micro units, we assume that in every period an establishment with capital K_{est} suffers a sales/purchase shock τ_{est} , so that its capital becomes $(1 + \tau_{est})K_{est}$. The τ 's are i.i.d. draws from a zero mean symmetric distribution with standard deviation σ_τ . We consider two possibilities, a standard normal distribution and a t -distribution with one degree of freedom. Since the sectors in our model are composed of a continuum of establishments, our choice of a distribution with zero mean for purchase/sales shocks ensures that sectoral and aggregate statistics are unaffected by this extension as well. We choose a symmetric distribution so that asymmetries in the histogram of investment rates cannot be

³⁸ $x_t \sim \text{AR}(1; \rho, \sigma)$ means that the process x_t follows an AR(1) with first order autocorrelation ρ and standard deviation of innovations equal to σ .

³⁹ Sectoral innovations are independent across sectors and independent from the innovations of the aggregate shock. Establishment level innovations are independent across establishments and independent from the innovations of the aggregate and sectoral shocks. Finally, unit level innovations are independent across units and independent from the innovations of the aggregate, sectoral and establishment-level shocks.

attributed to this choice.

We denote by \tilde{i}_{est} the investment rate for a given establishment according to our model, and by i_{est} the corresponding investment rate recorded by the LRD. The latter differs from the former in that it includes the sale/purchase of units from other establishments, which is ignored in our original model. We then have:

$$i_{est} = (1 - \tau_{est})\tilde{i}_{est} - \tau_{est}(1 - \delta). \quad (18)$$

Summing up, our (admittedly simple) extension introduces two parameters over which we can optimize to fit establishment level moments without affecting the match of sectoral and aggregate statistics. These parameters are the degree to which productivity shocks are correlated across units within an establishment, and the average magnitude of sales and purchases of micro units across establishments.

E.2 Matching Establishment Level Statistics

We work with $\chi = 0.5$. For a fixed value of ζ , we generate a histogram with 2,500 realizations of establishment level I/K using our model.⁴⁰

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% , respectively. And denote by $\pi_i(\sigma_\tau)$ the fraction of units with adjustment in the previous bins after applying the transformation described in (18). We choose the value of σ_τ that minimizes $\sum_i |f_i - \pi_i(\sigma_\tau)|/f_i$, that is, we minimize the absolute relative error.

Table 20 presents our results. It also presents the values obtained by KT.⁴¹ We consider all possible combinations ($2 \times 2 \times 2 = 8$) of calibration strategies (the one described in section 3 imposing $EIS=1$ and the one that calibrates the EIS used in section 5.1), correlation among adjustment costs across units within a firm (0 or 1) and the shape of the distribution of sales and acquisition of units (normal or t-distribution). As can be seen, our model does a reasonable job matching the micro statistics which have been considered earlier in the literature. In fact, our fit is within the ballpark of the fit obtained by KT. Also, even though we optimize over $\zeta \in [0, 1]$, the statistics we obtain vary relatively little with ζ (the maximum, over ζ , average absolute deviation exceeds the minimum value by more than 20% for only one of the eight cases considered).

⁴⁰We compute these investment rates using the approximation described in Appendix D.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 .

⁴¹Recall that f_1, f_2, f_3, f_4 and f_5 represent the fraction of establishments with I/K less than -20% , between -20 and -1% , between -1 and 1% , between 1 and 20% and above 20% , respectively.

Table 20: MATCHING LRD MOMENTS: NEW

Calib.	Model		f_1	f_2	f_3	f_4	f_5	Avge. abs. dev.
	corr.	distr.						
3	0	n	.000	.090	.046	.729	.135	.3770
3	0	t	.022	.040	.015	.833	.090	.4823
3	1	n	.000	.090	.079	.726	.105	.3306
3	1	t	.031	.081	.043	.696	.149	.3045
5	0	n	.000	.090	.046	.728	.135	.3776
5	0	t	.022	.040	.015	.833	.090	.4823
5	1	n	.000	.094	.160	.655	.091	.5138
5	1	t	.019	.088	.097	.684	.113	.1451
<i>Data</i>			.019	.090	.082	.622	.187	.0000
KT			.010	.165	.073	.567	.185	.3032

F The Responsiveness Index

Given an economy characterized by a distribution μ_t and aggregate productivity level z_t we denote the resulting aggregate investment rate by $\frac{I}{K}(\mu_t, \log z_t)$ and define

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

where σ_A is the standard deviation of the aggregate innovation.

Following Caballero and Engel (1993b) we define the Responsiveness Index $F(\mu_t, \log z_t)$ for $\frac{I}{K}$ as:

$$F_{k,t} \equiv 0.5(1 - \theta - \nu) [\mathcal{I}^+(\mu_t, \log z_t) + \mathcal{I}^-(\mu_t, \log z_t)]. \quad (19)$$

The factor $(1 - \theta - \nu)$ is included so that the index is approximately one when no sources of smoothing are present. More precisely, in a static, partial equilibrium setting, with no time-to-build, micro units solve:⁴²

$$\max_{k,n} z k^\theta n^\nu - \omega n - k,$$

leading to the following optimal capital target level as a function of aggregate technology:

$$k^* = C z^{1/(1-\theta-\nu)},$$

where C is a constant that depends on the wage and the technology parameters. Taking logs

⁴²For notational simplicity we leave out idiosyncratic and sectoral shocks.

and first differences leads to

$$\Delta \log k^* = \frac{1}{1 - \theta - \nu} \Delta \log z,$$

thereby justifying the normalization constant.