



WORKING PAPERS

RESEARCH DEPARTMENT

WORKING PAPER NO. 07-24
IDIOSYNCRATIC SHOCKS AND THE ROLE OF
NONCONVEXITIES IN PLANT AND
AGGREGATE INVESTMENT DYNAMICS

Aubhik Khan
Federal Reserve Bank of Philadelphia

Julia K. Thomas
Federal Reserve Bank of Philadelphia and NBER

August 8, 2007

RESEARCH DEPARTMENT, FEDERAL RESERVE BANK OF PHILADELPHIA

Ten Independence Mall, Philadelphia, PA 19106-1574 • www.philadelphiafed.org/econ/index.html

Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics

Aubhik Khan*

Federal Reserve Bank of Philadelphia

Julia K. Thomas

Federal Reserve Bank of Philadelphia and NBER

August 08, 2007

ABSTRACT

We study a model of lumpy investment wherein establishments face persistent shocks to common and plant-specific productivity, and nonconvex adjustment costs lead them to pursue generalized (S,s) investment rules. We allow persistent heterogeneity in both capital and total factor productivity alongside low-level investments exempt from adjustment costs to develop the first model consistent with the cross-sectional distribution of establishment investment rates. Examining the implications of lumpy investment for aggregate dynamics in this setting, we find that they remain substantial when factor supply considerations are ignored, but are quantitatively irrelevant in general equilibrium.

The substantial implications of general equilibrium extend beyond the dynamics of aggregate series. While the presence of idiosyncratic shocks makes the time-averaged distribution of plant-level investment rates largely invariant to market-clearing movements in real wages and interest rates, we show that the dynamics of plants' investments differ sharply in their presence. Thus, model-based estimations of capital adjustment costs involving panel data may be quite sensitive to the assumption about equilibrium. Our analysis also offers new insights about how nonconvex adjustment costs influence investment at the plant. When establishments face idiosyncratic productivity shocks consistent with existing estimates, we find that nonconvex costs do not cause lumpy investments, but act to eliminate them.

*email: mail@aubhik-khan.net, mail@juliathomas.net. We thank Robert King, John Leahy, David Levine, Robert Lucas and Marcelo Veracierto, three anonymous referees, seminar participants at Columbia, Michigan, Pompeu Fabra, Stanford, Wharton and the Federal Reserve Banks of Minneapolis, Philadelphia, Richmond and St. Louis, participants in the April 2006 Penn State Conference and the October 2006 Study Center Gerzensee Conference, as well as session participants at the 2004 Midwest Macro and SED meetings, the 2007 NBER Summer Institute, for useful comments and suggestions. Thomas thanks the Alfred P. Sloan Foundation and the National Science Foundation (grant #0318163) for research support. The views expressed here are those of the authors and do not represent the views of Federal Reserve Bank of Philadelphia or the Federal Reserve System. This paper is available free of charge at www.philadelphiafed.org/econ/wps/index.html

1 Introduction

Over the past fifteen years, an influential body of research has developed the argument that, in order to understand cyclical fluctuations in aggregate investment, we must explicitly examine changes in the underlying distribution of capital across establishments. This growing literature challenges the usefulness of smooth aggregative models for business cycle research, emphasizing that there are important nonlinearities in aggregate investment originating from nonconvexities at the establishment level. In particular, it has been argued that nonconvex costs of adjustment lead establishments to adjust capital infrequently in the form of *lumpy investments* and that occasional synchronization in the timing of establishments' investments can sharply influence the dynamics of the aggregate series. As explained by Caballero and Engel (1999), a large aggregate shock in such a setting may lead to a substantial change in the number of establishments undertaking capital adjustment. This, in turn, implies a time-varying elasticity of aggregate investment demand with respect to shocks. The further claim is that such nonlinearities help explain the data.

The substantial heterogeneity that characterizes (S, s) models of capital adjustment has largely dissuaded researchers in the lumpy investment literature from undertaking general equilibrium analysis.¹ One early exception was the dynamic stochastic general equilibrium model of Khan and Thomas (2003), where the aggregate nonlinearities predicted by previous studies were seen to disappear in general equilibrium. Proponents of lumpy investment remained unconvinced by the finding, however, partly because of important discrepancies in the model's microeconomic implications relative to the data. The distribution of investment arising in our early model differed sharply from that in the data, implying a similar mismatch for the distribution of capital across establishments. But this distribution lies at the heart of the debate.

One important limitation of the first-generation DSGE lumpy analysis was an assumption that differences in capital were the sole source of heterogeneity across plants. Moreover, as is the convention throughout the lumpy literature, there was a stark assumption that

¹Examples of partial equilibrium (S, s) models include Caballero and Engel (1999), Caballero, Engel and Haltiwanger (1995), Cooper, Haltiwanger and Power (1999) and Cooper and Haltiwanger (2006). Veracierto (2002) provides a general equilibrium analysis of plant-level (S, s) policies caused by irreversible investment. Bachmann, Caballero and Engel (2006) study a version of our general equilibrium lumpy investment model, but follow a different calibration strategy, as is discussed in section 6. Gourio and Kashyap (2007) modify the Thomas (2002) environment to include random machine breakdowns and less uncertainty in capital adjustment costs.

nonconvex adjustment costs applied to all capital adjustments irrespective of their size. These two abstractions prevented the theory from usefully addressing a richer set of facts on investment at the establishment-level recently documented by Cooper and Haltiwanger (2006).

Confronting these issues, this paper reconsiders whether plant-level nonconvexities are an essential element lacking in our standard models of the business cycle, this time in a more realistic setting. We take two important steps away from the first-generation model to construct a model that is quantitatively consistent with the available evidence on the cross-sectional distribution of capital adjustment. First, we allow for persistent differences across plants not only in their capital stocks but also in their total factor productivities. Next, we further generalize the model to permit plants to undertake low levels of investment without incurring adjustment costs.

Before exploring aggregate results, we verify that our model satisfies two prerequisites for a useful study. We begin by confirming its consistency with the features of the plant investment distribution that motivate our current work. Next, we confirm that the lumpy investment technology therein has a substantial influence on dynamics that does not evaporate with aggregation. We show that, in partial equilibrium (that is, with real wage and interest rates fixed at their steady-state values), nonconvex adjustment costs at the plant sharply increase skewness and kurtosis in the distribution of aggregate investment rates. Having established our model's consistency with previous predictions of aggregate nonlinearities throughout the partial equilibrium lumpy literature, we trace these features to large changes in the target capital stocks adopted by adjusting plants (changes along the *intensive margin*) in response to aggregate shocks that, in turn, induce large changes in the numbers of plants actively undertaking these adjustments (changes along the *extensive margin*).

Despite the sharp improvement in our current model's ability to reproduce investment patterns in the microeconomic data, our aggregate findings here reinforce those isolated in our previous studies. Microeconomic lumpiness continues to have perceptible effects on aggregate investment dynamics only when equilibrium factor supply considerations are ignored. General equilibrium itself matters tremendously in shaping aggregate dynamics. First, it is extremely effective in smoothing changes in aggregate series, yielding investment rates close to the postwar US data in both their cyclical variability and persistence, irrespective of capital adjustment costs. Second, it dampens much of the movement along the extensive margin that would otherwise distinguish the lumpy investment economy, leaving no trace of aggregate nonlinearities.

Whether nonconvex adjustment costs cause only a fraction of all plants to (unconstrainedly) adjust their capital stocks in any period (in the lumpy model), or whether all

plants adjust their stocks in every period (in a control model without adjustment frictions), households' preferences for consumption smoothing imply equilibrium movements in relative prices that sharply restrain changes in the target capitals to which these plants adjust in response to aggregate disturbances. These dampened changes along the intensive margin immediately imply reduced changes in the value that plants in the lumpy investment economy place on adjustment, thus eliminating large extensive margin changes. Without these large movements in the numbers of adjusting plants, the nonlinearities distinguishing the lumpy investment economy disappear. As a result, our model economy achieves through modest movements along two margins the same aggregate investment dynamics that appear in the frictionless control model through the intensive margin alone. Moreover, we find that the near-zero skewness and excess kurtosis in our general equilibrium models' aggregate investment rate series matches the third and fourth moments of postwar U.S. aggregate investment rates reasonably well, and far more closely than does the partial equilibrium lumpy model.

Our development of a DSGE model consistent with richer aspects of the establishment-level data has led us to some additional new results regarding *microeconomic* investment behavior that may be of independent interest. First, we find that intertemporal changes in the cross-sectional distribution of plant investment rates are sharply dampened in general equilibrium. Although the reallocation of investment goods across plants in response to idiosyncratic shocks is unaffected, we show that the micro-elasticity of response to aggregate shocks is greatly reduced when the associated equilibrium movements in relative prices are permitted to feed back into establishments' decisions. This indicates that equilibrium analysis is essential for understanding investment dynamics even at the most disaggregated level, particularly in times of large aggregate disturbances - those episodes where previous partial equilibrium studies (e.g., Caballero, Engel and Haltiwanger (1995), Caballero and Engel (1999), Cooper and Haltiwanger (2006)) have argued that synchronization in establishment-level investments is critical in shaping aggregate dynamics.

Second, we find that the microeconomic role of nonconvex capital adjustment costs can change substantially in the presence of persistent idiosyncratic risk. When plant-specific productivity shocks are volatile relative to aggregate shocks, as suggested by the data, nonconvex costs no longer *cause* the plant-level investment spikes that are the hallmark of lumpy investment. Rather, their primary role shifts to one of *eliminating* some of these spikes, as is necessary to move the model-implied average cross-sectional distribution in line with the data. These costs now take on a secondary role in reproducing a stark asymmetry in the occurrence of positive versus negative investment. Noting that each of these effects also arises in the presence of investment irreversibilities or combinations of irreversibilities

and convex capital adjustment costs, this finding suggests that there actually may be no distinct role for nonconvex costs even at the microeconomic level. That in turn suggests it may be extremely difficult to correctly disentangle the importance of nonconvex costs, relative to other adjustment frictions, from the establishment-level investment data.

2 Model

In our model economy, there are both fixed costs of capital adjustment and persistent differences in plant-specific productivity, which together lead to substantial heterogeneity in production. In this section, we describe the economy beginning with production units, then follow with households and equilibrium. Next, using a simple implication of equilibrium, we characterize the capital adjustment decisions of production units as a two-sided generalized (S, s) policy. This decision rule for investment is what distinguishes the model from the stochastic neoclassical growth model.

2.1 Production and capital adjustment

We assume a large number of production units. Each establishment produces its output using predetermined capital stock k and labor n , via an increasing and concave production function, F :

$$y = z\varepsilon F(k, n).$$

Here, z reflects stochastic total factor productivity common across plants, while ε is plant-specific productivity. For convenience, we assume that z is a Markov chain, $z \in \{z_1, \dots, z_{N_z}\}$, where

$$\Pr(z' = z_j \mid z = z_i) \equiv \pi_{ij} \geq 0,$$

and $\sum_{j=1}^{N_z} \pi_{ij} = 1$ for each $i = 1, \dots, N_z$. Similarly, we assume that $\varepsilon \in \{\varepsilon_1, \dots, \varepsilon_{N_\varepsilon}\}$, where

$$\Pr(\varepsilon' = \varepsilon_m \mid \varepsilon = \varepsilon_l) \equiv \pi_{lm}^\varepsilon \geq 0,$$

and $\sum_{m=1}^{N_\varepsilon} \pi_{lm}^\varepsilon = 1$ for each $l = 1, \dots, N_\varepsilon$.

In each period, a plant is defined by its predetermined stock of capital, $k \in \mathbf{R}_+$, its idiosyncratic productivity level, $\varepsilon \in \mathcal{E} \equiv \{\varepsilon_1, \dots, \varepsilon_{N_\varepsilon}\}$, and its fixed cost associated with capital adjustment, $\xi \in [0, \bar{\xi}]$, which is denominated in units of labor. Given the current aggregate state of the economy, the plant chooses its current level of employment, production occurs, and its workers are paid. Next, the plant chooses its investment; in doing so, it chooses whether it will pay or avoid its current adjustment cost.

