

Federal Reserve Bank of Minneapolis
Research Department Staff Report XXX

April 2002

**Nonconvex Factor Adjustments in Equilibrium
Business Cycle Models: Do Nonlinearities Matter?**

Aubhik Khan*

Federal Reserve Bank of Philadelphia

Julia K. Thomas

University of Minnesota

and Federal Reserve Bank of Minneapolis

ABSTRACT

Recent empirical analysis has found nonlinearities to be important in understanding aggregated investment. Using an equilibrium business cycle model, we search for aggregate nonlinearities arising from the introduction of nonconvex capital adjustment costs. We find that, while such costs lead to nontrivial nonlinearities in aggregate investment demand, equilibrium investment is effectively unchanged. Our finding, based on a model in which aggregate fluctuations arise through exogenous changes in total factor productivity, is robust to the introduction of shocks to the relative price of investment goods.

*Research Department, Federal Reserve Bank of Philadelphia, 10 Independence Mall, Philadelphia, PA 19106; aubhik.khan@phil.frb.org. This paper is the result of a conversation with John Leahy; we are grateful to him for suggesting the topic to us. We would also like to thank Ricardo Caballero and Tony Smith for a series of helpful discussions. In addition we owe thanks to seminar participants at the University of California-Riverside, Penn State, the Iowa City Midwest Macro meetings, the Society for Economic Dynamics meetings in San Jose, the 2000 NBER Summer Institute, the Penn Meetings of the Rogerson-Wright NBER group, The Federal Reserve Bank of Minneapolis and, in particular, Marcelle Chauvet, Larry Christiano, Russ Cooper, Martin Eichenbaum, Jonas Fisher, John Haltiwanger, Tim Kehoe and Lee Ohanian. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Banks of Minneapolis or Philadelphia, nor of the Federal Reserve System.

1 Introduction

We evaluate the aggregate implications of discrete and occasional capital adjustment in an equilibrium business cycle model. In our model economy, nonconvex costs of capital adjustment vary across establishments and lead to periods of investment inactivity. Thus, the model generates a distribution of plants over capital. This distribution evolves over the business cycle in response to changes in productivity that affect not only the levels of investment undertaken by investing plants, but also the number of plants actively engaged in adjusting their capital stocks. Our objective is to evaluate the contribution of such distributional changes to the business cycle.

Recent studies of establishment-level investment provide evidence of lumpy capital adjustment. Examining a 17-year sample of large, continuing U.S. manufacturing plants, Doms and Dunne (1999) find that typically more than half of a plant's cumulative investment occurs in a single episode. Long periods of relatively small changes are interrupted by investment spikes. This has been widely interpreted as evidence of (S,s) type investment decisions at the establishment level. Perhaps due to nonconvexities in the costs of capital adjustment, plants invest only when their actual capital stock deviates sufficiently far from a target value. Supporting evidence is provided by Cooper, Haltiwanger and Power (1999), who find that the probability of an establishment undergoing a large investment episode is rising in the time since its last such episode.

Exploring the aggregate implications of establishment-level lumpy investment, Caballero, Engel and Haltiwanger (1995) focus on the effect of interaction between the rising adjustment hazard, the probability of capital adjustment as a function of an establishment's gap between actual and target capital stocks, and the resultant distribution of capital. They argue that shifts in the hazard, in response to large shocks to demand or productivity, magnify fluctuations in aggregate investment demand and cause a time-varying elasticity of aggregate investment with respect to shocks. This emphasis on aggregate nonlinearities arising through micro-level nonconvexities is also found in Cooper, Haltiwanger and Power (1999), who stress that movements in the distribution of capital are important in explaining unusually large deviations in total investment. Further, Caballero and Engel (1999) document skewness and excess kurtosis in aggregate investment rates, as well as a procyclical elasticity

of aggregate investment demand to shocks. Exploring the ability of the lumpy investment model to explain these features of the data, they write “the nonlinear model we estimate has the potential to generate brisker expansions than its linear counterparts. It is also this feature that largely explains its enhanced forecasting properties.” (p. 785, paragraph 1)

These and related papers, all of which abstract from the effect of equilibrium price changes, suggest a potentially important role for lumpy investment in propagating the business cycle.¹ However, when Thomas (2001) evaluates the effects of nonconvex capital adjustment costs in an equilibrium business cycle model, she finds that standard price movements offset the tendency for large changes in the distribution of capital. Solving the model using a system of linear difference equations, she finds that aggregate quantity responses are virtually unaffected by the presence of lumpy investment patterns.²

Noting the above emphasis on aggregate nonlinearities, we re-evaluate the equilibrium lumpy investment model of Thomas (2001) using a solution method designed to preserve such phenomena. Our first step is to fix prices and confirm that the introduction of nonconvex capital adjustment costs does indeed imply aggregate nonlinearities in the model. Next, we explore whether these nonlinearities in aggregate investment demand survive equilibrium price determination. Finally, we analyze the aggregate implications of lumpy investment in the context of an equilibrium business cycle model containing an additional source of cyclical fluctuations. In addition to the conventional exogenous changes in total factor productivity, we allow for movements in the productivity of investment itself. The recent work of Christiano and Fisher (1998) and Greenwood, Hercowitz and Krusell (2000) suggests that such investment-specific productivity shocks are an important source of cyclical fluctuations. Since, in the context of a model of lumpy investment, transitory movements in the benefit from investment expenditures are more likely to shift the adjustment hazard than shocks to total factor productivity, we explore their contribution to the generation of aggregate nonlinearities.

The economies we study involve state vectors that are sufficiently large to make unmodified nonlinear solution methods impractical. Therefore, we approximate the aggregate state vector, which involves a distribution of plants across capital, with a smaller object and solve the model using a method closely related to the approaches of den Haan (1996, 1997) and

¹ See Caballero (1999) for a survey.

² Veracierto (1998), examining investment irreversibilities, finds similar results.

Krusell and Smith (1997, 1998). In our context, the method itself presents information on the importance of changes in the distribution for the overall business cycle. Despite such efforts, our results provide little support for the importance of discrete and occasional investment in shaping model dynamics; the inclusion of lumpy investment does not perceptibly alter the cycle. As a result, the equilibrium model is unable to generate the skewness and excess kurtosis of empirical investment emphasized by Caballero and Engel (1999) and others, and discussed above. Alternatively, when subject to equilibrium, the lumpy investment model fails to generate the brisker expansions that have been emphasized as the basis for its relative empirical success. This finding holds for both the original model of Thomas (2001) and the model with separate shocks to output and investment.

2 The Model

The model, taken from Thomas (2001) extends the basic equilibrium business cycle model to include costs associated with undertaking capital adjustment. To match the observed empirical distribution of investment rates across establishments, we assume a large number of production units, each of which faces time-varying costs of undertaking capital adjustment. Within any period, these costs are fixed at the plant level, being independent of the level of capital adjustment. Given differences in fixed costs across production units, some plants will adjust their capital, while others will not, at any point in time. As a result, there is heterogeneity across production units, and the model is characterized by a distribution of plants over capital.³

At any date, a production unit is identified by its capital stock, k , and its current fixed cost of capital adjustment, $\xi \in [0, B]$. This fixed cost is denominated in hours of labor and drawn from a time-invariant distribution $G(\xi)$ common across plants. Capital and labor, n , are the sole factors of production, and output at the plant is determined by

$$y = zF(k, n),$$

where z is stochastic total factor productivity. For convenience, we assume that productivity follows a Markov Chain, $z \in \{z_1, \dots, z_J\}$, where

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

³ Given that most available data on establishment-level capital adjustment focus on continuing plants, we abstract from entry and exit by assuming a constant unit measure of production units.

and $\sum_{j=1}^J \pi_{ij} = 1$ for each $i = 1, \dots, J$. Note that both z and F are common across plants; the only source of heterogeneity in production arises from differences in plant-level capital.⁴

After production, the plant must decide whether to absorb its current cost, in which event it is able to adjust capital. However, it may avoid this cost by setting investment to 0 and passively allowing capital to depreciate. We denote investment by i and, measuring adjustment costs in units of output using the wage rate ω , summarize the salient features of this choice below.

$i \neq 0,$	cost = $\omega\xi,$	$\gamma k' = (1 - \delta)k + i$
$i = 0,$	cost = 0,	$\gamma k' = (1 - \delta)k$

Throughout the paper, primes indicate one-period-ahead values, and all variables measured in units of output are deflated by the trend level of labor augmenting technological progress. This leads to the familiar introduction of γ in the capital accumulation equation above.⁵

Let capital be defined on $\mathcal{K} \subseteq \mathbf{R}_+$ and let $\mu : \mathcal{B}(\mathcal{K}) \rightarrow [0, 1]$ be a Borel measure that represents the distribution of plants over capital in the current period. The aggregate state of the economy is described by (z, μ) , and the distribution of plants evolves over time according to a mapping, Γ , which varies with the aggregate state of the economy, $\mu' = \Gamma(z, \mu)$. We will define this mapping below.

In addition to the aggregate state, an establishment is affected by its individual level of capital and adjustment cost. Let $v^1(k, \xi; z, \mu)$ represent the expected discounted value of a plant having current capital k and fixed adjustment cost ξ when the aggregate state of the economy is (z, μ) .

We state the dynamic optimization problem for the typical plant using a functional equation, which is defined by (1) and (2) below. 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 $(k; z, \mu)$.

⁴ Additional sources of heterogeneity, for example persistent differences in productivity across plants, are unlikely to contribute to the nonlinearities we isolate in section 5.1, as further explained therein.

⁵ We assume that efficiency units of labor grow at the exogenous rate $\gamma^{1-\theta} - 1$, where θ is capital's share of output; this implies a trend growth rate for output of $\gamma - 1$. Exogenous technological progress requires that next period's capital be measured relative to the efficiency units of labor available at that time. For additional details, see King and Rebelo (1999).

$$v^0(k; z, \mu) \equiv \int_0^B v^1(k, \xi; z, \mu) G(d\xi) \quad (1)$$

Assume that $d_j(z_i, \mu)$ is the discount factor applied by plants to their next period expected discounted value if productivity at that time is z_j and current productivity is z_i . (Except where necessary for clarity, we suppress the index for current productivity below.) Their profit maximization problem, which takes as given the evolution of the distribution of plants over capital, $\mu' = \Gamma(z, \mu)$, is then described by the following functional equation.

$$\begin{aligned} v^1(k, \xi; z, \mu) = & \max_n \left(zF(k, n) - \omega(z, \mu)n + (1 - \delta)k \right. \\ & + \max \left\{ -\xi\omega(z, \mu) + \max_{k'} \left(-\gamma k' + \sum_{j=1}^J \pi_{ij} d_j(z, \mu) v^0(k'; z_j, \mu') \right), \right. \\ & \left. \left. - (1 - \delta)k + \sum_{j=1}^J \pi_{ij} d_j(z, \mu) v^0\left(\frac{(1 - \delta)}{\gamma}k; z_j, \mu'\right) \right\} \right) \end{aligned} \quad (2)$$

Given (k, ξ) and the equilibrium wage rate $\omega(z, \mu)$, the plant chooses current employment n . Next it must choose whether to adjust capital, the value of which is represented by the first term in the internal binary maximum choice above, or avoid its current fixed cost by setting investment to 0. Instead of merely subtracting the cost of investment from current profits, we adopt an equivalent but notationally more convenient approach. The value of nondepreciated capital augments current profits, and the plant is seen to repurchase its entire capital stock each period. Since adjustment costs do not affect the choice of current employment, we denote using $n^f(k; z, \mu)$ the common choice of employment by all type k plants. Further, let $k^f(k, \xi; z, \mu)$ denote the choice of capital for next period by plants of type k with adjustment cost ξ .

The economy is populated by a unit measure of identical households. Households' wealth is held as one-period shares in plants, which we denote using the measure λ . They determine their current consumption, C , hours worked, N , as well as what number of new shares, $\lambda'(k)$, to purchase at price $\rho(k; z, \mu)$. Their lifetime expected utility maximization problem is described below.

$$W(\lambda; z, \mu) = \max_{C, N, \lambda'} \left(U(C, 1 - N) + \beta \sum_{j=1}^J \pi_{ij} W(\lambda'; z_j, \mu') \right) \quad (3)$$

subject to

$$C + \int_{\mathcal{K}} \rho(k; z, \mu) \lambda'(dk) \leq \omega(z, \mu) N + \int_{\mathcal{K}} v^0(k; z, \mu) \lambda(dk)$$

Let $c(\lambda; z, \mu)$ describe their choice of current consumption, $n^h(\lambda; z, \mu)$ their current allocation of time to working and $\Lambda(k, \lambda; z, \mu)$ the quantity of shares they purchase in plants that begin the next period with k units of capital.

