American Economic Journal: Macroeconomics 2013, 5(4): 29–67
http://dx.doi.org/10.1257/mac.5.4.29

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

By Rüdiger Bachmann, Ricardo J. Caballero, and Eduardo M. R. A. Engel*

The sensitivity of US aggregate investment to shocks is procyclical. The response upon impact increases by approximately 50 percent 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* E13, E22, E32)

United States nonresidential, private fixed investment exhibits conditional heteroscedasticity. Figure 1 depicts a smooth, nonparametric, normalized estimate of the squared residual from fitting an autoregressive process to quarterly aggregate investment rates from 1960 to 2005, as a function of the average recent investment rate. This figure suggests that investment is significantly more responsive to shocks in times of high investment.

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

The left and center panels in Figure 2 depict the response over five quarters to a one standard deviation shock taking place at selected points of the US investment cycle, for an ARCH-type time series model and our calibrated lumpy investment DSGE model, respectively.

*Bachmann: RWTH-Aachen University, Templergraben 64, Rm. 513, 52062 Aachen, Germany and CEPR, IFO Institute, and CESifo, (e-mail: ruediger.bachmann@rwth-aachen.de); Caballero: Massachusetts Institute of Technology, Department of Economics, 50 Memorial Drive, Building E52, Room 373A, Cambridge, MA 02142-1347 and National Bureau of Economic Research (e-mail: caball@mit.edu); Engel: University of Chile, Department of Economics, Diagonal Paraguay 257 Piso 14, Santiago, Chile and National Bureau of Economic Research (e-mail: emraengel@gmail.com). We are grateful to Fabian Duarte, David Rappoport, and Jose Tessada for excellent research assistance, and to Olivier Blanchard, William Brainard, Jordi Galí, Pete Klenow, Giuseppe Moscarini, Anthony Smith, Julia Thomas, two anonymous referees, and workshop participants at the American Economic Association annual meeting, University of Bonn, Cornell University, the Econometric Society, Karlsruhe, Mainz, NBER-EFG, New York University, SITE, Stanford University, University of Toulouse, U. de Chile-CEA, and Yale University for helpful comments. Financial support from the National Science Foundation (SES 0550134) is gratefully acknowledged.

[†]Go to http://dx.doi.org/10.1257/mac.5.4.29 to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

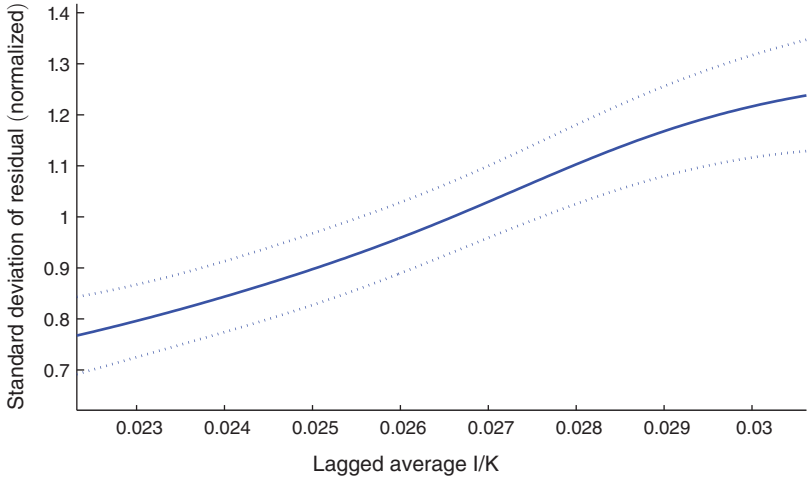

FIGURE 1. CONDITIONAL HETEROSCEDASTICITY OF THE AGGREGATE INVESTMENT RATE

Notes: This figure depicts a smooth, nonparametric, normalized estimate of the squared residual from fitting an autoregressive process to quarterly aggregate investment rates from 1960 to 2005, as a function of the average recent investment rate. We use a Gaussian kernel and determine the bandwidth via cross-validation. Both the autoregressive process and the average of recent investment rates consider six lags. These choices follow from the results we present in Section I. The dotted lines depict one-standard error bands.

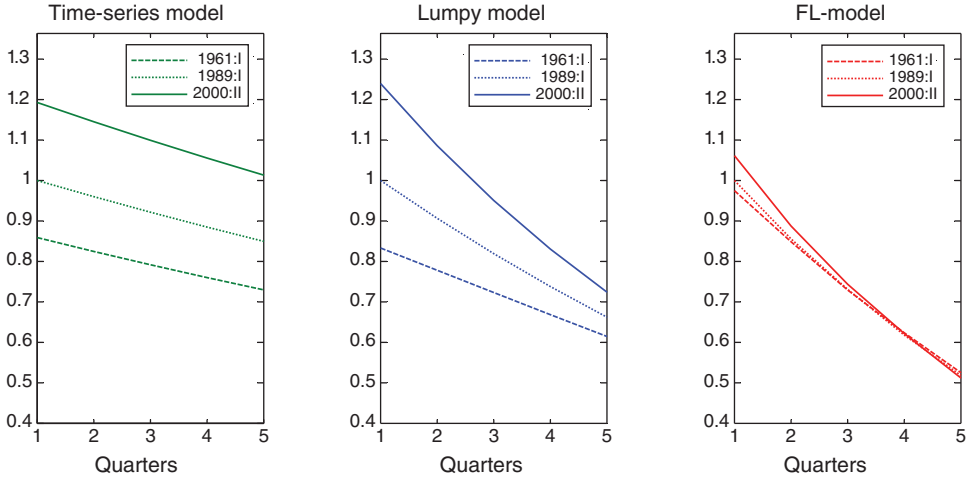

FIGURE 2. IMPULSE RESPONSE IN DIFFERENT YEARS–TIME SERIES, LUMPY AND FRICTIONLESS MODELS

Notes: This figure depicts the response over five quarters to a one standard deviation shock taking place at selected points of the US investment cycle: a trough in 1961, a period of average investment activity in 1989, and a peak in 2000. The figures in the three panels are normalized so that the impulse response in 1989:I (normal investment activity) is one upon impact. It does so for an ARCH-type time series model (left panel), our calibrated lumpy investment DSGE model (center panel), and a frictionless investment model (right panel).

The periods considered are the trough in 1961, a period of average investment activity in 1989, and the peak in 2000. The differences in the impulse response functions are due to differences in the distribution of productivity levels and capital

stocks across productive units. Following a sequence of above average productivity shocks, these units concentrate in a region of the state-space where they are more responsive to any additional shock. The variability of these impulse responses is large and similar in the left and center panels. For example, the immediate response to a shock in the trough in 1961 and to the peak in 2000 differs by roughly 50 percent. 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, Veracierto (2002), Thomas (2002), and Khan and Thomas (2003, 2008) have provided examples in which general equilibrium undoes the partial equilibrium features.

Why do we reach a different conclusion? Because, implicitly, earlier calibrations imposed that the bulk of investment dynamics was determined by general equilibrium price responses 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 price responses, play a relevant role.

Our calibration begins by noting that the objective in any dynamic macroeconomic model is to trace the impact of aggregate shocks on aggregate endogenous variables (investment in our context). The typical response of the endogenous variable is attenuated and spread over time by both microeconomic frictions and aggregate price responses. We refer to this process as *smoothing*, and decompose it into its adjustment cost (AC) and price response (PR) components.

In the context of nonlinear lumpy-adjustment models, AC-smoothing does *not* merely refer to the existence of microeconomic inaction and lumpiness per se, but to their impact in smoothing the response of aggregates. This is a key distinction in this class of models, as in many instances microeconomic inaction translates into limited aggregate inertia (recall the classic Caplin and Spulber 1987 result, where price-setters follow *Ss* rules but the aggregate price level behaves as if there were no microeconomic frictions).

In a nutshell, our key difference with the previous literature is that the latter explored combinations of parameter values that implied microeconomic lumpiness but left almost no role for AC-smoothing, thereby precluding the possibility of fitting facts, such as the conditional heteroscedasticity of aggregate investment rates depicted in Figures 1 and 2.

Table 1 illustrates our model's decomposition into AC- and PR-smoothing. The lower entry shows the volatility of quarterly aggregate investment rates in our model with adjustment costs and price responses. The upper entry reports this statistic when neither smoothing mechanism is present, that is, when adjustment costs are set to zero and prices to their average value in our model with both sources of smoothing. The intermediate entries consider only one source of smoothing at a time, for example, "only AC-smoothing" retains adjustment costs but sets prices

TABLE 1—CONTRIBUTION OF AC AND PR FORCES TO SMOOTHING OF I/K

	No smoothing (0.0425) 0 percent	
Only AC smoothing (0.0040) 81.0 percent		Only PR smoothing (0.0036) 84.6 percent
	AC and PR smoothing (0.0023) 100 percent	

Notes: This table shows the quarterly volatility of the aggregate investment rate from four models. The upper entry refers to the case when adjustment costs are set to zero and prices to their average value in our baseline lumpy investment model. The intermediate entries consider only one source of smoothing at a time, for example, "only AC-smoothing" retains adjustment costs, but sets prices at their average values in the economy that leads to the lower entry. The lower entry refers to our baseline lumpy investment model with adjustment costs and prices that adjust to clear markets (AC and PR smoothing).

at their average values in the economy that leads to the lower entry. The reduction of the standard deviation of the quarterly aggregate investment rate achieved by AC-smoothing alone amounts to 81.0 percent of the reduction achieved by the combination of both smoothing mechanisms. At the other extreme, the additional smoothing achieved by AC-forces, beyond what PR-smoothing achieves by itself, is 15.4 percent of the total, since PR-smoothing can account for 84.6 percent of total smoothing.

It is clear from Table 1 that both sources of smoothing do not enter additively, so some care is needed when quantifying their relative importance. Averaging the upper and lower bounds mentioned above suggests roughly similar roles for both. By contrast, as discussed in detail in Section III, the contribution of AC-smoothing is typically much smaller in the recent DSGE literature.

Our calibration strategy is designed to capture the role of AC-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 price responses have a chance to play a significant smoothing role. Specifically, we argue that the response of semiaggregated (e.g., three-digit) investment to corresponding sectoral shocks is less subject to general equilibrium price responses, and, hence, serves to identify the relative importance of AC-smoothing.

The first row in Table 2 shows the observed volatility of annual sectoral and aggregate investment rates, and their ratio. The second and third rows show these values for our baseline lumpy model and the model with no adjustment costs in investment, respectively. The fourth row reports these statistics for the model in Khan and Thomas (2008). It is apparent from this table that the frictionless model fails to match the sectoral data (it was never designed to do so). In contrast, by reallocating smoothing from the PR- to the AC-component, the lumpy investment model is able to match both aggregate and sectoral volatility.

