NBER WORKING PAPER SERIES

ON THE NATURE OF CAPITAL ADJUSTMENT COSTS

Russell W. Cooper
John C. Haltiwanger

Working Paper 7925
http://www.nber.org/papers/w7925

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
September 2000

The authors thank the National Science Foundation for financial support. Andrew Figura, Chad Syversons and Jon Willis provided excellent research assistance in this project. We are grateful to Andrew Abel, Victor Aguirregabiria, Ricardo Caballero, Fabio Canova, V.V. Chari, Jan Eberly, Simon Gilchrist, George Hall, Adam Jae, Pat Kehoe, John Leahy, David Runkle and Jon Willis for helpful discussions in the preparation of this paper. Helpful comments from seminar participants at Boston University, Brandeis, Cen-TER, Columbia University, the 1 998 Winter Econometric Society Meeting, the University of Bergamo and IFS Workshop on Applied Economics, the University of Texas at Austin, UQAM, Federal Reserve Bank of Cleveland, Federal Reserve Bank of Minneapolis, Federal Reserve Bank of New York, Pennsylvania State University, Wharton School, McMaster University, Pompeu Fabra, Yale and the NBER Summer Institute are greatly appreciated. Financial support from the NSF is gratefully acknowledged. The views expressed are those of the authors and not necessarily those of the National Bureau of Economic Research.

On the Nature of Capital Adjustment Costs
Russell W. Cooper and John C. Haltiwanger
NBER Working Paper No. 7925
September 2000

ABSTRACT

This paper studies the nature of capital adjustment at the plant-level. We use an indirect inference procedure to estimate the structural parameters of a rich specification of capital adjustment costs. In effect, the parameters are optimally chosen to reproduce the nonlinear relationship between investment and profitability that we uncover in the plant-level data. Our findings indicate that a model which mixes both convex and nonconvex adjustment costs with irreversibility fits the data best.

Russell W. Cooper
Department of Economics
Boston University
270 BSR
Boston, MA 02215
and NBER
rcooper@bu.edu

John C. Haltiwanger
Department of Economics
University of Maryland
College Park, MD 20742
and NBER
halt@cyclops.umd.edu

1 Motivation

The goal of this paper is to understand the nature of capital adjustment costs. This topic is central to the understanding of investment, one of the most important and volatile components of aggregate activity. Moreover, understanding of the nature of adjustment costs is vital for the evaluation of policies, such as tax credits, that attempt to influence investment and thus aggregate activity. Despite the obvious importance of investment to macroeconomics, it remains an enigma.

Costs of adjusting the stock of capital reflect a variety of interrelated factors that are difficult to measure directly or precisely.[1] Changing the level of capital services at a business generates disruption costs during installation of any new or replacement capital and costly learning must be incurred as the structure of production may have been changed. Installing new equipment or structures often involves delivery lags and time to install and/or build. The irreversibility of many projects caused by a lack of secondary markets for capital goods acts as another form of adjustment cost.

Some industry case studies (e.g., Holt et al. [1960], Peck [1974], Ito, Bresnahan and Greenstein [1998]) provide a detailed characterization of the nature of the adjustment costs for specific technologies. A reading of these industry case studies suggest that there are indeed many different facets of adjustment costs and that, in terms of modeling these adjustment costs, both convex and nonconvex elements are likely to be present.[2]

Despite this perspective from the industry case studies, the workhorse model of the investment literature has been a standard neoclassical model

[1]As direct measurement of these many factors is difficult, for the most part the study of capital adjustment costs has been indirect through studying the dynamics of investment itself.

[2]Holt et. al. [1960] found a quadratic specification of adjustment costs to be a good approximation of hiring and layoff costs, overtime costs, inventory costs and machine setup costs in the selected manufacturing industries. These components of adjustment costs for changing the level of production are relevant here but are by no means the only relevant costs. In terms of changes in the level of capital services, Peck [1974] studies investment in turbogenerator sets for a panel of 15 electric utility firms and found that "The investments in turbogenerator sets undertaken by any firm took place at discrete and often widely dispersed points of time."In their study of investment in large scale computer systems, Ito, Bresnahan and Greenstein [1998] also find evidence of lumpy investment. Their analysis of the costs of adjusting the stock of computer capital include items which they term "... intangible organization capital such as production knowledge and tacit work routines."

with convex costs (often approximated to be quadratic) of adjustment.[3] This model has not performed that well even at the aggregate level (see Caballero [1999]) but the recent development of longitudinal establishment databases has raised even more questions about the standard convex cost model.

An alternative approach, highlighted in the work of Doms-Dunne [1994], Cooper, Haltiwanger and Power [1999], Abel-Eberly [1994, 1996], and Caballero, Engel and Haltiwanger [1995], argues that nonconvexities and irreversibilities play a central role in the investment process. The primary basis for this view, reviewed in detail below, is plant-level evidence of a nonlinear relationship between investment and measures of fundamentals, including investment bursts (spikes) as well as periods of inaction.

One limitation of this recent empirical literature is that it has focused primarily on reduced form implications of nonconvex vs. convex models. The results that emerge reject the reduced form implications of a pure convex model and are consistent with the presence of nonconvexities. The reduced form nature of the results have left us with several important, unresolved questions: what is the nature of the capital adjustment process at the micro level? Does the micro evidence support the presence of both convex and nonconvex components of adjustment costs as might be expected based upon the limited number of industry case studies? More specifically, what are the structural estimates of the convex and nonconvex components of adjustment costs that are consistent with the micro evidence? Finally, what are the aggregate and policy implications of the estimated investment model?

To address these questions, this paper considers a rich model of capital adjustment which nests alternative specifications. To do so, we specify a dynamic optimizing problem at the plant-level which incorporates both convex and nonconvex costs of adjustment as well as irreversible investment. The model's implications are matched with plant-level observations from the Longitudinal Research Database (LRD) as part of an estimation routine based upon the indirect inference procedure advanced by Gourieroux, Monfort and Renault [1993] and Smith [1993]. We recover structural estimates of the convex and nonconvex components of adjustment costs.

The key link between the theory and the plant-level data is the estimated relationship between investment rates and fundamentals, measured as profitability shocks that we infer from plant-level observations. This relationship

[3] Hamermesh and Pfann [1996] provide a detailed review of convex adjustment cost models and numerous references to the motivation and results of that lengthy literature.

is highly nonlinear: investment is relatively insensitive to small variations of profitability but responds quite strongly to large shocks. Further, there is an empirically important asymmetry present in this estimated relationship: the response to large positive shocks is much more pronounced than is the response to negative shocks. These key features of the data have been pointed out in the recent empirical literature. Our value added is that we take these prominent features of the data and, through the indirect inference procedure, recover the underlying structural parameters.

Our results can be summarized by reference to extreme models with only one form of adjustment cost. The convex cost of adjustment model can not match the periods of inactivity in capital adjustment. Further, that model cannot reproduce the observed nonlinear relationship between investment and profitability. Both the nonconvex and the irreversibility models are able to produce nonlinear relationships between investment and fundamentals which are much closer to the data. Further both of these models imply inactivity and investment bursts. Interestingly, irreversibility creates an asymmetry as well since the loss from capital sales is more relevant when profitability shocks are below their mean. Combinations of these types of adjustment fit the data best. A general theme that emerges from our analysis is that the investment dynamics are much better described with models that have convex and nonconvex costs rather than either convex or nonconvex costs of adjustment alone.

In terms of macroeconomic implications, the natural question is whether these nonconvexities "matter" for aggregate investment. Our findings indicate that at the plant-level, the nonconvexities identified in our estimation are important: a model with only convex adjustment does poorly at the plant level. However, a model with only convex adjustment costs fits the aggregate data created by our estimated model reasonably well though, as reported independently by Cooper, Haltiwanger and Power [1999], hereafter CHP, the convex models tend not to track investment well at turning points. We also find, not surprisingly, that the nonconvexities are less important at the aggregate level than they are for understanding plant level observations.

2 Facts

2.1 Data Set

Our data are a balanced panel from the Longitudinal Research Database consisting of approximately 7,000 large, manufacturing plants that were continually in operation between 1972 and 1988.[4] This particular sample period and set of plants is drawn from the dataset used by Caballero, Engel and Haltiwanger [1995], hereafter CEH. The unique feature of this data relative to other studies that have used the LRD to study investment is that information on both gross expenditures and gross retirements (including sales of capital) are available for these plants for these years (Census stopped collecting data on retirements in the late 1980s in its Annual Survey of Manufactures which is why our sample ends in 1988). Incorporating retirements (and in turn sales of capital) is especially important in this exploration of adjustment costs and frictions in adjusting capital at the micro level. Investigating the role of transactions costs and irreversibilities is quite difficult with the use of expenditures data alone.

