Disentangling employment and wage rigidity*

Luigi Guiso
Chigaco GSB, Univ. di Sassari and CEPR

Luigi Pistaferri
Stanford University, SIEPR and CEPR

Fabiano Schivardi
Bank of Italy and CEPR

Abstract

We use matched employer-employees data for Italy to study the joint response of wages and employment to firm-level shocks. We construct a simple dynamic general equilibrium model of labor demand and supply that allows us to identify separately firing (or internal) and mobility (or external) adjustment costs. We show that the two type of costs cannot be discriminated empirically by looking at labor or wage adjustment separately. Mobility costs have distinctive implications on wage response to firm-level employment changes but they can only be identified with worker-level information on wages. We find that both types of costs are present, but the internal component accounts for a large share of total adjustment costs. Our results are consistent with a labor market where workers are fairly mobile within locations but scarcely mobile across them.

JEL classification numbers: C33, D21, J63
Keywords: Adjustment costs, mobility costs, matched employer-employees data

*We thank Sam Bentolilla, Tito Boeri, Ricardo Caballero, Robert Hall, Peter Hansen, Michele Pellizzari, John Pencavel, Fabio Schiantarelli, Robert Shimer and Ed Vytlacil and seminar participants at the Bank of Italy, Bocconi, Boston College, Johns Hopkins and Stanford for useful comments and suggestions. We are responsible for any mistake. The views expressed here are our own and do not necessarily reflect those of the Bank of Italy.

1 Introduction

The idea that labor markets deviate substantially from the walrasian competitive allocation mechanism has a long history in economics. Indeed, much of the macroeconomic debate on the business cycle originates from it. More recently, differences in the functioning of the labor market have been indicated as one of the main factors behind the diverging economic performances of continental Europe and the US. As a consequence, over the past fifteen years an enormous amount of research has been devoted to understanding the microeconomics of the labor market, focusing on the features that make the exchange of labor services different from other economic transactions, and better characterized in terms of employment relationships in contrast to anonymous, spot exchanges. The modern analysis of the employment relationship has the existence of frictions in the "creation" and "destruction" of employment at the very center of its research agenda.

Alongside theoretical developments, over the last decade the availability of data at the micro level has spurred a number of studies on the costs of adjusting labor, documenting the existence of nonlinearities and of non-convex costs in the adjustment policies of individual units (see Hamermesh and Pfann (1996) for a survey). This body of work, following the seminal theoretical work of Bentolilla and Bertola (1990), is cast in a partial equilibrium framework where wage-taking firms face an infinitely elastic labor supply curve, so that labor adjustments at the firm level can be studied in isolation from wage adjustments.

On the other hand, the literature on search and matching in the labor market (Mortensen and Pissarides 1999) has shown that, in the presence of search frictions, employment and wages are jointly determined even at the level of the firm. In fact, research in the labor literature has documented that wages too respond to firm-level conditions. In particular, recent work based on matched employer-employees data has shown that shocks to the firm are partially transmitted to the compensation of its employees (Bronars and Famulari 2001, Guiso, Pistaferri and Schivardi 2005), in contrast with the hypothesis that worker's compensation is insulated from idiosyncratic changes in business conditions.

This paper argues that the joint consideration of labor and wages response to firm-level shocks can help shed light on the nature of the frictions that characterize the employment relation. In particular, our approach allows for the separate identification of adjustment costs internal to the firm - such as firing and hiring costs - and external to it, i.e. those borne by the workers due to costly mobility across firms, such as for search, geographical mobility and re-training needs. This is clearly an important distinction because such costs have different implications for the functioning of the labor market and, more importantly, for the design of policies aimed at improving it, particularly when it comes to the lifting of obstacles to employment or wage adjustment induced by institutional factors, such as employment protection legislation in Europe.

We adapt a general equilibrium model developed by Bertola (2004) with firing costs on the firm side, mobility costs on the workers' side, and idiosyncratic shocks to labor demand. The model features patterns of adjustments that deviate from the frictionless paradigm in important ways. Most importantly, it shows that both internal and external costs can generate the type of employment response to shocks that have been traditionally associated with firing costs, that is non-adjustment in response to small shocks and lumpiness of labor adjustment. This implies that the interpretation of the results of the previous literature as evidence of internal adjustment costs alone might be unwarranted.

