On the Nature of Capital Adjustment Costs *

Russell W. Cooper
Department of Economics, Boston University,
270 BSR, Boston, Mass. 02215, USA

John C. Haltiwanger
Department of Economics, University of Maryland,
College Park, Md. 20742, USA

October 12, 1998

Abstract

This paper studies the nature of capital adjustment at the plant level. We consider a variety of models of the adjustment process allowing for both convex and nonconvex costs of adjusting the capital stock. Using plant level data, we attempt to match the moment implications of the models with those found in the data. Tentatively, we find that a model which mixes both types of adjustment costs fits the data better.

1 Motivation

The goal of this paper is to understand the nature of capital adjustment costs and thus investment activity. This topic is central to the evaluation of

*This paper is preliminary and incomplete. The authors thank the National Science Foundation for financial support. Andrew Figura and Jon Willis provided excellent research assistance in this project. We are grateful to Ricardo Caballero, Fabio Canova, Jan Eberly, Zvi Eckstein, Simon Gilchrist, John Leahy and Jon Willis for helpful discussions in the preparation of this paper. Helpful comments from seminar participants at Boston University, the Federal Reserve Bank of Minneapolis, Pompeu Fabra and the NBER Summer Institute are greatly appreciated. Financial support from the NSF is gratefully acknowledged.

policies, such as investment tax credits, that attempt to influence aggregate activity. Further, investment is one of the most volatile components of output and thus understanding the determinants of investment is indispensable to any analysis of aggregate fluctuations.

Despite the obvious importance of investment to macroeconomics, it remains an enigma. The standard neoclassical model based upon convex costs of adjustment has difficulty matching many features of the data. An alternative approach, highlighted in the work of Doms-Dunne [1994], Cooper, Haltiwanger and Power [1995], Abel-Eberly [1994, 1996], and Caballero, Engel and Haltiwanger [1995], argues that nonconvexities play a central role in the investment process. The primary basis for this view, reviewed in detail below, is plant-level evidence of both large investment bursts (spikes) as well as periods of relatively low investment.

The approach of this paper is to study competing models of investment within a common structure. To do so, we specify a dynamic optimizing problem at the plant level. The model incorporates both convex and nonconvex costs of adjustment at the plant-level. The model's implications are matched with plant-level observations from the Longitudinal Research Database (LRD) as part of an estimation routine.

Our results can be summarized by reference to two extreme models, one with convex adjustment costs and one with nonconvex adjustment costs. Neither of these extremes match the data well. In particular, the data indicate that plants experience periods of inactivity as well as periods of large investment activity. The convex cost of adjustment model can not match the periods of inactivity while the nonconvex cost of adjustment model generally has trouble matching the serial correlation of investment as well as its correlation with "profitability shocks". Combinations of these two types of adjustment, either by allowing for differences across plants in terms of their adjustment costs or by allowing a mixture of the two adjustment cost models within a plant, say due to capital heterogeneity, fit the data much better.[1]

2 Evidence

Before proceeding to a detailed examination of competing models, it is useful to have an understanding of some basic features of the data. For this discus-

[1] Ultimately, we will estimate both parameters of the adjustment process for different types of plants as well as the distribution of plant types.

sion, our data is a balanced panel from the Longitudinal Research Database consisting of approximately 10,000 large, manufacturing plants that were continually in operation between 1972 and 1991. The LRD provides detailed information on output flows, employment and investment at the plant level that enables us to evaluate the nature of capital adjustment. Our analysis generally focuses on investment rates, defined as gross investment in machinery and equipment divided by the real capital stock at each of our plants.[2] As discussed in Cooper, Haltiwanger and Power [1995], the investment activity from our sample of plants is representative of investment for total manufacturing.

2.1 Moments of the Data

Some of the main features of the data are summarized in Table 1. First, note that 10.4% of the (plant,year) observations entail zero investment (investment rate less than 1%) in equipment and machinery. Relatedly, about 10% of the observations indicate zero adjustment in employment at a plant: i.e. employment growth less than 1% in absolute value. Thus the data exhibit significant inaction in both the adjustment of capital and labor.

These observations of inaction are complemented by periods of rather intensive adjustment in both the capital and labor inputs. In the analysis that follows we term episodes of investment rates or employment rates in excess of 20% **spikes**.[3] For example, investment rates (measured by the ratio of investment expenditures on equipment and machinery to the stock of this capital) exceed 20% in about 14% of our sample observations. On average these large bursts of investment account for about 38.5% of total investment activity. Further, the absolute value of employment adjustment rates exceeds 20% in 16% of the observations.

Observations of this nature are found in data from other countries as well.[4] For example, Nilsen and Schiantarelli [1998] study investment in Norwegian

[2] Cooper, Haltiwanger and Power [1995] describes the measurement of the real capital stock at the plant level. Future work will look at investment in structures as well as retirements.

[3] 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.

[4] Related evidence on the lumpy nature of investment for Colombia is provided by Huggett and Ospina [1998].