The plant's capital stock evolves according to $\gamma k' = (1 - \delta)k + i$, where i is its current investment and $\delta \in (0, 1)$ is the rate of capital depreciation.² The plant can undertake an *unconstrained* investment only upon payment of its fixed adjustment cost, ξ . Specifically, by forfeiting $\omega\xi$ units of current output, where ω denotes the real wage rate, the plant can invest to any future capital, $k' \in \mathbf{R}_+$. Alternatively, because fixed costs do not apply to adjustments that are sufficiently minor relative to the existing capital stock, the plant can avoid its fixed cost by selecting a *constrained* investment, $i \in [ak, bk]$, where $a \leq 0 \leq b$. In this case, the plant achieves future capital $k' \in \Omega(k) \subseteq \mathbf{R}_+$, where

$$\Omega(k) \equiv \left[\frac{1 - \delta + a}{\gamma}k, \frac{1 - \delta + b}{\gamma}k \right]. \quad (1)$$

For the plant, capital adjustment involves a nonconvexity; conditional on undertaking an unconstrained adjustment, the cost ξ incurred is independent of the scale of adjustment. At the same time, we assume that ξ varies across plants and over time for any given plant. Each period, every plant draws a cost from the time-invariant distribution $G : [0, \bar{\xi}] \rightarrow [0, 1]$. As a result, given its end-of-period stock of capital, a plant's current adjustment cost has no implication for its future adjustment. Thus, it is sufficient to describe differences across plants by their idiosyncratic productivity, ε , and capital, k . We summarize the distribution of plants over (ε, k) using the probability measure μ defined on the Borel algebra, \mathcal{S} , for the product space $\mathbf{S} = \mathcal{E} \times \mathbf{R}_+$. The aggregate state of the economy is then described by (z, μ) , and the distribution of plants evolves over time according to a mapping, Γ , from the current aggregate state; $\mu' = \Gamma(z, \mu)$. We will define this mapping below.

Let $v^1(\varepsilon_l, k, \xi; z_i, \mu)$ represent the expected discounted value of a plant entering the period with (ε_l, k) and drawing an adjustment cost ξ , when the aggregate state of the economy is (z_i, μ) . We state the dynamic optimization problem for the typical plant using a functional equation defined by (2) - (4). First we define the beginning of period expected value of a plant, prior to the realization of its fixed cost draw, but after the determination of $(\varepsilon_l, k; z_i, \mu)$:

$$v^0(\varepsilon_l, k; z_i, \mu) \equiv \int_0^{\bar{\xi}} v^1(\varepsilon_l, k, \xi; z_i, \mu) G(d\xi). \quad (2)$$

Assume that $d_j(z_i, \mu)$ is the discount factor applied by plants to their next-period expected value if aggregate productivity at that time is z_j and current productivity is z_i . (Except where necessary for clarity, we suppress the indices for current aggregate and plant productivity below.) The plant's profit maximization problem, which takes as given the evolution

²Throughout the paper, primes indicate one-period-ahead values, and all variables measured in units of output are deflated by the level of labor-augmenting technological progress, which implies output growth at the rate $\gamma - 1$ along the balanced growth path.

of the plant distribution, $\mu' = \Gamma(z, \mu)$, is then described by

$$v^1(\varepsilon, k, \xi; z, \mu) = \max_{n, k^*, k^C} \left[z\varepsilon F(k, n) - \omega(z, \mu)n + (1 - \delta)k \right. \\ \left. + \max \{ -\xi\omega(z, \mu) + r(\varepsilon, k^*; z, \mu'), r(\varepsilon, k^C; z, \mu') \} \right] \quad (3) \\ \text{s.t. } n \in \mathbf{R}_+, k^* \in \mathbf{R}_+ \text{ and } k^C \in \Omega(k),$$

where $r(\varepsilon, k'; z, \mu')$ represents the continuation value associated with any future capital stock:

$$r(\varepsilon, k'; z, \mu') \equiv -\gamma k' + \sum_{j=1}^{N_z} \pi_{ij} d_j(z, \mu) \sum_{m=1}^{N_e} \pi_{lm}^\varepsilon v^0(\varepsilon_m, k'; z_j, \mu'). \quad (4)$$

Given (ε, k, ξ) and the equilibrium wage rate $\omega(z, \mu)$, the plant chooses current employment n . Next, the plant decides upon either an unconstrained or a constrained choice of its capital stock for next period. The unconstrained choice, in the first term of the maximum operator, requires payment of the fixed labor cost of capital adjustment. However, if $k' \in \Omega(k)$ is selected, the second term applies, and this cost is avoided. Rather than subtracting investment from current profits, we adopt an equivalent but notationally more convenient approach in (3); there, the value of undepreciated capital augments current profits, and the plant is seen to repurchase its capital stock each period. Since adjustment costs do not affect the choice of current employment, we denote the common employment selected by all type (ε, k) plants using $N(\varepsilon, k; z, \mu)$. We let $K(\varepsilon, k, \xi; z, \mu)$ represent the choice of capital for the next period by plants of type (ε, k) with adjustment cost ξ .

2.2 Households

The economy is populated by a unit measure of identical households. Household wealth is held as one-period shares in plants, which we denote using the measure λ .³ Given the prices they receive for their current shares, $\rho_0(\varepsilon, k; z, \mu)$, and the real wage they receive for their labor effort, $\omega(z, \mu)$, households determine their current consumption, c , hours worked, n^h , as well as the numbers of new shares, $\lambda'(\varepsilon', k')$, to purchase at prices $\rho_1(\varepsilon', k'; z, \mu)$. The

³Households also have access to a complete set of state-contingent claims. However, as there is no heterogeneity across households, these assets are in zero net supply in equilibrium. Thus, for brevity, we do not explicitly model them.

lifetime expected utility maximization problem facing each of them is listed below.

$$W(\lambda; z, \mu) = \max_{c, n^h, \lambda'} \left[U(c, 1 - n^h) + \beta \sum_{j=1}^{N_z} \pi_{ij} W(\lambda'; z_j, \mu') \right] \quad (5)$$

subject to

$$c + \int_{\mathbf{S}} \rho_1(\varepsilon', k'; z, \mu) \lambda'(d[\varepsilon' \times k']) \leq \omega(z, \mu) n^h + \int_{\mathbf{S}} \rho_0(\varepsilon, k; z, \mu) \lambda(d[\varepsilon \times k]).$$

Let $C(\lambda; z, \mu)$ describe the household choice of current consumption, $N^h(\lambda; z, \mu)$ the current allocation of time to working, and $\Lambda^h(\varepsilon', k', \lambda; z, \mu)$ the quantity of shares purchased in plants that begin the next period with productivity ε' and k' units of capital.

2.3 Recursive equilibrium

A *recursive competitive equilibrium* is a set of functions,

$$\left(\omega, (d_j)_{j=1}^{N_z}, \rho_0, \rho_1, v^1, N, K, W, C, N^h, \Lambda^h \right),$$

that solve plant and household problems and clear the markets for assets, labor and output:

1. v^1 satisfies (2) - (4), and (N, K) are the associated policy functions for plants.
2. W satisfies (5), and (C, N^h, Λ^h) are the associated policy functions for households.
3. $\Lambda^h(\varepsilon_m, k', \mu; z, \mu) = \mu'(\varepsilon_m, k')$, for each $(\varepsilon_m, k') \in \mathbf{S}$.
4. $N^h(\mu; z, \mu) = \int_{\mathbf{S}} \left[N(\varepsilon, k; z, \mu) + \int_0^{\bar{\xi}} \xi \mathcal{J} \left(\frac{\gamma K(\varepsilon, k, \xi; z, \mu) - (1 - \delta)k}{k} \right) G(d\xi) \right] \mu(d[\varepsilon \times k])$,
where $\mathcal{J}(x) = 0$ if $x \in [a, b]$; $\mathcal{J}(x) = 1$ otherwise.
5. $C(\mu; z, \mu) = \int_{\mathbf{S}} \left[z \varepsilon F(k, N(\varepsilon, k; z, \mu)) - \int_0^{\bar{\xi}} \left(\gamma K(\varepsilon, k, \xi; z, \mu) - (1 - \delta)k \right) G(d\xi) \right] \mu(d[\varepsilon \times k])$.
6. $\mu'(\varepsilon_m, B) = \int_{\{(\varepsilon_l, k, \xi) \mid K(\varepsilon_l, k, \xi; z, \mu) \in B\}} \pi_{lm}^\varepsilon G(d\xi) \mu(d[\varepsilon_l \times k])$, for all $(\varepsilon_m, B) \in \mathcal{S}$, defines Γ .

2.4 (S, s) decision rules

Using C and N to describe the market-clearing values of household consumption and hours worked satisfying conditions 4 and 5 above, it is straightforward to show market-clearing requires that $\omega(z, \mu) = \frac{D_2 U(C, 1 - N)}{D_1 U(C, 1 - N)}$ and that $d_j(z, \mu) = \frac{\beta D_1 U(C'_j, 1 - N'_j)}{D_1 U(C, 1 - N)}$. We may

then compute equilibrium by solving a single Bellman equation that combines the plant-level profit maximization problem with these equilibrium implications of household utility maximization. Defining $p(z, \mu)$ as the price plants use to value current output, we have the following two conditions.

$$p(z, \mu) = D_1 U(C, 1 - N) \quad (6)$$

$$\omega(z, \mu) = \frac{D_2 U(C, 1 - N)}{p(z, \mu)} \quad (7)$$

A reformulation of (2) - (4) then yields an equivalent description of a plant's dynamic problem. Suppressing the arguments of the price functions and exploiting the fact that the choices of n and k' are independent, we have

$$V^1(\varepsilon, k, \xi; z, \mu) = \max_{n \in \mathbf{R}_+} \left(z\varepsilon F(k, n) - \omega n + (1 - \delta)k \right) p \quad (8)$$

$$+ \max \left\{ -\xi \omega p + \max_{k' \in \mathbf{R}_+} R(\varepsilon, k'; z, \mu'), \max_{k' \in \Omega(k)} R(\varepsilon, k'; z, \mu') \right\},$$

where

$$R(\varepsilon, k'; z, \mu') \equiv -\gamma k' p + \beta \sum_{j=1}^{N_z} \pi_{ij} \sum_{l=1}^{N_e} \pi_{lm}^{\varepsilon} V^0(\varepsilon_m, k'; z_j, \mu'), \quad (9)$$

$$V^0(\varepsilon, k; z, \mu) \equiv \int_0^{\bar{\xi}} V^1(\varepsilon, k, \xi; z, \mu) G(d\xi). \quad (10)$$

Equations (8) - (10) will be the basis of our numerical solution of the economy. This solution exploits several results that we now derive. First, note that plants choose labor $n = N(\varepsilon, k; z, \mu)$ to solve $z\varepsilon D_2 F(k, n) = \omega(z, \mu)$. Next, we examine the capital choice of a type (ε, k) plant drawing adjustment cost ξ . Define the value associated with the unconstrained capital choice, $E(\varepsilon, z, \mu)$, and that associated with the constrained choice, $E^C(\varepsilon, k, z, \mu)$, as follow:

$$E(\varepsilon, z, \mu) \equiv \max_{k' \in \mathbf{R}_+} R(\varepsilon, k'; z, \mu') \quad (11)$$

$$E^C(\varepsilon, k, z, \mu) \equiv \max_{k' \in \Omega(k)} R(\varepsilon, k'; z, \mu'). \quad (12)$$

Next, define the plant's *target capital* as the unconstrained choice of k' solving the right-hand side of (11).

