

THE COST OF LABOR ADJUSTMENT: INFERENCES FROM THE GAP

Russell Cooper and Jonathan L. Willis

DECEMBER 2002; LAST REVISED JULY 2004

RWP 02-11

Research Division
Federal Reserve Bank of Kansas City

Russell Cooper is a professor of economics at the University of Texas and, when this paper was written, was a visiting scholar at the Federal Reserve Bank of Minneapolis. Jonathan L. Willis is an economist at the Federal Reserve Bank of Kansas City. The authors would like to thank the NSF for financial support. Cooper thanks the Federal Reserve Bank of Minneapolis for its hospitality during revisions of this manuscript. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City, Federal Reserve Bank of Minneapolis, or the Federal Reserve System.

Cooper email: cooper@eco.utexas.edu
Willis email: jonathan.willis@frb.kc.org

Abstract

This paper studies labor adjustment costs. Our focus is on inferring the structure of adjustment costs at the micro-level from aggregate observations of employment and hours growth. We specify a dynamic optimization problem at the plant level, allowing for both convex and non-convex adjustment costs. We estimate the parameters of the adjustment process using an indirect inference procedure in which simulated moments are matched with data moments. For this study we use estimates of reduced-form adjustment functions obtained by the "gap methodology" reported in Caballero and Engel (1993) as data moments. Contrary to evidence at the micro level in support of non-convex adjustment costs, our findings indicate that piecewise quadratic adjustment costs are sufficient to match these aggregate moments.

JEL classification: E24, J23, J6

Keywords: Aggregate Employment, Employment, Adjustment Costs

1 Introduction

This paper studies labor adjustment costs. Our focus is on inferring the structure of adjustment costs at the micro-level from aggregate observations of employment and hours growth. In doing so, we consider models with both quadratic and non-convex adjustment costs. We find that a model with piece-wise quadratic adjustment costs fits the facts best.

The analysis of labor adjustment in Hamermesh (1989) and Caballero, Engel, and Haltiwanger (1997) presents a serious challenge to the quadratic adjustment cost model and supplies evidence in favor of non-convex adjustment costs at the plant level. Caballero and Engel (1993), hereafter CE, and Caballero, Engel, and Haltiwanger (1997), hereafter CEH, go further and argue that aggregated employment growth is a nonlinear function of employment gaps, defined as the difference between current and target employment. This conclusion is viewed as: (i) evidence against the quadratic adjustment cost model and (ii) suggestive that plant-level non-convexities have aggregate implications.

Cooper and Willis (2003) challenges this conclusion by analyzing the implications of a model of dynamic labor demand with quadratic adjustment costs. While this model is unable to match the rich patterns of labor adjustment at the plant level, Cooper and Willis (2003) find that it can produce a nonlinear relationship between aggregate employment growth and employment gaps using the methodologies of CE and CEH. Thus the findings of CE and CEH do not necessarily reflect aggregate effects of plant-level non-convexities.

But, Cooper and Willis (2003) do not uncover the structural model that would generate the CE findings. That is, their analysis does not answer the following question: what plant-level labor adjustment costs underlie the results reported in CE? Do the results in CE reflect the aggregate implications of non-convex adjustment costs at the plant level?

Answering these questions is the goal of this paper. By doing so, we are taking some leading “facts” about labor dynamics, neatly summarized by CE, and indirectly inferring the underlying structure of adjustment costs. Thus our results are informative about the nature of labor adjustment costs as well as the inferences one can draw from the CE gap methodology.

In this analysis, we specify a dynamic optimization problem at the plant level allowing quadratic, two types of non-convex (fixed and disruption), and piece-wise linear costs

of adjusting labor. We use the reduced-form estimates of CE to estimate the structural parameters of adjustment costs. **Our main finding is that a piecewise quadratic adjustment cost function fits the data best.**¹ Thus the aggregate evidence used in CE is consistent with the implications of a model with piece-wise quadratic adjustment costs without non-convexities.²

2 A Dynamic Optimization Framework

Our approach begins with the specification of a dynamic optimization problem at the plant level. We introduce a variety of adjustment costs into our model.³ The optimal decision rules are characterized and used to create a simulated data set which provides a basis for estimation of structural parameters. As the estimation uses the gap approach of CE and CEH, we display properties of these alternative models from that perspective.

Following the notation and presentation in CEH, the gap between the desired employment and the actual employment (in logs), at the start of period t for plant i is defined in the following equation.⁴

$$z_{it} \equiv \tilde{e}_{it}^* - \tilde{e}_{it-1} \quad (1)$$

Here \tilde{e}_{it}^* is the target level of employment given the realization of all period t random variables and \tilde{e}_{it-1} is the (log) level of employment at the start of period t prior to any period t adjustments. Thus, z_{it} is a gap between the target and actual employment levels.⁵

A key is measuring the target level of employment, \tilde{e}_{it}^* . For the analysis of the models in this section, we define that target as the level of employment that would be chosen if all adjustment costs were removed for a single period. Following CE we term this the **frictionless target**. We discuss how this is constructed for the different specifications of adjustment costs.

¹By piecewise quadratic we mean a quadratic form for adjustment costs in which the scalar parameter depends on the sign of employment growth.

²The complementary analysis of plant-level data is reported in Cooper, Haltiwanger, and Willis (2003).

³Hamermesh (1993) provides a detailed discussion of various models of labor adjustment.

⁴Whenever we refer to these gap measures and relate them to employment, all variables are in logs. To distinguish logs from levels, we denote the log of employment by \tilde{e}_{it} .

⁵This measure of the gap does not correspond to that in Section II of CE. We are careful in the estimation phase to convert to the CE definitions but find the specification in (1) easier for the analysis. See Cooper and Willis (2003) for further discussion of these measures.

Given a measure of the employment target, the change in employment, $\Delta\tilde{e}_{it}$, for a plant is expressed as a function of the employment gap, $H(z_{it})$, where $H(\cdot)$ is the hazard or adjustment rate function.⁶ In the following sections, we illustrate the mapping from alternative specifications of adjustment costs to $H(\cdot)$.

2.1 Basic Optimization Problem

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

$$V(A, e_{-1}) = \max_{e, h} R(A, e, h) - \omega(e, h) - C(e, e_{-1}) + \beta E_{A'|A} V(A', e). \quad (2)$$

Here h represents the input of hours per worker, e_{-1} is the inherited stock of workers and e is the stock of current workers. Note the timing assumption of the model: workers hired in a given period become productive immediately.

Let $e = \phi_e(A, e_{-1})$ and $h = \phi_h(A, e_{-1})$ be the policy functions. Since we observe variations in employment and hours, these policy functions, which depend on the underlying parameters of the functions in (2), provide the link between the model and the data.

As the interesting aspect of this problem is the cost of labor adjustment, $C(e, e_{-1})$, we simplify the analysis and assume there are no adjustment costs for capital.⁷ As the plant optimally chooses over the rental of capital, we specify a reduced expression of current revenues (net of capital costs) given by $R(A, e, h)$. For the case of a Cobb-Douglas production function in which the labor input is simply the product eh , the revenue function is

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

where the parameter α is determined by the shares of capital and labor in the production function as well as the elasticity of demand.⁸

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 key for

⁶As in CEH, at this level of generality, we can not distinguish between a partial adjustment structure, in which $H(\cdot)$ measures the rate of adjustment, and a non-convex environment, in which $H(\cdot)$ measures the probability of full adjustment.

⁷CEH do not include capital in their formulation either though it could be implicit in their target and hence in their gap measure. See Shapiro (1986) for a study which includes both capital and labor adjustment.

⁸See the discussion in CE and Appendix A of Cooper and Willis (2001).

generating movements in both hours and the number of workers. If compensation is proportional to hours, employers will simply modify hours in reaction to shocks and avoid costly adjustment in the number of workers. Thus the elasticity of the compensation function with respect to hours interacts with the costs of adjusting workers to produce the joint response of hours worked and employment to profitability shocks.