Table 2—Volatility and Aggregation

Model	3-digit	Aggregate	3-digit/aggregate − ratio
Data	*0.0163*	*0.0098*	*1.66*
This paper	0.0163	0.0098	1.66
Frictionless	0.1839	0.0098	18.77
Khan-Thomas (2008)	0.4401	0.0100	44.01

Notes: This table compares annual sectoral and aggregate investment rates, and their ratio, for the data and three models. Sectoral 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.B and A.C. The volatility of aggregate investment rates in the Khan and Thomas (2008) entry of this table is taken from table III in their paper. The volatility of sectoral investment rates is based on our calculations.

The calibration strategy described above goes a long way toward capturing the heteroscedasticity present in aggregate investment data. It accounts for more than 60 percent of this heteroscedasticity. To match all the heteroscedasticity in aggregate investment data, we introduce maintenance and replacement investment as an essential feature of production units, and assume that some within-period maintenance is necessary to continue operation. Even though there is evidence on the quantitative relevance of maintenance and replacement investment (e.g., McGrattan and Schmitz 1999 and Verick, Letterie, and Pfann 2004), there is a lack of microeconomic studies to help gauge the extent to which these forms of investment are needed to continue operation. We therefore use aggregate statistics, prominent among them a conditional heteroscedasticity measure, to help us identify the maintenance parameter.

The remainder of the paper is organized as follows. In the next section, we provide additional evidence on conditional heteroscedasticity in aggregate investment data. Section II presents our dynamic general equilibrium model. Section III discusses the calibration method in detail. Section IV presents the main macroeconomic implications of the model. Section V concludes and is followed by several appendices.

I. Conditional Heteroscedasticity

In this section, we present time series evidence for conditional heteroscedasticity in aggregate US investment-to-capital ratios. We also explain why we prefer nonlinearity measures based on conditional heteroscedasticity rather than the skewness and kurtosis measures commonly used in the investment literature.

A. *Time Series Models*

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

$$x_t = \sum_{j=1}^{p} \phi_j x_{t-j} + \sigma_t e_t, \tag{1}$$

where $x_t \equiv I_t/K_t$ denotes the investment to capital ratio, the e_t are independently and identically distributed with zero mean and unit variance, and σ_t is a simple function of recent values of x_t as summarized by the following index:[1]

$$\overline{x}^k_{t-1} \equiv \frac{1}{k}\sum_{j=1}^{k} x_{t-j}. \tag{2}$$

For model 1, we stipulate

$$\sigma_t = \alpha_1 + \eta_1 \overline{x}^k_{t-1}, \tag{3}$$

while for model 2, we posit

$$\sigma_t^2 = \alpha_2 + \eta_2 \overline{x}^k_{t-1}. \tag{4}$$

It follows from (1) that the impulse response of x to e upon impact at time t, denoted by $\text{IRF}_{0,t}$, is equal to σ_t. Hence,

$$\text{IRF}_{0,t} = \begin{cases} \alpha_1 + \eta_1 \overline{x}^k_{t-1}, & \text{for model 1;} \\ \sqrt{\alpha_2 + \eta_2 \overline{x}^k_{t-1}}, & \text{for model 2.} \end{cases} \tag{5}$$

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

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

B. *Estimation and Results*

Assume observations for x_t are available for $t = 1, \ldots, T$, and denote by p_{max} and k_{max} the largest values considered for p and k in (1) and (2), respectively. For all pairs

[1] This index can be viewed as a parsimonious and robust approximation to the sequence of innovations up to period $t - 1$.

TABLE 3—EVIDENCE OF HETEROSCEDASTICITY—US INVESTMENT TO CAPITAL RATIO

Series	All		Equipment		Structures	
Model	1	2	1	2	1	2
p^*	6	6	7	7	6	6
k^*	6	6	8	8	2	2
$\eta \times 10^3$	45.93	0.03731	30.62	0.05380	39.95	0.02581
$t-\eta$	3.121	2.496	2.089	1.724	4.097	3.245
p-value$(\eta > 0)$-bootstrap	0.0088	0.0236	0.0375	0.0742	0.0033	0.0094
$\pm \log(\sigma_{max}/\sigma_{min})$	0.7367	0.5933	0.5521	0.4395	1.1167	1.1169
$\pm \log(\sigma_{95}/\sigma_{5})$	0.6118	0.4816	0.4520	0.3547	0.9194	0.8894
$\pm \log(\sigma_{90}/\sigma_{10})$	0.5203	0.4082	0.3355	0.2646	0.8003	0.7403
Skewness	0.1574	0.1574	0.3759	0.3759	−0.1051	−0.1051
Excess kurtosis	−0.9803	−0.9803	−0.1401	−0.1401	−0.9864	−0.9864
First-order autocorrelation e_t	−0.0452	−0.0412	−0.0151	−0.0131	−0.0823	−0.0826
Observations	172	172	172	172	172	172

(p, k) with $p \leq p_{max}$ and $k \leq k_{max}$, we estimate an AR(p) using OLS, and then use the residuals from this regression, denoted ϵ_t, to estimate α and η via OLS from:[2]

$$(6) \qquad \text{Model 1:} \qquad |\varepsilon_t| = \sqrt{\frac{2}{\pi}}\left(\alpha_1 + \eta_1 \overline{x}^k_{t-1}\right) + \text{error},$$

$$\text{Model 2:} \qquad \varepsilon_t^2 = \alpha_2 + \eta_2 \overline{x}^k_{t-1} + \text{error}.$$

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

Table 3 presents the estimates obtained for both models; for US private, fixed, nonresidential investment; and for equipment and structures separately. The frequency is quarterly, from 1960:I to 2005:IV. We use $p_{max} = k_{max} = 12$.

The first and second rows report the optimal values for p and k. The following seven rows report statistics related to the magnitude and significance of the parameter that captures heteroscedasticity and time variation in impulse responses, η. The third row has the point estimate for η, and the fourth row has the corresponding t-statistic, obtained from OLS estimates for (6). The latter may overstate the significance of η, since it ignores variations in the first-stage regressions that determine the autoregressive order, p^*. For this reason, we use 10,000 bootstrap simulations for the investment rate series, starting from our estimates for e_t in (1), to provide an alternative measure of the precision of our estimates for η.[3] The fifth row presents the p-values we obtain for $\eta > 0$ via bootstrap simulations. We

[2] The first equation is based on

$$\mathrm{E}\left[\,|\epsilon_t|\,|\overline{x}^k_{t-1}\right] = \sqrt{\tfrac{2}{\pi}}\left(\alpha_1 + \eta_1 \overline{x}^k_{t-1}\right),$$

while the second equation comes from

$$\mathrm{E}\left[\epsilon_t^2|\overline{x}^k_{t-1}\right] = \alpha_2 + \eta_2 \overline{x}^k_{t-1}.$$

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

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

report one-sided p-values since Ss-type models predict $\eta > 0$. The next four rows present measures for the range of values taken by the estimated impulse response upon impact; σ_{max}, and σ_{min} denote the largest and smallest heteroscedasticity estimates over the sample considered (172 observations), σ_p the p-th percentile. We sign the range estimates by the estimated sign of η. The last row reports the first-order autocorrelation for the estimated innovations (the e_t in (1)), which are consistent with the independently and identically distributed assumption.

Table 3 shows that nonresidential investment exhibits significant (both statistically and economically) heteroscedasticity for both models. This is also the case for structures, and for equipment under Model 1. The range of heteroscedasticity values implied by the estimated models is large. For example, the estimates for model 2 imply that the ninety-fifth percentile is 61.9 percent larger ($e^{0.4816} \simeq 1.619$) than the fifth percentile.

It also follows from Table 3 that nonlinearities in aggregate investment are much larger (and more significant) for structures than for equipment. This is consistent with a more prominent role for lumpy adjustment in the case of structures.[4]

Appendix B provides additional evidence supporting our heteroscedasticity finding for aggregate investment. We apply the methodology described above to the cyclical component of TFP and find no significant heteroscedasticity. This suggests that the heteroscedastic behavior we find in aggregate investment does not come from the underlying shocks and lends credibility to the explanation we explore in this paper, based on lumpy investment behavior. We also show that there is no significant heteroscedasticity in the cyclical component of GDP, while Berger and Vavra (2012) apply the methodology developed here to durables consumption, which is likely to be subject to similar nonconvex adjustment costs as business investment, and find significant heteroscedasticity.

C. *Choosing a Nonlinearity Measure*

In this paper, we introduce a new measure to capture nonlinear relations between shocks and the endogenous aggregate of interest (the investment-to-capital ratio). It is worth comparing this measure with measures that have been used previously in the lumpy investment literature.

Caballero, Engel, and Haltiwanger (1995) and Caballero and Engel (1999) used skewness and kurtosis of the aggregate investment rate to capture nonlinear behavior, as did Thomas (2002) and Khan and Thomas (2003, 2008). This approach is justified as follows: If the model's driving force is Gaussian and the relation between aggregate investment and shocks is linear, the investment rate will also be Gaussian. Finding skewness and kurtosis measures for aggregate investment that differ from those expected under normality can then be interpreted as evidence in favor of a nonlinear relationship.

[4] Similarly, when comparing lumpy investment models with linear time-series models, Caballero and Engel (1999) find a much larger reduction in out-of-sample forecast errors for structures than for equipment.

A problem with this approach is that mapping skewness and kurtosis of aggregate investment into quantities of interest in macroeconomics is far from obvious. Beyond what threshold do departures from normality in the skewness and kurtosis coefficients become relevant from a macroeconomic perspective? By contrast, the reduced-form time-series models introduced in Section IA establish a direct relation between the nonlinearity measure introduced in this paper and the impulse response function for aggregate investment (see equation (5)). Furthermore, a simple function of the parameters of our time-series model measures the time-variation of the impulse response upon impact. This is, we believe, an important advantage when it comes to assessing the macroeconomic relevance of nonlinearities.

A second advantage of the nonlinearity measure we advocate in this paper is that its statistical power is significantly higher than that of the skewness and kurtosis measures. As shown in Appendix B, statistical tests that detect departures from a frictionless RBC-type model using skewness and kurtosis statistics have considerably lower power than tests based on estimates of η (or a function thereof) using the simple time-series models presented in this section.[5] It is not surprising that a statistic especially tailored to capture the specific type of nonlinearity characteristic of *Ss* models does a better job, as reflected in higher statistical power.

II. 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 Khan and Thomas (2008), both in terms of substance and notation. Aside from parameter differences, we have three main departures from Khan and Thomas (2008). First, production units face persistent sector-specific productivity shocks, in addition to aggregate and idiosyncratic shocks. Second, production units undertake some within-period maintenance investment, which is necessary to continue operation; some parts and machines that break down need to be replaced. Third, the distribution of aggregate productivity shocks is continuous rather than a Markov discretization, which allows us to back out the aggregate shocks that are fed into the model to produce Figures 2 and 3.