Note that the solution to the unconstrained problem in (11) is independent of both k and ξ , but not ε , given persistence in plant-specific productivity. As a result, all plants sharing the same current productivity ε that pay their fixed costs to make unconstrained capital adjustments will choose a common target capital for the next period, $k' = k^*(\varepsilon, z, \mu)$, and

achieve a common gross value of unconstrained adjustment, $E(\varepsilon, z, \mu)$. By contrast, plants that do not pay adjustment costs, instead undertaking constrained capital adjustments solving (12), will choose future capital that may depend on their current capital, $k' = k^C(\varepsilon, k, z, \mu)$. (The exception occurs for plants with $k^*(\varepsilon, z, \mu) \in \Omega(k)$; for such plants, the constraint in (12) does not bind, and the target capital is achieved without incurring an adjustment cost.)

Referring again to the functional equation in (8), it is clear that a plant will absorb its fixed cost to undertake an unconstrained adjustment if the net value of achieving the target capital, $E(\varepsilon, z, \mu) - \xi\omega p$, is at least as great as the continuation value under constrained adjustment, $E^C(\varepsilon, k, z, \mu)$. It follows immediately that a plant of type (ε, k) will undertake unconstrained capital adjustment if its fixed cost, ξ , lies at or below some (ε, k) -specific threshold value. In particular, let $\widehat{\xi}(\varepsilon, k; z, \mu)$ describe the fixed cost that leaves a type (ε, k) plant indifferent between these investment options:

$$-p(z, \mu)\widehat{\xi}(\varepsilon, k; z, \mu)\omega(z, \mu) + E(\varepsilon, z, \mu) = E^C(\varepsilon, k, z, \mu). \quad (13)$$

Next, define $\xi^T(\varepsilon, k; z, \mu) \equiv \min\{\bar{\xi}, \widehat{\xi}(\varepsilon, k; z, \mu)\}$, so that $0 \leq \xi^T(\varepsilon, k; z, \mu) \leq \bar{\xi}$. Any plant with an adjustment cost at or below its type-specific threshold, $\xi^T(\varepsilon, k; z, \mu)$, will pay the fixed cost and adjust to its target capital.

Using the target and constrained capital choices identified above, alongside the threshold adjustment costs, the plant-level decision rule for capital may be conveniently summarized as follows. Any establishment identified by the plant-level state vector $(\varepsilon, k, \xi; z, \mu)$ will begin the subsequent period with a capital stock given by

$$k' = K(\varepsilon, k, \xi; z, \mu) = \begin{cases} k^*(\varepsilon, z, \mu) & \text{if } \xi \leq \xi^T(\varepsilon, k; z, \mu) \\ k^C(\varepsilon, k, z, \mu) & \text{if } \xi > \xi^T(\varepsilon, k; z, \mu). \end{cases} \quad (14)$$

Thus, within each group of plants sharing a common (ε, k) , fraction $G(\xi^T(\varepsilon, k; z, \mu))$ pay their labor-denominated fixed costs to undertake an unconstrained capital adjustment. It then follows that the market-clearing levels of consumption and work hours required to determine p and ω using (6) and (7) are given by:

$$C = \int_{\mathbf{S}} \left(z\varepsilon F(k, N(\varepsilon, k; z, \mu)) - G(\xi^T(\varepsilon, k; z, \mu)) \left[\gamma k^*(\varepsilon, z, \mu) - (1 - \delta)k \right] - \left[1 - G(\xi^T(\varepsilon, k; z, \mu)) \right] \left[\gamma k^C(\varepsilon, k, z, \mu) - (1 - \delta)k \right] \right) \mu(d[\varepsilon \times k]), \quad (15)$$

$$N = \int_{\mathbf{S}} \left[N(\varepsilon, k; z, \mu) + \int_0^{\xi^T(\varepsilon, k; z, \mu)} \xi G(d\xi) \right] \mu(d[\varepsilon \times k]). \quad (16)$$

Finally, based on (14), we can now describe the evolution of the plant distribution,

$\mu' = \Gamma(z, \mu)$. First, define the indicator function $\overline{\mathcal{J}}(x) = 1$ for $x = 0$; $\overline{\mathcal{J}}(x) = 0$ for $x \neq 0$. Informally, for each $(\varepsilon_m, \widehat{k}) \in \mathcal{S}$,

$$\begin{aligned} \mu'(\varepsilon_m, \widehat{k}) = & \sum_{l=1}^{N_\varepsilon} \pi_{lm}^\varepsilon \left[\overline{\mathcal{J}}(\widehat{k} - k^*(\varepsilon_l, z, \mu)) \int G(\xi^T(\varepsilon_l, k; z, \mu)) \mu(\varepsilon_l, dk) \right. \\ & \left. + \int [1 - G(\xi^T(\varepsilon_l, k; z, \mu))] \overline{\mathcal{J}}(\widehat{k} - k^C(\varepsilon_l, k, z, \mu)) \mu(\varepsilon_l, dk) \right]. \end{aligned} \quad (17)$$

Consider the cases of $\widehat{k} = k^*(\varepsilon_l, z, \mu)$, for each given ε_l , $l = 1, \dots, N_\varepsilon$. The first line of equation (17) represents those plants (ε_l, k) that pay their fixed costs to adjust to this target. However, our law of motion must also reflect those plants that reach $\widehat{k} = k^*(\varepsilon_l, z, \mu)$ without paying fixed costs. For such plants, $k^C(\varepsilon_l, k, z, \mu) = k^*(\varepsilon_l, z, \mu)$, so $\xi^T(\varepsilon_l, k; z, \mu) = 0$. Thus, they are a subset of the plants avoiding fixed costs in the second line of (17), those with current capital such that $k^*(\varepsilon_l, z, \mu) \in \Omega(k)$. Next, consider the cases of $\widehat{k} \neq k^*(\varepsilon_l, z, \mu)$. Those plants reflected in the second line for which $k^*(\varepsilon_l, z, \mu) \notin \Omega(k)$ are plants that face either a binding upper constraint on their capital choice (with $k < \frac{\gamma}{1-\delta+b} k^*(\varepsilon_l, z, \mu)$) or a binding lower constraint (with $k > \frac{\gamma}{1-\delta+a} k^*(\varepsilon_l, z, \mu)$). Of this group, those with $k^C(\varepsilon_l, k, z, \mu) = \widehat{k}$ begin the next period with \widehat{k} .

3 Calibration

We evaluate the plant-level and aggregate implications of nonconvex capital adjustment costs using several numerical experiments across which we vary the stochastic process for idiosyncratic shocks to plants' total factor productivity and the parameterization of capital adjustment costs. All other production parameters, as well as preferences, are held constant throughout. Each experiment is based on a 10,000-period model simulation, and the same random draw of aggregate productivity is used in each. Below, we discuss functional forms and parameter values for technology and preferences that are identical across models. Thereafter, in section 3.2, we explain the choice of idiosyncratic shocks and the distribution of capital adjustment costs. The description of our numerical method is provided in Appendix A.

3.1 Common parameters

Across our model economies, we assume that the representative household's period utility is the result of indivisible labor (Hansen (1985), Rogerson (1988)): $u(c, L) = \log c + \varphi L$, and the establishment-level production function takes a Cobb-Douglas form, $z\varepsilon F(k, N) =$

$z\varepsilon k^\alpha N^\nu$. We fix the length of a period to correspond to one year, allowing us to use evidence on establishment-level investment in selecting parameters governing the distributions of adjustment costs and idiosyncratic productivities below.

Model parameters, other than those involving idiosyncratic shocks and adjustment costs, are selected to ensure agreement with observed long-run values for key postwar U.S. aggregates in a nested *frictionless* version of our model without capital adjustment costs described in Appendix B. As proven in lemma 2 of this appendix, macroeconomic aggregates are insensitive to the presence of idiosyncratic productivity differences in the absence of capital adjustment costs. This allows us to choose parameter values for technology and preferences that are consistent with empirical counterparts before specifying an idiosyncratic shock process. For these parameters, we apply the same values to the lumpy investment model. We are able to use this approach because the aggregate first moments across our model economies are extremely similar.

The mean growth rate of technological progress is chosen to imply a 1.6 percent average annual growth rate of real per capita output, and the discount factor, β , is then set to imply an average real interest rate of 4 percent. Given the rate of technological progress, the depreciation rate, δ , is selected to match an average investment-to-capital ratio of 10 percent, corresponding to the average value for the private capital stock between 1954 and 2002 in the U.S. Fixed Asset Tables. Labor’s share is then set to 0.64 as in Prescott (1986); given this value, capital’s share of output is determined by targeting an average capital-to-output ratio of 2.353 as in the data. Next, the parameter governing the preference for leisure, φ , is taken to imply an average of one-third of available time spent in market work. Table 1 lists the resulting parameter values.

In specifying our exogenous stochastic process for aggregate productivity, we begin by assuming a continuous shock following a mean zero AR(1) process in logs: $\log z' = \rho_z \log z + \eta'_z$ with $\eta'_z \sim N(0, \sigma_{\eta_z}^2)$. Next, we estimate the values of ρ_z and σ_{η_z} from Solow residuals measured using NIPA data on US real GDP and private capital, together with the total employment hours series constructed by Prescott, Ueberfeldt, and Cociuba (2005) from CPS household survey data, over the years 1959-2002. Finally, we discretize the resulting productivity process using a grid with 11 shock realizations; $N_z = 11$.

3.2 Plant-specific shocks and adjustment costs

The remaining parameters involve the distribution of plant-specific productivity and the adjustment costs facing plants in the lumpy investment economy. We determine idiosyncratic shocks $(\varepsilon_i)_{i=1}^{N_\varepsilon}$ and the Markov Chain determining their evolution $(\pi_{ij}^\varepsilon)_{i,j=1}^{N_\varepsilon}$ by discretizing a log-normal process, $\log \varepsilon' = \rho_\varepsilon \log \varepsilon + \eta'_\varepsilon$ using 15 values ($N_\varepsilon = 15$). To main-

tain controlled comparisons, the same stochastic process is applied to both the frictionless and the lumpy investment models. In the latter, fixed costs of investment are assumed to be drawn from a uniform distribution, $G(\xi) = \xi/\bar{\xi}$, and the range of investment rates that do not incur such costs is assumed to be symmetric around 0; in other words, $|a| = b$.

There is little agreement about the persistence of the idiosyncratic shock process, ρ_ε . (Compare, for example, the values in Comin and Phillipon (2005) to those of Cooper and Haltiwanger (2006).) Given this, we simply set it equal to the persistence of the aggregate shock, $\rho_\varepsilon = \rho_z$.⁴ Next, the remaining plant-level parameters ($\sigma_{\eta_\varepsilon}$, $\bar{\xi}$, b) are selected to best match the empirical average distribution of plant investment rates, as summarized by Cooper and Haltiwanger (2006).

Constructing their own plant capital series using data on retirements and investment from the Longitudinal Research Database, Cooper and Haltiwanger provide a detailed set of time-averaged moments on plants' investment rates, which are reproduced in the shaded row of Table 2. They define any plant with an investment rate (ratio of investment to capital) less than 1 percent in absolute value as *inactive*. Positive investment rates are those at or exceeding 1 percent, while negative investment rates are those falling at or below -0.01 . Finally, positive *spikes* are investment rates exceeding 0.2, and negative spikes are observations of $\frac{i}{k} < -0.2$.

Several features of the time-averaged plant data are prerequisites for our study. First, investment inactivity is relatively rare, occurring among only 8 percent of plants on average. Next, there is a sharp asymmetry in positive versus negative investment rates; in the average year, roughly 82 percent of plants actively raise their capital stocks. Finally, the columns summarizing observations of investment spikes indicate not only extreme investment episodes occurring among a nontrivial fraction of establishments (roughly 20 percent) in the tails of the average plant distribution, but also right skewness. Here again we see a sharp asymmetry; positive spikes are observed 10 times as often as negative spikes.

