

**DYNAMICS OF LABOR DEMAND: EVIDENCE
FROM PLANT-LEVEL OBSERVATIONS
AND AGGREGATE IMPLICATIONS**

**Russell Cooper, John Haltiwanger,
and Jonathan L. Willis**

DECEMBER 2003

RWP 03-12

Research Division
Federal Reserve Bank of Kansas City

Russell Cooper is a professor of economics at the University of Texas at Austin. John Haltiwanger is a professor of economics at the University of Maryland. Jonathan Willis is an economist at the Federal Reserve Bank of Kansas City. The authors are grateful to participants at the 2003 Madrid Conference on Lumpy Investment, Durable Purchases and Technical Change, the 2003 Society of Economic Dynamics conference and the 2003 ITAM-University of Texas Conference. The authors thank the NSF for financial support. The views expressed herein are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System.

Cooper email: cooper@eco.utexas.edu

Haltiwanger email: haltiwan@econ.umd.edu

Willis email: Jonathan.Willis@kc.frb.org

Abstract

This paper studies the dynamics of labor demand at the micro and aggregate level. The correlation of hours and employment growth is negative at the plant level and positive in aggregate time series. Further, hours and employment growth are about equally volatile at the plant level while hours growth is much less volatile than employment growth in the aggregate data. Given these differences, we specify and estimate the parameters of a plant-level dynamic optimization problem using simulated method of moments to match plant-level observations. Our findings indicate that non-convex adjustment costs are critical for explaining plant-level moments on hours and employment. Aggregation generates time-series implications which are broadly consistent with observation. Further, we find that a model with quadratic adjustment costs alone can also broadly match the aggregate facts.

JEL classification: E24, J23, J6

Keywords: Adjustment Costs, Employment, Aggregate Employment

1 Introduction

This paper studies dynamic labor demand at the micro and aggregate level. The work is motivated by observations on the labor input (employees and hours) at both the plant level and in the aggregate. As examined in some detail below, the micro and the aggregate behavior of the labor input is quite different. At the plant-level, there is about equal adjustment in hours per worker and employees and the co-movement of these variables is negative. In the aggregate time series, adjustment is largely in the number of workers rather than hours and hours growth is positively correlated with employment growth.

Understanding plant-level **and** aggregate observations on the labor input requires a model with rich labor adjustment costs. In particular, specifications which ignore labor adjustment costs, such as the standard stochastic growth model, or assume quadratic adjustment costs, as in the linear-quadratic model of Sargent (1978), are unable to match the plant-level observations. The importance of non-convexities in adjustment costs for explaining plant-level observations has been pointed out by other researchers, such as Hamermesh (1989), Caballero and Engel (1993) and Caballero, Engel, and Haltiwanger (1997). Our goal is to go beyond those studies by estimating these adjustment costs with a fully specified structural model at the micro level.¹

Our approach is to specify a dynamic optimization problem at the plant-level. We allow for a fairly rich specification of adjustment costs, including both quadratic and non-convex costs of changing the number of workers. Importantly, the estimated parameters include those which govern the process of the exogenous shocks. Since movements in employment and hours jointly reflect adjustment costs and the shock process, it is critical to distinguish these factors.

We estimate the parameters of the plant-level problem with a simulated method of moments procedure using moments calculated from plant-level data on the number of workers and hours. In particular, we use the standard deviations of employment and hours, the correlation of these two inputs, and the coefficients estimated from a VAR of plant-level employment and hours growth as moments.

For the plant-level data, we find empirical support for adjustment costs which disrupt the production process: the adjustment of the stock of employees entails a reduction in plant-level production. The costs of adjustment are identified from two key features of

¹These other studies do not use plant-level observations along with a fully specified structural model at the micro level to estimate the costs of adjustment.

labor dynamics at the plant-level: the standard deviations of hours and employment growth are about the same and the correlation of hours and employment growth is **negative**. Non-convex adjustment costs are needed to match salient features of plant-level adjustment. Importantly, a model with quadratic adjustment costs is unable to reproduce these moments.

Aggregate moments are quite different from their plant-level counterparts. If we aggregate our plant-level data, we find that hours growth is much less volatile than employment growth and the correlation of hours and employment growth is **positive**. Interestingly, if we aggregate plant-level data produced by a simulation using the estimated adjustment costs that yield a negative micro correlation, we find a positive aggregate correlation. Thus it appears that smoothing over heterogeneous plants is quite important for understanding the difference between plant-level and aggregate observations. While a quadratic adjustment cost model fails to match the plant-level moments, it does better in terms of match of the aggregate data.² An open question, however, is the reliability of using such a quadratic or convex model for policy analysis or for business cycle analysis, as the recent literature suggests that the role of lumpiness at the micro level for aggregate dynamics is especially important in response to large shocks (see Caballero, Engel, and Haltiwanger (1997) and Cooper, Haltiwanger, and Power (1999)).

The paper begins by reviewing some facts about labor adjustment at the plant-level. We then specify and simulate a number of dynamic labor demand models to better understand the mapping from adjustment costs to employment dynamics. Next we present estimates of labor adjustment costs, using the basic facts to guide our empirical exercise. We conclude by looking at aggregate implications.

2 Facts on Labor Adjustment

Our analysis relies on observations of labor input, employees, and average hours at the plant-level. We exploit the Longitudinal Research Database (LRD), a manufacturing plant-level data set for the US. We study a group of plants which are in the sample from 1972 to 1980. Our relatively short time-series sample is due to data limitations. In the LRD, plant-level hours were only collected on a quarterly basis from 1972 to 1980. The sample and data we use are identical to those used for Caballero, Engel, and Haltiwanger

²This ability to match aggregate time series is consistent with findings reported in Sargent (1978).

(1997).³ For these plants the LRD contains information on the number of production workers in each quarter and their total hours worked during the quarter. From this information we can measure hours per worker (average) at the plant for each quarter in our sample.

In what follows, we focus on the growth (log first difference) of hours per worker and employment at the plant level. Focusing on growth rates is natural in this context given the emphasis on adjustment costs. Moreover, the plant-level numbers of workers vary dramatically across plants given the very skewed nature of the size distribution of employment. Our model and analysis have little to say about the size distribution.