Our specification of the compensation function is adopted from CE. We assume

$$\omega(e, h) = e (w_0 + w_1 h^\zeta). \quad (3)$$

This compensation function features a constant marginal wage elasticity of $\zeta - 1$.

We study the policy functions and hazard functions for alternative specifications of adjustment costs. We choose parameters for the adjustment costs which are reasonable in light of other studies or evidence. As discussed in section 3.1, we calibrate the other parameters of the dynamic programming problem in creating these simulations.

2.2 Quadratic Adjustment Costs

A quadratic specification is the traditional model of adjustment costs. With any convex adjustment cost structure, adjustment will be partial, reflecting the increased cost associated with rapid adjustment. The quadratic specification shares this property and is more tractable, as seen in Sargent (1978).

We specify quadratic adjustment costs as

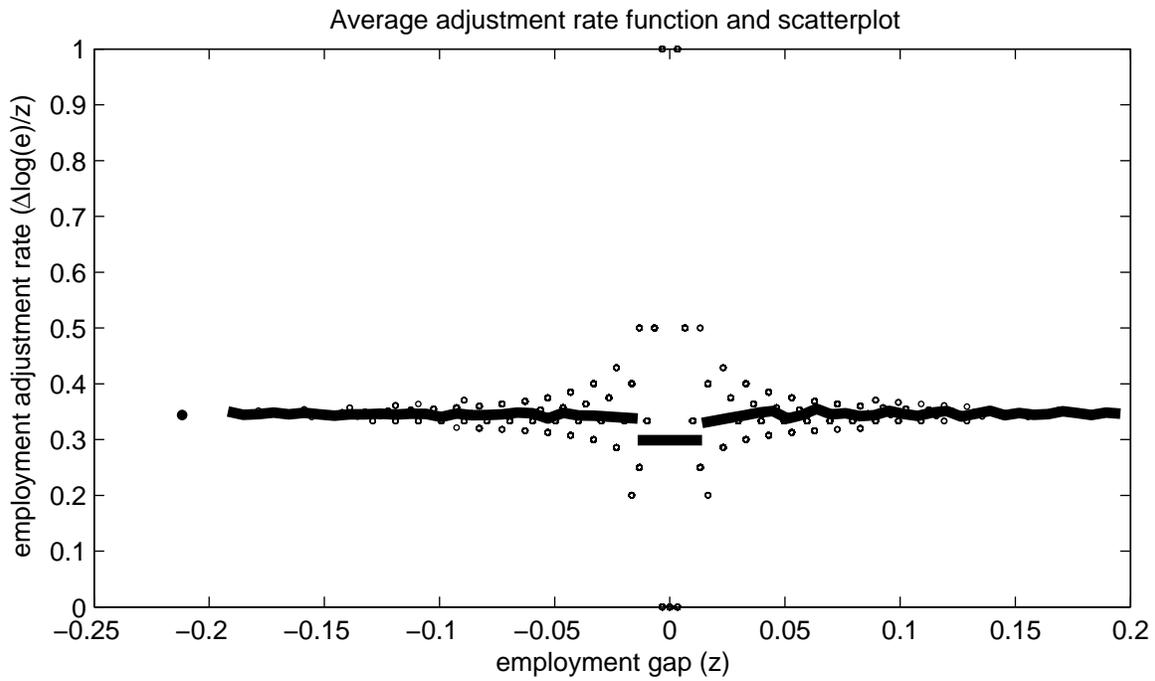
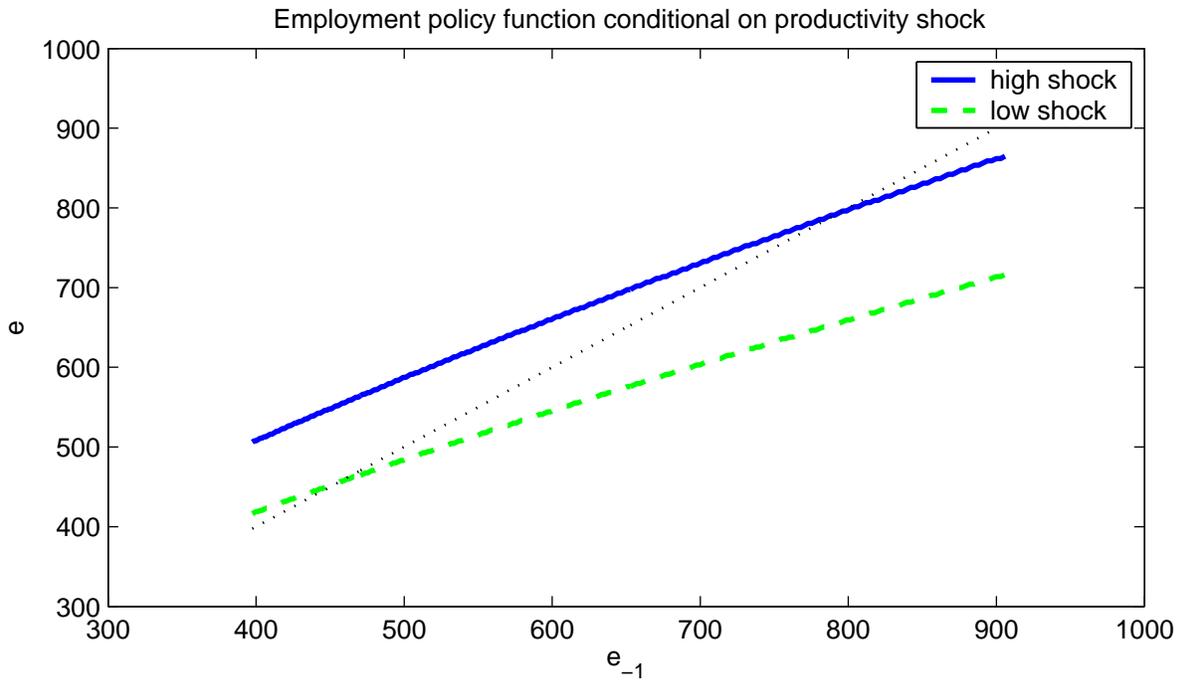
$$C(e, e_{-1}) = \begin{cases} \frac{\nu^+}{2} \left(\frac{e - e_{-1}}{e_{-1}} \right)^2 e_{-1} & \text{if } e > e_{-1} \\ \frac{\nu^-}{2} \left(\frac{e - e_{-1}}{e_{-1}} \right)^2 e_{-1} & \text{if } e \leq e_{-1}. \end{cases} \quad (4)$$

The costs of adjustment are assumed to be a quadratic function of the difference between the stock of workers in the current period (e) and those from the previous period. The plant pays the adjustment cost for net, *not* gross, hires. That is, if the plant hires enough workers to offset those who quit, then there are no adjustment costs. Here we have allowed the quadratic specification to be asymmetric.

Figure 1 shows the policy functions and hazard function for a parameterization in which $\nu^+ = \nu^- = 2$.⁹ There are two policy functions shown in the top panel: one for

⁹This parameterization is purely for illustration though, as our estimates indicate, a value of $\nu = 2$ is not unreasonable. All other parameters were set at values described below in Section 3.1 with the addition of $\rho = 0.95$, $\sigma = 0.022$, and $\zeta = 1.9$.

Figure 1: $\phi_e(A, e_{-1})$ and $H(z)$ with Quadratic Adjustment Costs



a high profitability state and one for a low profitability state. The policy functions are upward sloping in lagged employment and $\phi_e(A, e_{-1})$ is increasing in A .

For this model, the frictionless target, denoted $e^*(A)$, is the solution to the optimization problem when $\nu = 0$ for a single period. Equivalently, $e^*(A)$ solves $e = \phi_e(A, e)$ and thus corresponds to the employment level where the employment policy function, given A , crosses the 45° line.