Before proceeding, we discuss our reasons for assuming a region of capital adjustment that is exempt from adjustment costs. Throughout the lumpy investment literature thus far, it has been assumed that all active adjustments to a plant's capital stock incur fixed costs. Given that assumption, we show in the next section that the inclusion of idiosyncratic productivity shocks is not sufficient to yield consistency with the average distribution of investment rates in the plant-level data. Specifically, the traditional lumpy investment model matches the average occurrence of investment spikes only by substantially exaggerating the frequency of inaction. One possible explanation for this tension in reconciling the theory

⁴In a previous version of this paper, we instead selected a much lower persistence, $\rho_\varepsilon = 0.53$, taken from Cooper and Haltiwanger (2002). Our findings here are entirely unaffected by the change.

with microeconomic data is that, in reality, fixed adjustment costs apply only to those investments that are comparatively large relative to a plant’s existing capital. Alternatively, it may be reasonable to suppose that the fixed costs associated with relatively large capital adjustments, such as building a new structure, are substantially greater than those associated with minor ones, such as installing a new computer. We adopt a rough proxy for these distinctions by permitting some low-level capital adjustments that are exempt from fixed costs. This generalization allows our model to overcome the tension noted above, making it the first to succeed in matching the available moments from the cross-sectional distribution of plant investment rates.⁵

4 Prerequisites

4.1 Consistency with microeconomic data

A prerequisite for our current study is that our environment reproduce the key aspects of the microeconomic data described above. In Table 2, we evaluate the microeconomic performance of our model, comparing it to that of the traditional lumpy investment model previously studied by Khan and Thomas (2003).⁶ There, all non-zero investment rates were subject to fixed adjustment costs ($b = 0$), and there were no plant-specific productivity disturbances ($N_\varepsilon = 1$, $\sigma_{\eta_\varepsilon} = 0$). Row 1 presents the results for this special case of our current model when the upper support of the adjustment cost distribution alone is selected to best match the LRD data.

The traditional lumpy model reproduces only one aspect of the micro data, the frequency of positive investment spikes. There, some plants repeatedly draw relatively high fixed costs, and hence forego capital adjustment, for several consecutive periods. We observe positive spikes when such plants finally take action, because their effective capital stocks lie far below the target to which they invest, a result of ongoing depreciation and technological progress. The trade-off in reproducing observations of investment spikes versus inactivity is a common difficulty among quantitative models of lumpy investment (see, for example,

⁵Cooper and Haltiwanger (2006) estimate adjustment costs using moments from the establishment data that is summarized in Table 2. They do not attempt to simultaneously reproduce the frequency of both investment spikes and inaction. Bayer (2006) estimates investment functions using firm-level data. Bloom (2007) estimates both capital and labor adjustment cost parameters, again using firm-level data, and evaluates the fit of his model against an alternative set of moments not directly addressing inaction or spikes.

⁶These moments from the cross-sectional distribution in each model’s steady state match closely with corresponding time-averages taken over long general equilibrium model simulations. Partial equilibrium simulations yield similar results when plants’ individual investment decisions are more influenced by idiosyncratic relative to aggregate disturbances.

Cooper and Haltiwanger (2006) or Gourio and Kashyap (2007)). Though not shown in the table, plant-specific productivity shocks alone do not solve the problem. With their introduction, inactivity continues to exceed 75 percent when observations of positive spikes match those in the data.

The sharp disparities between the moments summarizing actual plant-level investment patterns and those in the traditional lumpy model of our previous work have motivated the extensions we have undertaken here. When plants face idiosyncratic productivity shocks, those shifting from high productivities to low ones can find themselves with too much capital and choose to undertake negative investment. Moreover, when these shocks are sufficiently large, we observe negative spikes. Next, the tension between reproducing the empirical observations of spikes versus inaction is resolved by allowing for the possibility that not all investment is subject to fixed costs. In this case, plants not paying their adjustment costs may exhibit active investments, while nonetheless having their activities sufficiently constrained that they will eventually undertake an investment spike.

Aside from the moments of the time-averaged investment rate distribution presented in table 2, we also report our model's fit to some moments of establishment-level investment dynamics that were not targets in our calibration. First, we find that the variability of plant investment rates is reasonably well reproduced by our general equilibrium model. Simulating 1000 plants over 10,000 periods, the standard deviation of the typical plant's investment rate is 24.4 percent in our model, while it is 33.7 percent in the LRD (Cooper and Haltiwanger (2006)). Next, we consider the measure developed by Gourio and Kashyap (2007) to gauge the importance of the extensive margin in explaining changes in investment spikes. Using the LRD, they find that the correlation between the number of plants experiencing positive investment spikes and the total investment in these plants (as a fraction of aggregate capital) is 87 percent. The corresponding correlation for our model is 66.9 percent.

4.2 Partial equilibrium aggregate nonlinearities and smoothing

We begin our study of the implications of nonconvex capital adjustment costs by confirming that, when real wages and interest rates are held fixed at their steady state values, our model of lumpy investment exhibits important nonlinearities that survive aggregation. We simulate a partial equilibrium version of the model and compare its results to those in an otherwise identical frictionless model (distinguished only by its upper support on adjustment costs, $\bar{\xi} = 0$) in panel A of Table 3. Both models are subject to the same 10,000 period random draw of aggregate shocks. In choosing a margin along which to compare them, we follow the empirical investment literature, which has focused on changes in investment rates (that is, movements in the unfiltered ratio of investment to capital).

4.2.1 Nonlinearities

The frictionless model serving as our control is an element of the set of models that Caballero, Engel and Haltiwanger (1995) and Caballero (1999) refer to as *linear*, in that it is a special case of a quadratic capital adjustment cost model. These authors term such models linear based on the result that, if shocks are normally distributed, then so too are investment rates. Consistent with this, our frictionless model generates approximately zero skewness and excess kurtosis in aggregate investment rates.⁷ In the lumpy investment model, by contrast, nonconvex capital adjustment at the plant-level leads to a distribution of aggregate investment rates that is both sharply right-skewed and fat-tailed. This is the central and well-known nonlinearity in models of lumpy investment that has motivated the interest in their aggregate implications, summarized by Caballero (1999, pages 841-2) as follows. *"What is the aspect of the data that makes these models better than linear ones at explaining aggregate investment dynamics? ... it is the flexible cyclical elasticity of the increasing hazard model which allows it to better capture the high skewness and kurtosis imprinted on aggregate data by brisk investment recoveries."*

In fact, lumpy investment in our model increases skewness roughly 3-fold and kurtosis more than 15-fold relative to the frictionless control. This vivid evidence of nonlinearity in panel A establishes that our model is capable of delivering an aggregate role for lumpy investment similar to that found in previous partial equilibrium studies and summarized in Caballero's (1999) survey.⁸ However, if one compares the two rows of this panel to the near-zero third and fourth moments in the shaded row representing postwar U.S. investment rates, it appears that the additional skewness and kurtosis generated by lumpy investment does not improve model fit, but instead moves the investment series further from the data.⁹

The issue of the data's higher moments warrants further discussion. Clearly, the skewness (0.008) and excess kurtosis (-0.715) in our postwar data provide no evidence for the type of nonlinearities that partial equilibrium lumpy investment models are known to exhibit. Indeed, based on our 52-year sample of the aggregate private investment rate, we

⁷The slight skewness in the frictionless model arises from the log-normal distribution of aggregate shocks and decreasing returns to scale in the aggregate production technology.

⁸See, for example, Caballero and Engel (1999), Caballero, Engel and Haltiwanger (1995) and Cooper, Haltiwanger and Power (1999).

⁹These aggregate investment rate moments are similar whether we use the private sector capital stock, as we do here, or the business capital stock. In that case, persistence and standard deviation are 0.745 and 0.010, respectively, while skewness and excess kurtosis are 0.068 and -0.990 . Nonetheless, these moments of the data do depend upon the level of aggregation. Examining investment rates from two-digit U.S. manufacturing industries, Caballero and Engel (1999) find skewness and kurtosis of 0.61 and 0.74, respectively, for equipment and 0.76 and 0.87 for structures.

cannot reject the null hypothesis that the underlying stochastic process is normal. Following Valderrama (2007), we apply the Jarque-Bera test for departures from normality and obtain a test statistic of 1.1092. Under the null hypothesis that our data are drawn from a normal distribution, this statistic is distributed $\chi^2(2)$, so we cannot reject normality at even the 50 percent confidence level.¹⁰

We must note, however, that this does not resolve the issue of whether aggregate investment rates exhibit nonlinearities, because higher moments of the underlying stochastic process are often poorly estimated using time series of our length.¹¹ To illustrate this point, we simulate an alternative model for 10,000 periods. While the resulting model-generated data have persistence and volatility similar to the postwar aggregate data, their skewness and excess kurtosis are significantly different: 0.315 and 0.033, respectively.¹² As we look across 50-period subsamples of this data, there is considerable noise in the sample third and fourth moments. For example, although the median skewness, 0.30, is close to that of the full sample, we observe values at or below 0.008 (the postwar U.S. value) in 24.5 percent of the subsamples. In summary, while our empirical moments certainly are not suggestive of nonlinearities in the aggregate data, these higher moments computed from annual postwar data must be interpreted with caution.

4.2.2 Smoothing

While the lumpy investment literature has primarily focused on aggregate nonlinearities, there is also a smoothing aspect associated with fixed capital adjustment costs that has been emphasized in the recent work of Bachmann, Caballero and Engel (2006). We conclude our discussion of the partial equilibrium panel in table 3 by noting from its first two columns that our model is consistent with this more recent focus. This may come as little surprise, given an analogous finding by Bertola and Caballero (1994) in the context of investment irreversibilities.

It is well-understood that, in partial equilibrium, the frictionless model lacks necessary smoothing. In contrast to the aggregate data, its aggregate investment rate is both far too volatile and negatively auto-correlated. With real wages and interest rates held fixed, the full response in aggregate investment demand takes place immediately at the date of an aggregate shock. Thus, while productivity and hence aggregate capital are persistent, investment is not. Lumpy investment improves model fit in these two columns by removing

¹⁰ A Cramer-von Mises test on the residuals from an AR(2) specification of the aggregate investment rate finds a 0.934 probability of normality.

¹¹ We thank an anonymous referee for pointing this out to us, and for providing a simple example that is the basis of the one presented here.

¹² This is the model of Bachmann, Caballero and Engel (2007) discussed in section 6.

a substantial fraction of the excess variability and introducing some persistence.

Nonconvex adjustment costs at the micro-level induce inaction among plants with relatively high current costs or capital close to their target value. Thus, in the aggregate, investment demand initially responds less to a change in aggregate productivity than in the frictionless model without adjustment costs. However, because aggregate productivity changes are very persistent, many initially inactive plants undertake capital adjustments in subsequent periods. Thus, in partial equilibrium, fixed capital adjustment costs make aggregate investment both less variable and more persistent. However, despite these improvements, the partial equilibrium lumpy model continues to exhibit more than 10 times the variability, and far too little persistence, relative to the aggregate data.

4.2.3 Extensive margin

Figure 1 provides further evidence of the substantial changes lumpy investment implies for our model's aggregate dynamics when relative prices are fixed at their steady state values. The top panel shows the histogram of aggregate investment rates over the partial equilibrium simulation in the lumpy model; the bottom panel shows the corresponding histogram for the control model without adjustment frictions. Note first the abruptness in the frictionless model's investment rate distribution. Looking then to the top panel, we see that fixed adjustment costs smooth away some of this abruptness, as consistent with our discussion above. Moreover, while the distribution in the lower panel appears roughly symmetric, the inclusion of lumpy investment in the upper panel causes the distribution to lean rightward, and shifts more mass into the tails.

The added kurtosis arises from the fact that aggregate investment in the partial equilibrium lumpy model is more responsive to large aggregate shocks than to small ones, as consistent with the time-varying elasticity of investment rates stressed by previous authors in this literature. This follows directly from the rising shape of the hazards that govern the fractions of plants undertaking (unconstrained) capital adjustment in a period. This shape implies that small shifts in the hazards yield minimal changes in the numbers of adjusting plants, while larger shifts can generate disproportionately large changes in these numbers. The increased skewness arises from the fact that the model's investment series is more responsive to large positive shocks than it is to large negative ones. As we will explain, this happens because there are usually more plants concentrated on the lower ramps of the adjustment hazards, carrying too little capital relative to their targets, versus the upper ramps associated with excess capital.

