ISSN 1936-5330

THE FEDERAL RESERVE BANK *of* KANSAS CITY
ECONOMIC RESEARCH DEPARTMENT

Euler-Equation Estimation for Discrete Choice Models: A Capital Accumulation Application

Russell Cooper, John Haltiwanger, Jonathan L. Willis

January 2010
RWP 10-04

RESEARCH WORKING PAPERS

Euler-Equation Estimation for Discrete Choice Models: A Capital Accumulation Application*

Russell Cooper[†], John Haltiwanger[‡], Jonathan L. Willis[§]

January 2010

RWP 10-04

Abstract

This paper studies capital adjustment at the establishment level. Our goal is to characterize capital adjustment costs, which are important for understanding both the dynamics of aggregate investment and the impact of various policies on capital accumulation. Our estimation strategy searches for parameters that minimize ex post errors in an Euler equation. This strategy is quite common in models for which adjustment occurs in each period. Here, we extend that logic to the estimation of parameters of dynamic optimization problems in which non-convexities lead to extended periods of investment inactivity. In doing so, we create a method to take into account censored observations stemming from intermittent investment. This methodology allows us to take the structural model directly to the data, avoiding time-consuming simulation-based methods. To study the effectiveness of this methodology, we first undertake several Monte Carlo exercises using data generated by the structural model. We then estimate capital adjustment costs for U.S. manufacturing establishments in two sectors.

JEL Classification: C13, D21, E22
Keywords: discrete choice, investment, nonconvex adjustment costs, Euler equation estimation, incomplete spells

---

* We are grateful to seminar participants at NYU, CREST, Shanghai University of Finance and Economics, the National University of Singapore, the Federal Reserve Banks of Dallas and Kansas City, the Norwegian School of Economics and Business Administration and the University of Texas. The authors thank the NSF under grant #0819682 for financial support. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the U.S. Census Bureau. All results have been reviewed to ensure that no confidential information is disclosed.
† Department of Economics, European University Institute, Florence, Italy, Department of Economics, University of Texas and NBER; Email: russellcoop@gmail.com
‡ Department of Economics, University of Maryland and the NBER; Email: haltiwan@econ.umd.edu.
§ Research Department, Federal Reserve Bank of Kansas City; Email: jonathan.willis@kc.frb.org.

1 Introduction

This paper estimates capital adjustment costs using an Euler-equation methodology. As in the recent literature, our model incorporates various forms of capital adjustment costs intended to capture the rich nature of capital adjustment at the plant level. Our goal here is to characterize these adjustment costs, which are important for understanding both the dynamics of aggregate investment and the impact of various policies on capital accumulation.

Our estimation strategy searches for parameters which minimize *ex post* errors in an Euler equation. This strategy is quite common in models for which adjustment occurs in consecutive periods. Here, following Pakes (1994) and Aguirregabiria (1997), we extend that logic to the estimation of parameters of dynamic optimization problems in which non-convexities lead to extended periods of investment inactivity. We do so in the context of the capital adjustment problem, taking into account the issue of censored observations stemming from intermittent investment.[1]

This paper thus makes two contributions. First, we obtain parameter estimates for capital adjustment costs. Second, we obtain these estimates using a novel methodology that is significantly less computationally intensive than existing estimation techniques such as simulated method of moments.

The paper begins by specifying the dynamic optimization problem at the plant level. This problem is used to generate the Euler equation that underlies our empirical analysis. The empirical strategy is then laid out in some detail. We provide results using simulated data to illustrate our contribution to resolving the problems of censored observations that commonly occur in large panel datasets with relatively few observations in the time dimension.

Finally, estimates of adjustment costs using plant-level data for two sectors (transportation and steel) from the Longitudinal Research Database (LRD) are reported. Like other methodologies for estimating adjustment costs, the Euler-equation-based approach used here finds evidence of both quadratic and non-convex adjustment costs. As in the simulated method of moments estimates reported in Cooper and Haltiwanger (2006) (hereafter CH),

[1] This is an extension of the capital adjustment problem considered in Pakes (1994) and follows the evidence on the nature of capital adjustment costs from Cooper and Haltiwanger (2006) and Caballero and Engel (1999). See Eberly (1994) for an application of similar ideas to the purchase of cars.

the estimated profit function exhibits significant curvature, reflecting market power, and quadratic adjustment costs are relatively small. The Euler-equation approach finds less irreversibility and smaller disruption costs than CH for the two sectors we study.

2 Model

The dynamic optimization model draws upon CH. The dynamic programming problem is specified as:

$$V(A,K) = \max\{V^i(A,K), V^a(A,K)\}, \qquad \forall (A,K) \tag{1}$$

where K represents the beginning of period capital stock and A is a profitability shock. The superscripts refer to active investment "a," where the plant undertakes investment to obtain capital stock K' in the next period, and inactivity "i," where no investment occurs. These options, in turn, are defined by:

$$V^i(A,K) = \Pi(A,K) + \beta E_{A'|A} V(A', K(1-\delta)) \tag{2}$$

and

$$\begin{aligned} V^a(A,K) \;=\; & \max_{K'} \{\lambda \Pi(A,K) - p_b(I>0)(K'-(1-\delta)K) \\ & + p_s(I<0)((1-\delta)K - K') - \frac{\nu}{2}\left(\frac{K'-(1-\delta)K}{K}\right)^2 K + \beta E_{A'|A} V(A',K')\}. \end{aligned} \tag{3}$$

Here $\Pi(.)$ represents profits and $I \equiv K' - K(1-\delta)$ is gross investment.

The model includes three types of adjustment costs that, as reported in CH, are the leading types of estimated adjustment costs. The first is a disruption cost parametrized by λ. If $\lambda < 1$, then any level of gross investment implies that a fraction of profits is lost.[2] The second is the quadratic adjustment cost parametrized by ν. The third is a form

[2] It is possible to modify the specification so that the non-convex adjustment costs apply once the investment rate exceeds an arbitrary threshold.

of irreversibility in which there is a gap between the buying, p_b, and selling, p_s, prices of capital. These are included in (3) by the use of the indicator function for the buying ($I > 0$) and selling of capital ($I < 0$).

Assume the profit function has the following form

$$\Pi(A, K) = \lambda^Z A K^\alpha \tag{4}$$

where Z is an indicator function equal to 1 if plant investment is non-zero and equal to 0 in periods of investment inactivity. This is a reduced-form profit function which can be derived from an optimization problem over flexible factors of production (labor, materials, etc.). The parameter α reflects factor shares as well as the elasticity of demand for the plant's output. Here A is a plant-specific profitability shock.[3]

The non-convex adjustment cost, parametrized by λ, is explicit in (4). The presence of λ means that measured profits during periods of adjustment will include this adjustment cost.

The first-order condition for the investment decision is

$$p(I) + \nu \left(\frac{K' - (1-\delta)K}{K} \right) = \beta E_{A'|A} V_2(A', K') \tag{5}$$

where $p(I) = p_b$ if $I > 0$ and capital is purchased and $p(I) = p_s$ if $I < 0$ and capital is sold. Here the expectation is with respect to A'. The uncertainty is thus over the future marginal profitability of capital as well as the likelihood of adjustment.

The left side of (5) is the marginal cost of adjustment. The right side is the expected marginal gain and includes the effects on both the intensive (the amount of capital) and extensive (to adjust or not) margins. Yet, the right side of (5) appears to ignore the effects of the choice of K' on the probability of adjustment. This is correct since the effect of capital adjustment on the probability of adjustment is evaluated just at a point of indifference between adjusting and not adjusting. That is, for each K', there are critical values of A which characterize the boundaries between adjustment and non-adjustment. Though variations in K' influence these boundaries, since the boundaries are points of indifference between

[3]In the empirical implementation, A will have both a plant-specific and a common component.

adjustment and no-adjustment, there is no effect on the value of the objective function.[4]

In models without non-convex adjustment costs, investment activity occurs each period. Estimation of adjustment cost and profit function parameters then follows the procedure introduced by Hansen and Singleton (1982). *Ex post* errors are calculated using observations on capital and profit flows. Then parameters are estimated using orthogonality conditions.