The bottom panel of Figure 1 characterizes the adjustment rate function. The adjustment rate is essentially flat for all values of the employment gap, except in the neighborhood of zero.¹⁰ In this region, the discrete nature of the solution method limits the employment choices of plants when they are very close to their target level.¹¹

2.3 Fixed Adjustment Costs

Alternatively, there may exist non-convex costs of adjusting the work force in the form of a fixed adjustment cost. In terms of training new workers, it is certainly reasonable that this process might entail increasing returns to scale: the resources devoted to a training class may be largely independent of the size of that class. Further, the costs of recruiting and interviewing may logically have a fixed cost component as well. As discussed in Hamermesh and Pfann (1996) and understood by most economists, the annual recruiting of new assistant professors entails a sizeable fixed cost stemming from advertising, endless department meetings, expenses of attending the convention, etc.

Empirically, CEH stress upward sloping hazards and (see their Figure 1b-1d) a mode of zero adjustment as well as one of full adjustment. These types of observations could reflect non-convexities in the adjustment process of the form introduced here. Thus, we add a fixed cost F to capture these effects. To match the observation of zero changes in the number of workers at the plant level, we assume that this fixed cost is borne for net adjustments in the work force.

Thus we consider a dynamic programming problem in which

$$V(A, e_{-1}) = \max[V^a(A, e_{-1}), V^n(A, e_{-1})] \quad (5)$$

¹⁰The adjustment rate function, illustrated by the solid line in the bottom panel, is constructed by dividing the employment gap distribution into small intervals and then connecting average values from each interval. The exception is that in the neighborhood of an employment gap of zero, the flat adjustment function line represents the average adjustment rates for observations where the gap is between -0.014 and 0.014.

¹¹As a robustness check of the estimation results, we double the number of points in the relevant region of state space to ensure that the discrete nature of the solution method is not biasing the results.

where $V^a(A, e_{-1})$ represents the value of adjusting employment and $V^n(A, e_{-1})$ represents the value of not adjusting employment. These are given by

$$V^a(A, e_{-1}) = \max_{h,e} R(A, e, h) - \omega(e, h) - C(e, e_{-1}) + \beta E_{A'|A} V(A', e) \quad (6)$$

$$V^n(A, e_{-1}) = \max_h R(A, e_{-1}, h) - \omega(e_{-1}, h) + \beta E_{A'|A} V(A', e_{-1}). \quad (7)$$

Note that non-adjustment implies that the plant hires enough workers to offset quits. For this case, the cost of adjustment function is

$$C(e, e_{-1}) = \begin{cases} F^+ & \text{if } e > e_{-1} \\ F^- & \text{if } e < e_{-1} \\ 0 & \text{if } e = e_{-1}. \end{cases} \quad (8)$$

Figure 2 shows the policy and hazard functions for this specification where the fixed cost is set at 1% of average profits.¹² The policy functions indicate that there are two regions: action and inaction. The inaction region is an interval of lagged employment in which $\phi_e(A, e_{-1}) = e_{-1}$. In this region, the adjustment of the labor input to variations in the state vector arise through $\phi_h(A, e_{-1})$. Hours will decrease with lagged employment and increase with the profitability shock.

As indicated in the figure, if there is adjustment, the level of employment chosen is independent of lagged employment. In fact, for this model, when adjustment occurs, the frictionless target is chosen: $e = \phi_e(A, e_{-1}) = e^*(A)$.¹³ The frictionless level of employment lies in the interval of inaction.

Note that there is a type of overshooting in this model. At the level of e_{-1} in which the plant is indifferent between increasing employment and non-adjustment, the level of employment chosen, if the plant adjusts, exceeds e_{-1} . This makes intuitive sense: if the plant is going to bear the fixed adjustment cost, the magnitude of adjustment must be large enough to create a profit gain sufficient to offset the cost.

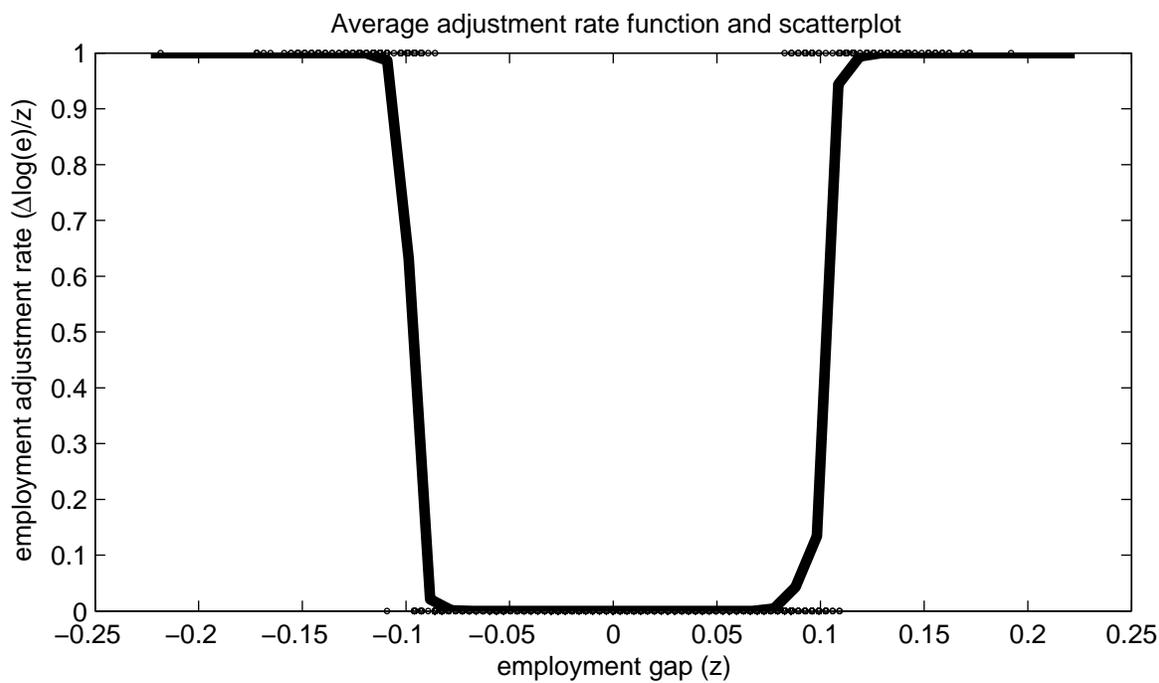
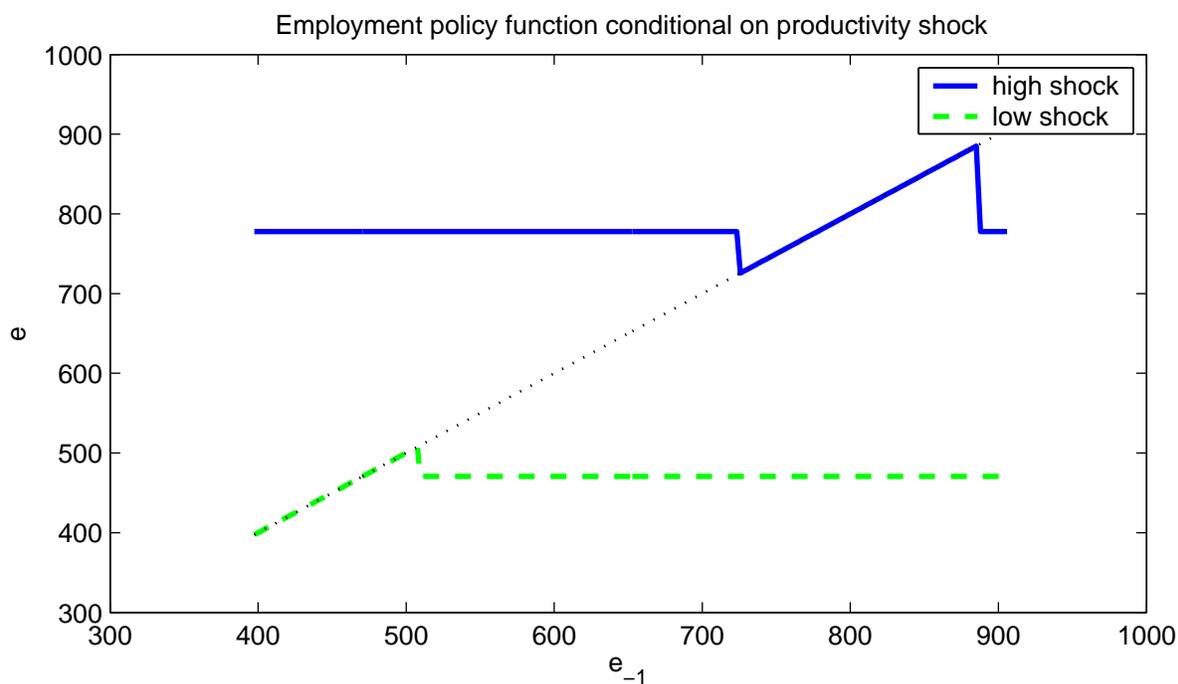
The bottom panel of Figure 2 shows the relationship between the adjustment rate and the gap. Here we see there is substantial inaction (approximately 93%) for $z \in [-0.08, 0.08]$. The adjustment hazard is essentially 0 or 100%.¹⁴ In contrast to the