Figure 2 illustrates the skewness arising in the partial equilibrium lumpy investment model by showing the responses in aggregate capital following a two standard deviation

positive shock to aggregate total factor productivity versus a same-sized negative shock. There, we plot capital's percent deviation from steady state in the lumpy investment and frictionless models under the assumption that the wage and real interest rate remain at their steady-state values. In response to the positive shock in period 20, the rise in the lumpy model's aggregate capital stock, 58 percent, is roughly the same as in the frictionless model, 59 percent. However, following the negative shock in period 40, the aggregate capital stock falls by 37 percent in the frictionless model, but by only 20 percent in the lumpy investment model. Thus, while nonconvex adjustment costs do smooth the responses in aggregate investment and thus capital to shocks, their effect is very nonlinear; responses to positive shocks are hardly affected, while responses to negative shocks are greatly dampened.

Skewness in the lumpy investment model's aggregate responses is caused by asymmetric changes in the numbers of plants undertaking (unconstrained) adjustments to their capital stocks. To explore this asymmetry, we must examine how the distribution of plants over capital evolves in response to aggregate disturbances. For expositional ease, we abstract from plant productivity shocks in this discussion to consider the effects of the two shocks above in a common productivity version of our model. In this case, there is a single adjustment hazard determining the fractions of plants that pay their fixed costs to adjust from each capital level to one common target. Figure 3 shows this adjustment hazard and the corresponding distribution in the model's steady state. The highest capital value at which the distribution has positive mass is the target, just below 1.38, which is adopted by all plants that pay their fixed adjustment costs. The dashed curve, which may be read off the right vertical axis, shows adjustment rates as a function of capital.

Note that the adjustment hazard rises in the distance between current capital and the capital stocks associated with the target (the capital stocks from which a plant can reach k^* for the next period without suffering an adjustment cost), because plants with capital further from the target are willing to suffer larger fixed costs to correct their stocks. Next, notice that, because both physical and economic depreciation continually erode nonadjusted capital stocks, plants enter the average period concentrated along the left ramp of the hazard with capital at or below the target.¹³ We define the aggregate adjustment rate in our model as the population-weighted sum of the fractions of plants adjusting to their target from each current capital; this is 0.223 in the steady state of the common productivity model shown in figure 3.

The left panel of figure 4 illustrates the extensive margin response to the two standard

¹³More generally, in our model with plant-specific productivity shocks, there is an adjustment hazard associated with each plant productivity level. Nonetheless, given mean-reversion in the shocks, the downward pressure of depreciation and technological progress continues to imply disproportionate concentrations of plants along the left ramps of the hazards.

deviation rise in productivity from above in figure 2, beginning with the bolded steady state distribution and hazard of figure 3, and continuing to hold relative prices fixed at steady state. Because changes in aggregate productivity are expected to persist, the positive shock causes a large rise in the expected marginal product of capital that in turn raises target capital sharply, to roughly 2.06. This leads the adjustment hazard to shift rightward, re-centering at the much higher capitals associated with the new target. Given that plants have all entered the period located along the left ramp of the steady state hazard, this shift increases the gap between actual and target capital for each of them, raising the value they place on capital adjustment. The total number of adjusting plants jumps to 0.986. This rise along the extensive margin reinforces the intensive-margin rise in the average investment undertaken by each adjusting plant. As a result, aggregate capital rises by far more than it would had adjustment rates remained unchanged. For the common productivity model, the 72 percent rise in the aggregate capital stock exceeds the 59 percent rise in the frictionless model (where idiosyncratic productivity differences do not affect aggregate responses).

By contrast, the right panel of figure 4 illustrates how the equivalent fall in aggregate productivity leads to a sharp decrease in adjustment rates, again beginning with the bolded steady state distribution and hazard. With the drop in expected future productivity, target capital falls to 0.926, shifting the hazard into the existing plant distribution. As a result, the fraction of plants for which adjustment is sufficiently valuable to offset the associated fixed costs declines markedly, and the number adjusting falls to 0.159. This offsets some of the decline in aggregate capital that would otherwise occur with the fall in target capital, leading the aggregate capital stock to fall by only 13 percent, far less than the 37 percent fall in the frictionless model.

In sum, we have seen that the lumpy investment model can exhibit asymmetry in its responses to shocks, and thus skewness in the distribution of aggregate investment rates, because rightward versus leftward shifts in its adjustment hazards generate asymmetric changes in the number of adjusting plants. These changes are sufficient to drive pronounced aggregate nonlinearities in partial equilibrium, because aggregate shocks are followed by sharp movements in target capital that cause large shifts in adjustment hazards. Returning to our full lumpy model with heterogeneity in both capital and productivity, this explains the sharp response following the positive shock in figure 2 (which, despite its adjustment frictions, very nearly reaches that in the frictionless control model) and its markedly dampened response following the negative shock, and thus the skewed distribution of partial equilibrium investment rates in figure 1.¹⁴

¹⁴Evidence of aggregate nonlinearity under partial equilibrium is even more extreme in the common productivity case of our model analyzed above. There, skewness and kurtosis are 1.90 and 5.29, respectively.

5 Aggregate Results

Having established our model’s consistency with essential features of the microeconomic data, as well as its nonlinearities in partial equilibrium, we now examine aggregate results. We find that, when relative prices vary over time to clear the markets for labor and goods, aggregate nonlinearities associated with lumpy investment disappear, leaving the economy’s dynamics virtually indistinguishable from those in the standard frictionless environment. Fluctuations in target capital stocks are smoothed dramatically, and thus so are the sharp changes along the extensive margin that generated the large positive skewness and excess kurtosis seen above. Moreover, by restraining changes along both the intensive and extensive margins, market-clearing relative price movements also smooth away the excessively large and abrupt fluctuations in aggregate investment demand we saw in partial equilibrium, thereby yielding an aggregate investment rate with persistence and volatility close to the data.

5.1 Model fit

Figure 5 presents the general equilibrium counterpart to the distribution of aggregate investment rates in the lumpy versus frictionless model examined in figure 1. Here, we report simulation results based on the same 10,000 period random draw of aggregate shocks as before, this time solving the models in equilibrium. Note that both model economies’ investment rates are greatly dampened by market-clearing movements in real wage and interest rates; thus, the range of investment rates in figure 5 is compressed nearly 10 times relative to figure 1. Nonetheless, differences in the equilibrium histograms appear minimal, with little evidence of added skewness or kurtosis in the lumpy investment panel.

These observations are confirmed by the second and higher moments presented in Table 3. There, we see an unambiguous improvement in model fit as we move from the partial equilibrium lumpy investment row of panel A to its general equilibrium counterpart in panel B. Persistence increases sharply, nearly reaching the empirical autocorrelation, while the excessive volatility in column 2 is virtually eliminated. Moreover, comparing the GE lumpy model to its PE counterpart, we see that equilibrium dramatically reduces the skewness and excess kurtosis in the distribution of aggregate investment rates. Viewing the four columns as a whole, the simulated aggregate investment rate series matches its counterpart in the

Plant-specific productivity shocks reduce the lumpy investment model’s potential for aggregate nonlinearities, because their presence implies more dispersion in the average distribution of plants over capital, as well as greater symmetry in the typical concentration of plants along the left versus right ramps of adjustment hazards.

data relatively well when the effects of equilibrium are included in the lumpy investment environment, and far less so when these effects are ignored.

Although our general equilibrium model yields empirically viable aggregate investment dynamics, comparison of the two rows within panel B reveals that the nonconvex investment technology faced by plants has no role in this success. Consistent with our observations in Figure 5, we see that differences in the aggregate dynamics of the lumpy investment versus frictionless model are largely eliminated in general equilibrium. Aggregate investment rates exhibit nearly identical volatility, skewness, and kurtosis across the two model economies. Moreover, there is virtually no difference in persistence, which again is far greater in both models than it was in panel A and very close to that in the data. This similarity in the models' aggregate investment rate dynamics extends to other key macroeconomic aggregates. The variabilities and contemporaneous correlations of aggregate output, employment, consumption and investment in Table 4 indicate that the aggregate business cycle is essentially unaffected by lumpy investment.

Given the differing investment technologies across these models, alongside the marked differences they implied in the partial equilibrium dynamics of section 4.2, the explanation for the similarities within panel B of Table 3 (as well as those in Figure 5 and Table 4) must be traceable to the influence of the representative household they share in common. Persistence in aggregate investment rates is an immediate result of households' preference for consumption smoothing. The omission of this channel in partial equilibrium places an emphasis on capital adjustment costs in Panel A to generate some of this otherwise absent persistence. In equilibrium, by contrast, adjustment costs are not necessary to smooth aggregate investment demand; this is achieved far more effectively through market-clearing changes in relative prices.

5.2 Why nonlinearities dissolve

As discussed above, the nonlinearities generated by lumpy investment in partial equilibrium arise because changes in aggregate productivity are followed by large movements in target capitals that can cause sharp, concurrent changes in the fractions of plants undertaking (unconstrained) capital adjustment. However, such synchronizations in the timing of large investment projects would in turn imply large movements in households' consumption. When we impose market-clearing, this volatility is sharply restrained by procyclical real wages and interest rates, which dampen the changes in target capitals arising from aggregate shocks, and thus dampen changes in adjustment rates.

Recall our earlier example in Figure 2, where we traced our model's responses to a large rise and fall in aggregate productivity. In contrast to the asymmetry there, where

aggregate capital rose far more sharply following a 2-standard deviation positive aggregate shock than it fell after the same-sized negative shock, the general equilibrium responses to aggregate shocks are far more symmetric. Following the positive shock, capital now rises by 2.86 percent, while it falls by 2.53 percent with the negative shock. This reflects no greater asymmetry than occurs in the near-linear frictionless model, where capital rises 2.86 percent and falls 2.6 percent in response to the two shocks.

The reason general equilibrium is so effective in eliminating the lumpy investment model's aggregate nonlinearities is that it smooths away much of the movement in target capitals that are essential in generating large changes along the extensive margin. For example, following the positive shock in figure 2, the average target capital stock (weighting the 15 productivity-specific targets by the ergodic distribution over productivities) rose more than 43 percent when our model was solved in partial equilibrium, thereby triggering sufficiently large rightward shifts in the adjustment hazards as to raise the total adjustment rate by 59 percentage points and the aggregate investment rate by 59 percentage points. In general equilibrium, the same shock causes only a 5.1 percent rise in the average target. As a result, shifts in the adjustment hazards are minimal, leading the aggregate adjustment rate to rise by only 3.6 percentage points and the aggregate investment rate to rise just 2.9 percentage points.

Similarly, following the large negative aggregate shock, the fall in average target capital is 5.5 percent in general equilibrium (versus 40.4 percent in partial equilibrium), yielding only a 3 percentage point fall in the number of adjustors, and a 2.6 percentage point decline in the aggregate investment rate (versus the 20.3 percentage point decline under partial equilibrium). Thus, the rise and fall in the aggregate investment rate following these aggregate shocks is an order of magnitude smaller in equilibrium and almost perfectly symmetric.¹⁵

In closing, we re-emphasize that the absence of aggregate nonlinearities must not be viewed as a failure of the equilibrium lumpy (or frictionless) model, as these features are not apparent in the aggregate data. With or without capital adjustment costs, our general equilibrium models match the higher moments of postwar aggregate investment rates reasonably well, far more so than the partial equilibrium lumpy model where skewness and kurtosis are sharply overstated.

¹⁵To appreciate the extent of this symmetry, one need only compare these outcomes with those in the frictionless model. There, through changes in the intensive margin alone, the aggregate investment rate rises by 2.91 percentage points and falls by 2.67 percentage points in response to the two shocks.

6 Robustness and reconciliation

While the calibration of most parameters of our lumpy investment model is standard and consistent with the method of Prescott (1986), we have selected the adjustment cost parameters distinguishing this model from its frictionless counterpart, as well as the volatility of plant productivity shocks, to match summary statistics taken from time-averages of the microeconomic investment data. This approach differs sharply from that taken in the recent work of Bachmann, Caballero and Engel (2006), and it implies far smaller fixed adjustment costs. These authors choose the size of fixed costs in their single sector model to match the volatility of sectoral investment rates from 3-digit manufacturing data. Defining a group of plants drawing a common sectoral shock as a sector, and assuming sectoral changes do not affect real wages or interest rates, Bachmann et al. require large adjustment costs to smooth sectoral capital reallocation sufficiently to match the variability in the data, since sectoral outputs are perfectly substitutable in their environment.