We begin with basic statistics on the growth rate of hours and employment. Strikingly, at the plant level the standard deviations of hours growth and employment growth are about the same: 0.180 for hours and 0.189 for employment. Moreover, hours growth and employment growth are inversely correlated at the plant level, with the correlation given by -0.290 . In what follows, these are the key moments of the plant level adjustment dynamics that we seek to match. Note that these statistics are for seasonally-adjusted plant-level variation and they change only slightly if the aggregate quarterly time effect (quarterly mean) is removed from the plant-level data. In particular, after removing a full set of aggregate time effects, the standard deviations of hours and employment are respectively 0.179 and 0.188, and the correlation between hours and employment growth at the micro level is -0.296 . It is striking that the micro moments are driven almost entirely by idiosyncratic effects. We exploit this latter point in our estimation strategy in the analysis that follows.

Beyond the most basic statistics, we also consider a parsimonious characterization of the dynamics of plant-level adjustment based upon a two-variable VAR ordered as hours growth (Δh) and employment growth (Δe). We estimate a VAR using quarterly plant-level data, and to control for both seasonal and aggregate cycle effects, we include a complete set of quarterly time dummies (one for each quarter, year observation). This allows us to estimate the relationships using the cross-sectional variation in the time series of plant-level adjustments as distinct from aggregate time-series variation.

The VAR results are presented in Table 1 below. Here the first (second) row contains the results from the hours (employment) growth regression. The first (second) column contains the coefficient on the lagged value of hours (employment) growth.

³Like Caballero, Engel, and Haltiwanger (1997), we work with a balanced panel and thus do not attempt to jointly study entry and exit along with variations in labor demand. This should be taken into account when considering the aggregate implications of our findings.

Both series exhibit some negative conditional serial correlation. The dependence of hours growth on lagged employment growth is relatively weak, whereas employment growth is positively related to lagged hours growth. Evidently employment growth is led by hours growth.

Table 1: Plant-Level VAR

	Δh_{-1}	Δe_{-1}
Δh	-0.417 (0.002)	0.010 (0.002)
Δe	0.127 (0.002)	-0.207 (0.002)

Note: All variables are growth rates. Standard errors in parentheses.

Putting the basic statistics together yields an interesting pattern at the micro level. Hours growth leads to employment growth, but contemporaneously, hours growth and employment growth are inversely related. In what follows, these patterns enable us to identify and estimate the structure of adjustment costs.⁴

These plant-level moments contrast rather sharply with the standard characterization of hours and employment growth in aggregate data. If we aggregate our plants, the standard deviation of hours growth is 0.011, and the standard deviation of employment growth is 0.02. The correlation of hours and employment growth for the aggregated data is 0.545.⁵

Note that the standard deviations of both aggregate hours growth and aggregate employment growth are considerably lower than the plant-level standard deviations. This reflects the fact that cross-sectional variation in the LRD at the plant level is much larger than the time-series variation. Further, in the time-series hours growth is about one-half as variable as employment growth. Finally, the correlation in the aggregate time series is strongly positive, while it is negative in the cross-section. Thus in the aggregate

⁴We are not the first to detect these basic patterns in the micro data. Caballero, Engel, and Haltiwanger (1997) note the pattern of co-movement between hours and employment and use it to motivate their estimation of adjustment hazards.

⁵Interestingly, when we obtained quarterly hours and employment growth rates for manufacturing production workers from the BLS for the 1972-80 period, we obtained a time-series correlation of 0.54 suggesting that our sample is representative in terms of its aggregate time-series properties.

data, adjustment occurs primarily via the number of workers, and hours growth appears to be complementary to employment growth.

These are interesting facts, particularly the differences between plant-level and time-series observations. Our strategy is to estimate the model at the plant level from the cross-sectional facts. We then see how well the model's predictions, aggregated over plants, can match the time-series evidence.

3 Model

Letting A represent the profitability of a production unit (e.g. a plant), we consider the following dynamic programming problem:

$$V(A, e_{-1}) = \max_{h,e} R(A, e, h) - \omega(e, h) - C(A, e_{-1}, e) + \beta E_{A'|A} V(A', e) \quad (1)$$

for all values of the state vector, (A, e_{-1}) . Here h represents the input of hours by a worker, e_{-1} is the inherited stock of workers before quits occur (at an exogenous rate of q) and e is the stock of current workers. Note the timing assumption of the model: workers hired in a given period become productive immediately. In (1), β is a fixed discount factor.⁶

For our analysis we assume a Cobb-Douglas production function in which the labor input is simply the product eh . Allowing for market power by the plant, we obtain

$$R(A, e, h) = A(eh)^\alpha \quad (2)$$

where the parameter α is determined by the shares of capital and labor in the production function as well as the elasticity of demand.⁷

⁶Since the estimation is off of plant-level rather than time-series variation, assuming a fixed discount factor is reasonable.

⁷The value of α is given by optimization of capital (K) in the fully specified production function, assuming no adjustment costs of investment:

$$\tilde{R}(A, e, h, K) = \left(\tilde{A} (eh)^{\alpha_e} K^{\alpha_K} \right)^{\frac{\eta-1}{\eta}} - rK$$

where α_e and α_K are the respective labor and capital shares, η is the price elasticity of demand, and r is the rental rate on capital. Maximization with respect to capital leads to the reduced form in (2) where

$$\alpha = \frac{\frac{\eta-1}{\eta} \alpha_e}{1 - \frac{\eta-1}{\eta} \alpha_K}.$$

The function $\omega(e, h)$ represents total compensation to workers as a function of the number of workers and their average hours. This compensation function is critical for generating movements in both hours and the number of workers. We adopt a constant elasticity formulation: $\omega(e, h) = e(w_0 + w_1 h^\xi)$.

The cost of adjustment function nests quadratic and non-convex adjustment costs.⁸ In its most general form, $C(A, e_{-1}, e)$ is given by

$$C(A, e_{-1}, e) = F + \frac{\nu}{2} \left(\frac{e - e_{-1}}{e_{-1}} \right)^2 e_{-1} + (1 - \lambda)R(A, e, h) \quad (3)$$

if $e \neq e_{-1}$. Here F represents a fixed cost of hiring or firing workers. In (3), ν parameterizes the level of the quadratic adjustment costs for the hiring and firing of workers. Note that quadratic adjustment costs are based upon *net*, not *gross*, hires. Finally, $(1 - \lambda)$ is the fraction of revenue lost during employment adjustment. This cost represents the disruption of the production and sales processes during periods of labor adjustment.

