

Federal Reserve Bank of Minneapolis
Research Department Staff Report 310

August 2002

The Economics of Labor Adjustment: Mind the Gap*

Russell Cooper

Boston University and Federal Reserve Bank of Minneapolis

Jonathan L. Willis

Federal Reserve Bank of Kansas City

ABSTRACT

We study inferences about the dynamics of labor adjustment obtained by the “gap methodology” of Caballero and Engel [1993] and Caballero, Engel and Haltiwanger [1997]. In that approach, the policy function for employment growth is **assumed** to depend on an unobservable gap between the target and current levels of employment. Using time series observations, these studies reject the partial adjustment model and find that aggregate employment dynamics depend on the cross-sectional distribution of employment gaps. Thus, nonlinear adjustment at the plant level appears to have aggregate implications. We argue that this conclusion is not justified: these findings of nonlinearities in time series data may reflect mismeasurement of the gaps rather than the aggregation of plant-level nonlinearities.

*We are grateful to seminar participants at Boston University, Emory University, the University of British Columbia, the Federal Reserve Bank of Boston, the University of Haifa, the University of Iowa, the University of Pennsylvania, the University of Texas at Austin and the 2000 CMSG conference at McMaster University for comments and suggestions. Discussions with John Haltiwanger, Daniel Hamermesh, Peter Klenow and Christopher Ragan were much appreciated. The authors thank the NSF for financial support. We also appreciate comments from a referee and the editor. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City, the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

In recent contributions, Caballero and Engel [1993], hereafter CE, and Caballero, Engel and Haltiwanger [1997], hereafter CEH, investigate labor dynamics using a methodology which postulates that employment changes depend on a gap between the actual and target levels of employment.¹ Both studies find evidence of nonlinearities in aggregate time-series data: employment growth depends on the cross-sectional distribution of employment gaps. This finding is taken as evidence that nonlinear adjustment at the microeconomic level “matters” for aggregate time series. This is important for business cycle and policy analysis as it implies macroeconomics must take plant-level distributions into account. This paper questions the methodology and thus the conclusions of these studies.² We argue that these reported aggregate nonlinearities may be the consequence of mismeasurement of the gap rather than nonlinearities in plant-level adjustment.³

Both CE and CEH rely upon a hypothesis that employment changes (Δe) respond to a gap (z) between the desired and actual number of workers at a plant. The advantage to the gap approach is that the choice of employment, an inherently difficult dynamic optimization problem, is characterized through a nonlinear relationship between (Δe) and (z). That is, the adjustment rate, $\Delta e/z$, is a nonlinear function of z . However, there is no “free lunch”: the desired number of workers, and hence the employment gap, is unobservable. Thus, in order to confront data, this approach needs an auxiliary theory to infer z from observed variables. Both CE and CEH use observed hours variations to infer the employment gap: this inference is one element of our critique of the gap methodology.

To assess this methodology, we construct a dynamic model of labor adjustment **assuming** quadratic adjustment costs and follow the approaches of CE and CEH to analyze their implications. The quadratic adjustment cost model is a useful benchmark for two reasons. First, it has served as the primary model for the study of aggregate employment dynamics. To quote CE (p. 365),

Since the latter [representative-agent framework with quadratic adjustment costs] is a specification often used by macroeconomists to characterize aggregate dynamics, it constitutes a convenient benchmark for a discussion of the more realistic increasing-hazard [adjustment rate] models.

¹Hamermesh [1989] uses a gap methodology as well but does not adopt the approach of estimating a nonlinear hazard function (explained below) to infer the nature of adjustment costs. Hence we focus on CE and CEH in this discussion of methodology.

²We do not contest the general view of nonlinear employment adjustment at the plant-level. This finding is consistent with other evidence that points to inactivity as well as bursts of employment adjustment at the plant-level. For example, Hamermesh [1989] provides a revealing discussion of lumpy labor adjustment at a set of manufacturing plants. Davis and Haltiwanger [1992] document large employment changes at the plant level. CEH also report evidence of inactivity in plant-level employment adjustment. There seems little doubt that an explanation of plant level employment dynamics requires a model of adjustment that is richer than the quadratic adjustment cost structure and includes some forms of non-differentiability and/or nonconvexity.

³To the extent that the gap approach is used in numerous other applications, our concerns may be relevant for those exercises as well.

Second, the quadratic adjustment cost model is nested in the employment gap approach. As a matter of theory, if adjustment costs are quadratic, shocks follow a random walk and the gap is correctly measured, then the adjustment rate is constant, implying that aggregate employment is independent of the cross-sectional distribution of employment gaps.⁴ From our simulations, this result holds when the gap is properly measured, even if shocks do not follow a random walk.⁵ Thus, if the CE and CEH procedures uncover aggregate nonlinearities from a data set created from a model with quadratic adjustment costs and stationary shocks, then this is a consequence of mismeasurement of the gap rather than economic fundamentals.

In our quadratic adjustment cost model we find the following:

- If the gap is correctly measured, the adjustment rate is essentially constant and the cross-sectional distribution of employment gaps is irrelevant for aggregate employment dynamics.
- If the employment gap is **mismeasured**, then
 1. A quadratic cost of adjustment model can generate a nonlinear adjustment rate ($\Delta e/z$ depends nonlinearly on z).
 2. Aggregate employment dynamics can depend on the cross-sectional distribution of the employment gap.
- The gap is mismeasured using either the CE or the CEH approaches.

Problems measuring the gap may be severe enough to create nonlinearities that otherwise would not be present. This is true for the simulated data and may well lie behind the results reported by CE and CEH for actual data.

We thus conclude that the time-series evidence of nonlinear hazards reported by CE and CEH does not necessarily imply that nonlinear adjustment at the plant-level has aggregate effects. A methodology which is unbiased under the null hypothesis of quadratic adjustment costs is needed to assess the aggregate implications of that model relative to a competing model with non-convex adjustment costs.

2. The Gap Approach: An Overview

As background for our analysis, we begin with a summary of the methodology employed by CE and CEH as well as a more precise statement of their findings.

⁴If the adjustment rate is independent of the gap, then the cross-sectional distribution of the gap is irrelevant for aggregate behavior. The fact that the partial adjustment model implies constant adjustment is essentially by construction. The link between the quadratic cost of adjustment structure and the partial adjustment model is more subtle and is discussed further below.

⁵The issue of the correlation of the shocks is important and one that we return to below. We are grateful to the referee for stressing this point.

A. Gap Methodology

We follow the notation and presentation in CEH. The gap between the desired employment and the actual employment (in logs) in period t for plant i is defined as

$$\tilde{z}_{i,t} \equiv e_{i,t}^* - e_{i,t-1}. \quad (1)$$

Here $e_{i,t}^*$ is the desired level of employment given the realization of all period t random variables and $e_{i,t-1}$ is the level of employment prior to any period t adjustments. Thus $\tilde{z}_{i,t}$ represents a gap between the state of the plant at the beginning of the period and the level of employment it would choose if it could “costlessly” alter employment.

CEH hypothesize a relationship between employment growth $\Delta e_{i,t}$ and $\tilde{z}_{i,t}$ given by

$$\Delta e_{i,t} = \phi(\tilde{z}_{i,t}). \quad (2)$$

A key issue is characterizing the policy function, $\phi(z_{i,t})$, and inferring properties of adjustment costs from it. In some cases, it is convenient to refer to an adjustment rate or hazard function:⁶

$$\Phi(\tilde{z}_{i,t}) \equiv \phi(\tilde{z}_{i,t})/\tilde{z}_{i,t}.$$

Specifying that employment adjustment depends only on the gap is an assumption: the validity of this approximation to the optimal policy function of the plant can be evaluated using our structural model.

Because the gap is central to this analysis, it is important to be very precise about how it is defined and measured. The key is the meaning of “costlessly adjusting employment.” In fact, there are two ways to characterize the target, and as we demonstrate in our quantitative analysis, the results depend on the definition.

First, one could define the target as the level of employment that would arise if there were **never** any costs of adjustment.⁷ This version of the target is quite easy to characterize since it solves a static optimization problem. This is termed the **static target** in the discussion that follows.

Second, one could construct a target measure in which the adjustment costs are re-

⁶From the discussion in CE and CEH, there are two interpretations of this function. $\Phi(z)$ represents either the magnitude of adjustment (e.g., the fraction of a gap that is closed) or a probability of full adjustment to the target. The interpretation, of course, depends on the nature of adjustment costs. For the quadratic adjustment cost case, this function represents the rate at which the gap is closed. For a model with stochastic adjustment costs taking the values zero and infinity, this is the probability of full adjustment. We use the terminology of a hazard function throughout, as do CE and CEH.