Table 1: Summary Statistics

Variable	LRD
Investment Rate	10.2%
Inaction Rate: Investment	10.4%
Spike Rate: Investment	13.9%
Inaction Rate: Employment	10.4%
abs(Employment Growth) >.2	16.0%

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.[5] Further they find that investment rates exceeding 20% arise in about 10% of their observations and account for about 38% of total equipment investment.

2.2 Hazard Functions

Another convenient way to summarize the data is through the computation of hazard functions which relate the probability of some activity to the state of the plant. As such, these functions are extremely useful in modeling the probability of some discrete event. While the empirical formulation taken in this paper does not rest upon the definition of discrete events, the hazard function approach will still be useful in analyzing our results. That is, the hazard functions are simply viewed as summarizing an important aspect of the data that our models should match.

Cooper, Haltiwanger and Power [1997] compute hazard functions for investment spikes, defined as observations in which the plant level investment (in equipment) rate exceeds 20%. While the choice of the actual critical investment rate is not mandated by the theory, the point is to capture the idea that plants experience large bursts of investment activity. The timing between these bursts is then characterized by the hazard function.

As discussed in some detail in Cooper, Haltiwanger and Power [1997], the presence of unobservable heterogeneity across plants significantly complicates

[5] Nilsen and Schiantarelli provide a detailed breakdown of investment patterns according to the function of the production unit and various firm characteristics.

the analysis of the hazard functions. The approach adopted in that paper is to allow for a small number of "types" and estimate the hazard for each type.

These hazards are presented in Cooper, Haltiwanger and Power [1997, Figure 11]. An important property of these hazards is that there is some evidence that they slope upwards. This is particularly true for two of the three types of adjustment processes. Since the theory model with nonconvex costs of adjustment generally predicts upward sloping hazards, these results are consistent with those models.

A second hazard function formulation is studied in Caballero, Engel and Haltiwanger [1995]. In that model, a plant-year specific gap variable is postulated which indicates the difference between the actual and desired capital stock of the plant. While the actual stock is observed, the desired capital stock is derived from an auxiliary model that links firm observables to this target capital stock. Their model then predicts a relationship between the size of this gap and the probability of acting.[6]

Caballero, Engel and Haltiwanger [1995, Figure 8] provides additional evidence on investment adjustments. In particular, larger values of the gap measure translate into higher probabilities of action. This is viewed as evidence of non-linearities in the adjustment process such that larger gaps are more likely to be filled; a prediction of models with nonconvexities in the costs of adjustment.[7]

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:

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

[6] In fact, the model is identical to one where the investment rate is a non-linear function of the gap.

[7] Caballero, Engel and Haltiwanger [1997] analyze employment adjustment using a similar structure and find evidence of non-linearities in the employment adjustment process as well.

where $K' = K(1-\delta) + I$. Here unprimed variables are current values and primed variables refer to the future.

In this problem, the agent chooses the level of investment, denoted I, which becomes productive with a one period lag. Current profits are given by $\Pi(A, K)$, where the labor input has been optimally chosen, a shock to profitability is indicated by A and K is the current stock of capital.[8]

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 2, 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 3. 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

The competing models of adjustment that we explore share certain important features. Throughout the analysis, the profit function is specified as

$$\Pi(A_t, K_t) = A_t K_t^{\theta}.$$

A key parameter is thus the curvature of the profit function.

This parameter was estimated from the our panel of plants from the LRD. In particular, we assume that there are both aggregate (a_t) and plant specific profitability shocks (ε_t),with $A_t = a_t \varepsilon_t$. Profits and capital stocks were calculated at the plant level.

[8]Note that this model does not explicitly include a labor input. For the purposes of our analysis, we are assuming that the optimal labor choice is reflected in the profit function.Clearly this formulation implies that there are no costs to adjusting the labor input. In the estimation of the model, we start with a more primitive profit function which allows us to infer the optimal labor choice from the variables describing the current state of the plant, (A, K). An interesting aspect of our quantitative analysis will be to see if it is possible to match observations on the distribution of microeconomics employment changes without costs of adjusting this input.

This analysis yields three important elements. First, θ is estimated at .69.[9] Second, this regression yields both an aggregate shock and idiosyncratic shock series at the plant level. These series provide important information for the solution of (1) since this dynamic programming problem requires the calculation of a conditional expectation of future profitability.[10] The discretization of these random variables, which is necessary for solving the dynamic programming problem, is constructed to reproduce both the standard deviation and serial correlation from these two sources.

In particular, 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 the a_h 10% above steady state and a_l 10% below and estimate the diagonal elements in T at .8. The idiosyncratic shocks take 7 possible values and are also serially correlated. The transition matrix for these shocks is chosen to reproduce statistics from the idiosyncratic profitability shock series.

For the remaining parameters, we set the annual discount rate (β) at .9 and the annual rate of depreciation at .15.[11] This depreciation rate may seem high. As argued in Cooper, Haltiwanger and Power [1997] this depreciation rate includes the rate of obsolescence created by deterministic technical progress.