The use of the retirements data requires a somewhat modified definition of investment. The definition of investment and capital accumulation that we use follows that of CEH and satisfies:

$$I_t = EXP_t - RET_t \tag{1}$$

$$K_{t+1} = (1 - \delta_t)K_t + I_t \tag{2}$$

where I_t is our investment measure, EXP_t is real gross expenditures on capital equipment, RET_t is real gross retirements of capital equipment, K_t is our measure of the real capital stock (generated via a perpetual inventory method at the plant level), and δ_t is the in-use depreciation rate. This measurement specification differs from the usual one that uses only gross expenditures data and the depreciation rate captures both in-use and retirements. Following the methodology used in CEH, we use the data on expenditures and retirements along with investment deflators and BEA depreciation rates to construct real measures of these series and also an estimate of the in-use

[4]While the balanced panel enables us to avoid modelling the entry/exit process there is undoubtedly a selection bias induced.

depreciation rate.[5] In what follows, we focus on the investment rate, I_t/K_t, which can be positive or negative.

2.2 Moments of the Data

The histogram of investment rates that emerges from this measurement exercise are reported in Figure 1. It is transparent that the investment rate distribution is non-normal having a considerable mass around zero, fat tails, and is highly skewed to the right (standard tests for non-normality yield strong evidence of skewness and kurtosis). Some of the main features of the distribution (and its underlying components in terms of gross expenditures and retirements) are summarized in Table 1. First, note that about 8% of the (plant, year) observations entail an investment rate near zero (investment rate less than 1% in absolute value). Of this inaction, about 6% of the observations indicate gross expenditures less than 1% of the plant capital stock and the retirement rate is less than 1% in 42.3% of our observations. Thus the data exhibit significant inaction in terms of capital adjustment. This is one of the driving observations for our analysis.[6]

These observations of inaction are complemented by periods of rather intensive adjustment of the capital stock. In the analysis that follows we term episodes of investment rates in excess of 20% **spikes**.[7] Investment

[5] A relevant measurement point here is that the retirement data are based upon sales/retirements of capital that yield a change in the book value of capital. Using a FIFO structure and the history of investment and retirements, CEH develop a method to convert this to a real measure of retirements. The methodology yields a measure of the real changes in the plant-level capital stock induced by retirements. In what follows, it is important to note that it does not already capture the difference between buying and selling prices of capital that may influence the adjustment process. We recover that difference as an estimate in our estimation.

[6] Observations of inaction and investment bursts are found in data from other countries as well. For example, Nilsen and Schiantarelli [1998] study investment in Norwegian manufacturing plants for the period 1978-91. For production units, they report that 21% of the units have zero investment expenditures over a given year. Further they find that investment rates exceeding 20% arise in about 10% of their observations and account for about 38% of total equipment investment. Related evidence on the lumpy nature of investment for Colombia is provided by Huggett and Ospina [1998].

[7] Of course, one strength of this approach relative to Cooper, Haltiwanger and Power is that we do not need to reduce our analysis to a dynamic discrete choice problem. Nonetheless looking at these extreme episodes is informative about both the data and the models.

rates exceed 20% in about 18% of our sample observations. On average these large bursts of investment account for about 50% of total investment activity. Decomposing the investment rate in terms of gross expenditures and retirements, there are gross expenditure rate spikes in approximately 23% of the plants while negative investment spikes occur in about 1.4% of the observations.[8]

A third related feature of our data is that investment rates are highly asymmetric. It is important to emphasize that our measurement of negative investment here is through a direct measurement of retirements reflecting purposeful selling or destruction of capital. We find a negative investment rate in roughly 10 percent of our observations, zero investment in almost 10 percent and positive investment rates in the remaining 80 percent of our observations. This striking asymmetry between positive and negative investment is an important feature of the data that our analysis seeks to match.

Variable	LRD
Average Investment Rate	12.2%
Inaction Rate: Investment	8.1%
Fraction of Observations with Negative Investment	10.4%
Spike Rate: Positive Investment	18%
Spike Rate: Negative Investment	1.4%

Table 1: Summary Statistics

2.3 Non-Linearities in the Relationship Between Investment and Fundamentals

A closely related aspect of recent empirical findings from micro data is the nonlinear relationship between investment and fundamentals.[9] This evidence, along with the observed periods of inactivity and investment bursts, is certainly suggestive of nonconvex costs of adjustment. However, one must be careful since these observations, particularly the investment bursts, may be

[8] Interestingly, there are retirement spikes in about 3% of our plants so that apparently some plants are ordering new capital when they are retiring capital.

[9] For example, CHP find that the probability of having a large investment episode is increasing in the time since the last episode and CEH find a highly nonlinear relationship between the rate of investment and a measure of the gap between desired and actual capital.

indicative of large shocks as well. Hence a key to our analysis is understanding the mapping from exogenous shocks to the profitability of enterprises to their capital adjustment. As explained below, by identifying shocks we can infer the nature of adjustment costs from observed investment behavior.

In our analysis, we seek a simple reduced form characterization of the mapping between shocks and investment behavior. We use this relationship in our structural analysis via indirect inference techniques.

2.3.1 Investment Profitability Relationship

Our indirect inference techniques rely heavily on exploiting a simple reduced form empirical relationship between investment and a measure of fundamentals. In our case, we specify a simple nonlinear reduced form relationship between investment and measures of shocks to shocks to plant level profitability. We first describe how profitability is measured at the plant level and then provide an empirical characterization of this relationship.

Estimation of Profit Functions Current profits, for given capital, are given by $\Pi(A, K)$, where the variable inputs (L) have been optimally chosen, a shock to profitability is indicated by A and K is the current stock of capital. That is,

$$\Pi(A, K) = \max_{L} R(\hat{A}, K, L) - Lw(L)$$

where $R(\hat{A}, K, L)$ denotes revenues given the inputs of capital (K) and labor (L) and a shock to revenues, denoted $\hat{A}$. Here $Lw(L)$ is total labor cost. Clearly this formulation implies that there are no costs of adjusting labor. Once we specify a revenue function, we can use this optimization problem to determine the labor input and to derive the profit function $\Pi(A, K)$,where A reflects both the shocks to the revenue function and variations in costs of labor.

Throughout the analysis, the plant level profit function is specified as

$$\Pi(A_{it}, K_{it}) = A_{it} K_{it}^{\theta}. \tag{3}$$

A key parameter is thus θ, the curvature of the profit function. This profit function can be derived from a model in which the production function is

Cobb-Douglas (CRS) and the plant sells a product in an imperfectly competitive market. If α_L denotes labor's coefficient in the Cobb-Douglas technology and ξ is the elasticity of the demand curve, then

$$\theta = ((1 - \alpha_L)(1 + \xi))/(1 - \alpha_L(1 + \xi)) \tag{4}$$

This parameter was estimated from the our panel of plants from the LRD. To do so, we assume that there are both aggregate (A_t) and plant specific profitability shocks (ε_{it}), with A_{it}=$A_t\varepsilon_{it}$. Real profits and capital stocks were calculated at the plant level as explained above. We then estimated θ from (3) using nonlinear least squares.[10]

From the plant-level data, θ is estimated at .51 (standard error is 0.01). Using the LRD plant-level data, we estimate $\alpha_L = .72$ using cost shares. This, in turn implies a demand elasticity of -4.8 and a markup of about 27 percent.

Based upon this estimate of θ, we can, in principle, use the profit and capital measures along with (3) to backout the profit shocks A_{it}. In practice, we generate the profit shock series in an indirect fashion. The reason is that we suspect that there is considerable measurement error in measured profits with a nontrivial number of outliers so that the implied distribution of profit shocks has an enormously large variance. To avoid this problem with measurement error in profit rates, we obtain A_{it} indirectly by using the first order condition for employment which depends upon A_{it}, capital and parameters such as those underlying θ for which we have estimates. Employment is measured with much less measurement error and accordingly we find a substantially lower variance of the profit shocks using this indirect method. Even with this indirect method, we remove fixed effects from the distribution of profit shocks. As discussed below, if there are some underlying structural differences across businesses (which there undoubtedly are) that yield permanent differences in profitability across businesses, then we need to remove them from impacting our analysis since such structural, permanent heterogeneity is outside the scope of our model.[11]

[10] We used plant level fixed effects in this specification and estimated θ using the methodology proposed by Kiviet [1995].

[11] If instead one used the direct measure of profit shocks from the residual profit/capital relationship, removing fixed effects still leaves an enormous variance with incredible outliers – again suggesting the presence of substantial measurement error in our direct measures of profit rates.

With the estimate of the profit shocks at the plant level, we decompose these shocks into aggregate and idiosyncratic components. The aggregate component is simply the yearly mean of the profit shocks; the idiosyncratic component is the deviation from that mean. These series provide the necessary information for the solution of the plant level optimization problem which requires the calculation of a conditional expectation of future profitability.[12]

The Investment Profitability Relationship Letting a_{it}=ln(A_{it}), we study the following relationship between investment and profitability:

$$i_{it} = \alpha_i + \psi_0 + \psi_1 a_{it} + \psi_2 (a_{it})^2 + u_{it} \tag{5}$$

where i_{it} is the investment rate at plant i in period t. Given our interest in understanding nonconvexities in the adjustment process, we allow the investment rate to be a nonlinear of profitability shocks.[13] The specification removes unobserved heterogeneity through the inclusion of fixed effects.