⁷This approach to approximating the dynamic optimization problem is applied extensively but, from our perspective, places too much emphasis on static optimization. Nickell [1978, pp. 329–30] says,

the majority of existing dynamic models of factor demand simply analyse the optimal adjustment of the firm towards a static equilibrium and it is very difficult to deduce from this anything whatever about optimal behaviour when there is no “equilibrium” to aim at.

moved for a single period. The target would correspond to the level of employment to which an optimizing agent would eventually adjust, *ceterus paribus*. This is termed the **frictionless target**. For the quadratic adjustment cost model, this target would represent the level of employment where the state-dependent policy choice for current employment, expressed as a function of previous employment, crosses the 45-degree line.

This hypothesized relationship between employment changes and the gap cannot be analyzed directly since $\tilde{z}_{i,t}$ is a theoretical construct that is not observed: there exists no data set which includes $\tilde{z}_{i,t}$. In the literature, various approaches have been pursued.

B. CEH Measurement of the Gap and Findings

CEH hypothesize a second relationship between another (closely related) measure of the gap, ($\tilde{z}_{i,t}^1$), and plant-specific deviations in hours:

$$\tilde{z}_{i,t}^1 = \theta(h_{i,t} - \bar{h}). \quad (3)$$

Here $\tilde{z}_{i,t}^1$ is the gap in period t **after** adjustments in the level of e have been made: $\tilde{z}_{i,t}^1 = \tilde{z}_{i,t} - \Delta e_{i,t}$.⁸

Intuition says θ should be positive. As profitability rises, hours and the desired number of workers will both increase. The gap decreases as workers (e) are added and hours fall closer to \bar{h} . Thus the supposed relationship between this measure of the gap and hours deviations seems reasonable, both in terms of the response of these variables to a shock and in terms of transition dynamics. Note, though, that the correlation between hours and employees is somewhat complicated: the shock leads to positive comovement between e and h , but in the adjustment process, the comovement is negative.

Rewriting this relationship in terms of the pre-adjustment gap leads to

$$\tilde{z}_{i,t} = \theta(h_{i,t} - \bar{h}) + \Delta e_{i,t}. \quad (4)$$

Hence, given an estimate of θ , one can infer $\tilde{z}_{i,t}$ from hours and employment observations.

The issue is estimating θ . Using (1) in (4) and taking differences yields

$$\Delta e_{i,t} = -\theta \Delta h_{i,t} + \Delta e_{i,t}^*. \quad (5)$$

Adding a constant (δ) and noting that $\Delta e_{i,t}^*$ is not observable, CEH estimate θ from

$$\Delta e_{i,t} = \delta - \theta \Delta h_{i,t} + \varepsilon_{i,t}. \quad (6)$$

⁸Implicitly this assumes that there is no lag between the decision to adjust employment and the actual adjustment. That is, unlike the time to build aspect of investment, employment adjustments take place immediately. We use this timing assumption in our structural model.

Further, we have removed the heterogeneity in \bar{h} and in θ that is important for the empirical implementation in CEH. Finally, note that by assumption \bar{h} is independent of any shocks to the profitability of employment. We will argue below that this is an important restriction.

As CEH note, estimation of this equation may yield biased estimates of θ since the error term (principally $\Delta e_{i,t}^*$) is likely to be correlated with changes in hours. That is, a positive shock to profitability may induce the plant to increase hours (at least in the short run) and will generally cause the desired level of employment to increase as well. CEH argue that this problem is (partially) remedied by looking at periods of large adjustment since then the changes in hours and employment will overwhelm the error.⁹ As we proceed, evaluating the implications of this bias will be important.

CEH use their plant-level measures of the gap in two ways. First, they analyze the relationship between employment adjustment and employment gaps at the plant level. Second, they investigate aggregate implications by estimating a reduced form hazard function from time series. Letting $f_t(z)$ be the period t probability density function of employment gaps across plants, the aggregate rate of employment growth is given by

$$\Delta E_t = \int_z z \Phi(z) f_t(z) dz. \quad (7)$$

As $\Phi(z)$ is the adjustment rate or hazard function indicating the fraction of the gap that is closed by employment adjustment, $z\Phi(z)$ is the size of the employment adjustment for plants with a gap of z . As in CEH [Section IV], simplification of (7) given the specification of a hazard function produces an aggregate relationship between employment changes and non-centered moments of the distribution of z .

The findings by CEH are summarized as follows:

- Using (6), CEH report a mean (across 2-digit industries) estimate of $\theta = 1.26$. Their estimate comes from using observations in which percent changes in both employment and hours exceed one standard deviation of the respective series.
- Using their estimates of θ to construct a gap measure ($\tilde{z}_{i,t}$), CEH (Figure 1a, p. 122) find a nonlinear relationship between the average adjustment rate, $\Phi(\tilde{z}_{i,t})$, and $\tilde{z}_{i,t}$.
- CEH specify that

$$\Phi(z) = \begin{cases} \lambda_0 + \lambda_1^- z & \text{for } z < 0 \\ \lambda_0 + \lambda_1^+ z & \text{for } z > 0 \end{cases} \quad (8)$$

and estimate $\lambda_1^- = 1.30$ and $\lambda_1^+ = 1.32$. Hence, employment growth, expressed by (7), depends on the second moment of the distribution of employment gaps.

C. CE Measurement of the Gap and Findings

In contrast to CEH, CE do not estimate θ but instead calibrate it from a structural model of static optimization by a plant with market power.¹⁰ The appendix characterizes the

⁹They also note the presence of measurement error, which they address through the use of a reverse regression exercise. We have not included measurement error in our simulated environment.

¹⁰Though CE do not have any microeconomic data, CEH work with plant-level data, and so we refer to these microeconomic units as *plants*.

mapping from the structural parameters of the quadratic adjustment cost model (presented in the next section) to θ .

An important element in their approach is the use of a static target. CE argue that the static targets are relevant benchmarks for measuring employment gaps if shocks follow a random walk because they will only differ from the frictionless targets by a constant. But, if the shocks are stationary, this simple relationship between frictionless and static targets will no longer hold. Instead of adhering to a fixed deviation from the static solution, plants will solve a dynamic optimization problem, explored below, taking into account conditional expectations of future shocks. Plants balance the gains from adjusting to productivity shocks against the costs imposed on employment adjustment in the future. We analyze the bias in the measurement of the gap stemming from the use of a static target in the presence of stationary shocks.¹¹

Because CE do not have plant-level data, their estimation uses aggregate observations on net and gross flows for U.S. manufacturing employment to estimate a hazard function. They create the following measure of the growth of the aggregate target using the calibrated value for θ :

$$\Delta E_t^* = \Delta E_t + \theta \Delta H_t. \quad (9)$$

This growth of the target, which is an aggregate version of (5), is then used in a specification of employment growth:

$$\Delta E_{t+1} = \int_{-\infty}^{\infty} (\Delta E_{t+1}^* - z) \Lambda(z - \Delta E_{t+1}^*) f_t(z) dz \quad (10)$$

which is similar to (7).¹²

CE consider both a constant and a quadratic specification for $\Lambda(\cdot)$. To obtain parameter estimates, they calculate the growth rate of the employment target from (9) using observations on employment and hours growth. This measure is then used in (10).

CE estimate a quadratic hazard:

$$\Lambda(z) = \tilde{\lambda}_0 + \tilde{\lambda}_2(z - z_0)^2 \quad (11)$$

where z_0 is a constant. Given a specification of the hazard function, they generate a predicted growth rate for employment and a predicted sequence of cross-sectional distributions of the gap. They choose parameter values for the hazard that minimize the sum of squared differences between the actual and predicted employment growth. CE (p. 375, Table I, BLS)

¹¹Determining whether shocks are best classified as stationary or a random walk is an unresolved empirical question. We will present evidence in later sections to support our assumption that shocks follow a stationary process.

¹²However, the notation and definitions of CEH differ from those used by CE. In particular, CE define the gap as $\tilde{z}_{i,t} \equiv e_{i,t} - e_{i,t}^*$. Accordingly, this expression for aggregate employment growth differs from that of CEH.

report the following estimates: $\tilde{\lambda}_0 = 0.02$; $\tilde{\lambda}_2 = 0.53$; $z_0 = -0.82$. CE conclude that a quadratic hazard specification fits the data better than the flat hazard.

3. A Dynamic Optimization Framework

Our analysis builds from the specification of a dynamic optimization problem at the plant level. Our structure is purposely close to that outlined in CE.¹³ We use the model as a data-generating mechanism to evaluate the CE and CEH methodologies.

A. Quadratic Adjustment Cost Model

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

$$V(A, E_{-1}) = \max_{H, E} R(A, E, H) - \omega(E, H) - \frac{\nu}{2} \left(\frac{E - E_{-1}}{E_{-1}} \right)^2 E_{-1} + \beta \xi_{A'|A} V(A', E) \quad (12)$$

for all (A, E_{-1}) . Here H represents the input of hours by a worker, E_{-1} is the inherited stock of workers, E is the stock of current workers, and ξ is the expectations operator.¹⁴ Note the timing assumption of the model: workers hired in a given period become productive immediately.