3.2 Convex Costs of Adjustment

The traditional investment model assumes that costs of adjustment are convex. 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

[9] Here profits are measured as the real value of output less costs of labor and materials. This estimate is essentially unchanged once fixed effects are introduced. Further, though capital is predetermined but not exogenous and thus potentially correlated with past profitability shocks, from simulation results (using the model with θ=.69 and then reestimating the profit function from simulated data) it appears that this bias is minimal.

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

[11] As this work proceeds, we will evaluate the robustness of our results to variations in these parameters.

function with respect to capital and the cost of capital (p). That is, the solution to (1) implies

$$i = (1/\gamma)[\beta EV_k(A', K') - p]$$

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 observed.

In the case of 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, $V_k(A, K) = V(A, K)/K$.[12]

To analyze the quantitative properties of this model, we parameterize this function with $\gamma = 5$.[13] With this specification, the dynamic programming problem is solved numerically using a value function iteration routine. The resulting policy rules are used to create a panel data set which, for each plant, includes a time series on investment, capital, output and profitability shocks.

For this model, we compute summary statistics for the same key moments as presented for the LRD. These appear in the second column of Table 3, where irate_{it} denotes the investment rate at plant i in period t and π_{it} is the profit rate.

Note that this model is unable to produce observations of zero investment; due to the convex cost the capital stock is continually being adjusted. The convex model is capable of producing large investment burst arising from realizations of large idiosyncratic shocks.

For this model, we can also create the analogues of the hazard function from Cooper, Haltiwanger and Power [1997] and the relationship between the disequilibrium in the capital stock and the investment rate studied by Caballero, Engel and Haltiwanger [1995].

For this specification, the Kaplan-Meier hazards are downward sloping: there is a high probability of an investment burst in the period following

[12] This point is made by Hayashi [1982]. Of course, given that the estimate of the curvature of the profitability function is .69 rather than 1, any q theory based investment regressions are misspecified. Cooper-Ejarque [forthcoming] are looking at the implications of this point.

[13] The choice of $\gamma = 5$ is consistent with evidence from q theoretic regressions, as in Gilchrist-Himmelberg [1988] though higher than the estimate of nearly 20 reported in Hayashi [1982]. At this point, this parameterization is used for illustration purposes. Estimation of this parameter is an important element of our structural estimation exercise.

a large investment episode and then the probability of another burst falls substantially as capital ages. As in Caballero, Engel and Haltiwanger, one can look at the hazard function computed using the difference between the optimal level of investment in the absence of any costs of adjustment (denoted by $k^*(A)$) and the actual capital. The resulting hazard is shown in Figure 1 and is relatively flat (away from the neighborhood of a zero gap).

While the convex cost of adjustment model is unable to explain either the preponderance of zero investment episodes or the estimated hazards, empirically there is ample evidence that plant level investment is often relatively small and that investment bursts "spillover" across periods. This is indicative of some form of convex adjustment that tends to smooth investment once the process of capital accumulation has begun. Hence it is important to include convex adjustment costs in our model of investment behavior.

3.3 Nonconvex costs of Adjustment

Following the specification explored in Cooper, Haltiwanger and Power [1997] , there are two forms of nonconvex costs of adjustment that directly affect profit flows. During a period of investment, the plant loses a fixed amount of output (F) and a fixed share of output ((1-λ)). The model allows for two types of nonconvex adjustment costs since the response of investment to shocks depends critically on whether the cost of adjustment is of the opportunity cost variety ($\lambda < 1$) or not.[14] From the analysis and results in Cooper, Haltiwanger and Power [1997], if all costs are nonconvex, then the plant will undertake large, infrequent bursts of investment.

These costs represent the need for plant restructuring, worker retraining and organizational restructuring during periods of intensive investment. Both of these costs are assumed to be independent of the level of capital adjustment undertaken by the firm. Generally, these nonconvex costs of adjustment are intended to capture indivisibilities in capital, increasing returns to the installation of new capital (often modelled as shutdowns) and increasing returns to retraining and restructuring of production activity.

For example, in their study of investment in large scale computer systems, Ito, Bresnahan and Greenstein [1998] find evidence of lumpy investment. Their analysis of the costs of adjusting the stock of computer capital

[14] See Cooper, Haltiwanger and Power [1997] for a discussion of the cyclical properties of investment and its relation to the specification of nonconvex adjustment costs.