A. *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 predetermined capital stock (k) and labor (n), according

[5] Caballero, Engel, and Haltiwanger (1995) and Caballero and Engel (1999) did not face the statistical power problem we highlight here because they worked with 20 sectoral investment series instead of one aggregate investment series, as is common in the DSGE literature.

to the following Cobb-Douglas decreasing-returns-to-scale production function ($\theta > 0, \nu > 0, \theta + \nu < 1$):

$$y_t = z_t \epsilon_{S,t} \epsilon_{I,t} k_t^{\theta} n_t^{\nu}, \tag{7}$$

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

We denote the trend growth rate of aggregate productivity by $(1 - \theta)(\gamma - 1)$, so that y and k grow at rate $\gamma - 1$ along the balanced growth path. From now on we work with k and y (and later C) in efficiency units. The detrended aggregate productivity level, which we also denote by z, evolves according to an AR(1) process in logs, with persistence parameter ρ_A and normal innovations with zero mean and variance σ_A^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 $[0, \overline{\xi}]$, 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 predetermined capital stock, the sector it belongs to and the corresponding sectoral productivity level, its idiosyncratic productivity, and its capital adjustment cost. Given the aggregate state, it decides its employment level, n, production occurs, workers are paid, and investment decisions are made. Upon investment the unit incurs a fixed cost of $\omega\xi$, where ω is the current real wage rate. Capital depreciates at a rate δ and a fraction of depreciated capital is replaced to continue operation. Then the period ends.

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

Note that $(i^M)/(k) \equiv \gamma - 1 + \delta$ is the investment rate needed to fully compensate depreciation and trend growth. The degree of necessary maintenance or replacement, χ, can then be conveniently defined as a fraction of $(i^M)/(k)$. If $\chi = 0$, no maintenance investment is needed; if $\chi = 1$, all depreciation and trend growth must be replaced for a production unit to continue operation. We can now summarize the

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

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

evolution of the unit's capital stock (in efficiency units) between two consecutive periods, from k to k', after nonmaintenance investment i and maintenance investment i^M take place, as follows:

	Fixed cost paid	Future capital k'	Total investment $i + i^M$
$i \neq 0$	$\omega\xi$	any $k' > 0$	$\gamma k' - (1 - \delta)k$
$i = 0$	0	$(1 - \chi)\dfrac{(1 - \delta)}{\gamma k} k + \chi k$	$\chi(\gamma - 1 + \delta)k$

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

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

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

Next, we describe the dynamic programming problem of each production unit. We take two shortcuts (details can be found in Khan and Thomas 2008, section 2.4). 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 independently and identically distributed nature of the adjustment costs, continuation values can be expressed without explicitly taking into account future adjustment costs.

We simplify notation by writing maintenance investment as

$$i^M = (\psi - 1)(1 - \delta)k, \tag{8}$$

with $\psi \in \left[1, \frac{\gamma}{1 - \delta}\right]$ defined via

$$\psi = 1 + \left(\frac{\gamma}{1 - \delta} - 1\right)\chi. \tag{9}$$

[8] We note that our version of maintenance investment differs from what Khan and Thomas (2008) call "constrained investment." Here, maintenance refers to the replacement of parts and machines without which production cannot continue, while in Khan and Thomas (2008) it is an extra margin of adjustment for small investment projects.

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

$$V^0(\epsilon_S, \epsilon_I, k; z, \mu) = \int_0^{\bar{\xi}} V^1(\epsilon_S, \epsilon_I, k, \xi; z, \mu)\, G(d\xi). \tag{10}$$

With this notation the dynamic programming problem is given by

$$V^1(\epsilon_S, \epsilon_I, k, \xi; z, \mu) = \max_n \left\{ \text{CF} + \max\left(V_I, \max_{k'} [-AC + V_A] \right) \right\}, \tag{11}$$

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,

$$\text{CF} = \left[z \epsilon_S \epsilon_I k^\theta n^\nu - \omega(z, \mu) n - i^M \right] p(z, \mu), \tag{12A}$$

$$V_I = \beta \text{E}\left[V^0(\epsilon_S', \epsilon_I', \psi(1 - \delta) k/\gamma;\, z', \mu') \right], \tag{12B}$$

$$AC = \xi \omega(z, \mu) p(z, \mu), \tag{12C}$$

$$V_A = -ip(z, \mu) + \beta \text{E}\left[V^0(\epsilon_S', \epsilon_I', k';\, z', \mu') \right], \tag{12D}$$

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 return 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 predetermined, the optimal employment decision is independent of the current adjustment cost draw.

B. *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. Following the argument in Khan and Thomas (2008, section 2.4), we focus on the first-order conditions of the households that determine the equilibrium wage and the intertemporal price.

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

$$U(C, N^h) = \log C - AN^h, \tag{13}$$

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

$$p(z, \mu) \equiv U_C(C, N^h) = \frac{1}{C(z, \mu)}, \tag{14}$$

and from the intratemporal first-order condition:

$$\omega(z, \mu) = -\frac{U_N(C, N^h)}{p(z, \mu)} = \frac{A}{p(z, \mu)}. \tag{15}$$

C. *Recursive Equilibrium*

A *recursive competitive equilibrium* is a set of functions

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

that satisfy

(i) *Production unit optimality*: Taking ω, p, and Γ as given, $V^1(\epsilon_S, \epsilon_I, k; z, \mu)$ solves (10) and the corresponding policy functions are $N(\epsilon_S, \epsilon_I, k; z, \mu)$ and $K(\epsilon_S, \epsilon_I, k, \xi; z, \mu)$.

(ii) *Household optimality*: Taking ω and p as given, the household's consumption and labor supply satisfy (13) and (14).

(iii) *Commodity market clearing*:

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

(iv) *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.

(v) *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 (i), (ii), (iii), and (iv) define an equilibrium given Γ, while step (v) specifies the equilibrium condition for Γ.

D. *Solution*

As is well-known, (11) 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:

$$\log \overline{k}' = a_k + b_k \log \overline{k} + c_k \log z, \tag{16A}$$

$$\log p = a_p + b_p \log \overline{k} + c_p \log z, \tag{16B}$$

where $\overline{k}$ denotes aggregate capital holdings. Given (15), 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 online Appendix D for details).

To implement the computation of sectoral investment rates, we simplify the problem further and impose two additional assumptions: (i) $\rho_S = \rho_I = \rho$ and (ii) enough sectors, so that sectoral shocks have no aggregate effects. Combining both assumptions reduces the state space in the production unit's problem further to a combined technology level $\epsilon \equiv \epsilon_S \epsilon_I$. Now, $\log \epsilon$ follows an AR(1) with first-order autocorrelation ρ and Gaussian innovations $N(0, \sigma^2)$, with $\sigma^2 \equiv \sigma_S^2 + \sigma_I^2$. Since the sectoral technology level has no aggregate consequences by assumption, the production unit cannot use it to extract any more information about the future than it has already from the combined technology level. Finally, it is this combined productivity level that is discretized into a 19-state Markov chain. The second assumption allows us to compute the sectoral problem independently of the aggregate general equilibrium problem.[9]

Combining these assumptions and substituting $\overline{k}$ for μ into (11) and using (16A)–(16B), we have that (12A)–(12D) become

$$\text{CF} = \left[z\epsilon k^{\theta} n^{\nu} - \omega(z, \overline{k}) n - i^M\right] p(z, \overline{k}), \tag{17A}$$

$$V_I = \beta \text{E}\left[V^0(\epsilon', \psi(1 - \delta) k/\gamma; z', \overline{k}')\right], \tag{17B}$$

$$AC = \xi \omega(z, \overline{k}) p(z, \overline{k}), \tag{17C}$$

$$V_A = -ip(z, \overline{k}) + \beta \text{E}\left[V^0(\epsilon', k'; z', \overline{k}')\right]. \tag{17D}$$

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

With the above expressions, (11) becomes a computable dynamic programming problem with policy functions $N = N(\epsilon, k; z, \overline{k})$ and $K = K(\epsilon, k, \xi; z, \overline{k})$. We solve this problem via value function iteration on V^0 and Gauss-Hermitian numerical integration over $\log(z)$ (see online 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:

$$N(\epsilon, k; z, \overline{k}) = \left(\frac{\omega(z, \overline{k})}{\nu z \epsilon k^{\theta}} \right)^{1/(\nu-1)}. \tag{18}$$

Next, we comment on the computation of the production unit's decision rules and value function, given the equilibrium pricing and movement rules (16A)–(16B). From (17D) 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 (17D) with (17B) shows that $V_A(\epsilon; z, \overline{k}) \geq V_I(\epsilon, k; z, \overline{k})$.[10] It follows that there exists an adjustment cost factor that makes a production unit indifferent between adjusting and not adjusting:

$$\hat{\xi}(\epsilon, k; z, \overline{k}) = \frac{V_A(\epsilon; z, \overline{k}) - V_I(\epsilon, k; z, \overline{k})}{\omega(z, \overline{k}) p(z, \overline{k})} \geq 0. \tag{19}$$

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

$$k' = K(\epsilon, k, \xi; z, \overline{k}) = \begin{cases} k^*(\epsilon; z, \overline{k}) & \text{if } \xi \leq \xi^T(\epsilon, k; z, \overline{k}), \\ \psi(1 - \delta) k / \gamma & \text{otherwise.} \end{cases} \tag{20}$$

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

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

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

Now we turn to the second step of the computational procedure that takes the value function $V^0(\epsilon, k; z, \overline{k})$ as given, and prespecifies a randomly drawn sequence of aggregate technology levels: $\{z_t\}$. We start from an arbitrary distribution μ_0, implying a value $\overline{k}_0$. We then recompute (11), using (17A)–(17D), at every point along the

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

sequence $\{z_t\}$, and the implied sequence of aggregate capital levels $\{\overline{k}_t\}$, *without* imposing the equilibrium pricing rule (16B):

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

with V_I defined in (12B) and evaluated at $\overline{k}' = \overline{k}'(k_t)$. This yields new "policy functions":

$$\tilde{N} = \tilde{N}(\epsilon, k; z_t, \overline{k}_t, p)$$

$$\tilde{K} = \tilde{K}(\epsilon, k, \xi; z_t, \overline{k}_t, p).$$

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 $\{\overline{k}_t\}$ endogenously, with which assumed rules (16A)–(16B) 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 online 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 (16A)–(16B). This is confirmed by the fact that adding higher moments of the capital distribution does not increase forecasting performance significantly.

III. Calibration

Our calibration strategy and parameters are standard with two additional features: we combine sectoral and aggregate investment rate volatilities and conditional heteroskedasticity of the aggregate investment rate in order to infer the relative importance of AC- and PR-smoothing as well as the maintenance parameter.

A. *Calibration Strategy*

The model period is a quarter. The following parameters have standard values: $\beta = 0.9942$, $\gamma = 1.004$, $\nu = 0.64$, and $\rho_A = 0.95$. The depreciation rate δ matches the average quarterly investment rate in the data, 0.026, which leads to $\delta = 0.022$. The disutility of work parameter, A, is chosen to generate an employment rate of 0.6.