While stylized, the model is flexible enough to allow for the structural identification of the adjustment cost parameters and distinguish between firing and mobility costs. It also allows for a clear and intuitive representation of our identification strategy. The idea behind the empirical test is simple. If mobility costs are important, then an expanding firm will need to compensate the new workers for the mobility cost they bear. As Joan Robinson (1933) pointed out seventy years ago, "*..there may be a certain number of workers in the immediate neighborhood and to attract those from further afield it may be necessary to pay a wage equal to what they can earn near home plus their fares to and fro*". Stated differently, mobility costs imply an upward sloping supply for labor at the firm level. When a firm changes the level of employment, the workers' compensation should also change in the same direction if mobility costs matter, while no change in wages should be observed if the firm faces only firing/hiring costs. We therefore supplement the employment adjustment equations on the extensive and intensive margin previously used in the literature with a wage adjustment equation, which allows us to separately identify internal and external adjustment costs.

The empirical problem with this approach is that it can hardly be implemented with a firm-level measure of the wage, such as the total wage bill divided by the labor force − the standard measure of wage used in the literature. This measure, in fact, is likely to be strongly influenced by changes in employment for reasons that have nothing to do with adjustment costs.[1] To overcome this problem we merge company-level data for a large sample of Italian firms with social security data on worker-level compensation available for a random sample of their employees for the 1982-1994 period. The detailed information at the firm level allows to compute measures of idiosyncratic shocks to the firm and then study the response to these shocks of firm-level employment and *individual compensations* after controlling for workers' and firms' characteristics. The fact that we use idiosyncratic shocks implies that we are abstracting from aggregate events that might change aggregate labor demand: rather, we identify the costs of mobility faced by firms in isolation from aggregate shocks.

We find that total adjustment costs are substantial. According to our preferred estimates, the per capita cost of changing employment is in the order of 13,000 euros, about 75 percent of the average gross annual compensation. This figure is of the same order of magnitude as that estimated by Abowd and Kramarz (2003) for France, a country with a labor market similar to the Italian one.

In terms of external and internal costs, we find that both components are present and statistically significant. The internal component accounts for about 85 percent of the total, indicating that internal costs are a more important impediment to labor adjustment than mobility costs, which are in the order of 1,700 euros. This result is robust to a number of extensions, such as accounting for heterogeneity in mobility costs across workers.

The relatively modest size of mobility costs is consistent with the fact that we consider idiosyncratic shocks to firms: given that we are abstracting from aggregate labor demand

[1] For example, an expanding firm may hire highly skilled workers and thus pay a wage skill premia which raises the average firm wage genrating a correlation between wage and employment adjustment, even with no mobility costs. Moreover, given that the number of employees would appear in the denominator of a measure of average compensation, any measurement error due to the timing with which employment is recorded would induce a spurious correlation between employment and wages adjustment.

changes, an increase in labor demand by a firm will most likely be satisfied on the local market, without resorting to long distance mobility. This hypothesis is supported by the analysis of mobility patterns in our dataset, that indicate that workers' mobility is mostly local. The characterization of the labor market that emerges therefore is one where workers are fairly mobile within locations but scarcely mobile across them.

Our results suggest that mobility costs faced by workers, though less important than internal costs, cannot nevertheless be neglected, implying that the assumption of fixed wages made in the literature is empirically misleading as it tends to overstate the role of hiring and firing costs internal to the firm in the process of labor adjustment. This has an important bearing for the debate on labor market flexibility, stressing the importance of clearly distinguishing between the two types of costs. Moreover, in terms of mobility costs, our findings suggest that a fruitful area for policy intervention would be the reduction of the segmentation across local labor markets rather than that within them.

Our work is connected to two main strands of literature. The first is the growing literature on estimating dynamic labor demand functions with non convex adjustment costs. Within this line of work, our paper is closer to Alonso-Borrego (1998) and Nilsen, Salvanes and Schiantarelli (2003) who, like us, account for the endogenous selection of the adjustment regime. They, however, ignore wage responses to shocks. The second is the literature on wage responses to labor demand shocks. Topel (1986) is the first to use this approach to identify mobility costs across local labor markets, defined in terms of US states. Our approach is similar to his, but more microeconomic in nature, as we conduct the analysis at the level of the firm: in fact, our focus is on the distinction between mobility on one side and firing and hiring costs on the other. More recently, Belzil (2000) also uses firm adjustment in a wage equation, finding that wages are correlated with measures of job creation and destruction at the plant level. Also related to our approach is Buchinsky, Fougère, Kramarz and Tchernis (2003), who estimate a wage equation together with a mobility and a participation equation, but they are interested in obtaining an unbiased estimate of the return to seniority rather than disentangling the nature of adjustment costs. To our knowledge, we are the first to use employment and individual-level wage data to jointly estimate wage and employment responses to firm-specific shocks to identify the nature of adjustment costs.