A Recursive Competitive Equilibrium is a set of functions

$$\left(\omega, (d_j)_{j=1}^J, \rho, v^1, n^f, k^f, W, c, n^h, \Lambda \right)$$

such that:

1. v^1 satisfies 1 - 2 and (n^f, k^f) are the associated policy functions for plants.
2. W satisfies 3 and (c, n^h, Λ) are the associated policy functions for households.
3. $\Lambda(k', \mu; z, \mu) = \mu'(k') = \int_{\{(k, \xi) \mid k' = k^f(k, \xi; z, \mu)\}} G(d\xi) \mu(dk)$.
4. $n^h(\mu; z, \mu) = \int_{\mathcal{K}} \left(n^f(k; z, \mu) + \int_0^B \xi \mathcal{J} \left(\frac{(1-\delta)}{\gamma} k - k^f(k, \xi; z, \mu) \right) G(d\xi) \right) \mu(dk)$,
where $\mathcal{J}(x) = 0$ if $x = 0$; $\mathcal{J}(x) = 1$ if $x \neq 0$.
5. $c(\mu; z, \mu) = \int_{\mathcal{K}} \int_0^B \left[z F(k, n^f(k; z, \mu)) + (1 - \delta) k - \gamma k^f(k, \xi; z, \mu) \right] G(d\xi) \mu(dk)$

Using C and N , as given by 4 and 5, to now describe the market-clearing values of consumption and hours worked by the household, it is straightforward to show that equilibrium requires $\omega(z, \mu) = \frac{D_2 U(C, 1 - N)}{D_1 U(C, 1 - N)}$ and that $d_j(z, \mu) = \frac{\beta D_1 U(C', 1 - N')}{D_1 U(C, 1 - N)}$. It is then possible to compute equilibrium by solving a single Bellman equation that combines plants' profit maximization problem with the equilibrium implications of household utility maximization. Let p denote the price plants use to value current output, where

$$p(z, \mu) = D_1 U(C, 1 - N), \quad (4)$$

$$\omega(z, \mu) = \frac{D_2 U(C, 1 - N)}{p(z, \mu)}. \quad (5)$$

A reformulation of (2) yields an equivalent description of a plant's dynamic problem. Suppressing the arguments of the price functions,

$$\begin{aligned}
V^1(k, \xi; z, \mu) = & \max_n \left([zF(k, n) - \omega n + (1 - \delta)k] p \right. \\
& + \max \left\{ -\xi \omega p + \max_{k'} \left(-\gamma k' p + \beta \sum_{j=1}^J \pi_{ij} V^0(k'; z_j, \mu') \right) \right. \\
& \left. \left. - (1 - \delta) k p + \beta \sum_{j=1}^J \pi_{ij} V^0\left(\frac{(1 - \delta)}{\gamma} k; z_j, \mu'\right) \right\} \right)
\end{aligned} \tag{6}$$

where

$$V^0(k; z, \mu) \equiv \int_0^B V^1(k, \xi; z, \mu) G(d\xi). \tag{7}$$

Equations 6 and 7 will be the basis of our numerical solution of the economy. This solution exploits several results which we now derive. First, note that plants choose labor $n = n^f(k; z, \mu)$ to solve

$$zD_2F(k, n) = \omega(\mu, z).$$

Next, we examine the capital choice of establishments undertaking active adjustment decisions. Define the gross value of undertaking such capital adjustment, the first choice in the internal binary maximum within (6), as

$$E(z, \mu) \equiv \max_{k'} \left(-\gamma k' p + \beta \sum_{j=1}^J \pi_{ij} V^0(k'; z_j, \mu') \right). \tag{8}$$

Note that the target capital stock solving this maximization problem is independent of both k and ξ . As a result, all plants that actively adjust their capital stock choose a common level of capital for the next period, $k' = k^*(z, \mu)$, which solves the right-hand side of (8). This independence of target capital from current capital implies that the value of adjustment, $E(z, \mu)$, is itself independent of current capital. Referring again to the functional equation in (6), it is now clear that a plant will absorb its fixed cost and adjust only if the net value of achieving the target capital, $E(z, \mu) - \xi \omega p$, exceeds its continuation value under nonadjustment (line three). It follows immediately that a plant of type k will undertake capital adjustment if its fixed adjustment cost, ξ , falls below some k -specific threshold value, $\bar{\xi}_k$.

Let $\widehat{\xi}_k = \widehat{\xi}(k; z, \mu)$ describe that level of ξ , given current k , that leaves a plant indifferent between active capital adjustment and simply allowing its capital stock to depreciate.

$$\begin{aligned} & -p(z, \mu) \widehat{\xi}_k \omega(z, \mu) + E(z, \mu) \\ & = -p(z, \mu) (1 - \delta) k + \beta \sum_{j=1}^J \pi_{ij} V^0 \left(\frac{(1 - \delta)k}{\gamma}; z_j, \mu' \right) \end{aligned} \quad (9)$$

Next, define $\bar{\xi}(k; z, \mu) \equiv \min \left\{ B, \max \left\{ 0, \widehat{\xi}_k \right\} \right\}$ so that $0 \leq \bar{\xi}(k; z, \mu) \leq B$. Plants with adjustment costs at or below $\bar{\xi}_k$ will adjust their capital stock. Thus, plants described by the plant-level state vector $(k, \xi; z, \mu)$ will begin the subsequent period with capital stock given by:

$$k' = k^f(k, \xi; z, \mu) = \begin{cases} k^*(z, \mu) & \text{if } \xi \leq \bar{\xi}(k; z, \mu), \\ \frac{(1 - \delta)k}{\gamma} & \text{if } \xi > \bar{\xi}(k; z, \mu). \end{cases} \quad (10)$$

Given (10), we now explicitly define the evolution of the distribution of plants over capital, $\mu' = \Gamma(z, \mu)$. For $k \in \mathcal{K}$ such that $k \neq k^*(z, \mu)$,

$$\mu'(k) = \left[1 - G \left(\bar{\xi} \left(\frac{\gamma}{1 - \delta} k; z, \mu \right) \right) \right] \mu \left(\frac{\gamma}{1 - \delta} k \right), \quad (11)$$

while for $k \in \mathcal{K}$ such that $k = k^*(z, \mu)$,

$$\begin{aligned} \mu'(k) & = \int_{\mathcal{K}} G(\bar{\xi}(k; z, \mu)) \mu(dk) \\ & \quad + \left[1 - G \left(\bar{\xi} \left(\frac{\gamma}{1 - \delta} k; z, \mu \right) \right) \right] \mu \left(\frac{\gamma}{1 - \delta} k \right). \end{aligned} \quad (12)$$

It then follows that the market-clearing levels of consumption and hours required to determine p and ω using (4) and (5) are given by

$$C = \int_{\mathcal{K}} \left(zF(k, n^f(k; z, \mu)) - G(\bar{\xi}(k; z, \mu)) \left[\gamma k^*(z, \mu) - (1 - \delta)k \right] \right) \mu(dk) \quad (13)$$

$$N = \int_{\mathcal{K}} \left[n^f(k; z, \mu) + \int_0^{\bar{\xi}(k; z, \mu)} \xi G(d\xi) \right] \mu(dk). \quad (14)$$

3 Model Solution

Given our focus on nonlinearities induced by nonconvex adjustment costs, we adapt existing nonlinear solution methods to solve the model. The solution algorithm involves repeated application of the contraction mapping implied by (6) and (7), given the price functions (4) - (5), to solve for V^0 . In implementing this algorithm, we reduce the large state vector of the model to allow for nonlinear approximation of the plant's value function.

We solve nonlinearly for V^0 across a multi-dimensional grid of points, using tensor product splines to interpolate function values at other points. Johnson et al. (1993) have shown this type of multivariate spline approximation to be relatively efficient in comparison with multilinear grid approximation. Our splines are generated as the product of univariate functions, with one of these corresponding to each argument of the value function. Each such univariate function is itself a spline constructed piecewise using a grid of values, or knots, on the space of its argument. Each piece of the spline is a polynomial, and adjacent pieces meet at the interior knot points. We use cubic splines constructed using third-order polynomials, and each univariate spline is determined as follows: (i) the spline is required to exactly equal the approximated function at each knot point, and (ii) it must be twice-continuously differentiable at each interior knot point. Two additional conditions, commonly referred to as endpoint conditions, are required to determine all 4 coefficients of each polynomial piece. We use the not-a-knot endpoint conditions that require thrice differentiability at the first and last interior knot.⁶ In using these tensor product splines, we increase the number of knots used for each variable until there is no noticeable change in the approximation.

A difficulty with using nonlinear methods is that the curse of dimensionality restricts the feasible number of variables; as a result we must reduce the state vector of our economy for computational tractability. We adopt the method of Krusell and Smith (1997,1998) to approximate the state vector of the economy (z, μ) , which contains a large object, the distribution of production units over capital, with a smaller object (z, m) where m is a vector of elements derived from μ . For example, Krusell and Smith use statistical moments derived from the distribution, in particular the mean and standard deviation. For our problem, we find it more efficient to use a set of conditional means. Specifically, when m has I elements, we derive these by partitioning the distribution μ into I equal-measure parts and then setting

⁶ Additional details on univariate splines are available in De Boor (1978) and Van Loan (2000). De Boor also provides details on implementing multivariate splines using the B -form; however, we implement these using the pp -form by developing the algorithm outlined in Johnson (1989).

$m = (m_1, \dots, m_I)$, where m_i is the mean of the i – th partition.

Nonlinear approximation and state space approximation are the two key ingredients of our solution method. We now turn to a brief overview of the details of the method. Given I , we assume functional forms that predict current equilibrium prices, p , and next period's proxy endogenous state, m' , as functions of the current state, $p = \hat{p}(z, m; \chi_l^p)$ and $m' = \hat{\Gamma}(z, m; \chi_l^m)$, where χ_l^p and χ_l^m are parameters that are determined iteratively using a procedure explained below, and l indexes these iterations. For the class of utility functions we use, the wage is immediate once p is specified; hence there is no need to assume a wage function. For any I , \hat{p} , and $\hat{\Gamma}$, we undertake the following steps.

1. The first step uses (χ_l^p, χ_l^m) , having replaced μ with m in (6) - (7), and Γ with $\hat{\Gamma}$, to solve for V^0 at each point on a grid of values for $(k; z; m)$.
2. The second step simulates the economy for T periods, recording, at each point in time, $t = 1, \dots, T$, the actual distribution of plants over capital, μ_t , which is a large but finite-dimensional object in our economy.⁷ To solve the model in each period, we begin by calculating m directly from the actual distribution, and then using $\hat{\Gamma}$ to specify expectations of m' , $m' = \hat{\Gamma}(z, m; \chi_l^m)$. This determines $\beta \sum_{j=1}^J \pi_{ij} V^0(k'; z_j, m')$ for any k' . Thus, given any \tilde{p} , we are now able to solve for $k^*(z, \mu)$ and $\bar{\xi}(k; z, \mu)$, as well as $n^f(k; z, \mu)$ using (8) - (10). The implied levels of consumption and total hours worked are then obtained using equation (13) and (14). The equilibrium price for current output, p , is that which leads to plant decision rules, k^* , n^f and $\bar{\xi}$ that in turn imply market-clearing levels of consumption and hours worked for the household: $p = D_1 U(C, 1 - N)$. (It is important to emphasize that we have not used \hat{p} , but rather have directly solved for equilibrium p here.) Thereafter equations (11) - (12) deliver μ_{t+1} , and we move into the next date of the simulation. After the completion of the simulation, the resulting data, $(p_t, m_t)_{t=1}^T$, are used to re-estimate $(\chi_{l+1}^p, \chi_{l+1}^m)$ using OLS.