¹²Throughout the analysis, the fixed costs are set in proportion to average profits. This is without loss of generality but gives these costs economic context and allows us to assess their magnitude.

¹³For this model, the frictionless target is the solution to the optimization problem underlying $V^a(A, e_{-1})$.

¹⁴There are some values of z where the adjustment rate is between 0 and 100%. This reflects the fact that z is not quite a sufficient statistic for the plant's policy function.

Figure 2: $\phi_e(A, e_{-1})$ and $H(z)$ with Fixed Adjustment Costs



adjustment hazards reported in Figures 1b-1d of CEH, there is no partial adjustment for this case.

For the fixed adjustment cost model, there is a more natural definition of the gap than the one specified above. To see this, suppose $F^+ = F^- = F$ and assume that F is an i.i.d. random variable with a cdf given by $G(F)$. As in Rust (1987), the probability of adjustment is given by

$$G(V^a(A, e_{-1}) - V^n(A, e_{-1}))$$

so that adjustment occurs iff the gains to adjustment exceed the realized cost F .¹⁵ So $(V^a(A, e_{-1}) - V^n(A, e_{-1}))$ is the natural measure of the gap for this problem.

Of course, this gap depends on the state vector (A, e_{-1}) . It is natural to ask how well the actual gap using the frictionless target relates to this difference in values. For the model specified so that the fixed costs are 1% of average profits, the correlation between these variables is 0.97.¹⁶ Thus, in this case, the actual gap is a very good proxy for the gap that actually determines the likelihood of adjustment.

2.4 Disruption Model

Another model of non-convex adjustment costs entails the disruption of production activities. This specification appears in work on capital adjustment, as in Cooper, Haltiwanger, and Power (1999) and Caballero and Engel (1999), where the introduction of new capital requires the shut-down of a production process. For employment adjustment, Cooper, Haltiwanger, and Willis (2003) find that this specification leads to a better match of employment dynamics in plant-level data than a model containing fixed or quadratic adjustment costs.

The dynamic programming problem for this case is characterized by equations (5), (6), and (7) with the the cost of adjustment function specified as

$$C(e, e_{-1}) = \begin{cases} (1 - \lambda^+)R(A, e, h) & \text{if } e > e_{-1} \\ (1 - \lambda^-)R(A, e, h) & \text{if } e < e_{-1} \\ 0 & \text{if } e = e_{-1}. \end{cases} \quad (9)$$

The fraction of revenue lost in adjustment process is parameterized (λ^+, λ^-) . This may, as in the capital adjustment model, reflect the actual shut-down of the production

¹⁵If the adjustment costs are not iid, then the value functions need to be recomputed adding F as a state variable, though the point of adjustment depending on the gap remains.

¹⁶This is the correlation between the differences in values and the absolute value of the gap.

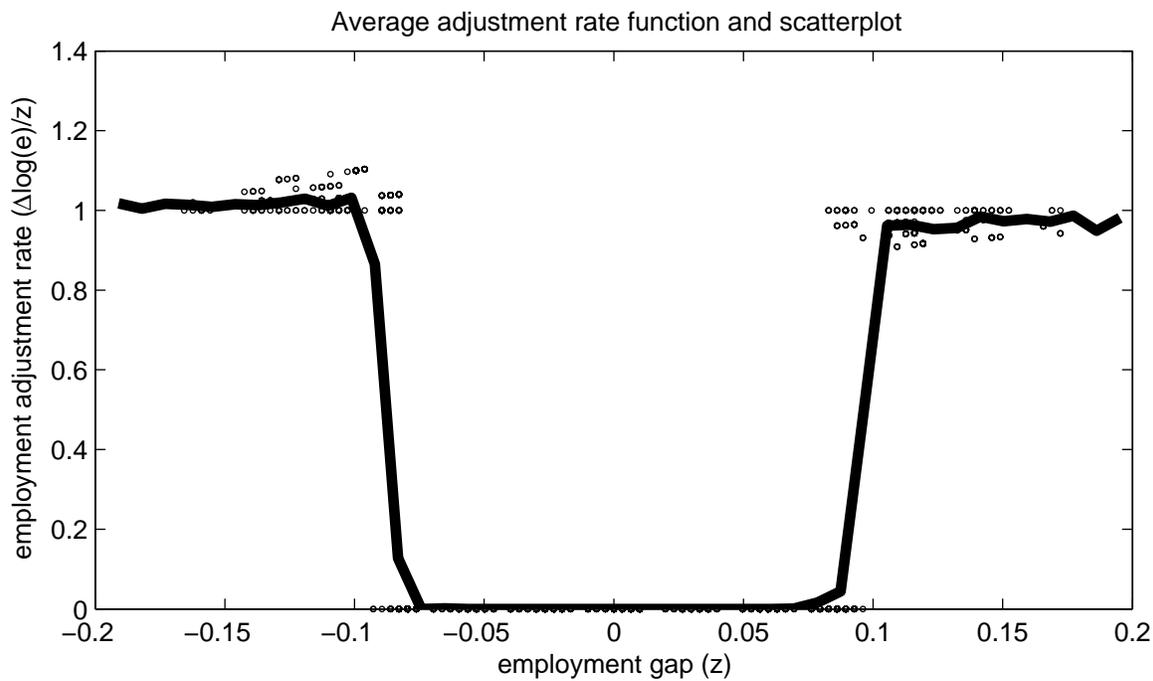
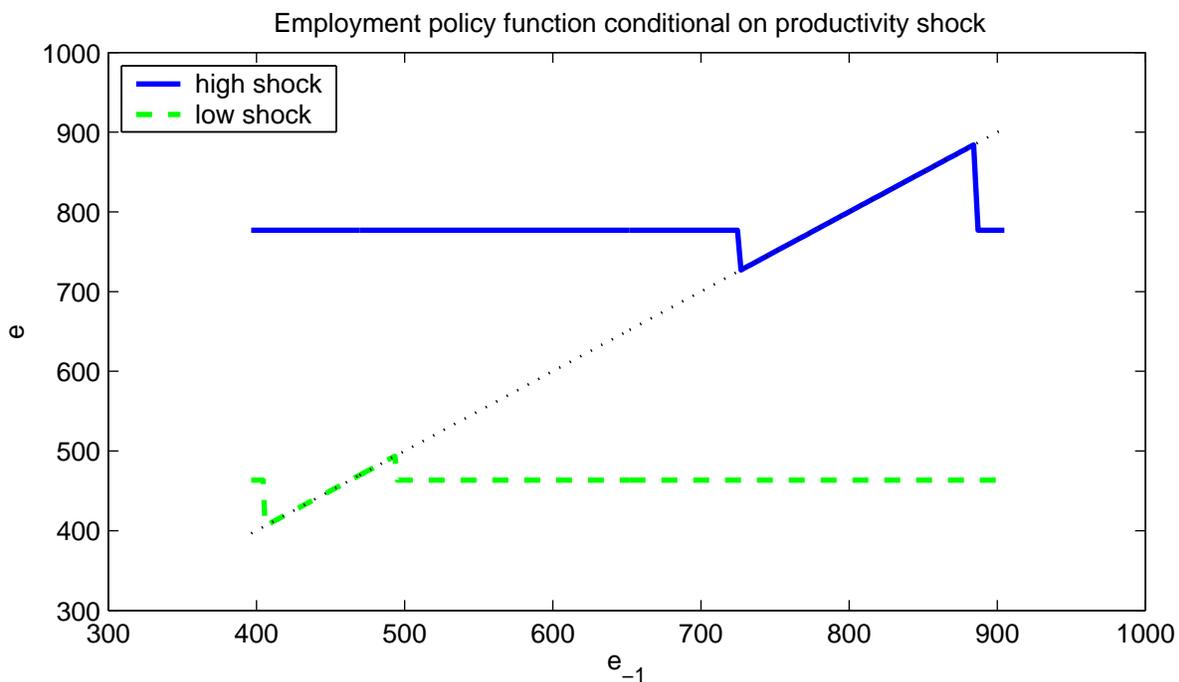
facility. More generally, these costs may reflect time lost due to training and reorganization of the labor force. This specification will also pick-up any loss in productivity associated with labor adjustment.