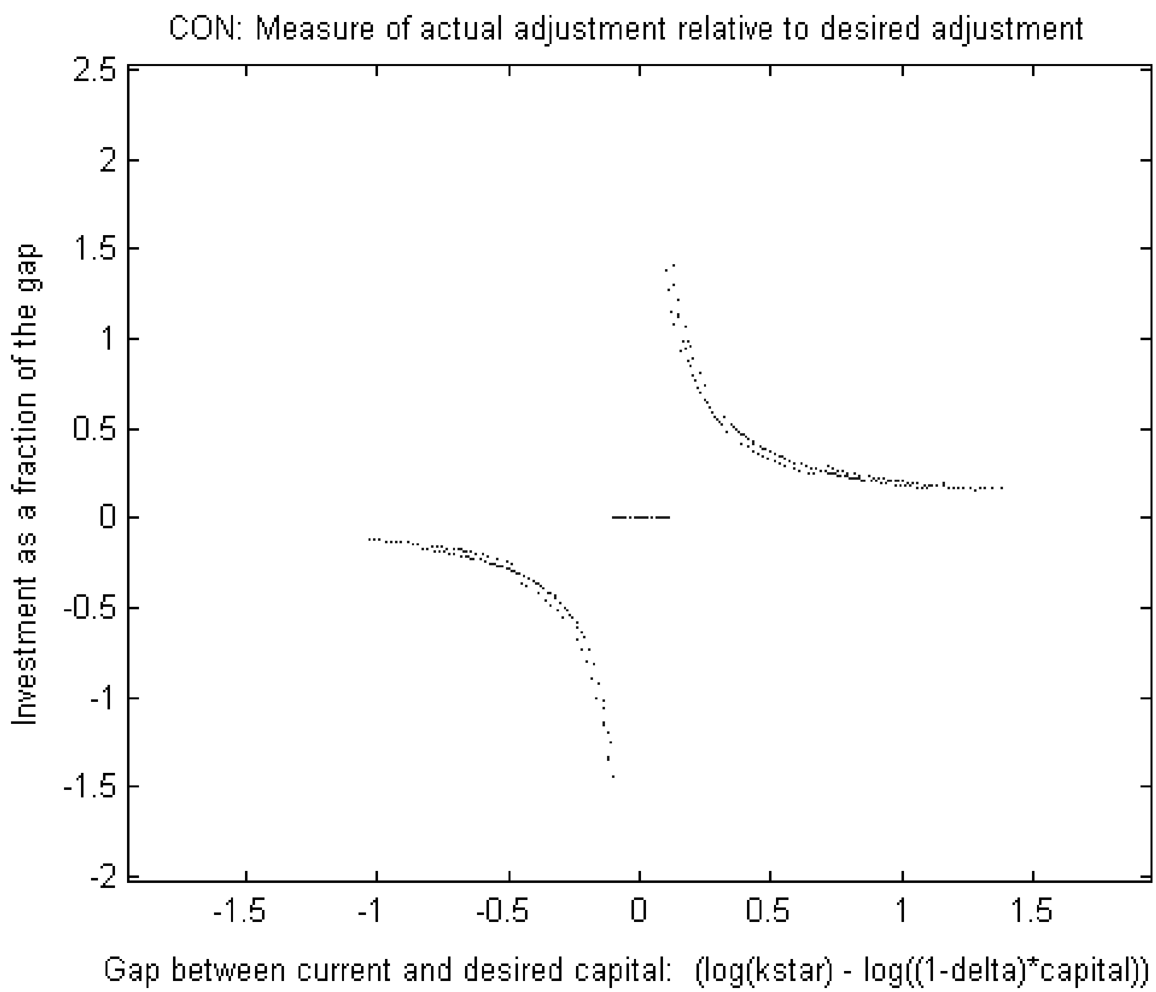
CON: Measure of actual adjustment relative to desired adjustment
2.5
2
1.5
1
0.5
0
-0.5
-1
-1.5
-2
Investment as a fraction of the gap
-1.5
-1
-0.5
0
0.5
1
1.5
Gap between current and desired capital: (log(kstar) - log((1-delta)*capital))

Figure 1:

include items which they term ".. intangible organization capital such as production knowledge and tacit work routines." To some extent, these costs are independent of the actual size of investment expenditures.[15]

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:

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

and

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

In this second optimization problem are the two costs of adjustment that are independent of the investment activity of agent as described above. Here the cost of new investment goods is normalized at 1. For the initial quantitative analysis of this model, we set F=0 and $\lambda = 0.9$.[16]

The moment implications for this specification are summarized in Table 3. Note that this form of adjustment cost produces both inaction (81% of the observations) and large bursts of activity (19% of the observations are investment spikes, i.e. irate_{it}>.2). In the absence of profitability shocks, there would be an underlying deterministic replacement cycle in which capital is replaced every T periods. The introduction of shocks essentially makes T stochastic but, as is apparent from Table 3, the investment patterns are produced by periods of inaction and then large bursts of investment activity as old capital is replaced.

Also, investment is "countercyclical" in the sense that the correlation between profitability and investment is negative, in contrast to the positive correlation found in the data. The model produces the negative correlation

[15] An alternative formulation appears in Gomes [1998] who argues that the nature of obtaining financing creates a nonconvexity in the investment process. Interestingly, this type of argument leads one to consider fixed costs at the firm rather than plant level.

[16] This choice is partly arbitrary and partly influenced by the fact that it is easier to construct a reasonable first guess on λ since it is a fraction and thus need not reflect plant size and so forth.

due to the opportunity cost associated with plant shut-downs. Finally, the counterpart of the upward sloping hazard (discussed below) is the negative serial correlation in investment. Again this is in contrast to observation.