There are two features of the adjustment cost structure to note. First, the costs are symmetric. Though we can add asymmetric adjustment costs to the model, the moments we have selected to match are symmetric.⁹ Second, the costs apply iff there is a net change in the number of employees.¹⁰

The model allows for two forms of non-convex adjustment costs in which the cost of adjustment is independent of the level of employment change. In one specification, hereafter termed a **fixed cost**, the firm incurs an additive cost of F each period there is a (net) hire or fire. In a second specification, hereafter termed a **disruption cost**, labor adjustment disrupts the production process so that the plant loses a fraction of revenue each time there is a net change in the number of workers.¹¹

There are some interesting differences between these versions of non-convex adjustment costs. While the fixed cost does not interact with the state of profitability, the

As an example, with η set equal to 5, corresponding to a markup of 25%, and assuming constant returns to scale in capital and labor with $\alpha_e = .65$, α is equal to 0.72.

⁸See Hamermesh (1993) for a lengthy discussion of models of labor adjustment.

⁹That is, in principle one could go further to study a VAR which separates job creation and job destruction.

¹⁰Due to data limitations, we are unable to detect gross hires and fires that leave the number of employees constant. Given the assumption of adjustment costs on *net*, not *gross*, changes in employment, the quit rate, plays no role in the firm's optimization problem.

¹¹Cooper and Haltiwanger (1993) discuss these types of adjustment costs. Caballero and Engel (1999) and the recent version of Cooper and Haltiwanger (2000) find support for this type of adjustment cost in their studies of capital adjustment.

disruption cost does: adjustment is more costly during periods of high profitability. In fact, since new workers are employed within the period, the actual adjustment cost in period t depends on the number of workers in that period. Thus, this is different from most typical non-convex adjustment cost specifications. Still, as demonstrated below, this specification does give rise to inactivity in employment adjustment. Further, we can alter the timing of the model by introducing a time-to-hire assumption, so that workers hired in period t become productive in period $t + 1$. In this case, the disruption cost is independent of the flow of new workers in a period.

This adjustment cost function yields the following dynamic optimization problem

$$V(A, e_{-1}) = \max \{V^n(A, e_{-1}), V^a(A, e_{-1})\} \quad (4)$$

for all (A, e_{-1}) , where the superscript n refers to the choice of no adjustment of employment and the superscript a refers to adjustment. These options are defined for all (A, e_{-1}) by

$$V^a(A, e_{-1}) = \max_{h,e} R(A, e, h)\lambda - \omega(e, h) - F - \frac{\nu}{2} \left(\frac{e - e_{-1}}{e_{-1}} \right)^2 e_{-1} + \beta E_{A'|A} V(A', e),$$

$$V^n(A, e_{-1}) = \max_h R(A, e_{-1}, h) - \omega(e_{-1}, h) + \beta E_{A'|A} V(A', e_{-1}).$$

The policy functions indicate the discrete choice of the type of adjustment (a, n) as well as the number of workers and their hours. These policy functions depend on the state vector, (A, e_{-1}) .

For a given parameterization of the basic functions, the optimization problem given in (4) is solved using value function iteration to obtain the policy functions. The state space of employment is discretized into a fine grid with approximately 230 points in the relevant portion of the state space. The idiosyncratic shock is assumed to follow a log-normal process with serial correlation ρ and an i.i.d. innovation with standard deviation σ . Using the procedure outlined by Tauchen (1986), this shock process is then transformed into a discrete state space representation, containing 21 points.

We next turn to a discussion of the properties of the policy functions generated by these adjustment costs. The subsequent section presents the estimation results.

4 The Economics of Adjustment Costs

This section provides some insights into the mapping between the adjustment costs and key aspects of dynamic labor demand. Cooper and Willis (2002) provides a detailed dis-

cussion of the properties of policy functions for the quadratic and fixed adjustment cost specifications. The discussion here goes further by studying the case of disruption costs ($\lambda < 1$) and by looking at particular moments of the data, particularly the variability of hours growth relative to employment growth and the correlation of these two sources of the labor input.

Table 2a-b summarizes the moment implications of three models: quadratic adjustment costs (**Quad**), fixed costs (**Fixed**), and disruption (**Disrupt**). These specifications are defined by the parameter restrictions given in Table 2a. The cost of adjustment parameters are chosen so that the total adjustment costs paid by a plant are roughly equal across adjustment cost models. Throughout we set other parameters of the plant-level dynamic optimization problem at levels which are in the neighborhood of the estimates reported in the next section.¹² For the exercises in this section, we deliberately keep all parameters the same except for the adjustment costs so that we can gain further insight into the impact of the adjustment costs on the covariance structure of hours and employment growth.

A key issue is the co-movement of hours and employees created by these competing models of adjustment costs. As indicated in Table 2b, the quadratic adjustment cost model creates positive co-movement between employment and hours growth. This reflects the response of these variables to a shock to profitability: both hours and employment rise together at the time of the innovation. Though the adjustment dynamics do create some negative co-movement (as the workforce expands, average hours falls), the response to the innovation drives the positive correlation.

In contrast, the fixed cost and disruption models create negative co-movement. Imagine a plant impacted by a series of relatively small profitability shocks. The plant may initially respond by increasing hours of its workers, keeping the number of employees fixed to avoid the fixed cost of adjustment. But if profitability rises enough, the firm will increase the number of workers **and** cut average hours. This can produce the negative co-movement observed in the plant-level data for these adjustment cost specifications.

A second key moment is the relative variability of hours and employment growth. As indicated in Table 2b, the quadratic adjustment cost model produces a low relative variability, indicating that employment growth varies much less than hours growth.¹³

¹²These are given in Table 2a. **These parameters are estimated along with the adjustment costs, as described in the next section.**

¹³Of course, this also reflects the elasticity of wages with respect to hours and the parameterization of the driving process.