In contrast to the model with fixed adjustment costs, the disruption cost varies with the state of profitability at the plant. This creates an interesting dilemma. In periods of high profitability, it becomes costly to adjust the labor force. But these are periods when the addition of workers is most valuable. In periods of low profitability there is no trade-off as adjustment costs are relatively low and the plant has a desire to reduce the work force.

As seen in the policy functions in the top panel of Figure 3, this specification does give rise to inactivity in employment adjustment similar to that in the fixed adjustment cost model. The region of inaction is the interval of lagged employment in which $\phi_e(A, e_{-1}) = e_{-1}$. This figure is created assuming that $\lambda^+ = \lambda^- = 0.98$. With this specification, the plant loses 2% of its revenues during an adjustment period. While this loss may seem small, the inaction rate is 92% with this parameterization. The impact of changes in cost is visible through differences in the size of the region of inaction. In the high shock state (solid line), the high levels of productivity and desired employment will result in a larger cost to be paid if plants choose to adjust employment than in the low shock state (dashed line). This results in a wider range of lagged employment values in the high shock state over which the plant does not change employment (remaining on the 45 degree line) in comparison to the low shock state.

The bottom panel of Figure 3 shows the adjustment rate function. This panel is similar to the adjustment rate function for the fixed cost case in Figure 2 in that the adjustment rate is either 0 or almost 100%. The main difference is that with a disruption cost there is some overshooting and undershooting of the frictionless target, depending on the direction of adjustment. When plants increase employment in response to an above-average shock, they undershoot the target in some cases because the frictionless target is based on a momentary elimination of the disruption cost, where as in the model plants pay a cost that is increasing in the level of employment they choose. Therefore, they choose a slightly lower level of employment relative to the target, trading off some production revenue in order to reduce the size of the disruption cost. Likewise, in response to a below-average shock, plants that decrease employment will choose a lower level of employment, or overshoot, relative to the frictionless target in order to reduce the disruption cost.

Figure 3: $\phi_e(A, e_{-1})$ and $H(z)$ with Disruption Costs



2.5 Piecewise Linear Model

One of the important features of non-convex adjustment costs is that they imply inactivity in employment adjustment when the gains to changing the number of workers is small. This is relevant since at the plant level there is indeed evidence of inaction. Here we consider piecewise linear adjustment costs which can also produce inaction.

The cost of adjustment function is:

$$C(e, e_{-1}) = \begin{cases} \gamma^+ \Delta \tilde{e} & \text{if } e > e_{-1} \\ -\gamma^- \Delta \tilde{e} & \text{if } e \leq e_{-1}. \end{cases} \quad (10)$$

For the purposes of estimation, this specification strikes a useful compromise between the quadratic and fixed cost models. There is a basis for inaction due to the lack of differentiability in the neighborhood of zero adjustment. Small adjustments will not occur since the marginal cost of adjustment does not go to zero as the size of the adjustment goes to zero. As in the quadratic adjustment case, there is a cost per worker perhaps reflecting processing fees for applications and so forth. But, this specification of adjustment costs implies that there is no partial adjustment. Since the marginal cost of changing employment is constant, there is no basis for smoothing adjustment.

The optimal policy rules are determined by solving (2) using this specification of $C(e, e_{-1})$. The optimal policy is characterized by two boundaries: $e^-(A)$ and $e^+(A)$. If $e_{-1} \in [e^-(A), e^+(A)]$, then there is no adjustment. In the event of adjustment, the optimal adjustment is to $e^-(A)$ if $e_{-1} < e^-(A)$ and to $e^+(A)$ if $e_{-1} > e^+(A)$.

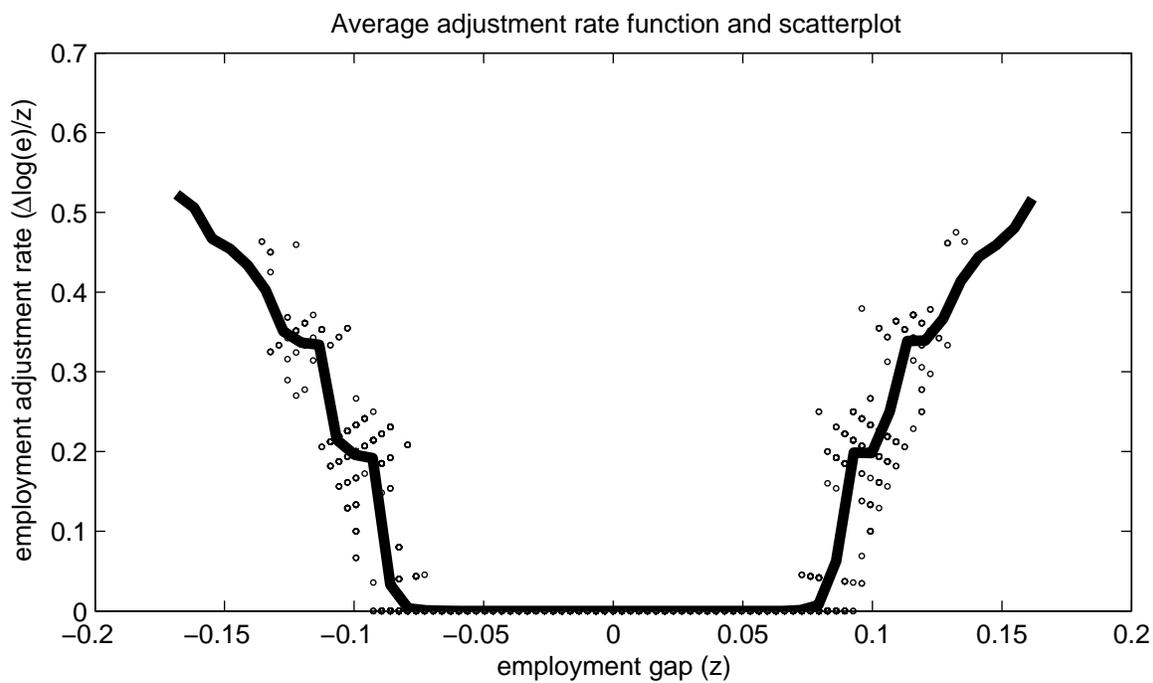
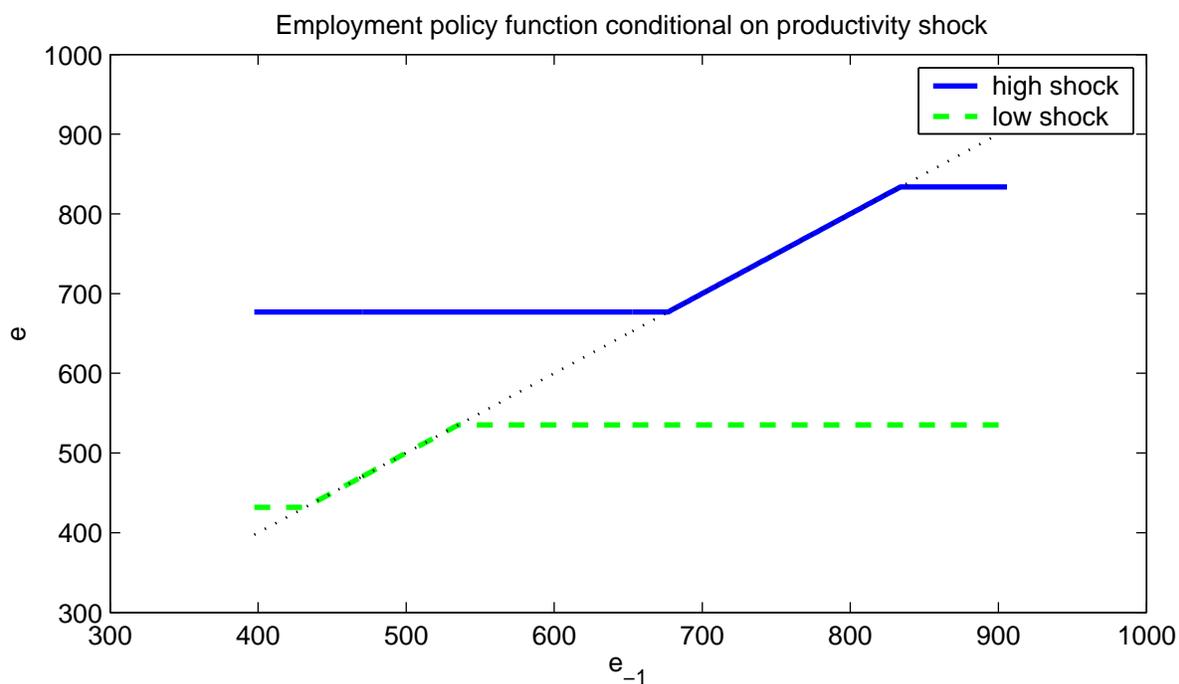
We can study the optimal policy function for this type of adjustment cost. Assume that $\gamma^+ = \gamma^- = 0.4$ which produces inaction at the plant level in 93% of the observations.¹⁷ Then (2) along with (10) can be solved using value function iteration and the resulting policy functions evaluated.