The hazard functions for this model look quite different from those generated by the quadratic cost of adjustment model. In particular, the Kaplan-Meier hazards look much closer to those reported in Figure 11 in Cooper, Haltiwanger and Power.[17] In particular, the evidence that large bursts of investment are more likely as capital ages found in the LRD reappears with this specification of adjustment costs.

Further, looking at the hazards created from the gap measure, Figure 2, there is a tendency for overshooting in this model. That is, once a plant invests, the level of capital accumulation is such that the capital stock is above its target level (i.e above k*(A)). In this way, the plant can then use depreciation to bring the capital stock back down towards the target.

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.[18] This is incorporated in the model by assuming that

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

where $p_b \geq p_s$. 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:

[17] Given that there is no unobserved structural heterogeneity in our model, it is appropriate to match the hazards created by purging the data of unobserved structural heterogeneity. Hence the attempt here is to match the hazards reported in Figure 11 of Cooper, Haltiwanger and Power rather those appearing in Figure 9.

[18] In fact, Abel and Eberly [1994] include other forms of nonconvex adjustment in their model.

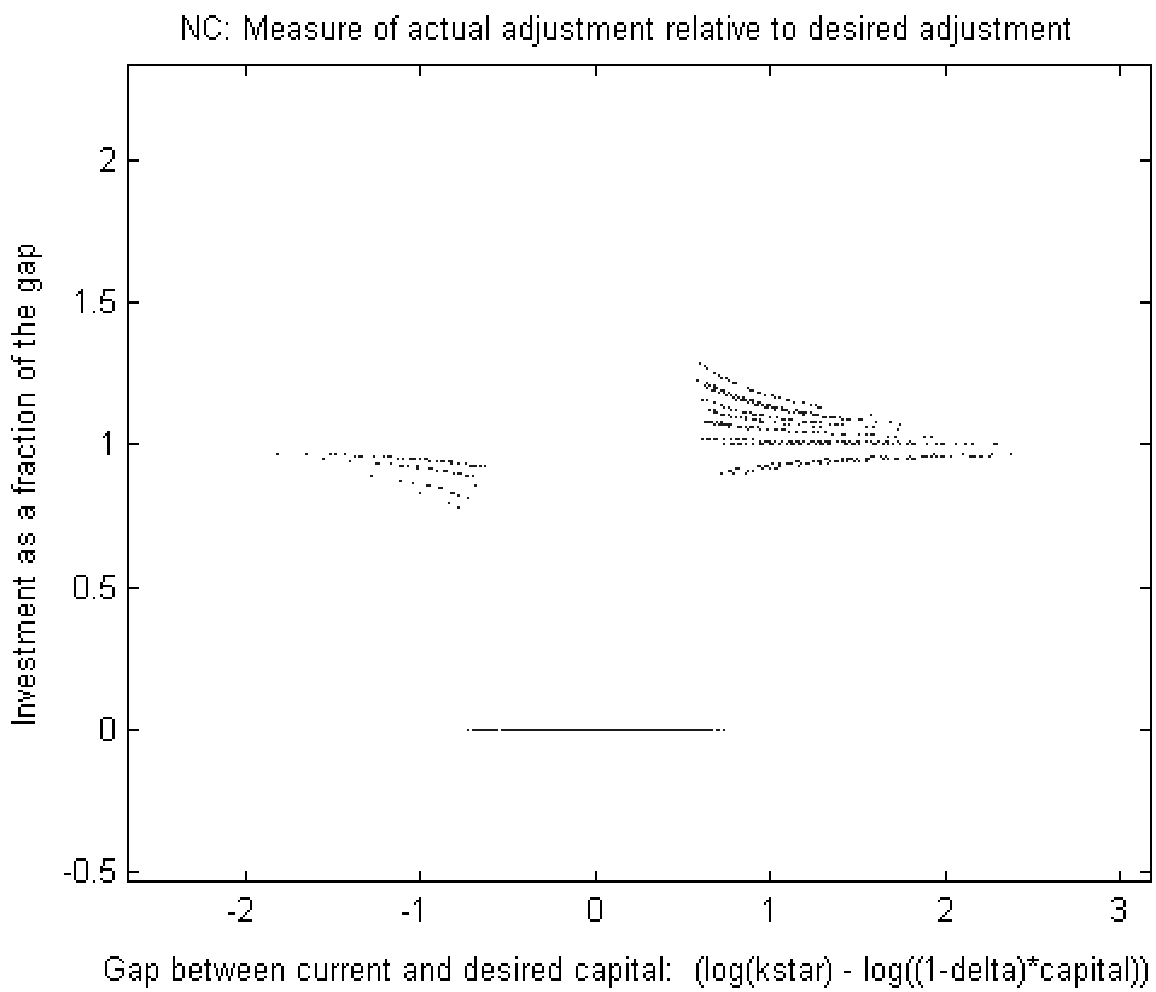
NC: Measure of actual adjustment relative to desired adjustment
Investment as a fraction of the gap
Gap between current and desired capital: (log(kstar) - log((1-delta)*capital))
2
1.5
1
0.5
0
-0.5
-2
-1
0
1
2
3