This relative variability is far below the estimate of 1.05 from the plant data in the LRD.

In contrast, the fixed cost and disruption models create more variability in employment for the same shock process and the same value of ζ . This increased variability of employment reflects the bunching of job creation stemming from these forms of adjustment costs.

There is also a sharp distinction between the quadratic adjustment cost model and the two non-convex adjustment cost models in terms of the parameters from the simulated plant-level VAR, displayed in Table 2b. In the quadratic adjustment cost model, there is a positive conditional serial correlation for hours growth, whereas the estimate from the LRD data displays negative serial correlation. The quadratic adjustment cost model also displays strong negative dynamics between hours growth and lagged employment growth, where the relationship is basically zero in the data. The other two adjustment cost models are qualitatively similar to the estimates from the LRD. As noted in the Cooper and Haltiwanger (1993) study of capital adjustment, both the fixed and disruption specifications produce bursts of adjustment followed by inactivity, leading to low or negative serial correlation, which contrasts with the positive serial correlation produced in a model of partial adjustment due to quadratic adjustment costs.

Table 2a: Structural Parameters

Model	Structural Parameters						
	ζ	α	ρ	σ	ν	F	λ
Quad	1.02	0.70	0.70	0.20	0.10		
Fixed	1.02	0.70	0.70	0.20		0.005	
Disrupt	1.02	0.70	0.70	0.20			0.996

Table 2b: Simulated Moments

Model	Moments						C/Π
	σ_e/σ_h	$corr(h, e)$	hh	he	eh	ee	
LRD	1.05	-0.30	-0.42	0.01	0.13	-0.21	
Quad	0.35	0.68	0.11	-1.49	0.04	0.26	0.09
Fixed	1.82	-0.11	-0.41	-0.06	0.33	-0.12	0.09
Disrupt	1.10	-0.16	-0.35	-0.15	0.26	-0.04	0.15

Note: e and h are growth rates of employment and hours. σ_i is the standard deviation of the growth rate of i , $i = e, h$. The VAR coefficients are labelled ij , for i & $j \in \{e, h\}$, where i is the dependent variable and j is the lagged regressor.

The final column in Table 2b is a measure of the size of the adjustment costs paid as a percentage of profits (C/Π), measured as revenues minus labor costs. This was calculated from the simulation of a single plant over 400 periods. The same sequence of shocks is used for each calculation so that they are directly comparable. So, for this parameterization, the adjustment costs in the disrupt case are the largest, at .15% of profit.

Further insights into these models of adjustment costs are seen in Figures 1-3. These figures are created using the parameter values indicated in Table 2a and an identical sequence of shocks.

Figure 1 illustrates the response of employment and hours growth to profitability variations in the quadratic adjustment cost case. Note that periods of large realizations of the shock leads to growth in **both** hours and employment. Further, hours growth is considerably more volatile than employment growth in the simulation.

Figure 2 illustrates the fixed cost case. Here, employment growth is frequently zero, with variations in hours used to respond to profitability shocks. There are infrequent bursts of job creation and job destruction in many periods, and in some instances, hours growth is of the opposite sign. These movements underlie the moments reported in Table 2b.

Figure 3 illustrates the disruption cost case. The inaction and negative co-movement from the fixed cost case are present in this specification as well. Clearly though, the variabilities of growth in hours and employment are about the same.

Comparing Figures 2 and 3 is instructive. Recall that the cost of adjustment in the disruption case depends on the state of profitability: all else the same, adjustment costs are larger if A is higher.¹⁴ Consequently, there are states in which there is adjustment in the presence of fixed costs but no adjustment with the disruption cost specification. Instead, the disruption cost case relies more on hours variation.

5 Estimation of Labor Adjustment Costs

Estimation of the parameters of the dynamic optimization problem follows a simulated method-of-moments procedure. We first discuss the procedure and then turn to results.

¹⁴These costs also depend on e_{-1} , which is not shown in the figure.

Figure 1:

Quadratic Adjustment Cost





Figure 2:

Fixed Adjustment Cost







Figure 3:

Disruption Adjustment Cost



5.1 Approach

The estimation entails finding the vector of structural parameters, Θ , which minimizes the weighted distance between moments from the data, M^d , and moments produced from a simulation of the model given a vector of parameters, $M^s(\Theta)$. Thus our estimate of Θ minimizes $\mathcal{L}(\Theta)$ where

$$\mathcal{L}(\Theta) \equiv (M^d - M^s(\Theta))W(M^d - M^s(\Theta))' \quad (5)$$

and W is a weighting matrix.

We solve this minimization problem by simulation since we have no analytic representation of the mapping from Θ to the moments. So, for a given vector Θ , we solve (4) to generate policy functions for employment and hours. Using these policy functions, we create a simulated data set and compute the moments, $M^s(\Theta)$.

For our analysis, the parameter vector may include all forms of adjustment costs: $\Theta = (\zeta, \alpha, \rho, \sigma, \nu, F, \lambda)$. The first parameter represents the responsiveness of compensation to variations in hours. The second one parameterizes the curvature of the profit function. The third and fourth parameters characterize the AR(1) process for the plant-specific profitability shocks.¹⁵ The final three parameters are the adjustment cost parameters. All of our specifications include the first four parameters. We do not estimate all of the adjustment costs simultaneously but rather focus on subsets.

As for the moments, we include the ratio of the standard deviation of employment growth to the standard deviation of hours growth, the correlation of hours and employment growth, and the four VAR coefficients from Table 1. These moments are at the center of the discussion of labor dynamics and, as indicated by the standard errors of the estimated structural parameters, are responsive to variations in Θ .

Importantly, these moments are constructed by removing time-series variation. In this regard, for the ratio of the standard deviation of employment to the standard deviation in hours growth and the correlation, we use plant-level growth rates where a full set of year-quarter time effects have been removed. In addition, the VAR coefficients from Table 1 are based on a VAR with a full set of year-quarter time effects. Thus the estimation relies on variations across plants. One virtue of exploiting the idiosyncratic variation across plants is that general equilibrium effects can be neglected. That is, the first-order effects from general equilibrium considerations presumably result from aggregate shocks that impact wages and interest rates. By exploiting cross-sectional variation,

¹⁵Since we have removed all time-series variation from the moments, the model only includes plant-specific shocks.

we have effectively abstracted from general equilibrium effects.¹⁶ The weighting matrix, W , is based on the variances of the moments estimated from the LRD.¹⁷ Assuming that the covariance between moments is zero, the weighting matrix is constructed as the inverse of a matrix in which the variances of the moments are on the diagonal and all off-diagonal elements are zero.