For our analysis we use a Cobb-Douglas production function in which the labor input is simply the product of employment and hours, EH . Allowing for market power of the plant, the revenue function is specified as

$$R(A, E, H) = A(EH)^\alpha \quad (13)$$

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

The costs of adjustment are assumed to be a quadratic function of the percent change in the stock of employed workers multiplied by the initial stock of employees.¹⁵ That is, the adjustment cost arises for net, *not* gross, hires. In (12), ν parameterizes the level of the adjustment cost function.

The function $\omega(E, H)$ represents total compensation to workers as a function of the number of workers and the average number of hours per worker. This compensation function is critical for generating movements in both hours and the number of workers.¹⁶ For our

¹³For example, we have not added stochastic adjustment costs since these would drive an immediate wedge between employment changes and any gap measure.

¹⁴Other inputs into the production function, such as capital and energy, are assumed, for simplicity, to be flexible. Maximization over these factors is thus subsumed by $R(A, E, H)$, and variations in inputs costs are part of A .

¹⁵The literature uses both a quadratic specification in which the cost is in terms of percent differences (Bils [1987]) and specifications in which adjustment costs are in terms of employment changes alone (Hamermesh [1989]).

¹⁶A simpler model with a production function, a fixed wage rate and an employment adjustment cost is not

analysis, we follow Bilts [1987] and Shapiro [1986] and assume that

$$\omega(E, H) = w \times E \times [w_0 + H + w_1 (H - 40) + w_2 (H - 40)^2] \quad (14)$$

where w is the straight-time wage.¹⁷

Using the reduced-form profit function and assuming quadratic costs of adjustment, we solve the dynamic programming problem using value function iteration. Let $E = \psi_E(A, E_{-1})$ be the policy function for employment. Employment is determined by a stochastic difference equation from the policy function.¹⁸ Let $H = \psi_H(A, E_{-1})$ be the policy function for hours.

The frictionless target, $E^*(A)$, is the solution to the optimization problem when $\nu = 0$ for one period. For this model, the frictionless target is equivalent to the solution to $E = \psi_E(A, E)$. The adjustment process, defined by iterations of $E = \psi_E(A, E_{-1})$ given A , converges to the frictionless target, $E^*(A)$. The frictionless hours target is denoted by $H^*(A) = \psi_H(E^*(A), A)$ and will generally depend on A .

The static target, used by CE, is defined as the solution to (12) when $\nu = 0$ in **all** periods. Thus employment and hours satisfy static first order conditions.

The top two panels of Figure 1 illustrate the policy functions and employment targets for two realizations of A . Both the frictionless and the static employment targets are indicated in the figure. Since plants take future adjustment costs into account in determining the frictionless target, this target is not as responsive as the static target to changes in the productivity shock. In general, the frictionless target is less than the static target for above average productivity shocks and vice versa for below average shocks.

As a result, the frictionless hours target for a given shock, $H^*(A)$, deviates from the static hours target, as shown in the bottom panel of Figure 1.¹⁹ If the frictionless employment target is below the static employment target for a given shock, then the frictionless hours target is above the static hours target to compensate for the lower level of employment.

B. Partial Adjustment Model

Within this model, one can be much more explicit about the partial adjustment structure and the resulting flat hazard specification. The partial adjustment model is a policy function defined by

$$e = \lambda e^* + (1 - \lambda)e_{-1} \quad (15)$$

sufficient as there is no “penalty” for overworking employees. Thus, as long as there is no cost to adjusting hours, firms will only modify hours in reaction to shocks. There will be no need to adjust employees.

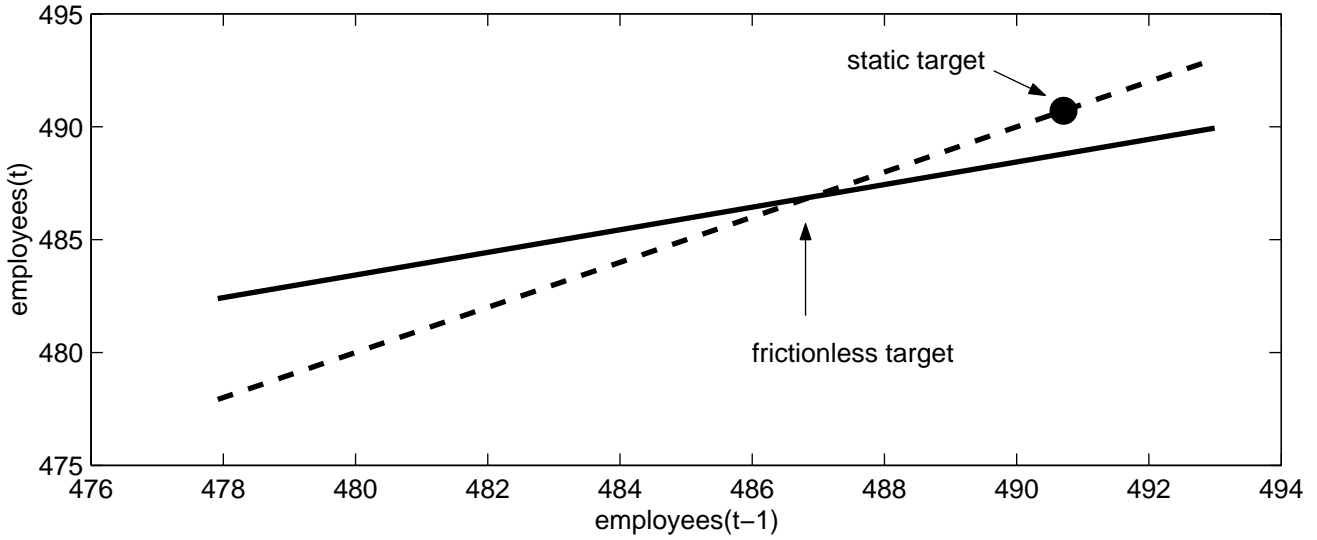
¹⁷In contrast to Sargent [1978] there is no exogenous component to wage variation. In Sargent’s study, variations in productivity were much larger than variations in wages. Further we follow CE and consider a wage function rather than a model with overtime hours as in Sargent.

¹⁸See Sargent [1978] for a further discussion of this problem and the solution methodology for finding the path of employment adjustment.


¹⁹See the appendix for a discussion of the static hours target. It is determined from the first-order condition for hours **if** employment is set at its static target. As discussed in the appendix, the static hours target is not state dependent.

Figure 1

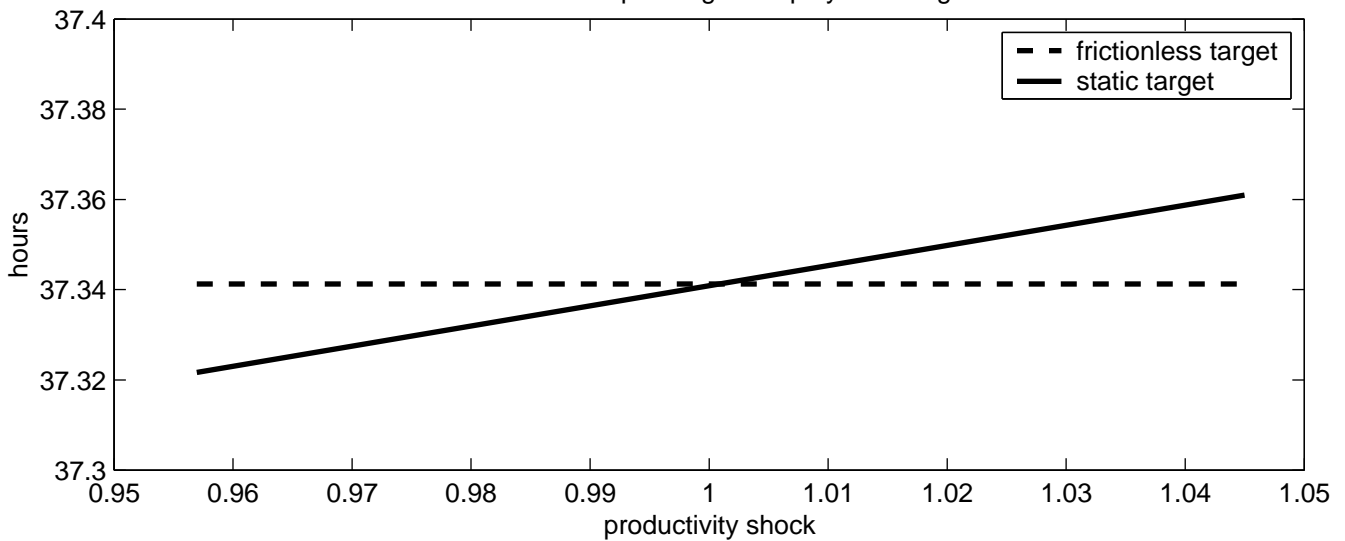
Employment policy function conditional on high productivity shock



Employment policy function conditional on low productivity shock



Hours worked corresponding to employment target levels



for $\lambda \in [0, 1]$, where e represents the log of the stock of current workers. The dependence of e on A comes from the specification of the log of the target, e^* .²⁰ If the optimal policy has this form, then the flat hazard implication is immediate:

$$\Phi(z) = \frac{e - e_{-1}}{e^* - e_{-1}} = \lambda.$$

But what is (15) a solution to? When does it solve (12)?