We repeat this two-step process, first determining V^0 given (χ_l^p, χ_l^m) , next using our solution for plants' value functions to determine equilibrium decision rules over a simulation, then aggregating these rules to obtain $(p_t, m_t)_{t=1}^T$, and updating χ^p and χ^l , until these parameters converge.

⁷ The method is easily extended to cases where μ is countable or larger using a polynomial approximation.

The simulation step may be used to compute errors implied by the use of the set of conditional means, m , instead of μ , and the functional forms \hat{p} and $\hat{\Gamma}$. In each period, we compare the equilibrium price to the forecasted price and the actual values of the conditional means to their predicted values. Given any functional form, we increase the number of partitions (the number of conditional means used to approximate the distribution of plants over capital) until these differences are small. We also experiment with different functional forms. Below, we report these expectational errors and use them to determine I .

4 Parameter Choices

We compare the lumpy investment economy to an economy without investment frictions that is otherwise identical, the standard business cycle model. This use of the frictionless neoclassical model as a reference model is appealing both due to its common usage in business cycle studies and because it provides a benchmark against which to measure nonlinearities, as it has been shown to respond approximately linearly to reasonable-sized shocks.⁸ Toward our comparison, we specify identical functional forms in utility and production across models. We follow Hansen (1985) and Rogerson (1988) in assuming indivisible labor, so that the representative household's momentary utility function is additively separable and linear in leisure: $u(c, L) = \log c + AL$. Establishment-level production functions take a Cobb-Douglas form, $zF(k, N) = zk^\theta N^\nu$, as consistent with the observation that capital and labor shares of output have remained roughly constant in U.S. time series.

Our solution of each model economy requires the specification of several parameters governing preferences and technology. We fix the length of a period to correspond to one year, allowing us to use evidence on establishment-level investment in the parameterization of the adjustment cost distribution below. Model parameters are selected to ensure agreement between the reference model and observed long-run values for key postwar US aggregates. In particular, we choose the mean growth rate of technological progress to imply a 1.6 percent average annual growth rate of real per capita output, the discount factor, β , to yield an average interest rate of 6.5 percent (King and Rebelo 1999), and the rate of capital depreciation to match an average investment-to-capital ratio of 7.6 percent (Cooley and Prescott 1995). Given these values, capital's share of output is determined such that the

⁸ This follows from Christiano (1990), who shows that the LQ approximation of Kydland and Prescott (1982) is highly accurate for this class of models.

average capital-to-output ratio is 2.6 (Prescott 1986). Labor’s share is consistent with direct estimates from postwar data, while the parameter governing the preference for leisure, A , is taken to imply an average of 20 percent of available time spent in market work (King, Plosser and Rebelo 1988).

To complete our calibration of the reference model, we next estimate parameters for a continuous shock and then assume an equivalent discretized shock process. Specifically, we assume an exogenous productivity process of the form,

$$z' = z^\rho e^{\varepsilon'}; \quad \varepsilon \sim n(0, \sigma_\varepsilon^2),$$

selecting the persistence term ρ and the variability of the log normal innovations, σ_ε , to be consistent with measured Solow residuals from the US economy 1953-1997, using the Stock and Watson (1999) data set. Next, we discretize the productivity process, using a grid of 5 possible shock realizations. We select this grid of values, along with the transition matrix Π (with typical element $\pi_{ij} \equiv \Pr(z' = z_j | z = z_i)$) to match the required shock persistence and variability, following a method developed by Rouwenhorst (1995).

Table 1 and equation (15) summarize the parameter set for the reference model.

$$Z = [.9328 \quad .9658 \quad 1.0000 \quad 1.0354 \quad 1.0720] \tag{15}$$

$$\Pi = \begin{bmatrix} 0.8537 & 0.1377 & 0.0083 & 0.0002 & 0.0000 \\ 0.0344 & 0.8579 & 0.1035 & 0.0042 & 0.0001 \\ 0.0014 & 0.0690 & 0.8593 & 0.0690 & 0.0014 \\ 0.0001 & 0.0042 & 0.1035 & 0.8579 & 0.0344 \\ 0.0000 & 0.0002 & 0.0083 & 0.1377 & 0.8537 \end{bmatrix}$$

As this set of parameters is also used for the lumpy investment model, only the properties of adjustment costs remain to be determined. We assume that adjustment costs are uniformly distributed, with cumulative distribution $G(\xi) = \frac{\xi}{B}$. The distribution’s upper support, B , is selected to maximize the model’s agreement with three results from Doms and Dunne’s (1998) study of establishment-level investment: (i) In the average year, plants raising their real capital stocks by more than 30 percent (lumpy investors) are responsible for 25 percent of aggregate investment, (ii) these lumpy investors constitute 8 percent of plants, while (iii) 80 percent of plants are low-level investors exhibiting annual capital growth below 10

percent. Setting $B = .002$, (a value corresponding to one percent of the average fraction of time devoted to employment,) roughly matches these observations, with lumpy investments comprising 27 percent of aggregate investment, and lumpy investors (low-level investors) representing 6 percent (78 percent) of plants.

We have chosen our baseline value for B to match the data we have on lumpy investment. Nonetheless, in the results that follow, we explore the sensitivity of model results to this choice by considering a ten-fold rise to $B = .02$. In this high cost case, the extent of lumpy investment is observed to be dramatically overstated. Model counterparts for the Doms and Dunne moments above become (i) lumpy investment as proportion of aggregate investment: 85 percent, (ii) lumpy investors as proportion of plants: 12 percent, (iii) low-level investors as proportion of plants: 85 percent.

5 Results

To isolate aggregate nonlinearities that may arise through the introduction of nonconvex adjustment costs, we now compare the behavior of the lumpy investment model with two benchmarks. The first, with no adjustment costs, is the standard equilibrium business cycle model. It will serve as our reference. The second, with convex adjustment costs, has been used extensively in the empirical investment literature. This, the partial adjustment model, assumes adjustment costs of the form $\frac{\phi}{2} \left(\frac{i}{k} - \lambda\right)^2 k$. Here, λ represents the economy's steady state investment-to-capital ratio, and deviations from this average investment rate entail the payment of a quadratic cost of capital adjustment. Following Kiyotaki and West (1996), we set the parameter ϕ governing the magnitude of this quadratic cost at $\phi = 2.2$, which implies a steady-state elasticity of the investment-to-capital ratio to Tobin's marginal q of 5.98. In all other respects, this partial adjustment model is identical to our reference model.

We begin by illustrating the distinguishing feature of the lumpy investment model, the rising adjustment hazard described by Caballero, Engel and Haltiwanger (1995). It is an immediate and important implication of our model that the probability of capital adjustment rises in the gap between target and actual capital. This result follows from the following observations. It is straightforward to show that V^0 is increasing in k , plant-level capital. Consequently,

$$E(z, \mu) + (1 - \delta)kp - \beta \sum_{j=1}^J \pi_{ij} V^0 \left(\frac{(1 - \delta)}{\gamma} k; z_j, \mu' \right)$$

is increasing in $\left| \frac{(1-\delta)}{\gamma}k - k^*(z, \mu) \right|$. In other words, the larger the difference between unadjusted capital and target capital, the greater the value of adjustment. It then follows from (9) that $\bar{\xi}(k; z, \mu)$ is also increasing in the gap between unadjusted and target capital. Hence the probability that a production unit of type k undertakes capital adjustment, $G(\bar{\xi}(k; z, \mu))$, is increasing in its capital deviation, as shown in the upper panel of figure 1. Notice that the hazard is centered at the current capital level that yields the target capital for next period without need of adjustment, $\frac{\gamma k^*(z, \mu)}{1-\delta}$. Probabilities of adjustment monotonically rise as capital deviates to the left or right of this value. For sake of exposition, we temporarily abstract from aggregate uncertainty and assume that the economy is in its steady state. In this case, all plants are positioned along the left ramp of the hazard, given depreciation and trend technological progress, having capital levels at or below that associated with the target. The implication of this is a monotonically rising steady state distribution of plants, as shown by the solid curve of the figure's lower panel. The lower, dashed, curve depicts the measure of plants at each capital level that do not adjust their capital stocks. Thus, the area between represents the steady state measure of adjusting plants, here roughly 30 percent.

5.1 Dynamics under fixed prices

We begin with a series of fixed price experiments designed to gauge lumpy investment's potential for nonlinearities. In these examples, we study aggregate factor demand responses to fluctuations in total factor productivity under the assumption that wages and interest rates faced by the economy's establishments remain fixed at their steady state values. We view this as a useful way of exploring the ability that our model has for producing the sorts of features uncovered by previous partial equilibrium studies, as discussed in section 1 above. Perhaps more importantly, this series of examples illustrates how shifts in the rising adjustment hazard, introduced above, may interact with the underlying distribution of plants to either amplify or dampen the effect of exogenous changes in productivity.

We begin by examining the effects of transitory rise in productivity, continuing to use the steady state as a starting point. The first panel of figure 2 displays an initial adjustment hazard for the lumpy investment model, centered at the capital value associated with steady state target capital. In the face of a one standard deviation rise in productivity that is expected to persist, establishments' target capital increases sharply. This re-centers the adjustment hazard, shifting it rightward. Recall that, on average, most plants are positioned along the left ramp of the hazard, due to capital depreciation and trend productivity growth.

When those plants associated with initial capital holdings below 1.18 experience a rise in target capital, they find that their current capital lies sufficiently far below their (raised) target that they are now willing to suffer large adjustment costs to correct the shortfall. In this particular example, all plants lie along the left ramp of the initial hazard. Here the rise in productivity generates such large rises in the differences between current and target capital that even the highest adjustment cost draw does not dissuade any plant from investing. As a result, the total measure of adjustors rises dramatically from .295 to 1 in the figure's lower panel.

Next, consider the converse, the effects of a one standard deviation drop in productivity, as depicted in figure 3. In this case, the fall in target capital implies a substantial leftward shift in the adjustment hazard. Those plants with very low capital holdings, initially associated with high adjustment probabilities, now find their current capital much closer to the target value and, hence, are less likely to undertake costly adjustment. At the same time, plants with current capital roughly between 1 and 1.2 find that, rather than having a minor capital shortfall, they now have substantial excess. For these plants, adjustment probabilities rise. On balance, the left-shifting adjustment hazard implies only a minor rise in the number of active capital adjustors, from .295 to .308, as depicted in the lower panel of the figure. While positive productivity shocks have the potential to generate substantial external-margin effects on aggregate investment demand, negative productivity shocks do not.

The significance of this distributional asymmetry becomes apparent in figure 4. There, we consider deviations from trend growth rates in response to the positive shock of figure 2, occurring in period 6, followed by 14 periods of average productivity during which the economy resettles, and then the negative shock of figure 3. In panel A, target capital's deviation from steady state behaves roughly symmetrically and matches the approximately linear reference model closely. However, in response to the first shock, the rise in target capital is substantially amplified by a large rise in the measure of investors, as indicated by panel B, where the growth rate of aggregate capital demand under lumpy investment rises roughly 18 percent more than in the reference model. By contrast, when the negative shock occurs, changes along the external margin play only a minor role. There, the fall in target capital is mitigated by the fact that only about 30 percent of establishments actually disinvest to the new target; consequently, the growth rate of aggregate capital demand exhibits less than half the decline seen in the reference model (where all plants disinvest). We conclude from this example that: (i) our model of lumpy investment does have the potential to generate

aggregate nonlinearities; (ii) these nonlinearities may take the form of asymmetric responses to shocks - sharper expansions and dampened contractions - as suggested by the findings of previous authors; (iii) these features result entirely from the asymmetric effects of rightward and leftward shifts in the adjustment hazard upon the total number of adjustors, and hence subsequent distributions of plants; (iv) the dynamics of adjustment along the intensive margin are roughly unaffected by the presence of nonconvexities in plant-level adjustment technologies.

Our next experiment establishes that the powerful distributional effects we have seen driving aggregate nonlinearities in investment demand are not merely the result of the artificial expositional example constructed above. Specifically, we subject all three models to the same realization of exogenous shocks, drawn from our calibrated distribution. Over this 2500 period simulation of each economy, illustrated in figure 5 and table 2, we again hold prices fixed. For each model's investment series we compute deviations relative to the simulation mean. In the upper panels of figure 5, we display the distribution of these aggregate investment deviations for the reference and lumpy investment models, respectively. The horizontal axes represent 5 broad categories of investment episodes, ranging from extremely low to extremely high. (With prices fixed, aggregate investment can vary as much as five-fold relative to its mean.) The vertical axis measures the fraction of dates spent in each of these ranges.