This very simple specification is motivated in a number of ways. First, the prior literature and our analysis of basic moments above suggests a nonlinear relationship between investment and fundamentals. In particular, we know from Table 1 that the investment distribution exhibits a relatively small share of negative rates, a mode at zero investment and then a distribution of positive rates which is very skewed to the right. Moreover, we know from Figure 8 of CEH that there is a highly nonlinear relationship between investment and a measure of the gap between desired and actual investment with modest negative rates even for large negative gaps and increasingly large investment rates for positive gaps. The above simple regression has the potential to capture key features of the asymmetric distribution of investment and the related nonlinear relationship between investment and fundamentals found in the prior literature. It is true that the above specification puts much less structure on this relationship than in the prior literature – but this is

[12]Clearly these series as well as those obtained from production function estimation at the plant level are of independent interest in terms of evaluating competing models of business cycles.

[13]In a previous version of this paper, we allowed cubic terms as well. We prefer the quadratic specification for a number of reasons although the resulting estimates of the structural parameters reported below are very close to those obtained with the cubic function. For reasons of parsimony, we do not split the shock into aggregate and idiosycratic components for this estimation.

intentional as our methodology seeks to identify basic features of the data that we can measure/estimate relatively precisely and then in turn relate that to the underlying structural model via indirect inference techniques.

A second related motivation for the above specification is that we are relatively confident of our ability to measure the investment rate and the profit rate shocks that are used in the above specification. Moreover, in our subsequent analysis using numerical value function iteration we can specify the underlying shock process in the simulated environment to mimic closely the distribution of the shocks in the actual data. Put differently, this reduced form specification has a great advantage in yielding a tight relationship between estimating this reduced form regression in the actual data and the simulated data.

Estimation of (5) at the plant level yields parameter estimates reported in Table 2. The relationship between the investment rate and profitability is shown in Figure 2. The domain of a_{it} reflects the underlying distribution of idiosyncratic profit shocks estimated above. It is not uncommon for profitability to be 50% above or below its average value.

Reduced Form Regression Results	
Coefficients	
ψ_0	-.013 (0.001)
ψ_1	.265 (0.008)
ψ_2	.20 (0.022)
R-squared	0.071
No. observations	96097
(Standard Errors in Parentheses)	

Table 2

There is a statistically significant and economically important nonlinearity in the relationship between investment rates and profitability. For values of the profitability shock near its mean, the investment rate is near zero. It rises rapidly at an increasing rate as the profitability shock increases. Thus, for positive values of a_{it}, the relationship is increasing and convex. In contrast, the response to reductions in profitability is not nearly as large. Thus there is an asymmetry between the response to positive and negative profitability shocks which mimics the basic features of the data emphasized in Table 1 and in the existing literature.

This nonlinear relationship will be used in our estimation as a means of discriminating across competing specifications of the capital adjustment process. To motivate that approach, we first turn to a characterization of leading specifications of the adjustment process.

3 Models and Quantitative Implications

Our most general specification of the dynamic optimization problem at the plant-level is assumed to have both components of convex and nonconvex adjustment costs. Formally, we consider variations of the following stationary dynamic programming problem:

$$V(A,K) = \max_I \Pi(A,K) - C(I,A,K) + \beta E_{A'|A} V(A',K') \tag{6}$$

where $\Pi(A,K)$ represents the (reduced form) profits attained by a plant with capital K, a profitability shock given by A, I is the level of investment and $K' = K(1-\delta) + I$. Here unprimed variables are current values and primed variables refer to future values. In this problem, the manager chooses the level of investment, denoted I, which becomes productive with a one period lag. The costs of adjustment are given by the function $C(I,A,K)$. This function is general enough to have components of both convex and nonconvex costs of adjustment as well as a variety of transactions costs.

This section of the paper provides an overview of the competing models of adjustment. The parameterizations are summarized in Table 3, at the end of this section. For each, we describe the associated dynamic programming problem and display some of the quantitative predictions of the models in Table 4. At this stage these quantitative properties are meant to facilitate an understanding of the competing models. The next section of the paper discusses estimation of underlying parameters.

3.1 Common Elements of the Specification

For the numerical analysis and subsequent estimation, we specify processes for the simulated shocks based upon the actual distributions uncovered from the estimation of the profit function. In the simulations, the aggregate shocks are represented by a first-order, two-state Markov process with $A_t \in \{A_h, A_l\}$ with a transition matrix given by T. For this analysis we set A_h 10% above steady state and A_l 10% below and estimate the diagonal elements in T at .8. The variance of the shocks as well as the degree of serial correlation are based upon the empirical analogues of A_t in the LRD (that is, we compute the empirical analogue of A_t by taking the yearly mean of the A_{it} series computed from the LRD). The idiosyncratic shocks take 11 possible values and are also serially correlated. The transition matrix for these shocks is computed directly from the empirical transitions observed at the plant-level and thus reproduces statistics from the idiosyncratic profitability shock series.

For the remaining parameters, we set the annual discount factor (β) at .95 and the annual rate of depreciation at 6.9%. This depreciation rate is consistent with the one used to create the capital stock series at the plant level less a retirement rate of 3.2%.

3.2 Convex Costs of Adjustment

The traditional investment model assumes that costs of adjustment are convex. Here we adopt a quadratic cost specification and consider the following specification of the adjustment function,

$$C(I, A, K) = pI + \frac{\gamma}{2}[I/K]^2 K$$

where γ is a parameter. The first-order condition for the plant level optimization problem relates the investment rate to the derivative of the value function with respect to capital and the cost of capital (p). That is, the solution to (6) implies

$$i = (1/\gamma)[\beta E V_k(A', K') - p] \tag{7}$$

where i is the investment rate and EV_k is the expectation of the derivative of the value function in the subsequent period. In practice, this derivative is not observable.

If profits are proportional to the capital stock, θ =1, the model reduces to the familiar "Q theory" of investment in which the value function is proportional to the stock of capital. Hence, the derivative of the value function can be inferred from the average value of a firm, $V_k(A, K) = V(A, K)/K$.[14]

[14] This point is made by Hayashi [1982]. Of course, given that the estimate of the curvature of the profitability function is significantly less than 1, any Q theory based investment regressions are misspecified. Cooper-Ejarque [2000] investigate the implications of this for the statistical significance of profits in investment regressions.

As suggested by (7), the investment policy has a partial adjustment structure. There is a gap between the value of a marginal increment to the capital stock and the price of capital. The optimal policy is to partially close this gap where the speed of adjustment is parameterized by γ.Clearly, this model implies continuous investment activity and thus will be unable to match observations of inactivity. Note though that large bursts of investment are possible within this framework as long as the shocks are sufficiently volatile and persistent.

We also study the special case of no adjustment costs, $C(I, A, K) \equiv 0$. In this case, the optimal capital stock for the plant satisfies:

$$\beta E V_k(A', K') = p$$

which comes directly from (6). In this specification, the future capital stock and thus investment are extremely responsive to variations in persistent movements in profitability.

3.3 Nonconvex costs of Adjustment

Building upon the analysis of Abel and Eberly [1999] and Cooper, Haltiwanger and Power [1999], during periods of investment plants incur a fixed cost which is proportional to their stock of capital.[15] These fixed adjustment costs represent the need for plant restructuring, worker retraining and organizational restructuring during periods of intensive investment. Generally, these nonconvex costs of adjustment are intended to capture indivisibilities in capital, increasing returns to the installation of new capital and increasing returns to retraining and restructuring of production activity.

For this formulation of adjustment costs, the dynamic programming problem is specified as:

$$V(A, K) = \max\{V^i(A, K), V^a(A, K)\}$$

where the superscripts refer to active investment "a" and inactivity "i". These options, in turn, are defined by:

$$V^i(A, K) = \Pi(A, K) + \beta E_{A'|A} V(A', K(1-\delta))$$

[15] That analysis also allowed for a loss proportional to current profits due to shutdowns and so forth. Here we do not allow that form of adjustment cost as it is impossible to separate it from the idiosyncratic profitability shocks we have estimated.

and

$$V^a(A,K) = \max_I \Pi(A,K) - FK - I + \beta E_{A^\theta|A} V(A', K')$$

In this second optimization problem, the cost of adjustment is independent of the investment activity of agent as described above. Here the cost of new investment goods is normalized at 1.

The intuition for optimal investment policy in this setting comes from CHP. In the absence of profitability shocks, the plant would follow an optimal stopping policy: replace capital iff it has depreciated to a critical level. Adding the shocks creates a state dependent optimal replacement policy but the essential characteristics of the replacement cycle remain: there is frequent investment inactivity punctuated by large bursts of capital purchases/sales. Relative to the partial adjustment of the convex model, the model with nonconvex adjustment costs provides an incentive for the firm to "overshoot its target" and then to allow physical depreciation to reduce the capital stock over time.