The standard partial adjustment structure is often “rationalized” by solving for the optimal transition path towards the target in the presence of quadratic adjustment costs and a quadratic loss function.²¹ Consider a dynamic programming problem given by

$$\mathcal{L}(e^*, e_{-1}) = \min_e \frac{(e - e^*)^2}{2} + \frac{\kappa}{2}(e - e_{-1})^2 + \beta E_{e^*|e^*} \mathcal{L}(e^*, e) \quad (16)$$

where the loss depends on the gap between the log-level of workers (e) and the log-level of the target (e^*). Here there is no model of the target; it is taken as an exogenous process. Assume that e^* follows an AR(1) process with serial correlation of ρ . From with this quadratic specification, it is straightforward to show that the optimal policy is linear in the state variables:

$$e = \lambda_1 e^* + \lambda_2 e_{-1}.$$

If shocks follow a random walk ($\rho = 1$), then partial adjustment is optimal ($\lambda_1 + \lambda_2 = 1$).²²

The optimal policy may not take the partial adjustment form for two reasons. First, (16) is an approximation of (12). Second, shocks may not follow a random walk. If ρ is less than one, the value of λ_1 is lower, implying that $\lambda_1 + \lambda_2$ will be less than 1.

4. Empirical Approach

Our goal is to consider the empirical implications of the quadratic adjustment cost model. To do so, we use our model to directly measure the employment gap at the plant

²⁰Clearly e^* ought to be the frictionless rather than the static target since adjustment will stop for a dynamically optimizing plant once that target is reached.

²¹Alternatively, consider a dynamic optimization framework, such as (12), and assume that the within-period return function is a quadratic function and that shocks follow a random walk. Then, the optimal employment level is a linear function of the static optimum and the lagged level of employment. This can be seen directly, for example, from the first-order conditions provided in Sargent [1978] in the linear quadratic framework.

²²Essentially, guess that the policy function is linear in the state variables and use that to solve the first-order condition from the dynamic programming problem. The solution has

$$\lambda_1 = \frac{1 + \beta\kappa\lambda_1\rho}{1 + \kappa - \beta\kappa(\lambda_2 - 1)}$$

and

$$\lambda_2 = \frac{\kappa}{1 + \kappa - \beta\kappa(\lambda_2 - 1)}.$$

level. We call this the **observed gap**. Corresponding to the **frictionless** and **static** targets are two measures of the observed gap: the **frictionless gap** and the **static gap**. We measure these directly using our model as a data-generating mechanism.

Of course, neither CE nor CEH directly observe these gaps. Thus, we follow CEH and infer the employment gap from observed hours variations, using (4) where θ is estimated from (6). We term this the **CEH gap**. Following CEH, we provide two measures of this gap based upon two estimates of θ . The first uses the full simulated panel, and the second uses a subsample comprised of observations entailing large changes in employment and hours, where large changes are defined as those greater than one standard deviation. Similarly, we use the CE procedure of estimating a hazard function from (9) and (10) with time-series data produced by our model.

To solve the dynamic programming problem given in (12), we need to specify functional forms and calibrate the parameters. We assume the following:

- The production function is Cobb-Douglas, where hours and workers are perfectly substitutable. Labor’s share is 0.65, and the markup is set at 25%.
- The compensation function uses the estimates of Bils [1987] and Shapiro [1986]: $\{w_0, w_1, w_2\} = \{1.5, 0.19, 0.03\}$ and the straight time wage, w , is normalized to 0.05. The elasticity of the marginal wage with respect to hours is close to 2 on average.
- We consider two values of the adjustment cost parameter, $\nu = 1$ and $\nu = 10$, so that the half-life of a gap is between 1 quarter and 1 year.²³
- We assume that the profitability shock consists of two multiplicative exogenous components, an aggregate shock (A_{agg}) and an idiosyncratic shock (A_{idio}), such that the profitability shock to plant j in period t is given by $A_{j,t} = A_{agg,t} \times A_{idio,j,t}$. We also assume that both exogenous components follow log-normal AR(1) processes of the following form:

$$\log A_{agg,t} = \rho_{agg} \log A_{agg,t-1} + \epsilon_{agg,t}$$

$$\log A_{idio,j,t} = \rho_{idio} \log A_{idio,j,t-1} + \epsilon_{idio,j,t}.$$

- In our benchmark case, we assume that the idiosyncratic and aggregate processes have the same serial correlation and innovation properties, $\rho_{agg} = \rho_{idio} = 0.95$ and $\sigma_{\epsilon_{agg}} = \sigma_{\epsilon_{idio}} = 0.007$.

Of the components of the basic parameterization, the last assumption is the most controversial for the evaluation of the CE and CEH results.²⁴ As noted above, the quadratic-loss partial adjustment model requires random-walk shocks. Further, both CE (explicitly) and CEH (implicitly) **assume** that shocks follow a random walk.

²³We are grateful to Dan Hamermesh for suggestions on this parameterization.

²⁴We are grateful to the referee for highlighting this important point.

In a complete empirical exercise, the stochastic process for the shocks would be jointly estimated with the adjustment cost process.²⁵ However, the methodologies of CE and CEH do not include estimation of the driving processes. Therefore, outside evidence must be used to calibrate the parameters $(\rho_{agg}, \rho_{idio}, \sigma_{\epsilon_{agg}}, \sigma_{\epsilon_{idio}})$.

Several studies of aggregate shocks suggest that this process is stationary. Sargent [1978] estimates the driving process for aggregate shocks in a quadratic labor adjustment cost model. He does not find a unit root. Hansen [1997] compares moments from a stochastic growth model to U.S. data under alternative assumptions for the serial correlation of the technology shock. He concludes that a model with random-walk shocks does a poorer job of matching observed business cycle features than a model with stationary shocks. His results indicate that a model with a technology shock autocorrelation between 0.9 and 0.99 would best fit the data. Additionally, the chosen parameterization closely approximates the serial correlation and the standard deviation of observed employment.²⁶ Clearly though, exploring the robustness of our results to alternative representations of the shocks is important.²⁷

There are only a few studies on the serial correlation of plant-level shocks, and none of them quite fits our framework. Using a stationary model of technology shocks, Olley and Pakes [1996] estimate production functions in their study of productivity in the telecommunications industry. However, this is a single industry, and they do not report the estimated parameters of the stochastic process. In an earlier version of this paper, we used the results of Cooper and Haltiwanger [2000], in which the profitability shocks are represented by a first-order Markov process and are decomposed into aggregate and idiosyncratic components as a baseline. Neither shock followed a random walk. However, Cooper-Haltiwanger obtain these estimates from a model in which there were, by assumption, no adjustment costs to labor. Thus, that representation of the shocks is not appropriate for this study.

Finally, there is indirect evidence on the sources of fluctuations from variations in job creation and job destruction. Davis et al. [1996, p. 18] say, “most of the job creation and destruction captured by the quarterly figures reflects plant-level employment adjustments that are reversed within a year.”²⁸ Within the context of the CEH study, note that the job creation and job destruction data are not seasonally adjusted. Further, fluctuations are largely driven by oil price shocks. These fluctuations in employment are usually not thought of as being the consequence of random walk disturbances.

Given this parameterization of the basic functions, the optimization problem given in (12) is solved using value function iteration to obtain policy functions. The state space of employment is discretized into a fine grid with 250 points in the relevant portion of the

²⁵See Sargent [1978] on this and the identification problems in distinguishing adjustment costs from the serial correlation of shocks.

²⁶The employment data represent BLS data on manufacturing production workers from 1972 to 1986, which correspond to one of the samples used by CE. For the (log) employment series, the estimated serial correlation is 0.96 and the standard deviation is 0.06.

²⁷See section 7.C.

²⁸Also see their Table 2.3.

state space. For the given values of the serial correlation and standard deviation for both the aggregate and idiosyncratic shocks, we use the procedure outlined in Tauchen [1986] to create a discrete state space representation for the shocks.²⁹ Using these policy functions, we create a simulated panel data set where the number of plants equals 1000 and the number of time periods is 1000.³⁰

5. Aggregate Implications

Given that both CE and CEH present quantitative results on the estimation of hazard functions from time series data, we begin by analyzing the aggregate implications of the quadratic adjustment model. We create a time series by aggregating across the plants in our simulated panel data set. Following CE and CEH, we investigate aggregate implications by looking at the relationship between aggregate employment changes and the cross-sectional distribution of the employment gap given by (7).