For the reference model, in the upper left panel, the distribution of investment deviations is symmetric. By contrast, the lumpy investment economy displays a disproportionate fraction of extremely high, relative to high, investment episodes and has fewer extremely low, relative to low, observations. Specifically, while times of near-average investment occur with roughly equal frequency, the figure's lower right panel reveals that the inclusion of nonconvex capital adjustments shifts 2.5 percent of very low investment realizations upward into the low range, while nearly 2 percent (50 periods) of high investment episodes are pushed into the extremely high range. At lower left, we align these histograms alongside results for the partial adjustment model, which by comparison displays far less dispersion in investment, given the convexifying force of quadratic adjustment costs.

For a more detailed examination of the fixed price simulation results, we next undertake a date-by-date comparison between the lumpy investment and reference models. Specifically, we construct a series representing the differences between the two models' investment deviations at each date in the simulation. The first row of table 2 summarizes the notable

features of this series, while the second row presents the analogous results for the partial adjustment versus reference models. Note that average differences from the reference economy are substantial, 46 percent for lumpy investment and 69 percent for partial adjustment.

We close our preliminary examination of the lumpy investment model under fixed prices with two observations. First, table 2 and figure 5 clearly indicate that the fixed costs calibrated in the previous section are more than sufficient to generate substantive differences between the lumpy investment and reference models. Second, having demonstrated the mechanism that yields the model’s potential for nonlinear behavior, we are now able to discuss our decision to abstract from additional sources of heterogeneity. Certainly the abstraction facilitates model solution. However, as the example of figures 2 - 4 indicates, it is also likely to bias our results toward larger nonlinearities than would otherwise occur. If, for example, there were idiosyncratic plant-specific differences in productivity, there would be an adjustment hazard and target capital associated with each level of productivity. Moreover, since generally there would be some plants distributed on the right half of these hazards, the asymmetry described here would be dampened though not, given technological progress and depreciation, eliminated. In this respect, our results in the next section may be interpreted as providing an upper bound to the potential for equilibrium nonlinearities in this class of model.

5.2 Equilibrium dynamics

We begin this section with a discussion of the accuracy of the forecasting rules, $\hat{\Gamma}$ and \hat{p} , used by agents in the model solution. Table 3 displays the equilibrium forecasting functions when the distribution is approximated using only a single partition.⁹ (Here we refer to I , the number of elements in m , as the number of partitions.) We condition the forecasts on current productivity; thus, as our discretization of stochastic productivity assumed 5 values, there are 5 regressions. The standard errors and R^2 s associated with each regression indicate that the statistical mean alone is an efficient proxy for the distribution. This is confirmed in table 4, where we re-solve the economy using two partitions to approximate the distribution and find only marginal reductions in the standard errors on equilibrium price regressions, indicating little additional relevant information. It is difficult to draw

⁹ We have experimented with a variety of functional forms, including, for example, higher order terms. These produce similar results to the log linear form reported here. In the extension of the model, in section 6 below, we use quadratic forms.

inferences from the relative magnitudes of the errors in forecasting future conditional means, as neither m'_1 nor m'_2 in table 4 corresponds to the mean in table 3. Instead, we use figure 6 to present the aggregate capital series from each lumpy investment economy over the same 2500 period history. We find no discernible difference; the maximum difference in the two series is 2.1×10^{-4} . Finally, to further emphasize plants' forecasting ability in using only the single mean of the distribution, we compare the results of table 3 with corresponding results from an economy whose distribution is *exactly* its mean. Specifically, when we solve for equilibrium forecasting functions in the reference economy, we find minimal changes in the regression coefficients and standard errors. As an illustration of this, the reference economy's standard errors for m'_1 and p are 1.22×10^{-4} and 2.63×10^{-5} , respectively, when productivity is at its highest value z_5 , and 2.49×10^{-4} and 5.36×10^{-5} for $z = z_3$. Comparing these values with the corresponding errors of table 3 foreshadows the remaining results of this section. In any case, we take these findings as strong evidence that we need not partition the distribution further. Thus, except where explicitly noted otherwise, the lumpy investment results below correspond to the 2- partition economy.

We now re-examine the productivity simulation of our fixed price experiments in general equilibrium. We begin with an overview of second moments in table 5.¹⁰ Panel A displays percentage standard deviations in the growth rates of output, investment, consumption, employment, wages and interest rates across model economies. From these results, it is clear that the variability under lumpy investment is virtually identical to the reference economy, regardless of whether we use one partition of the distribution (row 3) or two (row 2). This similarity is further emphasized by comparison with the partial adjustment model, where the cycle is dampened by sluggish responsiveness of investment demand. The similarities between Lumpy Investment and Reference economies are also evident in the comovements with output reported in panel B. By contrast, aggregate quantities move more closely with the cycle in the partial adjustment results.

Higher adjustment costs do not induce more investment volatility in the lumpy investment model. In the final row of table 5, we explore the robustness of its similarities to the reference by raising the magnitude of the adjustment cost's support ten-fold above the value consistent with the empirical observations discussed in section 4, to $B = .02$. We have already noted that this value leads to a dramatically overstated proportion of lumpy investment as a share

¹⁰ For tables 5-7, simulated data are logged (with exception of interest rates) and HP-filtered using a weight of 100.

of aggregate investment. Nonetheless, in table 5 we see that the results remain close to the reference. Moreover, the high cost parameterization dampens the cycle slightly, shifting it in the direction of the partial adjustment model. This may be explained as follows. First, as the larger upper support corresponds to 10 percent of the average time devoted to labor supply, the high B induces much longer average delays between investments at the plant level, thus exaggerating the presence of lumpy investment. (Recall that, while Doms and Dunne (1998) find that lumpy investment is 25 percent of total, it is 85 percent on average in this model with high costs.) More lumpy investment is associated with a flattening of the average adjustment hazard; with this, shifts in the hazard imply smaller changes in the number of adjusting plants, reducing the effects of changes in target capital.

From table 5, it is evident that lumpy investment fails to reshape the aggregate cycle in equilibrium. We explore this further, in figure 7, by presenting histograms of the relative deviations in investment from trend, the equilibrium counterpart to figure 5. Two features of this figure are noteworthy. First, investment in both the lumpy investment and reference economies exhibits far less dispersion than was evident in figure 5, as changes in factor prices largely offset the swings in investment demand seen under fixed prices. Second, while the reference economy's investment series continues to be approximately symmetric around zero, the distribution is now closer to the Normal. Here again, price movements offset plants' desires for large capital adjustments, shifting substantial mass away from extreme investment episodes inward toward more moderate changes. This same force removes the lumpy investment economy's tendency for sharp expansions, shifting mass from the highest investment deviations downward. As a result, the differences in these two histograms essentially disappear in equilibrium; the largest difference is in the zero band, where the lumpy investment economy displays about 0.5 percent fewer realizations than the reference economy.

From the results presented so far, it is apparent that lumpy investment does not produce the stronger expansions and dampened recessions suggested by the fixed price results of figure 5, at least on average. Table 6, the equilibrium version of table 2, confirms that the differences between the Lumpy Investment and Reference investment series are never of quantitative significance; the maximum difference at any date is 0.3 percent. We also see that the gaps present in the second row are reduced when price changes are present to dampen fluctuations in the benchmark investment series.

Based on the discussion above, it would appear that changes in extensive-margin capital adjustment within the lumpy investment economy must be minor in equilibrium. After all,

it was such changes that distinguished the lumpy investment model from the reference under fixed prices. Perhaps surprisingly, while the aggregate cyclical behavior of the nonconvex adjustment cost model is essentially identical to that of the reference model, this does not imply a lack of movement in the distribution of plants. The fraction of plants engaging in capital adjustment, 0.295 in the steady state, is strongly procyclical. Isolating, as above, cyclical components using the Hodrick-Prescott filter with a weight of 100, this series has a percentage standard deviation of 4.25, more than twice that of output. Furthermore, the contemporaneous correlation of the adjustment rate with output is 0.88, and with investment it is 0.96. In equilibrium, there are changes in both the measure of adjusting plants and their capital targets. It is this interplay between the extensive and intensive margins of capital adjustment that allows an approximate reproduction of the aggregate dynamics of the reference economy.

Changes in adjustment rates are important in our model. However, since production units are owned by the household, these changes act to reduce consumption volatility to the level of the reference economy.¹¹ They do not generate brisker expansions. In figure 8, we reconsider the asymmetric shock history that illustrated lumpy investment’s potential for nonlinearities when real wages and interest rates were constant. We find that, even in this example, the equilibrium lumpy investment economy exhibits no greater evidence of an asymmetric response than does the approximately linear reference economy.

6 Investment Shocks

The results of the preceding section indicate that nonconvex capital adjustments can generate large nonlinearities in an environment with unchanging prices, but fail to do so when markets clear. It is tempting, then, to conclude that lumpy investment is not particularly important to the business cycle. However, this conclusion may rely on the assumption that business cycles are generated by a single driving force - an aggregate productivity shock that affects all production units in the economy.

Recent work by Christiano and Fisher (1998) and Greenwood, Hercowitz and Krusell (2000) suggests that, in fact, fluctuations in the price of investment goods may explain a substantial portion of the business cycle. Identifying the relative price of new equipment as a measure of the price of investment goods, Greenwood et al. present evidence that shocks

¹¹ Consumption fluctuates considerably more in the partial adjustment model. See table 5.

shifting the price of investment above and below its long-run downward trend can account for 30 percent of the cyclical variation in output. Measuring investment good prices more broadly, Christiano and Fisher find that investment-specific shocks explain 75 percent of output fluctuations at business cycle frequencies.¹²

The importance of investment-specific shocks in the economy raises questions about the generality of our results in the previous section. We reason as follows. All plants benefit from the effects of a positive total factor productivity shock, regardless of whether they expand their factors of production; to better exploit these benefits, some plants increase capital. By contrast, a positive investment-specific shock provides a more direct incentive for capital adjustment, since it benefits only those establishments that invest. Thus such shocks have the potential to yield much larger shifts in the economy’s adjustment hazard, perhaps sufficient to overcome the convexifying forces of equilibrium. To explore this possibility, we now extend our previous description of the lumpy investment model (and the reference model) to allow for exogenous fluctuations in the productivity of investment.

Our extension of the model is related to the approaches taken by Christiano and Fisher (1998) and Greenwood et al. (2000) and involves the following modifications to our previous specification. We assume that investment-specific productivity follows a first order Markov process with average growth rate $\mathcal{G} - 1$. Plant-level capital accumulation is now governed by

$$\chi k' = (1 - \delta)k + \zeta i,$$

where ζ denotes the current level of detrended investment-specific productivity, and $\chi - 1$ denotes the long-run growth rate of aggregate capital, which is $\gamma \mathcal{G}^{\frac{1}{1-\theta}} - 1$.¹³ The exogenous aggregate state is given by (z, ζ) , and we follow the previous authors in assuming that shocks to total factor productivity and investment-specific productivity are independently distributed. Transition probabilities are

$$\begin{aligned} \pi_{ij} &\equiv \Pr(z' = z_j \mid z = z_i) \\ \tau_{ls} &= \Pr(\zeta' = \zeta_s \mid \zeta = \zeta_l). \end{aligned}$$

¹² Fluctuations in the price of investment goods may be interpreted as the result of shocks to the productivity of investment, or *investment-specific* technology shocks. Throughout this section, we follow this interpretation.

¹³ As before, all variables denominated in units of output are growth-deflated. With the inclusion of investment-specific productivity growth, trend output now grows at rate $\gamma \mathcal{G}^{\frac{\theta}{1-\theta}} - 1$, rather than $\gamma - 1$. We recalibrate the household’s discount factor β to maintain the steady state interest rate at 6.5 percent.

