NBER WORKING PAPER SERIES

THE ECONOMICS OF LABOR ADJUSTMENT: MIND THE GAP

Russell Cooper
Jonathan L. Willis

The Economics of Labor Adjustment: Mind the Gap
Russell Cooper and Jonathan L. Willis
NBER Working Paper No. 8527
October 2001
JEL No. E24, J41, J6

ABSTRACT

We study the inferences about labor adjustment costs obtained by the "gap methodology" of Caballero and Engel [1993] and Caballero, Engel and Haltiwanger [1997]. In that approach, the policy function of a manufacturing plant is assumed to depend on the gap between a target and the current level of employment. Using time series observations, these studies reject the quadratic cost of adjustment model and find that aggregate employment dynamics depend on the cross sectional distribution of employment gaps. We argue that these conclusions may not be justified. Instead these findings may reflect difficulties measuring the gap. Thus it appears that the gap methodology, as currently employed, may be unable to: (i) identify the costs of labor adjustment and (ii) assess the aggregate implications of labor adjustment costs.

Russell W. Cooper
Department of Economics
270 Bay State Road
Boston University
Boston, MA 02215
and NBER
rcooper@bu.edu

Jonathan L. Willis
Research Department
Federal Reserve Bank of Kansas City
Kansas City, KS 64198

1 Introduction

In recent contributions, Caballero and Engel [1993], hereafter CE, and Caballero, Engel and Haltiwanger [1997], hereafter CEH, investigate labor adjustment using a methodology, reviewed extensively below, that encompasses both convex and nonconvex adjustment processes. With this methodology, employment changes are postulated to depend on the gap between the actual and target levels of employment.[1] Both studies reach very similar conclusions on the nature of labor adjustment. The relationship between employment adjustment and the employment gap is **nonlinear**: the response to a gap is increasing in the absolute value of the gap. More importantly, both studies find evidence of these nonlinearities in time series data.

This paper questions the methodology and thus the conclusions of these studies, particularly the aggregate implications of the nonlinear plant level adjustment.[2] We argue that these reported nonlinearities may reflect difficulties in measurement of the gap rather than economic fundamentals.

The approach taken by CE and CEH relies upon a hypothesis that employment changes (Δe) respond to a gap (z) between the desired and actual number of workers at a plant. Throughout we refer to z as the **employment gap** and $\frac{\Delta e}{z}$ as the **adjustment rate**. The gain to the gap approach is that the choice of employment, an inherently difficult dynamic optimization problem, can be characterized through a nonlinear relationship between (Δe) and (z). That is, the adjustment rate can be 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 observables. We argue that these measurement problems may be severe enough to bias the inferences from these studies.

This paper constructs a dynamic model of labor adjustment **assuming** quadratic adjustment costs and uses the employment gap approach to analyze its implications. While both CE and CEH do extensive data analysis, they do not provide a mapping from their reduced form estimates to parameters of the adjustment cost function.[3] We use simulated data from our dynamic model to characterize the mapping from the structural parameters of adjustment costs to the adjustment functions and aggregate employment relationships characterized by CE and CEH.

---

[1] Hamermesh [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.

[2] To the extent that this approach is used in numerous other applications, our concerns are relevant for those exercises as well.

[3] CE do provide a static model which they use to generate a target level of employment. We embed their static model into a dynamic optimization problem.

We do not contest the general view of non-linear 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.[4]

But, does this plant-level nonlinearity matter for aggregate employment dynamics? CE and CEH, using the gap methodology, conclude that indeed it does. This finding is a major contribution of these papers. It is potentially important for business cycle and policy analyzes as it implies macroeconomics must take plant-level distributions into account. The finding implies that linear time series representations of aggregate employment, as in Sargent [1978], are misspecified.

We find:

- if the gap is correctly measured, the adjustment rate is essentially constant and the cross sectional distribution of employment gaps is irrelvant for aggregate employment dynamics

- if the employment gap is **mismeasured**, then

    1. a quadratic cost of adjustment model can generate a nonlinear adjustment rate ($\frac{\Delta e}{z}$ depends nonlinearly on z)

    2. aggregate employment dynamics can depend on the cross sectional distribution of the employment gap

- the gap measures created by CE and CEH do not correspond very closely to the actual gap measures in the quadratic cost of adjustment model.

To interpret these results, note that inferences from the gap approach are based upon the following argument: if adjustment costs are quadratic, then the adjustment rate is constant implying that aggregate employment is independent of the cross sectional distribution of employment gaps.[5] This proposition seems valid when the gap is properly measured. But it fails using the procedures of CE and CEH to measure these gaps. We find that both the CE and CEH procedures will reject the null hypothesis of quadratic adjustment costs

---

[4]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 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.

[5]Clearly, if the hazard function 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 a flat hazard 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.

even if that hypothesis is true. A methodology that is unbiased under the null hypothesis of quadratic adjustment costs is needed to assess that model.

We thus conclude that the time series evidence of nonlinear hazards reported by CE and CEH should not be taken as a finding against the quadratic adjustment cost model. Nor do their results provide evidence that nonlinear behavior at the plant level has aggregate effects.

2 The Gap Approach: An Overview

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

2.1 Gap Methodology

We follow the notation and presentation in CEH.[6] The gap between the desired employment and the actual employment (in logs) in period t for plant i is defined as:

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

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:

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

Thus 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:[7]

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

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.

---

[6] The notation and definitions in 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 their expression for aggregate employment growth differs from that in CEH.

[7] There are two interpretations of this function. Either $\Phi(z)$ represents the magnitude of adjustment (e.g. the fraction of a gap that is closed) or a probability of adjustment. The interpretation, of course, would depend on the nature of adjustment costs.

As 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.[8] 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 removed for a single period. The target would correspond to the level of employment to which an optimizing agent would eventually adjust to in the absence of any changes in the stochastic variables. This is termed the **frictionless target**. For the quadratic adjustment model, this would be the level of employment where the state dependent policy function crosses the 45 degree line.

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

2.2   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:

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

Here $\tilde{z}_{i,t}^1$ is the gap in period t **after** adjustments in the level of e have been made: $\tilde{z}_{i,t}^1 = \tilde{z}_{i,t} - \Delta e_{i,t}$. [9]

---

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

> "... the majority of existing models of factor demand simply analyze the optimal adjustment of the firm towards a static equilibrium and it is very difficult to deduce from this anything whatever about optimal behavior when there is no 'equilibrium' to aim at."

[9]Implicitly 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.

Intuitively, 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:

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

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:

$$\Delta e_{i,t} = -\theta \Delta h_{i,t} + \Delta e_{i,t}^*$$

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

$$\Delta e_{i,t} = \alpha - \theta \Delta h_{i,t} + \varepsilon_{i,t}. \tag{5}$$

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 can be (partially) remedied by looking at periods of large adjustment since then the changes in hours and employment will overwhelm the error.[10] 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:

$$\Delta E_t = \int_z z \Phi(z) f_t(z). \tag{6}$$

---

[10]They 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.

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 based upon the given specifications of a hazard function produces an aggregate relationship between employment changes and non-centered moments of z.