There are some parameters we do not estimate. We set $\beta = 0.99$, reflecting our use of quarterly observations. Further, the remaining parameters of the compensation function, w_0 and w_1 , were chosen so that steady-state hours were 40 and steady-state employment at each plant was 600.¹⁸

5.2 Results

Our findings are summarized in Tables 3a and 3b. There are four models estimated. The first three are the respective adjustment cost models, estimated separately.¹⁹ These special cases are included to make clear the significant role of fixed and disruption costs at the plant-level relative to the quadratic adjustment cost model. The fourth model allows for the presence of both quadratic and disruption costs. This joint adjustment cost model was the only combination of the three special cases that resulted in a superior fit of the model to the data.

¹⁶While we think it is reasonable to argue that the variation we are exploiting here is driven by the idiosyncratic shock distribution (as opposed to aggregate shocks), an argument can be made that some aspects of the idiosyncratic variation we are exploiting is actually driven by aggregate shocks. That is, in principle, aggregate shocks may impact not only the first moments of the hours and employment growth rate distributions but higher moments as well. Following this logic, the changes in wages and interest rates induced by aggregate shocks may in turn have an impact on higher moments of the plant-level distributions. This issue is one we will address in future research. We are grateful to Victor Aguirregabiria and Simon Gilchrist for discussions on this point.

¹⁷The respective variances of the six LRD moments displayed in Table 2b are (0.02077, 0.0000030, 0.0000030, 0.0000028, 0.0000036, 0.0000034).

¹⁸This is the average number of workers per plant in the balanced panel taken from the LRD.

¹⁹In this table, F is expressed as a fraction of average profits, gross of adjustment costs. Of course, it is treated as a fixed cost in the estimation, as in (4).

Table 3a: Structural Parameter Estimates

Model	Structural Parameters						
	ζ	α	ρ	σ	ν	F	λ
Quad	1.014 (0.001)	0.86 (0.001)	0.996 (0.000)	0.01 (0.005)	0.007 (0.000)		
Fixed	1.034 (0.003)	0.66 (0.01)	0.65 (0.004)	0.33 (0.009)		0.09 (0.002)	
Disrupt	1.009 (0.005)	0.69 (0.22)	0.48 (0.007)	0.47 (0.15)			0.988 (0.005)
Quad & Disrupt	1.008 (0.014)	0.76 (0.07)	0.54 (0.01)	0.48 (0.14)	0.00003 (0.0002)		0.983 (0.019)

Note: Standard errors in parentheses.

Table 3b: Plant-Level Moments

Model	Moments							$\mathcal{L}(\Theta)$
	σ_e/σ_h	$corr(h, e)$	hh	he	eh	ee	C/Π	
LRD	1.05	-0.2960	-0.4173	0.0096	0.1267	-0.2066		
Quad	33.15	0.15	-0.49	-0.002	0.36	0.001	0.000002	144941
Fixed	0.83	-0.14	-0.28	-0.048	0.12	-0.172	0.804	16822
Disrupt	1.11	-0.32	-0.39	-0.039	0.15	-0.242	0.152	1686
Quad & Disrupt	1.14	-0.29	-0.38	-0.003	0.17	-0.244	0.236	1275

Note: e and h are growth rates of employment and hours. σ_i is the standard deviation of the growth rate of i , $i = e, h$. The VAR coefficients are labelled ij , for $i \& j \in \{e, h\}$, where i is the dependent variable and j is the lagged regressor.

Here we see the prominent role of disruption costs in matching observations at the plant-level, as measured by the relatively low value of $\mathcal{L}(\Theta)$ in Table 3b. This specification of adjustment costs alone comes quite close to the observations in many respects. In particular, the disruption case model mimics the negative correlation between employment and hours growth along with the coefficients of the plant-level VAR. It also generates variability in employment growth slightly greater than the variability of hours growth.

The fixed adjustment cost model qualitatively matches many of the moments. In particular, the estimated fixed cost model generates the observed negative correlation between hours and employment growth though it is weaker for the model than in the data. Further, this model also broadly matches the VAR coefficients. A weakness of this estimated model is its failure to generate employment growth that is more variable than hours growth.

The estimated quadratic adjustment cost parameter is very small. This specification roughly matches the conditional serial correlation of hours growth, (hh) , but fails to match the other moments. The quadratic adjustment cost model is unable to generate the negative co-movement of hours and employment growth at the plant level. It also fails to closely capture the relative standard deviations of employment growth and hours growth.²⁰ A partial explanation for the lack of fit of this moment is the low weight placed on the relative standard deviation compared to weights on the other moments, as determined by the inverse of the variances of the respective moments in the LRD. Note also that for this model, the parameter for the serial correlation of the profitability shock is estimated to be very close to one. This is a sharp contrast to the relatively low serial correlations in the other model and relates to the differences between a convex adjustment cost model, where plants are almost always adjusting employment, and a non-convex adjustment cost model, where there are often longer periods of inaction.

The model allowing for both disruption costs and quadratic costs comes closest to the VAR results. It is important to note that the parameter for the quadratic adjustment cost is extremely small, 0.00003. However, in comparing this joint adjustment cost model with the disruption cost special case, it is clear that the small amount of quadratic adjustment cost does improve the fit of the simulated moments to the actual data.

In studying these results, recall that the estimation includes the parameters governing the shocks, the curvature of the profit function and the compensation function as well as the adjustment costs. Generally, we find support for a strictly concave profit function $\alpha < 1$ and relatively little sensitivity of compensation to hours variation, ζ near unity. Except for the quadratic adjustment case, the serial correlation of the plant-level shocks is around 0.50 and the standard deviation is relatively large (particularly in comparison to aggregate time series variation as described later). In the quadratic case, the estimated shock series exhibits a substantial amount of serial correlation and a relatively low standard deviation. Among other things, this exercise indicates that adjustment costs and the driving process must be jointly estimated if the estimation is based on the

²⁰The large value of σ_e/σ_h for the quadratic case reflects the very low estimate of ν .

observed outcomes of employment and hours.