Figure 2:

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

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

and

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

For the purposes of model evaluation, we set p_s=.75 so that the sale of used capital entails a loss of 25% in value. Relative to some existing studies of particular sectors, such as the aerospace industry by Ramey and Shapiro [1998], this fraction of lost value is reasonable.[19]

The moment implications are summarized in Table 3. This parameterization creates substantial investment inactivity (39%) as well as bursts of investment (nearly 24%). Further, in contrast to the NC specification, the fraction of observations in which the investment rate exceeds 40% (60%) is lower than the fraction exceeding 20% percent. Finally, this form of adjustment cost generates both positive serial correlation in investment (reflecting the underlying driving processes) and positive correlation between investment and profitability.

The relationship between investment rates and the capital gap, Figure 3, is different from the earlier models, in an interesting way. Outside of the area of inaction, there is substantial investment. However, plants have an incentive to undershoot their target since they pay a cost (through the irreversible nature of investment) of having "too much" capital. In this way, these adjustment hazards, computed using the difference between k*(A) and the actual capital stock are quite different from those reported for the NC model.

3.5 Summary

Our models and quantitative findings are summarized in the following tables. The first displays the parameter values chosen for each of the three competing models:

[19] As noted in Ramey and Shapiro, in a more complete model, p_s is endogenous and thus dependent on cyclical and sectoral factors.

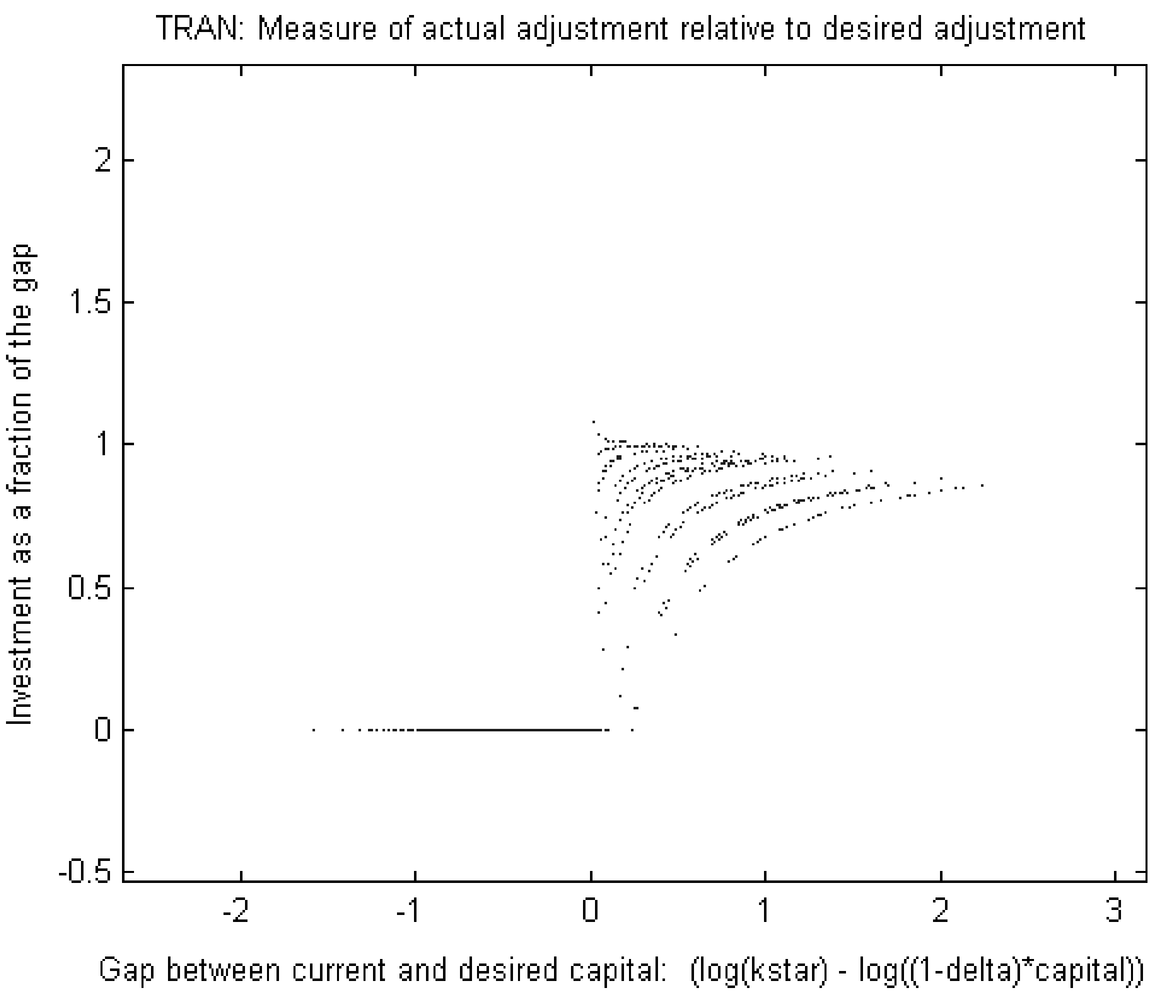

TRAN: Measure of actual adjustment relative to desired adjustment
Investment as a fraction of the gap
Gap between current and desired capital: (log(kstar) - log((1-delta)*capital))
2
1.5
1
0.5
0
-0.5
-2
-1
0
1
2
3