With these alterations, equations (6) and (7) describing the plant's dynamic problem in section 2 change as shown below.

$$\begin{aligned}
V^1(k, \xi; z, \zeta, \mu) &= \max_n \left(\left[zF(k, n) - \omega n + \frac{(1-\delta)k}{\zeta} \right] p \right. \\
&+ \max \left\{ -\xi \omega p + \max_{k'} \left(-\frac{\chi k'}{\zeta} p + \beta \sum_{j,s} \pi_{ij} \tau_{ls} V^0(k'; z_j, \zeta_s, \mu') \right) \right. \\
&\left. \left. - \frac{(1-\delta)k}{\zeta} p + \beta \sum_{j,s} \pi_{ij} \tau_{ls} V^0 \left(\frac{(1-\delta)k}{\chi}; z_j, \zeta_s, \mu' \right) \right\} \right) \quad (16)
\end{aligned}$$

$$V^0(k; z, \zeta, \mu) \equiv \int_0^B V^1(k, \xi; z, \zeta, \mu) G(d\xi). \quad (17)$$

Equations (8) and (9), determining the target capital and threshold adjustment costs, change accordingly, z being replaced with (z, ζ) , and χ replacing γ , and the evolution of the distribution of plants over capital $\mu' = \Gamma(z, \zeta, \mu)$ is given by the following.

$$\mu'(k) = \left[1 - G \left(\bar{\xi} \left(\frac{\chi}{1-\delta} k; z, \zeta, \mu \right) \right) \right] \mu \left(\frac{\chi}{1-\delta} k \right); \quad k \neq k^*(z, \zeta, \mu) \quad (18)$$

$$\begin{aligned}
\mu'(k) &= \int_{\mathcal{K}} G(\bar{\xi}(k; z, \zeta, \mu)) \mu(dk) \\
&+ \left[1 - G \left(\bar{\xi} \left(\frac{\chi}{1-\delta} k; z, \zeta, \mu \right) \right) \right] \mu \left(\frac{\chi}{1-\delta} k \right); \quad k = k^*(z, \zeta, \mu)
\end{aligned} \quad (19)$$

The equations describing equilibrium consumption and hours, (13) - (14), are similarly modified. Finally, we assume that shocks to the investment productivity process take the form

$$\zeta' = \zeta^{\rho_\zeta} e^{\varepsilon'_\zeta}; \quad \varepsilon_\zeta \sim n(0, \sigma_{\varepsilon_\zeta}^2), \quad (20)$$

then discretize the exogenous state space (z, ζ) on a 3×3 grid of values, using the procedure outlined above in section 4.

Note that the growth-deflated relative price of a unit of investment is simply $\frac{1}{\zeta}$. As such, we are able to use a relative price series for aggregate investment, based on US data from 1982 through 1998, to directly estimate the parameters of (20), ρ_ζ and $\sigma_{\varepsilon_\zeta}$, as well as the trend parameter \mathcal{G} .¹⁴ This yields $\mathcal{G} = 1.022$, $\hat{\rho}_\zeta = .706$, and $\hat{\sigma}_{\varepsilon_\zeta} = .017$.

¹⁴ We follow Christiano and Fisher (1998), section 2.2, closely in constructing this price series, adapting their method only as required to translate the quarterly series to an annual frequency.

We now examine the effect of lumpy investment in an economy subject to both fluctuations in aggregate productivity and investment-specific shocks.¹⁵ In table 7, note that adding the investment shock raises overall volatility in both the lumpy investment and reference economies. This is particularly true for investment and employment and consequently interest rates, whose relative standard deviations rise substantially, as is consistent with our reasoning above. We do find a slightly more pronounced difference between the two economies; however, the results continue to be quite close. The same is true for the output correlations shown in the table’s lower panel. Next, we examine the histograms of investment deviations relative to simulation mean in figure 9. In comparison to figure 7, where only the TFP shock was present, there are no more pronounced differences with the inclusion of the investment shock, as is evident from the scale of the lower right panel.

We have hypothesized that the investment shock is more likely to affect the distribution of plants than does the total productivity shock. Given this, it is possible that the latter is acting to dampen nonlinearities. In table 8 and figure 10, we explore this possibility by eliminating variation in total factor productivity. Beginning with the table, it is apparent that the investment shock alone is insufficient to drive the cycle; output variability is reduced nearly half relative to the results of table 5. However, now relative volatilities in investment and employment are at their highest; in the case of employment, the rise is dramatic. In both panels, we begin to see larger differences between R and L rows, particularly for investment, as expected. Nonetheless, these differences remain negligible. It is worthy of note that, beyond the difficulty of reduced output volatility in both models, the consumption, wage and interest rate series have become countercyclical in absence of the TFP shock. Hence, while this example may be useful in studying the aggregate effects of lumpy investment, it is not a plausible model for business cycle analysis.¹⁶

Finally, examining figure 10, we find that the histograms for investment deviations do exhibit greater differences when the productivity shock is removed. This is clearest when they are viewed together in the lower left panel and confirmed by the differences plotted at lower right. Nonetheless, the variations across the lumpy investment versus reference economies continue to be small, with only about 2.8 percent more dispersion away from near-average

¹⁵ As standard errors in forecasting regressions continue to be small, averaging roughly 6×10^{-4} for m_1 and 8×10^{-4} for p , we present results only for the $I = 1$ partition economy.

¹⁶ Christiano and Fisher (1998) avoid these problems by allowing for two sectors in the economy and assuming that labor input must be determined prior to the investment shock’s realization.

investment episodes in the former than the latter. From this and the previous set of results, we conclude that the conjecture that prompted our inclusion of an investment-specific driving process was correct but quantitatively irrelevant.

7 Concluding Remarks

In the preceding sections, we solved an equilibrium business cycle model where, at the individual level, investment is subject to nonconvex costs of adjustment. Calibrating these costs to reproduce the empirical regularities found in establishment-level capital adjustment data, we found that when the sole source of cyclical fluctuation is changes in total factor productivity, and when wages and interest rates are held constant, the lumpy investment model exhibits nonlinearities in aggregate investment demand that drive sharper expansions and dampened recessions relative to a reference model without adjustment costs. However, in equilibrium, the cyclical behavior of the lumpy investment model is remarkably close to that of this reference model. Moreover, additional sources for the business cycle, investment-specific productivity shocks, fail to deliver interesting differences, qualitatively or quantitatively, between the lumpy investment and the reference models.

In developing a business cycle economy characterized by lumpy microeconomic investment, we have generalized production but, to allow the clearest comparison with the standard model, we have maintained the assumptions of a representative household, complete markets and perfect competition. Relaxing some of these assumptions to reduce the influence of the household in determining equilibrium allocations, for example, through the introduction of imperfect competition, may allow for larger aggregate effects of lumpy investment in equilibrium. Alternatively, one might expect that the changes lumpy investment implies for aggregate investment demand would have a larger equilibrium impact in an open-economy setting. As an extreme example, consider a small open economy facing a perfectly elastic supply of savings at an exogenous world interest rate. Under this assumption, comparisons between the lumpy investment economy and the reference would correspond much more closely to those of our fixed price examples than to the results seen here. More generally, lumpy investment could have more important effects in an intermediate setting where the link between savings and investment is relaxed but not removed. It is well-known that, absent additional frictions, simple two-country equilibrium business cycle models suffer from dramatically excessive volatility in investment. While one common remedy for this is the

assumption of convex capital adjustment costs, the inclusion of lumpy investment could represent an alternative means of smoothing the series.

References

- Caballero, R. J., 1999, Aggregate investment, in: M. Woodford and J. Taylor, eds., *Handbook of Macroeconomics IB*. (Elsevier Science, Amsterdam), 813 - 862.
- 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.
- 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, 1-54.
- Christiano, L. J., 1990, Linear-quadratic approximation and value-function iteration: A comparison, *Journal of Business and Economic Statistics* 8, 99-113.
- Christiano, L. J. and J. D. M. Fisher, 1998, Stock market and investment good prices: Implications for macroeconomics, Working paper (Federal Reserve Bank of Chicago, Chicago, IL).
- Cooley, T. F. and E. C. Prescott, 1995, Economic growth and business cycles, in: T. F. Cooley, ed., *Frontiers of business cycle research* (Princeton University Press, Princeton, NJ), 1-38.
- 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.
- De Boor, C., 1978, *A practical guide to splines*. (Springer-Verlag, New York, NY).
- Doms, M. and T. Dunne, 1998, Capital adjustment patterns in manufacturing plants, *Review of Economic Dynamics* 1, 409-430.
- Greenwood, J., Z. Hercowitz and P. Krusell, 2000, The role of investment-specific technological change in the business cycle, *European Economic Review* 44, 91-115.
- den Haan, W.J., 1996, Heterogeneity, aggregate uncertainty, and the short-term interest rate, *Journal of Business and Economic Statistics* 14, 399-411.
- den Haan, W.J., 1997, Solving dynamic models with aggregate shocks and heterogeneous agents, *Macroeconomic Dynamics* 1, 355-386.
- Hansen, G. D., 1985, Indivisible labor and the business cycle, *Journal of Monetary Economics* 16, 309-327.
- Johnson, S. A., 1989, Spline approximation in discrete dynamic programming with application to stochastic multi-reservoir systems, Unpublished dissertation (Cornell, Ithaca, NY).
- Johnson, S. A., J. R. Stedinger, C.A. Shoemaker, Y. Li, and J.A. Tejada-Guibert, 1993, Numerical solution of continuous-state dynamic programs using linear and spline interpolation, *Operations Research* 41, 484-500.

- King, R. G., C. I. Plosser and S. T. Rebelo, 1988, Production, growth and business cycles: I. The basic neoclassical model, *Journal of Monetary Economics* 21, 195-232.
- King, R. G. and S. T. Rebelo, 1999, Resuscitating real business cycles, in: M. Woodford and J. Taylor, eds., *Handbook of Macroeconomics IB* (Elsevier Science, Amsterdam), 927 - 1007.
- Kiyotaki, N. and K. D. West, 1996, Business fixed investment and the recent business cycle in Japan, in: B. S. Bernanke and J. J. Rotemberg, eds., *NBER Macroeconomics Annual*, 277-323.
- Krusell, P. and A.A. Smith Jr., 1997, Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns, *Macroeconomic Dynamics* 1, 387-422.
- Krusell, P. and A.A. Smith Jr., 1998, Income and wealth heterogeneity in the macroeconomy, *Journal of Political Economy* 106, 867-896.
- Kydland, F. and E. C. Prescott, 1982, Time to build and aggregate fluctuations, *Econometrica* 50, 1345 - 1370.
- Prescott, E. C., 1986, Theory ahead of business cycle measurement, *Quarterly Review* 10, Federal Reserve Bank of Minneapolis, 9-22.
- Rogerson, R., 1988, Indivisible labor, lotteries and equilibrium, *Journal of Monetary Economics* 21, 3-16.
- Rouwenhorst, G. K., 1995, Asset pricing implications of equilibrium business cycle models, in: T. F. Cooley, ed., *Frontiers of business cycle research*. (Princeton University Press, Princeton, NJ), 294-330.
- Stock, J. H. and M. W. Watson, 1999, Business cycle fluctuations in U.S. macroeconomic time series, in: M. Woodford and J. Taylor, eds., *Handbook of Macroeconomics IA* (Elsevier Science, Amsterdam), 3-64.
- Thomas, J., 2002, Is lumpy investment relevant for the business cycle?, *Journal of Political Economy*, forthcoming.
- Van Loan, C. F., 2000, *Introduction to scientific computing: A matrix-vector approach using MATLAB*® 2nd ed. (Prentice Hall, Upper Saddle River, NJ).
- Veracierto, M., 1998, Plant level irreversible investment and equilibrium business cycles, Working paper (Federal Reserve Bank of Chicago, Chicago, IL).