Next we explain our choices for θ and the parameters of the sectoral and idiosyncratic technology process (ρ_S, σ_S, ρ_I, and σ_I). The output elasticity of capital, θ,

is set to 0.18, in order to capture a revenue elasticity of capital, $\frac{\theta}{1-\nu}$, equal to 0.5, while keeping the labor share at its 0.64-value.[11] We determine σ_S and ρ_S by a standard Solow residual calculation on annual three-digit manufacturing data, taking into account sector-specific trends and time aggregation. This leads to values of 0.0273 for σ_s and 0.8612 for ρ_S.[12] For computational convenience we set $\rho_I = \rho_S$, and σ_I to 0.0472, which leads to an annual standard deviation of the sum of sectoral and idiosyncratic shocks equal to 0.10.[13]

We turn now to the joint calibration of the two key parameters of the model, the adjustment cost parameter, $\overline{\xi}$, and the maintenance parameter, χ, together with the volatility of aggregate productivity shocks.

With the availability of new and more detailed establishment level data, researchers have calibrated adjustment costs by matching establishment level moments (see, e.g., Khan and Thomas 2008). This is a promising strategy in many instances, however, there are two sources of concern in the context of this paper's objectives. First, one must take a stance regarding the number of productive units in the model that correspond to one productive unit in the available micro data. Some authors assume that this correspondence is one-to-one, while others match a large number of model-micro-units to one observed productive unit.[14]

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

[11] In a world with constant returns to scale and imperfect competition this amounts to a markup of approximately 22 percent. The curvature of our production function lies between the values considered by Khan and Thomas (2008) and Gourio and Kashyap (2007). Cooper and Haltiwanger (2006), using LRD manufacturing data, estimate this parameter to be 0.592; Henessy and Whited (2005), using Compustat data, find 0.551.

[12] See Appendix A, Section C for details and online Appendix C for robustness checks. We note that both this paper and Cooper and Haltiwanger (2006) use direct evidence on the driving forces to calibrate the parameters for these processes; in our case, evidence at the sectoral level, in the case of Cooper and Haltiwanger (2006), evidence at the plant level. By contrast, Khan and Thomas (2008) treat parameters of the idiosyncratic driving force as a free parameter to match the plant-level investment rate histogram. This may explain why Khan and Thomas (2008) obtain a smaller value for the volatility of idiosyncratic productivity shocks than other studies as well as why their adjustment costs are much smaller than those obtained in other papers (see Table 4). We thank a referee for pointing out this insight.

[13] In Table 11 in online Appendix C we consider values of 0.075 and 0.15 for the annual total standard deviation, with no significant changes to our baseline calibration.

[14] See Cooper and Haltiwanger (2006) and Khan and Thomas (2008) for an example of the former, and Abel and Eberly (2002) and Bloom (2009), who respectively assume that a continuum and 250 model micro units correspond to one observed plant or firm, for examples of the latter.

We believe that ultimately information about investment rates and shock processes at many levels of aggregation, including the plant or the production unit level, should be brought to bear in order to evaluate richer nonlinear models of investment dynamics. Nevertheless, because of the aforementioned concerns, we follow an approach here where we use three-digit sectoral rather than plant level data to calibrate adjustment costs. More precisely, we choose σ_A, $\overline{\xi}$, and χ jointly to match three statistics in the data: the volatility of the aggregate US investment rate, the volatility of sectoral US investment rates, and the logarithm of the ratio between the ninety-fifth and the fifth percentile of the estimated values for the conditional heteroscedasticity of a simple ARCH model (see Section IA). We refer to this statistic as the *heteroscedasticity range* in what follows.

The novelty in our calibration strategy is that it focuses on matching the relative importance of AC- and PR-smoothing directly. This approach assumes that the sectors we consider are sufficiently disaggregated so that general equilibrium price responses 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 three-digit level.[15]

Given a set of parameters, the sequence of sectoral investment rates is generated as follows. First, the units' optimal policies are determined as described in Section IID, 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.[16] The value of sectoral volatility of annual investment rates we match is 0.0163. To obtain this number we compute the volatilities of the linearly detrended three-digit sectoral investment rates and take a weighted average. As noted in the introduction, this number is one order of magnitude smaller than the one predicted by the frictionless model. To match this annual sectoral volatility in the model simulations, we aggregate over time the quarterly investment rates generated by the model.

As shown in Figure 1 and Section IA, the residuals from estimating an autoregressive process for aggregate US investment exhibit time-varying heteroscedasticity. We use this information as follows. Given a quarterly series of aggregate investment-to-capital ratios, x_t, the moment we match is obtained—both for actual and model-simulated data—by first regressing the series on its lagged value and then regressing the squared residual from this regression, $\hat{e}_t^2$, on x_{t-1}. Denoting by σ_{95} and σ_5 the

[15] Table A3 in Appendix A, Section C provides information on the average number of establishments per two-digit, three-digit, and four-digit sectors, both in absolute terms as well as in relation to the whole US economy. Table 11 in online Appendix C shows that our calibration results do not change significantly if we work with two- or four-digit sectors.

[16] Online Appendix 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. Finally, it should be noted that our calibration abstracts from relative price effects between sectors that may dampen sectoral investment rate volatilities in response to sectoral shocks.

TABLE 4—THE ECONOMIC MAGNITUDE OF ADJUSTMENT COSTS—ANNUAL

Model	Adjustment costs/ unit's output (in percent) (1)	Adjustment costs/ unit's wage bill (in percent) (2)
This paper ($\chi = 0$)	38.9	60.9
This paper ($\chi = 0.25$)	12.7	19.8
This paper ($\chi = 0.50$)	3.6	5.6
Caballero-Engel (1999)	16.5	—
Cooper-Haltiwanger (2006)	22.9	—
Bloom (2009)	35.4	—
Khan-Thomas (2008)	0.5	0.8
Khan-Thomas (2008) "Huge Adj. Costs"	3.7	5.8

Notes: This table displays the average adjustment costs paid, conditional on adjustment, as a fraction of output (left column) and as a fraction of the wage bill (right column), for various models. Rows 4–6 are based on table IV in Bloom (2009). For Cooper and Haltiwanger (2006) and Bloom (2009) we report the sum of costs associated with two sources of lumpy adjustment: fixed adjustment costs and partial irreversibility. The remaining models only have fixed adjustment costs.

ninety-fifth and fifth percentile of the fitted values from the latter regression, the heteroscedasticity range is equal to $\log(\sigma_{95}/\sigma_5)$. The target value for the heteroscedasticity range in the data is 0.3021.[17]

B. *Calibration Results*

The upper bound of the adjustment cost distribution, $\overline{\xi}$, the maintenance parameter, χ, and the volatility of aggregate productivity innovations, σ_A, that jointly match the sectoral and aggregate investment volatilities as well as the conditional heteroscedasticity statistic are $\overline{\xi} = 8.8$, $\chi = 0.50$, and $\sigma_A = 0.0080$.[18] The average cost actually paid is much lower than the average adjustment cost, $\overline{\xi}/2$, as shown in Table 4, since productive units wait for good draws to adjust. The third row shows that, conditional on adjusting, in our calibrated model a production unit pays 3.6 percent of its annual output (column 1) or, equivalently, 5.6 percent of its regular wage bill (column 2).[19] These costs are at the lower end of previous estimates, as shown by comparing them with rows 4–6.

The first two rows in Table 4 report the magnitude of adjustment costs for $\chi = 0$ and $\chi = 0.25$. When calibrating these models, we no longer match the heteroscedasticity range in the data, but continue to match both sectoral and aggregate investment

[17] See Appendix B, Section C for further details on our calibration strategy.

[18] Cooper 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 would suggest a maintenance parameter just below 40 percent. Alternatively, McGrattan and Schmitz (1999) show, for Canadian data, that maintenance and repair expenditures for equipment and structures amounts to roughly 30 percent of expenditures on new equipment and structures. This would suggest just below 25 percent maintenance as a fraction of overall investment. And Verick, Letterie, and Pfann (2004) report that replacement investment in Germany accounts for 66 percent of all investment, which would suggest a value for χ of 0.66. Notice that we do not necessarily need these types of investment to be frictionless, as long as their adjustment costs are much smaller than those for large and lumpy investment projects and are required for the continuing operation of a production unit.

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

Table 5—Heteroscedasticity Range

Model	$\log(\sigma_{95}/\sigma_5)$
Data	*0.3021*
This paper ($\chi = 0$)	0.1830
This paper ($\chi = 0.25$)	0.2173
This paper ($\chi = 0.50$)	0.2901
Quadratic adj. costs ($\chi = 0$)	0.0487
Quadratic adj. costs ($\chi = 0.25$)	0.0411
Quadratic adj. costs ($\chi = 0.50$)	0.0321
Frictionless	0.0539
Khan-Thomas (2008)	0.0468

Notes: This table displays heteroscedasticity range ($\log(\sigma_{95}/\sigma_5)$) for the data (row 1) and various model specifications that vary in terms of the maintenance parameter χ and the adjustment technology for capital: fixed adjustment costs (rows 2–4), quadratic adjustment costs (rows 5–7), a frictionless model, and the Khan-Thomas (2008) model. The adjustment costs for the models in rows 2–7 have been calibrated to match aggregate and sectoral investment rate volatilities.

volatilities. For $\chi = 0.25$, the magnitude of adjustment costs lies slightly below the average of those estimated in the literature, for $\chi = 0$ they are slightly above the maximum value, but still within the ballpark.

Table 4 also shows that, ultimately, the main difference between our calibration and Khan and Thomas (2008) is the size of the adjustment cost. As the next to last row indicates, the average adjustment costs paid by firms in the Khan and Thomas (2008) baseline economy are small, compared to the estimates in the literature and our calibration. This is true even for Khan and Thomas' "huge adjustment costs" calibration (see Khan and Thomas 2008, section 6), namely a 25 fold increase in $\overline{\xi}$. Since the option value of waiting is higher when $\overline{\xi}$ is larger, actual adjustment costs are still only one-tenth of what we obtain for the same maintenance parameter, $\chi = 0$. Conversely, for $\chi = 0$ we find that approximately a 400 fold increase in the value of $\overline{\xi}$ used by Khan and Thomas (2008) in their benchmark model is needed to jointly match the aggregate and sectoral volatility of investment rates.

The first two rows of Table 2 and Table 5 show that our model fits both the sectoral and aggregate volatility of investment, as well as the range of conditional heteroscedasticity in aggregate data. This is not surprising, since our calibration strategy is designed to match these moments. In contrast, the bottom two rows in each of these tables show that neither the frictionless counterpart of our model nor the Khan and Thomas (2008) model match these features.[20]

The first row in Table 5 shows the values obtained directly from the data using our ARCH model. The second and third rows show the range of heteroscedasticity values for versions of our model with values of χ smaller than in the benchmark case. Even though these ranges now are smaller than those in the data, they continue being significantly larger than those implied by a frictionless model. The model with $\chi = 0$ has a heteroscedasticity range three times as large as in the frictionless model, and it accounts for more than 60 percent of the conditional heteroscedasticity

[20] As noted earlier, Khan and Thomas (2008) exhibits slightly lower nonlinearity than our calibration of a frictionless model because of differences in the curvature of the revenue function.