The challenge is to use that approach when investment activity does not occur each period. It is not possible to use (5) directly since the marginal value of capital is not observable.

To evaluate (5) *ex post*, we expand the $E_{A'|A}V_2(A',K')$ term until the plant's next episode of capital adjustment is observed. With non-convex adjustment costs, $\lambda < 1$, adjustment will generally not occur each period. We then replace expectations with realizations to calculate the *ex post* errors from the Euler equation.

To understand the procedure, consider a plant that adjusts in two consecutive periods, t and $t+1$. Then the *ex post* error, denoted $\varepsilon_{t,t+1}(\theta)$, from (5) is

$$\varepsilon_{t,t+1}(\theta) = \nu \frac{I_t}{K_t} + p(I_t) - \beta \left[\Pi_2(A_{t+1}, K_{t+1}) + p(I_{t+1})(1-\delta) + \nu(1-\delta)\frac{I_{t+1}}{K_{t+1}} + \frac{\nu}{2}\left(\frac{I_{t+1}}{K_{t+1}}\right)^2 \right] \tag{6}$$

where $I_t = K_{t+1} - K_t(1-\delta)$.

This error, of course, depends on the parameter vector $\theta \equiv \{\alpha, \lambda, \nu, p_s, p_b, \delta, \beta\}$. The first two terms in the *ex post* error are the marginal costs of capital in period t and the remaining terms are the marginal gains for the next period, including the marginal profitability and the marginal effects on adjustment costs next period. To be clear, since the plant adjusts in period $t+1$, from (4), $\Pi_2(A_{t+1}, K_{t+1}) = \lambda\alpha A_{t+1}K_{t+1}^{\alpha-1}$.

Of course, not all plants adjust every period when faced with significant nonconvex adjustment costs. It is not appropriate to estimate parameters based solely on the *ex post*

[4]We thank Borghan Nezami Narajabad, Jean-Michel Grandmont and Guy Laroque for questions which lead to this explanation of (5). While the policy function, $K'(A,K)$ is not continuously differentiable at a point of indifference between activity and inactivity, the right side of (5) is a conditional expectation of the marginal value of capital and thus these points of non-differentiability are of measure zero. Put differently, $E_{A'|A}V_2(A',K') = \int_{A\in adjust} V_2^a(A',K') + \int_{A\in inactive} V_2^i(A',K')$. The effect of changes in K' on the boundaries of the sets of action and inaction disappear as the values of action and inaction are equal at these boundary points.

error from plants that choose to adjust in consecutive periods because such estimates would suffer from selection bias. Thus we need a more general condition which allows estimation of the structural parameters.

In general, if the plant adjusts in period t and subsequently in period $t+\tau$, then the *ex post* error, denoted $\varepsilon_{t,t+\tau}(\theta)$, from the first-order condition is

$$\begin{aligned}\varepsilon_{t,t+\tau}(\theta) &= \nu\frac{I_t}{K_t} + p(I_t) - \sum_{i=1}^{\tau}\beta^i\Pi_2(A_{t+i},K_{t+i})(1-\delta)^{i-1} \\ &\quad -\beta^\tau\left[p(I_{t+\tau})(1-\delta)^\tau + \nu(1-\delta)^\tau\frac{I_{t+\tau}}{K_{t+\tau}} + \frac{\nu}{2}\left(\frac{I_{t+\tau}}{K_{t+\tau}}\right)^2(1-\delta)^{\tau-1}\right].\end{aligned} \tag{7}$$

From this general expression, the first terms on the right are the marginal costs of adjustment and the remaining terms are the gains in profitability between the periods of adjustment. During the periods between adjustment, there is an effect of capital accumulation on marginal profitability. Finally, in the subsequent period of adjustment, i.e. when the spell of inactivity ends, there is a final term reflecting the effects of K_{t+1} on the marginal adjustment cost in period $t+\tau$. As before, $\Pi_2(A_{t+\tau},K_{t+\tau}) = \lambda\alpha A_{t+\tau}K_{t+\tau}^{\alpha-1}$. The inclusion of λ in this marginal profit term will allow us to estimate it from orthogonality conditions involving $\varepsilon_{t,t+\tau}$.

As in the estimation of quadratic adjustment cost models, the *ex post* errors should not be predictable. In the next section we discuss estimation of all parameters, including the non-convex adjustment cost parameter λ, using the orthogonality restrictions generated by optimization.

3 Euler-equation Estimation

Pakes (1994) argues that the logic of Hansen and Singleton (1982) can be applied to the estimation of the structural parameters in dynamic, discrete-choice problems. The application in Pakes (1994) is investment coupled with an exit decision. Aguirregabiria (1997) considers a dynamic labor-demand model. We estimate the parameters of a capital accumulation problem that differs from the one specified in Pakes (1994). Further, an important component of our contribution is to compute *ex post* errors for plants that adjust in some period t but do not adjust again within the sample.

3.1 Complete Spells

Suppose that a plant adjusts its capital stock in some period t and then adjusts again in period $t+\tau$ within the sample. We term the interval between capital adjustments in periods t and $t+\tau$ as a complete spell. In contrast, if a plant adjusts in period t but does not adjust again in the sample, then the spell is termed incomplete. Focusing first on complete spells, we discuss how to estimate the parameters of the capital adjustment model.

Using equation (7), we can compute the *ex post* errors between adjustment periods. The optimization condition of the firm given by (5) imposes structure on these errors. Optimality implies that the period t expectation of the *ex post* errors should be zero at the true value of the parameters, θ^*:

$$E_{\tau|t}[\varepsilon_{t,t+\tau}(\theta^*)] = 0 \tag{8}$$

for all t.[5] Here $\varepsilon_{t,t+\tau}(\theta^*)$ is the *ex post* error calculated from equation (7) at the true value of the parameters.

In equation (8), the expectation is conditional on all variables known in period t. The variable $t+\tau$ indicates the period of the first active capital adjustment after period t. Of course, τ is not known in period t since the adjustment decision following period t is state dependent.

The estimation of the structural parameters comes from the condition that $\varepsilon_{t,t+\tau}$ ought to be uncorrelated with period t and prior variables. This orthogonality comes from expanding the right side of (5) to incorporate the uncertainty over the future realizations of the shocks, the future discrete choices of whether to adjust, and the future intensive-margin choices of how much to invest conditional on adjusting.

[5] To construct (5) from (8) requires the use of (7) for all τ along with the associated probabilities that adjustment occurs in period $t+\tau$.

Using a vector of N variables predetermined in period t, z_t, the following orthogonality condition can be used in an estimation procedure.

$$E_{\tau|t}[z_t \varepsilon_{t,t+\tau}(\theta^*)] = 0. \tag{9}$$

The sample analog of this condition is

$$m = \frac{1}{n} Z' \varepsilon(X, \theta) = m(\theta) \tag{10}$$

where n is the number of observations (investment spells), Z is the matrix of N variables over T periods, and $\varepsilon(X, \theta)$ are the *ex post* errors calculated using the sample data, X, and the parameter vector of interest, θ.