Table 1 presents estimates of (7) for three specifications of a hazard function ($\Phi(z)$): constant, piecewise linear and quadratic.³¹ More precisely, we specify

$$\Phi(z) = \begin{cases} \lambda_0 + \lambda_1^- z + \lambda_2 z^2 & \text{for } z < 0 \\ \lambda_0 + \lambda_1^+ z + \lambda_2 z^2 & \text{for } z > 0 \end{cases} \quad (17)$$

which nests the different specifications of the aggregate hazard function used in CE and CEH. As described in those papers, (17) is substituted into the aggregate growth equation (7) yielding the following equation in which aggregate employment growth depends on the parameters of (17) and the moments of the cross-sectional distribution of the gaps:

$$\Delta E_t = const + \lambda_0 m_{1,t} - \lambda_1^- F_t(0) m_{2^-,t} + \lambda_1^+ (1 - F_t(0)) m_{2^+,t} + \lambda_2 m_{3,t} + \varepsilon_t \quad (18)$$

where $m_{i,t}$ is the i^{th} uncentered moment of the cross-sectional distribution of the gap in period t , an index of $+$ ($-$) indicates that the moment only includes observations with a positive (negative) gap, and $F_t(0)$ represents the fraction of plants with a negative gap at time t .

²⁹The aggregate and idiosyncratic profitability shocks are each represented by an 11-point state space equally divided between two standard deviations of their respective means. Given the mean, standard deviation, and serial correlation parameters, an 11×11 transition matrix is created for each process. In a simulated sample of 1000 plants over 1000 periods, the serial correlation and variance properties of these generated series are very close to the parameters used to designate the state space and transition matrix. Adding points to the state space does not meaningfully change the results.

³⁰CEH have a panel with 36 quarters and 10,000 plants. Our results are robust to adding more plants. We analyze only 1000 plants to reduce computation time. The number of time periods is set at 1000 to minimize simulation error.

³¹Note that this hazard function is imposed on the aggregate data which themselves come from a panel created by the optimal decisions at the plant level. These optimal decisions will not necessarily obey any of these simple hazard specifications.

A. Frictionless Target

The results for the estimation of (18) with the frictionless target computed using the observed gap are reported in Table 1a. When the appropriate target is used, the results are consistent with intuition: the estimated hazard is flat with an adjustment rate that is nearly 0.48 when $\nu = 1$ and 0.17 when the adjustment cost is larger, $\nu = 10$. There is essentially no evidence of any economically significant nonlinearity.³² The R^2 for this specification is virtually 1: the model with a constant hazard fits quite well.³³ So, even though our driving process is not a random walk, the flat hazard prediction of the quadratic adjustment cost model seems to work well **using the observed frictionless gap**.

There are three deviations from this benchmark associated with three potential “errors” in measuring the gap. First, as in CE, the static target, which is easy to compute in our simulated environment, may be used instead of the frictionless target. Second, the estimation procedure used by CE relies on an artificial measure of the static target. The third is the CEH measure of the gap. For each of these measures, we consider the specifications of (18).

B. Static Target

Table 1a also shows the results obtained when (18) is estimated using the **observed static gap**. Using this measure, one would strongly **reject** the hypothesis that the hazard function is flat in favor of either the piecewise linear or the quadratic case. For example, in the quadratic specification, we find that when $\nu = 1$, λ_2 is estimated at 0.62 with a standard error of 0.26. Further, the coefficients in the piecewise linear specification ($\lambda_1^+ = 0.1, \lambda_1^- = 0.12$) are also statistically and economically significant. Note, though, that here the R^2 for the constant hazard model is essentially 1, so that adding these higher moments of the cross-sectional distribution, though they are significant, does not lead to increases in R^2 . The nonlinearity is **not** statistically significant when $\nu = 10$ for either the piecewise linear or the quadratic cases.³⁴

The table also includes the quadratic specification given in (18), where $\lambda_1^+ = \lambda_1^-$ has been imposed. The mapping between these estimates $\{\lambda_1, \lambda_2\}$ and those reported in CE $\{\tilde{\lambda}_0, \tilde{\lambda}_2\}$ is given by $\lambda_0 = \tilde{\lambda}_0 + \tilde{\lambda}_2 z_0^2$, $\lambda_1 = -2z_0 \tilde{\lambda}_2$ and $\lambda_2 = \tilde{\lambda}_2$.

In the $\nu = 1$ case, the estimate of λ_2 equals 0.61 and is significantly different from zero. In fact, this estimate of the nonlinearity in the hazard is not far from the estimate of $\tilde{\lambda}_2$, 0.53, reported by CE. The constant terms (λ_0) in the hazard functions are quite close as well: the CE specification yields a constant term of 0.38 while we report a constant of 0.47. However, CE find z_0 equal to -0.82 while our estimate of z_0 is approximately 0.

The difference in results between using the frictionless and static targets to determine

³²Though the regression coefficients on some of the nonlinear pieces are statistically significant, they are not economically significant in the observed distribution of the gaps, $z \in [-0.5, 0.5]$.

³³This high value of R^2 partly reflects the limited nature of the model: there are no other factors of production with adjustment costs, no shocks to the adjustment costs directly, no measurement error, etc.

³⁴As we shall see in section 7.C, the nonlinearity can be present for $\nu = 10$ for alternative parameterizations of the shocks.

the employment gap can be viewed as the introduction of measurement error into the regression. If the static target is equal to the frictionless target, we should not see any change in results. Figure 1, however, illustrates the difference between the two targets. Switching to the static target is likely to lead to a bias in the estimate, as there is not a constant difference between these targets.

Using the hazard given in (17) with the restrictions to produce the CE quadratic hazard, one can rewrite the aggregate employment growth equation, (18), as follows:

$$\Delta E_t = \lambda_0 m_{1,t}^s + \lambda_1 m_{2,t}^s + \lambda_2 m_{3,t}^s + \overbrace{\varepsilon_t + \lambda_0 \left(m_{1,t}^f - m_{1,t}^s \right) + \lambda_1 \left(m_{2,t}^f - m_{2,t}^s \right) + \lambda_2 \left(m_{3,t}^f - m_{3,t}^s \right)}^{\text{error term}}. \quad (19)$$

The error term contains three mismeasurement terms in addition to ε_t . If any of these measurement errors is correlated with the moments of the static employment gap, then a bias in the estimates will be present.

To study this bias, we regress the measurement error in the first uncentered moment, $m_{1,t}^f - m_{1,t}^s$, on the three moments of the static gap $\{m_{1,t}^s, m_{2,t}^s, m_{3,t}^s\}$ using data from the $\nu = 1$ case.³⁵ We estimate $\{-0.02, 0.03, 1.26\}$ as the coefficients on the three moments with standard errors of $\{0.003, 0.03, 0.5\}$. These results indicate that the error in (18) is related to the static gap in a nonlinear way, thus leading to the nonlinear estimates of the adjustment function.³⁶

C. CE Measure of the Static Gap

The results presented in the above subsection assume that the static target is observed. Of course, CE do not observe this and must instead infer the growth of the employment target using (9). We could use our simulated data to create this measure and then estimate the hazard function in (10). The question would be whether the CE procedure introduces nonlinearities in addition to those reported in Table 1a.

In fact, we can address this point without replicating the entire CE methodology. From the simulated data, we directly create an aggregate measure of growth in the static employment target. Using aggregate employment growth and hours growth from the same simulated data set, we also use the CE methodology, as in (9), to create their measure of the employment target using CE's assumed value of $\theta = 5$. In the simulated data, the correlation of these two measures of the growth in the employment targets was 0.9996. Therefore, the CE procedure of looking at the time series does not, it appears, produce additional nonlinearities since the only input into (10) is the growth rate of the employment target. The key to the nonlinearity seems to be the substitution of the static for the frictionless gap as indicated in

³⁵Thanks to Peter Klenow for discussions on this characterization of the effects of the measurement error.

³⁶Regressions of the higher-order measurement error terms on the three moments of the static gap yield qualitatively similar results.

Table 1a.

D. CEH Measure of the Gap

Alternatively, the frictionless target could be inferred from variations in observed hours at the plant-level, as in CEH, opening the possibility of additional measurement error. The results for this case are in Table 1b. The different sections refer to alternative treatments of the data. “Full sample” means that we use the complete sample while “big change” refers to a sample constructed by including only observations in which the employment and hours changes exceed one standard deviation, as in the sample splits of CEH.

For both specifications ($\nu \in \{1, 10\}$), the flat hazard specification yields rather non-sensical results: the adjustment rate is in excess of 100% for both samples. The constant hazard hypothesis is not rejected for both full samples but would be rejected for the big change samples. Evidently, the sample selection of CEH creates nonlinearity in the quadratic adjustment cost case. The source of the misspecification in this case is discussed in the next section.