The CEH findings can be summarized as:

- using (5), 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) find a nonlinear relationship between the average adjustment rate, $\Phi(\tilde{z}_{i,t})$, and $\tilde{z}_{i,t}$

- CEH specify that $\Phi(z)$ is piece-wise linear. Table 3 in CEH summarizes aggregate implications and indicates that employment growth depends on the second moment of the distribution of employment gaps.

2.3   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. Appendix A characterizes the mapping from the structural parameters of the quadratic adjustment 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 the appropriate benchmarks for measuring employment gaps if shocks follow a random walk. But, if the shocks are stationary, then this measure will not provide the relevant employment target for a plant. Instead of adhering to 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.

As CE do not have plant level data, their estimation uses aggregate observations on net and gross flows for US manufacturing employment to estimate a hazard function. CE consider both a constant and a quadratic specification for $\Phi(z)$. They find that a quadratic hazard specification fits the data better than the flat hazard.[11]

---

[11]There is a difference then between the CE and CEH approaches to characterizing aggregate employment growth: CE impose a quadratic hazard while CEH work with a piecewise linear adjustment function. Our analysis will use both of these specifications.

3 A Dynamic Optimization Framework

Our analysis builds from the specification of a dynamic optimization problem at the plant level. Our structure is purposefully close to that outlined in CE. [12] We use the model as a data generating mechanism to evaluate the CE and CEH methodologies.

3.1 Quadratic Adjustment Cost Model

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

$$V(A, e_{-1}) = \max_{h,e} R(A, e, h) - \omega(e, h) - \frac{\nu}{2}\left(\frac{e - e_{-1}}{e_{-1}}\right)^2 e_{-1} + \beta E_{A'|A} V(A', e). \qquad (7)$$

Here h represents the input of hours by a worker, e_{-1} is the inherited stock of workers before quits occur (at an exogenous rate of q) and e is the stock of current workers.[13] Note the timing assumption of the model: workers hired in a given period become productive immediately.

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

$$R(A, e, h) = A(eh)^\alpha \qquad (8)$$

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

The costs of adjustment are assumed to be a quadratic function of the percent change in the stock of workers that are employed (e) and the number of workers at the start of the current period.[14] That is, the adjustment cost arises for net *not* gross hires. In (7), ν parameterizes the level of the adjustment cost function.

---

[12]For example, we have not added stochastic adjustment costs since these would drive an immediate wedge between employment changes and any gap measure. CE also include a idiosyncratic shock to the plant's gap that has no apparent counterpart in the optimization model. We did not know how to include this in our formulation.

[13]Note that there is a slight change in notation here as e and h both refer to levels and not log levels. Other 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.

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

The function $\omega(e, h)$ represents total compensation to workers as a function of the number of workers and their average hours. This compensation function is critical for generating movements in both hours and the number of workers.[15] For our analysis, we follow Bils [1987] and Shapiro [1986] and assume:

$$\omega(e, h) = w * e * \left[w_0 + h + w_1 (h - 40) + w_2 (h - 40)^2 \right]$$

where w is the straight-time wage.[16]

Using the reduced form profit function and assuming quadratic costs of adjustment, the dynamic programming problem can be solved using value function iteration. Let $e = \zeta(A, e_{-1})$ be the policy function for employment. Using this policy function, employment is determined by a stochastic difference equation.[17] Let $h = H(A, e_{-1})$ 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 $e = \zeta(A, e)$. The adjustment process, defined by iterations of $e = \zeta(A, e_{-1})$ given A, converges to the frictionless target, $e^*(A)$. Denote by $h^*(A) = H(e^*(A), A)$ the frictionless hours target. Note that this target will generally be a function of A.

The static target, used by CE, is defined as the solution to (7) when $\nu = 0$ in **all** periods. Thus employment and hours simply 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 static employment targets are indicated in the figure. Since plants take future adjustment costs into account in determining the frictionless target, they will not be as responsive as the static target to changes in the productivity shock. In general, the frictionless target will be 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, $h^*(A)$, will also deviate from the static hours target, as shown in the bottom panel of Figure 1.[18] If the frictionless employment target is below the static employment target for a given shock, then the

---

[15]A simpler model with a production function, a fixed wage rate and an employment adjustment cost is not 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.

[16]In contrast to Sargent [1978] there is no exogenous component to wage variation. In his study, variations in productivity were much larger than variations in wages.

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

[18]See Appendix A 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 Appendix A, the static hours target is not state dependent.

frictionless hours target will be above the static hours target to compensate for the lower level of employment.

3.2 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:

$$e = \lambda e^* + (1 - \lambda)e_{-1} \tag{9}$$

for $\lambda \in [0, 1]$. The dependence of e on A comes through the specification of the "target," e^*.[19] 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 (9) a solution to? When does it solve (7)?

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.[20] Consider a dynamic programming problem given by:

$$\pounds(e^*, e_{-1}) = \min_e \frac{(e - e^*)^2}{2} + \frac{\kappa}{2}(e - e_{-1})^2 + \beta E_{e^{*\prime}|e^*} \pounds(e^{*\prime}, e). \tag{10}$$

where the loss depends on the gap between the current stock of workers (e) and 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. Working 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 the shocks follow a random walk ($\rho = 1$), then partial adjustment is optimal ($\lambda_1 + \lambda_2 = 1$).[21]

The optimal policy may not take the partial adjustment form for two reasons. First, (10) is an approximation to (7). Second, shocks may not follow a random walk.

---

[19]Clearly 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.

[20]Alternatively, consider a dynamic optimization framework, such as (7), and assume that the within period return function can be written as 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.

[21]Essentially guess that the policy function is linear in the state variables and use that to solve the first

4   Empirical Implications

Our goal is to consider the empirical implications of the quadratic adjustment cost model. We do so first by looking at the aggregate implications, as in the regression results reported in CE and CEH. We then characterize the microeconomic hazards in terms of the response of employment to the employment gap and finally study the determination of the employment gap.

In the following analysis, we use our model to directly measure the employment gap at the plant level. We call this the **observed gap**. As noted above, there are two commonly used notions of targets: the **frictionless** and **static targets**. Corresponding to these two measures of the target are thus two measures of the observed gap: the **frictionless gap** and the **static gap**. We can measure these directly using our model as a data generating mechanism.

Also, we can follow CEH and try to infer the employment gap from observed hours variations, using (4) where θ is estimated from (5). 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 is defined as a change greater than one standard deviation.

4.1   Parameterization

To solve the dynamic programming problem given in (7), we need to calibrate a number of parameters and specify functional forms. We assume:

- a Cobb-Douglas production function in which 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 wage with respect to hours is close to 1 on average

---

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

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

- as estimated in Cooper-Haltiwanger [2000], the profitability shocks are represented by a first-order Markov process and are decomposed into aggregate (A) and idiosyncratic components (ε). We assume that $A \in \{0.9, 1.1\}$ and ε takes on 15 possible values. The serial correlation for the plant-level shocks is 0.83 and it is 0.8 for the aggregate shocks. The standard deviation of the plant-level shocks is set at 0.3 as estimated in Cooper and Haltiwanger [2000].[22]

- We consider two values of the adjustment cost parameter, $\nu = 1$ and $\nu = 10$, as these lead to adjustment rates such that the half-life of a gap is between 1 quarter and 1 year.[23]

Given this parameterization of the basic functions, the optimization problem given in (7) is solved using value function iteration to obtain policy functions. 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.[24]

4.2 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 can investigate aggregate implications by looking at the relationship between aggregate employment changes and the cross-sectional distribution of the employment gap given by (6).