As our equilibrium model does not show evidence of aggregate nonlinearities, while that in Bachmann et al. does, we begin this section by considering whether the disparity in our findings may arise from the differing size of our fixed costs, or from our lesser volatility in plant-specific productivities. In row 1 of table 5, we return to the same 10,000 period simulation used in our baseline results, this time reporting the resulting aggregate investment rate moments when the upper support of the fixed cost distribution is raised 10-fold to imply much larger adjustment costs in the model. In the next row, we make these costs even larger, raising the upper support 25-fold, and simultaneously increase the variability of idiosyncratic shocks to three times that of the aggregate shock. While these changes do slightly reduce aggregate volatility, the third and fourth moments indicate that they have little effect in generating nonlinearities.

Given that large increases in the size of our fixed adjustment costs and plant-specific productivity variations fail to alter our baseline results, we next consider the second distinguishing feature of our model relative to that examined by Bachmann et al., the incidence of adjustment costs. Our baseline calibration allows for a range of investment rates around zero that are exempt from fixed costs. Bachmann et al. allow no such interval; instead, they assume that, if a plant chooses not to pay its fixed cost, it must replace 50 percent of the capital that it would otherwise passively shed through depreciation. In the third row of table 5, we re-examine our model's aggregate results when we modify the interval of investment rates exempted from adjustment costs to allow plants to replace all of their depreciated capital without incurring fixed costs; however, we do not force our plants to undertake this investment. Again, we find negligible changes in the model's aggregate dy-

namics relative to our baseline results, with no greater evidence of skewness or kurtosis. From this, we conclude that it is not plants' ability to undertake maintenance investment without incurring adjustment costs that explains the nonlinearities uncovered in Bachmann et al. and absent in our results.¹⁶

In the remainder of this section, we take a different approach toward isolating the sources of the disparate findings across these two equilibrium studies. We begin by imposing each of the assumptions made by Bachmann et al. to reproduce their aggregate nonlinearities, using essentially their parameter values (with an innocuous adjustment to allow for a balanced growth path). Next, to gauge the importance of each, we remove one assumption at a time so long as evidence of nonlinearities remains.

A final assumption that will be important in this exercise is that regarding the representative household's attitude toward risk and consumption smoothing. Across our lumpy investment and frictionless model economies, the representative household has a unit elasticity of intertemporal substitution. While this may be considered somewhat high, it lies within the range of standard values applied in quantitative dynamic stochastic general equilibrium studies.¹⁷ By contrast, the household in the Bachmann et al. lumpy investment economy is almost risk-neutral, with this elasticity exceeding 9. Such preferences imply far greater tolerance for fluctuations in consumption, thus making the household supply of investment goods very flexible and encouraging larger movements in target capital stocks and thus adjustment hazards.

When we combine the large adjustment costs, high variance in plant-level productivity, and mandatory maintenance investment described above together with the near risk-neutrality assumption, we obtain the Bachmann et al. nonlinearities in aggregate investment rates; skewness rises from the 0.067 of our model nearly 5-fold, to 0.315. After reducing the elasticity of intertemporal substitution to our value, $\sigma = 1$, we find that roughly half

¹⁶Motivated by arguments in Gourio and Kashyap (2007), we have also explored cases where adjustment costs are both very large and highly predictable for plants, as well as cases with quite low returns to scale in production, where the overall incidence of fixed costs is magnified by difficulties in concentrating production among few plants. In each case, the equilibrium lumpy model exhibits no evidence of nonlinearities. For example, when we assume that fixed costs are drawn from a beta distribution that is sharply right-skewed with most probability concentrated at its upper support ($\alpha = 3$ and $\beta = \frac{1}{3}$), and we select the upper support at $\frac{1}{2}$ to imply total output lost to adjustment costs averages 3.5 percent of total investment (as in Gourio and Kashyap (forthcoming)), the resulting third and fourth moments for the model's aggregate investment rate are 0.077 and -0.062 respectively. Continuing to assume the beta distribution, these moments are 0.080 and -0.065 when we simultaneously reduce returns to scale from 0.90 to 0.65 (while lowering φ to maintain total hours averaging $\frac{1}{3}$ and raising σ_{η_z} to hold the variability of output in the frictionless model unchanged relative to the baseline results).

¹⁷The same value is chosen by Bachmann et al. for the household residing in their frictionless economy.

of this skewness remains, 0.151. Next, we eliminate the assumption that plants must replace one-half of their depreciated capital in any period they do not pay their fixed costs, instead applying the more traditional assumption that such plants undertake zero investment.¹⁸ This removes virtually all remaining skewness, despite the large adjustment costs and volatile plant productivity shocks still in place. The result, at 0.071, is indistinguishable from that in our model. Thus, we find that the aggregate results of Bachmann et al. are reconciled to ours with the removal of their assumptions regarding household preferences and mandatory maintenance. The question of whether nonconvex capital adjustment costs cause aggregate nonlinearities then appears simply a question about the plausibility of these two assumptions.

7 Additional Results

We now turn to examine our model’s predictions involving plant-level investment. In this section of additional results, we explore how plants’ investments are influenced by the interaction of idiosyncratic shocks, fixed adjustment costs, and equilibrium price movements.

7.1 Role of general equilibrium in microeconomic dynamics

Given the large influence of general equilibrium in our findings involving aggregate dynamics, we begin by considering how much information is lost in abstracting from market-clearing price adjustments if our interest lies instead in the dynamics of highly disaggregated series, such as the investments undertaken by individual production units. Naturally, when idiosyncratic productivity shocks give rise to a nontrivial distribution of plants over investment rates, and the effect of these plant-level shocks is large relative to that of the calibrated aggregate shock, the time-averaged cross-sectional distribution is relatively unaffected by equilibrium. Differences in plants’ investment rates, on average, represent largely a reallocation of capital from one production unit to another, and such reallocations have no effect on aggregate investment. However, this does not imply that the *dynamics* of plant-level investment are independent of equilibrium.

The period-by-period distribution of plant investment rates changes over time in response to aggregate shocks. Because the associated movements in relative prices feed back into

¹⁸In a sense, required maintenance investment has the effect of imposing a partial investment irreversibility. Under this assumption, plants experiencing rises in their relative productivity are reluctant to adopt a high capital stock, as they may be forced to retain much of it in the future until they pay a fixed adjustment cost. This then compresses the steady-state distribution of plants over capital. As a result, any given shift in adjustment hazards implies greater changes in overall adjustment rates.

plant-level decisions, the extent of these changes may be sharply distorted in a study that omits market-clearing relative price adjustments. For example, when the lumpy investment model is simulated for 10,000 periods in partial equilibrium, the standard deviation of the fraction of the economy's plants exhibiting positive investment spikes in each period is 0.12, and the standard deviation of the size of these spikes is 0.08. When the same simulation is undertaken in general equilibrium, each of these standard deviations falls to 0.01.

We further illustrate this point through the comparison of some simple panel regressions in Table 6. In row 1, we regress plant investment rates on changes in aggregate productivity, Δz , and changes in plant-specific productivity, $\Delta \varepsilon$, using simulated data from our general equilibrium lumpy investment model. In row 2, we repeat this same exercise using data from the lumpy model simulated in partial equilibrium. As expected, the coefficients on $\Delta \varepsilon$ across these two rows are large and essentially identical. However, plant investments are far less responsive to changes in aggregate total factor productivity when the resulting market-clearing price movements are included than they are when these restraints are ignored. The general equilibrium coefficient on Δz reflecting plants' investment elasticity to an aggregate shock, 0.423, is exaggerated 13-fold when real wage and interest rates are not allowed to respond to the change in aggregate conditions.

The sharply differing micro-level elasticity to changes in aggregate total factor productivity naturally implies that a given plant will invest quite differently when it faces equilibrium prices. Over those dates when changes in aggregate productivity are large, the errors introduced by ignoring endogenous movements in relative prices will be substantial. From this, we conclude that equilibrium analysis is essential for understanding the dynamics of investment even at the most disaggregated level, most particularly in times of large aggregate disturbances. It may be useful to recall that it is precisely during such times that existing partial equilibrium studies have found substantial differences between lumpy investment models and standard linear models with convex adjustment technologies (see, for example, Caballero, Engel and Haltiwanger (1995) and Caballero and Engel (1999)).

Our results in Table 6 suggest that structural parameters cannot be safely inferred from partial-equilibrium models, even using the most disaggregated data, when the moments used to identify these parameters are not drawn from time-averaged cross-sectional data. To see this, consider the following simple exercise loosely based upon Cooper and Haltiwanger's (2006) estimation of capital adjustment costs using indirect inference. Suppose a researcher observes establishment-level data on $\frac{i}{k}$ and a , where $a \equiv z\varepsilon$, drawn from the general equilibrium simulation of our lumpy investment economy. Regressing the plant investment rate on $\log(a_t) - \log(a_{t-1})$, alongside a constant and a time dummy, he finds the resulting coefficient on the composite productivity shock is $\beta_1 = 3.66$. Next, aware that the economy has

a single type of capital adjustment friction, a fixed cost distributed $U(0, \bar{\xi})$, and perfectly informed about all other structural parameters, the researcher estimates $\bar{\xi}$ using the partial equilibrium model. A standard simulated method of moments approach targeting β_1 leads him to an estimate of 0.0371. This is roughly 5 times the actual value, $\bar{\xi} = 0.00825$. In lieu of market-clearing relative price adjustments, the size of capital adjustment frictions must be exaggerated to reign in plants' excessive responses to changes in aggregate productivity.

7.2 Role of nonconvexities in microeconomic investment

Throughout the literature, it has often been suggested that nonconvex capital adjustment costs cause lumpy investments at the plant.¹⁹ A brief comparison of the average plant investment distribution in our model economy relative to that arising in absence of adjustment costs reveals that this is not necessarily true. In fact, if the effect of plant-level productivity shocks is large relative to aggregate shocks, as consistent with recent estimates, their presence can completely overturn the previously understood role of nonconvex adjustment costs at the plant.

Absent differences in plants' productivities, the standard model with frictionless investment implies a continuum of identical plants that, on average, undertake modest positive investments in every period to replace their depreciated capital. The introduction of nonconvex capital adjustment costs to this environment necessarily generates the trademark features of lumpy investment, inaction and spikes, since they lead some plants to delay adjustment sufficiently that their capital stocks drift far from that to which they eventually adjust. However, consider instead the frictionless model with idiosyncratic productivity shocks reported in row 1 of Table 7. Here, before the inclusion of capital adjustment costs, there is already a nontrivial cross-sectional distribution of plant investment rates determined by the distribution of plant-specific productivity shocks. Notice that these volatile idiosyncratic shocks on their own cause both positive and negative investment spikes in the average year; in fact, these observations are overstated relative to the LRD. When fixed adjustment costs are added to this environment, moving from row 1 to row 2, we observe that fixed costs do not *cause* additional investment spikes, but instead *eliminate* some of their occurrences.

Under frictionless capital adjustment, plants' investments are, on average, extremely responsive to changes in their individual productivities. Through a pure reallocation of aggregate investment, the economy in row 1 exhibits positive investment spikes among plants

¹⁹It is well-understood that partial irreversibilities cause inactivity in investment, but not spikes (for an excellent analysis see Caballero (1999)). Among capital adjustment frictions, this characteristic of lumpy investment can only be attributed to fixed costs, given the increasing returns adjustment technology they imply.

experiencing large increases in relative productivity, as well as negative spikes among those suffering large decreases in relative productivity. Nothing restrains this reallocation, because it can be costlessly reversed in any subsequent date. However, such reversals are not costless in the economy with capital adjustment frictions. A plant in row 2 realizing raised relative productivity will be more cautious in selecting its new target capital, knowing that, in some nearby date when its productivity may fall, it will face a high fixed cost to re-adjust its capital to a lowered target. Thus, even if it pays its current fixed cost to adopt a new target, the plant's investment is tempered by an effort to avoid future adjustment costs. A similar restraint applies to negative investment in response to a fall in relative productivity. For this reason, fixed adjustment costs act to reduce the volatility of plant investments, yielding fewer investment spikes and more episodes of inactivity.