As with the static target, we can study the correlation between the measurement error and the moments of the cross-sectional distribution of the employment gaps when the CEH procedure is used to measure those gaps. This is an alternative version of (19) where m_i^s is the i^{th} moment of the gap distribution derived using the CEH procedure. We again regress the measurement error in the first uncentered moment on the three moments of the CEH gap using data from the $\nu = 1$ case. We estimate $\{0.001, 0.01, 4.19\}$ as the coefficients on the three moments with standard errors of $\{0.0001, 0.007, 1.52\}$.³⁷ These results again indicate that the error in (18) is related to the CEH gap in a nonlinear way, thus leading to the nonlinear estimates of the adjustment function.³⁸

E. Summary

Thus, from the aggregate estimation results, we find that the hazard function is essentially flat *iff* the gap is properly measured. When the static target is used to construct the gap, we find that the flat hazard model may be rejected. Using either the CE or the CEH procedure for measuring the gap, one would reject the flat hazard specification and conclude that adjustment costs were not quadratic. Here we have seen that this conclusion is not valid: the measurement of the gap, not economic behavior, introduces nonlinearities.³⁹

6. Determination of θ

A key element in the CE and CEH procedures is the calibration/estimation of θ . This is the main link between the observable variable (hours) and the unobservable variable (the

³⁷These are the actual coefficients divided by 1000. We discuss below why the coefficients from the CEH approach are so large.

³⁸As with our discussion of the static target, regressions of the higher order measurement error terms on the three moments of the static gap yield qualitatively similar results.

³⁹An earlier version of this paper included evidence on the plant-level hazards. We found the presence of nonlinearities when the gap was not measured properly there as well.

gap). Once this parameter is determined, CE calculate the aggregate targets from (9) and CEH use (4) to construct the plant-level gaps. The logic in both cases is to infer movements in the gap from variations in hours. Accordingly, the final step in our evaluation of the gap methodology is to explore the estimates of θ using these two procedures.

Table 2 summarizes the estimates of this parameter for a number of different specifications. The first two rows correspond to the value of θ estimated from (3) using the actual gap that we construct in our simulated environment. Of these rows, the first measure uses the frictionless target to create the gap, and the second measure uses the static target. The other rows use the CEH approach to estimate θ . Note that the CEH results do not depend on the definition of the target since it is not observable. Here, as in Table 1, results are reported for the two parameterizations of the quadratic adjustment cost model, $\nu = 1$ and $\nu = 10$.

A. CE approach

To evaluate the CE observed gap approach, consider the top part of Table 2. The estimate of θ obtained by using the frictionless and static gap measures differ. In fact, using the static target, as in CE, produces an estimate of θ that is exactly equal to the one obtained analytically from the firm's optimal choice of hours worked per employee.⁴⁰ However, the gap measure produced by using this estimate of θ **does not** correspond with the relevant measure for a dynamically optimizing firm, the frictionless gap. The difference is due to the dependence of the frictionless hours target on the productivity shock. As shown in Figure 1, the static hours target is independent of the shock.

This distinction between the two hours targets has important implications for the measurement of the gaps. From a log-linearization of the first-order conditions from (12), the relationship between the employment gap and hours deviations can be written as

$$\tilde{z}_{i,t} = \theta (h_{i,t} - h^*(A_{i,t})). \quad (20)$$

Using the correct target for hours and the frictionless employment gap, we do obtain the analytically calculated value of θ .

The problem for the CE methodology is that there are two unobservables in (20). The hours target cannot be approximated by a constant mean, as was assumed in the construction of (3). Even if an estimate for θ is available, the employment gap cannot be accurately constructed without observing the hours target. The errors caused by having the correct θ and using the mean level of hours to approximate the hours target are illustrated precisely by the observed static target results above for the aggregate hazards.

⁴⁰Using (26) from the appendix and the given parameterization, CE would find θ is equal to 8.8.

B. CEH approach

To evaluate the CEH approach, recall their regression equation, (6):

$$\Delta e_{i,t} = \delta - \theta \Delta h_{i,t} + \varepsilon_{i,t}.$$

From Table 2, the sign of the estimated value of θ from (6) is opposite that obtained when the observed gap is used in the regression, as in (3). Since the CEH methodology relies on θ to construct a measure of the gap, this difference is important to understand.


The error term in (6) contains the change in the employment target level. If changes in hours are uncorrelated with changes in employment targets, the sign on θ will be determined by the unconditional correlation between changes in hours and changes in employment. In the simulated data for $\nu = 1$, this latter correlation is 0.51. The driving force behind this positive correlation is the partial adjustment to changes in employment targets. When plants experience productivity shocks, they respond to changes in employment targets by changing both hours and employment in the same direction. This positive correlation between hours and employment implies the negative sign on θ in (6), as reported in Table 2.

But there is no rationale for the assumption that changes in hours are uncorrelated with changes in employment targets. Because CEH acknowledge that hours and employment target changes could be correlated, they use only observations in which there are large changes in both hours and employment to estimate θ . They argue that in these periods, the changes in employment targets are swamped by the effects of large changes in hours and employment. But in a model of convex adjustment costs, the only periods in which there are large changes in hours and employment are periods in which there are large changes in employment target levels.


This is evident in the simulated data: the correlation between changes in hours and changes in employment target levels is 0.88 in the full sample and 0.96 in the CEH-criterion subsample. Therefore, the CEH methodology produces a biased estimate of θ . To obtain an unbiased estimate of θ in a model of quadratic costs of adjustment, it is essential to control for changes in employment target levels.

The implications of the sign reversal are displayed in Figure 2, which shows a sample of employment changes, deviations in hours, and various measures of the employment gap from a simulation of the model. The upper panel displays the two measures of the actual gap, and the lower panel displays two measures of the gap constructed from CEH estimates of θ . The differences between the gap measures are readily apparent once the scales of the two panels are taken into account. The series for employment changes and hours deviation are identical in the two panels. In the upper panel, the gap measures have a higher degree of variability than employment changes, indicative of the expected plant behavior of partial adjustment when faced with convex costs of adjustment. In the lower panel, employment changes greatly exceed the CEH gap measures; hence the large parameter estimates in Table 1b. Since hours and employment are positively correlated, the negative sign on θ causes the constructed

Figure 2
Simulated employment gaps: measured directly
($v = 1$)



Simulated employment gaps: measured using CEH methodology



employment gap to be a dampened version of the change in employment. The actual measures of the gap and the CEH gap measures are positively correlated (approximately 0.77 for the big change subsample at $\nu = 1$), but the conclusions to be drawn from analysis of these series are very different.

7. Robustness

The conclusions we have reached concerning the inferences from the gap methodology are, admittedly, based upon the selection of parameters for the plant-level optimization problem and for the driving processes. It is natural to explore the robustness of these findings.

A. Specification of Optimization Problem

With regard to the specification of the plant level optimization problem, we consider two variations. First, our production function assumes that the labor input is the product of hours and the number of employees. Yet, CE, citing Bills [1987], analyze a model in which

$$R(A, e, h) = A(e^{\alpha_e} h^{\alpha_h}) \quad (21)$$

with $\alpha_e = 0.72$, $\alpha_h = 0.77$.⁴¹ In this case, our conclusions on the methods of CE and CEH do not change: nonlinearities remain in the aggregate regressions.

Second, as noted earlier, the literature is somewhat mixed on the specification of the quadratic adjustment cost model. In our model, we assume that the cost depends on the rate of change in employment, not the change alone. Instead we could consider

$$\frac{\nu}{2}(e - e_{-1})^2. \quad (22)$$

Using this specification of the adjustment cost function does not have a significant effect on our conclusions: nonlinearities remain in the aggregate regressions.

B. Variability of Hours

The relative variability of hours and employment is largely determined by the elasticity of compensation with respect to hours and the employment adjustment costs. In our baseline parameterization the standard deviation of (the log of) hours is about 0.0018 while the standard deviation of (the log of) employment is 0.07. In contrast, from the BLS data set that we used to measure the serial correlation and standard deviation of employment, the standard deviation of (log) hours is 0.013 while the standard deviation of (log) employment is 0.06. Clearly our model exhibits too little variability in hours. This is potentially important for the results as both the CE and CEH procedures infer movements in targets from movements in hours. So, from (9), low variability in hours is comparable to a value of θ near 0, in that the growth in the target is just the observed value of employment growth so that the estimation

⁴¹The values for α_e and α_h are produced by assuming constant returns to scale in capital and employment and a markup of 25% and using the production relationship between hours and employment reported in CE.

of (8) is less interesting.⁴²

We explore the robustness of our results to a parameterization in which the elasticity of the marginal wage is much lower. Specifically, we replace the compensation schedule in (14) with

$$\omega(E, H) = E(w_0 + w_1 H^\zeta) \quad (23)$$

so that the elasticity of the marginal wage schedule is $(\zeta - 1)$.⁴³ With $\zeta = 1.1$ we are able to reproduce the BLS standard deviations of hours and employment.⁴⁴ At this elasticity, our results remain: the estimated hazard function using the static gap measure is nonlinear.

C. Shocks

To explore the sensitivity of the results to the idiosyncratic and aggregate shock processes, we compute results for the aggregate regressions over a range of plausible alternative parameter settings. We focus on the estimation of (18) with the restriction that $\lambda_1^+ = \lambda_1^- = 0$, which allows us to focus on the quadratic hazard specification. Tables 3–5 provide estimates of λ_2 from (18) for various parameterizations of the aggregate and idiosyncratic shock processes. For these results, the gap is constructed using the **observed static target**.