Table 1 presents estimates for three specifications of a hazard function ($\Phi(z)$): constant, piece-wise linear and quadratic.[25] More precisely, we specify

$$\Phi(z) = \begin{cases} \lambda_0 + \lambda_1^- z + \lambda_2 z^2 \text{ for } z < 0 \\ \lambda_0 + \lambda_1^+ z + \lambda_2 z^2 \text{ for } z > 0 \end{cases} \tag{11}$$

---

[22]Cooper-Haltiwanger obtain these estimates from a model in which there were, by assumption, no adjustment costs to labor. Thus we view this parameterization as a starting point and explore the robustness of our findings to variations in these parameters of the distributions. The shocks do not follow a random walk. Relatedly, in Sargent [1978] all stochastic processes are found to be stationary.

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

[24]CEH 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.

[25]To be clear, this hazard function is imposed on the aggregate data which itself comes 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.

which nests different specifications of the hazard function.

CEH restrict λ_2 to be zero and estimate $\lambda_1^- = 1.30$ and $\lambda_1^+ = 1.32$. In contrast, CE estimate a quadratic hazard given by:

$$\Phi(z) = \tilde{\lambda}_0 + \tilde{\lambda}_2(z - z_0)^2 \tag{12}$$

where z_0 is a constant.[26] Expanding this hazard, the parameters $(\tilde{\lambda}_0, \tilde{\lambda}_2, z_0)$ can be estimated from (11) with the restriction that $\lambda_1^+ = \lambda_1^-$.[27] CE (Table 2, BLS) report $(\tilde{\lambda}_0 = 0.02, \tilde{\lambda}_2 = 0.53, z_0 = -0.82)$.We provide estimates for both specifications of the hazard function using both the observed gap and the CEH gap for the two types of targets.

4.2.1   Frictionless Target

The results for the frictionless target computed using the observed gap are reported at the top of 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 0.5 when $\nu = 1$ and 0.19 when the adjustment cost is larger, $\nu = 10$. There is essentially no evidence of any economically significant nonlinearity: the model with a constant hazard fits quite well.[28] The R^2 for this specification is essentially 1.[29]

There are two deviations from this benchmark associated with two potential "errors" in measuring the gap. First, as in CE, the static target, which is easy to compute, may be used instead of the frictionless target. The second is the CEH measure of the gap.

4.2.2   Static Target

Using the static target one would strongly **reject** the hypothesis that the hazard function is flat in favor of either the piecewise linear or quadratic cases as shown in the lower portion of Table 1a. For example, in the quadratic specification, we find that when $\nu = 1$, λ_2 is estimated at 0.33 with a standard error of 0.01. Further, the coefficients in the piecewise linear specification ($\lambda_1^+ = 1.46, \lambda_1^- = 1.16$) are also statistically and economically significant. The nonlinearity is also statistically significant when $\nu = 10$ for both the piecewise linear and quadratic cases. Note though that here the R^2 for the constant hazard

---

[26]CE introduce additional features that we have avoided. As discussed in their Section IV.2, they apply an idiosyncratic shock to the distribution of plant deviations. We are not sure what this transformation represents in our structural model and thus we have excluded it from our analysis.

[27]Thus when we refer to estimation of (12), we are doing so by estimating (11) with the restriction that $\lambda_1^+ = \lambda_1^-$.

[28]Though the regression coefficients on some of the nonlinear pieces are statistically significant, they add essentially nothing to the goodness of fit and the estimated coefficients are very small.

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

model is quite high (0.91) so that adding these higher moments of the cross sectional distribution, while significant, do not lead to the large increases in R^2 reported by CE.

The table also includes the quadratic specification given in (12) in the bottom row of Table 1a, where $\lambda_1^+ = \lambda_1^-$ has been imposed. From this regression, we estimate ($\tilde{\lambda}_0 = -0.35, \tilde{\lambda}_2 = 0.54, z_0 = -0.15$) when $\nu = 1$ and ($\tilde{\lambda}_0 = -0.22, \tilde{\lambda}_2 = 0.17, z_0 = -0.23$) when $\nu = 10$. The coefficients are all significantly different from zero. [30]

The difference in results between using the frictionless and static targets to determine 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, illustrated the difference between the two targets. Therefore, 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 (12), one can rewrite the aggregate employment growth equation, (6), as:

$$\Delta E_t = \lambda_0 m_{1,t}^s + \lambda_1 m_{2,t}^s + \lambda_2 m_{3,t}^s + \varepsilon_t + \lambda_0 \left(m_{1,t}^f - m_{1,t}^s \right) + \lambda_1 \left(m_{2,t}^f - m_{2,t}^s \right) + \lambda_2 \left(m_{3,t}^f - m_{3,t}^s \right)$$

where $m_{i,t}^s$ is the i^{th} uncentered moment of the cross-sectional distribution of the static gap in period t and m_i^f is the corresponding measure for the frictionless gap. The error term contains three measurement error terms in addition to ε_t. If any of these measurement errors are 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 on the three moments of the static gap.[31] We estimate $(-1.66, 0.37, 1.07)$ as the coefficients on the three moments. The standard errors are all less than 0.04. These results indicate that the error is related to the static gap in a nonlinear way, thus leading to the nonlinear estimates of the adjustment function.

4.2.3 CEH Measure of the GAP

Second, the frictionless target could be inferred from variations in observed hours, as in CEH, opening the possibility of 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 in-

---

[30]At least with regards to the quadratic term and $\nu = 1$, our estimates are not at odds with the CE findings. Again, our methodology does not include the extra step of randomization across the gaps. Also, our model does not imply that z_0 should be present: if there is a zero gap, there should be no adjustment.

[31]Thanks to Peter Klenow for discussions on this characterization of the measurement error.

cluding only observations in which the employment and hours changes exceed one standard deviation, as in the sample splits of CEH.

Note that for both specifications ($\nu \in \{1, 10\}$), the constant hazard hypothesis is rejected for both the full and big change samples. In fact, the flat hazard specification yields rather nonsensical results: ie. the adjustment rate is in excess of 100% for the full sample and is actually negative ($\nu = 10$) once we concentrate on the large changes in employment and hours. Further, there is a nontrivial increase in the R^2 associated with adding these terms to the hazard function, particularly for the big change sample with $\nu = 1$.

4.2.4 Summary

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

We now turn to further explorations of these rejections of the constant hazard model. To what extent do they reflect aggregation? To what extent do they reflect mismeasurement of the gap created in the estimation of auxiliary model that links hours variations to the target? We address these questions in turn.