The minimum distance estimator is the $\hat{\theta}$ that minimizes

$$\begin{aligned} s &= m(\hat{\theta})' W^{-1} m(\hat{\theta}) \\ &= \frac{1}{n^2} [\varepsilon(X, \hat{\theta})' Z] W^{-1} [Z' \varepsilon(X, \hat{\theta})]. \end{aligned} \tag{11}$$

where W is the weighting matrix for the estimator. We use the nonlinear two-stage least squares (NL2SLS) estimator as described by Amemiya (1985), where the covariance matrix of the instruments is used as the weighting matrix.

$$W_{NL2SLS} = \frac{1}{n} Z' Z \tag{12}$$

This estimator is consistent and asymptotically normal.[6] The estimated asymptotic covariance matrix of the NL2SLS estimator is

$$V(\hat{\theta}) = \hat{\sigma}^2 \left(G\left(\hat{\theta}\right) \left(Z \left(Z' Z \right)^{-1} Z' \right) G\left(\hat{\theta}\right)' \right)^{-1} \tag{13}$$

where $G(\hat{\theta}) = \frac{\partial m(\hat{\theta})}{\partial \hat{\theta}}$ is numerically computed and $\hat{\sigma}^2 = \frac{1}{n} \left(\varepsilon(X, \hat{\theta})' \varepsilon(X, \hat{\theta}) \right)$.

[6] We initially used the GMM estimator following Hansen (1982), which is based on the optimal choice of the weighting matrix, $W_{GMM} = \text{Var}\left(Z' \varepsilon(X, \theta)\right)$. However, the convergence properties of this estimator were very poor in our Monte Carlo exercises so we switched to the NL2SLS estimator, which yielded much better convergence properties.

3.2 Incomplete Spells

If the number of time periods is large enough and thus the number of complete spells is large relative to the number of incomplete spells, then the estimation procedure described in the previous section will produce a consistent estimate of θ^*.[7] But in general, panel datasets are short on the time dimension. This implies that the fraction of incomplete spells is substantial. In that case, estimation using only complete spells is not consistent. The moment condition underlying this Euler-equation estimation is based upon the expectation taken across *ex post* errors for all plants that adjusted in a given period. The omission of plants with incomplete spells constitutes a form of *ex post* selection and thus a source of inconsistency.

An alternative to excluding periods containing incomplete spells is to use all of the available data and attempt to control for the bias created by incomplete spells. This section provides a methodology and simulation results for doing so by "completing spells". To do so, we approximate the unobserved portion of the incomplete spell by estimating the marginal value of capital from the dynamic programming problem.[8]

We consider a multi-stage method for estimating the structural parameters, retaining the assumption that A is observed. In the first stage of this methodology, parameter estimates are obtained from (11) by including all complete spells as observations in the estimation. We denote these first stage estimates as Θ_1. Assuming that we have obtained all of the other structural parameters of the model from other sources, we then solve the dynamic programming problem using Θ_1.

From this solution, we can compute the expected derivative of the value function that appears in the first-order condition of the investment decision expressed in (5). This expected derivative is a function of the current profitability shock and the capital stock resulting from the investment decision in the current period, conditional on the parametrization Θ_1,

$$\psi(A, K'; \Theta_1) = E_{A'|A} V_2(A', K'; \Theta_1). \tag{14}$$

[7] Thus it is important that the asymptotic properties of the estimator be determined as the number of periods gets large not as the number of plants gets large.

[8] Up to this point, we have not had to solve the dynamic programming problem.

This function can then be evaluated using observations of A and $K' = (1-\delta)K + I$ from the final period of the sample, and *ex post* errors for all incomplete spells can be computed using the following specification:

$$\begin{aligned} \varepsilon_{t,incomplete} &= \nu \frac{I_t}{K_t} + p(I_t) - \sum_{i=1}^{T-t} \beta^i \Pi_2(A_{t+i}, K_{t+i})(1-\delta)^{i-1} \\ &\quad -\beta^{T-t+1}(1-\delta)^{T-t}\psi(A_T, K_{T+1}, \Theta). \end{aligned} \tag{15}$$

A second stage estimation then includes all complete and incomplete spells by combining the *ex post* errors from (7) and (15). Denote the resulting parameter estimates as Θ_2. This process is repeated by computing $\psi(A, K'; \Theta_2)$ and obtaining a third stage estimate, Θ_3. Additional repetitions are computed until estimates of Θ converge.

4 Monte Carlo

Before estimating this model, we construct a simulation-based exercise. There are a number of goals of this experiment. First, there is the verification of methodology to demonstrate it can consistently estimate the parameters of interest. Second, the simulations allow us to gauge the magnitude of the bias associated with focusing on complete spells and to evaluate our proposed solution for completing spells.

4.1 Methodology

We first describe the methodology used for this exercise. We then turn to the results with complete and then incomplete spells.

4.1.1 Creation of simulated dataset

A simulated data set is constructed in the following steps. First, the structural parameters of the model are chosen and the investment policy functions of the dynamic programming problem are obtained through value function iteration. The parameters of interest in this exercise are those that can be estimated via the Euler equation: $\theta = \{\alpha, \nu, \lambda, p_s\}$.[9]

We consider two different parametrizations of θ in order to assess the properties of the estimation procedure. The first case, $\theta_q = \{0.6, 2, 1, 1\}$, includes only a quadratic cost of adjustment. The second case, $\theta_n = \{0.6, 0.2, 0.8, 0.98\}$ adds asymmetry between the buying and selling prices of capital and a disruption cost of investment. This parametrization results in inactivity of capital adjustment due to the introduction of non-convex costs and most closely matches the specification of CH.

The other structural parameters of the model are chosen to be similar to those used by CH.[10] The frequency of the model is annual, so the discount rate, β, is set at 0.95. The profitability shock, A, consists of an aggregate shock and an idiosyncratic shock. Each of these shocks follows a log-normal autoregressive process. The aggregate shock process has a persistence of 0.85 and the innovation to this process has a standard deviation of 0.05. The idiosyncratic shock process has a persistence of 0.85 and the standard deviation of the innovation is 0.3. The depreciation rate, δ, is 0.07.

The model is solved using value function iteration. The capital state space is discretized onto a grid with over 1000 points that are equally spaced in log terms. The autoregressive aggregate and idiosyncratic shock processes are transformed into first-order Markov processes using Tauchen (1986). The grid for the aggregate shock state space has 9 states, and the grid for the idiosyncratic shock state space has 25 states.[11] Using the structural parameters described above, the investment policy function is obtained by iterating the value function until convergence.

The simulated panel data set is created using the investment policy functions in conjunction with the randomly drawn innovations to the aggregate and idiosyncratic profitability shock processes. The data consist of observations on the shocks, investment, capital and profits. This allows us to construct the *ex post* errors in (7).

For these exercises, we study results from three data sets. The first contains 200 plants and has a sample length of 15 periods. This is approximately the size of an average manu-

[9]In this exercise, we normalize $p_b = 1$.

[10]These estimates are discussed in CH.

[11]The grid is set such that each point in the state space has an equal probability of being selected from the ergodic distribution of the process.

facturing sector in the actual data. The second and third data sets are much larger, either with more plants or more time periods. These other data sets are useful for evaluating the behavior of our estimates in larger samples. For our Monte Carlo results, we estimate parameters from 1000 simulated datasets and report the mean and standard deviation of parameter estimates.

4.1.2 Parameter estimation

The parameter vector θ is estimated by minimizing the weighted sum of squared moments statistic in (11). The instruments are observed in the actual data and include a constant, the investment rate ($\frac{I}{K}$), the profit rate ($\frac{\Pi}{K}$) and the capital stock (K); $Z_t = \{1, \frac{I_t}{K_t}, \frac{I_{t-1}}{K_{t-1}}, \frac{\Pi_t}{K_t}, \frac{\Pi_{t-1}}{K_{t-1}} K_t, K_{t-1}\}$.[12] As discussed previously, the covariance matrix of the instruments is used as the weighting function in the nonlinear two-stage least squares estimation procedure.

4.2 Results: Complete Spells Only

For the first set of estimates, we use a sample consisting of only complete spells, which are defined by observed investment in period t and in period $t+\tau$. Other observations, which are incomplete spells, are excluded from the construction of the *ex post* errors.

This exercise serves a couple of purposes. First, with a large enough sample length, almost all spells are complete and thus this approach will converge on consistent parameter estimates. Second, with a relatively small sample, the procedure will not uncover the true parameters, illustrating a selection bias from a sample of complete spells.