Figure 1A: The Stationary Adjustment Hazard

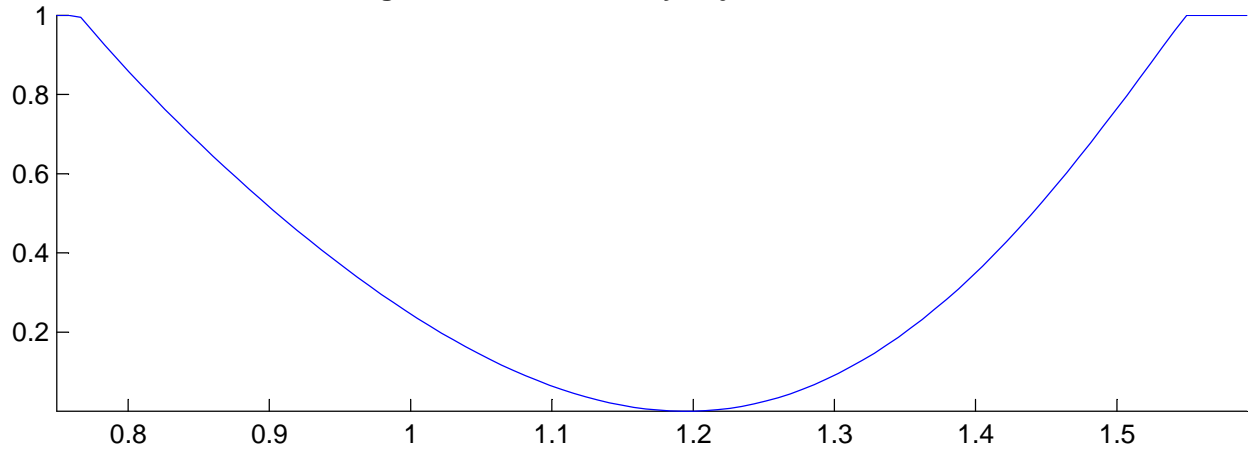


Figure 1B: Adjustment Across the Distribution (total = 0.295)

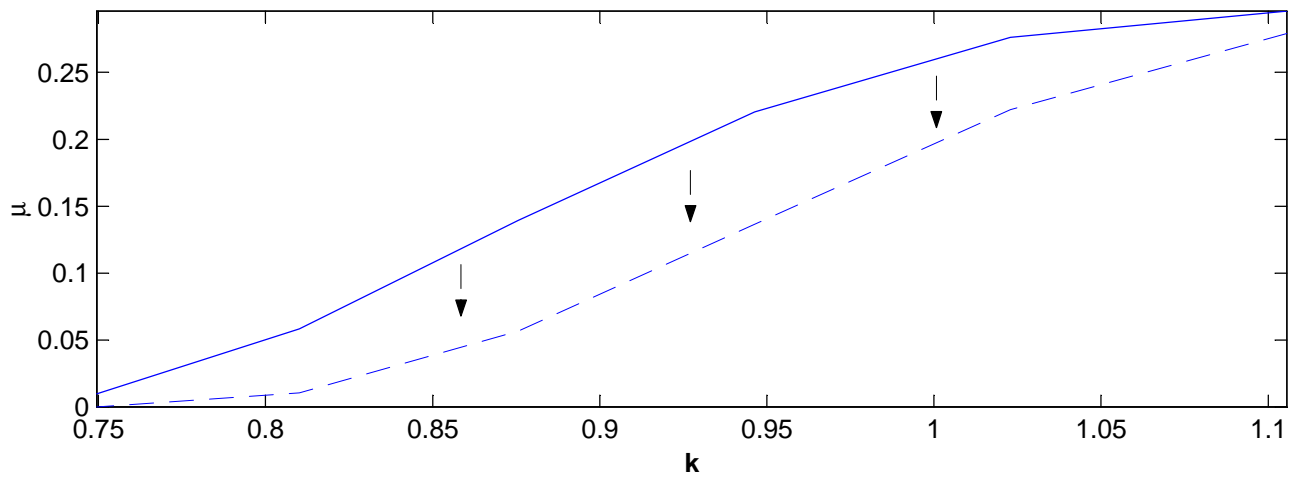


Figure 2A: Rise in Productivity: Hazard Shifts Right

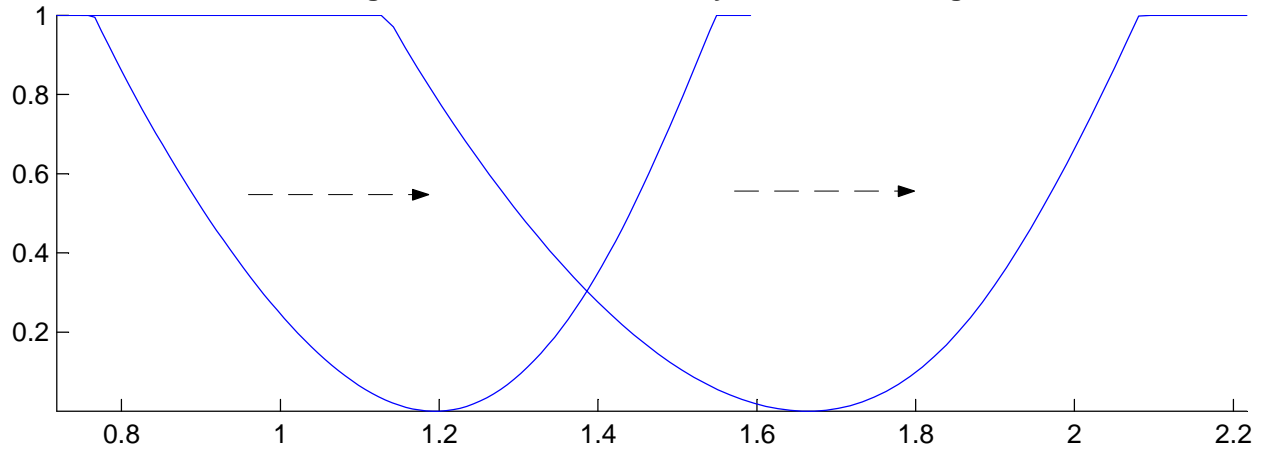


Figure 2B: Adjustment across the Distribution (total = 1)




Figure 3A: Fall in Productivity: Hazard Shifts Left

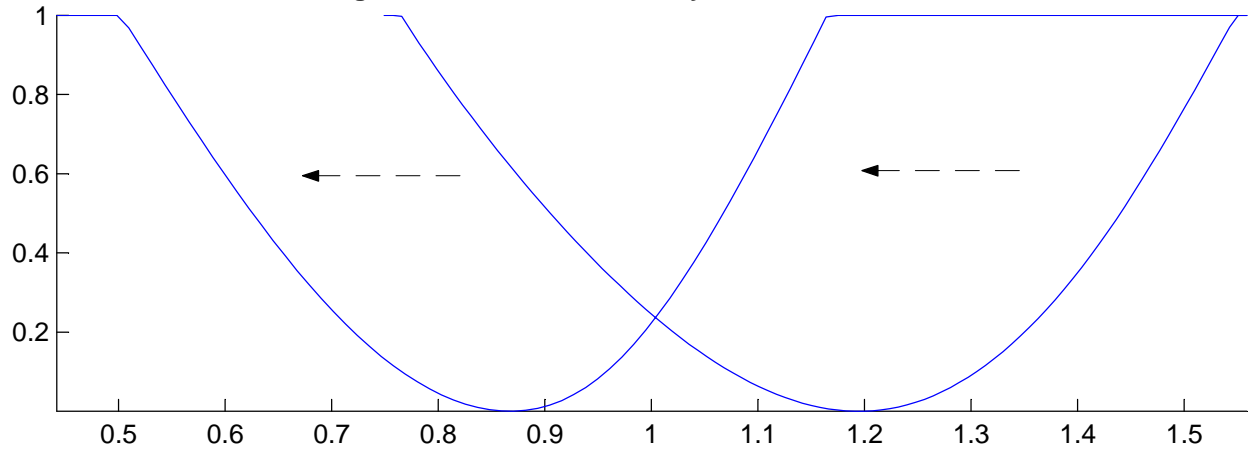


Figure 3B: Adjustment across the Distribution (total = 0.308)

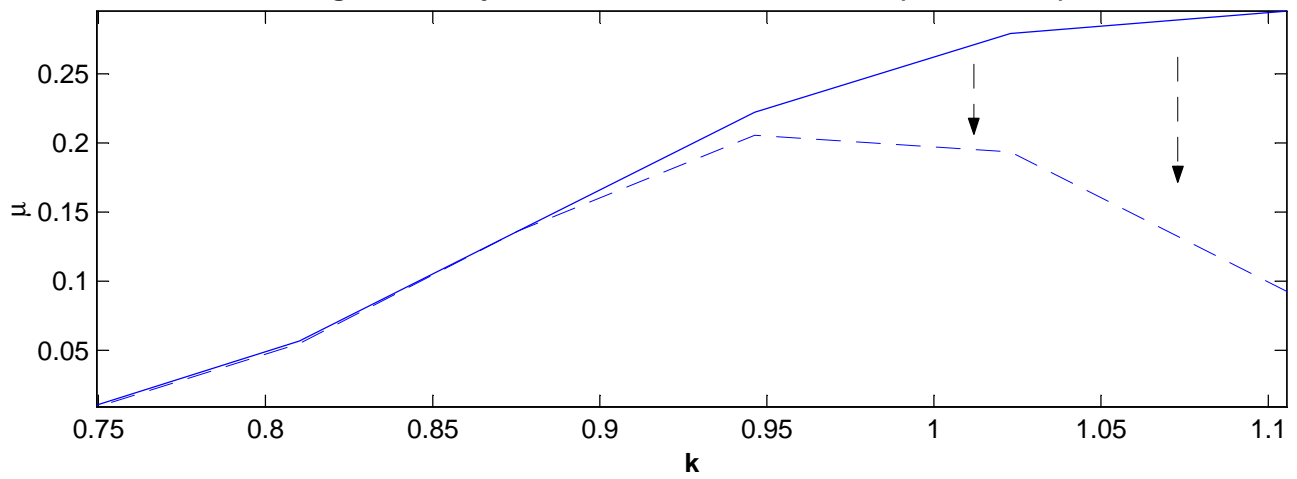


Figure 4A: Target Capital in Reference and Lumpy Inv. models




Figure 4B: Aggregate Capital in Reference and Lumpy Inv. models




Figure 5: Distribution of Investment Deviations

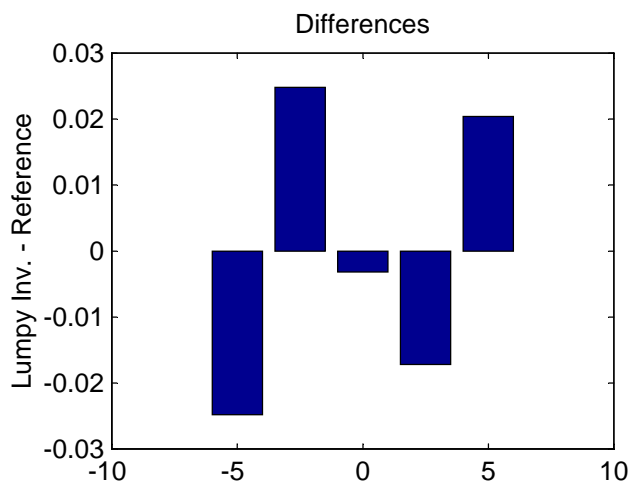
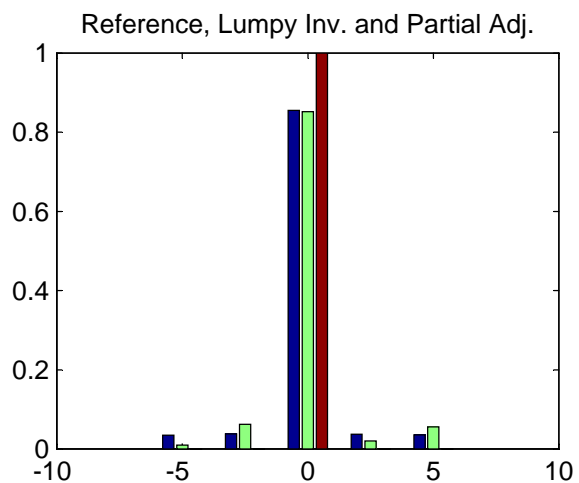
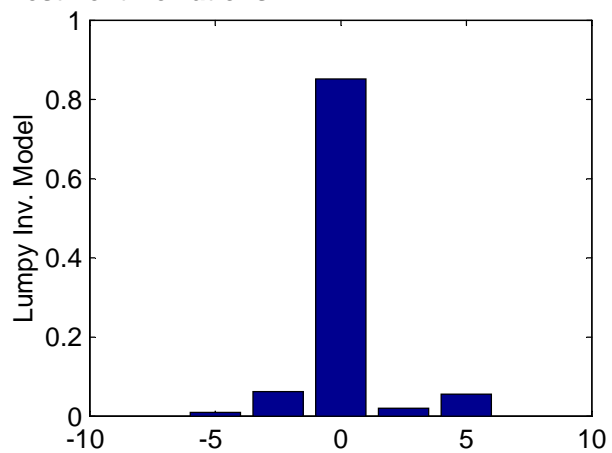
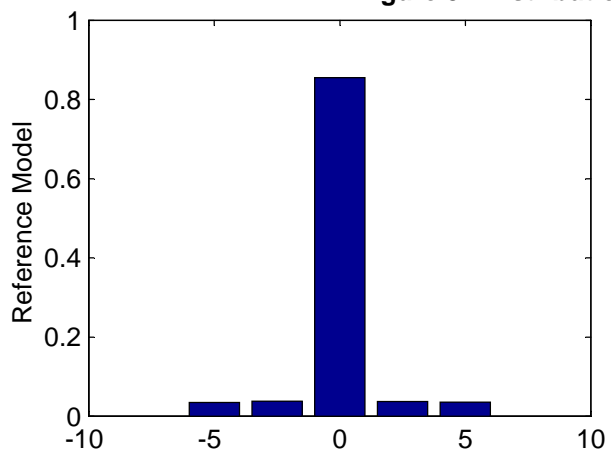