4.3 Micro Hazards

For the analysis of plant level hazards, we consider an expanded specification of the hazard function used in the aggregate analysis:

$$\Phi(z) = \begin{cases} \lambda_0 + \lambda_1^- z + \lambda_2 z^2 + \xi_0 \frac{1}{z} \text{ for } z < 0 \\ \lambda_0 + \lambda_1^+ z + \lambda_2 z^2 + \xi_0 \frac{1}{z} \text{ for } z > 0. \end{cases}$$

The additional $\frac{1}{z}$ term in the hazard function is added to reflect the possible presence of a constant term in the policy function ($\phi(z)$). For this specification, the constant hazard prediction is that all coefficients except for λ_0 should be insignificantly different from zero. If only λ_0 and λ_1^-/λ_1^+ are significant, then the hazard function is linear but not flat. The estimates (standard errors in parenthesis) for this exercise are presented in Table 2a (observed gap) and 2b (CEH gap).

4.3.1 Frictionless Target

Using the frictionless target measure, the hazard rate (λ_0) is very close to the estimates obtained from aggregate data for both values of ν. For example at $\nu = 1$, the value of

λ_0 from the micro hazard is 0.50, the same as the value estimated from the aggregate data. Evidently, there is no bias introduced by the aggregation per se as long as the appropriate measure of the target is being used. Using this gap measure there is no evidence of nonlinearity at the plant level as λ_2 is essentially zero. While there is some evidence for a linearly decreasing hazard as a function of the employment gap as $\lambda_1^- + \lambda_1^+$ sum to zero. Adding these linear terms $\left(\lambda_1^-, \lambda_1^+\right)$ into the estimation can only account for eight percent of the total sum of squared deviations from a constant hazard hypothesis. Also important to note is that the inverse of z has no explanatory power, $\xi_0 = 0$. These same points hold for $\nu = 10$.

For the observed frictionless gap, Figure 2 presents the simulated data when $\nu = 1$. This is intended as a counterpart to Figure 1a in CEH. The top part of the figure contains a scatterplot of the observed frictionless gap and the rate of employment adjustment, expressed as the change in employment divided by the employment gap, along with a line depicting a constructed average adjustment rate. Note that this relationship is a bit "cloudy" as the gap measure is not quite a sufficient statistic for employment adjustment.

The bottom panel shows the cross sectional distribution of the employment gaps produced by the model. This distribution is slightly more diffuse than the corresponding picture in Figure 1a in CEH. Figure 5 presents the same panels for the case of $\nu = 10$. Here the average adjustment rate line has a slightly positive slope.

4.3.2 Static Target

For the case of the static target we find evidence of significant nonlinearity at the plant-level. Figure 3 displays the scatterplot of adjustment rates and a constructed average adjustment measure using the static target.

Starting first with a simple constant hazard specification for $v = 1$, the estimate for λ_0 reported in Table 2a is 0.35, much lower than the 0.5 estimate using the frictionless target. This reflects the fact that the static target does not incorporate the cost of adjustment as does the frictionless target. As described earlier, Figure 1 illustrates how measures of static employment targets corresponding to above and below-average shocks will be exaggerated relative to the frictionless target. Therefore, the static target on average will indicate that more adjustment is required, i.e. a larger gap exists, and the adjustment rate for a given change in employment will be smaller.

The more general hazard specification (including all but $\frac{1}{z}$) indicates marginally significant nonlinearities, but the R^2 on this regression is 0. However, when the inverse of z is added to a constant hazard specification, its coefficient, ξ_0, is strongly significant and the R^2 increases to 0.42. This result indicates that there is not a proportional relationship between the static gap and changes in employment.[32] The results corresponding to $\nu = 10$

---

[32]This result can again be understood from the perspective of the measurement error induced by using

are similar. Though the nonlinearities are significant, the only change in the R^2 coincides with the inclusion of $\frac{1}{z}$.

4.3.3 CEH Measure of the Gap

For the hazards estimated at the plant level using the CEH measure of the gap, there is little to be discerned from the results in Table 2b. The hazard rate (λ_0) is estimated to be *much* greater than the value associated with the frictionless target, often in excess of 100% adjustment, depending on the value of ν. As we shall see, the problems with this measure of the gap are a direct consequence of severely biased estimates of θ.

4.4 Estimates of θ

The final step in understanding the aggregate results and the CEH methodology is to explore the estimates of θ, the relationship between the gap and hours. The issue is to determine if problems with the CEH gap measure stem from the estimation of θ. As CE did not estimate θ, this issue does not arise directly in their analysis. However, there is still the issue of whether the θ they impose from their static model generates a reasonable gap measure.

Table 3 summarizes the estimates of this parameter for a number of different specifications. The first two rows correspond to the estimated value of θ 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 while the second measure uses the static target. The other rows use the CEH approach to estimate θ. Note that their results do not depend on the definition of the target since it is not observed to them. Results are again reported for the two different parameterizations of the quadratic adjustment cost model, $\nu = 1$ and $\nu = 10$.

First note that the sign on θ from the CEH regression is opposite that obtained when the observed gap is used in the regression, as in (3). Since their methodology relies on θ to construct a measure of the gap, this difference is important to understand.

Recall that in (5), the error term contains the change in the employment target level. Assuming that 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, this correlation is 0.69, indicating that the sign on θ in (5) will be negative. 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.

---

the static rather than the frictionless employment target.

CEH acknowledge that hours and employment target changes could be correlated, so they only use 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 will be swamped by the effects of large changes in hours and employment. But in a model of convex adjustment, the only periods in which there will be 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 0.95 in the full sample and 0.99 in the CEH-criterion subsample. To obtain an unbiased estimate of θ in a model of quadratic costs of adjustment, controlling for changes in employment target levels is essential.

The implications of the sign reversal are displayed in Figure 4, which shows a sample of employment changes, deviation 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 contains 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 both 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. Since hours and employment are positively correlated, the effect of the negative sign on θ is essentially to force the constructed employment gap to be a dampened version of the change in employment. The correlation between the actual measures of the gap and the CEH gap measures are positively correlated (approximately 0.52 for the big change subsample at $\nu = 1$), but the conclusions to be drawn from analysis of these series are very different.

As for the CE approach, the estimate of θ obtained from using the frictionless and static gap measures differ. In fact, the estimate using the static target, as in CE, produces an estimate of θ that is exactly equal to the one obtained analytically.[33] However, the gap measure produced by using this estimate of θ **does not** correspond with the frictionless gap measure. The difference is due to the fact that the frictionless hours target is dependent upon the productivity shock, whereas the static hours target is independent of the shock.

This distinction between the two hours targets has important implications for the measurment of the gaps in CEH, represented by (3). A complicated log-linearization in the form of (18) may be constructed using the frictionless targets. The relationship between

---

[33]Using (18) from Appendix A and the given parameterization, CE would find θ is equal to 8.8.

the employment gap and hours deviations used by CEH can be written as

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

Using the correct target for hours and the frictionless employment gap, we do obtain the analytically calculated value of θ. But the problem for the CEH methodology is that there are now two unobservables since the hours target cannot be approximated by a constant mean. 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 is illustrated precisely by the observed static target results above for the aggregate and micro hazards.

4.5 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.

4.5.1 Specification of Optimization Problem

With regards 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 Bils [1987], analyze a model in which:

$$R(A, e, h) = A \left(e^{\alpha_e} h^{\alpha_h} \right) \tag{13}$$

with $\alpha_e = 0.72, \alpha_h = 0.77$.[34] In this case, our conclusions on the methods of CE and CEH do not change: nonlinearities remain in the aggregate regressions, using (12) as the hazard function, as shown in the second row of Table 4.