Table 3 displays estimates of λ_2 from (18) assuming there is no difference between the parameters of the idiosyncratic and aggregate shocks: $\rho_{agg} = \rho_{idio}$ and $\sigma_{\epsilon_{agg}} = \sigma_{\epsilon_{idio}}$ throughout. There are numerous parameter values, including both values of ν , where there is significant nonlinearity (λ_2 significantly different from zero) in the regressions using the static target. Interestingly, when the process is closest to a random walk ($\rho = 0.99$), there is significant nonlinearity for both target measures.

To distinguish the source of the nonlinearities, Tables 4 and 5 provide results from additional simulations, in which we relax the restriction that the aggregate and idiosyncratic shocks are from the same stochastic process. In Table 4, results are presented for values of ρ_{agg} between 0.6 and 0.99 and for values of $\sigma_{\epsilon_{agg}}$ between 0.02 and 0.3. For both values of ν , there is evidence of nonlinearity, particularly for low levels of $\sigma_{\epsilon_{agg}}$. There is significant nonlinearity even when the aggregate process is closest to being a random walk ($\rho_{agg} = 0.99$). For all of these specifications, the aggregate time series of employment exhibits serial correlation and volatility in the neighborhood of the observed employment process.

Table 5 explores the relationship between aggregate nonlinearity and the parameteri-

⁴²As CE note, this is why one cannot estimate θ in their model.

⁴³There seems to be some confusion between Bils and CE over the value of the marginal wage elasticity. Bils (p. 848) finds that an increase in weekly hours from 40 to 41 is associated with a 4.6 percent increase in the marginal wage. This translates into a marginal wage elasticity of 1.84. CE (p. 372) specify a wage function “where $(\mu - 1)$ is the elasticity of the marginal wage schedule with respect to average hours worked.” They then set μ , the *wage elasticity*, equal to Bils’ estimate of the *marginal wage elasticity*. Hence, the marginal wage elasticity used by CE (0.9) is lower than the estimate by Bils (1.84). Using Bils’ estimate, the calibrated value of θ would be approximately 10 instead of the value of 5 used by CE.

⁴⁴Our marginal wage elasticity differs from Bils’ estimate in part due to the fact that we are using the null quadratic model that admittedly is not the true model for the data. A marginal wage elasticity lower than Bils’ estimate is necessary to generate sufficient volatility in hours.

zation of the idiosyncratic shock. As indicated in Table 5a, there is again a range of values of ρ_{idio} and $\sigma_{\epsilon_{idio}}$ where the aggregate nonlinearity is present when $\nu = 1$. There are no such values when $\nu = 10$. Table 5b indicates that this finding is not general: there are certainly values of ρ_{idio} and $\sigma_{\epsilon_{idio}}$ for which λ_2 is significant even with $\nu = 10$.

In sum, this section displays the sensitivity of results to the parameterization of the shocks. The results support the view that the parameters of the process governing the shocks ought to be estimated along with the adjustment parameters. Assuming that shocks follow a random walk, as in CE and CEH, not only is inconsistent with (the meager) existing evidence on plant-level shocks but also may weaken the validity of inferences about the nature of adjustment costs and their significance for aggregate time series.

8. Conclusions

The point of this paper is to assess the findings of CE and CEH that aggregate employment dynamics depend upon the cross-sectional distribution of employment gaps. We argue that due to measurement problems, a researcher might indeed find that the cross-sectional distribution matters for aggregate time series even if adjustment costs are quadratic. Thus, the conclusion of CE and CEH that nonlinear adjustment at the plant-level is present in aggregate time series is not based on convincing evidence. So, despite the overwhelming evidence that plant-level adjustment is nonlinear, the question of whether this matters for aggregate employment dynamics remains an open issue.

Can we do better? Within the gap methodology, it is apparent that the CEH methodology is inferior to that employed by CE.⁴⁵ However, even the CE approach falls short, due primarily to state-contingent differences between the frictionless and static employment targets. We have seen that adding a state-dependent hours target to the model yields the appropriate frictionless target, though implementing this procedure with actual data is less clear.

There are competing approaches to estimating a parameterized version of an adjustment cost function nesting both convex and nonconvex costs that do not rely on gap measures. Examples of this, which now exist in the literature on investment, durables and price-setting, involve using indirect inference techniques to match the moments produced by simulations of a structural model with those observed.⁴⁶ Clearly, labor is next.

⁴⁵We understand that data limitations led CEH to their formulation.

⁴⁶We have tried without success to use that approach to match the aggregate regression results reported in CEH. An alternative is to use VARs, following Sargent [1978], and also to structure the estimation around plant-level reduced form regressions as in Cooper-Haltiwanger [2000]. This is in process.

9. References

- Bils, M., "The Cyclical Behavior of Marginal Cost and Price," *American Economic Review*, 77 (1987), 838–55.
- Caballero, R. and E. Engel, "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics*, 108 (1993), 359–83.
- Caballero, R., E. Engel and J. Haltiwanger, "Aggregate Employment Dynamics: Building From Microeconomic Evidence," *American Economic Review*, 87 (1997), 115–37.
- Cooper, R. and J. Haltiwanger, "On the Nature of Capital Adjustment Costs," NBER Working Paper # 7925, September 2000.
- Davis, S. and J. Haltiwanger, "Gross Job Creation, Gross Job Destruction and Employment Reallocation," *Quarterly Journal of Economics*, 107 (1992), 819–63.
- Davis, S., J. Haltiwanger and S. Schuh, *Job Creation and Destruction*, Cambridge Mass.: MIT Press, 1996.
- Hamermesh, D., "Labor Demand and the Structure of Adjustment Costs," *American Economic Review*, 79 (1989), 674–89.
- Hansen, G., "Technical Progress and Aggregate Fluctuations," *Journal of Economic Dynamics and Control*, 21 (1997), 1005–23.
- Nickell, S., "Fixed Costs, Employment and Labour Demand over the Cycle," *Economica*, 45 (1978), 329–45.
- Olley, S. and A. Pakes, "The Dynamics of Productivity in the Telecommunications Equipment Industry," *Econometrica*, 64 (1996), 1263–97.
- Sargent, T., "Estimation of Dynamic Labor Demand Schedules under Rational Expectations," *Journal of Political Economy*, 86 (1978), 1009–44.
- Shapiro, M., "The Dynamic Demand for Capital and Labor," *Quarterly Journal of Economics*, 101 (1986), 513–42.
- Tauchen, G., "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economics Letters*, 20 (1986), 177–81.

Appendix: CE Gap and Target

CE use the following basic framework to generate some analytic results on θ , the parameter that connects variations in hours with variations in the number of workers and the target employment level. Their approach is completely static. They maximize

$$R(A, E, H) - \omega(E, H)$$

for the optimal choice of hours given (A, E) , where A represents the profitability of a production unit, E is the level of employment, and H is the input of hours by a worker. This maximization yields a first order condition of

$$R_H(A, E, H) = \omega_H(E, H). \quad (24)$$

The level of hours satisfying this first order condition is also appropriate in a dynamic setting since the hours choice entails no costs of adjustment. Similarly, CE optimize over the number of workers setting hours at \bar{H} , implying that

$$R_E(A, E, \bar{H}) = \omega_E(E, \bar{H}). \quad (25)$$

This first-order condition is intended to characterize a target level of employment as hours are set at their optimal level. We let $E^{**}(A)$ denote the solution to (25). This is the **static target**, and it is, by construction, independent of the specification of the adjustment cost function. Given $E^{**}(A)$ and the specifications above for the compensation and production functions, plants will always choose the same steady-state level of hours per worker, $H^{**}(A) = \bar{H}, \forall A$.

Log-linearizing (25), given the functional forms assumed earlier, yields

$$\hat{A}_t + (\alpha - 1) \hat{E}_t + \alpha \hat{H}_t = \frac{w'(\bar{H}) \bar{H}}{w(\bar{H})} \hat{H}_t$$

where “ \hat{X}_t ” is the percent deviation from steady state in period t . Since the static target for hours is independent of deviations in the productivity shock, we can express the relationship between the static employment target and the productivity shock from (25) as

$$\hat{A}_t = (1 - \alpha) \hat{E}_t^{**}.$$

Substitution of this relationship into the log-linearized version of (24) yields

$$(1 - \alpha) \hat{E}_t^{**} + (\alpha - 1) \hat{H}_t + \alpha \hat{E}_t = \hat{E}_t + \xi_w \hat{H}_t$$

where ξ_w is the marginal wage elasticity with respect to hours.⁴⁷ This can be rewritten as

$$\hat{E}_t^{**} - \hat{E}_t = \frac{1 - \alpha + \xi_w}{1 - \alpha} \hat{H}_t. \quad (26)$$

Using the mean level of observed hours as an approximation for \bar{H} , equation (3) denotes the same relation as (26) with θ equal to $\frac{1-\alpha+\xi_w}{1-\alpha}$.⁴⁸

Relative to the parameterization of our model, CE would set $\theta = 8.8$ using the following analysis. The marginal wage elasticity is evaluated at the static steady state level of 37.3 hours. From this, $\xi_w = 2.19$.