It is again instructive to consider the magnitude of these adjustment costs. In comparison to the existing literature, these adjustment costs are relatively small.²¹ In the best-fit model, the adjustment costs paid in a 400 period simulation are only 0.24% of profits. Despite this, both Tables 2b and 3b indicate that variations in these adjustment costs can produce large variations in the simulated moments.

From a statistical perspective, the models are all rejected since the reported values of $\mathcal{L}(\Theta)$ are quite high. However, in this setting, this reflects the fact that the moments are calculated from a very large panel data set, implying very small standard deviations of the moments (and a very large W). Given how precisely the micro moments are estimated from the actual data, virtually any model would be formally rejected with even very modest deviations of the simulated moments from the actual moments. As we have emphasized above, the fit of the model in the last line of Table 3b is actually quite good in terms of matching the data moments on both a qualitative and quantitative basis.

5.3 Sectoral Results

The above estimates of adjustment costs are based on the assumption that the parameters are common to all plants in our panel. In this subsection, we report some sectoral estimates using plants from the auto and steel sectors. The motivation for undertaking this exercise is to see whether our results are sensitive to disaggregation by industry.

Table 4a-4b summarizes our findings. Here we only estimate an adjustment cost model that includes both the quadratic and disruption cost specifications.²² Qualitatively, the results are quite similar to those reported for all of manufacturing. The estimate of ζ is again near unity implying that compensation is not that sensitive to variations in hours. Further, there is again support for strict concavity of the profit function as well as modest serial correlation of the shocks.

In terms of the adjustment costs, we again see that disruption costs are present along with relatively small quadratic adjustment costs. Interestingly, these disruption costs are much more prominent for autos than steel. The higher disruption estimated for autos is partially driven by the substantially higher standard deviation of employment growth relative to hours in autos relative to steel.

²¹Hall (2002) is an exception and finds relatively small quadratic adjustment costs.

²²As in the overall manufacturing estimates, the model with fixed costs did not perform as well and hence is not shown.

Finally, the sector specific models appear to fit much better. This reflects two factors. First, there is more flexibility here in that the parameters are allowed to be sector specific and thus we can fit the moments better. Second, the sectoral moments are not as precisely estimated from the data as there are obviously fewer plants in a sector relative to all of the plants. This reduction in precision is reflected in the weighting matrix, W , used to compute $\mathcal{L}(\Theta)$. Overall, though, the results at the sectoral level are quite similar to those for the U.S. manufacturing panel.

Table 4a: Structural Parameter Estimates

Model	Structural Parameters					
	ζ	α	ρ	σ	ν	λ
Autos	1.028 (0.06)	0.79 (0.14)	0.59 (0.03)	0.50 (0.06)	0.00002 (0.0001)	0.92 (0.12)
Steel	1.006 (0.04)	0.76 (0.02)	0.61 (0.05)	0.46 (0.54)	0.000006 (0.0003)	0.98 (0.07)

Note: Standard errors in parentheses.

Table 4b: Plant-Level Moments

Sector		Moments						$\mathcal{L}(\Theta)$
		σ_e/σ_h	$corr(h, e)$	hh	he	eh	ee	
Autos:	LRD	1.25	-0.44	-0.42	0.039	0.18	-0.24	14.0
	Est.	1.36	-0.43	-0.38	0.039	0.20	-0.26	
Steel:	LRD	0.95	-0.41	-0.41	0.007	0.16	-0.14	32.0
	Est.	1.16	-0.41	-0.38	-0.033	0.20	-0.18	

Note: e and h are growth rates of employment and hours. σ_i is the standard deviation of the growth rate of i , $i = e, h$. The VAR coefficients are labelled ij , for $i \& j \in \{e, h\}$, where i is the dependent variable and j is the lagged regressor.

6 Aggregate Implications

In this section, we study the aggregate implications of our model. As noted earlier, the aggregate moments are quite different from the plant-level observations. This leads us to inquire whether the model estimated using plant-level observations can, through smoothing by aggregation over heterogeneous plants, produce the moments observed in aggregate data. The first sub-section studies this question and indicates that we can mimic some but not all of the aggregate moments. The second sub-section explores how well we can match the aggregate moments by variations in plant-level adjustment costs.

6.1 Smoothing By Aggregation

The time-series moments obtained by aggregating over the plants are given in Table 5. The stark contrast between time-series and plant-level variation in the LRD is evident from comparing Tables 3 and 5.

In particular, for the time-series moments the ratio of standard deviations, σ_E/σ_H , is much higher than at the plant-level. Further, the correlation of hours and employment growth is positive for the aggregate time series but was negative at the plant level.

The question is whether the models estimated using plant-level data can reproduce these time-series moments. To study this question, we supplement our simulated environment by adding in aggregate shocks at the plant level and then aggregating over the plants using exactly the same procedure as we employed in working with the LRD. Since we have identified the adjustment cost parameters and other relevant micro parameters from the cross-sectional variation, adding aggregate shocks in this manner allows us to evaluate how closely we can match the aggregate moments through the smoothing over heterogeneous plants. A primary item missing is that wages and interest rates do not respond to the aggregate shocks. We comment further on this point below.

To be specific, for each of the adjustment cost models, an aggregate source of variation has been introduced through a common component in A , the profitability shock. The process for the shocks $(\rho_A, \sigma_{\epsilon_A}) = (0.95, 0.05)$ is chosen so that the model reproduces the basic serial correlation and standard deviation from aggregated (log) employment.²³ Note that for all of the adjustment cost specification except for the Quad case, the standard deviation of the aggregate time series shocks is considerably less than the cross-

²³This is aggregate employment growth from the BLS manufacturing employment from 1972-80. We match this series by aggregating over simulated choices at the plant level using the estimated parameters for the Quad & Disrupt case reported in Table 3a.

sectional variation.²⁴ With this additional source of uncertainty, we simulate a panel using the estimated parameters for the plant-level optimization problem. The resulting data was aggregated to create time-series moments using the same methodology as in the aggregation of the LRD panel.