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:

$$\frac{\nu}{2} \left(e - e_{-1} \right)^2. \tag{14}$$

Using this specification of the adjustment cost function does not have a significant effect on any of our conclusions: nonlinearities remain in the aggregate regressions (see row 3 of Table 4).

---

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

4.5.2 Shocks

Of particular concern are the specifications of the stochastic processes as these parameterizations came from Cooper and Haltiwanger [2000] who investigated a model without costs of adjusting labor in a panel of approximately 7,000 plants over the 1972-88 period. Thus it is important to consider robustness with respect to the process governing both the aggregate and plant specific shocks.

For the idiosyncratic shocks, our analysis uses the standard deviation and serial correlation reported by Cooper-Haltiwanger and creates a state space with 15 shocks and a transition matrix that mimics this process assuming that the shocks are normally distributed.[35] To explore the sensitivity of the results to the idiosyncratic shock parameterization, we compute results over a range of plausible alternative parameter settings. Table 5 displays the corresponding results for the aggregate regression, where the static targets are used to measure the employment gaps. The serial correlation values range from 0.7 to 0.99 and the standard deviations range from 0.1 to 0.4. The actual values used in the baseline estimates are (0.83,0.3). In the left side of the table, where $\nu = 1$, increases in the serial correlation lead to higher estimates of λ_2, the quadratic coefficient of the adjustment function. When the standard deviation is increased, the linear coefficient, λ_1, increases, but the effect on λ_0 and λ_2 depends on the level of serial correlation. For high levels of serial correlation, an increase in the standard deviation leads to small increases in the nonlinearity, but the level of nonlinearity decreases for lower levels of serial correlation. In all of the results, however, the hypothesis of a constant or linear hazard function would be rejected. The right side of Table 5 shows results when the scalar on the adjustment cost function, ν, is increased to 10. The conclusions are very similar.

We also consider variations in the representation of the aggregate shock process and here find that the specification can matter. Cooper-Haltiwanger represent the aggregate shock process as a two-state Markov process where the shock has values of $\{1.1, .9\}$ and a transition matrix of

$$\begin{pmatrix} .8 & .2 \\ .2 & .8 \end{pmatrix} \tag{15}$$

which reproduces the serial correlation and variance of aggregate shocks. By construction, the distribution of shocks is uniform. Note that there are two aspects of the approximation: the number of elements in the state space and the elements in the matrix.

An alternative estimate of the variance and serial correlation comes from Sargent [1978] who estimates an aggregate shock process in his model of dynamic labor demand. Sargent reports an aggregate shock process in which the serial correlation is 0.94. Row 4 of Table

---

[35]Thus we follow the procedure outlined in Tauchen [1986].

4 shows our results for that specification of the aggregate shocks.[36] Again our findings are robust to this alternative specification of the aggregate shocks.

However, changes in the fineness and the spread of the state space can influence our findings with regards to the CE aggregate results (the hazard given in (12)). Suppose that we increase the number of elements in the aggregate state space from 2 to 11.[37] In this case, for the aggregate regression we find that the level of nonlinearity is lower (row 5 of Table 4), but is still strongly significant. For $\nu = 10$, our results are similar.

To study the spread of the state space, we set the endpoints one standard deviation above and below the mean, where the endpoints were previously set 0.4 standard deviations from the mean. We use an 11-point state space representation where the aggregate shocks here go from 0.77 to 1.3. This change produces a substantial decrease in the estimated degree of nonlinearity in the aggregate regression (row 6 of Table 4). The estimate of λ_0 is now close the the constant hazard estimate of 0.4 for this specification. Similar results are reported for $\nu = 10$. In both cases, one would still reject a flat hazard, but the magnitude of the degree of nonlinearity is much smaller than baseline estimates. Interestingly, the degree of nonlinearity in plant-level employment adjustment actually increases when the endpoints of the aggregate state space are expanded. Further, even for the aggregate regressions, there is evidence of significant nonlinearity in the piecewise linear specification. Note further that these aggregate results hold when there is an implausibly large domain for the aggregate shocks.

5   Conclusions

The point of this paper was to assess the inferences of CE and CEH that aggregate employment dynamics depend upon the cross sectional distribution of employment gaps. The paper argues that due to measurement problems, a researcher might find that the cross sectional distribution matters for aggregate time series even if adjustment costs are quadratic due to measurement problems. Thus the time series evidence presented by CE and CEH is not convincing. So, despite the overwhelming evidence that plants 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.[38] However, even the CE approach falls short due,

---

[36]Sargent did not estimate the mean value of the shock, so we are unable to normalize the estimate of the variance of the innovations for use in our model. Therefore, we set the variance at 0.008, which corresponds to the estimates by Cooper-Haltiwanger. Here shocks are drawn from a normal distribution.

[37]In order to solve the model we also have to reduce the number of idiosyncratic shocks from 15 discrete points to 5 points. This reduction leads to only a marginal change in results.

[38]We understand that data limitations led CEH to their formulation.

primarily due 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. However, implementing this procedure with actual data is less clear.

There are, however, 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 now exist in the literature on investment, durables and price setting. These involve using indirect inference techniques to match the moments produced by simulations of a structural model with those observed.[39] Clearly, labor is next.

6   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), 313-58.

Caballero, R. Engel, E. and J. Haltiwanger, "Aggregate Employment Dynamics: Building From Microeconomic Evidence," *American Economic Review,* 87 (1997),115-137.

Cooper, R. and J. Haltiwanger, "On the Nature of the Capital Adjustment Process," NBER Working Paper # 7925, October 2000.

Davis, S. and J. Haltiwanger, "Gross Job Creation, Gross Job Destruction and Employment Reallocation," *Quarterly Journal of Economics*, 107 (1992), 819-64.

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

Hamermesh, D. *Labor Demand,*Princeton, NJ: Princeton University Press, 1993.

Nickell, S. "Fixed Costs, Employment and Labour Demand over the Cycle," *Economica*, 45 (1978),329-45.

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 Labor and Capital," *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.

---

[39]We 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.

Appendix A: CE Gap and Target

CE use this basic framework to generate some analytic results on θ, the parameter that connects variations in hours with variations in the number 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) yielding a first order condition of:

$$R_h(A, e, h) = \omega_h(e, h). \tag{16}$$

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, they optimize over the number of workers setting hours at $\bar{h}$ implying:

$$R_e(A, e, \bar{h}) = \omega_e(e, \bar{h}). \tag{17}$$

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 (17). 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 (17) given the functional forms 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 a 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 (17) as

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

Substitution of this relationship into the log-linearized version of (16) 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.[40] This can be rewritten as

$$\hat{e}_t^{**} - \hat{e}_t = \frac{1 - \alpha + \xi_w}{1 - \alpha}\hat{h}_t. \tag{18}$$

---

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

Using the mean level of observed hours as an approximation for $\bar{h}$, equation (3) denotes the same relation as (18) with θ equal to $\frac{1-\alpha+\xi_w}{1-\alpha}$.[41]

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, $\xi_w = 2.18$.

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 (8) 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 $\alpha_L = .65$, α is equal to 0.72. Using these calculation, θ can be determined from $\theta = \frac{1-\alpha+\xi_w}{1-\alpha}$.