TABLE 6—SMOOTHING DECOMPOSITION

Model	AC smoothing/total smoothing (in percent)		
	LB	UB	Average
Khan-Thomas-lumpy annual	0.0	16.1	8.0
Khan-Thomas-lumpy annual, our $\overline{\xi}$	8.1	59.2	33.7
Our model annual ($\chi = 0$), Khan and Thomas' $\overline{\xi}$	0.8	16.0	8.4
Our model annual ($\chi = 0$)	18.9	75.3	47.0
Our model annual ($\chi = 0.25$)	19.1	75.7	47.4
Our model annual ($\chi = 0.50$)	19.9	76.6	48.3
Our model quarterly ($\chi = 0$)	14.5	80.9	47.7
Our model quarterly ($\chi = 0.25$)	15.4	80.9	48.2
Our model quarterly ($\chi = 0.5$)	15.4	81.0	48.2

found in the data, for the model with $\chi = 0.25$ it is four times as large. Both are much closer to the value in the data than the frictionless model.

Table 5 also shows (in rows 5–7) that convex, quadratic adjustment costs do not generate significant conditional heteroscedasticity for any level of the maintenance parameter, and are, in fact, close to the frictionless model in this respect. This is not surprising, given that quadratic adjustment cost models lead to partial adjustment in the aggregate and thus to essentially linear aggregate investment rate dynamics.

A final way to see the difference between our calibration and Khan and Thomas (2008) is given by a smoothing decomposition, similar to Table 1. Table 6 shows this smoothing decomposition, by reporting upper and lower bounds for the contribution of AC-smoothing to total smoothing, for several models, at different frequencies. The upper and lower bounds for the contribution of AC-smoothing are calculated as follows:

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

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

where σ denotes the standard deviation of aggregate investment rates, NONE refers to the model with fixed prices and with no microeconomic frictions, BOTH to the model with both micro frictions and endogenous price movements, AC to the model that only has microeconomic frictions so that prices are fixed at their average levels of the BOTH specification, and PR to the model with aggregate price responses and no adjustment costs.

The main message can be gathered from the first two rows of this table. By changing the adjustment cost distribution in Khan and Thomas' (2008) model for ours,[21] its ability to generate substantial AC-smoothing rises significantly. Conversely, introducing Khan and Thomas (2008) adjustment costs into an annual version of

[21] Since Khan and Thomas (2008) 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 Khan and Thomas (2008) adjustment costs into our model, we double it.

Table 7—Volatility of Aggregates in Percent

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

Table 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	0.94	0.96

our lumpy model with zero maintenance (third row) leads to a similarly small role of AC-smoothing as in their model. Rows four to nine show the much larger role for AC-smoothing under our calibration strategy, robustly for annual and quarterly calibrations and low and high values of the maintenance parameter.

C. *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 (we label it RBC). As with all models, the volatility of aggregate productivity shocks is chosen to match the volatility of the aggregate investment rate.[22]

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 nonfarm payroll employment from the BLS), the lumpy model improves significantly when matching the volatility of consumption.[23] The lumpy model also slightly increases the persistence of most aggregate variables, bringing these statistics closer to their values in the data.

[22] The value of σ_A required for the frictionless model is $\sigma_A = 0.0051$, the one for the RBC model is 0.0058. This shows that lumpy microeconomic adjustment also dampens conventional second moments in our calibration, thereby providing an additional dimension in which nonconvex adjustment costs have macroeconomic implications. Nonetheless, since the focus of this paper is on aggregate nonlinearities (in the relation between the aggregate investment rate and aggregate productivity shocks), we recalibrate σ_A for each model so as to match aggregate investment volatility. Also note that for the lumpy model, the employment statistics are computed from total employment, that is, including workers that work on adjusting the capital stock. We work with all variables in logs and detrend with an HP-filter using a bandwidth of 1,600.

[23] Consistent 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.

IV. 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.

A. *Understanding Time-Varying Impulse Responses*

In order to understand the different behavior of the frictionless model and the baseline lumpy model ($\chi = 0.50$), we first, following Caballero and Engel (1993b), define the so-called responsiveness index at time t. Given an economy characterized by a joint distribution of capital and productivity μ_t, and an aggregate productivity level z_t, we denote the resulting aggregate investment rate by $(I/K)(\mu_t,\ \log z_t)$ and define the normalized response of this economy to a positive and negative one standard deviation aggregate productivity shock, respectively, as

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

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

where σ_A is the standard deviation of the aggregate innovation. The responsiveness index at time t then is defined as

$$F_t \equiv 0.5\left[\mathcal{I}^{+}(\mu_t,\ \log z_t) - \mathcal{I}^{-}(\mu_t,\ \log z_t)\right]. \tag{21}$$

That is, this index captures the response upon impact of the aggregate investment rate to an aggregate productivity innovation, conditional on the current state of the economy.[24]

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

To generate Figure 3 we back out, for each of the two DSGE models, the aggregate shock that matches the aggregate quarterly investment rate at each point in time over the sample period. Using these shocks, we run the models to compute the responsiveness index. We initialize the process with the economy at its steady state in the fourth quarter of 1959.[25]

[24] Using both $\mathcal{I}^{+}$ and $\mathcal{I}^{-}$ to define F is motivated by the possibility of asymmetric responses to positive and negative shocks. This concern turns out to be unwarranted, as the actual time-series behavior of both series is very similar.

[25] By "steady state" we mean the ergodic (time-average) distribution, which we calculate as follows: starting from an arbitrary capital distribution and the ergodic distribution of the idiosyncratic shocks, we simulate the

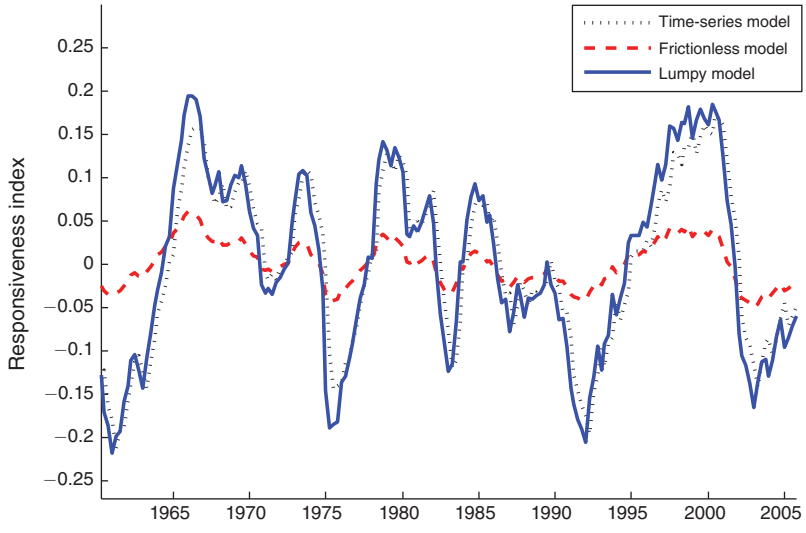

FIGURE 3. TIME PATHS OF THE RESPONSIVENESS INDEX

Notes: This figure plots the evolution of the quarterly responsiveness index for the 1960–2005 period (in log deviations from its average value). The solid and dashed lines represent the index for the lumpy ($\chi = 0.50$) and frictionless models, while the dotted line represents the index for the ARCH-type time series model.

The figure confirms the statement in the introduction, according to which, in the lumpy capital adjustment model, the initial response to an aggregate shock varies significantly more over time than in a frictionless model. The responsiveness index grows by 50.9 percent between trough and peak, and is considerably larger than the 11.6 percent variation implied by the frictionless model.

To understand how lumpy adjustment models generate time-varying impulse responses, two features of the time paths of the responsiveness index are important. Note first that the index fluctuates much less in the frictionless economy than in the lumpy economy. Recall also that the frictionless economy only has general equilibrium price responses to move this index around. From these two observations we can conjecture that the contribution of the price responses 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 AC-smoothing mechanism that is responsible for the rise in the index during the boom phase. When this mechanism is weakened, the responsiveness index in the lumpy economy grows by more than that of the frictionless economy in a boom.

Figure 4 illustrates why the AC-smoothing mechanism weakens as the boom progresses. The figure shows the cross section of mandated investment (and the probability of adjusting, conditional on mandated investment) at two points in the US business cycle: a period of booming aggregate investment, the second quarter of 2000 (dashed line); and a period of depressed aggregate investment in the first

development of an economy with no aggregate innovations for 300 periods, but using the policy functions under the assumption of an economy subject to aggregate shocks.

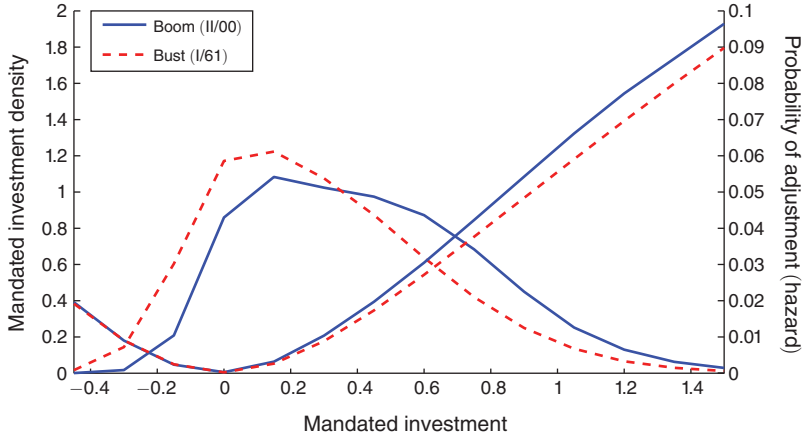

FIGURE 4. INVESTMENT BOOM-BUST EPISODE: CROSS-SECTION AND HAZARD

Notes: This figure shows the cross section of mandated investment (the hump-shaped curves), and the probability of adjusting, conditional on mandated investment (the U-shaped curves), at two points in the US business cycle: a period of booming aggregate investment, the second quarter of 2000 (dashed line), and a period of depressed aggregate investment in the first quarter of 1961 (dotted line). All graphs are computed from the baseline lumpy investment model ($\chi = 0.50$).

quarter of 1961 (dotted line).[26] These cross-sectional densities of mandated investment are computed from our baseline lumpy investment model.

The concept of mandated investment used here extends, to a general equilibrium setting, the analogous partial equilibrium concept that played a central role in Caballero, Engel, and Haltiwanger (1995) and Caballero and Engel (1999). Even though the general equilibrium counterpart does not fully characterize a production unit's state, it continues being useful when describing the mechanism through which lumpy investment models lead to aggregate nonlinearities.

It is apparent from Figure 4 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 percent increases significantly. Also note that the fraction of units in the region where mandated investment is negative decreases during the boom, since the sequence of positive shocks moves units away from this region.

The convex curves in Figure 4 depict the state-dependent adjustment hazard; that is, the probability of adjusting conditional on mandated investment. These adjustment hazards are computed from our baseline lumpy investment model. 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

[26] See Section IID for the formal definition of mandated investment. See Appendix A, Section B, Figure A1 for a time path of the quarterly aggregate investment rate in the United States.