Table 5: Time-Series Moments

	Moments					
Model	σ_E/σ_H	$corr(H, E)$	HH	HE	EH	EE
LRD-agg.	1.81	0.55	0.06	-0.08	0.56	0.40
	Aggregation with Aggregate Shocks					
Quad	13.57	0.67	-0.12	-0.04	-1.28	0.15
Fixed	1.47	0.20	-0.28	-0.08	0.29	0.02
Disrupt	0.47	0.16	-0.29	0.06	0.07	0.01
Quad & Disrupt	0.38	0.16	-0.28	0.22	0.03	0.04

Note: E and H are growth rates of aggregate employment and aggregate hours. σ_I is the standard deviation of the growth rate of I , $I = E, H$. The VAR coefficients are labelled IJ , for $I \& J \in \{E, H\}$, where I is the dependent variable and J is the lagged regressor.

The results are summarized in Table 5. By introducing aggregate shocks and then aggregating, a different pattern of variation in hours and employment growth emerges relative to the model without aggregate shocks.

The most dramatic changes occur in the correlation between hours and employment and in the VAR coefficients, especially for the fixed and disruption adjustment cost cases. In particular, for the fixed and disruption models, hours and employment growth are positively correlated, as is the case for the aggregate data. Recall that these correlations were negative for the models and data at the plant-level. Thus, the presence of the aggregate shocks combined with the aggregation across heterogeneous units switches the sign of the correlation in both the data and the models. The quadratic model displayed a positive correlation at the plant level, and this correlation is even higher for the aggregated data.

²⁴There is an inherent apples and oranges interpretation problem in the the degree of smoothing by aggregation over these different models given the differences in the estimated idiosyncratic shock processes across these models. Our objective here is to take these estimated models at face value and to consider the smoothing by aggregation in the context of the estimated distribution of idiosyncratic shocks.

In the aggregated data, employment growth is considerably more variable than hours growth. Comparing Table 3b and 5, the relative standard deviation moment is lower in the aggregate time series relative to the plant-level data for all models except the fixed cost specification.

As for the VAR, these coefficients move in the direction of those found in the actual aggregate data. The ee coefficient was negative in the plant-level data and for the fixed and disrupt cases. The time-series counterpart, the EE coefficient, is positive in the data and also much larger for the models than in the plant-level data. In terms of accounting for the serial correlation in aggregate employment growth, the quadratic model comes closest. The positive interaction between current employment growth and lagged hours, the EH coefficient, is reproduced by the fixed case and to some degree by the disruption case as well. The quadratic model yields a large negative interaction between current employment growth and lagged hours which is at odds with the actual data.

Aggregation of the fixed-cost model yields results that mimic the change in the moments created by aggregating the plant-level data. Interestingly, the disruption cost model which fits the micro data the best does not match the aggregate data as well, especially in terms of the ratio of standard deviations of employment and hours growth. Instead, the fixed-cost model does better in matching the aggregate moments. It is also interesting that the quadratic model, which did so poorly at the micro level, does a much better job of matching many (although not all) of the aggregate moments.

Smoothing by aggregation is clearly a factor that can help reconcile differences in micro and macro moments. However, the micro models that best fit the micro moments do not fully capture the aggregate movements. In the next section, we explore a calibration sensitivity analysis to shed further light on these issues.

6.2 Are there adjustment cost parameters that match the aggregate time-series moments?

There may be many forces to explain why the aggregated versions of the best fitting micro models miss some important aspects of the aggregate moments. Such possibilities include time variation in factor prices induced by general equilibrium effects and/or a richer characterization of the frictions of employment adjustment that induce interactions between plant-level and aggregate dynamics (e.g., search frictions that depend upon aggregate fluctuations). Exploring these forces in a complete fashion is beyond the scope of this paper although we discuss some ideas for future work along these lines in

the conclusion.

In this section, we conduct a calibration exercise that searches for parameterizations of the adjustment costs that yield aggregate moments that match the actual moments. Our motivation for this calibration exercise is as follows. A common finding in the literature is that a quadratic adjustment costs specification fits the aggregate data reasonably well. We want to know if this finding holds in our setting – i.e., is there a specification of the quadratic cost model in this context that captures the aggregate moments reasonably well? More generally, we want to know whether there are parameterizations involving either convex or non-convex adjustment costs (or both) that can match the quantitative actual aggregate moments. Put differently, we want to know how different the adjustment cost parameters need to be from those that best fit the plant-level moments to match the aggregate moments.

As a starting point and to provide some discipline on this exercise, we start with the best fitting micro model. As indicated in Table 3a, this involves a quadratic component and a disruption component to adjustment costs. For this calibration exercise we consider variations in these two parameters alone, keeping all other parameters (e.g., shock processes, elasticities, etc.) the same as those for the best fitting micro model. We also specify the same aggregate shock process as in the previous sub-section.

We search across the adjustment cost parameters (ν, λ) for the values that yield the best matches to the aggregate moments. In doing so, we use an identity matrix for W .

Before proceeding, we emphasize that we regard this exercise as an exploratory sensitivity analysis rather than an alternative estimation exercise. Estimating the structural adjustment cost parameters using plant-level moments is preferable on many grounds. For one, these moments are much more precisely estimated. For another, the plant-level moments we use for estimation abstracted from aggregate variation and as such we could avoid the complexities of the forces alluded to above that influence the aggregate moments (these include smoothing by aggregation, time series variation in factor prices, frictions that are a function of aggregate fluctuations, etc.). Still, in spite of these limitations, we regard this calibration exercise as interesting because it potentially sheds light on the relationship between the results in this paper and the literature that estimates or calibrates adjustment costs using aggregate data only. Moreover, one might think of the parameterization of adjustment costs for a macro model somewhat differently than for a micro model – that is, adjustment costs in a macro model might be specified to encompass forces that generate smoothing in the aggregate.