The policy functions are shown in Figure 4 for two values of the profitability shocks. There is no adjustment for values of e_{-1} in the interval where the employment policy function coincides with the 45° line. Outside of that interval there are two targets: $e^-(A)$ and $e^+(A)$. In contrast to the non-convex cases, there is no overshooting here. Further, in contrast to the quadratic adjustment cost case, there is no partial adjustment.

The bottom panel of Figure 4 shows the adjustment rate as a function of the gap. Recall that the target is defined as level of employment that would be chosen if all adjustment costs were removed for a single period. For the piecewise linear adjustment

¹⁷This inaction rate is too high relative to observation: the parameterization is for illustration only.

Figure 4: $\phi_e(A, e_{-1})$ and $H(z)$ with Piecewise Linear Adjustment Costs



cost model, the target is determined by solving (2) where the adjustment cost for the current period is zero, but expectations are taken over a value function that incorporates future adjustment costs.

In this figure, there is again a sizeable inaction region. In contrast to Figure 2, however, the region in which adjustment occurs is characterized by an upward sloping hazard function. This feature is similar to the adjustment rate function in Figure 1a of CEH, which was estimated from plant-level manufacturing data.

The upward sloping nature of the adjustment rate function, however, is primarily a result of using a target definition that is irrelevant for plants facing this particular form of adjustment costs. With piecewise linear adjustment costs, the only circumstance in which a plant would ever choose the target level is if the plant entered the period at the target level.

Returning to the top panel of Figure 4, the target level would lie somewhere between the two adjustment boundaries defined earlier. If adjustment occurs, the plant adjusts directly to the relevant boundary threshold. In that sense, the adjustment rate is 100% – there is no partial adjustment. As a result, plants in the extremes of the employment distribution appear to be adjusting at higher rates in the bottom panel simply because they require larger adjustment to reach the relevant boundary. In the adjustment rate calculation, this translates into a higher adjustment rate because they are adjusting by a larger fraction of the gap than a plant that lies closer to the optimal boundary threshold.

3 Estimation

Our primary goal is to estimate the adjustment costs for labor. In doing so, we consider specifications which mix the extreme cases highlighted above:

$$C(e, e_{-1}) = \begin{cases} \frac{\nu^+}{2} \left(\frac{e - e_{-1}}{e_{-1}} \right)^2 e_{-1} + F^+ + (1 - \lambda^+) R(A, e, h) + \gamma^+ \Delta e & \text{if } e > e_{-1} \\ \frac{\nu^-}{2} \left(\frac{e - e_{-1}}{e_{-1}} \right)^2 e_{-1} + F^- + (1 - \lambda^-) R(A, e, h) - \gamma^- \Delta e & \text{if } e < e_{-1} \\ 0 & \text{if } e = e_{-1}. \end{cases} \quad (11)$$

The corresponding dynamic optimization problem combines this specification of adjustment costs with (2).

Our estimation strategy follows the indirect inference approach, as in Gourieroux and Monfort (1996). In that approach, the researcher specifies an auxiliary model to summarize key elements from the data. For expositional convenience, we will refer to the

auxiliary model as a collection of moments from the data. The structural parameters are then estimated by matching these moments from the simulated data as close as possible to the data moments. We discuss the choice of moments in detail below. To estimate the structural parameters using the indirect inference procedure, we solve the dynamic programming problem given a specification of adjustment costs and given a vector of the parameters we are estimating. We focus on estimating the parameters of the adjustment cost function and set other parameters outside of the dynamic programming/estimation loop. The next subsection discusses the parameterization and simulation of the problem.

3.1 Parameterization and Simulation

To solve the various dynamic programming problems specified above, we need to calibrate a number of parameters and specify function forms. We assume:

- a Cobb-Douglas production function in which hours and employees are perfectly substitutable. Labor’s share is 0.65 and the markup is set at 25%.
- the discount factor (β) is set at 0.99 reflecting our use of quarterly data
- two parameters of the compensation function, w_0 and w_1 , were chosen so that steady state hours were 40 and that steady state employment at each plant was 600.¹⁸ The remaining parameter, ζ , is estimated.
- Following Cooper and Willis (2003), the profitability shock consists of two multiplicative exogenous components: an aggregate shock (A_{agg}) and an idiosyncratic shock (A_{idio}). Both exogenous components follow independent log-normal AR(1) processes. We assume that both processes have the same serial correlation, ρ , and innovations with the same standard deviation, σ . These two parameters are estimated in the indirect inference procedure.

After solving the dynamic programming problem, we then create simulated data. We simulate four panel data sets each consisting of 1000 plants for 200 periods. For each dataset, aggregate data series are created by summing employment data across plants and averaging hours data across plants in each period. The use of multiple datasets helps control for the exogenous choice of an initial distribution of “gaps” in the CE estimation procedure described in the next subsection.

¹⁸This is the average number of workers per plant in the balanced panel that underlies CEH. This figure comes from conversations with John Haltiwanger.

3.2 The CE Methodology

Our indirect inference approach uses the findings reported in CE as characterizing key moments of employment dynamics. We use these moments to infer the structural adjustment parameters. Thus it is important to understand the gap methodology.