3.4 Transactions Costs

Finally, as emphasized most recently by Abel and Eberly [1994,1996], it is reasonable to consider the possibility that there is a gap between the buying and selling price of capital, reflecting, inter alia, capital specificity and a lemons problem.[16] This is incorporated in the model by assuming that

$$C(I,A,K) = pI \text{ where p=p}_\text{b} \text{ if I>0 and p=p}_\text{s} \text{ if I<0}$$

where $1 = p_b \geq p_s$. In this case, the gap between the price of new and old capital will create a region of inaction.

The value function for this specification is given by:

$$V(A,K) = \max\{V^b(A,K), V^s(A,K), V^i(A,K)\}$$

where the superscripts refer to the act of buying capital "b", selling capital "s" and inaction "i". These options, in turn, are defined by:

[16] In fact, Abel and Eberly [1994] include other forms of nonconvex adjustment in their model. Part of the point of looking at retirements (i.e. sales of capital) is to better evaluate this model.

$$V^b(A,K) = \max_I \Pi(A,K) - I + \beta E_{A'|A} V(A', K(1-\delta) + I),$$

$$V^s(A,K) = \max_R \Pi(A,K) + p_s R + \beta E_{A'|A} V(A', K(1-\delta) - R)$$

and

$$V^i(A,K) = \Pi(A,K) + \beta E_{A'|A} V(A', K(1-\delta))$$

Here we distinguish between the purchase of new capital (I) and retirements of existing capital (R). As there are no vintage effects in the model, a plant would never simultaneously purchase and retire capital.

The presence of irreversibility will have a couple of implications for investment behavior. First, there is a sense of caution: in periods of high profitability, the firm will not build its capital stock as quickly since there is a cost of selling capital. Second, the firm will respond to an adverse shock by holding on to capital instead of selling it in order to avoid the loss associated with $p_s < 1$.

3.5 Evaluation of Competing Models

As indicated by Table 3, we explore the quantitative implications of four models. While these parameterizations are not directly estimated from the data, they provide some interesting benchmark cases that highlight the key issues arising between models with convex and nonconvex costs of adjustment. The first, denoted "No AC" is the extreme model in which there are no adjustment costs. The second row, denoted CON, corresponds to a specification in which there are only convex costs of adjustment. The case labeled "NC" assumes that there are only nonconvex costs of adjustment with F>0. Finally, the case labeled "TRAN" imposes a gap of 25% between the buying and selling price of capital.

Model	γ	F	p_s	p_b
No AC	0	0	1	1
CON	2	0	1	1
NC	0	0.05	1	1
TRAN	0	0	.75	1

Table 3

Our quantitative findings for the specifications in Table 3 along with data from the LRD are summarized in Table 4. Figure 3 shows the estimated relationships between investment rates and profitability for these models. The associated regression coefficients are reported in Table 4 as well.

Moment	LRD	No AC	CON	NC	TRAN
investment inaction	.081	0.001	0	.917	.511
investment bursts	.18	.308	0	.081	.153
Corr(i_{it},i_{it-1})	.007	-.167	.479	-.083	.002
Corr(i_{it}, a_{it})	.245	.442	.858	.246	.499
Coefficients					
ψ_0	-.013	-.028	-.002	-.022	-.016
ψ_1	.265	.542	.053	.339	.246
ψ_2	.20	.377	.023	.303	.212

Table 4

As noted earlier, there is evidence of lumpiness and inaction in the LRD. In addition, there is essentially no autocorrelation in plant-level investment and a nontrivial positive correlation between investment and profitability. The lack of autocorrelation is noteworthy given that idiosyncratic shocks to profitability exhibit a correlation of 0.14.

Comparing the columns of Table 4 pertaining to the extreme models with the column labeled LRD, none of the models alone fits these key moments from the LRD. The extreme case of no adjustment costs (labelled No AC) is given in the second column. This model produces no inaction but is capable of producing bursts in response to variations in the idiosyncratic profitability shocks. Note that this model actually creates negative serial correlation in investment rates reflecting the lack of a motivation for smoothing investment expenditures.

The quadratic adjustment cost model (labeled CON) adds convex adjustment costs to the No AC model. This specification cannot capture either the inactivity or investment bursts. Further, it yields much higher autocorrelations and investment/profit correlations than are observed in the data. In fact, the convex cost of adjustment model, through the smoothing of investment, creates serial correlation in investment relative to the shock process.

Non-convex costs of adjustment (NC) and/or the model with irreversibility (TRAN) are able to create investment inactivity at the plant level. However, the pure non-convex model creates modest negative serial correlation in investment data and a lower correlation between investment rates and profitability. The negative serial correlation of the NC model is analogous to the upward sloping hazards characterized by CHP.

Looking at Figure 3, note that at $\gamma = 2$, the convex model produces a very flat relationship between the investment rates and profitability shocks. However, nonlinearities are certainly apparent for the other models, even that without any adjustment costs whatsoever. Again, these nonlinear responses reflect the curvature in the profit function, the adjustment costs and the nonlinearities induced by our specification which uses the investment rate as a dependent variable.[17]

In particular, both the NC and TRAN models create relationships between investment rates and profitability shocks that mimic important aspects of the data: an increasing, convex response to shocks above average and a more muted response to adverse profitability shocks. For the TRAN specification, this dampened response to adverse shocks seems warranted since selling off capital when profitability falls is "expensive". For the NC specification, the bunching of investment activity, reflected in the frequent periods of inactivity followed by bursts of investment, implies that the investment rate will rise quickly as profitability differs from its mean but investment rates will generally be near zero when profitability is near its mean.[18]

4 Estimation

None of these extreme models is rich enough to match key simple properties of the data. Our approach is to consider a hybrid model with all forms of

[17]So, for example, if there are no adjustment costs then there is a log-linear relationship between the future capital stock and the current profitability shock but this implies a nonlinear relationship between the log of this shock and the investment rate.

[18]Of course, even in the absence of any shocks, there will be infrequent bursts of investment followed by periods of inaction.

adjustment costs and, in turn, to estimate the key parameters of this hybrid model by matching the implications of the structural model with key features of the data.

4.1 Indirect Inference

The methodology that we use for this purpose is the indirect inference method of Gourieroux et al. [1993] and Smith [1993]. There are two key elements in the implementation of this methodology. First, there is the question of selecting a reduced form regression for the indirect inference procedure. Second, there is the issue of unobserved heterogeneity. We discuss these in turn. In the course of this presentation, it will become clear why the indirect inference methodology has particular advantages in this setting.

The key to this methodology is a regression (hereafter termed the reduced form regression) which is run on both actual and simulated data. The simulated data set is created by solving the dynamic programming problem given a vector of parameters. The resulting policy functions are then used to create a panel data set comparable to the LRD. The structural parameters are chosen so that the coefficients of the reduced form regression from the simulated data are "close" to the estimates from the actual data.

The choice of a reduced form regression is a crucial piece of the analysis. For our purposes, the reduced form regression needs to satisfy two criteria. First, the parameters of the regression should be "informative" about the underlying structural parameters. That is, as the structural parameters are varied, the regression coefficients should be responsive.[19] Second, the reduced form regression should summarize relevant aspects of the investment decision. As emphasized above, one of the basic insights of the recent theoretical and empirical literature on nonconvexities is that they imply nonlinearities in the relationship between investment and fundamentals.[20]

[19] More formally, a sufficient condition for identification of the parameters (as in Assumption 2 of Smith [1993]) is that there exists a one-to-one mapping between the structural parameters and the moments calculated from the data. The sensitivity of the reduced form coefficients to variations in the structural parameters is a property of the regression that can be evaluated in simulations and reappears in determining the magnitude of the standard errors.

[20] For example, the specification and findings in CEH can be interpreted in this fashion. More directly, Barnett and Sakellaris [1999] explicitly fit a flexible functional form allowing for a nonlinear relationship between investment and Q and find evidence of significant nonlinearities. In this earlier work, the precise link between the underlying structural

This point was used to motivate our analysis of the nonlinear relationship between investment rates and profitability and we continue to use that relationship as a focus for our estimation. Thus our specification remains:

$$i_{it} = \alpha_i + \psi_0 + \psi_1 a_{it} + \psi_2 (a_{it})^2 + u_{it} \quad (8)$$

Note that we have included a fixed effect in this regression. Clearly there is unobserved heterogeneity in the LRD. To match our model with data thus requires us to either build the unobserved differences across plants into our analysis or to purge the LRD of these differences. We have chosen the latter approach which seems quite natural within the regression oriented indirect inference methodology.[21]

4.2 Structural Estimation of a Mixed Model

The model we estimate includes convex and nonconvex adjustment processes as well as irreversible investment.[22] Specifically, assume that the dynamic programming problem for a plant is given by:

$$V(A, K) = \max\{V^b(A, K), V^s(A, K), V^i(A, K)\}$$

where, as above, the superscripts refer to the act of buying capital "b", selling capital "s" and inaction "i". These options, in turn, are defined by:

$$V^b(A, K) = \max_I \Pi(A, K) - FK - I - \frac{\gamma}{2}[I/K]^2 K + \beta E_{A'|A} V(A', K(1-\delta) + I),$$

model and the nonlinear empirical specifications is not specified. In many ways, the value-added of our approach and analysis is that we make this link explicit and in turn we can recover the underlying structural parameters.