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

$$\tilde{R}(A, E, H, K) = \left(\tilde{A} (EH)^{\alpha_L} K^{\alpha_K} \right)^{\frac{\eta-1}{\eta}} - rK$$

where α_L 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 (13), where

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

With η set equal to 5, corresponding to a markup of 25%, and assuming constant returns to scale in capital and labor with $\alpha_L = 0.65$, α is equal to 0.72. Using these calculations, θ can be determined from $\theta = \frac{1-\alpha+\xi_w}{1-\alpha}$.

⁴⁷The marginal wage elasticity can be expressed as $\xi_w = \frac{2w_2\bar{H}}{1+w_1+2w_2(\bar{H}-40)}$.

⁴⁸We are grateful to Robert King for pushing us to make this connection.

Table 1a: Aggregate Implications, $\nu = 1$

Observed Gap	λ_0	λ_1^+	λ_1^-	λ_2	R^2
Frictionless	0.48				1.00
	(0.000)				
	0.48	-0.02	0.03		1.00
	(0.000)	(0.004)	(0.004)		
	0.48			0.04	1.00
	(0.000)			(0.026)	
Static	0.47				1.00
	(0.001)				
	0.46	0.10	0.12		1.00
	(0.00)	(0.04)	(0.04)		
	0.48			0.62	1.00
	(0.002)		(0.26)		
	0.47	-0.01	-0.01	0.61	1.00
	(0.002)	(0.014)	(0.014)	(0.23)	

Table 1a: Aggregate Implications, $\nu = 10$

Observed Gap	λ_0	λ_1^+	λ_1^-	λ_2	R^2
Frictionless	0.17				1.00
	(0.000)				
	0.17	0.02	-0.01		1.00
	(0.000)	(0.002)	(0.003)		
	0.17			0.043	1.00
	(0.000)			(0.013)	
Static	0.15				0.97
	(0.001)				
	0.15	0.03	-0.04		0.97
	(0.004)	(0.04)	(0.04)		
	0.15			-0.01	0.97
	(0.002)		(0.19)		
	0.15	0.04	0.04	0.01	0.97
	(0.002)	(0.013)	(0.013)	(0.19)	

Notes: Results from estimation of (18). Standard errors in parentheses.

Table 1b: Aggregate Implications, $\nu = 1$

CEH Gap	λ_0	λ_1^+	λ_1^-	λ_2	R^2
Full sample	2.01				0.99
	(0.01)				
	2.04	-1.29	-2.19		0.99
	(0.02)	(1.46)	(1.34)		
Big change	2.02			-50.47	0.99
	(0.01)			(31.93)	
	5.88				0.76
	(0.001)				
Big change	4.45	310.93	225.04		0.76
	(0.56)	(118)	(94.17)		
	5.28			16562	0.76
	(0.26)			(6703)	

Table 1b: Aggregate Implications, $\nu = 10$

CEH Gap	λ_0	λ_1^+	λ_1^-	λ_2	R^2
Full sample	2.06				0.95
	(0.02)				
	2.12	-12.78	-1.94		0.95
	(0.08)	(10.79)	(12.70)		
Big change	2.11			-788.15	0.95
	(0.04)			(662.9)	
	0.87				0.03
	(0.16)				
Big change	10.34	-2642.79	-3018.89		0.09
	(1.17)	(357.98)	(358.32)		
	5.695			-368956	0.09
	(0.614)			(45481)	

Notes: Results from estimation of (18). Standard errors in parentheses.

Table 2: Estimate of θ

	$\nu = 1$		$\nu = 10$	
	θ	R^2	θ	R^2
Observed Gap				
Frictionless target	8.59 (0.001)	0.99	7.42 (0.002)	0.96
Static target	8.87 (0.000)	1.00	8.87 (0.000)	1.00
CEH Gap				
Full sample	-4.01 (0.007)	0.26	-0.85 (0.002)	0.12
Big change	-6.90 (0.001)	0.72	-1.51 (0.005)	0.44

Notes: Results from estimation of (18). Standard errors in parentheses.

Table 3: Robustness with Identical Idiosyncratic and Aggregate Shock Parameters

Parameters				Estimate of λ_2	
ρ_{idio}	$\sigma_{\varepsilon_{idio}}$	ρ_{agg}	$\sigma_{\varepsilon_{agg}}$	$\nu = 1$	$\nu = 10$
0.70	0.02	0.70	0.02	1.22 (0.38)	0.12 (0.05)
0.80	0.02	0.80	0.02	0.54 (0.28)	0.21 (0.07)
0.90	0.02	0.90	0.02	0.88 (0.41)	0.13 (0.13)
0.95	0.02	0.95	0.02	0.62 (0.23)	-0.01 (0.19)
0.99	0.02	0.99	0.02	8.13 (2.74)	0.87 (0.28)

Note: $\rho_{agg} = \rho_{idio}$ and $\sigma_{\varepsilon_{agg}} = \sigma_{\varepsilon_{idio}}$ throughout this table.

Table 4: Robustness with Varying Aggregate Shock Parameters

Parameters		Estimate of λ_2	
ρ_{agg}	$\sigma_{\varepsilon_{agg}}$	$\nu = 1$	$\nu = 10$
0.60	0.02	0.67 (0.27)	0.24 (0.07)
0.60	0.10	0.05 (0.02)	0.01 (0.004)
0.60	0.30	0.007 (0.002)	-0.0000 (0.0005)
0.75	0.02	0.81 (0.34)	0.24 (0.09)
0.75	0.10	0.06 (0.02)	0.015 (0.01)
0.75	0.30	0.01 (0.002)	0.0002 (0.001)
0.90	0.02	0.88 (0.39)	0.14 (0.13)
0.90	0.10	0.09 (0.03)	0.01 (0.01)
0.90	0.30	0.02 (0.004)	-0.0003 (0.001)
0.95	0.02	0.62 (0.23)	-0.01 (0.19)
0.95	0.10	0.09 (0.03)	-0.003 (0.01)
0.95	0.30	0.03 (0.005)	-0.004 (0.002)
0.99	0.02	5.41 (2.20)	0.13 (0.21)
0.99	0.10	0.22 (0.09)	0.03 (0.02)
0.99	0.30	0.08 (0.02)	0.0001 (0.003)

Notes: $\rho_{idio} = 0.95$ and $\sigma_{\varepsilon_{idio}} = 0.02$. Standard errors in parentheses.

Table 5a: Robustness: With Varying Idiosyncratic Shock Parameters

Parameters		Estimate of λ_2	
ρ_{idio}	$\sigma_{\varepsilon_{idio}}$	$\nu = 1$	$\nu = 10$
0.60	0.02	1.03 (0.44)	0.001 (0.15)
0.60	0.10	0.13 (0.06)	0.09 (0.03)
0.60	0.30	0.01 (0.003)	0.01 (0.002)
0.75	0.02	0.84 (0.36)	-0.003 (0.146)
0.75	0.10	0.12 (0.09)	0.07 (0.03)
0.75	0.30	0.01 (0.004)	0.01 (0.002)
0.90	0.02	0.64 (0.26)	-0.06 (0.14)
0.90	0.10	0.09 (0.15)	-0.08 (0.06)
0.90	0.30	0.01 (0.01)	-0.01 (0.003)
0.95	0.02	0.62 (0.23)	-0.01 (0.19)
0.95	0.10	0.36 (0.22)	0.01 (0.02)
0.95	0.30	0.03 (0.01)	-0.02 (0.01)
0.99	0.02	0.46 (0.21)	0.01 (0.22)
0.99	0.10	1.27 (0.49)	0.06 (0.25)
0.99	0.30	0.22 (0.08)	0.08 (0.02)

Notes: $\rho_{agg} = 0.95$ and $\sigma_{\varepsilon_{agg}} = 0.02$. Standard errors in parentheses.

Table 5b: Robustness with Varying Idiosyncratic Shock, $\nu = 10$

Parameters				Estimate of λ_2
ρ_{idio}	$\sigma_{\varepsilon_{idio}}$	ρ_{agg}	$\sigma_{\varepsilon_{agg}}$	$\nu = 10$
0.60	0.02	0.60	0.02	0.14 (0.05)
0.60	0.02	0.99	0.02	0.71 (0.23)
0.60	0.10	0.99	0.02	0.15 (0.02)
0.75	0.02	0.60	0.02	0.19 (0.06)
0.75	0.02	0.95	0.02	0.49 (0.23)
0.95	0.02	0.60	0.02	0.24 (0.07)
0.95	0.02	0.75	0.02	0.24 (0.09)
0.99	0.02	0.75	0.02	0.28 (0.11)
0.99	0.10	0.75	0.02	0.32 (0.13)

Note: Standard errors in parentheses.