---

[41]We are grateful to Robert King for pushing us to make this connection.

Table 1a: Aggregate Implications

Specification	$\nu = 1$						$\nu = 10$					
	const	λ_0	λ_1^+	λ_1^-	λ_2	R^2	*const*	λ_0	λ_1^+	λ_1^-	λ_2	R^2
Observed gap:												
frictionless target	-0.012	0.50				1.00	0.001	0.19				1.00
	(0.000)	(0.001)					(0.000)	(0.000)				
	-0.001	0.51	-0.04	0.03		1.00	-0.001	0.20	0.002	-0.02		1.00
	(0.000)	(0.001)	(0.002)	(0.002)			(0.000)	(0.001)	(0.001)	(0.001)		
	-0.012	0.54			-0.04	1.00	0.001	0.19			0.003	1.00
	(0.000)	(0.004)			(0.005)		(0.000)	(0.001)			(0.002)	
static target	0.03	0.34				0.91	0.01	0.07				0.86
	(0.001)	(0.003)					(0.000)	(0.001)				
	-0.03	-1.04	1.46	1.16		0.96	-0.02	-0.35	0.39	0.29		0.89
	(0.00)	(0.04)	(0.05)	(0.04)			(0.00)	(0.03)	(0.03)	(0.02)		
	0.03	-0.10			0.33	0.95	0.01	-0.12			0.10	0.88
	(0.00)	(0.02)			(0.01)		(0.00)	(0.01)			(0.01)	
	-0.03	-0.34	0.16	0.16	0.54	0.96	-0.03	-0.21	0.08	0.08	0.17	0.90
	(0.00)	(0.02)	(0.01)	(0.01)	(0.02)		(0.00)	(0.01)	(0.01)	(0.01)	(0.01)	

Table 1b: Aggregate Implications

Specification	$\nu = 1$						$\nu = 10$					
	const	λ_0	λ_1^+	λ_1^-	λ_2	R^2	*const*	λ_0	λ_1^+	λ_1^-	λ_2	R^2
CEH gap:												
full sample	0.00	1.53				0.51	0.00	1.43				0.34
	(0.00)	(0.05)					(0.00)	(0.06)				999
	0.30	9.48	-62.53	-18.74		0.76	-0.06	-2.50	107.94	26.55		0.40
	(0.01)	(0.48)	(2.98)	(1.67)			(0.01)	(1.19)	(20.19)	(21.98)		(999.00)
	0.03	-0.28			30.29	0.58	0.00	-2.42			993.57	0.37
	(0.00)	(0.16)			(2.51)		(0.00)	(0.61)			(157.71)	(999.00)
big change	0.00	0.56				0.06	0.00	-0.61				0.12
	(0.00)	(0.07)					(0.00)	(0.05)				999
	0.39	0.62	-27.87	31.16		0.76	0.02	-8.96	144.87	184.75		0.15
	(0.01)	(0.39)	(2.47)	(1.82)			(0.00)	(1.70)	(32.74)	(34.85)		999
	0.02	-4.06			118.87	0.32	0.00	-4.88			1500.83	0.14
	(0.00)	(0.25)			(6.21)		(0.00)	(0.83)			(293.07)	(999.00)

Table 2a: Estimates of Policy Function

Specification	$\nu = 1$						$\nu = 10$					
	λ_0	λ_1^+	λ_1^-	λ_2	ξ_0	R^2	λ_0	λ_1^+	λ_1^-	λ_2	ξ_0	R^2
Observed gap:												
frictionless target	0.50					0.00	0.19					0.00
	(0.000)						(0.000)					
	0.50	-0.04	0.04			0.08	0.17	0.05	0.04			0.03
	(0.000)	(0.000)	(0.000)				(0.000)	(0.000)	(0.000)			
	0.50	-0.06	0.03	0.01		0.08	0.15	0.20	0.19	-0.19		0.07
	(0.000)	(0.001)	(0.001)	(0.001)			(0.000)	(0.001)	(0.001)	(0.001)		
	0.50				0.00	0.08	0.19				0.00	0.00
	(0.000)				(0.000)		(0.000)				(0.000)	
	0.50	-0.04	0.04		0.00	0.14	0.17	0.05	0.04		0.00	0.03
	(0.000)	(0.000)	(0.000)		(0.000)		(0.000)	(0.000)	(0.000)		(0.000)	
	0.50	-0.05	0.02	0.01	0.00	0.14	0.15	0.20	0.19	-0.19	0.00	0.07
	(0.000)	(0.001)	(0.001)	(0.001)	(0.000)		(0.000)	(0.001)	(0.001)	(0.001)	(0.000)	
static target	0.35					0.00	0.10					0.00
	(0.01)						(0.00)					
	0.34	0.13	-0.04			0.00	0.14	-0.04	-0.07			0.00
	(0.01)	(0.02)	(0.01)				(0.00)	(0.00)	(0.00)			
	0.36	0.04	-0.13	0.06		0.00	0.20	-0.27	-0.31	0.15		0.00
	(0.01)	(0.04)	(0.04)	(0.03)			(0.00)	(0.01)	(0.01)	(0.01)		
	0.32				0.05	0.42	0.09				0.01	0.12
	(0.00)				(0.000)		(0.00)				(0.000)	
	0.23	0.13	0.19		0.05	0.42	0.10	-0.01	-0.03		0.01	0.12
	(0.01)	(0.01)	(0.01)		(0.000)		(0.00)	(0.00)	(0.00)		(0.000)	
	0.13	0.56	0.64	-0.31	0.05	0.42	0.13	-0.10	-0.12	0.06	0.01	0.12
	(0.01)	(0.03)	(0.03)	(0.02)	(0.000)		(0.00)	(0.01)	(0.01)	(0.01)	(0.000)	

Table 2b: Estimates of Policy Function

Specification	$\nu = 1$						$\nu = 10$					
	λ_0	λ_1^+	λ_1^-	λ_2	ξ_0	R^2	λ_0	λ_1^+	λ_1^-	λ_2	ξ_0	R^2
CEH gap:												
full sample	1.64					0.00	1.62					0.00
	(0.12)						(0.06)					
	1.73	-1.00	-0.77			0.00	1.62	-0.20	-0.14			0.00
	(0.19)	(1.86)	(1.58)				(0.10)	(3.01)	(3.14)			
	1.93	-4.80	-5.31	13.67		0.00	1.78	-13.07	-12.51	162.12		0.00
	(0.25)	(3.55)	(3.95)	(10.90)			(0.13)	(8.04)	(7.82)	(93.88)		
	1.65				-0.01	0.00	1.60				-0.01	0.02
	(0.12)				(0.000)		(0.06)				(0.000)	
	1.75	-0.59	-1.25		-0.01	0.00	1.59	4.56	-3.93		-0.01	0.02
	(0.19)	(1.86)	(1.57)		(0.00)		(0.09)	(2.98)	(3.11)		(0.00)	
	1.98	-5.12	-6.66	16.28	-0.01	0.00	1.70	-4.38	-12.53	112.68	-0.01	0.02
	(0.25)	(3.55)	(3.94)	(10.88)	(0.00)		(0.13)	(7.96)	(7.75)	(93.01)	(0.00)	
big change	1.20					0.00	0.55					0.00
	(0.18)						(0.10)					
	1.40	-2.23	-2.15			0.00	0.70	-5.31	-5.64			0.00
	(0.31)	(3.20)	(3.12)				(0.18)	(6.39)	(6.63)			
	1.96	-15.55	-16.48	54.75		0.00	0.86	-21.48	-21.33	248.70		0.00
	(0.42)	(7.24)	(7.66)	(26.72)			(0.26)	(18.69)	(18.28)	(270.11)		
	1.18				0.03	0.02	0.55				0.00	0.00
	(0.18)				(0.000)		(0.10)				(0.000)	
	1.36	-4.46	0.59		0.03	0.02	0.72	-1.53	-10.85		0.00	0.00
	(0.31)	(3.16)	(3.09)		(0.00)		(0.18)	(6.37)	(6.62)		(0.00)	
	1.81	-15.15	-10.91	43.94	0.03	0.02	0.88	-18.00	-26.82	253.25	0.00	0.00
	(0.41)	(7.16)	(7.57)	(26.43)	(0.00)		(0.25)	(18.65)	(18.24)	(269.48)	(0.00)	