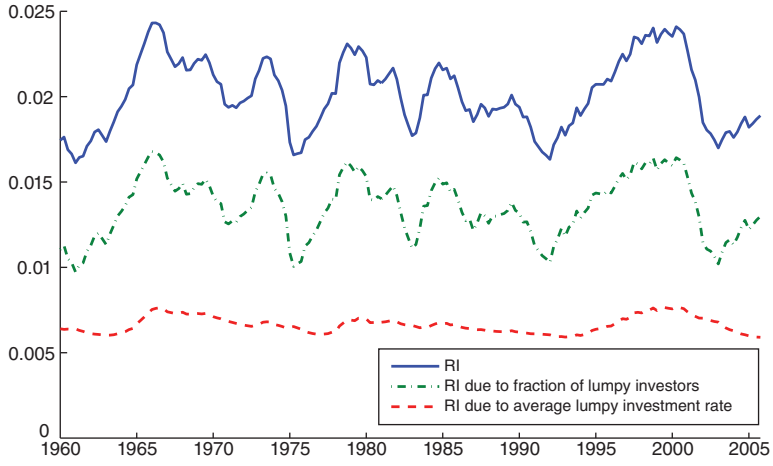

FIGURE 5. DECOMPOSITION OF RESPONSIVENESS INDEX: INTENSIVE AND EXTENSIVE MARGINS

units with larger values of mandated investment. Since units move into the region with a higher slope of the adjustment hazard during the boom, aggregate investment becomes more responsive. This effect is further compounded by the fact that the adjustment hazard shifts upward as the boom proceeds, although this effect is small.

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

The importance of fluctuations in the fraction of adjusters is also apparent in the decomposition of the path of the aggregate investment rate into the contributions from the fluctuation of the fraction of adjusters and the fluctuation of the average size of adjustments, as shown in Figure 6. Both series are in log-deviations from their average values. This is consistent with what Doms and Dunne (1998) documented for establishment-level investment in the United States and Gourio and Kashyap (2007) for the United States and Chile, where the fraction of units undergoing major investment episodes accounts for a much higher share of aggregate (manufacturing in their case) investment than the average size of their investment.[27]

[27] Doms and Dunne (1998) show that the number of plants that have their highest investment in a given year has a correlation with aggregate investment of roughly 60 percent. Gourio and Kashyap (2007, figure 2) show that aggregate investment is mainly driven by investment spikes, and those, to a large degree, are accounted for by the fraction of units undergoing major investment episodes.

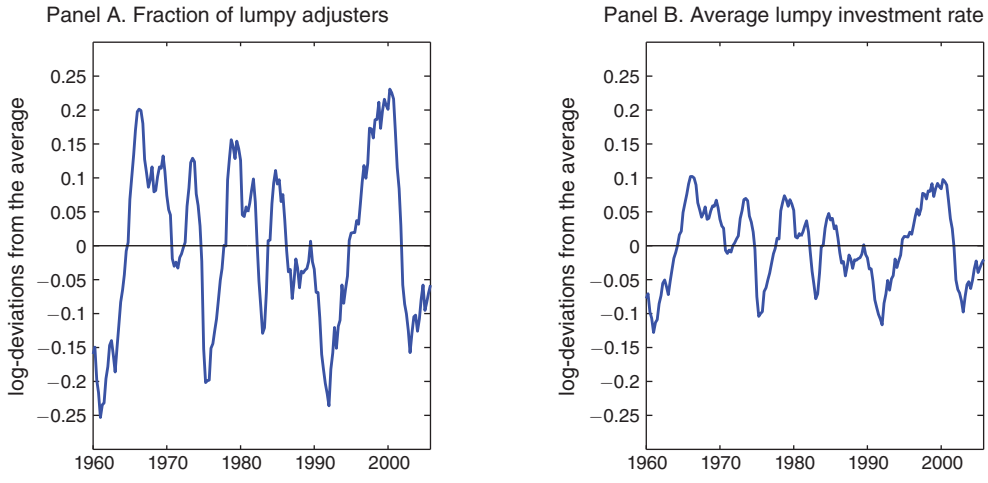

FIGURE 6. DECOMPOSITION OF I/K INTO INTENSIVE AND EXTENSIVE MARGINS

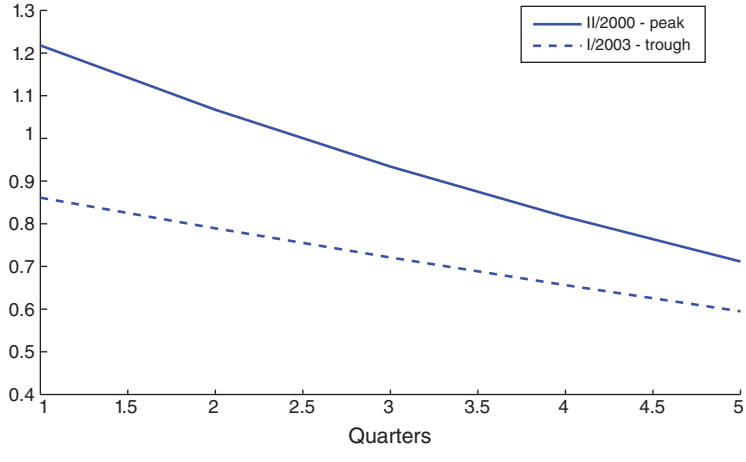

FIGURE 7. IMPULSE RESPONSES OF THE AGGREGATE INVESTMENT RATE IN THE 2000 BOOM-BUST CYCLE

Notes: This figure depicts the responses over five quarters of the baseline lumpy model ($\chi = 0.50$) to a 1 standard deviation shock taking place, respectively, during the second quarter of 2000 (solid line) and the trough in the first quarter of 2003 (dashed line), normalized by the average impulse response upon impact over the entire sample.

B. *The Turn-of-the-Millennium Boom-Bust Cycle—A Case Study*

Next we illustrate the time variation of the investment response during the turn-of-the-millennium boom-bust cycle. Figure 7 depicts the responses over five quarters of the baseline lumpy model to a one standard deviation shock taking place during the peak of this cycle in the second quarter of 2000 and the trough in the first quarter of 2003, normalized by the average impulse response upon impact over the entire sample. The response of investment to a stimulus (e.g., an investment credit) varies systematically over the cycle, being least responsive during a slowdown.

Using a linear model to gauge the effect of a stimulus is likely to overestimate the investment response during a downturn, by approximately 20 percent. This is because the response to a sequence of average shocks, which corresponds to the standard impulse response function calculated for a linear model, is in between both cases and fails to capture the significant time variation of the impulse responses in a world with lumpy investment.

C. *The Role of the Maintenance Parameter*

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

The negative correlation between adjustment costs and the maintenance parameter follows from noting that a higher maintenance parameter lowers the effective drift of mandated investment, defined as depreciation that is not necessarily undone in a given period.[28] Without maintenance, the effective drift is large and dominates over microeconomic uncertainty shocks, resulting in a cross section of mandated investment that is close to the Caplin and Spulber extreme where there is no AC-smoothing (see Caballero and Engel 2007).

Figure 8 shows the average cross-sectional distribution of mandated investment for our baseline model and for the model with $\chi = 0$. The former is clearly farther away from the Caplin-Spulber uniform limit, leaving more space for AC-smoothing.[29] It follows that we need to increase the adjustment cost parameter in order to keep AC-smoothing constant when we reduce the value of the maintenance parameter.

Note, however, that the compensation via an increase in adjustment costs is not enough to preserve the volatility of the impulse response as we drop maintenance (see Table 5). Nonetheless, Figure 9 shows that even with zero maintenance, the responsiveness index of the lumpy economy varies considerably more than in the frictionless economy.

V. Final Remarks

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

Next we argue that it is important to identify the relative contribution of microeconomic adjustment costs and general equilibrium price responses toward the smoothing of the impact of shocks on aggregate variables. In particular, in the case

[28] In the partial equilibrium *Ss* literature an effective drift of zero leads to a triangular ergodic distribution for mandated investment, while a large effective drift, relative to the standard deviation of shocks, leads to an approximately uniform ergodic distribution, as in Caplin and Spulber (1987).

[29] Our result from Table 3 in Section IB, that structures with a smaller effective drift display higher and statistically more significant nonlinearity, is consistent with this mechanism.

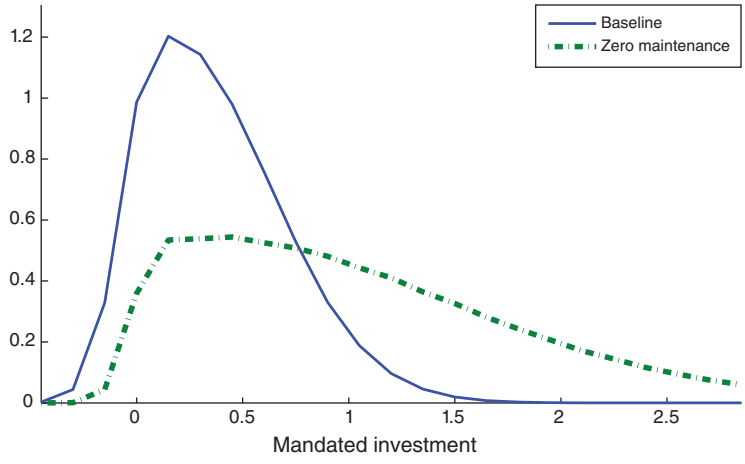

FIGURE 8. ERGODIC CROSS SECTION: ZERO AND BASELINE MAINTENANCE

Notes: See notes to Figure 4. It compares the average cross-sectional distribution of mandated investment for our baseline model ($\chi = 0.50$, solid line) and for the model with $\chi = 0$ (dashed-dotted line).

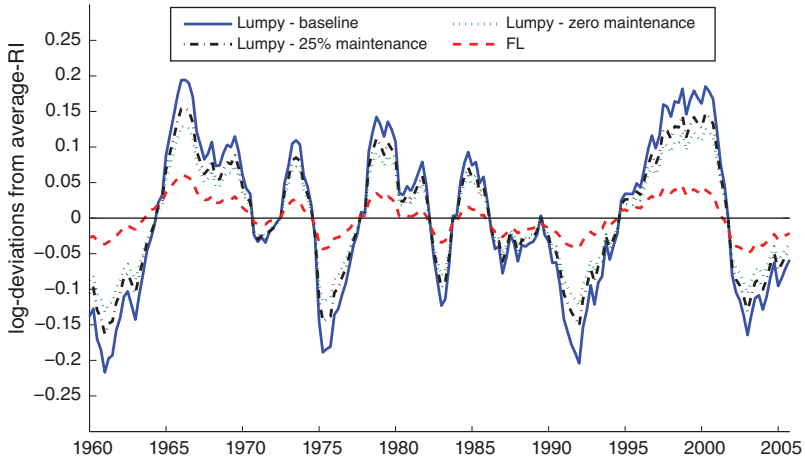

FIGURE 9. TIME PATHS OF THE RESPONSIVENESS INDEX—LOWER MAINTENANCE

Notes: See notes to Figure 3. This figure plots the evolution of the quarterly responsiveness index for the 1960–2005 period (in log deviations from its average value). The solid and dashed lines represent the index for the lumpy ($\chi = 0.50$) and frictionless models, while the dashed-dotted line represents the index for the $\chi = 0.25$ case, and the dotted line represents the index for the $\chi = 0$ case.

of investment models with lumpy capital adjustment, we find that only models that allow for a nontrivial role for adjustment cost smoothing can match the time series evidence on history dependent impulse responses.