Beyond eliminating excess spikes, nonconvex costs can take on a secondary role that further reduces the distance between model and actual plant data, but has not been emphasized in previous studies. Comparing the ratio of positive investment (spike) observations to negative investment (spike) observations in row 1 versus row 2, we find that the adjustment costs boost asymmetries in plant-level investment. Recall our explanation above for why these costs eliminate excess spikes. Because depreciation and technological progress automatically erode inactive plants' effective capital stocks, the reduction in a plant's value caused by a high future adjustment cost will be greater if a plant finds itself with too little capital, rather than too much capital. As a result, plants are less cautious in raising their capital stocks than in lowering them, so that the presence of fixed costs increases the fraction of positive investment (spike) observations relative to negative ones.

We emphasize the changed role of nonconvex costs in the presence of idiosyncratic shocks because it is essential that we understand what these costs actually do if we are to establish their importance in explaining establishment-level investment. If their role is to remove investment spikes and cause asymmetry, as we have seen here, then the same effect might be similarly achieved by either investment irreversibilities or by combinations of fixed costs, irreversibilities and convex costs. Thus there may be no aspect of microeconomic investment behavior that is uniquely explained by the presence of nonconvex adjustment costs. Viewed another way, if idiosyncratic shocks are the primary force explaining plant investment differences, it may be virtually impossible to use the establishment investment data to infer the relative size or importance of these costs relative to other frictions.

8 Concluding Remarks

We have studied a dynamic stochastic general equilibrium model with nonconvex capital adjustment costs and plant-specific differences in productivity. By introducing persistent plant-level productivity shocks and allowing small investments exempt from adjustment costs, our model reproduces essential empirical regularities involving establishment-level investment. In this environment, equilibrium movements in real interest rates and wages play an indispensable role in adding persistence to aggregate investment rates, bringing the model closer to the data. An additional consequence of such movements in relative prices is that they eliminate the implications of plant-level nonconvexities for aggregate dynamics. In partial equilibrium, these nonconvexities lead to aggregate nonlinearities through large changes in plants' target capital stocks that drive large changes in the fractions of plants adjusting to these targets. Such nonlinearities disappear in general equilibrium, however, because procyclical movements in real wages and interest rates substantially dampen the changes in plants' target capital stocks that follow an aggregate shock.

Throughout our analysis, we have assumed that output is perfectly substitutable across production units. This makes the reallocation of resources, in response to idiosyncratic shocks, optimal from the perspective of the representative household. Conversely, it also encourages the avoidance of capital adjustment costs through the concentration of investment. If we instead considered an environment where firms produced distinct goods, such disparities in the distribution of inputs would be more costly. This suggests that, for a given distribution of idiosyncratic shocks and capital adjustment costs, more firms would undertake unconstrained capital adjustments in each period. With the average fraction of firms adjusting nearer one, the model with fixed adjustment costs would more closely resemble environments with convex adjustment technologies.

References

- [1] Bachmann, R., R. J. Caballero and E. M. R. A. Engel (2006) “Lumpy Investment in Dynamic General Equilibrium,” Working Paper, Yale University.
- [2] Bayer, C. (2006) “Investment Dynamics with Fixed Capital Adjustment Cost and Capital Market Imperfections” *Journal of Monetary Economics* 53, 1909-1947.
- [3] Bertola, G. and R. J. Caballero (1994) “Irreversibility and Aggregate Investment,” *Review of Economic Studies* 61, 223-246.
- [4] Bloom, N. (2007) “The Impact of Uncertainty Shocks,” Working Paper, Stanford University.
- [5] Caballero, R. J. (1999) “Aggregate Investment,” chapter 12 in M. Woodford and J. Taylor (eds.) *Handbook of Macroeconomics*, vol. IB. Amsterdam: North Holland.
- [6] Caballero, R. J and E. M. R. A. Engel (1999) “Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach,” *Econometrica* 67, 783-826.
- [7] Caballero, R. J., E. M. R. A. Engel and J. C. Haltiwanger (1995) “Plant-Level Adjustment and Aggregate Investment Dynamics,” *Brookings Papers on Economic Activity* 2, 1-39.
- [8] Comin, D. and T. Philippon (2005) “The Rise in Firm-Level Volatility: Causes and Consequences,” NBER Working Paper No. 11388.
- [9] Cooper, R. W., J. C. Haltiwanger and L. Power (1999) “Machine Replacement and the Business Cycle: Lumps and Bumps,” *American Economic Review* 89, 921-946.
- [10] Cooper, R. W. and J. C. Haltiwanger (2002) “On the Nature of Capital Adjustment Costs,” Working Paper, University of Texas at Austin.
- [11] Cooper, R. W. and J. C. Haltiwanger (2006) “On the Nature of Capital Adjustment Costs,” *Review of Economic Studies* 73, 611-633.
- [12] Gourio, F. and A. K. Kashyap (2007) “Investment Spikes: New Facts and a General Equilibrium Exploration,” *Journal of Monetary Economics* 54, S1-S22.
- [13] Hansen, G. D. (1985) “Indivisible Labor and the Business Cycle,” *Journal of Monetary Economics* 16, 309-327.

- [14] Khan, A. and J. K. Thomas (2003) “Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?” *Journal of Monetary Economics* 50, 331-360.
- [15] Krusell, P. and A.A. Smith Jr. (1997) “Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns,” *Macroeconomic Dynamics* 1, 387-422.
- [16] Prescott, E. C. (1986) “Theory Ahead of Business Cycle Measurement,” *Federal Reserve Bank of Minneapolis Quarterly Review* 10, 9-22.
- [17] Rogerson, R. (1988) “Indivisible Labor, Lotteries and Equilibrium,” *Journal of Monetary Economics* 21, 3-16.
- [18] Thomas, J. K. (2002) “Is Lumpy Investment Relevant for the Business Cycle?” *Journal of Political Economy* 110, 508-534.
- [19] Valderrama, D. (2007) “Statistical Nonlinearities in the Business Cycle: A Challenge for the Canonical RBC Model” *Journal of Economic Dynamics & Control* 31, 2957–2983.
- [20] Veracierto, M. L. (2002) “Plant-Level Irreversible Investment and Equilibrium Business Cycles,” *American Economic Review* 92, 181-197.

Figure 1: Distribution of aggregate investment rates in partial equilibrium





Figure 2: Partial equilibrium aggregate capital responses




Figure 3: Steady state adjustment in common productivity lumpy model




Figure 4: Adjustment responses in the common productivity model




Figure 5: Distribution of aggregate investment rates in general equilibrium





Table 1. Baseline parameter values

γ	β	δ	α	ν	φ	ρ_z	$\sigma_{\eta z}$	ρ_ε	$\sigma_{\eta\varepsilon}$	b	$\bar{\xi}$
1.016	0.977	0.069	0.256	0.640	2.400	0.859	0.014	0.859	0.022	0.011	0.0083

Table 2. Matching the time-averaged plant investment data

<i>average annual fraction of plants exhibiting:</i>	Inaction	Positive Spike	Negative Spike	Positive Invest.	Negative Invest.
LRD data	0.081	0.186	0.018	0.815	0.104
(1) traditional lumpy	0.789	0.187	0.000	0.211	0.000
(2) extended lumpy with plant-specific TFP	0.073	0.185	0.010	0.752	0.175

DATA: Cooper and Haltiwanger (2005). Traditional model: $\sigma_{\eta\varepsilon} = b = 0$, upper support = 0.014 to fit positive spike obs.

Table 3. Role of nonconvexities in aggregate investment rate dynamics

	Persistence	Standard Deviation	Skewness	Kurtosis
postwar U.S.	0.695	0.008	0.008	- 0.715
A. partial equilibrium models				
(1) PE frictionless	- 0.069	0.128	0.358	0.140
(2) PE lumpy investment	0.210	0.085	1.121	2.313
B. general equilibrium models				
(1) GE frictionless	0.659	0.010	0.048	0.048
(2) GE lumpy investment	0.662	0.010	0.067	- 0.074

DATA: Annual private investment-to-capital ratio, 1954 – 2005, computed using BEA tables.

Table 4. Aggregate business cycle moments*

	Y	Z	N	C	I	K
A. Standard Deviations Relative to Output						
(1) GE Frictionless	2.277	0.602	0.645	0.429	3.562	0.494
(2) GE Lumpy	2.264	0.605	0.639	0.433	3.539	0.492
B. Contemporaneous Correlations with Output						
(1) GE Frictionless		1.000	0.955	0.895	0.976	0.034
(2) GE Lumpy		1.000	0.956	0.900	0.976	0.034

* The logarithm of each series is HP-filtered using a weight of 100. The first column of panel A reports percent standard deviations of output.

Table 5. Robustness

	Persistence	Standard Deviation	Skewness	Kurtosis
postwar U.S.	0.695	0.008	0.008	- 0.715
baseline lumpy	0.662	0.010	0.067	- 0.074
(1) big fixed costs	0.677	0.009	0.071	- 0.066
(2) huge fixed costs and big σ_ε	0.681	0.009	0.071	- 0.064
(3) no fixed costs for $i \in [0, \delta k]$	0.665	0.010	0.070	- 0.057

Table 6. Equilibrium and the dynamics of plant investment rates

<i>Panel regressions from simulated lumpy models.</i> plant i/k on:	constant	Δz	$\Delta \varepsilon$	adj. R-squared	S.E. of regression
(1) General equilibrium data	0.107	0.423 (0.047)	5.126 (0.032)	0.227	0.215
(2) Partial equilibrium data	0.112	5.464 (0.051)	5.414 (0.034)	0.289	0.233

Table 7. Role of fixed costs in plant-level investments

<i>average annual fraction of plants exhibiting:</i>	Inaction	Positive Spike	Negative Spike	Positive Invest.	Negative Invest.
LRD data	0.081	0.186	0.018	0.815	0.104
(1) frictionless model	0.032	0.204	0.028	0.611	0.356
(2) lumpy investment model	0.073	0.185	0.010	0.752	0.175

Appendices

A Numerical method

Solving the frictionless model is fairly straightforward, even in the presence of persistent plant-level shocks. Despite a distribution of plants over capital and productivities, the endogenous aggregate state vector may be characterized by total capital and a time-invariant distribution of plants' shares of the aggregate capital stock that are functions of their idiosyncratic productivity levels (as explained in section B). Given the invariance in this distribution of capital shares, the aggregate state vector contains only two time-varying elements, total capital and aggregate productivity, and standard methods may be used to solve the model. The one novelty in our approach is that we apply a nonlinear solution method using piecewise polynomial cubic spline interpolation of the planner's value function. In partial equilibrium, the same nonlinear approach is applied to solving plants' value functions for the lumpy investment models. Note that uncertainty in adjustment costs implies that value functions are smooth objects.

In contrast to its frictionless counterpart, the equilibrium lumpy investment model's aggregate state vector involves a nontrivial distribution of plants, which makes the computation of equilibrium more challenging. Our solution algorithm involves repeated application of the contraction mapping implied by (8) - (10) to solve for plants' start-of-period value functions V^0 , given the price functions $p(z, \mu)$ and $\omega(z, \mu)$ and the laws of motion implied by $\Gamma(z, \mu)$, (π_{ij}) and (π_{lm}^ε) . This recursive approach is complicated in two ways. First, recalling that a primary focus throughout this literature has been on lumpy investment's potential for generating aggregate nonlinearities, we must use a solution method that does not rule them out. Moreover, we must adopt a non-local method, because plants' (S, s) decision rules can sometimes hit corners (when interior values for their threshold costs shift to the boundaries of the cost distribution). As explained below, this implies that the number of capital stocks with a positive measure of plants changes over time.