CE consider the static optimization problem of a plant (firm) facing random walk shocks to both productivity and demand.¹⁹ With this specification, they replace the frictionless employment target with the state-contingent levels of employment from the static optimization problem.²⁰

Their estimation uses aggregate observations on net and gross flows for US manufacturing employment growth (denoted ΔE_{t+1}) to estimate a hazard function from

$$\Delta E_{t+1} = \int_{-\infty}^{\infty} (\Delta E_{t+1}^* - \tilde{z}) \Lambda(\tilde{z} - \Delta E_{t+1}^*) f_t(\tilde{z}) dz. \quad (12)$$

Here \tilde{z} represents the gap after the plant has been subjected to any idiosyncratic shocks.²¹ So $\Delta E_{t+1}^* - \tilde{z}$ is the size of the gap after any common shocks to the desired level of employment.

Here the left-side is period $t + 1$ employment growth. This reflects two factors: the fraction of plants adjusting and the magnitude of their adjustment. The function $\Lambda(\tilde{z} - \Delta E_{t+1}^*)$ is the aggregate hazard reflecting either the probability or the rate of adjustment. The argument of this function is the size of the plant-specific gap **after** the aggregate growth in the frictionless employment target, denoted ΔE_{t+1}^* , has taken place. In a non-convex adjustment cost environment, the function $\Lambda(\tilde{z} - \Delta E_{t+1}^*)$ represents the fraction of plants that adjust by the amount of $\Delta E_{t+1}^* - \tilde{z}$. In a quadratic adjustment cost environment, all plants are adjusting by a fraction, $\Lambda(\tilde{z} - \Delta E_{t+1}^*)$, of the total desired adjustment, $\Delta E_{t+1}^* - \tilde{z}$. In either case, the final term in (12) is $f_t(\tilde{z})$, the distribution of the gap prior to the aggregate shock to the employment target.

Of course, ΔE_{t+1}^* is not observable. CE create it from

$$\Delta E_t^* = \Delta E_t + \kappa \Delta H_t \quad (13)$$

where ΔE_t is observed aggregate employment growth and ΔH_t is observed aggregate

¹⁹Since they work with aggregated data, the distinction between a plant and a firm is irrelevant.

²⁰They argue this is legitimate due to the random walk shocks. Cooper and Willis (2003) argue that if shocks do not follow a random walk, then the static targets may not be proportional to the frictionless targets so that this procedure may be problematic.

²¹Thus \tilde{z} and the gap measure defined in (1) reflect the same concept measured at two different points in time.

hours growth. CE show this relationship comes from the aggregation over plants' static optimization problems. The parameter κ is calibrated from the static optimization problem of the plant.²²

As CE do not have plant-level data, they can not directly measure the distribution $f_t(z)$. This is created in their estimation by introducing plant-specific variations in z after the common variations. In particular, following the common shocks to the employment target, one-half of the plants have their gap increased (decreased) by an amount denoted σ_I .²³ This parameter is estimated as well.

CE consider a quadratic specification for $\Lambda(\cdot)$

$$\Lambda(\tilde{z}) = \tilde{\lambda}_0 + \tilde{\lambda}_2(\tilde{z} - z_0)^2 \quad (14)$$

where z_0 is a constant. To obtain parameter estimates, they use observations on employment and hours growth in (13) to calculate the growth rate of the employment target. This measure is then used in (12) along with (14). CE find parameter values for the hazard that minimize the sum of squared differences between the actual and predicted employment growth.²⁴

CE (Table 2, BLS) report the following estimates (standard errors in parentheses): $\tilde{\lambda}_0 = 0.02$ (0.01); $\tilde{\lambda}_2 = 0.53$ (0.01); $z_0 = -0.82$ (0.01) and $\sigma_I = 0.059$ (0.015). CE conclude that a quadratic hazard specification fits the data better than the flat hazard.

3.3 Estimation Results

The estimation of the structural parameters, denoted θ , involves solving the following minimization problem

$$\mathcal{L}(\theta) = \min_{\theta} [\Psi^d - \Psi^s(\theta)]' W [\Psi^d - \Psi^s(\theta)]. \quad (15)$$

where the vector Ψ^d represents the moments to be matched and W represents the weighting matrix.

²²Cooper and Willis (2003) provide a complete discussion of this derivation and its link to the approach of CEH which argues that κ represents the sensitivity of hours to the gap after all employment adjustment.

²³There is no direct counter-part of this shock in the dynamic optimization problem. It should not be confused with the plant-specific component of the profitability shock. Still, we can introduce this as noise in our simulation, following the procedure of CE.

²⁴The procedure for the estimation is laid-out in section IV.2 of CE. Our estimation mimics their procedure.

Seven moments are chosen for the estimation procedure. The first four moments are the results reported in CE corresponding to the following parameters of the hazard function: $(\lambda_0, \lambda_2, z_0, \sigma_I)$. The fifth element in the vector of moments is the sum of squared residuals (SSR) for aggregate employment in the CE procedure. The final two moments are the serial correlation, ρ_E , and standard deviation, σ_E , of the log of manufacturing employment from the BLS sample used in CE.

The CE hazard moments are chosen for two reasons. First, the results and conclusions in CE are frequently taken as evidence of the aggregate importance of non-convex adjustment at the micro-level. Thus these are “topical” moments. Second, in simulation, these moments are informative for determining the underlying structural parameters. Thus while Cooper and Willis (2003) point out that the results of CE can not be directly used to evaluate any particular model of adjustment costs, these regression results are informative enough to be useful as a basis for structural estimation via indirect inference.

The sum of squared residuals from the CE procedure is included to minimize the SSR between actual and predicted aggregate employment growth in our simulated environment. The motivation for including this statistic comes from its prominent role in CE. CE use this statistic not only to determine the best estimate for the parameters of a given hazard model, but also to compare different hazard specifications. If we had the true model, this sum of squared residuals should be 0. We include this moment in Ψ^d so that our selection of structural parameters will, in part, minimize this goodness of fit measure.²⁵

The final two moments come from the aggregated BLS data for the 1972-86 period used in CE. In particular, to capture the aggregate shock process, we include the serial correlation and standard deviation of manufacturing employment at a quarterly frequency. These moments help to determine the driving process as well as the elasticity of compensation with respect to hours.

Based upon values from CE and BLS data, the vector of moments for estimation is $\Psi^d = (0.019, 0.53, -0.816, 0.059, 0, 0.96, 0.0607)$. The weighting matrix, W , is constructed as the inverse of a diagonal matrix consisting of the empirical variances of the moments in Ψ^d . The standard errors for the serial correlation and standard deviation of log employment are $(0.042, 0.032)$. The standard errors for the CE hazard parameters are as reported earlier. The standard error of BLS manufacturing employment growth,

²⁵A similar approach of mixing moments and nonlinear least squares is presented in Adda and Cooper (2000).

0.019, is used as a measure of the standard error of the sum of squared residual.

The vector of structural parameters which we estimate is given by

$$\theta = (\rho, \sigma, \zeta, \nu^+, \nu^-, F^+, F^-, \gamma^+, \gamma^-, \lambda^+, \lambda^-).$$

The first two components characterize the process for both the aggregate and idiosyncratic shocks.²⁶ The third parameter, ζ , controls the sensitivity of wages to hours. The last eight components are the parameters for the adjustment costs.²⁷

We estimate θ by simulation. For a given vector of parameters, we solve our dynamic optimization problem, and we simulate using the policy functions. We aggregate the data and follow the CE procedure to estimate the parameters of the hazard function. We also compute the other moments from the simulated data. Thus these moments depend on the parameter vector: $\Psi^s(\theta)$.

Our main findings are reported in Tables 1 and 2. The first table reports structural parameter estimates and standard errors for a number of specifications of adjustment costs: quadratic, fixed, piece-wise linear, disruption and mixtures of quadratic with each of the others.²⁸ The second table reports the moments (i.e. the estimated hazards) for each set of structural parameters.

A couple of points stand out. First, the best fit is obtained with a piece-wise quadratic adjustment cost model. In particular, the parameter determining the cost of hiring, $\nu^+ = 2.77$, is significantly larger than the parameter determining the cost of firing, $\nu^- = 1.14$.²⁹ Second, the estimated elasticity of the wage function, $\zeta = 2.03$, is close to the calibrated value of 1.9 used by CE.