More precisely, we find that calibrating a standard lumpy investment model to match the volatilities of aggregate and sectoral investment delivers a parameter configuration, implying a strongly procyclical impulse response function that captures

more than 60 percent of the time variation of this response implied by the time series evidence. The resulting model has adjustment costs that are much larger, and more in line with previous estimates, than those from papers that find no aggregate implications for lumpy investment. We also find that introducing an additional parameter that captures maintenance investment necessary to continue operation leads to a parameter configuration with an impulse response function that accounts for the entire time-variation of this response suggested by an ARCH time-series model.

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

Appendix

In addition to the Appendices A and B here, there are Appendices C, D, and E available online.

Appendix A

A. *Parameters*

Table A1 summarizes the *common* parameters of the models explored in the paper ($\gamma = 1.0040$ for the quarterly calibration and $\gamma = 1.0160$ for the yearly calibration):

Table A1—Common Parameters

Calibration	ρ_A	$\rho_S = \rho_I$	σ_S	σ_I	δ	β	θ	ν
Quarterly	0.9500	0.8612	0.0273	0.0472	0.0220	0.9942	0.1800	0.6400
Yearly	0.8145	0.5500	0.0501	0.0865	0.0880	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: $\sigma_q = (\sigma_y)/(\sqrt{1 + \rho_q + \rho_q^2 + \rho_q^3})$. 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, σ_A, χ, $\overline{\xi}$ and A is explained in Section III. 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 $\overline{\xi}$ to match jointly the standard deviation of the quarterly aggregate investment rate and the standard deviation of the yearly sectoral investment rate, which is aggregated up

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

B. *Aggregate Data*

Since they are not readily available from standard sources, we construct quarterly series of the aggregate investment rate using investment and capital data from the national account and fixed asset tables, available from the Bureau of Economic Analysis (BEA). The time horizon is 1960:I–2005:IV. The quarterly aggregate investment rate in period t is defined as $I_t^{Q,real}/K_{t-1}^{Q,real}$, where the denominator is the real capital stock at the end of period $t-1$ and the numerator is real investment in period t.

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

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

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

$$D_4 = D_3(1+\pi_4) = D_2(1+\pi_3)(1+\pi_4) = D_1(1+\pi_2)(1+\pi_3)(1+\pi_4),$$

$$D^Y = D_1 + D_2 + D_3 + D_4,$$

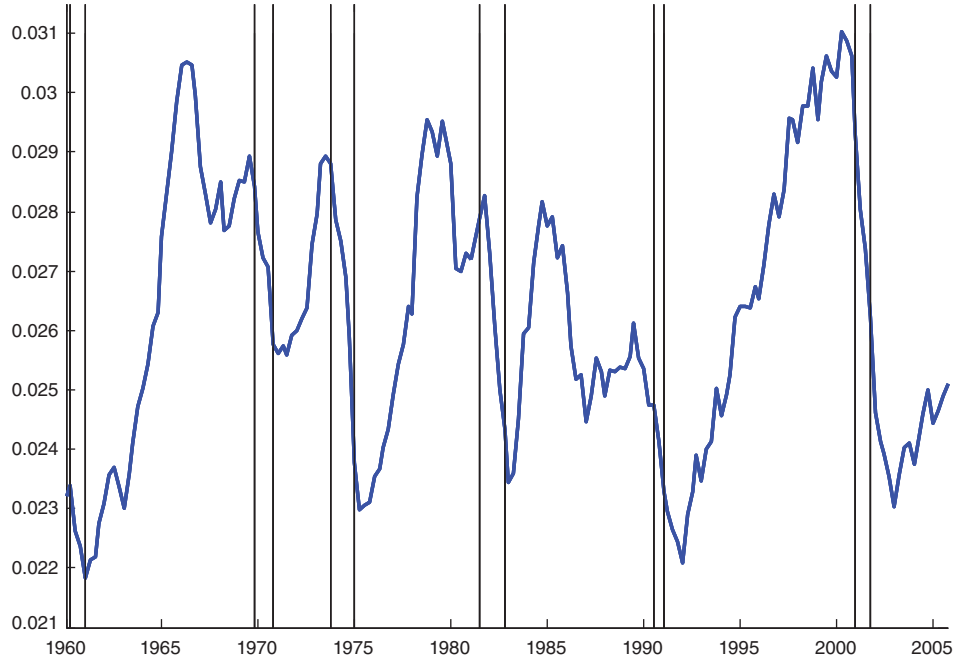

FIGURE A1. THE QUARTERLY US AGGREGATE INVESTMENT RATE

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

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

where all variables are nominal. For fourth quarter capital stocks we use the annual end-of-year data. Year-end prices reported by the BEA are the average of fourth-quarter prices in the current year and first-quarter prices in the following year, thus, nominal end-of-year capital, K_4^Q, for any given year is obtained from $K_4^Q = 2P_4^Q\tilde{K}^Y/(P_4^Q + P_{1'}^Q)$, where $P_{1'}^Q$ corresponds to the nominal price of investment in the first-quarter of next year. Real capital is then calculated as $K_t^{Q,real} = K_t^Q/P_t^Q$.

As Figure A1 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).

Table A2 summarizes statistics of the aggregate investment rate.[30]

TABLE A2—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

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

C. *Sectoral Data*

For sectoral data the best available source is the NBER manufacturing dataset, publicly available on the NBER website. It contains yearly four-digit industry data for the manufacturing sector, according to the SIC-87 classification. We look at the years 1960–1996. We take out industry 3,292, the asbestos products, because this sector essentially dies out in the nineties. This leaves us with 458 four-digit industries altogether.

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

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

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

[31] Aggregating to the two-digit levels leaves us with 20 industries.

[32] Indeed, using a weighted sum of three-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.

deviation of the sectoral shock innovation from 0.55 to 0.65, and from 0.0501 to 0.0518, respectively. The residuals of this trend regression are then taken as the pure sectoral Solow residual series. By construction, they are uncorrelated with the cross-sectional average series. We then estimate an AR(1)-specification for each of these series, and, to come up with a single value for σ_S and ρ_S, set σ_S equal to the value-added-weighted average of the estimated standard deviations of the corresponding innovations, which results in $\sigma_S = 0.0501$ (annual), and ρ_S equal to the value-added-weighted average of the estimated first-order autocorrelation, which leads to $\rho_S = 0.55$ (annual).

Since this computation is subject to substantial measurement error and somewhat arbitrary choices, we perform a number of robustness checks:

- We fix the employment share and capital share to $\nu = 0.64$ and $\theta = 0.18$, as in our model parametrization for all industries;
- Instead of using an OLS projection onto the cross-sectional mean, we simply subtract the latter;
- We look at unweighted means;
- We look at medians instead of means, again weighted and unweighted.

The resulting numbers remain in the ballpark of the parameters we use (see Table 11 in online Appendix C for a robustness analysis with some of these alternative choices).

Calculation of I/K-Moments.—To extract a pure sectoral component of the time series of the industry investment rate, which like the aggregate data includes equipment and structures, we perform the same regressions that were used for TFP-calculation, except that we use a deterministic linear trend to extract sector specific effects. A quadratic detrending of the driving force and a linear detrending for the outcome variable is a conservative approach, as it will make calibrated adjustment costs and aggregate nonlinearities smaller. We do not log or filter the investment rate series. The common component we regress the sectoral investment rate series on is now a capital-weighted average of the industry investment rates. Again, we perform robustness checks with fairly stable results. The resulting standard deviation of sectoral investment rates—our target of calibration—is 0.0163.[33]

Data for Different Digit Levels.—Finally, Table A3 provides information on the number of establishments per sector and the size of each sector within the US economy for the two-digit, three-digit, and four-digit levels.[34] It justifies our choice to use three-digit data in the baseline calibration.

The mean and median (across industries) number of establishments in the three-digit industries are 2,671 and 1,147, respectively. While at the four-digit level industries still contain a fairly large number of establishments on average, the continuum

[33] Their persistence is 0.55.

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

Table A3—Summary Statistics for Manufacturing Establishments

	Mean# est.	Median# est.	Mean frac. est.	Median frac. est.	Max. frac. est.
Two-digit	19,041	14,455	0.28%	0.21%	0.94%
Three-digit	2,671	1,147	0.04%	0.02%	0.51%
Four-digit	780	333	0.01%	0.00%	0.38%

assumption is certainly more justifiable for the three-digit level. Conversely, the across industries average fraction of industry establishments over the number of total establishments in the United States is 0.04 percent at the three-digit level, the median 0.02 percent and the maximum 0.51 percent. In other words, the manufacturing industry with the largest number of establishments has a share of half a percent in the total number of US establishments.[35] The table thus shows that the choice of the three-digit level is a good compromise between our two assumptions: small enough to not have general equilibrium impacts and large enough to justify the assumption of a large number of units. Nevertheless, in online Appendix C we report calibration results also for the two-digit and four-digit levels and show that our results do not hinge on this choice.

Appendix B. Conditional Heteroscedasticity

Here, we elaborate further on the nonlinearity measure introduced in Section I.

A. *GDP and TFP*

The first row in Table B1 shows the *t*-statistic we obtain when applying the methodology described in Section I to the cyclical component of log-GDP. We consider three commonly used filters to detrend GDP: the HP-1600 filter, the Baxter-King's bandpass filter, and a deterministic filter.[36] The second and third rows consider two measures for TFP, a standard Solow-type measure, and the utilization-adjusted Fernald (2012) measure. The following three rows report results for filtered versions of the investment-to-capital ratio.

Various conclusions can be drawn from examining Table B1. First, there is no evidence of heteroscedasticity in the GDP series and (almost) no evidence of heteroscedasticity for TFP (only 1 out of 12 *t*-statistics is significant at the 5 percent level). As mentioned in Section I, this suggests that the evidence of nonlinearities obtained in Table 3 does not come from the shocks but from the transmission mechanism from the shocks to aggregate investment.

We did not detrend the investment-to-capital series when computing the statistics in Table 3 in Section I because, by contrast with the GDP and TFP series, the

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

[36] To allow for changes in trend growth rates, we work with a cubic trend. With the bandpass filter, we isolate frequencies from 6 to 32 quarters, and use a moving average of 12 quarters.