Figure 3:

Model	γ	ϕ	λ	p_s	p_b
CON	5	2	1	1	1
NC	0	2	.95	1	1
TRAN	0	2	1	.75	1

Table 2

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. As discussed above, the quantitative analysis of these models yields the following moments.[20]

Moment	LRD	CON	NC	TRAN
irate_{it}<.01: frac0	.104	0	.8099	.3911
irate_{it}>.2	.139	.0317	.1901	.2412
irate_{it}>.4	.044	0	.1901	.1559
irate_{it}>.6	.022	0	.1901	.0437
Corr(irate_{it},irate_{it-1})	.1328	.9948	-.1570	.3948
Corr(irate_{it}, π_{it})	.1142	.9260	-.0218	.3625

Table 3

Clearly, none of the models alone fits these key moments from the LRD. Non-convex costs of adjustment and/or irreversibility are required to match the observed investment inactivity at the plant level. However, the pure non-convex model does not create the positive correlation between profitability and investment nor the serial correlation in investment observed in the plant-level data.

4 Estimation

[20] A more complete defense for selecting these moments is given below.

The logical approach to estimating the parameters of the model is through structural estimation. With this methodology one takes a vector of parameters and solves the dynamic programming problem given in (1). The resulting policy functions are then used to create moments, either aggregate or at the plant level, which can then be matched against the comparable moments from the LRD.

An important element in our empirical approach is dealing with unobserved heterogeneity. Here we consider two forms that this heterogeneity may take.

First, as suggested by the hazards estimated in Cooper, Haltiwanger and Power [1997], there may be multiple types of plants.[21] We take as given three models of the capital adjustment process which is consistent with the evidence presented by Cooper, Haltiwanger and Power.[22] In fact, our first estimation exercise assumes that there are three types of plants as described in the previous subsection. Given the parameterization of these three types, we then estimate the fraction of each type in our sample from the LRD.[23]

The second form of heterogeneity that we consider is across types of capital goods. It is reasonable to conjecture that the nature of the capital adjustment process is itself capital specific. Unfortunately, the LRD provides only a measure of total expenditures without an adequate breakdown into types of capital goods. Thus, our measure of investment activity is an aggregate over expenditures on heterogeneous capital goods. From this perspective, it is reasonable to model the plant level capital adjustment problem using a combination of convex and nonconvex adjustment costs. This is pursued here in the second subsection.

4.1 Estimation of Fractions of Types

Table 3 above summarizes the aggregate moment implications of the three models discussed above relative to the observations from the LRD. While there are obviously other moments to consider, those included in the table highlight some of the economically relevant predictions of these models.

[21] This is partly an interpretation of the finding that there are three types of spells between investment bursts.

[22] However, at this stage, it is not clear that the types specified here match those found in the data.

[23] The logical continuation of this exercise, of course, is to estimate both the structural parameters and the fractions of each type.

Given these three types of plants, we now turn to the estimation of the fraction of each type. This is a simple estimation routine in which the fraction of each type is chosen to bring the aggregated simulated moments close to the actual moments reported in the LRD. Since there are no interactions across our plants (as there would be in a general equilibrium model with endogenous wages, prices and interest rates), we are left to select two positive proportions whose sum is less than 1 such that the convex combination of the moments in Table 3 are close to those from the LRD.

This exercise leads to an estimate that 81.9% of the plants are of type TRAN and the remainder are type NC. Note that with these fractions the simulated aggregate moments more closely matches the aggregate moments computed from the LRD, as reported in Table 3. Still, the convex combination yields too much inaction (46%) as well as correlations between investment and profitability that are too high relative to the data.

4.2 Structural Estimation of a Mixed Model

The final phase of this analysis is to consider a model in which both convex and nonconvex adjustment processes are feasible within a given plant instead of across plants as in the previous subsection. In this case, consider a model which combines the specification labelled CON and NC above.[24] The intuition is to capture the lumpiness and the inactivity found in the NC specification with the smoothing of investment created by the convex aspect of adjustment costs. Combining these processes may be appropriate due to mixtures of capital types within a plant and also reflecting a mixture of adjustment costs for even a single type of capital.

Specifically, assume that the dynamic programming problem for a plant is given by:

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

where the superscripts again refer to two choices of acting "a" and inaction "i". These options, in turn, are defined by:

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

[24] Clearly, combining this with a gap between the buying and selling price of capital is of interest as well. Goolsbee-Gross [1997] provide some evidence for a mixed model in their discussion of investment activity in the airlines industry.

and

$$V^a(A,K) = \max_I \Pi(A,K)\lambda - I - \frac{\gamma}{2}(I/K)^2 K + \beta E_{A'|A} V(A',K')$$

Our analysis of this model proceeds as before. We have specified some parameters of the model (β, δ, θ) for the functionsal forms discussed above. Further, we retain our simple specification of the profit function, $\Pi(A,K) = AK^\theta$ with θ=.69.

For the structural estimation, we focus on two parameters, (λ, γ),which characterize the magnitude of the nonconvex and the convex components of the adjustment process. These parameters are estimated using the following routine. For arbitrary values of (λ, γ),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. Moments are then calculated from this simulated data set in exactly the same fashion as in the actual data.[25] The parameter vector that minimizes the sum of squared differences between actual and simulated moments represents yields our estimates.[26]

The key in this procedure is the choice of moments. For this exercise, we have focused on two moments from the LRD: frac0 and corr($\text{irate}_{it}, \pi_{it}$). Here frac0 measures investment inaction, defined earlier as the fraction of observations in which the rate of purchase of new equipment is within 1% of zero. This moment is clearly economicially of interest since it is precisely the inaction on the investment side that is a distinguishing feature of the competing models of investment. Further, through our experimentation, it is clear that this moment is sensitive to the key parameter of the non-convex cost of adjustment,λ.

The second moment, denoted corr($\text{irate}_{it}, \pi_{it}$) is the correlation between the plant specific profitability shock and the plant specific investment rate.[27] This moment was chosen due to its prominance in the q-theory literature as a measure of the marginal return to investment. Further, this correlation is quite sensitive to the convex cost of adjustment parameter (γ). Finally,

[25] The first 20 simulated periods was dropped to remove artificial effects of initial conditions.

[26] This minimization is performed using a simulated annealing routine to avoid local minima.

[27] This correlation is calculated in both the LRD and the simulated data. Specifically, this is the correlation between plant specific profitability shocks and idiosyncratic investment once year and plant specific means have been eliminated.

this moment seems key in distinguishing the "procyclical" aspect of adjustment due to convex costs from the "countercyclical" tendency of undertaking lumpy investment when the opportunity cost is low, as in recessions.

Using these moments, the estimated parameters are: ($\lambda = .9046, \gamma = 4.8756$).[28] This estimate of γ is certainly consistent with recent estimates from convex cost of adjustment models. In terms of matching observed moments from the LRD, we find:

Moment	LRD	Estimated Model
irate_{it}<.01	.104	.0317
irate_{it}>.2	.139	.063
irate_{it}>.4	.044	0
irate_{it}>.6	.022	0
Corr(irate_{it},irate_{it-1})	.1328	.9582
Corr(irate_{it}, π_{it})	.1142	.1081

Table 4

The estimated model captures the correlation between plant-level investment and profitability observed in the data and comes close to matching the observed inaction. Of course, the parameters were chosen to match these moments as closely as possible. The estimated model clearly fails to capture the observed serial correlation in investment.

5 Conclusions (tentative)

The goal of this paper is to analyze capital dynamics through competing models of the investment process. The methodology is to take each model of the capital adjustment process and solve the dynamic optimization problem at the plant level. Moments created by simulations using the resulting policy functions are then used to fit the models to the data.

Our empirical results point to the mixing of models of the adjustment process. The LRD indicates that plant exhibit periods of inactivity as well as large investment bursts. While these observations certainly motivate a

[28]We have not yet obtained standard errors for these estimates nor determined their sensitivity to assumed values for β and δ.

nonconvex adjustment cost model, elements of convex cost of adjustment seem warranted in order to better match other moments of the data. In particular, investment displays positive serial correlation and is positively correlated with plant-specific profitability.

From our perspective, this mixing of adjustment processes reflects both heterogeneity across plants in terms of their adjustment processes and the aggregation of different types of capital within a plant. Further analysis of plant level data as well as the use of other data sets that are more disaggregated in terms of capital type seems warranted.

Our estimation strategy requires the specification particular moments. As always in these exercises, the choice of moments is open to discussion and other moments will be explored. Further, there is the possibility of using the simulation/estimation approach to match other aspects of the data, such as the reduced form hazard functions.

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 useful to extend 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.

6 References

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

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

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

[7] Cooper, R., J. Haltiwanger and L. Power, "Machine Replacement and the Business Cycle: Lumps and Bumps," NBER Working Paper #5260, 1995 (revised 1997).

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

[9] Gilchrist, S. and C. Himmelberg, "Investment, Fundamentals and Finance," mimeo presented at the 1998 NBER Macroeconomics Annual Conference, Boston University, 1998.

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

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

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

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

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

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

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

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

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

[19] Rothschild, M. "On the Cost of Adjustment," *Quarterly Journal of Eco*nomics 85 (1971), 605-22.

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