Table 3: Estimate of theta

Specification	$v = 1$				$v = 10$			
	Estimate	Std. Error	R^2	Obs.	Estimate	Std. Error	R^2	Obs.
Observed gap:								
frictionless target	5.00	(0.002)	0.82	983000	3.37	(0.002)	0.82	983000
static target	8.95	(0.000)	1.00	983000	9.01	(0.001)	0.99	983000
CEH gap:								
full sample	-3.49	(0.004)	0.48	983000	-0.57	(0.001)	0.30	983000
big change	-4.66	(0.004)	0.88	165939	-0.90	(0.001)	0.82	134407

Table 4

Aggregate Implications for Alternative Specifications of the Model Using the Static Target Measure*

		$\nu = 1$					$\nu = 10$			
	const	λ_0	λ_1	λ_2	R^2	*const*	λ_0	λ_1	λ_2	R^2
(1) Static Baseline	-0.03	-0.34	0.16	0.54	0.96	-0.03	-0.21	0.08	0.17	0.90
	(0.00)	(0.02)	(0.01)	(0.02)		(0.00)	(0.01)	(0.01)	(0.01)	
(2) CE/Bils	-0.03	-0.33	0.16	0.54	0.96	-0.03	-0.21	0.08	0.17	0.90
	(0.00)	(0.02)	(0.01)	(0.02)		(0.00)	(0.01)	(0.01)	(0.01)	
(3) $(\nu/2)(e-e_{-1})^2$	-0.03	-0.33	0.12	0.52	0.96	-0.03	-0.18	0.07	0.15	0.90
	(0.00)	(0.02)	(0.01)	(0.02)		(0.00)	(0.01)	(0.01)	(0.01)	
(4) Sargent (2 point dist.)	-0.02	-0.20	0.12	0.43	0.96	-0.04	-0.28	0.10	0.22	0.93
	(0.00)	(0.02)	(0.01)	(0.01)		(0.00)	(0.01)	(0.00)	(0.01)	
(5) Sargent (11 point dist.)	0.01	0.01	0.04	0.32	0.93	-0.01	-0.08	0.03	0.11	0.87
	(0.00)	(0.03)	(0.01)	(0.03)		(0.00)	(0.02)	(0.01)	(0.01)	
(6) Sargent (11 point dist.)	0.04	0.31	0.01	0.07	0.92	0.00	0.07	0.02	0.03	0.87
	(0.00)	(0.01)	(0.01)	(0.01)		(0.00)	(0.01)	(0.00)	(0.00)	

*The regression results are for the aggregate employment growth equation (6) using the CE hazard function (12). In simplified form,

the estimation equation is $\Delta e_t = const + \lambda_0 m_{1,t} + \lambda_1 m_{2,t} + \lambda_2 m_{3,t} + \varepsilon_t$ where $m_{i,t}$ is the i^{th} uncentered moment of the cross-sectional

distribution of the employment gap.

Table 5

Aggregate Implications for Alternative Settings of the Idiosyncratic Shock Process Using the Static Target Measure*

ρ	σ	$v = 1$					$v = 10$				
		const	λ_0	λ_1	λ_2	R^2	*const*	λ_0	λ_1	λ_2	R^2
0.70	0.10	0.00	0.06	0.05	0.60	0.97	-0.01	-0.13	0.06	0.43	0.97
		(0.00)	(0.01)	(0.01)	(0.01)		(0.00)	(0.00)	(0.00)	(0.01)	
0.90	0.10	0.00	0.14	0.04	0.60	0.97	0.00	-0.08	0.06	0.47	0.98
		(0.00)	(0.01)	(0.01)	(0.01)		(0.00)	(0.00)	(0.00)	(0.01)	
0.99	0.10	0.01	0.19	0.02	0.63	0.97	0.00	-0.01	0.05	0.50	0.98
		(0.00)	(0.00)	(0.01)	(0.01)		(0.00)	(0.00)	(0.00)	(0.01)	
0.70	0.30	-0.05	-0.43	0.16	0.44	0.95	-0.04	-0.19	0.07	0.13	0.89
		(0.01)	(0.03)	(0.01)	(0.02)		(0.00)	(0.02)	(0.01)	(0.01)	
0.90	0.30	-0.02	-0.20	0.13	0.58	0.97	-0.03	-0.22	0.09	0.22	0.92
		(0.00)	(0.01)	(0.01)	(0.02)		(0.00)	(0.01)	(0.01)	(0.01)	
0.99	0.30	0.01	0.12	0.05	0.66	0.97	-0.01	-0.09	0.08	0.47	0.98
		(0.00)	(0.01)	(0.01)	(0.01)		(0.00)	(0.00)	(0.00)	(0.01)	
0.70	0.40	-0.06	-0.44	0.15	0.33	0.94	-0.04	-0.14	0.05	0.08	0.88
		(0.01)	(0.03)	(0.01)	(0.02)		(0.01)	(0.02)	(0.01)	(0.01)	
0.90	0.40	-0.03	-0.33	0.16	0.52	0.96	-0.04	-0.20	0.08	0.15	0.90
		(0.00)	(0.02)	(0.01)	(0.02)		(0.00)	(0.01)	(0.01)	(0.01)	
0.99	0.40	0.01	0.08	0.06	0.66	0.97	-0.01	-0.13	0.10	0.45	0.98
		(0.00)	(0.01)	(0.01)	(0.01)		(0.00)	(0.00)	(0.00)	(0.01)	

*The regression results are for the aggregate employment growth equation (6) using the CE hazard function (12). In simplified form, the estimation equation is $\Delta e_t = const + \lambda_0 m_{1,t} + \lambda_1 m_{2,t} + \lambda_2 m_{3,t} + \varepsilon_t$ where $m_{i,t}$ is the i^{th} uncentered moment of the cross-sectional distribution of the employment gap.

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

Figure 2
Average adjustment rate function and scatterplot
($\nu = 1$)

Figure 3
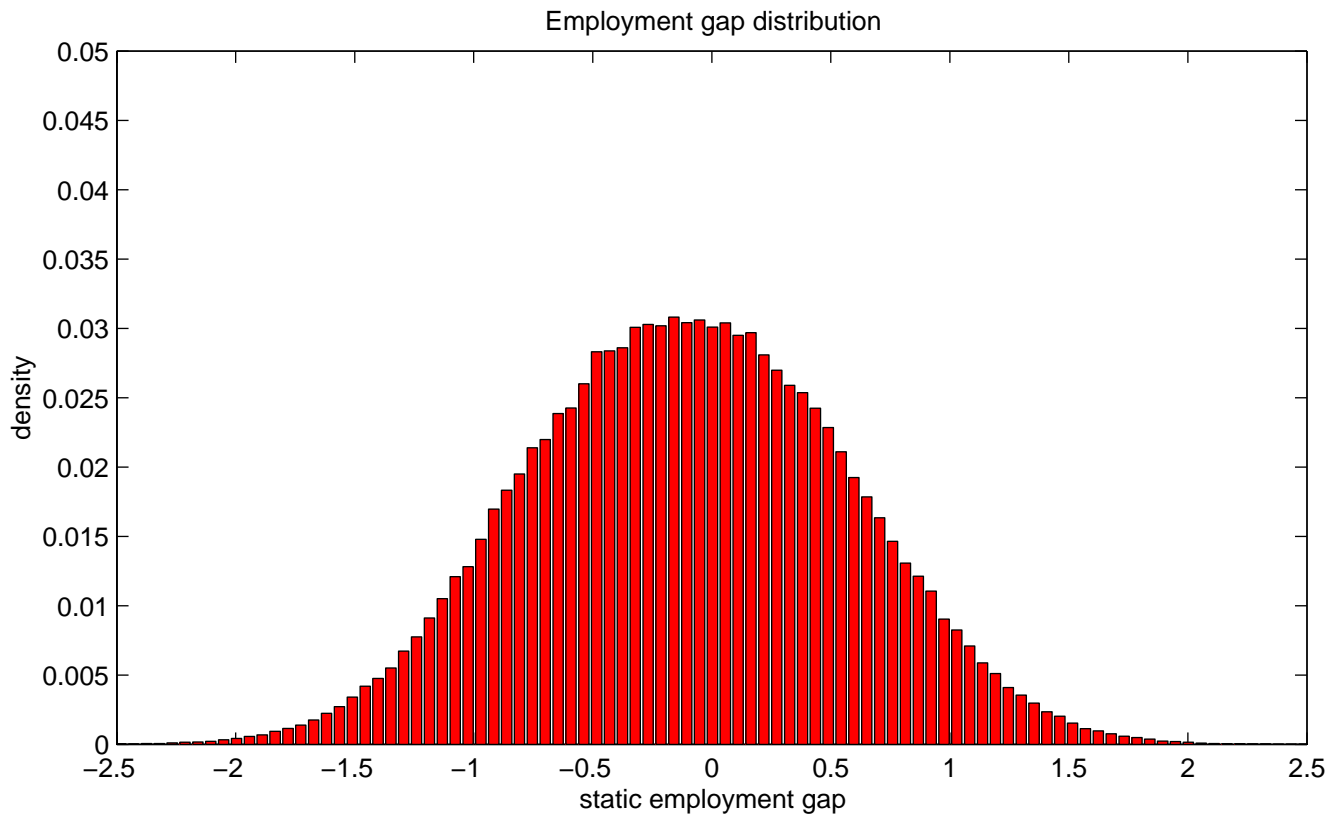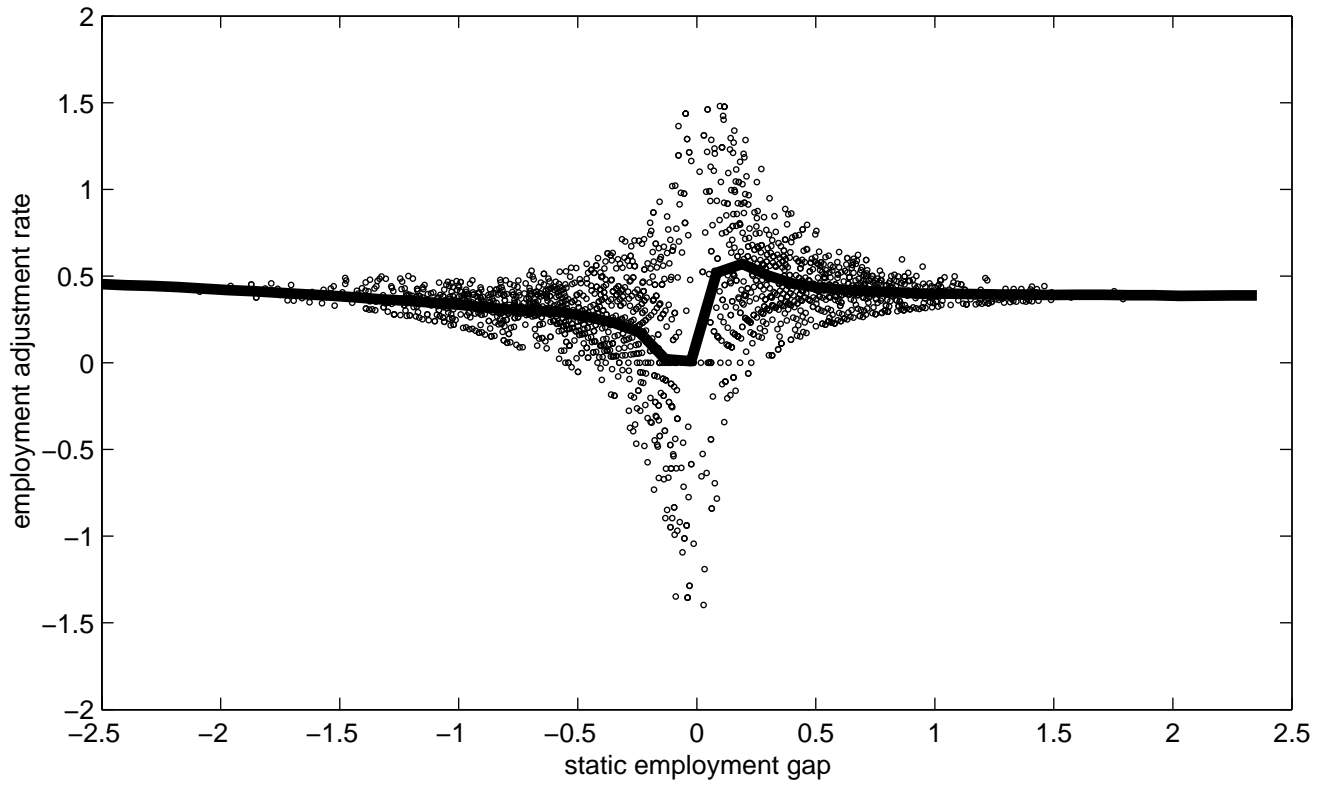Average adjustment rate function and scatterplot
($\nu = 1$)

Employment gap distribution

Figure 4
Simulated employment gaps: measured directly
($\nu = 1$)

Figure 5
Average adjustment rate function and scatterplot
($\nu = 10$)

Employment gap distribution