TABLE B1—t-STATISTIC FOR η. US GDP, TFP, AND FILTERED INVESTMENT RATE SERIES

Detrending Model	HP 1	HP 2	BK 1	BK 2	Determ. 1	Determ. 2
GDP	1.3937	1.4729	1.7997	1.0623	−1.0709	−1.2266
TFP-Solow	−0.5558	0.5310	1.8628	2.2088	−0.7714	0.5420
TFP-Fernald	0.9677	1.0346	−1.455	−1.0658	−1.5468	−1.1086
I/K-nonresidential	1.6172	1.5443	1.0452	3.3073	2.9889	2.2564
I/K-equipment	1.6353	1.8223	1.3536	1.3640	2.2832	2.1366
I/K-structures	2.4503	2.9450	2.1058	1.6605	3.4165	2.9534

investment-to-capital ratio series is stationary. We explore the extent to which our heteroscedasticity findings remain valid when we work with a detrended investment series in the last three rows of Table B1. The t-statistics we report for the nonlinearity parameter η show that our results are robust to deterministic detrending and, in the case of structures, to stochastic detrending as well. As we show in the following subsection, the weaker evidence we find for nonlinearities when working with stochastically detrended investment series is due to a loss of statistical power associated with stochastic detrending.

B. *Statistical Power*

Tables B2 and B3 show the statistical power of various tests for detecting nonlinearities in aggregate investment rate series. The null hypothesis, H_0, is that the true model is the frictionless RBC-type model, the alternative model, H_1, is the model calibrated in Section III ($\chi = 0.50$). The statistics considered are skewness and kurtosis of the investment series, as well as the conditional heteroscedasticity parameter η and the fifth-to-ninety-fifth percentile range for conditional heteroscedasticity. Table B2 considers tests of size 0.10 while Table B3 reports the power of tests of size 0.05.

Given a statistic, X, we determine the threshold x_α that defines the size-α test, by solving $\Pr\{X > x_\alpha | H_0\} = \alpha$. The statistical power for the test $X > x_\alpha$, reported in both tables, is calculated as $\Pr\{X > x_\alpha | H_1\}$. We simulated 500 series of length 172 for the frictionless model to calculate the x_α and then used 500 simulated series for our calibrated lumpy model, also each of length 172, to compute the statistical power.

The first row in Tables B2 and B3 show that, when working with the unfiltered investment rate series, the statistical power of kurtosis is extremely low, only slightly above the size of the test. Even though somewhat higher, the statistical power of skewness is considerably lower than that of both tests based on the conditional heteroscedasticity measures introduced in Section I. For example, the probability of rejecting the null when the alternative model is true, with a test of size 0.10, takes values between 0.47 and 0.51 for the tests based on the nonlinearity measures introduced in this paper, while their values are 0.27 for skewness and 0.11 for kurtosis.

The second row shows that the large advantage of the tests considered in this paper is robust to deterministic detrending. Even though somewhat smaller, statistical

TABLE B2—STATISTICAL POWER OF TESTS DETECTING NONLINEARITIES—SIZE OF TEST: 0.10

	Statistic					
			$\pm \log(\sigma_{95}/\sigma_5)$		η	
Detrending method	Skewness	Kurtosis	Model 1	Model 2	Model 1	Model 2
None	0.2667	0.1148	0.4926	0.5111	0.4741	0.5074
Deterministic	0.2259	0.1296	0.4000	0.4074	0.4167	0.4241
HP	0.2759	0.1704	0.2648	0.2852	0.2796	0.2593
BK	0.1944	0.1352	0.2259	0.2111	0.2037	0.1870

TABLE B3—STATISTICAL POWER OF TESTS DETECTING NONLINEARITIES—SIZE OF TEST: 0.05

	Statistic					
			$\pm \log(\sigma_{95}/\sigma_5)$		η	
Detrending method	Skewness	Kurtosis	Model 1	Model 2	Model 1	Model 2
None	0.1278	0.0630	0.3463	0.3407	0.3630	0.3648
Deterministic	0.1593	0.0519	0.2981	0.2778	0.2907	0.2870
HP	0.1963	0.0889	0.1556	0.1704	0.1741	0.1926
BK	0.1333	0.0796	0.1074	0.1259	0.1148	0.1093

power continues being much larger for tests based on the conditional heteroscedasticity of the aggregate investment rate series.[37]

The third and fourth rows in Tables B2 and B3 show that the power advantage of tests based on conditional heteroscedasticity measures essentially disappears, when compared with the skewness statistic, once we work with a stochastic detrended investment-to-capital series, both when using the HP filter, and when using the Baxter-King (BK) filter. A test based on kurtosis continues having lower power.

Summing up, combining the insights from Appendix B, Section A and Appendix B, Section B, we conclude that tests to detect nonlinearities based on the conditional heteroscedasticity measure highlighted in this paper have considerably more statistical power than tests based on statistics used earlier in this literature, such as skewness and kurtosis. The reason for this finding is that these tests are tailored to capture the particular nonlinearity present in *Ss*-type models. Furthermore, since the nonlinearities introduced by nonconvex adjustment costs are unlikely to be limited to business cycle frequencies, this power advantage is significantly reduced when working with stochastically detrended investment series. This justifies using the actual investment-to-capital series, which is stationary, to establish our heteroscedasticity findings in Section I.

C. *Using Conditional Heteroscedasticity in the Calibration*

To choose parameter values that match the heteroscedasticity present in aggregate US investment, it is useful to summarize the estimated conditional heteroscedasticity schedules (3) and (4) by one statistic. We do this via the signed log-ratio of

[37] We assume a cubic trend, as in Table B1. For a linear trend there is no loss of power at all.

Table B4—Heteroscedasticity Statistic and Maintenance Parameter

			χ					
Statistic	US I/K	Frictionless	0	0.10	0.20	0.30	0.40	0.50
$\pm \log(\sigma_{95}/\sigma_5)$	0.3021	0.0539	0.1830	0.1955	0.2095	0.2261	0.2539	0.2901

the ninety-fifth and fifth percentile of the fitted values for σ. Our calibration strategy is akin to the indirect inference approach proposed by Smith (1993), since we match a time-series moment informed by our DSGE model. We work with model 2 and consider $p = 1$, because the shocks in the DSGE model are AR(1), and $k = 1$. This is the time series model in Figures 2 and 3.

Table B4 reports estimates of the heteroscedasticity statistic for the frictionless model and models with various values for the maintenance parameter χ. In each case the simulated model matches the volatility of sectoral and aggregate investment. The first column reports the value for the range statistic in the actual US investment series. The second column reports the value for a model where capital can be adjusted at no cost ("frictionless model"). Columns 3–9 consider various values for the maintenance parameter. For each value of χ we generated a large number of time series of aggregate investment to capital ratios of the same length as the US investment series in our data. We then estimated the range statistic for these series—Table B4 reports the average values.

It follows from the first row of Table B4 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 value of 0.3021 for the range statistic, which is closest to the estimated value of 0.2901. We therefore choose $\chi = 0.50$ for our DSGE model with lumpy capital adjustment (the value for $\chi = 0.60$ would be 0.3207).

REFERENCES

Abel, Andrew, and Janice Eberly. 2002. "Investment and q with Fixed Costs: An Empirical Analysis." http://www.kellogg.northwestern.edu/faculty/eberly/htm/research/invest13.pdf.

Bachmann, Rüdiger, Ricardo J. Caballero, and Eduardo M. R. A. Engel. 2013. "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model: Dataset." *American Economic Journal: Macroeconomics*. http://dx.doi.org/10.1257/mac.5.4.29.

Berger, David, and Joseph Vavra. 2012. "Consumption Dynamics During Recessions." http://econweb.umd.edu/~davis/eventpapers/VavraConsumption.pdf.

Bloom, Nicholas. 2009. "The Impact of Uncertainty Shocks." *Econometrica* 77 (3): 623–85.

Caballero, Ricardo J., and Eduardo M. R. A. Engel. 1993a. "Microeconomic Adjustment Hazards and Aggregate Dynamics." *Quarterly Journal of Economics* 108 (2): 359–83.

Caballero, Ricardo J., and Eduardo M. R. A. Engel. 1993b. "Microeconomic rigidities and aggregate price dynamics." *European Economic Review* 37 (4): 697–711.

Caballero, Ricardo J., and Eduardo M. R. A. Engel. 1999. "Explaining Investment Dynamics in US Manufacturing: A Generalized (S, s) Approach." *Econometrica* 67 (4): 783–826.

Caballero, Ricardo J., and Eduardo M. R. A. Engel. 2007. "Price stickiness in *Ss* models: New interpretations of old results." *Journal of Monetary Economics* 54 (S1): 100–121.

Caballero, Ricardo, Eduardo M. R. A. Engel, and John C. Haltiwanger. 1995. "Plant-Level Adjustment and Aggregate Investment Dynamics." *Brookings Paper on Economic Activity* 25 (2): 1–54.

Caplin, Andrew S., and Daniel F. Spulber. 1987. "Menu Costs and the Neutrality of Money." *Quarterly Journal of Economics* 102 (4): 703–25.

Cooper, Russell W., and John C. Haltiwanger. 2006. "On the Nature of Capital Adjustment Costs." *Review of Economic Studies* 73 (3): 611–33.

Doms, Mark, and Timothy Dunne. 1998. "Capital Adjustment Patterns in Manufacturing Plants." *Review of Economic Dynamics* 1 (2): 409–29.

Fernald, John. 2012. "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity." Federal Reserve Board San Fransisco Working Paper Series 2012–19. http://www.frbsf.org/publications/economics/papers/2012/wp12–19bk.pdf.

Gourio, François, and Anil K Kashyap. 2007. "Investment spikes: New facts and a general equilibrium exploration." *Journal of Monetary Economics* 54 (S1): 1–22.

Henessy, Christopher A., and Toni M. Whited. 2005. "Debt Dynamics." *Journal of Finance* 60 (3): 1129–65.

Judd, Kenneth L. 1998. *Numerical Methods in Economics*. Cambridge, MA: MIT Press.

Khan, Aubik, and Julia K. Thomas. 2003. "Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?" *Journal of Monetary Economics* 50 (2): 331–60.

Khan, Aubik, and Julia K. Thomas. 2008. "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics." *Econometrica* 76 (2): 395–436.

Krusell, Per, and Anthony A. Smith, Jr. 1997. "Income and Wealth Heterogeneity, Portfolio Choice and Equilibrium Asset Returns." *Macroeconomic Dynamics* 1 (2): 387–422.

Krusell, Per, and Anthony A. Smith, Jr. 1998. "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy* 106 (5): 867–96.

McGrattan, Ellen R., and James A. Schmitz, Jr. 1999. "Maintenance and Repair: Too Big to Ignore." *Federal Reserve Bank of Minneapolis Quarterly Review* 23 (4): 2–13.

Smith, A. A., Jr. 1993. "Estimating nonlinear time-series models using simulated vector autoregressions." *Journal of Applied Econometrics* 8 (S1): S63–84.

Tauchen, George. 1986. "Finite state markov-chain approximations to univariate and vector autoregressions." *Economics Letters* 20 (2): 177–81.

Thomas, Julia K. 2002. "Is Lumpy Investment Relevant for the Business Cycle?" *Journal of Political Economy* 110 (3): 508–34.

Veracierto, Marcelo L. 2002. "Plant-Level Irreversible Investment and Equilibrium Business Cycles." *American Economic Review* 92 (1): 181–97.

Verick, Sher, Wilko Letterie, and Gerard Pfann. 2004. "Non-Linearities in the Expansion of Capital Stock." Institute for the Study of Labor (IZA) Discussion Paper 1132.