The distribution μ in the aggregate state vector is a large object, made larger by plants' ability to make small adjustments without paying fixed costs. In general, discrete choices imply that this distribution is highly non-parametric. For each level of plant productivity, we store the conditional distribution using a fine grid defined over capital. However, plants' choices of investment are not restricted to conform to this grid.

The nonlinear solution method predicated by our focus on aggregate nonlinearities makes this numerically intractable, so we use selected moments as a proxy for the distribution in

the aggregate state vector.

More specifically, our solution adapts the method of Khan and Thomas (2003) to allow for a two-dimensional distribution of plants over capital and idiosyncratic productivity. Thus, we assume that agents use a smaller object in proxy for the distribution as they forecast the future state to make decisions given current prices. In choosing this proxy, we apply a variation on the method of Krusell and Smith (1997), assuming that agents approximate the distribution in the aggregate state vector with a vector of moments, $m = (m_1, \dots, m_J)$, drawn from the true distribution. Because our work involves discrete choices among producers, we find that using the conditional means from I equal-sized partitions of the distribution is efficient, implying small forecasting errors.

Our solution algorithm iterates between an inner loop and an outer loop, as in Krusell and Smith (1997). In the inner loop, agents' value functions are solved based upon a given set of forecasting rules. Given these value functions, the economy is simulated in the outer loop, where p is endogenously determined in each date. Throughout our simulations, we use the actual distribution of plants over capital and productivity in each period, alongside plants' value functions (derived using the forecasting rules described above), to determine equilibrium prices and quantities, and thus the subsequent period's distribution. Next, the resulting simulation data are used to update the forecasting rules, with which we return to the inner loop, this two-step process continuing until the forecasting rules converge.

Table A1 presents agents' forecasting rules for the special case of our model without plant-level differences in productivity. In determining their current decisions, agents forecast the future proxy state, m'_1 , assumed to be the logarithm of the first moment of the distribution of plants over capital, using the logarithm of the mean of the current distribution, m_1 (and current aggregate productivity). Similarly, when solving for agents' value functions, we have them assume that the valuation of current output, p , is a log-linear function of this mean. Note that adjusted R-squareds are very high, and standard errors are small; almost all the true variation in the mean of the distribution, and in the relative price of output, may be anticipated using these simple forecasting rules.

In our full lumpy investment model, there is a two-dimensional distribution of plants over capital and idiosyncratic productivity. Nonetheless, we find that the solution method described above is robust to this additional source of heterogeneity. The equilibrium forecasting rules are presented in Table A2. Note that there is no loss of accuracy in the forecasting rules with the introduction of persistent differences in plant-specific productivity, though we continue to use only the unconditional mean of the distribution of capital as a proxy for the aggregate endogenous state. This suggests that our general equilibrium solution method may be applied to a broad class of models currently studied in partial equilibrium.

B A characterization of the frictionless model

In this appendix, we derive several analytical results for the frictionless model. In lemma 1, under the assumption of Cobb-Douglas production, we establish that the plant decision rule for next period's capital stock may be expressed as the product of two functions whose arguments are the current plant-specific productivity term and the aggregate state, respectively. Thus a plant's decision rule for future capital is independent of its current capital. It is then immediate that, given any initial distribution of plants, future distributions involve only N_ε time-varying values of capital with positive mass. Furthermore, the separability of plants' capital stock decision rules into a plant-specific and an aggregate component implies that the shares of the aggregate capital stock across plant types are time-invariant. This, in turn, implies that the aggregate capital stock is sufficient to fully characterize variation in the endogenous state vector, as is stated in lemma 2. It follows that all aggregate dynamics of the full model may be recovered using a representative firm approach, although we omit the details here.

We begin our analysis of the frictionless model by describing the problem of a plant. In the absence of capital adjustment costs, the value of any plant of type (ε_l, k) will solve the following functional equation:

$$v^1(\varepsilon_l, k; z_i, \mu) = \max_{n, k'} \left[z_i \varepsilon_l F(k, n) - \omega(z_i, \mu) n - \gamma k' + (1 - \delta) k \right. \\ \left. + \sum_{j=1}^{N_z} \pi_{ij} d_j(z_i, \mu) \sum_{m=1}^{N_\varepsilon} \pi_{lm}^\varepsilon v^1(\varepsilon_m, k'; z_j, \mu') \right], \quad (18)$$

subject to $\mu' = \Gamma(z_i, \mu)$. Let $N(\varepsilon_l, k; z_i, \mu)$ describe the plant's employment choice and $K(\varepsilon_l, k; z_i, \mu)$ its decision rule for next period's capital stock. The description of households in section 2.2 of the text is unchanged.

A *recursive competitive equilibrium* is a set of functions

$$\left(\omega, (d_j)_{j=1}^{N_z}, \rho_1, \rho_0, v^1, N, K, W, C, N^h, \Lambda^h, \Gamma \right)$$

that solve plant and household problems and clear the markets for assets, labor and output:

1. v^1 satisfies (18) and (N, K) are the associated policy functions for plants.
2. W satisfies (5) and (C, N^h, Λ^h) are the associated policy functions for households.
3. $\Lambda^h(\varepsilon_m, k', \mu; z, \mu) = \mu'(\varepsilon_m, k')$, for each $(\varepsilon_m, k') \in \mathbf{S}$.
4. $N^h(\mu; z, \mu) = \int_{\mathbf{S}} N(\varepsilon, k; z, \mu) \mu(d[\varepsilon \times k])$.

5. $C(\mu; z, \mu) = \int_{\mathbf{S}} \left(z \varepsilon F(k, N(\varepsilon, k; z, \mu)) - \gamma K(\varepsilon, k, \xi; z, \mu) + (1 - \delta)k \right) \mu(d[\varepsilon \times k])$.
6. $\mu'(\varepsilon_m, B) = \int_{\{(\varepsilon_l, k) \mid K(\varepsilon_l, k; z, \mu) \in B\}} \pi_{lm}^\varepsilon \mu(d[\varepsilon_l \times k])$, for all $(\varepsilon_m, B) \in \mathcal{S}$, defines Γ .

We now specialize to the case of Cobb-Douglas production and characterize plant decision rules. Let $\alpha \in (0, 1)$ represent capital's share of production and $\nu \in (0, 1)$ be labor's share, where $\alpha + \nu < 1$. The choice of employment, n , solves $\max_n (sk^\alpha n^\nu - \omega n)$, where $s = z\varepsilon$ and ω is the real wage. This yields the employment decision rule $n = \left(\frac{\nu sk^\alpha}{\omega}\right)^{\frac{1}{1-\nu}}$, allowing us to express production as $y = s^{\frac{1}{1-\nu}} k^{\frac{\alpha}{1-\nu}} \left(\frac{\nu}{\omega}\right)^{\frac{\nu}{1-\nu}}$. Production net of labor costs is then given by the following:

$$y - \omega n = (1 - \nu) s^{\frac{1}{1-\nu}} k^{\frac{\alpha}{1-\nu}} \left(\frac{\nu}{\omega}\right)^{\frac{\nu}{1-\nu}}. \quad (19)$$

Substituting (19) into (18), we remove the static employment decision:

$$\begin{aligned} v^1(\varepsilon_l, k; z_i, \mu) &= \max_{k'} \left[(1 - \nu) [z_i \varepsilon_l]^{\frac{1}{1-\nu}} k^{\frac{\alpha}{1-\nu}} \left(\frac{\nu}{\omega(z_i, \mu)}\right)^{\frac{\nu}{1-\nu}} \right. \\ &\quad \left. - \gamma k' + (1 - \delta)k \right] + \sum_{j=1}^{N_z} \pi_{ij} d_j(z_i, \mu) \sum_{m=1}^{N_\varepsilon} \pi_{lm}^\varepsilon v^1(\varepsilon_m, k'; z_j, \mu'). \end{aligned} \quad (20)$$

The first-order condition is

$$-\gamma + \sum_{j=1}^{N_z} \pi_{ij} d_j(z_i, \mu) \sum_{m=1}^{N_\varepsilon} \pi_{lm}^\varepsilon D_2 v^1(\varepsilon_m, k'; z_j, \mu') = 0.$$

Combining this with the Benveniste-Scheinkman condition below,

$$D_2 v^1(\varepsilon_l, k; z_i, \mu) = \alpha [z_i \varepsilon_l]^{\frac{1}{1-\nu}} k^{\frac{\alpha}{1-\nu} - 1} \left(\frac{\nu}{\omega(z_i, \mu)}\right)^{\frac{\nu}{1-\nu}} + (1 - \delta),$$

we have a stochastic Euler equation for capital:

$$\gamma = \sum_{j=1}^{N_z} \pi_{ij} d_j(z_i, \mu) \sum_{m=1}^{N_\varepsilon} \pi_{lm}^\varepsilon \left(\alpha [z_j \varepsilon_m]^{\frac{1}{1-\nu}} (k')^{\frac{\alpha}{1-\nu} - 1} \left(\frac{\nu}{\omega(z_j, \mu')}\right)^{\frac{\nu}{1-\nu}} + (1 - \delta) \right). \quad (21)$$

Define the following terms:

$$L_0(\varepsilon_l) = \left(\sum_{m=1}^{N_\varepsilon} \pi_{lm}^\varepsilon (\varepsilon_m)^{\frac{1}{1-\nu}} \right)^{\frac{1-\nu}{1-(\alpha+\nu)}} \quad (22)$$

$$L_1(z_i, \mu) = \left(\frac{\gamma - (1 - \delta) \sum_{j=1}^{N_z} \pi_{ij} d_j(z_i, \mu)}{\alpha \sum_{j=1}^{N_z} \pi_{ij} d_j(z_i, \mu) z_j^{\frac{1}{1-\nu}} \left(\frac{\nu}{\omega(z_j, \mu')}\right)^{\frac{\nu}{1-\nu}}} \right)^{\frac{1-\nu}{\alpha+\nu-1}}. \quad (23)$$

Simplification of (21) and use of the definitions in equations (22) - (23) proves the following.

Lemma 1 *The capital decision rule for a plant, $K(\varepsilon_l, k; z_i, \mu)$, is independent of k and takes the form $L_0(\varepsilon_l) L_1(z_i, \mu)$.*

Turning to aggregation, let $H = (h_1, \dots, h_{N_\varepsilon})^T$ be the vector representing the time-invariant distribution of plants over idiosyncratic shock values. Using lemma 1, define the future aggregate capital stock as:

$$K' = \sum_{l=1}^{N_\varepsilon} h_l L_0(\varepsilon_l) L_1(z_i, \mu). \quad (24)$$

Toward establishing a time-invariant distribution of plants over shares of the aggregate stock, it is useful to define the vector of share terms, $\chi = (\chi_1, \dots, \chi_{N_\varepsilon})$, where:

$$\chi_m = \frac{L_0(\varepsilon_m)}{\sum_{l=1}^{N_\varepsilon} h_l L_0(\varepsilon_l)}, \quad m = 1, \dots, N_\varepsilon. \quad (25)$$

While all plants with the same current idiosyncratic shock value will choose a common capital stock for next period, their subsequent idiosyncratic productivities will differ. Let \tilde{H} describe the two-dimensional distribution of plants over ε_{t-1} and ε_t , with elements:

$$\tilde{h}_{l,m} = \pi_{l,m} h_l, \quad \text{for } l = 1, \dots, N_\varepsilon \text{ and } m = 1, \dots, N_\varepsilon. \quad (26)$$

In any period $t+1$, where $t \geq 0$, the distribution of plants is then completely characterized by \tilde{H} and χ together with the aggregate capital stock, K_{t+1} . This establishes lemma 2 below.

Lemma 2 *Define $k_l \equiv \chi_l K$, $l = 1, \dots, N_\varepsilon$. For each ε_m , $m = 1, \dots, N_\varepsilon$, $\mu(\varepsilon_m, k_l) = \tilde{h}_{l,m} \geq 0$, and elsewhere $\mu = 0$.*