[21]The treatment of unobserved heterogeneity is an important issue. Our approach is to extract this from this data by included fixed effects in our reduced form regressions. It is useful to consider the implications of unobserved heterogeneity for the underlying structural model. As shown by Gilchrist-Himmelberg, unobserved heterogeneity in the adjustment function appears in the intercept of the standard Q regression. While our reduced form regression has profitabilty measures rather than average Q, we demonstrated through simulations that the reduced form coefficients (other than the intercept) were independent of fixed effects entered into the adjustment cost function.

[22]This combining of adjustment cost specifications may be appropriate for a particular type of capital (with say installation costs and some degree of irreversibility) and/or may also reflect differences in adjustment cost processes for different types of capital. Our data is not rich enough to study a model with heterogeneous capital.

$$V^s(A,K) = \max_R \Pi(A,K) + p_s R - FK - \frac{\gamma}{2}[R/K]^2 K + \beta E_{A'|A} V(A', K(1-\delta) - R)$$

and

$$V^i(A,K) = \Pi(A,K) + \beta E_{A'|A} V(A', K(1-\delta)).$$

We have specified some parameters of the model ($\beta = .95, \delta = .069$) for the functional forms discussed above. Further, we retain our specification of the profit function, $\Pi(A,K) = AK^\theta$ with θ=.51. For the structural estimation, we focus on three parameters, $\Theta \equiv (F, \gamma, p_s)$, which characterize the magnitude of the nonconvex and the convex components of the adjustment process and the size of the irreversibility of investment.

These parameters are estimated using the following routine. For arbitrary values of the vector of parameters (Θ),the dynamic programming problem is solved and policy functions are generated. Using these policy functions, the decision rule is simulated given arbitrary initial conditions. The simulation creates a version of the LRD. We then estimate the reduced form investment regression, (5), on the simulated version of the LRD. Let Ψ^d represent the estimates of $[\psi_1, \psi_2, \psi_3]$ from (8) reported in Table 2. Further, let $\Psi^s(\Theta)$ denote the vector of regression coefficients from estimation of (5) on the simulated data set, ignoring the constant term. Note that this vector of reduced form parameters depends on the vector of structural parameters, Θ, in a nonlinear way.

The estimate $\hat{\Theta}$ minimizes the weighted distance between the actual and simulated regression coefficients. Formally, we solve

$$\$ = \min_\Theta [\Psi^d - \Psi^s(\Theta)]' W [\Psi^d - \Psi^s(\Theta)]$$

where W is a weighting matrix. We use the optimal weighting matrix given by the inverse of the variance-covariance matrix of the regression coefficients.[23]

Of course, the $\Psi^s(\Theta)$ function is not analytically tractable. Thus, the minimization is performed using numerical techniques. Given the potential for discontinuities in the model and the discretization of the state space, we used a simulated annealing algorithm to perform the optimization.

Table 5 reports our results for four different models along with standard errors. The first row estimates the complete model with three structural

[23]See Smith [1993].

parameters used to match three reduced form coefficients. Here we find that we are able to come fairly close to matching the regression coefficients with a parameter vector given by: $\Theta = [0.043, 0.00039, 0.967]$. These parameter estimates imply relatively modest convex and nonconvex adjustment costs (but non-zero) and a relatively substantial transaction cost. Restricted versions of the estimated model are also reported for purposes of comparison. Clearly the mixed model does better than any of the restricted models. In particular, F is significantly different from zero, p_s is significantly different from 1 and γ is significantly different from zero.

Spec.	Structural Parm. Estimates (s.e.)			parm. est. for (5)		
	γ	F	p_s	ψ_0	ψ_1	ψ_2
LRD				-.013	.265	.20
all	.043 (0.00224)	.00039(.0000549)	.967(.00112)	-.013	.255	.171
F only	0	.0333(.0000155)	1	-.02	.317	.268
γ only	.125(.000105)	0	1	-.007	.241	.103
p_s only	0	0	.93(.000312)	-.016	.266	.223

Table 5

Figure 4 shows the empirical relationship between investment rates and profitability for the estimated model and Figure 4 shows that relationship using the estimates of the alternative models. Clearly, the full model does capture the nonlinear relationship between investment and fundamentals found in the data better than any of the restricted models.

Using our estimates, we can return to the moments reported in Table 4 which were not explicitly included in the our estimation procedure. With regard to our motivating themes of inaction and bursts of investment, the estimated model has both features reflecting both the nonconvex adjustment and the irreversibility of investment. In particular, the estimated model exhibits investment spikes (investment rates in excess of 20%) in about 13.7% of the plant year observations compared to the 18% spike rate in the LRD. Further, the investment inaction rate is 52.7% which is substantially higher than that found in the data. With these estimates, the correlation between profitability and investment is 0.586 while the serial correlation of investment rates is 0.089. Both of these correlations are much closer to the data than the pure convex model and reflect the fact that the introduction of nonconvex adjustment costs and irreversibilities reduces both the responsiveness of investment to shocks and its serial correlation.

4.3 Evaluation

Are these results reasonable? Of interest relative to other studies is the relatively low estimated value of the convex cost of adjustment parameter, γ. This parameter has received enormous attention in the literature since a regression of investment rates on the average value of the firm (Q) will identify this parameter when the profit function is proportional to K and the cost of adjustment function is convex and homogenous of degree one. Using the Q-theoretic approach, estimates of γ range from over 20 (Hayashi [1982]) to as low as 3 (Gilchrist-Himmelberg [1995], unconstrained subsamples, bond rating).

Relative to these results, our estimates of $\gamma = .043$ for the full model and $\gamma = .125$ in the pure convex model appear extremely low. However, following Cooper-Ejarque [2000], the misspecification of Q-theory based models may explain these differences. In particular, when there is curvature in the profit function (we estimate $\theta = .51$), the assumptions underlying Q-theory do not hold: i.e. the substitution of average for marginal Q produces a measurement error. Could this explain the differences in findings about the magnitude of γ?

To study this point, we simulated a panel data set using our estimates. From this data set and the associated values from the dynamic programming problem, we constructed measures of expected discounted average Q.[24] We then regressed investment rates on these measures of average Q and used then inferred the value of γ from the regression coefficient on average Q. When $\gamma = .043$ is used in the simulation, the coefficient on average Q in a regression of investment rates on a constant and average Q is very precisely estimated at .1427 implying an estimate of $\gamma = 7.01$! Thus, the measurement error induced by replacing marginal with average Q biases the coefficient on average Q downwards enough to create an inferred value of the convex adjustment cost parameter that is well within the range of conventional estimates.[25]

[24]Given the one-period time to build, the convex model relates investment to the expected discounted value of the derivative of the value function, as in (7).

[25]Essentially the substitution of average for marginal Q creates a negative correlation close to unity between the "error" and average Q. Of course, this correlation goes to 0 if

Further, while others have considered models with nonconvex costs of adjustment, there are no estimates comparable to our estimate of the fixed cost.[26] This estimate implies that the fixed cost of adjustment is about 0.04% of average plant level profits. Despite its modest magnitude, this fixed cost matters for investment decisions. Forcing this parameter to zero yields nontrivial changes in the reduced form parameter estimates.

Put differently, even though this fixed cost may seem small relative to profits, it can nonetheless be quite influential in terms of decisions if it is large relative to the gain from adjustment. Formally, define

$$\Delta V(A,K) = V^{act}(A,K) + FK - V^{i}(A,K)$$

where V^{act} is the value of action (regardless of whether it involves buying or selling capital) and V^{i} is the value of inaction. So, $\Delta V(A,K)$ measures the state contingent gain to action excluding the fixed cost of adjustment (which has been added back). Figure 6 shows $\Delta V(A,K)$ and the nonconvex adjustment cost (FK) for a particular value of A as a function of K using the estimates of the full model. Note that $\Delta V(A,K)$ is less than FK over a substantial range of the capital stock state space. Thus the fixed cost is large enough to generate inaction over this region.

Our estimates of the degree of irreversibility is within reason, though it is considerably lower (i.e. p_s is close to 1) than the estimate provided by Ramey and Shapiro [1998] for some plants in the aerospace industry. As noted earlier, this level of transactions cost is enough to generate inactivity instead of sales of capital. Holding the other parameter estimated fixed, if we remove the irreversibility by setting $p_s = 1$, the rate of inaction falls to about 36% and the incidence of investment bursts rises to about 20%. In this sense, the presence of even modest transactions costs creates inaction and also dampens the responsiveness of investment to shocks.

the profit function is proportional to K. Cooper-Ejarque develop this point to argue that the same measurement error can explain the significance of profit rates in Q regressions in the absence of capital market imperfections.