Our findings are summarized in Table 6 where the rows refer to alternative adjust-

ment cost parameterizations. The specification which fits the aggregate moments best is ($\nu = 0.004, \lambda = 0.998$). Relative to the best fitting model for the plant-level data, this specification entails larger quadratic adjustment costs (by a factor of 100) and smaller disruption costs. Still, these adjustment costs are relatively small. Further, the model is still not capable of capturing the full dynamics of hours and employment, particularly the positive EE coefficient.

Interestingly, the model with quadratic adjustment costs alone also is capable of matching both the relatively volatility and the correlation between hours and employment growth in the aggregate data. This specification too fails to capture the dynamics between these variables. It is worth noting that the magnitude of ν in the model with quadratic adjustment costs alone is still very small relative to the magnitude of quadratic adjustment costs that have been used in the empirical literature to capture aggregate fluctuations. Part of the reason for this is likely that we focus on a substantial number of aggregate moments while the literature focuses mostly on the dynamics of employment itself (so the EE coefficient gets much more weight in the existing literature).

Table 6: Robustness

Model	Parameters		Moments					
	ν	λ	σ_E/σ_H	$corr(H, E)$	HH	HE	EH	EE
LRD-agg.			1.81	0.55	0.06	-0.08	0.56	0.40
Quad	0.005		1.82	0.56	-0.28	-0.14	0.15	-0.01
Disrupt		0.993	1.82	0.12	-0.41	-0.08	0.25	-0.15
Quad & Disrupt	0.004	0.998	1.80	0.53	-0.25	-0.16	0.29	-0.07

Note: E and H are growth rates of aggregate employment and aggregate hours. σ_I is the standard deviation of the growth rate of I , $I = E, H$. The VAR coefficients are labelled IJ , for $I \& J \in \{E, H\}$, where I is the dependent variable and J is the lagged regressor.

7 Conclusions

This paper provides evidence on the nature and significance of labor adjustment costs. Evidently, the principle cost of labor adjustment is the disruption of the production process. These costs are needed to match plant-level observations on the covariance structure of hours and employment growth. For plant-level data, quadratic adjustment costs do not appear to be significant. Put differently, a model with only quadratic

adjustment costs yields micro predictions that are very much at odds with the micro data. A model with fixed adjustment costs does substantially better than a quadratic model but does not match the plant-level data as well as a model which combines disruption costs with (a little bit of) quadratic adjustment costs. The finding that non-convex rather than quadratic adjustment costs are needed to match plant-level observations is consistent with the evidence presented by Hamermesh (1989) and the summary of plant-level behavior in Caballero, Engel, and Haltiwanger (1997).

We also explore aggregate time series evidence with particular emphasis on the fact that hours and employment growth exhibit positive correlation in the time series but not in the cross-section. Further, hours and employment growth are about equally volatile in the cross section while much more of the labor adjustment takes the form of variations in employment growth in the aggregate time series. We indicate that the aggregation across plants with non-convex adjustment costs and idiosyncratic shocks can match some of these aggregate time series patterns. It is striking that a non-convex model can yield such different micro and macro behavior and that the potential key in accounting for such different behavior is smoothing by aggregation.

Still, the structural model estimated at the plant-level misses many of the aggregate time series features. We explored how variations in the adjustment cost parameters might better enable us to match the aggregate time series moments. Though we were able to get closer to the aggregate moments, it does seem that no combination of the adjustment costs parameters is capable of reproducing the aggregate moments. This exercise was largely intended to produce intuition for future directions of research. One important avenue is to introduce additional smoothing through relative price movements. A related avenue is to consider a richer characterization of adjustment costs where the frictions of employment adjustment between firms (e.g., search frictions) vary with aggregate conditions. At this point in the development of the literature on adjustment costs, no one has considered the fully specified model that considers the role of all of these factors in aggregation and in turn the differences in the dynamics at the micro and the macro level. There is a good reason for this because having a rich model of convex and non-convex adjustment costs, idiosyncratic shocks, aggregate shocks, endogenously determined factor prices and micro frictions that vary with aggregate fluctuations in a fully specified dynamic general equilibrium model is a major undertaking.

Of course, one view is that at least some of these forces can be abstracted from in considering macroeconomic fluctuations. We are skeptical of this view for a number of reasons. First, as we have shown key micro and the macro moments on the

covariance structure of hours and employment movements differ dramatically. Second, understanding the evolution of the economy in response to policy or other shocks, it may be important to understand how the policy impacts the cross sectional distribution of activity. In short, it is presumably important to have a structural model of adjustment that is consistent with both micro and macro dynamics. In this paper (and in the related recent literature), we have made more progress on the former than the latter. Clearly, making progress on both micro and macro dynamics simultaneously is an important step for future research.

References

- CABALLERO, R., AND E. ENGEL (1993): "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics*, 108, 313–58.
- (1999): "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S,s) Approach," *Econometrica*, 67, 783–826.
- CABALLERO, R., E. ENGEL, AND J. HALTIWANGER (1997): "Aggregate Employment Dynamics: Building From Microeconomic Evidence," *American Economic Review*, 87, 115–37.
- COOPER, R., AND J. HALTIWANGER (1993): "On the Aggregate Implications of Machine Replacement: Theory and Evidence," *American Economic Review*, 83, 360–82.
- (2000): "On the Nature of the Capital Adjustment Process," NBER Working Paper #7925.
- COOPER, R., J. HALTIWANGER, AND L. POWER (1999): "Machine Replacement and the Business Cycle: Lumps and Bumps," *American Economic Review*, 89, 921–946.
- COOPER, R., AND J. WILLIS (2002): "The Cost of Labor Adjustment: Inferences from the Gap," Federal Reserve Bank of Kansas City, Working Paper #02-11.
- HALL, R. (2002): "Industry Dynamics with Adjustment Costs," NBER Working Paper #8849.
- HAMERMESH, D. (1989): "Labor Demand and the Structure of Adjustment Costs," *American Economic Review*, 79, 674–89.
- (1993): *Labor Demand*. Princeton University Press.
- SARGENT, T. (1978): "Estimation of Dynamic Labor Demand Schedules under Rational Expectations," *Journal of Political Economy*, 86(6), 1009–1044.
- TAUCHEN, G. (1986): "Finite State Markov-Chain Approximation to Univariate and Vector Autoregressions," *Economics Letters*, 20, 177–81.