The model with piece-wise linear adjustment costs also does well in terms of matching the estimated hazard function but does relatively poorly in terms of fitting the employment series in the CE procedure. To be precise, the sum of squared residuals (SSR) in the CE procedure is 0.0055 for the quadratic model. For the piece-wise linear cost specification it is 50 percent larger, $SSR = 0.0088$. This contributes to the loss function, $\mathcal{L}(\theta)$, being more than 50 percent greater than for the quadratic model.

²⁶Recall that for parsimony, we restrict the aggregate and idiosyncratic processes to have the same values of ρ, σ .

²⁷Throughout the fixed costs, $\{F^+, F^-\}$, are normalized to be expressed as a fraction of average profits. These profit measures do not include adjustment costs.

²⁸Standard errors are calculated as described by Gourieroux and Monfort (1996).

²⁹Pfann and Palm (1993) also find evidence of asymmetry in their study of labor adjustment in the Netherlands and the U.K. Interestingly, for production workers (as in our study), hiring costs are more substantial than firing costs.

The fixed adjustment cost model does the best in terms of fitting the employment series with an SSR of 0.003. It does very poorly, however, in terms of matching the estimated hazard function. Accordingly, this model has a very large value of $\mathcal{L}(\theta)$.

The model with disruption costs, like the piece-wise linear model, is able to match the parameters of estimated hazard function but has a relatively high value of the SSR statistic. Note too that the estimates of (λ^+, λ^-) are both very close to one, indicating very small disruption costs.

Thus, the feature distinguishing these specifications is their ability to match both the hazard estimates from CE and aggregate employment growth. Evidently, the hazard estimates from CE alone are not sufficient to distinguish the models of adjustment costs. Once both the hazard and SSR measures are properly taken into account, the quadratic adjustment cost model fits best.

The other specifications included in Tables 1 and 2 entail a mixing of adjustment costs. None of these specifications leads to a significant improvement over the quadratic adjustment cost model. In particular, the parameter estimates for the fixed costs and piece-wise linear parameters are negligible. Thus again, the model which best fits the CE reduced form estimates and SSR measure is a piecewise quadratic structure.

4 Conclusions

This paper provides a structural interpretation of the regression results reported in CE. Those results, using the gap methodology, indicated that aggregate employment growth depended on the cross sectional distribution of the employment gap. That analysis left open an interpretation of the finding.

By using an indirect inference methodology, this paper finds that the results reported in CE are consistent with a model in which there are asymmetric quadratic adjustment costs at the individual plant level. Evidently, this adjustment cost structure is rich enough to generate the nonlinearities reported by CE.

This conclusion may appear surprising since the CE evidence is usually viewed as supporting some form of non-convex adjustment costs. Yet, as pointed out in Cooper and Willis (2003), using the methodology of CE, even a quadratic adjustment cost model can produce the nonlinearities that they uncover. This paper goes beyond that point by finding the adjustment cost structure which generates the regression results of CE.

Note too that this conclusion is about the aggregate regression results reported by CE. Our finding does not imply that non-convexities are not relevant at the plant level.

Cooper, Haltiwanger, and Willis (2003) finds evidence of non-convex adjustment costs (through disruption) using plant-level data. Instead, our results simply imply that non-convexities at the plant level are not needed to explain the aggregate employment dynamics examined by CE.

Finally, one concern with estimated models of adjustment costs is interpretation: what is captured by this black box? One interpretation, following Hamermesh and Pfann (1996) is that these adjustment costs represent search frictions.

Expanding the labor force requires the plant to post vacancies which are filled through a matching process. A plant that posts ten vacancies may, in the first year, fill seven of them. In the second year, two more may be filled. Of course the match rate is stochastic. Intuitively, the model ought to deliver a partial adjustment type outcome.

This interpretation is also consistent with the asymmetry in the adjustment costs. Clearly, there is no search friction associated with firing workers. Formalizing the search structure and evaluating its implications from the perspective of labor adjustment costs is an important extension of the model. This structure also suggests that current estimates of labor adjustment costs ought to include economy-wide variables, such as unemployment rates and vacancies rates, insofar as these variables influence matching probabilities.

References

- ADDA, J., AND R. COOPER (2000): “Balladurette and Juppette: A Discrete Analysis of Scrapping Subsidies,” *Journal of Political Economy*, 108(4), 778–806.
- 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., J. HALTIWANGER, AND L. POWER (1999): “Machine Replacement and the Business Cycle: Lumps and Bumps,” *American Economic Review*, 89, 921–946.
- COOPER, R., J. HALTIWANGER, AND J. WILLIS (2003): “Dynamics of Labor Demand: Evidence from Plant-level Observations and Aggregate Implications,” Federal Reserve Bank of Kansas City, Research Working Paper No. 03-12.
- COOPER, R., AND J. WILLIS (2001): “The Economics of Labor Adjustment: Mind the Gap,” NBER Working Paper # 8527.
- (2003): “The Economics of Labor Adjustment: Mind the Gap,” Federal Reserve Bank of Kansas City, Research Working Paper 03-05.
- GOURIEROUX, C., AND A. MONFORT (1996): *Simulation-Based Econometric Methods*. Oxford University Press.
- HAMERMESH, D. (1989): “Labor Demand and the Structure of Adjustment Costs,” *American Economic Review*, 79, 674–89.
- (1993): *Labor Demand*. Princeton University Press.
- HAMERMESH, D., AND G. PFANN (1996): “Adjustment Costs in Factor Demand,” *Journal of Economic Literature*, 34, 1264–92.
- PFANN, G., AND F. PALM (1993): “Asymmetric Adjustment Costs in Non-linear Labour Models for the Netherlands and U.K. Manufacturing Sectors,” *Review of Economic Studies*, 60, 397–412.

- RUST, J. (1987): “Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher,” *Econometrica*, 55(5), 999–1033.
- SARGENT, T. (1978): “Estimation of Dynamic Labor Demand Schedules under Rational Expectations,” *Journal of Political Economy*, 86(6), 1009–1044.
- SHAPIRO, M. (1986): “The Dynamic Demand for Labor and Capital,” *Quarterly Journal of Economics*, 101, 513–42.

Table 1: Structural Parameter Estimates

Model	Structural Parameters (std errors below)										
	ν^+	ν^-	F^+	F^-	γ^+	γ^-	λ^+	λ^-	ζ	ρ	σ
Quadratic	2.77 (0.59)	1.14 (0.03)							2.03 (0.15)	0.73 (0.022)	0.028 (0.003)
Fixed			0.0088 (0.017)	0.014 (0.035)					2.00 (0.43)	0.97 (0.009)	0.020 (0.002)
P-wise linear					-0.028 (0.173)	0.295 (0.108)			2.36 (0.13)	0.84 (0.014)	0.033 (0.001)
Disruption							0.991 (0.036)	0.996 (0.024)	2.50 (0.17)	0.83 (0.04)	0.025 (0.003)
Quad./P-wise	2.77 (5.72)	1.14 (0.12)			0 (3.395)	0 (3.548)			2.03 (0.64)	0.73 (0.17)	0.028 (0.014)

Table 2: Moments

Model	Reduced-Form Estimates							
	λ_0	λ_1	z_0	σ_I	SSR	ρ_E	σ_E	$\mathcal{L}(\theta)$
CE results*	0.019	0.53	-0.816	0.059	0.0028	0.96	0.061	na
Quadratic	0.018	0.53	-0.81	0.069	0.0055	0.915	0.049	17.45
Fixed	-0.010	0.44	-0.43	0.094	0.0030	0.994	0.049	1300.10
P-wise linear	0.019	0.52	-0.80	0.061	0.0088	0.954	0.057	28.96
Disruption	0.013	0.53	-0.81	0.105	0.0096	0.943	0.050	37.51
Quad./P-wise	0.018	0.53	-0.81	0.069	0.0055	0.915	0.049	17.45

*Table 1, p. 375.