[26] In particular, neither CHP nor CEH estimate adjustment costs directly. Further, while fixed costs of adjustment are present in the Abel and Eberly [1999] model they do not appear to be estimated either.

5 Aggregate Implications

The estimation results reported in Table 5 indicate that a model which mixes both convex and nonconvex adjustment processes can match moments calculated from plant level data quite well. An issue for macroeconomists, however, is whether the presence of nonconvexity at the microeconomic level "matters" for aggregate investment. CEH find that introducing the nonlinearities created by nonconvex adjustment processes can improve the fit of aggregate investment models. CHP find that the convex adjustment cost model fits aggregate investment (across their manufacturing plants) reasonably well on average. However, there are years where the interaction of an upward sloping hazard (investment probability as a function of age) and the cross sectional distribution of capital vintages does matter for aggregate investment.

To study the contribution of nonconvex adjustment costs to aggregate investment (defined by aggregating across the plants in our sample), we compare the aggregate implications of our estimated model, termed the **best overall fit model**, against two models with convex costs of adjustment. The first, termed the **estimated convex model**, is the best fit of the pure convex model to the properties of the micro data as reported in Table 5 (with $\gamma = .125$). The second, termed the **alternative convex model** imposes $\gamma = 2$,which is closer to that found in Q-based investment models, described in Gilchrist-Himmelberg [1995] for example. We simulated all three models for 100 periods using the same draw of aggregate and idiosyncratic shocks. The three simulated times series are shown in Figure 7.

For this exercise, we treat the aggregate time series created by the best overall fit model as "truth" and consider how well the alternative convex adjustment costs models match this time series. It is apparent that the alternative convex model does very poorly in terms of its aggregate dynamics. Aggregate investment from this specification is much smoother than that implied by the best overall fit model. Note, however, that the aggregate investment implied by the estimated convex model does a reasonably good job of matching the aggregate investment behavior of the best overall fit model. For this latter comparison, there are differences at "turning points" when the aggregate shock changes from one state to another. In particular, when the aggregate shock switches from below to above its mean, the resulting burst of investment is much higher for the best overall fit model than for the estimated convex model.

Using these three series, we computed a pseudo R^2 measure defined as the fraction of the fluctuations in the aggregate investment rate that are explained by the alternative convex models.[27] For the series shown in Figure 7, this measure of goodness of fit for the alternative convex model is 0.22 while it is 0.93 for the estimated convex model. These findings along with Figure 7 suggest that the alternative convex model, parameterized at "conventional" levels, yields aggregate dynamics that are far from those associated with the best overall fit model. However, the estimated convex model (which has quite weak convexity) yields aggregate patterns that are not too far from those implied by the best fit model.[28]

In our simulated environment, we can explore the factors that yield increases or decreases the fraction that can be accounted for by the convex model. If we double the variance of aggregate shocks for example this pseudo R^2 measure for the best fit convex model drops to 91.7%.[29] This sensitivity to the magnitude of the aggregate shocks is intuitive as larger aggregate shocks push the distribution of plants into different regions of the nonlinear relationship between investment and fundamentals that are a key feature of models with nonconvexities. This finding is important in practice because it suggests that the micro nonconvexities that we estimate as being important in the micro data are likely to be most important empirically for aggregate investment at times of especially large shocks.

It is also of interest to ask the question how well does the estimated convex model fit the overall distribution of investment rates. One way of answering this question is to compute the pseudo R^2 using the simulated plant level data for the pure and mixed cases. Computing the pseudo R^2 measure in this manner for the estimated convex model yields 0.79. Not surprisingly, smoothing by aggregation implies that the convex model does relatively better in matching aggregate relative to plant level data.

This analysis of the aggregate implications is potentially incomplete in that we do not explicitly consider general equilibrium considerations. As

[27] Letting ε be the difference between aggregate investment from the estimated model (est) and the convex model, we defined our goodness of fit measure as 1-var(ε)/(var(est)).

[28] This finding that approximately ninety percent of the aggregate fluctuations in investment can be accounted for with convex cost elements alone is roughly consistent with the findings in CEH and CHP. For example, in CEH, a regression of aggregate investment on the first moment of the gap between desired and actual capital yielded a R^2 of roughly .65 while adding higher moments raised the pseudo R^2 measure to 0.80.

[29] That is, the aggregate shocks take the values $\{.8,1.2\}$ instead of $\{.9,1.1\}$.

emphasized by Veracierto [1998], Caballero [1999] and Thomas [2000], it is likely that there is further smoothing by aggregation due to the congestion effects that are potentially present in the capital goods supply industry. Our focus here is more of a pure aggregation exercise considering the aggregate implications of the alternative mixed and pure specifications of structural models given the distributions of aggregate and idiosyncratic shocks.

6 Conclusions

The goal of this paper is to analyze capital dynamics through competing models of the investment process: what is the nature of the capital adjustment process? The methodology is to take a model of the capital adjustment process with a rich specification of adjustment costs and solve the dynamic optimization problem at the plant level. Using the resulting policy functions to create a simulated data set, the procedure of indirect inference is used to estimate the structural parameters.

Our empirical results point to the mixing of models of the adjustment process. The LRD indicates that plants exhibit periods of inactivity as well as large investment bursts. Further, the relationship between investment rates and measures of profitability at the plant level is highly nonlinear. A model which incorporates both convex and nonconvex aspects of adjustment, including irreversibility, fits these observations best.

In terms of further consideration of these issues, we plan to continue this line of research by introducing costs of employment adjustment. This is partially motivated by the ongoing literature on adjustment costs for labor as well as the fact that the model without labor adjustment costs implies labor movements that are not consistent with observation.

Further, it would be insightful to utilize this model to study the effects of investment tax subsidies. Here those subsidies enter quite easily through policy induced variations in the cost of capital.

References

[1] Abel, A. and J. Eberly, "A Unified Model of Investment under Uncertainty," American Economic Review, 84 (1994), 1369-84.

[2] Abel, A. and J. Eberly, "Optimal Investment with Costly Reversibility," Review of Economic Studies,"63 (1996), 581-93.

[3] Abel, A. and J. Eberly, "The Mix and Scale of Factors with Irreversibility and Fixed Costs of Investment," NBER Working Paper #6148, 1997.

[4] Abel, A. and J. Eberly, "Investment and q with Fixed Costs: An Empirical Analysis," mimeo, April 1999.

[5] Caballero, R., "Aggregate Investment," in Handbook of Macroeconomics, (edited by J. Taylor and M. Woodford), Amsterdam: North Holland, 1999.

[6] Caballero, R., E. Engel and J. Haltiwanger, "Plant-level Adjustment and Aggregate Investment Dynamics," Brookings Papers on Economic Activity, 2 (1995), 1-39.

[7] Caballero, R., E. Engel and J. Haltiwanger, "Aggregate Employment Dynamics: Building from Microeconomic Evidence," American Economic Review 87 (1997), 115-37.

[8] Cooper, R. and J. Ejarque, "Exhuming Q:Market Power vs.Capital Market Imperfections" mimeo, Boston University, August 2000.

[9] Cooper, R. and J. Haltiwanger, "The Aggregate Implications of Machine Replacement: Theory and Evidence," American Economic Review, 83 (1993), 360-82.

[10] Cooper, R., J. Haltiwanger and L. Power, "Machine Replacement and the Business Cycle: Lumps and Bumps," American Economic Review, 89 (1999), 921-946.

[11] Doms, M. and T. Dunne,"Capital Adjustment Patterns in Manufacturing Plants," Center for Economic Studies, U.S. Bureau of the Census, 1993.

[12] Gilchrist, S. and C. Himmelberg, "The Role of Cash Flow in Reduced-Form Investment Equations," *Journal of Monetary Economics, 1995.*

[13] Gilchrist, S. and C. Himmelberg, "Investment, Fundamentals and Finance," *The NBER Macroeconomics Annual*, 1998.

[14] Gomes, J. "Financing Investment," mimeo, Wharton School, University of Pennsylvania, January 1998.

[15] Goolsbee, A. "The Business Cycle, Financial Performance and the Retirement of Capital Goods," NBER Working Paper #6392, February 1998.

[16] Goolsbee, A. and D. Gross, "Estimating Adjustment Costs with Data on Heterogeneous Capital Goods," NBER Working Paper #6342, December 1997.

[17] Hayashi, F. "Tobin's Marginal and Average q: A Neoclassical Interpretation," *Econometrica* 50 (1982), 213-24

[18] Hamermesh, D. "Labor Demand and the Structure of Adjustment Costs," *American Economic Review*, 79 (1989), 674-89.

[19] Hamermesh, D. and G. Pfann, "Adjustment Costs in Factor Demand," *Journal of Economic Literature*, 34 (1996), 1264-92.

