Identifying Adjustment Costs of Net and Gross Employment Changes

João M. Ejarque and Øivind A. Nilsen *,†

January 26, 2007

Abstract

A relatively unexplored question in dynamic labour demand regards the source of adjustment costs, whether they depend on net or gross changes in employment. We simulate a structural model of dynamic labor demand where the firm faces adjustment costs related to gross and net changes in its workforce. These simulations allow us to select several moments of hiring and of net changes in employment that provide identification for our estimation with actual data. We then estimate adjustment costs in a panel of establishments on quarterly data. The main component of adjustment costs in our panel is quadratic adjustment costs to gross changes in employment. We find no significant evidence of fixed costs in our data.

JEL Classification: C33, C41, E24, J23

Keywords: Employment, adjustment costs, firm level data, structural estimation.

1 Introduction

In recent years there has been renewed interest in labour demand at the firm and plant level. A largely unexplored question in this area regards the source of adjustment costs, whether they depend on net or gross changes in employment. Thus, despite existing empirical work documenting a variety of different adjustment costs, it is unclear how the dynamics of employment at the firm and plant level are affected by a mix of the various adjustment costs. Understanding the adjustment costs structure is of great importance

*Ejarque: University of Essex, jejarque@essex.ac.uk. Nilsen: oivind.nilsen@nhh.no.

†Very preliminary and incomplete. Do not cite. Comments welcome.

for understanding firms' response to exogenous shocks, such as changes in labour market policies and demand fluctuations together with price variations in input factors. It should be noted that firms' response to exogenous shocks is crucial for policy interventions to fight unemployment. This paper studies the impact of the various adjustment cost structures using firm level data using a structural approach. Our dynamic optimization problem incorporates the different types of adjustment costs (convex and fixed, net and gross). We use a minimum (weighted) distance method to recover the structural parameters of the adjustment costs function. It should be noted that with the existence of fixed costs, there is no simple Euler equation linking the marginal cost of additional labour today with future marginal benefits. Thus, in most empirical papers there is only a weak or somewhat reduced form relationship between the structural mode and the empirical model. Contrary to these more reduced form econometric models, a structural model seems very suitable for the dynamic optimization problem. By using this structural approach we hope to get a better understanding of the costs associated with firms' path of labour demand over the business cycles.

The literature on net adjustment costs is vast, going back to Oi (1962) who used sectoral employment data. Most literature have assumed convex adjustment costs.[1] In micro level data the observed inaction in employment adjustments is consistent with non-convexities in adjustment costs (see for instance Caballero and Engel (1993) and Caballero et al. (1997). Rota (2004) uses annual firm level data from Italian manufacturing industry and finds that net fixed adjustment costs to be amazingly 15 months of labour costs. Nilsen et al. (2006), using annual data from Norwegian metal products and machinery industry, find that the fixed costs is likely to be the most important component of the net adjustment costs for small plants (less than 25 employees), especially when reducing the labour stock.

The empirical literature on gross flows is scarce. Among the few, there are the seminal studies by Hamermesh (1989, 1992, 1995). The data used in these latter studies are based on samples of very few units only. Thus it is hard to generalise from these studies and to be sure that the findings are not driven by some anomalies. Abowd and Kramarz (2003) analyze the costs of gross employment changes using a cross section of French manufacturing firms. They find the fixed costs to be very large and statistically significant, and about three times larger for separations than for hires. Hiring costs for skilled employees have both a fixed and a *concave* component, while hiring

[1]See Nickell (1986), Hamermesh and Pfann (1996a) for surveys of dynamic labour demand models including non-convexities in the adjustment costs function. See Rota (2004) and Cooper and Willis (2004).

costs for unskilled workers are largely fixed. Interesting though, in a follow-up study Michaud and Kramarz (2004) find that when using a panel, the fixed costs are negligible, but that there are still concavities in the adjustment costs functions. Thus, it seems like the estimated significance and magnitude of fixed adjustment costs depend crucially on the econometric model and structure of the dataset.

One of the few papers to study adjustment costs of both gross and net employment changes simultaneously is Hamermesh and Pfann (1996b).[2] They use particular assumptions regarding functional forms in a problem of the dynamic employment adjustment of the firm, but mainly they use aggregate data. Their findings indicate that quits do not seem to matter much for explaining employment movements.

In our model a firm with homogeneous labour face adjustment costs associated with both net and gross employment flows. The firm experiences stochastic demand such that its revenues fluctuate over time, and also faces voluntary quits from its workforce. Quits have two characteristics: they are voluntary and they are directly costless. In the absence of quits, net and gross flows are the same in every model period. Typically in the data we are able to observe only total exit (quits plus fires) and we therefore need indirect measures to allow us to separate quits from other exit. We first generate artificial data for a calibrated model under different costs structures (net or gross, fixed or convex adjustment costs) to understand their different impact on the dynamics of labour adjustment. We then explore a set of moments from these artificial data, looking for variation which will allow for identification in the real data. Identification in our problem is conditioned by the unobservability of voluntary exit (quits).

Once we find the identifying variation, we use the simulated method of moments to estimate the structural parameters of interest on actual panel data. We use a cut from a panel containing quarterly establishment level information for Portugal with information on employment stocks, entry and exit. Having access to quarterly data for a panel is quite unique. This also allows us to analyse the effect of time aggregation, from quarterly to annual data. We strongly believe that the answer to our question requires appropriate firm level data. The measure of time, in this case one quarter, is the same in the model and in the data, and is the frequency where information arises and hiring or firing decisions are taken. We find that, for our dataset, the firm faces quadratic costs of both net and gross adjustments, with an emphasis on gross adjustment. There are virtually no fixed costs.

[2] They examine US manufacturing quarterly data from 1960 to 1981, and find that quits are exogenous to the log of employment, despite a strong comovement.

The next section contains the model. Section III contains simulation results of a calibrated model. Section IV section discusses the moments we use to achieve identification of costs structures together with a description of the Portuguese data and the estimation of our model. Section V discusses the results. Section VI offers concluding remarks.

2 Model

Consider a model of a firm that decides on the optimal employment of homogeneous labour. Capital and other inputs are assumed to be maximized away, and the only decision taken by the firm is that of employment. The profit function of the firm is given by

$$\Pi_t\left(A_t, Q_t, L_{t-1}\right) = R_t - wL_t - AC\left(.\right)$$

where $R_t = A_t L_t^{\alpha}$ is the revenue function, L_t is the employment level used in production in period t, A_t is a demand or technology shock which may contain a permanent and a transitory component such that $A_t = A_t^P A_t^T$, and w is the constant wage rate per worker. We assume that R_t is concave in labour, $0 < \alpha < 1$, where the α reflects labour's share in the production function as well as the elasticity of the demand curve faced by the firms.

The following equation links employment, hiring and separations;

$$\Delta L_t = L_t - L_{t-1} = M_t - Q_t = (H_t - F_t) - Q_t$$

Quits (voluntary separations), Q_t, are assumed to have a deterministic and a stochastic component. The deterministic component is a constant fraction δ of the beginning of period labour stock, L_{t-1}, constrained to be an integer. In addition there is an iid quit shock, z, defined over a small set of integers $[0, 1, 2, ..., z^{\max}]$, uncorrelated with either the employment level or the technology shock. This shock embodies the idiosyncratic nature of each firm's Labour force composition. The total number of quits is of course bounded below by zero and above by initial employment:

$$Q_t = \min(\delta L_{t-1} + z_t, L_{t-1})$$

Adjustment costs $AC(.)$ are assumed to be a function of both net, ΔL_t, and gross, M_t, employment changes. The gross change to the labour stock is defined as $M_t = H_t - F_t$, which can be negative (fires) or positive (hires). In the model hiring and firing cannot occur in the same period (a quarter). This matches the fact that, according to most countries's labour employment legislation, one cannot hire and fire at the same time unless in case of serious misconduct. Therefore $M_t > 0$ implies both $F_t = 0$ and $M_t = H_t > 0$, while if downscaling, $M_t < 0$ implying both $H_t = 0$ and $M_t = -F_t < 0$. Exit in the model is given by $EX_t = Q_t - M_t \times I(M_t < 0)$, and entry by $EN_t = M_t \times I(M_t > 0)$.[3]

We furthermore assume that new hires become productive immediately and that the firm observes (A_t, Q_t, L_{t-1}), and then decides on its current labour force L_t.

We consider fixed and quadratic costs and both types of costs have gross flows (M_t) and net flows (ΔL_t) components:

$$\begin{aligned} AC(.) &= C^n(\Delta L_t) + C^g(M_t) \\ C^n(\Delta L_t) &= \frac{\gamma^n}{2}\left[1 + \frac{L_{t-1} + L_t}{2}\right]^{-1}[L_t - L_{t-1}]^2 + F^n \times I(\Delta L_t \neq 0) \\ C^g(M_t) &= \frac{\gamma^g}{2}\left[1 + \frac{L_{t-1} + L_t}{2}\right]^{-1}[M_t]^2 + F^g \times I(M_t \neq 0) \end{aligned}$$

where $I(.)$ is an indicator function. The parameters $(F^n, F^g, \gamma^n, \gamma^g)$ define adjustment costs. Also, we note that these functions are symmetric with respect to upward and downward adjustment, and this is a convenient benchmark to use.

3 Simulated Model

The model is simulated to identify which features of the artificial data respond to changes in the cost structure. This will provide the identifying variation necessary to the estimation exercises. Features such as the average labour force in the panel will be around 100 workers by construction, and should not be significantly affected by changes in the cost structure. It stands to reason that we should focus on features of net (ΔL_t), and gross (M_t), changes in employment, since they relate directly to the different cost functions we use. Finally, we look exclusively at quarterly data.

3.1 Parameter Values

The first order condition of the frictionless model is used to calibrate and center the model. We set the value of α at 0.66, which is the labour share in the frictionless model. The value of the wage is just a scale parameter and we set it at $w = 0.66$. The state space for L is the set of integers

[3] In the data we observe simultaneous entry and exit. This is due to either time aggregation or heterogeneity, which we do not discuss here. Otherwise, it is voluntary exit.

including zero up to a reasonable upper bound and we choose to center or model at $\bar{L} = 100$. The mean level of the shocks is then determined by $\bar{A} = \left(\bar{L}\right)^{1-\alpha}\left(\frac{w}{\alpha}\right) = \left(\bar{L}\right)^{1-\alpha}$. This has the further advantage that it sets $\bar{R} = \bar{L}$. This parameterization makes it easy to understand the magnitude of adjustment costs. Consider the adjustment of one worker around $\bar{L} = 100$, on quadratic costs associated with net changes:

$$C\left(\Delta L = 1\right) = \frac{1}{2}\frac{\gamma}{(101.5)} \approx 0.005\gamma$$

If we want the adjustment of one worker to cost 1% of quarterly revenues $\bar{R} = 100$, we have to set $\gamma = 200$. It turns out this creates too much inaction so we run the model with $\gamma^g = \gamma^n = 20$, implying the adjustment of one worker costs the firm one thousandth of revenues.

The discount factor $\beta = 1/(1+r)$ contains a 1.5% interest rate, which corresponds to 6.1% annual interest rate. The fixed cost is set at one monthly wage, $F^g = F^n = 0.22$. The deterministic quit parameter is set at $\delta = 0.01$. This implies 1% of labour force turnover due to quits every quarter. In our numerical model it is rounded to the nearest integer so these quits will on average equal 1. The iid stochastic component of quits has support on the set of integers $z \in [0, 1, 2]$, with discrete density $p\left(z\right) = [0.34, 0.33, 0.33]$, which yields an expected value of 1 worker per quarter.

The permanent technology shock A_t^P has a 3 point support implied by the employment values $\bar{L} = [90, 100, 110]$ and from the first order condition of the frictionless model. This process has a symmetric transition matrix with a value $q = 0.96$ on the diagonal and values $(1-q)/2$ in the off diagonal elements. Finally, the temporary process A_t^T has log normal distribution and a 9 point support and transition matrix obtained by applying the gaussian quadrature method of Tauchen (1986), with persistence and standard deviation values $(\rho = 0.15, \sigma(\varepsilon) = 0.04)$.[4]

3.2 Exercise

We generate data for a panel of 500 firms. The only thing that distinguishes these firms from each other is the realization of the shocks. We initialize one third of the firms at each permanent shock with their corresponding employment value. Then we simulate 141 periods and delete the first 100 periods of the panel and keep the last 41 quarters of artificial data. We then compute moments and statistics on the panel. After all this, we change the

[4]This shock is written: $A_t^T = e^{a_t}$, with $a_t = \rho a_{t-1} + \varepsilon_t$, with $\varepsilon_t\tilde{} N(0, \sigma)$.

cost structure and do it again, while keeping the same realizations of the random variables. We look at the following cases.

Experiments	
Only Quad. Costs	$\gamma^g > 0, \gamma^n = 0$ $\gamma^n > 0, \gamma^g = 0$
Only Fixed Costs	$F^g > 0, F^n = 0$ $F^n > 0, F^g = 0$
Estimation	$F \neq \gamma \neq 0$

Before discussing these simulation outcomes, it is important to point out the key difference between fixed and quadratic costs. Fixed costs induce a mass of probability **at** zero, while quadratic costs induce probability mass **around** zero.[5] Fixed costs also empty the density **around** zero, and shift this density either to zero or to the tails. In short, fixed costs induce a spike at zero **and** fat tails. Quadratic costs induce a bell shape with no fat tails.

3.3 Quadratic Costs

Quadratic adjustment costs have a dramatic impact on the model.[6] The measured costs are almost negligible, but their impact on employment dynamics is large.[7] The following two graphs show the distributions of quarterly entry (left plot) and of ΔL (right plot) for three cases, i) the case with convex net adjustment costs, $\gamma^g = 0$, $\gamma^n > 0$ (taller for ΔL, flatter for entry, no tails), ii) the case with convex gross adjustment costs, $\gamma^g > 0$, $\gamma^n = 0$ (flatter for ΔL, taller for entry, no tails), and for comparison iii) the frictionless case (thin line, tall tails),:

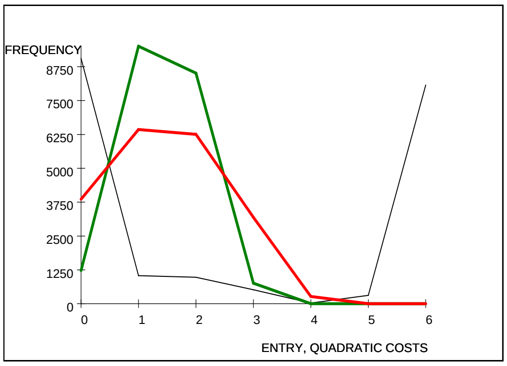

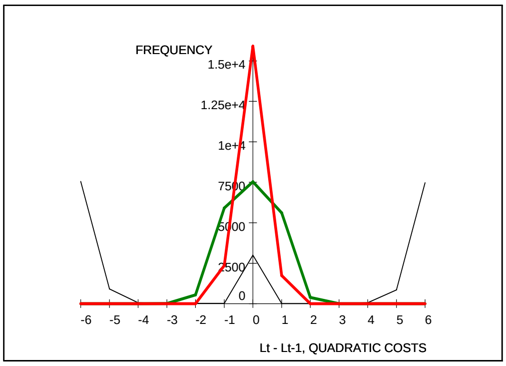

[5]Given the discrete nature of our data (number of employees) this distinction between at zero and around zero can be blurred for small firms.

[6]See the appendix for the simulation of the frictionless model, i.e. $F = \gamma = 0$.

[7]We stress the term *measured* costs, because total revenues are now on average one half of one percent lower, and this is a non negligible cost too.

Figure 1a and 1b: Entry, EN, and net employment changes, ΔL

It is the fact that for $\gamma^n > 0$, we have a narrower distribution of ΔL, and a flatter distribution of entry that provides the identifying variation since net and gross objects move in opposite directions as costs impact on net or on gross changes.[8]

We are also interested in how long inaction lasts. We want to know how long and how often firms remain inactive. Inaction is defined as the zeros of a particular variable, in this case ΔL, EN and EX. For example, under costs to net changes, there are 15908 observations of $\Delta L = 0$, while under costs of gross changes there are 7538 observations of $\Delta L = 0$. We study the distribution of these zeros. We measure the length of every sequence of zeros, and count the sequences, and look at what happens to this pattern as we change the cost structure.

Histogram of Inaction Sequences: Convex adjustment costs											
ΔL	1	2	3	4	5	6	7	8	9	Longest	Aver
$\gamma^g > 0$	3056	1070	444	131	49	18	7	4	1	13	1.57
$\gamma^n > 0$	517	322	207	146	97	98	107	60	56	40	7.40
EN	1	2	3	4	5	6	7	8	9	Longest	Aver
$\gamma^g > 0$	491	189	57	31	3	5	1	1	0	8	1.57
$\gamma^n > 0$	1620	583	193	73	17	11	5	1	1	9	1.54
EX	1	2	3	4	5	6	7	8	9	Longest	Aver
$\gamma^g > 0$	1684	530	164	43	14	7	2	0	0	7	1.44
$\gamma^n > 0$	1427	463	133	46	15	4	1	1	1	9	1.46

Table 1: Histogram of Inaction Sequences: Convex adj. costs

As we move from gross to net costs, we get many more longer sequences of zeros of ΔL. In particular the longest sequence of zeros of ΔL lasts 40 periods if costs are a function of net changes, while it is only 13 periods if costs are a function of gross changes. There are more zeros but the duration pattern is clearly different, and we have more consecutive periods of $\Delta L = 0$. The impact on duration in entry is existent, while the impact on exit is harder to see. In general (in other simulations) entry does tend to get

[8] We must be careful here: all else equal, the same value of gamma implies a different size of adjustment costs as a fraction of revenues for gross versus net functions. However, it is still the case that how smooth these two distributions are, and on how tall their peaks are, will allow identification of the two gammas.

shorter, and exit is less affected than entry, most likely because exit is more directly affected by quits which are exogenous to the cost structure.

3.4 Fixed costs

Our simulations based on fixed adjustment costs are given in the following figure.

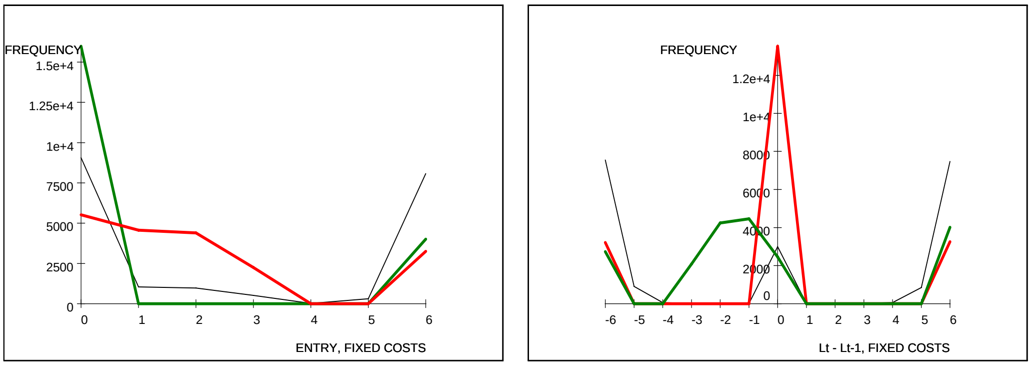

Figure 2a and 2b: Entry, EN, and net employment changes, ΔL

Here we see that when fixed costs are associated with net adjustments we obtain the high peak on the right hand side graph, and the smoother entry line on the left hand side one. The frictionless case has the low peak on the right hand side graph, and the high tails. The negative bias induced by quits on the distribution of ΔL when costs impact on gross variations provides identifying variation for gross versus net.[9] Additionally, the tails are still fat, since the main impact of the fixed costs is to generate a lot of zeros by taking density away from small variations.

In the following table we describe the inaction sequences in the presence of fixed adjustment costs.

[9] We assume that quit behaviour is exogenous to the cost structure, and in our empirical exercise the distribution of quits is fixed and not estimated. Actually, this is in fact not strictly true. If quits depend deterministically on previous employment, L_{t-1}, and employment is affected by the cost structure, we do not have exogeneity.

Histogram of Inaction Sequences: Fixed adjustment costs											
ΔL	1	2	3	4	5	6	7	8	9	Longest	Aver
$F^g > 0$	1372	373	81	17	4	2	0	0	0	6	1.33
$F^n > 0$	1487	892	525	376	232	168	111	107	80	29	3.23
EN	1	2	3	4	5	6	7	8	9	Longest	Aver
$F^g > 0$	701	848	678	544	398	268	190	148	94	23	3.92
$F^n > 0$	3323	690	171	51	12	6	1	0	0	7	1.30
EX	1	2	3	4	5	6	7	8	9	Longest	Aver
$F^g > 0$	1940	512	108	25	5	1	1	0	0	7	1.32
$F^n > 0$	1732	462	107	25	7	2	1	0	0	7	1.34

Table 2: Histogram of Inaction Sequences: Fixed adj. costs

Moving from gross to net costs has a clear opposite implication on the duration patterns of the zeros of entry and of ΔL: average duration increases clearly for ΔL and falls clearly for entry. This opposite effect again provides identifying variation. Again, exit is mostly determined by quits so that there is little variation. This justifies looking mainly at entry and at ΔL in our empirical exercise.

4 Estimation

We now take our model to actual data using the simulated method of moments, and we use a variety of data moments to estimate six parameters.[10] The data comes from "Inquerito ao Emprego Estruturado" (IEE). We have quarterly data from the first quarter of 1991 to the last quarter of 1995 on employment (end of quarter head count), and on hirings and separations during the quarter, at the establishment level, for Portugal.[11] A very detailed description of this data can be found in Varejão (2003). We select data for firms that report an initial employment value of between 80 and 120 workers, and use only the firms that report all the 20 consecutive quarters of information. This delivers a sample of 135 firms.

[10] In this exercise the period in the data (one quarter) is assumed to be the same period underlying the dynamic decision in the model. There is no time aggregation and therefore the difference between net and gross changes in only due to voluntary exit.

[11] For establishments with less than 100 workers the panel contains only a sample. For establishments with more than 100 workers it contains all establishments in the country. After cleaning the data, a significant fraction of them is naturally eliminated.

We act in line with the literature in setting several parameters which we do not estimate. Since we only have observations of employment, and not of revenues or other firm level information, complete identification of the separate effects of shocks versus cost structures is impossible. Our assumed parameters are $\alpha = w = 0.66$, and $\bar{L} = 100$. The permanent and transitory components, A_t^P and A_t^T are merged into a single process because there does not seem to be strong evidence of significant persistence at different points in the panel.[12] This merged process contains two parameters that we estimate, (ρ, σ), and, just like in our simulation experiments above, is discretized using Tauchen's method into a Markov process, although here with 49 points in the state space. The shock A_t here contains also a stochastic trend. For our 135 firm sample, aggregate employment is declining and this must be tackled in the model. The trend is introduced as a Markov process where all firms start at the high state such that during the 20 quarters aggregate employment evolves approximately as it does in the data. This has three points in the state space.[13,14]

The model has one more random variable. Voluntary separations are modelled as an iid random variable, and we set $\delta = 0$. The appendix contains details of this process. Finally, the four parameters left which we estimate, are the parameters of the cost function, namely $(\gamma^n, \gamma^g, F^n, F^g)$.

There are 24 moments in the data that are used in the estimation.[15] These are described here, and are highlighted in bold type in the tables below. We first compute a variety of averages (μ), standard deviations (σ), and first order serial correlation measures (ρ), for our data. We construct these moments first in a time series dimension using the 20 observations for each firm, and then take cross sectional averages $\mu^{cs}(.)$ and cross sectional

[12] For example, an establishment with 10 workers in the first quarter in this wider dataset is extremely unlikely to have 30 workers at the end of the 20 quarters. But fluctuations between 80 and 120 are common, as are large fluctuations within a firm.

[13] The data does not lend itself easily to detrending due to its discrete nature. The support of this process is [0.9 0.95 1] and the transition matrix has small probabilities of moving down and vanishingly small probabilities of moving up. This process was purely a pragmatic way of fixing the aggregate trend.

[14] Aggregate employment shrinks by 15% over the sample. There are 52 firms that end the sample 10% below trend relative to their initial employment (below $0.9 \times 0.85 \times L(1)$), and 47 firms that finish the sample 10% above trend (above $1.1 \times 0.85 \times L(1)$). 36 firms are within a 10% interval of to each side of the trend. This seems quite symmetric but with low persistence in the sense that most firms in the end of the sample have employment values outside a reasonable range of where they started. This outcome of [52, 36, 47] changes to [34, 77, 24] if the window is 20% around trend, and this would indicate some more persistence for L.

[15] See the Data Appendix for more details. We experimented with more, but the bootstrapped variance covariance matrix became close to singular.

standard deviations $\sigma^{cs}(.)$, of these firm specific means and firms specific standard deviations. We do this for four variables, employment levels (L), net changes (ΔL), entry and exit. We show below a table with these 6 measures for each of the four variables. Out of this table we pick 9 measures to use in our estimation.

We also use the frequencies of ΔL (-2,-1, 0, 1, 2), and the frequencies of EX (0, 1, 2). We choose to exclude exit moments because exit is more directly determined by the distribution of quits which is exogenous. [16] We also use as moments figures describing the inaction of ΔL and EN (the number of single-period and two-period sequences of inaction).

In addition to the described 21 moments, we use 3 more moments. We know that fixed costs create a mass of observations *at* zero and at the tails, eliminating mass *close to* zero. We include observations *at* and *around* zero. In fact, most of the moments selected above focus on the center of the various distributions. However, our data has significant variability, as we can see by the long tails of the different distributions. There is one moment so far that explicitly measures dispersion, and that is $\sigma^{cs}(\mu(L)) = 17.7$. We now want other statistics for the tails of the distribution. We choose arbitrarily the 10% cutoff of the distributions of *rates*, that is, the distribution of $\Delta L_t/L_{t-1}$, EN_t/L_{t-1}, and EX_t/L_{t-1}. These cut off values in our data are (-0.0645, 0.0403) for the distribution of $\Delta L_t/L_{t-1}$, 0.0862 for the distribution of entry rates, and 0.0968 for the distribution of exit rates. This means that between -0.0645 and 0.0403 we have 80% of the observations of $(L_t - L_{t-1})/L_{t-1}$, and to the left of 0.0862 (0.0968) we have 90% of the observations of the entry (exit) rate. We use the first three numbers just described, again leaving out exit.

All this comes down to 24 numbers we want our algorithm to match using only six parameters. We build a quadratic form which will be our distance measure that is minimized:

$$\begin{aligned} D &= [Y(\theta) - X]\, V^{-1} [Y(\theta) - X]' \\ \theta^* &= \arg\min D(\theta) \end{aligned}$$

where $Y(\theta)$ is the vector of simulated moments which is a function of the parameter vector θ, and X is the vector of moments in the data. The covariance matrix V is obtained by bootstrapping the panel, resampling the 135 firms one thousand times. This will provide us with a matrix with one thousand draws of a 24 long vector, and that provides us with a 24 squared

[16] Note also that increasing the number of moments makes the covariance matrix close to singular.

covariance matrix. The quadratic form D has a chi-squared distribution with 18 degrees of freedom.[17]

5 Results

We obtain the following estimates:

γ^n	γ^g	F^n	F^g	ρ	σ	D
0.8895	2.2181	0.0037	0.0013	0.0074	0.1182	310

Table 3: Estimated results

We see that quadratic costs have both net and gross components, while fixed costs are close to zero. For example, F^g is about one to 100 thousand of quarterly revenues. The shock process is virtually iid. This is a nice outcome as it means that the model is closest to the data when most dynamics are endogenously generated by the shape of production and cost functions. Recall that the quit shock is also iid. Clearly, the smooth shapes of **both** the distributions of entry and of ΔL, force the model to use both types of quadratic costs, while the fact that there is always a smooth and significant density at $\Delta L = 1$, and at entry equal to 1, implies fixed costs must be virtually non existent.

These parameters generate the following behavior:[18]

[17] A chi squared with 20 degrees of freedom at 5% has value 31.41, so that our goal is a value of our distance measure D (a quadratic form) in this neighbourhood.

[18] The numbers in the first column is mean of firm averages, the second st.dev of firm averages, the third means of firm specific st.dev., the forth st.dev. of firm specific st.dev. Furthermore, $\mu^{cs}(\rho_1)$, is the cross sectional mean (μ) of the 135 firm specific first order serial correlation (ρ_1) coefficients, each computed over the 20 observations for a given firm.

IEE Moments. 135 firms, 80-120 workers.						
	$\mu^{cs}(\mu)$	$\sigma^{cs}(\mu)$	$\mu^{cs}(\sigma)$	$\sigma^{cs}(\sigma)$	$\mu^{cs}(\rho_1)$	$\sigma^{cs}(\rho_1)$
L	92.3	**17.7**	**9.75**	6.27	**0.804**	0.215
ΔL	**-0.80**	1.45	**5.07**	4.46	**-0.058**	0.234
EN	**3.05**	2.56	**3.62**	3.65	**0.068**	0.268
EX	3.85	2.56	3.85	2.58	0.021	0.235

Estimated Moments. 135 firms. Matching IEE Data.						
	$\mu^{cs}(\mu)$	$\sigma^{cs}(\mu)$	$\mu^{cs}(\sigma)$	$\sigma^{cs}(\sigma)$	$\mu^{cs}(\rho_1)$	$\sigma^{cs}(\rho_1)$
L	92.0	6.11	6.10	2.69	0.84	0.14
ΔL	-0.61	0.69	2.76	0.45	0.02	0.23
EN	2.51	0.71	2.26	0.42	0.12	0.25
EX	3.12	0.61	2.69	0.59	-0.03	0.22

Table 4: Actual and estimated cross sectional moments

A salient feature of this previous table is that while we fit reasonably many of the numbers in columns 1, 3, and 5, we cannot do a good job on the cross sectional dimension reported in columns 2, 4 and 6. This cross sectional "heterogeneity"is notoriously difficult to match with a very stylized model such as the one we use, that does not contain any heterogeneity apart from the realization of the shocks. We can literally see that by the absence of fat tails in the estimated data:

Estimated model				Actual data			
	ΔL	EN	EX		ΔL	EN	EX
<-5	121	-	-	<-5	292	-	-
-5	104	-	-	-5	68	-	-
-4	175	-	-	-4	128	-	-
-3	254	-	-	-3	173	-	-
-2	**337**	-	-	-2	**286**	-	-
-1	**360**	-	-	-1	**388**	-	-
0	**461**	**752**	487	0	**560**	**1015**	602
1	**283**	**359**	436	1	**240**	**395**	447
2	**247**	**392**	436	2	**161**	**306**	376
3	167	361	346	3	123	234	285
4	98	292	281	4	57	158	206
5	57	211	198	5	59	126	172
>5	36	333	516	>5	165	466	612

Table 5: Frequencies of employment and employment changes

The distribution of ΔL is reasonable but its tails are too small. The fixed cost, despite being tiny, helps in matching the number of zero entries. However, it also emphasizes the fact that the number of single entries 352 is counterfactually *less* than the number of entries equal to 2 (389).[19] Nevertheless, the model is in a rough ballpark of the data. The inaction data are shown in the following table.

Actual data												
	1	2	3	4	5	6	7	8	9	Longest	Aver	n
ΔL	**305**	**62**	23	6	1	1	2	0	0	13	1.39	560
EN	**292**	**112**	44	16	17	4	8	2	2	20	2.00	1015
EX	276	68	28	6	4	3	2	1	1	13	1.54	602

Estimated model									
	1	2	3	4	5	6	7	8	9
ΔL	**330**	**48**	10	1	0	0	-	-	-
EN	**302**	**91**	38	19	6	3	3	1	-
EX	347	50	10	2	1	0	-	-	-

Table 6: Inaction sequences

From the table, we see that we fit reasonably for ΔL and EN, while exit matches worse. Finally, our measure of the tails yields:

	10% cut-off rates			
	$\Delta L/L^-$	$\Delta L/L^+$	EN/L	EX/L
Model	-0.0448	0.0331	0.0655	0.0780
Data	-0.0645	0.0403	0.0862	0.0968

Table 7: Cut off values of rate distributions

Again we see that we fall short on the dispersion we see in the data.

For our 135 firm sample, aggregate employment is declining as shown in the left hand side graph below. We tackled this problem successfully using a stochastic trend, and the artificial data is the smoother and more straight line which ends with slightly higher employment. But our data has a rich cross sectional variation which our estimated model does not match

[19] If we kill these tiny fixed costs leaving all else constant, the entry distribution numbers on [0,1,2] become [735, 377, 393] and for exit they become [473, 448, 437]. The D statistic rises to 324.

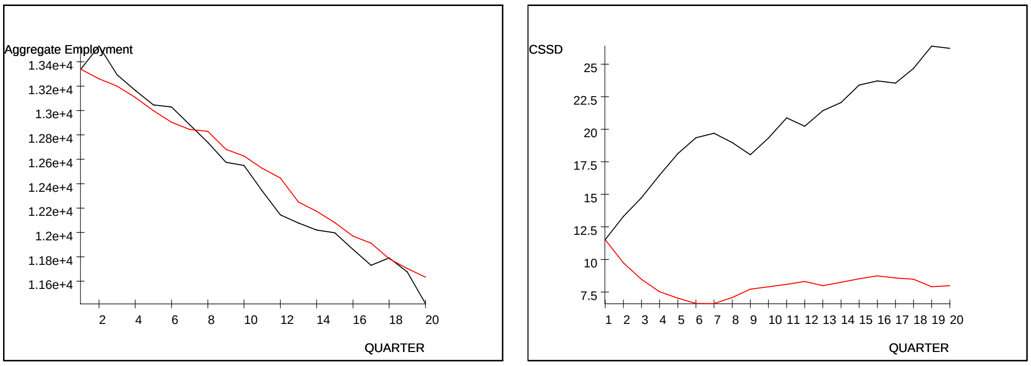

Figure 3: Aggregate employment and cross sectional st. dev.

The right hand side graph shows the evolution of the cross sectional standard deviation of employment levels: while it is increasing in the actual data (reflecting the arbitrary selection of firms by initial size), the estimated model starts by imposing the exact initial employment distribution but then the estimated stochastic process (because it is almost iid) drives the artificial panel to a smaller dispersion.

We also have somewhat mixed results on the dynamics of employment:

OLS: IEE quarterly data[20]				
X_t	const.	X_{t-1}	X_{t-2}	R^2
ΔL	-0.9625 (7.17)	-0.1066 (5.34)	-0.0277 (1.41)	0.012
EN	1.9952 (15.5)	0.1476 (7.83)	0.1385 (7.36)	0.056

OLS: artificial data				
X_t	const.	X_{t-1}	X_{t-2}	R^2
ΔL	-0.5718 (10.0)	$+0.1180$ (6.04)	$+0.0987$ (5.24)	0.030
EN	1.6001 (22.0)	0.2219 (11.3)	0.1024 (5.37)	0.076

Table 8: Dynamics in employment

Despite the fact that these regressions all have very low R^2, we do not fit well the dynamics of ΔL, while we do a reasonable job on the dynamics of entry. We can argue that the serial correlation is low on both cases, despite the fact that in the model it is positive while in the data it is negative, but the difference is too big to brush aside.

We suspect that much of the inability of the model to fit the data comes from one single data moment, the 1015 realizations of zero entry. This number, relative to the number of zero exits and zero net changes is too large such that the standard model has difficulty matching all zeros. The data also has a degree of variation in employment that is hard to replicate.

Nevertheless we succeed in picking moments that identify gross versus net adjustment costs and clearly estimate the different contributions of the two. Given the class of models and the particular richness of the data, even if we do not replicate all of the data features, we manage quite well on the proposed research question.

An additional issue that complicates the identification of adjustment costs is the frequency at which we observe the data. This matters if the interval at which the data is observed differs from the frequency of the arrival of information and of decision making in the firm. A possible inference bias may occur if we observe firms at annual intervals, for example on December 31st, and their workforce is of about the same size in consecutive Decembers, but we also observe significant entry and exit during the year. These observations may lead to the conclusion that net variations are the source of costs whereas the truth might be that entry and exit are costly but the firm incurred shocks at high frequency - for example quarterly - that induced the adjustment, and due to stationarity we end up more or less where we started.

We tackle this problem of time aggregation by performing two controlled experiments. We first aggregate to annual frequencies the actual quarterly data we have and generate artificial data also at quarterly frequencies, and also aggregate it. In this experiment we acknowledge the existence of an aggregation problem, and try to replicate the construction of the actual data in the artificial data. We then estimate the model again – as we did directly on the quarterly data – but using moments from the time aggregated data. This we compare to our quarterly estimates. In our second experiment, we simulate a model at annual frequency, and match the annual artificial data against the time aggregated data from our quarterly panel. In this way, we create a bias by allowing the model to have a time frequency different from that at which the data was originally observed. This again we compare to our two previous estimation exercises, allowing us to quantify for this data the impact of time aggregation. Finally, we explore the time aggregation issue more generally by running our two experiments only on simulated data under a variety of controlled "truths", where we generate artificial data at quarterly frequency given a set of parameters, and then aggregate it, and perform, for each value of parameters used – the two experiments described above. This will inform on what types of shapes of production and cost functions are more conducive to inference bias due to time aggregation.

6 Conclusion

This paper has given attention to whether employment adjustment costs are related to net or gross employment changes, or both. We have used a structural model of the dynamic employment decision of the firm. This allows us to directly analyses the importance of fixed adjustment costs since with such costs the future value of additional labour today depends on future choice with respect to adjustments. With the existence of fixed costs there is no simple Euler equation linking the marginal cost of additional labour today with future marginal benefits since it is no guarantee that this plant will be adjusting its labour stock in the next period. Our findings indicate that fixed costs for both net and gross employment changes are virtually zero. On the other hand there are convex costs associated with both net and gross changes. This implies that it is important to depart from a standard model with just adjustment costs associated with just net changes.

There are several issues we have not addressed and which need to be explored in future work. First would be to see the importance of time aggregation, for instance going from quarterly to annual data. Secondly one should allow for asymmetries, i.e. whether firms are increasing or decreasing in size. Furthermore, allowing for heterogeneous labour should be given attention.

References

[1] Abowd, J., and Kramarz, F. (2003), "The costs of hiring and separations", Labour Economics 10, 499-530

[2] Caballero and Engel (1993) "Microeconomic adjustment hazards and aggregate dynamics" Quarterly Journal of Economic, 108, 359-383

[3] Caballero Engel and Haltiwnager. (1997) "Aggregate employment dynamics: Building from microeconomic evidence", American Economic Review.

[4] Cooper, R., Haltiwanger, J., and Willis, J., (2004), "Dynamics of labour demand: evidence from plant-level observations and aggregate implications", NBER Working Paper 10297

[5] Ejarque, J. M., and Portugal P. (2006), "Labor adjustment costs in a panel of establishments", mimeo, University of Essex.

[6] Hamermesh, D. S. (1989), "Labour demand and the structure of adjustment costs", American Economic Review 79(4), 674–689.

[7] Hamermesh, D. (1992), "A general model of dynamic labor demand", Review of Economics and Statistics 74(4), 733–737

[8] Hamermesh, D. (1995), "Labor demand and the source of adjustment costs, Economic Journal 105, 620–634.

[9] Hamermesh, D., and Pfann, G. (1996a), "Adjustment costs in factor demand", Journal of Economic Literature 34, 1264-1292.

[10] Hamermesh, D., and Pfann, G. (1996b), "Turnover and the Dynamics of Labour Demand", Economica 63, 359-367

[11] Ingram, B. F., and Lee, B. (1991), "Simulation estimation of time series models", Journal of Econometrics 47, 197-205

[12] Michaud, M.-L., and Kramarz, F. (2004), "The Shape of Hiring and Separation Costs", IZA DP 1170

[13] Nickell, S. J. (1986), Dynamic Models of Labour Demand, in O. C. Ashenfelter and R. Layard (eds.) Handbook of Labour Economics vol. 1 473-522, Amsterdam: North-Holland Press

[14] Nilsen, Ø. A., Schiantarelli, F., and Salvanes, K. G. (2007), "Employment changes, the structure of adjustment costs, and plant size", forthcoming European Economic Review

[15] Oi, W. Y. (1962), "Labor as a Quasi-fixed Factor", Journal of Political Economy 70(6), 538-555

[16] Portugal, P., and Varejao, J. (2003), "Employment dynamics and the structure of labor adjustment costs", Bank of Portugal WP 12-03

[17] Portugal, P., and Varejao, J. (2004), "Matching workers to jobs in the fast lane: the operation of fixed term contracts", Bank of Portugal Working Paper 10-04

[18] Rota, P. (2004), "Estimating labour demand with fixed costs", International Economic Review 45(1), 25-48

[19] Tauchen, G. (1986), "Finite state Markow-chain approximation to univariate and vector autoregressions" Economics Letters 20, 177-181

[20] Varejao, J. (2003), "Job and worker flows in high adjustment cost settings", Portuguese Economic Journal 2, 37-51

Appendix 1: Simulated Frictionless Model

The first object we will look at is the distribution of observations of net and gross adjustment for the panel. Quits have no impact on net adjustment since the frictionless problem is a static optimization problem. The quit shock (z) has, for all cases in this section, the following distribution of observations on $[0, 1, 2]$: $[6831, 6588, 6581]$. The distributions for quarterly frequency are:

	ΔL	M	Q	F	EN	EX
<-5	7549	6075	-	-	-	-
-5	911	883	-	-	-	-
-4	61	897	-	-	-	-
-3	12	582	-	-	-	-
-2	21	84	-	-	-	-
-1	15	25	-	-	-	-
0	**3000**	514	3294	11454	**9060**	2491
1	12	1038	6767	25	1038	3920
2	14	975	6480	84	975	3691
3	15	**516**	**3459**	582	**516**	**1377**
4	63	22	0	897	22	61
5	851	311	0	883	311	911
>5	7476	8078	0	6075	8078	7549

The shock distribution by itself induces 3000 observations (15%) of zero net changes, with few small net changes in employment. The distribution of net and gross changes has **long** tails. Here, because we aggregate all variations above 5 into one statistic, they look like **fat** tails, but that would be an inaccurate description for our exercise.

Appendix 2: Data Appendix

In addition to the sequences of inaction, there are some interesting indicators are: the number of observations where ΔL_t, Entry, and Exit are simultaneously zero is 312, and of course the number of observations where $\Delta L_t = 0$ because entry equals exit is 248.

There is one important part of the data which is modelled exogenously, namely the stochastic part of quits. Here we show a truncated distribution of exit, conditional on positive entry. In the model, the only exit that can occur simultaneously with entry is voluntary exit. This distribution thus gives us a flavour of the distribution of quits. There are 1685 observations:

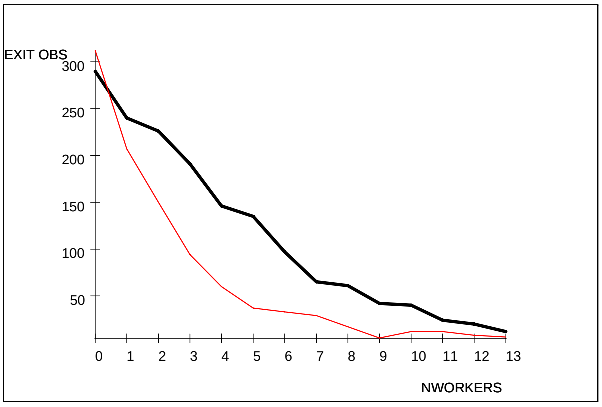

Exit given positive entry										
Number	0	1	2	3	4	5	6	7	8	9
Density	290	240	226	191	146	135	97	65	61	42
Number	10	11	12	13	14	15	16	17	18	19
Density	40	24	20	12	15	12	12	6	11	6

Under the null that the model is correct these numbers should be quits. However, this is at best a biased distribution of quits. For larger numbers of quits it is natural to observe some simultaneous counteracting entry. This implies the distribution above underestimates the frequency of small numbers of quits. If we look at what happens given zero entry we see exactly this:

Exit given zero entry										
Number	0	1	2	3	4	5	6	7	8	9
Density	312	207	150	94	60	37	33	29	17	5
Number	10	11	12	13	14	15	16	17	18	19
Density	12	12	8	6	1	5	4	1	3	1

We use the simple *average* of the two tables to construct our density, on the first 12 numbers (from 0 to 11). This accounts for 0.92% of all "quits".[21] Note that we must model quits somehow or else the model would always fail. Introducing iid random quits is as simple a model as we can have.[22]

[21] The 290 observations of zero exit given positive entry, and the 312 observations of zero exit given zero entry account respectively for 17.21% and 30.74% of observations.

[22] We later test our modelling of quits by looking at the correlation between current period *exit* and previous period employment levels. This correlation has a cross sectional average of 0.3318 and standard deviation of 0.2942. The **long version** of this correlation which uses all the panel in a single measure has the lower value 0.1971.

Appendix 3: Dynamic characteristics of actual and simulated data.

Here are some results of OLS regressions of an AR2 process for ΔL and Entry, for the simulations we used above in the text to illustrate behavior.

OLS: IEE quarterly data				
X_t	c	X_{t-1}	X_{t-2}	R^2
ΔL	-0.9625 (7.17)	-0.1066 (5.34)	-0.0277 (1.41)	0.0120
Entry	1.9952 (15.5)	0.1476 (7.83)	0.1385 (7.36)	0.0560

OLS. Artificial data. Quadratic Costs.				
	Costs of Gross		Costs of Net	
X_t	ΔL	Entry	ΔL	Entry
c	-0.027 (4.3)	0.307 (36.3)	-0.011 (4.1)	0.827 (58.1)
X_{t-1}	$+0.095$ (13.1)	0.503 (72.3)	0.416 (58.6)	0.245 (34.4)
X_{t-2}	$+0.055$ (7.6)	0.280 (40.2)	0.205 (28.9)	0.193 (27.1)
R^2	0.0131	0.5291	0.3061	0.1255

OLS. Artificial data. Fixed Costs.				
	Costs of Gross		Costs of Net	
X_t	ΔL	Entry	ΔL	Entry
c	-0.027 (0.3)	5.219 (68)	-0.026 (0.3)	5.467 (73)
X_{t-1}	-0.418 (59)	-0.185 (25)	-0.414 (59)	-0.166 (23)
X_{t-2}	-0.257 (36)	-0.097 (13)	-0.249 (35)	-0.071 (10)
R^2	0.1697	0.0375	0.1553	0.0289

These numbers also have identifying properties. In the quadratic costs simulation we see that entry and ΔL move in opposite directions with the cost structure. The simulation with fixed costs does not provide any variation for the move between gross and net adjustment costs. However, there is a strong variation of fixed relative to quadratic costs. Finally, there is the absolute value of both the coefficients and the R squared, which also provide identification.