The results in Table 1 show that the Euler-equation estimation procedure performs well in the case with only quadratic adjustment costs, θ_q, even in the smallest sample exercise. In this case, the true value of the production function parameter, α, is 0.6, and the scalar on the quadratic adjustment cost, ν, is set at 2.0. The means of the parameter estimates across the 1000 samples of 200 plants and 15 periods are $\{\bar{\alpha}, \bar{\nu}\} = \{0.597, 1.942\}$. The respective standard deviations across the 1000 parameter estimates are $\{0.027, 0.121\}$. Due to the

[12] In previous analysis we also used the state vector of the problem as the set of instruments, but this is not feasible unless A_t is observed.

discrete nature of the value function iteration solution, there are some situations where firms choose to remain inactive. Increasing the sample size, either by adding more plants or more time periods, reduces the standard errors.

Table 1: Euler-equation Estimates of θ_q Using All Completed Spells

α	ν	λ	p_s	plants	T	spells
0.6	2.0	1.0	1.0			
0.597	1.942	1.0	1.0	200	15	2551.0
(0.027)	(0.121)					(7.9)
0.599	1.952	1.0	1.0	500	15	6381.8
(0.023)	(0.067)					(11.5)
0.598	1.981	1.0	1.0	200	40	7463.8
(0.014)	(0.069)					(10.6)

Reported estimates are the mean value across estimates from 1000 simulated datasets. Standard deviations of the estimates are reported in parentheses. T denotes the length of the sample period. The last column reports the mean number of observations (spells) across the 1000 datasets.

Table 2 shows results based on the parametrization that most closely matches the estimates of CH, θ_n. This parametrization includes a disruption cost, $\lambda = 0.8$ and a lower value of the quadratic adjustment cost parameter than in the previous case. The disruption cost leads to more inactivity and longer observed incomplete spells, which translates into greater imprecision of the estimates due to the number of periods that must be excluded from the estimation.

For the smallest sample, the mean parameter estimates are far from θ_n, particularly for ν and p_s. Further, the standard errors on the estimates of (ν, λ, p_s) are quite large. Increasing the sample size by adding more plants does not noticeably improve the estimation of ν and p_s. However, once additional time periods are added so that $T = 100$, the parameter estimates are much closer to θ_n, and the standard errors are much lower than in the $T = 15$

Table 2: Euler-equation Estimates of θ_n Using All Completed Spells

α	ν	λ	p_s	plants	T	spells
0.6	0.2	0.80	0.98			
0.573	0.080	0.799	0.864	200	15	208.8
(0.018)	(0.090)	(0.217)	(0.111)			(22.1)
0.572	0.095	0.796	0.874	2000	15	2087.3
(0.008)	(0.033)	(0.079)	(0.034)			(172.6)
0.598	0.187	0.783	0.968	200	100	2528.2
(0.006)	(0.041)	(0.071)	(0.043)			(38.0)

Reported estimates are the mean value across estimates from 1000 simulated datasets. Standard deviations of the estimates are reported in parentheses. T denotes the length of the sample period. The last column reports the mean number of observations (spells) across the 1000 datasets.

case.

The last column labeled "spells" indicates the average number of complete spells used in the analysis. When the sample is increased from $T = 15$ to $T = 100$, the number of complete spells increases by a factor of 10. This is informative: our inability to estimate θ_n is not coming from our methodology but rather from the short sample.

These results make clear the bias induced by the selection of a sample of completed spells when T is small. To be more precise about this bias, we compute *ex post* errors from a sample comprised of completed spells from a simulation with 15 periods and 200 plants. These errors were computed using the true parameter vector, θ_n. The distribution of these errors is shown in Figure 1. In addition, we computed the errors for these plants using the necessary simulated data beyond $T = 15$ needed to complete **all** spells.[13] This distribution

[13] To understand this exercise, consider plants that do not complete a spell within a set length of the sample, such as $T = 15$. Since our simulated data now extends beyond 15 periods, it is possible to construct their *ex post* errors based upon their actual experience after $T = 15$. In other words, while the panel used in the estimation ends at $T = 15$, we can simulate the plants beyond this horizon until they next choose to invest, thus completing their spell. By doing so, we can calculate all of the *ex post* errors in the original sample.

is labeled "from extended simulation" in Figure 1. Through this exercise, we can see what the distribution of the *ex post* errors for these plants would have been had we enough data to follow them beyond the first 15 periods.

A comparison of these two distributions is informative about the bias induced by the selection of complete spells. Note that the distribution of errors from the complete-spells-only sample is shifted to the left. The means of the errors are indicated by the two vertical lines: the mean for the sample with extended simulation is essentially zero while the mean for the sample of completed spells is negative.

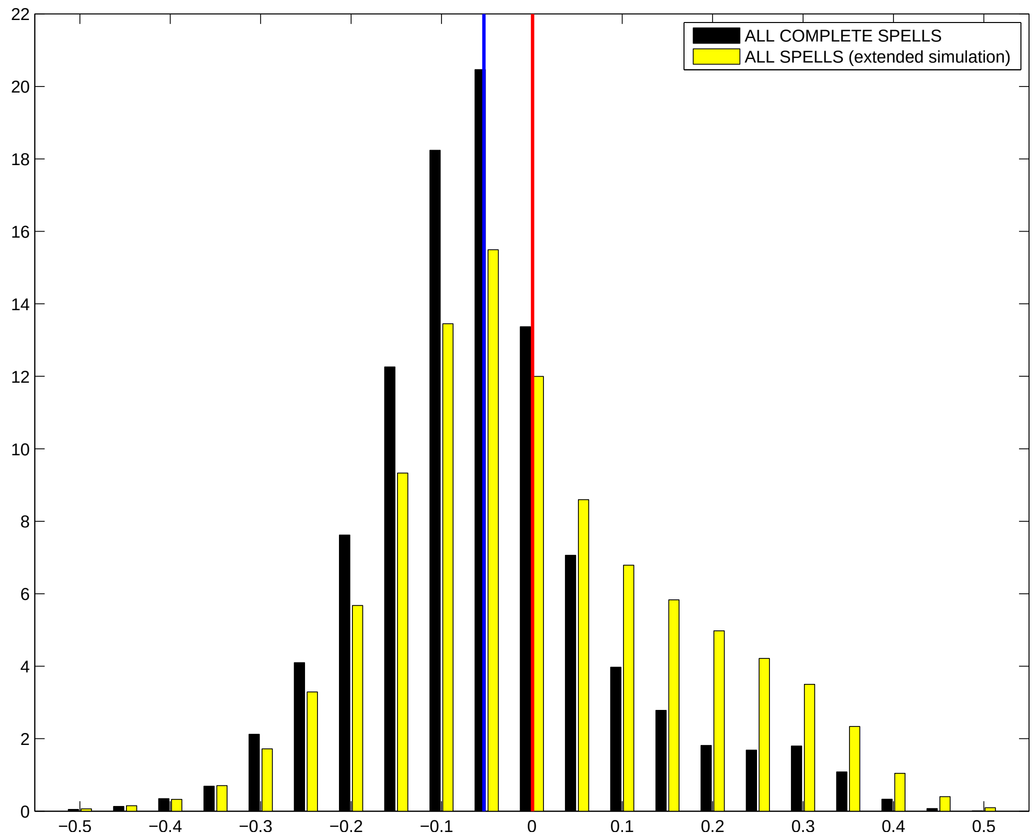

Figure 1: Distribution of *ex post* errors

To understand the source of the asymmetry in the error distribution, is it useful to return to the capital adjustment policy function. Figure 2 shows the regions of action and inaction. In this figure, the profitability shock is on the vertical axis. The horizontal axis shows the probability of adjustment. Displayed in the figure are three policy functions for different values of the capital stock. For the mean value of capital, the probability of adjustment is higher for large values of the profitability shock relative to small ones. This reflects the asymmetries in the adjustment costs from p_s and the interaction of λ with profitability. Since the shocks are distributed symmetrically around the mean, the asymmetry in the policy function translates into a selection bias on *ex post* adjustment. Since adjustment is more likely in high profitability states, the sample of completed spells has an excessive frequency of high shocks and hence negative *ex post* errors, calculated as the marginal cost less the *ex post* marginal benefit of adjustment. This leads to the differences in distributions shown in Figure 1.[14]

4.3 Results: Controlling for Incomplete Spells

The simulation exercise leads to two conclusions. First, the methodology works: for large enough samples, the minimization of (11) reproduces the structural parameters underlying the simulated data. Second, having a long sample is necessary if we are looking only at complete spells. This is an issue for empirical application since the plant-level data we use for the estimation of the model has only 19 years.

Our approach follows the one outlined in Section 3.2. We first present our estimation results and then turn to an evaluation of our technique for completing spells.

We summarize our results in a series of tables. These tables are intended to illustrate both our approach to dealing with incomplete spells and to indicate its success. All of the simulation results were obtained using θ_n, since this is the parametrization of the model that produces inaction.

Table 3 is created from a panel of plant-level data in which T is extended by simulation and thus is effectively very large. This exercise indicates that even for the smallest sample, the mean parameter estimates are quite close to θ_n. Relative to the results reported for the

[14]Linking this to the biases in coefficient estimates in Table 2 is made difficult by the inherent non-linearities of the estimation process.

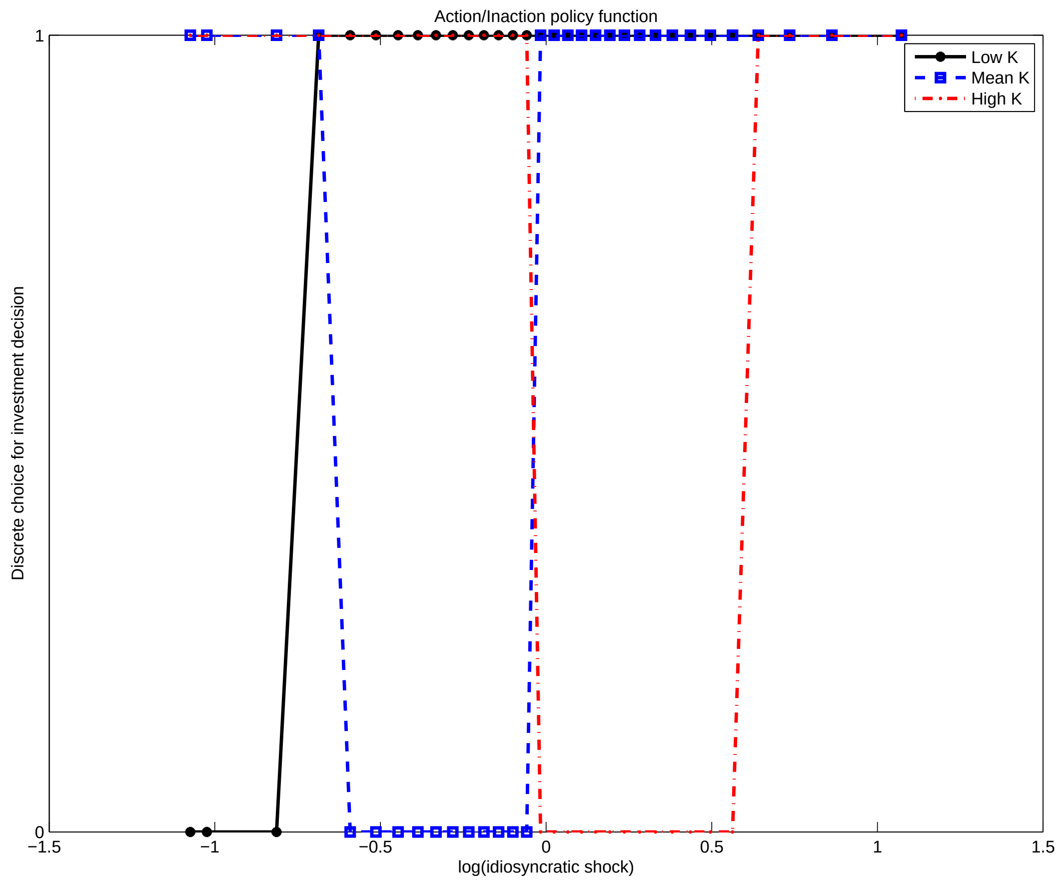

Figure 2: Action and inaction regions of the policy function

Table 3: Euler-equation Estimates of θ_n When Incorporating All Spells through Extended Simulation

α	ν	λ	p_s	plants	T	spells
0.6	0.2	0.80	0.98			
0.597	0.166	0.826	0.952	200	15	386.4
(0.013)	(0.088)	(0.170)	(0.092)			(24.8)
0.599	0.187	0.804	0.969	2000	15	3864.6
(0.005)	(0.030)	(0.052)	(0.032)			(204.3)
0.600	0.199	0.797	0.981	200	100	2728.2
(0.006)	(0.040)	(0.069)	(0.042)			(38.0)

Reported estimates are the mean value across estimates from 1000 simulated datasets. Standard deviations of the estimates are reported in parentheses. T denotes the length of the sample period. The last column reports the mean number of observations (spells) across the 1000 datasets.

small sample in Table 2, the estimates of both ν and p_s are closer to their true values. The point estimates are slightly closer to θ_n for the larger samples.

The next two tables illustrate our approach to completing spells when the simulated data used to extend the sample, and thus construct Table 3 estimates, are not available. For these tables, we only use the base sample with the short time period of $T = 15$.

To complete the spells for Table 4, we used the true parameter vector, θ_n in (14) as θ_1. Thus, there is no issue of convergence in the calculation of the *ex post* errors since there is no updating of the parameters in (14). For the short sample, there is clearly an improvement of results in estimating ν and p_s relative to the Table 2 findings. It is true, though, that the estimates of λ are not as close as in the results with completed spells only.

Table 5 displays our results using the full methodology outlined in section 3.2. In this case, we do not make use of θ_n in the calculation of (14). Instead, an initial estimate is used followed by an iteration procedure until θ converges. In practice, convergence was achieved within 10 iterations. The results indicate that our approach to dealing with incomplete

Table 4: Euler-equation Estimates of θ_n Using Correction for Incomplete Spells (correction based on true parameters)

α	ν	λ	p_s	plants	T	spells
0.6	0.2	0.80	0.98			
0.588	0.146	0.897	0.954	200	15	386.4
(0.023)	(0.109)	(0.227)	(0.103)			(24.8)
0.595	0.180	0.855	0.982	2000	15	3864.6
(0.008)	(0.031)	(0.060)	(0.032)			(204.3)

Reported estimates are the mean value across estimates from 1000 simulated datasets. Standard deviations of the estimates are reported in parentheses. T denotes the length of the sample period. The last column reports the mean number of observations (spells) across the 1000 datasets.

spells does indeed produce mean estimates close to the true value of θ_n with relatively small standard errors, particularly when the number of plants is large.

The key to understanding the performance of our procedure is to compare the distribution of *ex post* errors using the extended simulation approach, which was shown in Figure 1, with the distribution of *ex post* errors created using the measure of the expected marginal value of capital in (14) to compute the *ex post* errors in (15).

These two distributions from simulations with 15 periods and 200 plants using the parameter vector θ_n are shown in Figure 3. The distribution based on the incomplete spell correction is a good approximation to the distribution based on using extended simulation to complete all spells. In this way, our approach to dealing with incomplete spells closely resembles the case where all spells are completed through extended simulation.

4.4 Dealing with Unobserved Shocks

Thus far our analysis assumes the profitability shock is observed. Observing the profitability shock is an important assumption for the identification of λ from the pattern of *ex post*

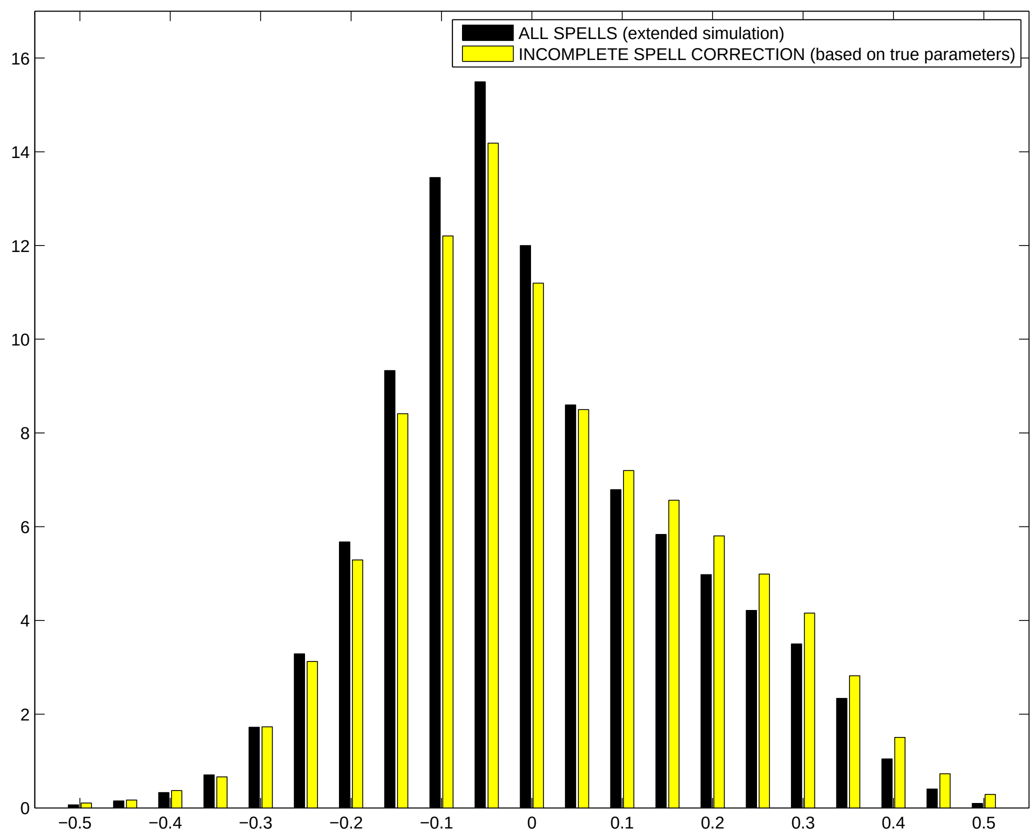

Figure 3: *Ex post* errors from a sample including all spells.

Table 5: Euler-equation Estimates of θ_n Using Correction for Incomplete Spells (using convergence criteria)

α	ν	λ	p_s	plants	T	spells
0.6	0.2	0.80	0.98			
0.593	0.141	0.890	0.943	200	15	386.4
(0.014)	(0.112)	(0.218)	(0.124)			(24.8)
0.597	0.180	0.851	0.9780	2000	15	3864.6
(0.006)	(0.029)	(0.056)	(0.034)			(204.3)

Reported estimates are the mean value across estimates from 1000 simulated datasets. Standard deviations of the estimates are reported in parentheses. T denotes the length of the sample period. The last column reports the mean number of observations (spells) across the 1000 datasets. Convergence was achieved within 10 iterations.

errors. It might appear that even if A_{it} is not directly observed, the analysis of (7) can proceed using observations on average profits as they are proportional to marginal profits, which are needed to compute the error. This does not work, however, because measures of average profits incorporate the disruption effects of λ. Hence, estimating λ from (7) is impossible if A_{it} is not observed. Our method for dealing with unobserved shocks is described below. We then report results from a Monte Carlo exercise before turning to estimation from plant-level data.

4.4.1 Methodology

The estimation procedure has multiple stages to deal with the fact that A_{it} is not observed. Building on prior work by CH, we estimate two structural parameters (α, λ) from the plant's profit function. Here we allow for disruption costs, λ, directly in the estimation. Variable profits at plant i in period t are given by

$$\Pi(A_{it}, K_{it}) = \begin{cases} A_{it}K_{it}^{\alpha} & \text{if } I = 0 \\ \lambda A_{it}K_{it}^{\alpha} & \text{if } I \neq 0. \end{cases} \tag{16}$$

Suppose that $a_{it} = \log(A_{it})$ has the following structure

$$a_{it} = b_t + \epsilon_{it} \tag{17}$$

where b_t is a common shock and ϵ_{it} is a plant-specific shock. Assume $\epsilon_{it} = \mu + \rho_\epsilon \epsilon_{i,t-1} + \eta_{it}$. Taking logs of (16) and quasi-differencing yields

$$\begin{aligned} \eta_{it} &= -\mu + \log\left(\Pi_{it}\right) - \rho_\epsilon \log\left(\Pi_{i,t-1}\right) - \log(\lambda) * (I_{it} \neq 0) \\ &\quad + \rho_\epsilon \log(\lambda) * (I_{i,t-1} \neq 0) - \alpha k_{it} + \rho_\epsilon \alpha k_{i,t-1} - b_t + \rho_\epsilon b_{t-1} \end{aligned} \tag{18}$$

where $(I_{it} \neq 0)$ is an indicator variable equal to 1 if $I_{it} \neq 0$.

The moment condition for estimation is

$$E_{t-1}\left[\eta_{it}\right] = 0. \tag{19}$$

For instruments, we can use any variables determined in period $t-1$ or earlier. We choose $\{\pi_{i,t-1}, k_{it}, k_{i,t-1}\}$ along with a constant, the lagged investment indicator variable, and annual dummy variables.

In addition to obtaining these parameter estimates, we also estimate the processes for the aggregate and idiosyncratic profitability shocks. These parameters are needed to solve the plant-level optimization problem which is used to create marginal values of capital for plant's with incomplete spells.

4.4.2 Simulation-based Results

Results for the two-stage estimation procedure using simulated data are displayed in Table 6. In the first stage, the curvature of the production function, α, the disruption cost, λ, and the serial correlation of the idiosyncratic profitability shock, ρ_ϵ, are estimated using the quasi-differenced profit function in (18). These parameter estimates are close to the true values for the base dataset, which consists of 200 plants and 15 periods. When the number of plants is increased from 200 to 500, the point estimates in the first stage do not change much, but the standard deviation of the estimates across the 1000 samples declines.

Table 6: Euler-equation Estimates of θ_n When Shocks are Unobserved

Profit Function (first stage)			Euler Equation (second stage)				
α	λ	ρ_ϵ	ν	p_s	plants	T	spells
0.6	0.80	0.85	0.2	0.98			
0.592	0.795	0.817	0.167	0.922	200	15	358.1
(0.067)	(0.030)	(0.019)	(0.126)	(0.087)			(23.0)
0.585	0.793	0.820	0.150	0.915	500	15	900.1
(0.049)	(0.022)	(0.012)	(0.112)	(0.080)			(56.7)

Reported estimates are the mean value across estimates from 1000 simulated datasets. Standard deviations of the estimates are reported in parentheses. T denotes the length of the sample period. The last column reports the mean number of observations (spells) across the 1000 datasets.

The first-stage estimates are used as inputs for the second stage estimation of the quadratic adjustment cost, ν, and the selling price of capital, p_s. To incorporate the incomplete spells into the estimation, the first-stage estimates along with initial guesses for $\{\nu, p_s\}$ are used to compute the expected derivative of the value function in (14). *Ex post* errors are then generated for complete and incomplete spells, and the parameters are estimated using the nonlinear two-stage least squares estimator in (11). The final estimates are obtained through an iteration process where the values of $\{\nu, p_s\}$ used to compute the expected derivative of the value function are updated using the Euler-equation estimates and then the estimation is repeated using the new derivative. The process continues until the estimates converge.

As shown in Table 6, the point estimates in the second stage are similar to the previous results in Table 4. As a reminder, the difference between the two tables is that we assume that the profitability shock is not observed for the estimates in Table 6, which then requires the two-stage estimation procedure. In comparing the estimates between Table 4 and Table

6, the estimate for the quadratic adjustment cost, ν, is closer to truth in the latter table while the estimate for the selling price of capital, p_s, is further from truth. Also note that the standard deviations are much larger in the second stage than the first stage. This is due in part to the number of observations used in each stage. The first stage includes observations for each period while the second stage includes observations for each spell. Due to the large number of observations where firms are inactive in respect to investment, the average number of spells (358) is much lower than the number of observations used in the first stage estimation (200 plants * 13 periods = 2600 observations).

5 Estimation

The estimation takes the procedures outlined above, with some modifications, to plant-level manufacturing data. A key issue for the estimation, discussed below, is the considerable noise in measured profits. We discuss our approach to dealing with that issue before presenting parameter estimates.

5.1 Data Description

The Longitudinal Research Database (LRD) plant-level data set is described in some detail in CH. Briefly, we note here that over the period 1972-88, the LRD includes measures of both capital expenditures as well as capital retirements and sales. After 1988, the Annual Survey of Manufactures (ASM), on which the LRD is based, dropped the questions on capital retirements and sales. For our analysis, we require both positive expenditures as well as retirements and sales so our analysis is restricted to the 1972-1988 sample period. Some pertinent aspects of the data for the entire U.S. manufacturing sector are summarized in Table 7. As discussed in CH, these moments vary across sectors although the qualitative features of these moments hold in every sector. Namely, we observe elements of lumpy investment behavior that is asymmetric – a substantial number of plants have zero investment, a substantial number have large, positive investment rates and only a small fraction of plants have negative investment. Moreover, the auto correlation in investment rates is low and close to zero and there is a modest positive correlation between estimated profit shocks and investment rates.

The approach in CH uses these moments in a minimum distance estimation exercise. In doing so, for each vector of structural parameters, the dynamic programming problem was solved through value function iteration, a data set was simulated and moments were calculated. In addition, a fixed discount factor was assumed through the analysis. As found in CH, the parameter estimates supported a mix of nonconvex and convex adjustment costs in order to match these moments. The small standard errors for the moments reported in Table 7 imply that the simulated method of moments (SMM) estimates are tightly estimated. This is one of the strengths of the SMM approach as such moments can be robustly estimated.

Table 7: Moments from the LRD

Variable	LRD
Average Investment Rate	12.2% (0.10)
Inaction Rate: Investment	8.1% (0.08)
Fraction of Observations with Negative Investment	10.4% (0.09)
Spike Rate: Positive Investment	18.6% (0.12)
Spike Rate: Negative Investment	1.8% (0.04)
Serial correlation of Investment Rates	0.058 (0.003)
Correlation of Profit Shocks and Investment	0.143 (0.003)

Standard errors in parentheses.

5.2 Euler-equation Estimation

The Euler-equation estimation approach taken here is computationally much faster as it does not require repeated solution of the dynamic programming problem. There is a considerable increase in the speed of the estimation exercise, though, in contrast to the approach of matching the moments in Table 7, the estimation requires access to the actual data rather than summary moments.

There are, however, a few difficulties posed by estimating the parameters of the model from the LRD. First, there is the relatively short sample length of 19 periods. Second, the analysis assumes A_{it} is observed. Section 4.4 described our procedure for dealing with the latter issue. However, even using the two stage method described in section 4.4 requires some further modifications in applying to the actual data. As discussed in CH and Gilchrist and Himmelberg (1995), profits suffer from considerable measurement error relative to revenues. Part of the challenge is to measure all of the components of costs and also to measure those costs accurately. This is especially a challenge for plants of multi-plant firms as some costs are covered by the parent firm. Accordingly, as described below we modify the method in 4.4 so that we can use the relationship between revenues and the capital stock to estimate the first stage parameters from the profit function.

Assuming a constant elasticity demand function and a Cobb-Douglas production technology implies that revenues are proportional to profits:

$$Rev(A_{it}, K_{it}) \quad = \quad \phi\Pi\left(\bar{A}_{it}, K_{it}\right) \tag{20}$$

where

$$\phi \quad = \quad \frac{1}{1-(1-\beta_K)\left(\frac{\eta-1}{\eta}\right)} \tag{21}$$

β_K is the capital elasticity of the underlying production function and η is the elasticity of demand. Using observations on revenues, it is possible to estimate (α, λ).

Taking logs of revenue and quasi-differencing as above yields a specification using revenues analogous to that above for profits. Even with the estimates of (α, λ) in hand, additional steps are required in the calculation of (7). From (4), average and marginal profits are proportional: $\pi_2(A,K) = \alpha AK^{\alpha-1} = \alpha\frac{AK^\alpha}{K} = \alpha\frac{\pi(A,K)}{K}$. Given $\pi(A,K) = \frac{Rev(A,K)}{\phi}$ we can calculate $\pi(A,K)$ from revenues if we know ϕ.[15] From the profit and revenue functions, we have $\frac{\eta-1}{\eta} = \frac{\alpha}{\beta_K+\alpha(1-\beta_K)}$ so that

[15]Gilchrist and Himmelberg (1995) calibrate the factor of proportionality relating the marginal product of capital to the revenue-capital ratio using the assumption that on average it should be approximately be true that the marginal product of capital is equal to the user cost of capital. We take an alternative approach given that we obtain an estimate of α and can estimate ϕ using α and an estimate of β_K.

$$\phi = \frac{1}{1-(1-\beta_K)\left(\frac{\alpha}{\beta_K+\alpha(1-\beta_K)}\right)}. \tag{22}$$

Hence, combined with the estimate of α we can obtain ϕ using an estimate of β_K. In what follows, we obtain the latter from the cost shares from the NBER productivity database (which is an industry-level dataset that is based on the ASM/LRD data used in this analysis).[16]

These objects are then used in the calculation of (7). In particular, we use the estimates of both λ and α from the revenue function. In addition, the $\pi_2(A_t, K_t)$ and $\lambda\pi_2(A_{t+\tau}, K_{t+\tau})$ terms are obtained from average revenues. There are two steps involved here. First, from (4), average and marginal profits are proportional: $\pi_2(A, K) = \alpha A K^{\alpha-1} = \alpha \frac{AK^\alpha}{K} = \alpha \frac{\pi(A,K)}{K}$. Second, as noted above, we can relate average profits and average revenues. Since the latter suffer less from measurement error, we use them in (7).

Even with this revenue-function approach, we are cognizant that the actual revenue and capital stock data likely are subject to classical measurement error not present in the simulated data. Such classical measurement error implies (see Griliches and Hausman (1986)) that once-lagged instruments are no longer valid in the quasi-differenced profit/revenue function estimation. Following the standard guidance of Griliches and Hausman (1986), we use twice lagged instruments in the estimation of the profit/revenue function.[17]

5.3 Motor Vehicles (371)

Our results for 3-digit SIC sector 371 are reported in Table 8. The results for both complete spells and all spells, thus combining complete and incomplete spells, are presented. In this sector, spells for more than 210 plants are used yielding 3082 spells of which 19 are

[16]For sector 331, the equipment cost share of gross output is materials is 0.03 and for sector 371 the equipment cost share is 0.02.

[17]CH also use twice lagged instruments in their estimation of the quasi-differenced profit/revenue function. We have implemented the method in section 4.4 using twice lagged instruments with the simulated data and obtain qualitatively similar results although less precisely estimated than in Table 6. The twice-lagged instruments perform well with the actual data. The reason that the twice-lagged instruments perform slightly less well with the simulated data as opposed to the actual data is that the simulated data has less cross-sectional variation and less persistence in profits, revenues and capital than the actual data.

incomplete.[18] Since α and λ are not estimated from the Euler equation, their estimates are independent of how we treat incomplete spells. In addition, the table includes the SMM results from CH for comparison.

Table 8: Results for Motor Vehicles (371)

Parameter	Euler-equation estimate (complete spells)	Euler-equation estimate (all spells)	SMM estimate
λ	0.75	0.75	0.68
	(0.383)	(0.383)	(0.040)
α	0.88	0.88	0.78
	(0.063)	(0.063)	(0.038)
ν	0.000	0.000	0.051
	(0.020)	(0.020)	(0.007)
p_s	0.956	0.947	0.81
	(0.061)	(0.062)	(0.022)

Standard errors in parentheses.

The first column of the table shows the results from the Euler-equation estimation. We find evidence of curvature in the profit function, a substantial disruption cost associated with changes in the capital stock, and evidence of irreversibility. The estimated quadratic adjustment cost is quite small. The estimates of ν and p_s do not change with the completion of spells since the inaction rate is relatively low in this sector (about 7 percent of plants in sector 371 have zero investment in any given year). As a result, the share of incomplete spells in the overall number of spells is relatively low. The low inaction rate also contributes to a large standard error on the disruption cost, λ, because this parameter is identified based on differences across active and inactive plants. Since there are relatively few inactive plants, there is only a limited amount of covariation to identify the disruption cost. Regarding the

[18] As described in CH we use a balanced panel of plants but note there are some periods for which there is missing data. Spells that include a period of missing data are omitted. We also note that the parameter δ is set equal to 0.063 and β to 0.95 consistent with CH.

standard errors, the estimated standard errors for ν and p_s have not been adjusted to take into account the uncertainty surrounding the first-stage estimates of λ and α.[19]

The last column shows the results reported in CH using a SMM approach. Qualitatively, the results are similar: low values of ν and non-convexities matter. For the SMM results, there is a larger estimated irreversibility and a greater opportunity cost effect, relative to the Euler-equation estimates. The SMM results have lower estimated standard errors reflecting the very robust estimates of the underlying moments used in SMM. While these alternative estimation methods exploit different variation in the data, the overall findings are qualitatively robust.

5.4 Steel (331)

The results for Steel, 3-digit SIC sector 331, are reported in Table 9.[20] In this sector, spells for more than 140 plants are used with 1842 total spells of which 27 are incomplete. As with the results for sector 371, we find substantial curvature in the profit function, relatively low quadratic adjustment costs, substantial disruption costs, and evidence of irreversibility. Once again, the Euler-equation approach finds similar qualitative patterns as the SMM approach but not as much irreversibility as in the SMM approach. Again, we find a low estimate of ν but this is similar to that found in the SMM approach. We find smaller standard errors with the SMM approach. In this sector, the inaction rate is higher than in sector 371 (about 17 percent of plants have zero investment in any given year in sector 331). As such, we find that incorporating incomplete spells has a more notable impact on the estimate of p_s.

6 Conclusions

This paper has two purposes. The first is to analyze a methodology for using the logic of Euler-equation estimation, as in Hansen and Singleton (1982), to settings in which adjust-

[19]An adjustment could be implemented which would compute second-stage standard errors based on bootstrap procedure. In such a procedure, inputs of λ and α for the second-stage procedure would be drawn from their distribution estimated in the first stage. Such a procedure would be computationally intensive for the estimation that incorporates the correction for incomplete spells as the underlying model would have to be solved for each bootstrap draw and the derivative of the value function computed.

[20]Consistent with CH, δ is set equal to 0.076 and $\beta = 0.95$ for sector 331.

Table 9: Results for Steel (331)

Parameter	Euler-equation estimate (complete spells)	Euler-equation estimate (all spells)	SMM estimate
λ	0.88	0.88	0.70
	(0.308)	(0.308)	(0.034)
α	0.688	0.688	0.66
	(0.036)	(0.036)	(0.027)
ν	0.000	0.000	0.015
	(0.066)	(0.066)	(0.004)
p_s	0.992	0.967	0.946
	(0.083)	(0.083)	(0.005)

Standard errors in parentheses.

ment is infrequent. Our analysis indicates how these procedures can estimate underlying adjustment costs, including those that create the inaction. We have used a simulation environment to identify powerful instruments and to guide us in the analysis of incomplete spells.

The second part of the paper takes this approach to plant-level data for U.S. manufacturers. There we are successful in estimating the parameters of the model. The parameter estimates are qualitatively similar to those reported in CH. One important difference is in the estimates of the irreversibility, which tend to be smaller in the Euler-equation-based approach. Still, both approaches yield evidence supporting substantial non-convexities operating both through irreversibilities as well as the disruption costs to profits during periods of investment. Our main conclusion then is that these alternative approaches applied to the same data yield the same basic insights.

As this research proceeds, we plan to supplement the estimation in two dimensions. First, it is possible to analyze a model in which the non-convex adjustment costs are incurred for investment rates above a critical value, $\bar{\imath}$. Thus far, we have set that value at 0. One way to proceed is to estimate the model for different values of $\bar{\iota}$ and compare the specifications by how well they match the moments.

Second, the empirical analysis has focused on the *ex post* Euler-equation errors. But we have ignored additional information contained in the fact that in some states, the optimizing plant chooses inaction over action, $V^i(A, K) > V^a(A, K)$. There are two ways to use the information contained in this inequality. One is to see how well the estimated model matches the data along this dimension. The second is to formally incorporate these inequalities into the estimation.[21]

Finally, there are numerous other applications of this methodology. One in particular arises in dynamic choice problems with occasionally binding constraints, such as borrowing restrictions. The Euler equation does not hold in periods where the borrowing constraint binds. By the logic of the approach taken in this paper, the parameters of the optimization problem can be estimated by looking over periods in which the constraint does not bind.

References

Aguirregabiria, V. (1997): "Estimation of Dynamic Programming Models with Censored Dependent Variables," *Investigaciones Economicas*, 21(2), 167–208.

Amemiya, T. (1985): *Advanced Econometrics*. Harvard University Press.

Caballero, R. J., and E. M. Engel (1999): "Explaining investment dynamics in US manufacturing: a generalized (S, s) approach," *Econometrica*, 67(4), 783–826.

Cooper, R., and J. Haltiwanger (2006): "On the Nature of the Capital Adjustment Costs," *Review of Economic Studies*, 73, 611–34.

Eberly, J. C. (1994): "Adjustment of Consumers' Durables Stocks : Evidence from Automobile Purchases," *Journal of Political Economy*, 102, 403–436.

Gilchrist, S., and C. Himmelberg (1995): "Evidence on the role of cash flow for Investment," *Journal of Monetary Economics*, 36, 541–72.

[21] This is related to the procedures described in Pakes, Porter, Ho, and Ishii (2006).

Griliches, Z., and J. Hausman (1986): “Errors in variables in panel data,” *Journal of Econometrics*, 31, 93–118.

Hansen, L. P. (1982): “Large Sample Properties of Generelized Method of Moments Estimators,” *Econometrica*, 50, 1029–1054.

Hansen, L. P., and K. J. Singleton (1982): “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models,” *Econometrica*, 50, 1269–1286.

Pakes, A. (1994): “Dynamic Structural Models, Problems and Prospects: Mixed Continuous discrete controls and market interactions,” in *Advances in Econometrics, Sixth World Congress*, ed. by C. Sims, vol. 2, pp. 171–259. Cambridge University Press.

Pakes, A., J. Porter, K. Ho, and J. Ishii (2006): “Moment Inequalities and Their Application,” http://www.economics.harvard.edu/faculty/pakes/papers/Moments_November-06.pdf, Harvard University, mimeo.

Tauchen, G. (1986): “Finite State Markov-Chain Approximation to Univariate and Vector Autoregressions,” *Economics Letters*, 20, 177–81.