[20] Holt, C., F. Modigliani, J. Muth, H. Simon, *Planning Production, Inventories, and Work Force*, Englewood Cliffs, N.J.: Prentice-Hall, 1960.

[21] Huggett, M and S. Ospina, "Does Productivity Fall After the Adoption of New Technology?"mimeo, September 1998.

[22] Nilsen, O. and F. Schiantarelli, "Zeroes and Lumps in Investment: Empirical Evidence on Irreversibilities and nonconvexities," mimeo, Boston College, March 1998.

[23] Peck, S. "Alternative Investment Models for firms in the Electric Utilities Industry," *Bell Journal of Economics and Management Science*, 5 (1974), 420-457.

[24] Ramey, V. and M. Shapiro, "Displaced Capital," mimeo, University of California, San Diego, March 1998.

[25] Rothschild, M. "On the Cost of Adjustment," Quarterly Journal of Economics 85 (1971), 605-22.

[26] Smith, A. "Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions," Journal of Applied Econometrics, 8 (1993), S63-S84.

[27] Thomas, J. "Is Lumpy Investment Relevant for the Business Cycle?,mimeo, Carnegie-Mellon University, August 2000. "

[28] Veracierto, M. "Plant Level Irreversible Investment and Equilibrium Business Cycles," Working Paper 98-1, Federal Reserve Bank of Chicago, February 1998.

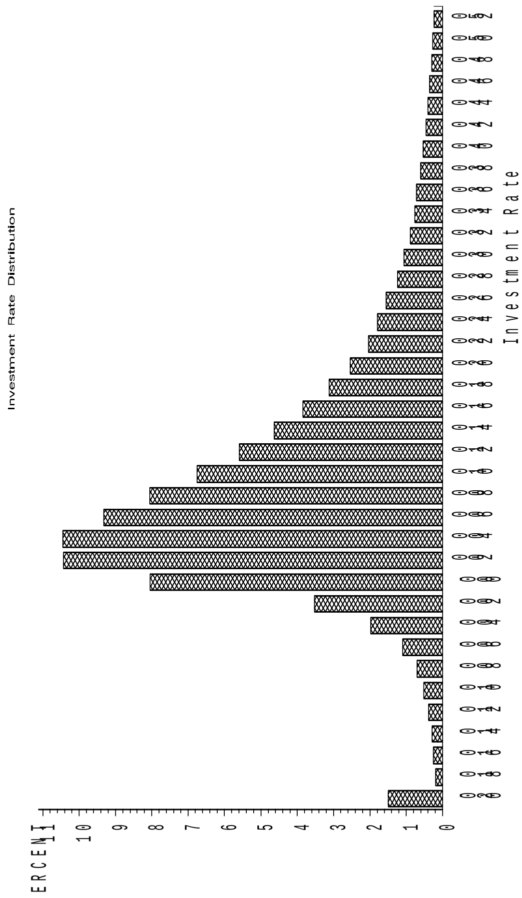

Figure 1:

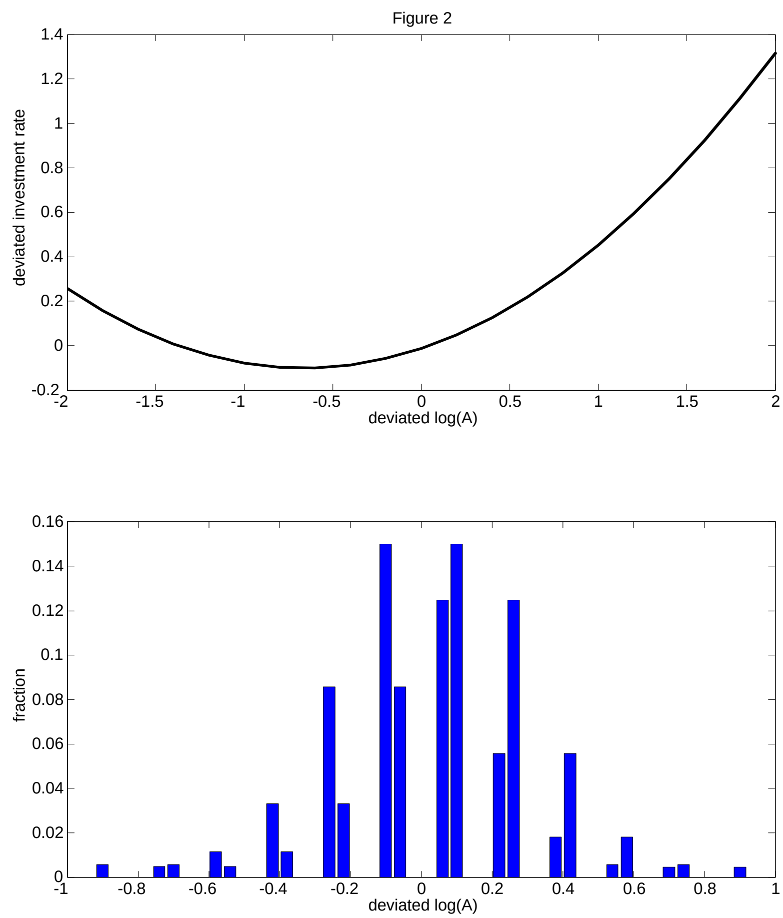
Figure 2
deviated investment rate
deviated log(A)
fraction
deviated log(A)

Figure 2:

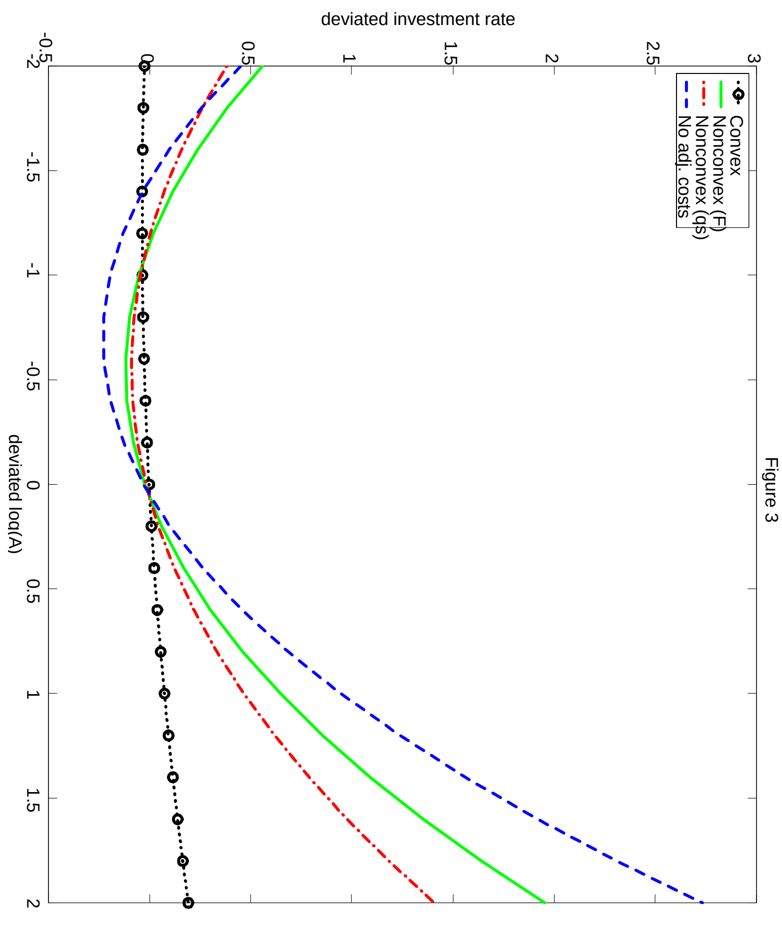
Figure 3
deviated investment rate
deviated log(A)
Convex
Nonconvex (F)
Nonconvex (qs)
No adj. costs
-0.5
0
0.5
1
1.5
2
2.5
3
-2
-1.5
-1
-0.5
0
0.5
1
1.5
2

Figure 3:

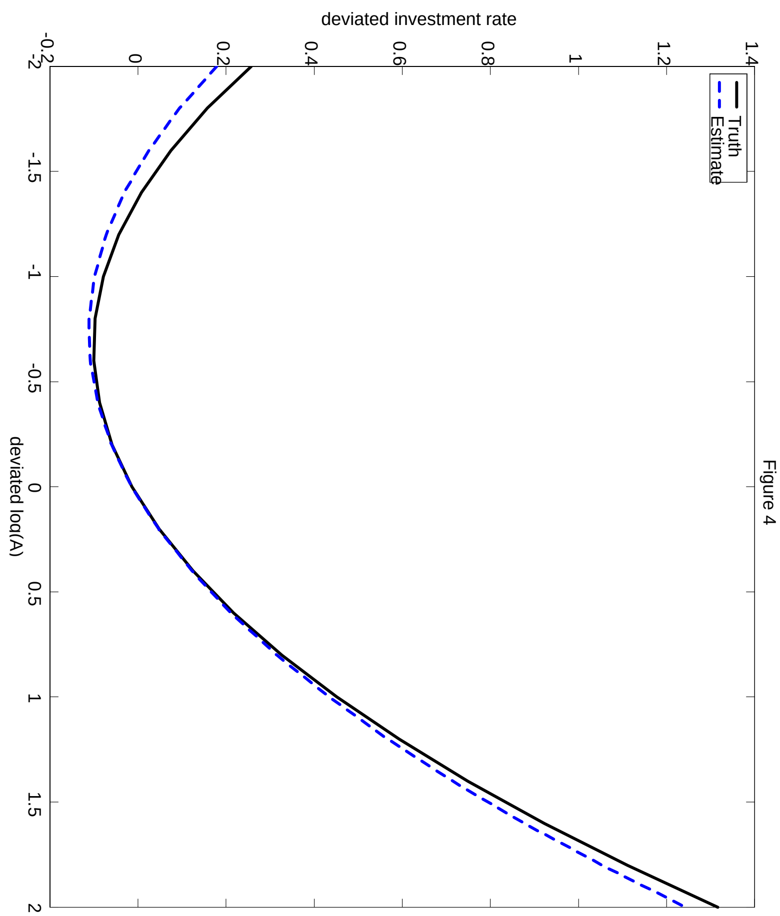
Figure 4
Truth
Estimate
deviated investment rate
deviated log(A)
-0.2
0
0.2
0.4
0.6
0.8
1
1.2
1.4
-2
-1.5
-1
-0.5
0
0.5
1
1.5
2

Figure 4:

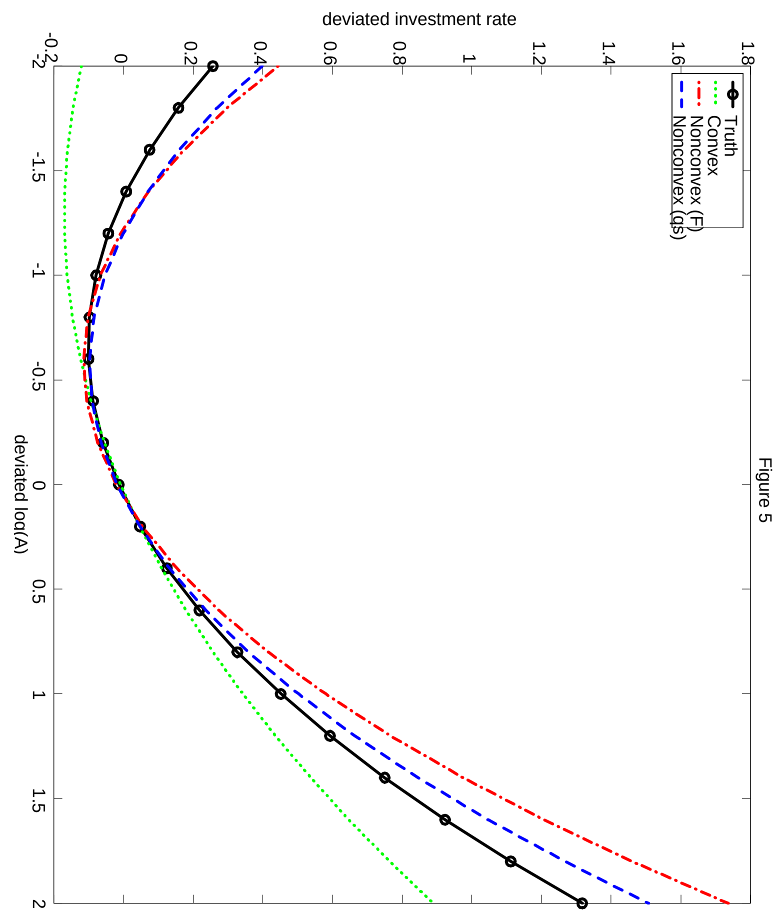
Figure 5
deviated investment rate
deviated log(A)
Truth
Convex
Nonconvex (F)
Nonconvex (φs)

Figure 5:

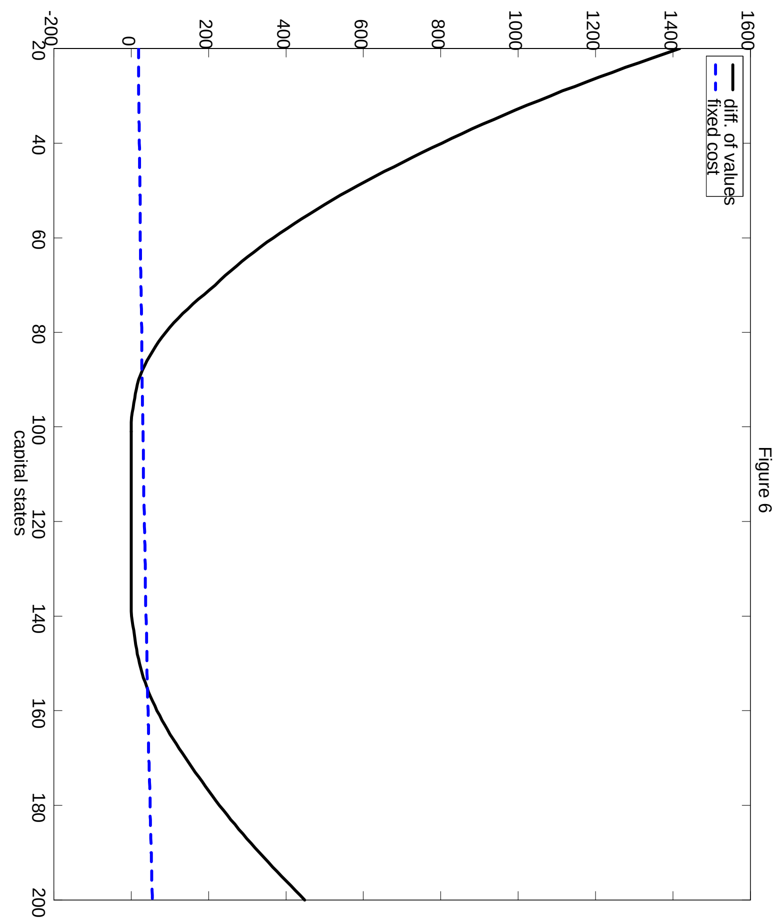
Figure 6
diff. of values
fixed cost
capital states
-200
0
200
400
600
800
1000
1200
1400
1600
20
40
60
80
100
120
140
160
180
200

Figure 6:

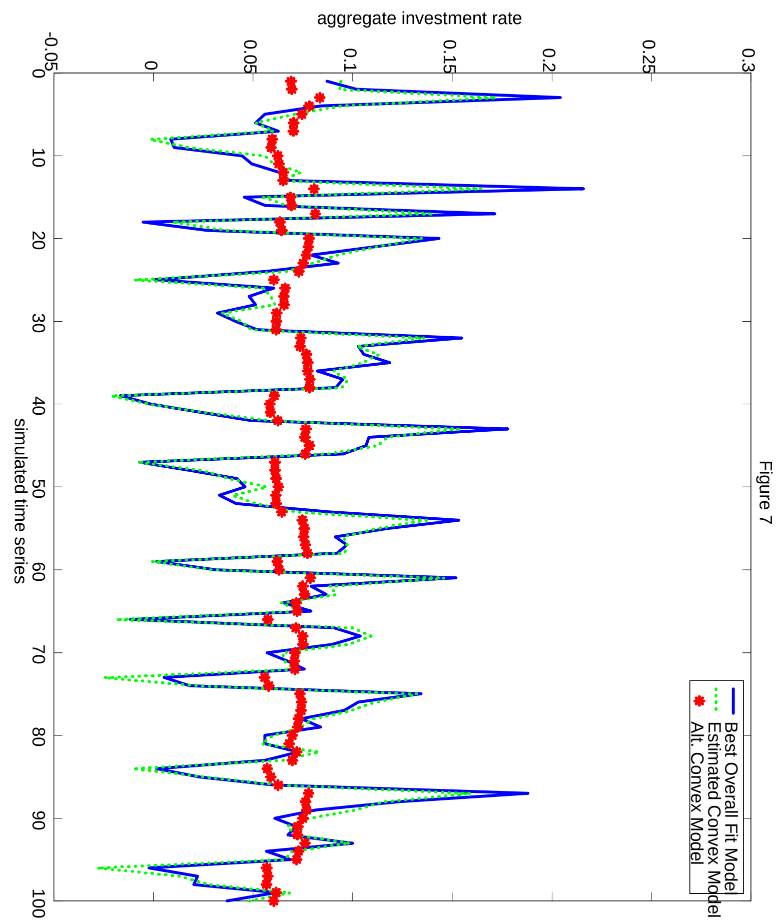
Figure 7
aggregate investment rate
simulated time series
Best Overall Fit Model
Estimated Convex Model
Alt. Convex Model

Figure 7: