DISCUSSION PAPER SERIES

IZA

IZA DP No. 3091

Labor Adjustment Costs in a Panel of Establishments: A Structural Approach

João Miguel Ejarque
Pedro Portugal

October 2007

Labor Adjustment Costs in a Panel of Establishments: A Structural Approach

João Miguel Ejarque
University of Essex

Pedro Portugal
Bank of Portugal, Universidade Nova de Lisboa and IZA

Discussion Paper No. 3091
October 2007

IZA

P.O. Box 7240
53072 Bonn
Germany

Phone: +49-228-3894-0
Fax: +49-228-3894-180
E-mail: iza@iza.org

IZA Discussion Paper No. 3091
October 2007

ABSTRACT

Labor Adjustment Costs in a Panel of Establishments: A Structural Approach*

This paper estimates a structural model of the employment decision of the firm. Our establishment level data displays an extreme degree of rigidity in that employment levels are largely constant throughout our sample. This can be due to the fact that establishments face large shocks but also large adjustment costs, or alternatively that they incur no adjustment costs but that shocks are negligible. Given our identifying assumptions, we find that rigidity is due to adjustment costs and not to the shock process. We further find that these costs reduce the value of the firm as much as 5%. Finally, small fixed costs of adjustment have a large impact on entry and exit job flows.

JEL Classification: C33, C41, E24, J23

Keywords: adjustment costs, employment, rigidity

Corresponding author:

Pedro Portugal
Banco de Portugal
Av. Almirante Reis 71
1150-165 Lisboa
Portugal
E-mail: pportugal@bportugal.pt

* We thank Russell Cooper, Dean Corbae, Martin Browning, Mette Ejrnaes, and seminar participants at the University of Copenhagen and the University of Texas at Austin.

1 Introduction

This paper uses a standard neoclassical model of the decision of the firm to estimate adjustment costs in a panel of firms operating in markets with large rigidities. Using an explicit structural approach allows us, conditional on the model, to disentangle the contribution of the structure of adjustment costs from that of the stochastic process driving the fluctuations in its employment.

Our model is close to that of Cooper, Haltiwanger and Willis (2004) and is constructed to explain employment observations from a panel of small establishments operating in Portugal. The Portuguese labor market is an interesting case study because of its high levels of job protection which result in very high firing costs. Collective dismissal rules involving a substantial amount of red tape apply to the dismissal of as few as two or five employees depending on the size of the establishment. These and other similar rules are responsible for the fact that the Portuguese labor market emerges as the most regulated in Europe in all existing rankings of indexes of employment protection (e.g., OECD, 1999; Varejão and Portugal, 2007). In this sense, the evidence we present and the estimates we obtain may be thought of as upper-bounds for the corresponding results in other countries.[1]

One key observation from this data is that of extreme levels of inaction, defined by firms having the same number of workers in many consecutive periods. Indeed, at quarterly frequencies, the magnitude of job flows in Portugal is much less than in other countries. Identification of what exactly lies beneath this inaction is fundamental for the understanding of labor market rigidity. This translates into the identification of the parameters of a structural model of the employment decision of the firm. To put it simply, a firm may decide not to change its labor force because its demand does not fluctuate, or because even though demand fluctuates, it has large costs associated with hiring and firing. The policy implications of these two alternatives are of course very different.

We solve the optimization problem of our artificial firm and use a simulated method of moments approach to the estimation of the structural parameters we are interested in. Our estimated model does a good job of matching the data and we find that quadratic adjustment costs impose a loss of about 5% of the value of the firm, while fixed costs of adjustment impose a loss of 0.65% of the value of the firm. More importantly, while the loss of value is moderate, the impact of these costs on the dynamics of the model is very large, as we can gauge by simulating the estimated model shutting down adjustment costs. This also shows that the observed rigidity is not likely to be driven by the shocks, but rather due to the costs of adjustment.

[1] For a comparison with the U.S., see Blanchard and Portugal (2001).

Our results are in line with those presented by Hammermesh (1989) who uses high frequency data to find that firms display a high level of rigidity in response to small shocks, but quick adjustment to large shocks in the sense that in our estimated model, inaction following large shocks is more frequent after larger than that following small shocks. In a broader sense, our results are consistent with the notion that, with stringent firing costs, labor adjustment is shaped by permanent product demand shocks, but not by transitory ones (Blanchard and Portugal, 2001).

2 Data

The data used in this study comes from the Inquérito ao Emprego Estruturado (IEE) which is a quarterly survey with detailed information on job and worker flows at the establishment level. The surveys is run by the Portuguese Ministry of Employment. Establishments of all sizes and in all industries are included in the IEE dataset.[2] The sample is drawn from the universe of the respondents to the 1990 spell of the Quadros de Pessoal, an annual survey which covers all establishments with at least one wage earner. Micro and quarterly (or more frequent) data are essential for studying the dynamics of factor adjustment because aggregation (spatial or temporal) smoothes away any nonconvexities present at the plant or firm levels (Hamermesh, 1993).

The probability of units with fewer than 100 employees being selected to the IEE sample is inversely related to the size of the establishment. Above that threshold, establishments are selected with certainty. The sample is statistically representative for three-digit industries (as defined by the SIC code), region and size class. For this purpose, seven regions - five in mainland Portugal and the islands of Madeira and the Azores - were considered and six size classes were defined. The IEE data used here span over twenty quarters, from the first quarter of 1991 until the last quarter of 1995. On average, 6,954 establishments respond to each spell of the survey.

A detailed description of this data can be found in Varejão (2003). We divide our panel in size categories determined by the initial size of the firm - the size of the labor force in the firm's first reported observation. In this paper we limit our attention to very small firms, those that report an initial number of workers of ten or less, and use only the firms that report all the 20 consecutive quarters of information. This delivers a sample of 538 firms and 10760 observations.[3] The main reason why we focus on small establishments is because of the large rigidity

[2] Only Agriculture, Fisheries, Public Administration and Household Services are excluded.

[3] An appendix, available upon request, details the data treatment and the characteristics of the sample.

we observe in them. In particular they should provide an extreme case of the importance of fixed costs of adjustment. In addition, by focusing on very small establishments we reduce dramatically the computational burden of our estimation procedure. We are convinced, nevertheless, that the scope of our analysis can be extended, *mutatis mutandis*, to larger establishments.

3 Empirical Moments

As in any other simulated method of moments exercise, one needs to choose carefully (revealing) identifying empirical moments to be matched by the model.

3.1 Relevant moments from the size distribution

The first set of observations we are interested in are the cross sectional averages and cross sectional standard deviations of three firm specific moments - mean, standard deviation, and first order serial correlation - of the *stock* of employment:

Sample moments	$\mu^{cs}(\mu)$	$\sigma^{cs}(\mu)$	$\mu^{cs}(\sigma)$	$\sigma^{cs}(\sigma)$	$\mu^{cs}(\rho)$	$\sigma^{cs}(\rho)$	$\rho^{cs}(\mu,\sigma)$
$N = 538$	**5.042**	2.769	**0.795**	0.897	**0.711**	0.275	0.4536

The mean (5.042) is used to center our model - a calibration moment - so is not a moment used in the estimation exercise. The average standard deviation (0.795) and serial correlation (0.711) are moments used to help estimate the properties of the stochastic shock in our model (as well as the other parameters, of course). The cross sectional covariances between firm level means and firm level standard deviations, $\rho^{cs}(\mu,\sigma)$ is a useful moment to describe size implications.

3.2 The action distribution

We want to study the hiring and firing behaviour of our firms. In our sample, the graphical depiction of the distribution of the $19 \times 538 = 10222$ (quarter-firm) observations, by the size of variation in the labor stock, $\Delta L_t = L_t - L_{t-1}$, is:[4]

[4]The number of observations for $\Delta L = [< -3, -3, -2, -1, 0, 1, 2, 3, > 3]$ are respectively $[25, 42, 144, 885, 8160, 771, 136, 33, 26]$.

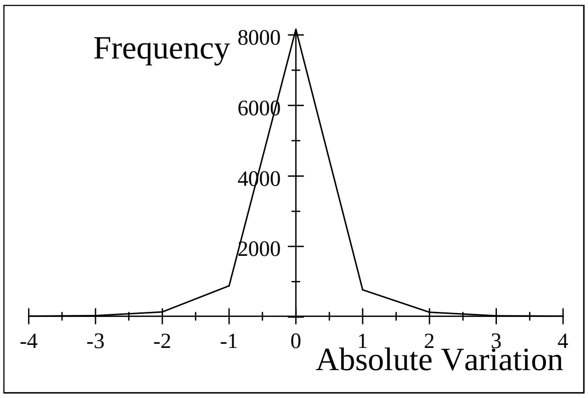

and one key property of this sample is the very high percentage of observations where employment remains unchanged (**79.8%**). The other important property of this sample is the fact that most non zero observations are of plus or minus one single worker, along with the smooth declining nature of the distribution. This may suggest that fixed costs may be of lesser importance. One final characteristic of this curve is the asymetry between -1 and +1. This may be due to the fact that aggregate employment falls over our sample.[5]

3.3 The inaction distribution

We also want to take advantage of the knowledge of how long and how often firms remain inactive. We look at episodes of inaction. An episode (a sequence) of inaction is the consecutive number of periods (quarters) a firm has the same number of workers. The histogram of inaction sequences is:

Inaction duration	1	2	3	4	5	6	7	...	17	18	19
Number of spells	**458**	**296**	188	117	94	68	73		18	15	79

This will prove to be useful information to help gauge the importance of adjustment costs and the frequency of the shocks.

3.4 Job flows

Job flows are standard measures in Labor Economics. Here we measure the observations of adjustment, $\Delta L_t \neq 0$, times the size of adjustment, when positive and when negative (providing the total number of *jobs created or destroyed*), and then just report the count of total number of observations, the total number of zeros,

[5] If we take the subset of firms that grow over the sample, they have more observations of $\Delta L = +1$, than observations of $\Delta L = -1$.

the total number of positives (observations of job creation) and the total number of negatives (observations of job destruction).

Number of firms	*Job creation*	*Job destruction*	N	$Zero$	$+\Delta L_t$	$-\Delta L$
538	1322	1448	10222	8160	966	1096

Here, approximately 12% of observations have job creation, and 15% have job destruction.

3.5 Worker flows

One additional information potentially revealing is the measure of entry and exit worker flows. Sometimes - not often - the stock of employment remains unchanged but there was actually entry and exit. Here we report the number of observations where entry is zero (8902), where exit is zero (8751), and where all $(\Delta L_t, E, X)$ are zero.

Number of firms	$Job\ flows = 0$	$Entry = 0$	$Exit = 0$	$Entry = Exit = 0$
538	8160	8902	8751	7874

How much of our inaction in the stock is then associated with inaction in flows? The answer is 7874/8160 = 0.965, or over 96 percent.

What about non zero entries and exits? As we can see around 96.7% of observations of entry and exit are either of zero or one worker:

538	0	1	2	3	> 3
entry	9325	**1092**	**238**	59	46
exit	9180	**1225**	**240**	74	41

3.6 Aggregate employment

Over the observation period, aggregate employment was declining as seen in the graph:

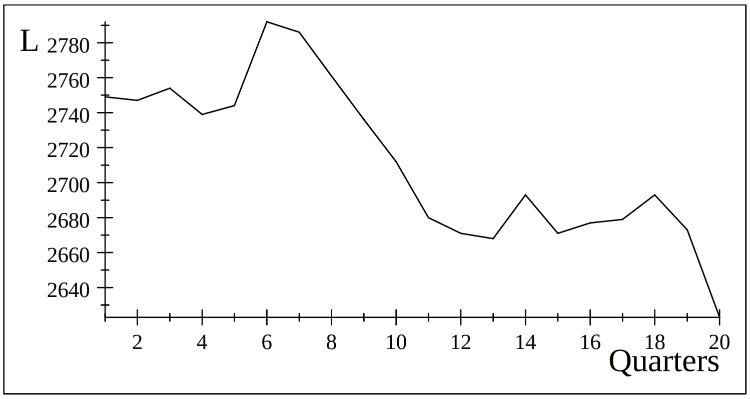

4 The model

We now propose a model of the employment decision of the firm to try and capture the rigidity as well as the action described above. The model follows Cooper, Haltiwanger and Willis (2004). Capital and other inputs are assumed to be maximized away. The profit function of the firm is given by

$$\Pi_t \left(A_t, \delta_t, L_{t-1}\right) = A_t L_t^{\alpha} - wL_t - C\left(L_{t-1}, L_t\right) - F\left(M_t \neq 0\right)$$

where revenues $A_t L_t^{\alpha}$ are concave in labor, $0 < \alpha < 1$. Employment obeys the law of motion $L_t = L_{t-1} - \delta_t + M_t$, where δ_t is the number of quits.[6] Accounting for quits allows the establishment to downsize via two distinct routes: firing workers and incurring in labor adjustment costs or waiting for the worker to leave voluntarillly, avoiding firing costs.[7] Quits is a non negative iid random variable defined over a small set of integers $[0, 1, 2, ..., \delta^{\max}]$. Finally, M_t is the net change in the stock of labor which can be negative.

The shock A_t - a demand or technology shock - has two components which we label persistent and transitory, $A_t = A_t^P A_t^T$. The transitory component follows a first order stochastic process, $A_t^T = \exp\left\{x_t^T\right\}$, with $x_t^T = \rho x_{t-1}^T + \epsilon_t$, which is discretized into a first order Markov process using Tauchen's method. The permanent component is a persistent Markov process aimed at matching cross sectional dispersion, and which we detail below.

[6] Quits are here defined in a broad sense. In essence, quits in the model account for natural friction such as that generated by retirements, voluntary exits, etc.

[7] This second option appears to be of relevance in high job protection countries such as Italy, Spain, or Portugal (Bover et al., 2000).

The firm faces quadratic costs of adjusting its labor stock:

$$C\left(L_{t-1}, L_t\right) = \frac{\gamma}{2}\left[1 + \frac{L_{t-1} + L_t}{2}\right]^{-1}\left[L_t - L_{t-1}\right]^2$$

and fixed costs

$$F\left(M_t \neq 0\right) = F^+\left(M_t > 0\right) + F^-\left(M_t < 0\right).$$

The firm pays fixed costs if it hires or fires anyone, but adjustment costs at the floor level only happen if it has less or more people to work with than it did the previous period, and so it has to adapt its processes to a different number of bodies. Fixed costs are split in two, distinguishing the events of hiring and firing. We interpret fixed costs of firing as being mainly regulation induced, whereas fixed costs of hiring are likely to be related to training and disruption costs.

The quadratic costs specification is adapted from the investment literature. The difference between what matters for fixed costs versus quadratic costs introduces a trade-off in the employment decision of the firm. On one hand the firm wants to have the same number of bodies as in the previous period not to disrupt the work flow which carries the quadratic costs, but on the other hand, as technology (or demand) shocks induce it to hire or fire workers, fixed costs induce large adjustments conditional on taking action.

The cost structure used here, with fixed costs extending the usual quadratic cost framework is known in the literature on labor and capital adjustment. Irrespective of its particular interpretation above, the main property is that it is standard and reasonably general given its fixed and quadratic components, and in this paper it is enough to describe the data very closely.

There is no time to build in the model. Even though labor is a state variable, and therefore a stock variable, new hires become productive immediately. We can think of a model where new hires go through a training process and only become productive in the following period, but we do not go that route here. Regarding the information set, the firm observes $(L_{t-1}, A_t^P, A_t^T, \delta_t)$, and then decides on its current labor force L_t.

This formulation of the problem encompasses shutdown and reopening of the establishment. The state space for employment in a set of non-negative integers and therefore includes zero. The costs of closing down and reopening the establishment are treated as any other adjustment costs of employment. In this sense, because the distribution of shocks is ergodic, the firm never dies in the model. In the data the firm exists already and does not die, and in the model, entry into the market is not defined.[8]

[8] It is possible to model entry and exit explicitly and then model the sample selection procedure that generates the actual data, but that would be another paper.

5 Calibration

We need to select some moments that will prove useful to help the model match the data both in the time series and the cross section dimension. Some parameters of the model are calibrated and some are estimated.

The **number of quits** is an iid random variable on the set of integers $\delta \in [0, 1, 2]$, with discrete space density $p(\delta) = [0.96, 0.03, 0.01]$. This has an expected value of around 0.05 quits per quarter, which *helps* to match the evidence that there is very little movement in entry and exit. Note that we cannot identify quits in the data. Exit in our model is given by

$$EX_t = \delta_t - M_t \times (M_t < 0)$$

and entry is given by $EN_t = M_t \times (M_t > 0)$, where δ_t are quits and M_t are net hires or fires. Note that this is the correct definition of exit for the model to match the data. In the data entry and exit can occur at the same time whereas that is ruled out in the model. Our data shows that most observations have zero exit, which is replicated by this process. About 99% of observations of *exit* are of zero, one or two workers. Therefore our state space for *quits* is limited to $[0, 1, 2]$.

The **discount rate** for the profits of the firm is 1.5% to match our observations of quarterly data. We have no reliable data on profitability of these establishments, which is information necessary to pin down alfa, so we fix the **wage** arbitrarily at $w = 0.5$, and we assume a value of $\boldsymbol{\alpha} = 0.7$ which is a curvature found in similar models and in the investment literature.[9]

Our data has a strong cross sectional dimension. The average number of workers for firms with up to 5 initial workers is 3.1576, and the average for the firms with 6 to10 initial workers is 7.5104. There are respectively 315 firms with up to 5 workers and 223 with 6 to 10 workers in their first reporting quarter, totalling 538 firms.

Employment means cannot be used for *estimation* as the state space of the model is *calibrated* using these values. We model the **size distribution** through the persistent shock (a quasi fixed effect) using a Markov process with four points in the support, $[2, 5, 8, 10]$, and symmetric transition matrix Π^{FE} with 0.985 in the diagonal. The support of the model is constructed by imposing these four center points for L, and then inverting the deterministic steady state expressions to impose the four "centers" for the conditional expectation of the A shock.[10]

We need also to account for the **aggregate decline in employment**. We do this via an additional Markov process with support $TR = [0.5, 0.75, 1]$, and

[9] Cooper, Haltiwanger and Willis estimate α to be between 0.66 and 0.86.

[10] We center the model by selecting $\bar{A}$ so that $\bar{L}$ is true in the deterministic steady state. When other parameters change, $\bar{A}$ changes but $\bar{L}$ does not.

transition matrix

$$\Pi^{AT} = \begin{bmatrix} 0.999 & 0.0005 & 0.0005 \\ 0.01 & 0.99 & 0 \\ 0.001 & 0.009 & 0.99 \end{bmatrix}$$

These two markov processes interact to generate our data. In the estimation part, we use ten panels of artificial data of the same size as our actual data panel, and average our resulting moments. But each of our artificial data panels is generated initializing 172 firms at $L = 2$, 202 firms at $L = 5$, 132 firms at $L = 8$ and 32 firms at $L = 10$, which is an approximate reproduction of our initial distribution of firms in the data. All our firms also start their time series at the high state in the TR process so that this process then simply generates the evolution of aggregate employment.[11]

The discrete state space for L is a cut from the natural numbers from 0 to 50, reflecting our data which has information on the number of workers in a firm, and not on hours worked. The number of points in the state space for the *conditional* (transitory) markov process of A is set at 11.

6 Matching moments

The remaining parameters, $(\rho, \sigma, \gamma, F^-, F^+)$, are estimated to get the model to fit the data as closely as possible, minimizing a quadratic distance measure explained below. But what data do we want to match? Here we focus on the following vector of moments, X:

Sample Moments	$\mu^{cs}(\sigma)$	$\mu^{cs}(\rho)$	$\Delta L = -1$	$\Delta L = 0$	$\Delta L = +1$	$I(1)$	$EX(1)$
Mean	0.795	0.711	885	8160	771	458	1225
Std deviation	0.039	0.012	33.275	73.943	31.504	27.300	44.356

As conventional, the first two moments are the cross sectional *mean* of firm specific first order standard deviation (0.795) and firm specific serial correlation (0.711) of the *stock* of labor. These are measures of the standard deviation and serial correlation of the *conditional* shock process. These measures isolate from the effect of heterogeneity across firms. A measure of such heterogeneity is contained in cross sectional standard deviations in these same firm specific moments. However, we have already calibrated the heterogeneity, and so we focus on the moments and parameters we have yet to identify.

[11] This is still an idiosyncratic shock, but the matrix is skewed so that there is more density at the bottom in the long run. This is a pragmatic solution to the problem posed by the aggregate employment movement.

The third, fourth and fifth moments are used to capture the dynamics of ΔL_t. The number of zeros (8160) are expected to track the magnitude of both fixed costs and quadratic costs.[12] The sixth and seventh moments are added as overidentifying restrictions. They contain potentially different information from the previous ones. 458 is the number of times we have single periods of inaction (the first element in the inaction distribution): of the 8160 periods of zero adjustment, 458 of those are preceded and succeeded by action. 1225 is the number of observations where exit is just one worker.

6.1 Model Estimates and bootstrapping.

The artificial numbers on our moments are averages of 10 replications with artificial data. We use bootstrapping on the actual data to generate the variance matrix of the moments. Bootstrapping is done by sampling firms, not observations. So, when one firm is picked, all 20 observations of that firm enter the bootstrap draw.[13]

Estimation is done using the simulated method of moments procedure following Ingram and Lee (1991). We search over $\theta = (\rho, \sigma, \gamma, F^{+}, F^{-})$ to find the combination that minimizes a quadratic form, Q. For a given parameter vector θ_0 we generate artificial data and compute the artificial counterpart $Y(\theta_0)$ of our actual data moments, X. Then we estimate the model by choosing the θ that minimizes $Q(\theta)$:

$$Q(\theta) = \left(1 + \frac{1}{k}\right)\left(\bar{Y}(\theta) - X\right) W^{-1} \left(\bar{Y}(\theta) - X\right)'$$

where

$$\bar{Y}(\theta^*) = \frac{1}{k} \sum_{j=1}^{k} Y_j(\theta^*)$$

and where $k = 10$ is the number of replications with artificial data. The weighting matrix is the inverse of the bootstrap matrix

$$W = \frac{1}{1000} \sum_{j=1}^{1000} \left[X_j - \bar{X}\right]\left[X_j - \bar{X}\right]'$$

The outcome of this quadratic form, $Q(\theta^*)$, has a chi squared distribution with degrees of freedom equal to the number of elements in X minus the number of parameters to be estimated, and this difference is two. The 5% confidence level

[12] In the absence of fixed costs, depreciation would be irrelevant, and any positive or negative action would be a result of shocks. As it is, asymetric fixed costs can play a part, but once we account for the fall in aggregate employment, this fixed cost asymetry is largely irrelevant.

[13] We construct a variance matrix for these moments by bootstrapping with one thousand draws.

stands at the value of 5.99.[14] This then provides an overidentifying restrictions test on the model.

Finally we need to compute the variance covariance matrix of the estimator. This is given by the following construction:

$$V = \left(1 + \frac{1}{k}\right) B'W^{-1}B, \qquad \text{where } B = \frac{\partial \bar{Y}(\theta)}{\partial \theta}$$

and in this way, the estimator has asymptotic distribution given by

$$\sqrt{T}\left(\hat{\theta} - \theta\right) \rightarrow N\left(0, V^{-1}\right)$$

The calculation of the matrix of derivatives B is done numerically with 1% deviations in each parameter, after we have converged in our estimator.

7 Empirical results

7.1 Estimation

7.1.1 Structural parameters

We obtain the following estimates for the structural parameters, (upper bound on) standard deviations, and measures of closeness:

ρ	σ	γ	F^+	F^-	$Qabs$	Q
0.1181	0.1095	0.4760	0.0269	0.0245	0.374	30.54
(0.0002)	(0.0008)	(0.15)	(0.0015)	(0.0006)		

The remarkable characteristic of these estimates is that fixed costs are very small ($F^- = 0.0245$ is about 15% of the monthly wage, or 5% of the quarterly wage which is 0.5). The statistic $Qabs$ is the sum of the absolute value of the percentage deviations between the actual data moments and the artificial data moments used in estimation.

7.1.2 Matching moments

[14] A chi squared with one (two) degree of freedom has a 10% threshold of 2.706 (4.61) and a 5% one of 3.841 (5.99).

Overall, our model does reasonably well in matching our moments, being off between 1.88% and 9.86%, being on average only 5.4% off.[15]

Sample Moments	$\mu^{cs}(\sigma)$	$\mu^{cs}(\rho)$	$\Delta L = -1$	$\Delta L = 0$	$\Delta L = +1$	$I(1)$	$EX(1)$
Data	0.795	0.711	885	8160	771	458	1225
Model	0.827	0.743	803	8346	728	449	1104
Difference (%)	3.99	4.52	**9.27**	2.28	5.62	1.88	**9.86**

A first set of moments illustrates how well the model does:

Sample Moments	$\mu^{cs}(\mu)$	$\sigma^{cs}(\mu)$	$\mu^{cs}(\sigma)$	$\sigma^{cs}(\sigma)$	$\mu^{cs}(\rho)$	$\sigma^{cs}(\rho)$	$\rho^{cs}(\mu,\sigma)$
Data	5.042	2.769	**0.795**	0.897	**0.711**	0.275	0.453
Model	4.980	2.560	**0.827**	0.840	**0.743**	0.285	0.289
Difference (%)	1.23	7.55	3.99	6.35	4.52	3.64	36.2

From a calibration perspective this is a reasonable success and only the cross sectional correlation between means and variances is substantially off. This moment shows that the model still lacks cross sectional variability. But the rest of the moments are a reasonable success. However, from a structural estimation point of view we do not have a success. We have two overidentifying restrictions which puts the Q statistic threshold at around 6, and we get a Q value of 30.

7.2 Implications for non-estimated moments.

7.2.1 The inaction distribution.

The model generates the following distribution of inaction sequences (average of 10 replications):

Inaction Duration	1	2	3	4	5	6	7	...	17	18	19
Data	**458**	296	188	117	94	68	73	...	18	15	79
Model	**449**	274	165	134	95	66	50	...	8	9	117

and we can see the plot of the two series:

[15] What makes this outcome more reasonable is that a broader set of 22 moments is off by 10.94% on average.

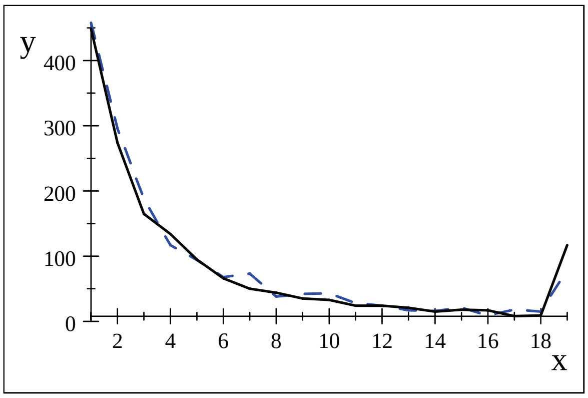

Inaction Sequences

The predicted frequency of duration of inaction is roughly in the ballpark.

7.2.2 The action distribution

The corresponding action distribution is also well determined:

ΔL_t	< -3	-3	-2	-1	0	1	2	3	> 3
Data	25	42	144	**885**	**8160**	**771**	136	33	26
Model	21	31	119	**803**	**8346**	**728**	154	20	0

7.2.3 The distribution of entry and exit

The model does also a good job matching entry and exit, again short of the tail of the distribution. The zeros (first column) are less than 2% off, but the ones (second column) are over 10% off. Nevertheless the pattern is similar to the data. It is also understandable that the tails should be harder to fit simply because they are thinner (have less observations).

Worker flows		0	1	2	3	> 3
Entry						
Data		9325	1092	238	59	46
Model		9491	945	283	39	1
Exit						
Data	entry	9180	**1225**	240	74	41
Model	exit	9372	**1104**	232	31	21

7.2.4 The distribution of quits

Not surprisingly, quits behave according to the Markov probabilities specified above.[16] The number of periods with quits equal to zero, one or two, and total quits, is computed for the average of ten artificial panels as:

Quits	0	1	2	Total
Model	10309	340	112	563

and we see that quits represent only about one third of total exits.

7.3 Welfare analysis

We know that small non-convex costs can change behaviour significantly. The value function allows us to measure the losses associated with them. We take our estimated model, and run a version of it with different adjustment cost parameters set at zero to evaluate the gains from having a flexible labor market. The gain is measured on the value function. The numbers below report the value function gains from eliminating costs measured over entire panel: our 10760 artificial observations generate 10760 values of the value function, $V(A_t, \delta_t, L_{t-1})$, and the values in the table are the averages over these 10760 realizations. The first value is the benchmark without shutting down any adjustment costs and the last column shuts down all the costs.

	V^0	$F^- = F^+ = 0$	$\gamma = F^- = F^+ = 0$
V	52.3594	52.6756	55.2606
$\frac{V-V^0}{V^0}$	0.0	0.60%	5.54%

The biggest gain comes from eliminating quadratic adjustment costs. Fixed costs are somewhat asymmetric but (not shown here) have almost the same small impact on the value of the firm.

But more important is perhaps what happens to the model when these three parameters are set at zero. The number of observations of zero adjustment falls from 8160 to 2654 and there is substantial action outside of +1 and -1. This suggests adjustment costs are fundamental for the understanding of inaction since we already introduced substantial persistence and rigidity in our treatment of heterogeneity, and that component of the model is still present. Going back to our identification exercise mentioned above, *rigidity is a result of costs rather than shocks*: without the adjustment costs most of the rigidity disappears. In fact, even if we shut down only the fixed costs, the number of observations of zero adjustment

[16] Moderate modifications of the quit probabilities do not disturb the results.

falls to 7317, which is a substantial change. This additional action appears as 1263 observations where $L_t - L_{t-1} = -1$, and 1252 observations where $L_t - L_{t-1} = +1$. This is a very significant change in the behaviour of the model. The number of observations with $L_t - L_{t-1} \neq 0$ increases by more than 50% from 10222 - 8346 = 1876 up to 2905 when we shut down the fixed costs alone, and the fixed costs are very small in value, and have an accordingly small impact on the value of the firm.

ΔL_t	< -3	-3	-2	-1	0	1	2	3	> 3
Data	25	42	144	**885**	**8160**	**771**	136	33	26
Model	21	31	119	**803**	**8346**	**728**	154	20	0
$\begin{cases} \gamma = 0 \\ F \neq 0 \end{cases}$	781	460	871	597	4541	465	811	515	757
$F^- = F^+ = 0$	11	31	158	1263	7317	1252	163	19	0
$\gamma = F^- = F^+ = 0$	895	335	1027	1331	2654	1310	1031	348	870

Our benchmark model is the best fit we could obtain. It contains a specification for adjustment costs which stems from the investment literature, in that it attaches quadratic costs to rates of change, and then scales it by a factor to set costs at the same order of magnitude as revenues:

$$\frac{\gamma}{2} L_{t-1} \left(\frac{L_t - L_{t-1}}{L_{t-1}} \right)^2$$

This type of function implies that the cost of adjusting any given number of workers falls with firm size. For larger firms, costs are smaller relative to revenues, even if the adjustment cost function are simply

$$\left(L_t - L_{t-1} \right)^2 ,$$

but with the function above even more so.

We therefore investigated a simpler cost structure, without scaling:

$$\frac{\gamma}{2} \left(L_t - L_{t-1} \right)^2 + \phi abs \left(L_t - L_{t-1} \right) + FC$$

First, we estimated only (γ, ϕ) jointly, keeping all other parameters at the values estimated in our benchmark. We obtained $(\gamma, \phi) = (0.052861, 0.0022781)$, with a J statistic of 57. With this structure, shutting down linear costs - while keeping all other coefficients at the estimated values - generates a J statistic of 59.7, while shutting down quadratic costs generates a J statistic of 20516. Other experiments and estimations with linear costs suggested also that the model with quadratic costs is far superior in matching the data.

Next we estimated this model without linear costs and with a single symmetric fixed cost:

$$\frac{\gamma}{2}(L_t - L_{t-1})^2 + F(M \neq 0)$$

and we obtained $(\gamma, F) = (0.0488, 0.0254)$, with a J statistic of 53.9. This estimated model generated 8098 zeros of ΔL. Shutting down quadratic costs here generated a J statistic of 23526 and 4469 zeros, while shutting down fixed costs generated a J statistic of 2156 - an order of magnitude smaller that 23 thousand, but still quite high - and a number of zeros of 5889.

It is useful to emphasize why quadratic costs generate most of the zeros in our model. It is still the case that fixed costs are the component that generates zeros, whereas quadratic costs generate density *around* zero. However, in this model, because of the discrete nature of the state space, and because of the small size of the establishments we are studying, adjusting a single worker is already a big adjustment. Therefore, the quadratic part of adjustment costs is responsible for most of the zeros. We expect that this peculiarity will fade if we take the model to larger establishments, and that for those, quadratic costs will generate density of adjustment **around** zero, while fixed costs generate density **at** zero. Nevertheless, these counterfactual experiments just shown suggest it is not the fact that we have a sample of small firms that produces these results of very small fixed costs. We saw that the fixed costs, despite being small still have a strong effect on the dynamics of the model. But more than that, it is the fact that the action distribution has a smooth normal looking density that suggests fixed costs are just a part of the story. There are no fat tails in the distribution of ΔL, and fundamentally, there is a smooth declining distribution of net changes in employment, where there are significant numbers of $\Delta L = \pm 1$, and that implies we must estimate small fixed costs. Fixed costs tend to eliminate density at $\Delta L = \pm 1$. This outcome is interesting also because it stands in stark contrast with the findings of Rota (2004), using data which has exactly the same pattern.

It is appropriate to note here that our identification of the impact of shocks versus costs is conditional on the model written (as in any exercise). We could have imposed that there are no adjustment costs and simply allow profits to be the intratemporal quantity $\Pi_t = A_t L_t^{\alpha} - wL_t$. Given our assumed value for wages and alfa, our data on L immediately implies data on the shocks. The exercise then would be to recover the process for the shock. However, this leaves us as ignorant as we started, and the stochastic process we would recover would be something very strange indeed.

7.4 Fitting the model to larger firms

One final exercise we performed was to put the predictive capacity of the estimated model to the test of seeing how well these parameters fare in replicating the behaviour of the sample of firms which report between 11 and 20 workers in their first quarter. We must modify other components of the model, such as the heterogeneity part of the shock process, to accommodate the wider state space. But other than that, everything else remains the same.[17]

The model does a good job of matching a wide variety of moments expect for the correlation between firm mean and firm standard deviations of the stock of employment, $\rho^{cs}(\mu, \sigma)$:

Sample moments	$\mu^{cs}(\mu)$	$\sigma^{cs}(\mu)$	$\mu^{cs}(\sigma)$	$\sigma^{cs}(\sigma)$	$\mu^{cs}(\rho)$	$\sigma^{cs}(\rho)$	$\rho^{cs}(\mu, \sigma)$
Data	14.334	3.994	1.**838**	1.447	**0.713**	0.238	0.162
Model	14.447	3.585	1.949	1.245	0.621	0.237	−0.144

Larger firms have less volatility of L, mostly because adjustment is less costly for larger firms in this functional form. Not surprisingly, this suggests the cost parameters estimated for the smaller firms are less able to deliver the dynamics of slightly larger firms, so that some additional treatment of heterogeneity is needed, perhaps experimenting with making adjustment costs more proportional to the size of the firm is needed to replicate a wider panel. Still, the model does reasonably well.

The model is also reasonably on target matching the number of zeros in the action distribution. But it is a bit off on the rest of the action distribution. The distribution of observations by size of variation in the labor stock, $\Delta L_t = L_t - L_{t-1}$ is:

ΔL_t	< -3	-3	-2	-1	0	1	2	3	> 3
Data	70	67	234	**811**	**3024**	**538**	181	58	71
Model	68	127	421	600	2741	459	444	147	42

7.4.1 Re-estimating the model

One way to shed a better light regarding the effect of firm size on our estimates is to reestimate the model for this subpanel of 266 firms between 11 and 20 workers.[18]

[17] Details of our results can be seen in the appendix, available uppon request.

[18] We changed the trend process to be equivalent to the trend process used in the smaller firms, and we changed the distribution of quits to [0.85, 0.10, 0.05], to adapt to the distribution of exit. These are rule of thumb changes, not estimated parameters, although they have a large impact on the result.

We did that with the following results.

Sample moments	$\mu^{cs}(\sigma)$	$\mu^{cs}(\rho)$	$\Delta L=-1$	$\Delta L=0$	$\Delta L=+1$	$I(1)$	$EX(1)$
Data	1.838	0.713	811	3024	538	498	1128
Model	1.572	0.708	821	3076	538	552	1100
Difference (%)	14.5	0.7	1.3	1.7	0	11	2.5

where once again we do a fair job of matching the data.

The estimated parameters are

Sample	ρ	σ	γ	F^{+}	F^{-}	$Qabs$	Q
11 to 20 workers	0.1202	0.0927	0.5891	0.0285	0.0227	0.317	19.03
1 to 10 workers	0.1181	0.1095	0.4760	0.0269	0.0245	0.374	30.5

and the interesting feature of these parameters is that they are very similar to our earlier estimates. In particular there is no big change in fixed costs, even though with an average firm size almost three times bigger the effects of quadratic and fixed costs must be much better disentangled at zero.[19]

The fit of the action distribution is quite good except for the tails. Again, what we find is that fixed costs are small, and it is the fact that we have smooth distributions, rather than the small size of firms that makes the model estimate small fixed costs. This of course does not mean fixed costs are without impact as we saw above. For this sample, setting fixed costs to zero delivers the following action distribution shown in the bottom row.[20]

ΔL_t	<-3	-3	-2	-1	0	1	2	3	>3
Data	70	67	234	**811**	**3024**	**538**	181	58	71
Model	22	54	285	821	3076	538	230	25	3
$F^{\pm}=0$	21	62	323	1196	2181	1003	238	27	3

8 Conclusion

In this paper we use the simulated method of moments to estimate a simple structural model of the employment decision of the firm. Under reasonable assumptions, this methodology allows us to separately identify the driving shocks from the adjustment cost parameters in our model.

[19] This also has to do with the choice of modelling the fixed costs in terms of net and gross.

[20] Hammermesh (1989) finds that firms display a high level of rigidity in response to small shocks, but quick adjustment to large shocks, and that is replicated by our estimated model: the number of zeros of ΔL following large shocks is substantially larger than that following small shocks.

Our estimated model does a good job of matching the Portuguese employment data. All outcomes are in the ballpark, and, in particular, the model replicates remarkably well the (employment) action and inaction distributions of the establishments.

The model shows that the extreme rigidity of employment outcomes in the Portuguese labor market should be attributed to adjustment costs. We find that different types of adjustment costs together take over 5% out of the value of the firm. Interestingly, eliminating fixed adjustment costs, while producing moderate gains in term of the value of firm, induces very large changes in the behaviour of labor demand. This may suggest that employment protection, as measured by adjustment costs generating low worker flows, can be achieved at a fairly small economic cost. If this trade-off is indeed feasible, it may help to explain why there is such disparity regarding employment protection institutional arrangements among developed countries.

References

[1] Blanchard, O. and Portugal, P. (2001), "What Hides behind an Unemployment Rate: Comparing Portuguese and US Labor Markets," *American Economic Review,* volume 91, number 1, 187-207.

[2] Bover, O. and Garcia-Perea, P. and Portugal, P. (2000) "Labour Market Outliers: Lessons from Portugal and Spain" *Economic Policy*, volume 15, number 31, 379–428.

[3] Brown, B. and Newey, W. (2002), "Generalized method of moments, efficient bootstrapping, and improved inference", *Journal of Business & Economic Statistics*, volume 20, number 4, 507-517.

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

[5] Hammermesh, D. (1989), "Labor Demand and the Strcture of Adjustment Costs", *American Economic Review* 79, number 4, 674-689.

[6] Hamermesh, D. (1993), Labor Demand, Princeton (NJ): Princeton University Press.

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

[8] Rota, P. (2004), "Estimating Labor Demand with Fixed Costs", *International Economic Review* 45, number 1, 25–48.

[9] Varejão J. e Portugal, P. (2005), "Matching workers to jobs in the fast lane: the operation of fixed term contracts," in *What Explains the Pattern of Labour Supply in Europe?* edited by R. Gomez Salvador, A. Lamo, B. Petrongolo, M. Ward, and E. Wasmer, Kluwer Academic Pres.

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

[11] Varejão, J. and Portugal, P. (2007), "Employment Dynamics and the Structure of Labor Adjustment Costs," *Journal of Labor Economics*, 25, number 1, 137-165.