Figure 6: Aggregate Capital in 1 Partition v. 2 Partition Lumpy Inv. Models




Figure 7: Distribution of Equilibrium Investment Deviations

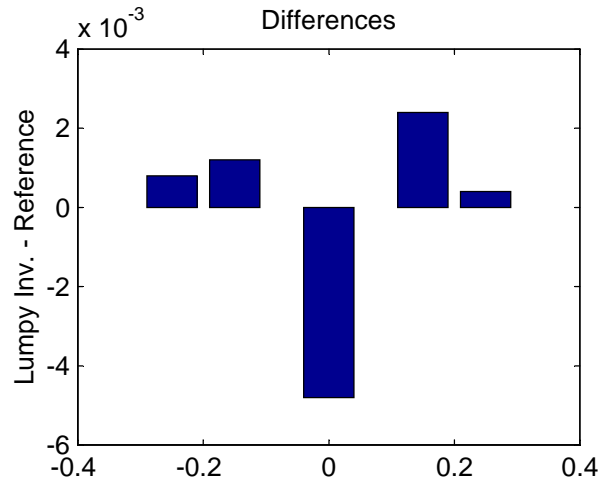
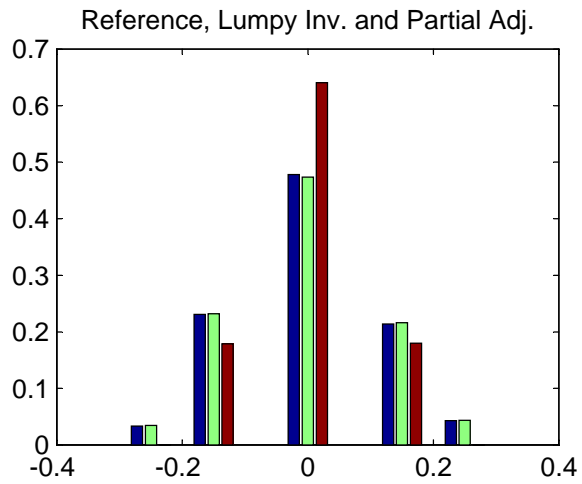
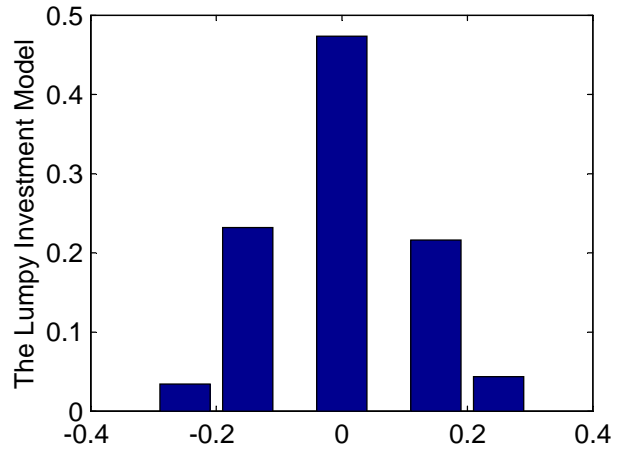
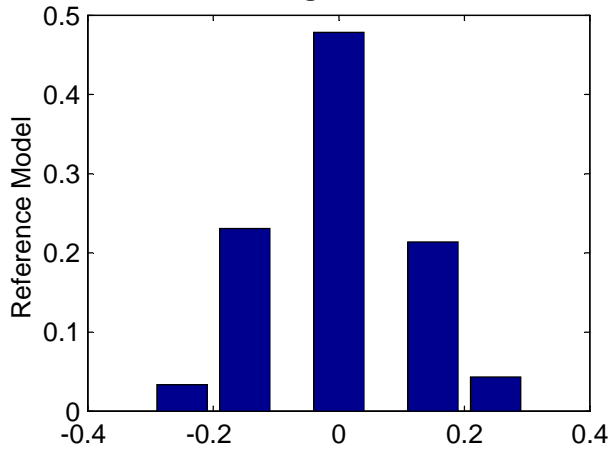


Figure 8: The Absence of Nonlinearities in Equilibrium

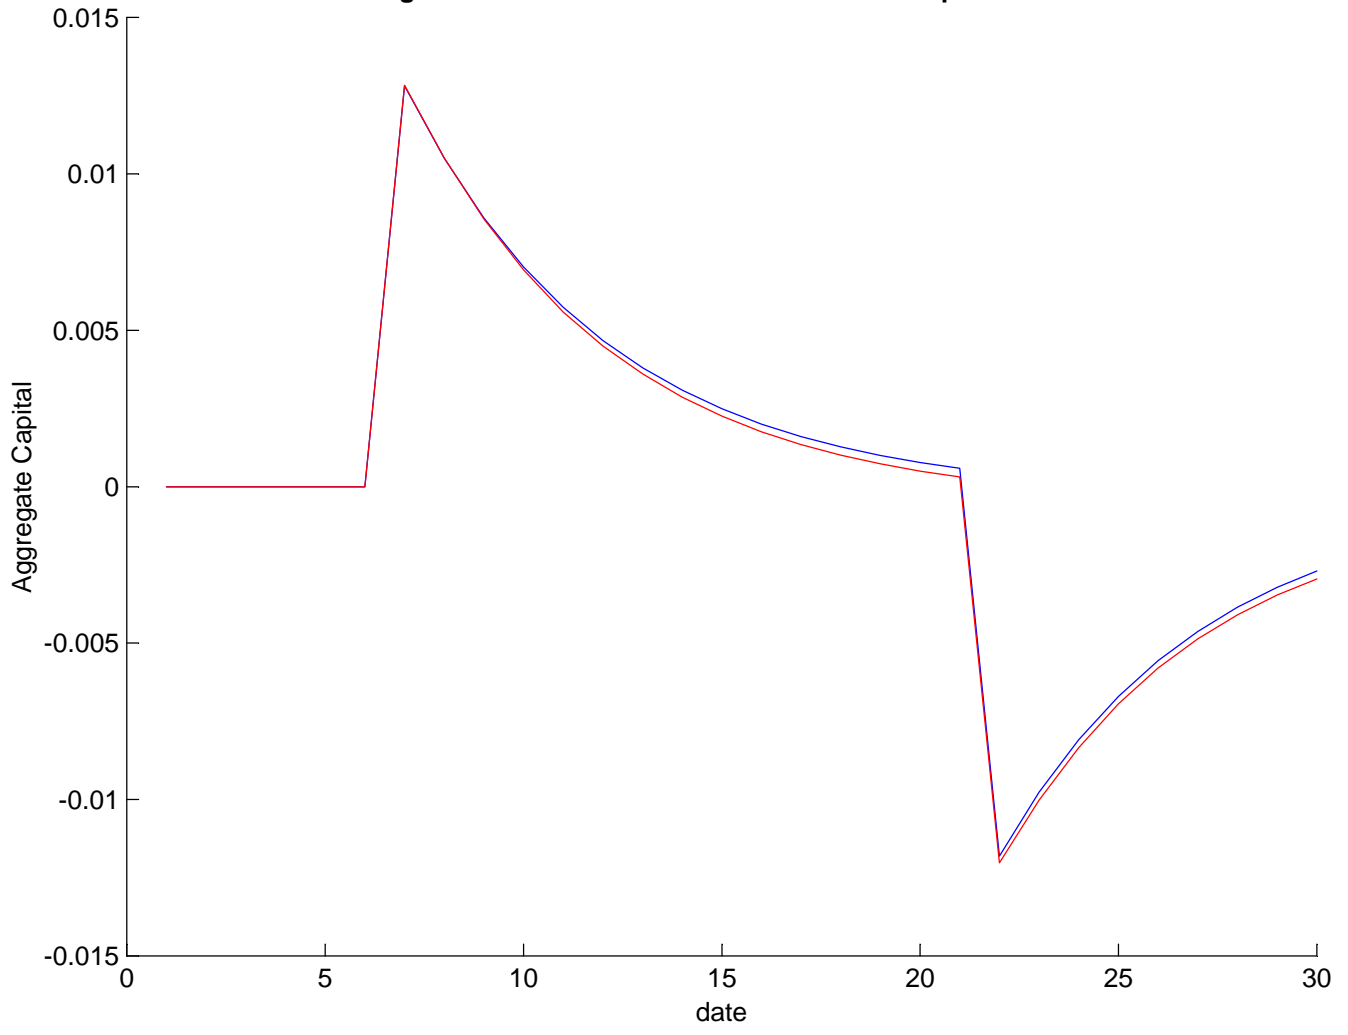


Figure 9: Distribution of Equilibrium Investment Deviations (w/ z and ζ shocks)

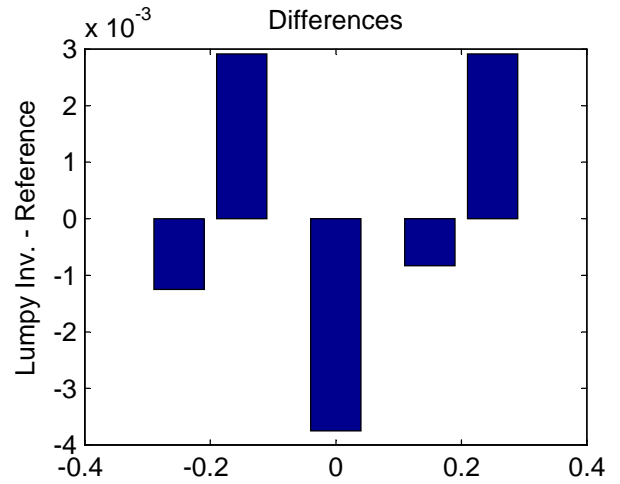
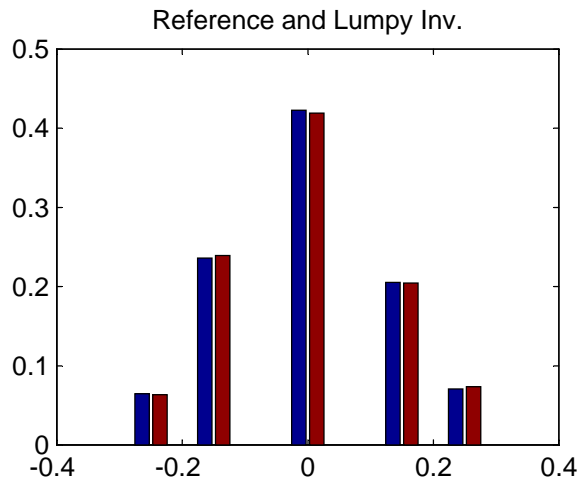
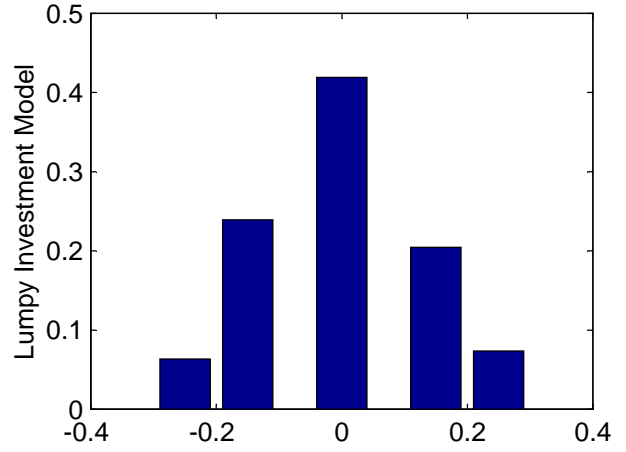
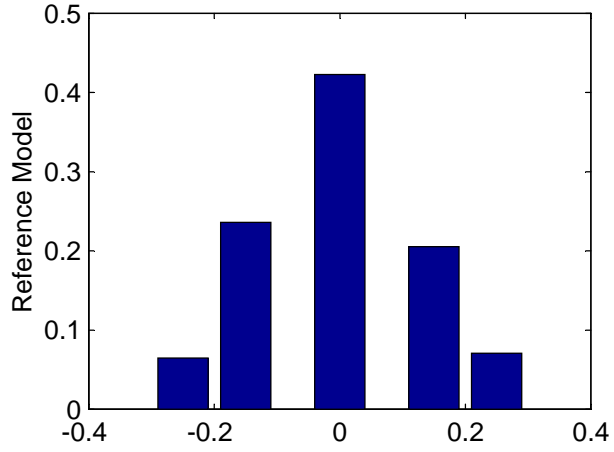


Figure 10: Distribution of Equilibrium Investment Deviations (w/ only ζ shock)

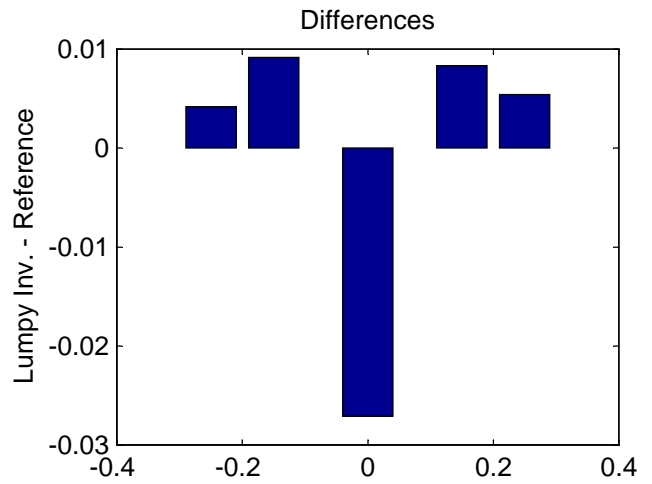
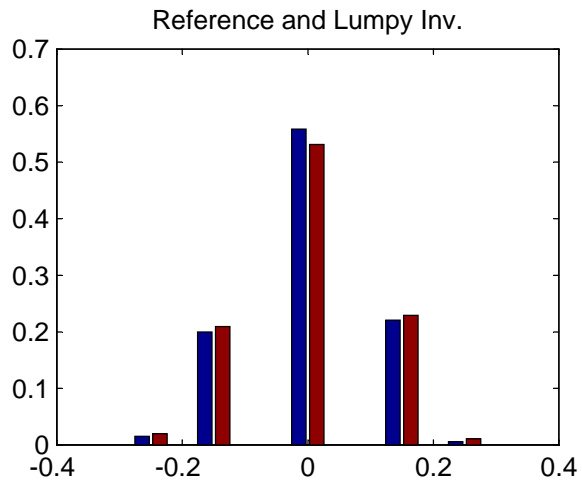
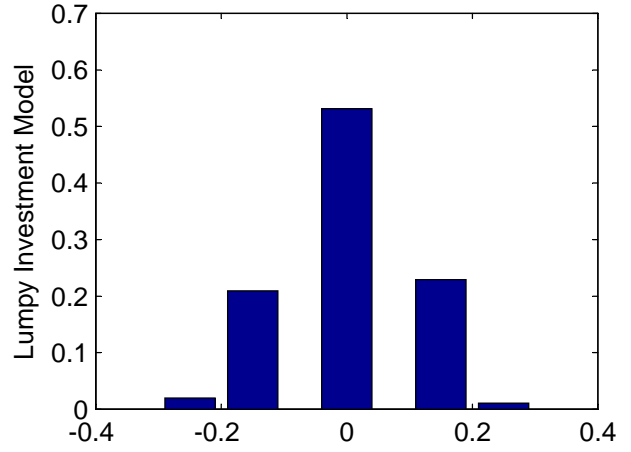
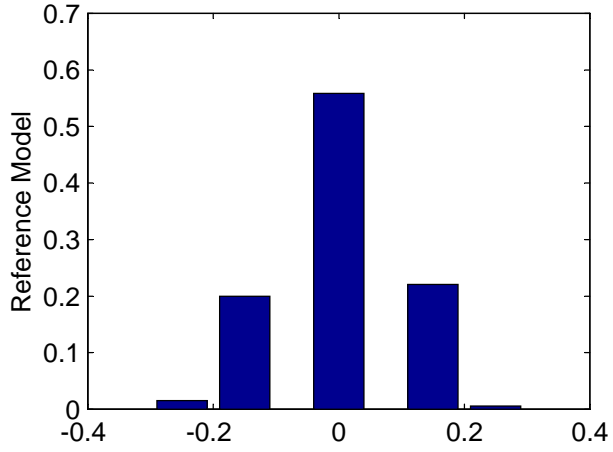


TABLE 1
PARAMETER CHOICES

γ	β	δ	θ	ν	A	ρ	σ_{ε}
1.016	0.954	0.060	0.325	0.580	3.614	0.9225	0.0134

TABLE 2
INVESTMENT DEMAND DEVIATIONS:
DIFFERENCES FROM REFERENCE

	minimum	mean	median	maximum
Lumpy Inv.	.00010	0.463	0.1590	2.053
Partial Adj.	.00008	0.690	0.0907	11.630

TABLE 3
FORECASTING RULES WITH ONE PARTITION

z_i (obs)	β_1	β_2	SE	R^2	
z_1 (161 obs)	m_1'	- 0.0195	+ 0.8255	1.874 e - 4	0.99995
	p	+ 1.2034	- 0.4855	1.000 e - 4	0.99996
z_2 (647 obs)	m_1'	- 0.0071	+ 0.8251	2.358 e - 4	0.99994
	p	+ 1.1810	- 0.4807	8.305 e - 5	0.99998
z_3 (903 obs)	m_1'	+ 0.0054	+ 0.8224	2.525 e - 4	0.99993
	p	+ 1.1584	- 0.4787	7.052 e - 5	0.99999
z_4 (626 obs)	m_1'	+ 0.0181	+ 0.8206	2.189 e - 4	0.99994
	p	+ 1.1355	- 0.4768	6.858 e - 5	0.99998
z_5 (163 obs)	m_1'	+ 0.0308	+ 0.8201	1.256 e - 4	0.99995
	p	+ 1.1123	- 0.4736	4.248 e - 5	0.99999

NOTE. – Forecasting rules are conditional on current productivity, z_i . Each regression is of the form $\log(y) = \beta_1 + \beta_2 \log(m_1)$.

TABLE 4
FORECASTING RULES WITH TWO PARTITIONS

z_1 (161 obs)	β_1	β_2	β_3	SE	R^2
m_1'	- 0.1041	+ 0.4827	+ 0.3115	1.029 e - 3	0.99836
m_2'	- 0.1864	+ 0.4305	+ 0.4314	1.066 e - 3	0.99851
p	+ 0.8664	- 0.2488	- 0.2368	7.261 e - 5	0.99998
z_2 (647 obs)	β_1	β_2	β_3	SE	R^2
m_1'	- 0.0750	+ 0.4252	+ 0.3749	8.894 e - 4	0.99903
m_2'	- 0.1904	+ 0.4670	+ 0.3864	1.069 e - 3	0.99877
p	+ 0.8474	- 0.2479	- 0.2330	7.401 e - 5	0.99998
z_3 (903 obs)	β_1	β_2	β_3	SE	R^2
m_1'	- 0.0632	+ 0.3389	+ 0.4391	7.254 e - 4	0.99939
m_2'	- 0.1804	+ 0.5580	+ 0.3145	8.884 e - 4	0.99926
p	+ 0.8263	- 0.2469	- 0.2318	6.312 e - 5	0.99999
z_4 (626 obs)	β_1	β_2	β_3	SE	R^2
m_1'	- 0.0528	+ 0.3625	+ 0.4141	9.151 e - 4	0.99887
m_2'	- 0.1715	+ 0.5659	+ 0.3072	7.980 e - 4	0.99931
p	+ 0.8053	- 0.2495	- 0.2272	5.423 e - 5	0.99999
z_5 (163 obs)	β_1	β_2	β_3	SE	R^2
m_1'	- 0.0241	+ 0.3216	+ 0.4653	9.068 e - 4	0.99874
m_2'	- 0.1763	+ 0.5945	+ 0.2640	8.883 e - 4	0.99896
p	+ 0.7844	- 0.2484	- 0.2252	3.238 e - 5	0.99999

NOTE. – Forecasting rules are conditional on current productivity, z_i . Each regression is of the form $\log(y) = \beta_1 + \beta_2 \log(m_1) + \beta_3 \log(m_2)$.

TABLE 5
BUSINESS CYCLE MOMENTS

	Output	Investment	Consumption	Employment	Wage	Interest Rate
Standard Deviations						
R	1.906	6.373	0.935	1.101	0.935	0.793
L2	1.906	6.386	0.933	1.102	0.933	0.793
L1	1.905	6.373	0.933	1.100	0.933	0.795
PA	1.547	3.458	1.094	0.473	1.094	1.068
High Costs	1.875	6.152	0.940	1.050	0.940	0.816
Contemporaneous Correlations with Output						
R	1.000	0.971	0.924	0.946	0.924	0.685
L2	1.000	0.972	0.925	0.947	0.925	0.683
L1	1.000	0.972	0.926	0.947	0.926	0.681
PA	1.000	0.990	0.995	0.972	0.995	0.545
High Costs	1.000	0.973	0.936	0.949	0.936	0.665

NOTE. – R: Reference; L2: Lumpy Investment with I=2 partitions; L1: Lumpy Investment with I=1 partition; PA: Partial Adjustment; High Costs: Lumpy Investment with B=.02. All series log HP-filtered.

TABLE 6
EQUILIBRIUM INVESTMENT DEVIATIONS:
DIFFERENCES FROM REFERENCE

	minimum	mean	median	maximum
Lumpy Inv.	1.6 e-7	5.9 e-4	4.8 e-4	0.003
Partial Adj.	3.6 e-6	3.3 e-2	2.5 e-2	0.167

TABLE 7
BUSINESS CYCLE MOMENTS WITH TFP & INVESTMENT-SPECIFIC SHOCKS

	Output	Investment	Consumption	Employment	Wage	Interest Rate
Standard Deviations						
R	2.172	8.575	1.374	2.251	1.374	1.410
L1	2.182	8.670	1.373	2.265	1.373	1.407
Contemporaneous Correlations with Output						
R	1.000	0.884	0.258	0.808	0.258	0.243
L1	1.000	0.885	0.253	0.810	0.253	0.242

NOTE. – R: Reference; L1: Lumpy Investment with I=1 partition.

TABLE 8
BUSINESS CYCLE MOMENTS WITH ONLY INVESTMENT-SPECIFIC SHOCKS

	Output	Investment	Consumption	Employment	Wage	Interest Rate
Standard Deviations						
R	1.128	7.037	1.042	2.021	1.042	1.140
L1	1.152	7.222	1.062	2.060	1.062	1.156
Contemporaneous Correlations with Output						
R	1.000	0.953	- 0.733	0.936	- 0.733	- 0.203
L1	1.000	0.952	- 0.732	0.936	- 0.732	- 0.200

NOTE. – R: Reference; L1: Lumpy Investment with I=1 partition.