The layout of the paper is as follows. Section 2 presents the institutional framework, while Section 3 introduces a simple general equilibrium model based on Bertola (2004). Section 4 details the data and Section 5 discusses estimation issues. The results are reported in Section 6, along with sensitivity analysis. Section 7 concludes.

2 Institutional aspects

Following the literature on adjustment costs at the firm level, we do not directly measure costs of hiring, firing and mobility, but rather infer them from the observed responses of employment and wages to shocks. Such costs to depend on the institutional features of the labor market: in fact, as other continental European countries, Italy features a fairly regulated labor market. We thus offer a brief sketch of its main institutional features.

According to Italian employment protection legislation (EPL), individual and collective dismissals of workers with open-end contracts are only allowed on a just cause basis. Workers can be fired for misbehavior (*giusta causa o giustificato motivo soggettivo*), or because of the firm's need to downsize or reorganize its activities (*giustificato motivo oggettivo*). Thus, it would not be possible to fire an employee with a long tenure and a high salary to replace her with a young worker paid the minimum contractual wage.

Workers can appeal in court against dismissal. Firing costs are nil when a dismissal is not contested or it is ruled to be fair, although firms may want to pay some form of compensation to the dismissed workers in order to avoid litigation (this is especially true in collective dismissals, when lump-sum payments are sometimes explicitly bargained with the unions). If the judge rules in favor of the worker, she is entitled to compensation that varies according to firm size. Firms with less than 16 employees must compensate unfairly dismissed workers with a severance payment that varies between 2.5 and 6 months of salary (*tutela obbligatoria*). Firms with more than 15 employees[2] have to compensate workers for

[2]More precisely, the rule refers to establishments with more than 15 employees, and to firms with more than 15 workers in the same municipality or with more than 60 employees in all establishments combined.

the loss of earnings from the date of the dismissal to the date of the ruling. Moreover, they are obliged to reinstate the worker, unless he or she opts for a further severance payment equal to 15 months worth of salary. Because of these differences, the costs of EPL have been traditionally thought to be substantially larger for firms above the 15 employees threshold. Recent studies that exploit the differential effects of EPL on the propensity to grow of firms just below the threshold have found significant but modest effects (Borgarello, Garibaldi and Pacelli 2003, Schivardi and Torrini 2004), an indirect evidence that the differential effects of EPL on small and large firms might be overstated.

In terms of wage setting, Italian industrial relations are based on multi-tier collective bargaining, with economy-wide, industry-wide and company-level agreements. In Guiso et al. (2005) we show that the latter provide sufficient room for wages potentially to respond to idiosyncratic firm shocks. According to data from the Bank of Italy survey on manufacturing firms with at least 50 employees, approximately 92% of workers were covered by a firm-level contract in 1994. Data for the Metal products, Machinery and Equipment sector, for which a breakdown of the wage bill into its various components is available, show that between one sixth and one fourth of the compensation was firm specific in the period covered by our sample (1982-1994). There is therefore room for wages to have an important firm specific component, possibly related to the firms' needs to attract or expel workers following shocks to labor demand.

There is a widespread consensus that geographical mobility in Italy is low because of high moving costs. For example, according to a 1995 survey of the National Institute of Statistics, more than 40% of unemployed were unwilling to take a job outside the municipality of residence and only 22% we ready to move anywhere (Faini, Galli and Rossi 1996). In fact, high unemployment rate in the South has persisted in the face of basically full employment in the rest of the country - that is large unemployment differentials persist due to low mobility rates. Most of this anecdotal evidence refers to geographical mobility, not mobility across firms. Since we focus on idiosyncratic shocks to firms, the most relevant concept of mobility for our exercise is across firms rather than across geographical areas. In other

The different provisions according to firm size are the subject of the hotly contested Art. 18 of the "*Statuto dei Lavoratori*".

words, in a certain local labor market firms that have received idiosyncratic positive shocks may coexist with firms that have received negative shocks, implying that long distance mobility on the workers' side will not be necessarily observed.

3 The Model

We adapt a general equilibrium model by Bertola (2004). Time is discrete. The economy is comprised of a continuum of infinitely lived firms and workers. Firms produce output using a decreasing return to scale technology with labor as the only input and stochastic productivity (or demand) shocks and face costly labor adjustments; we only consider idiosyncratic firms' shocks, i.e. shocks that do not change aggregate productivity and labor demand. Workers supply one unit of labor inelastically; they can pay a mobility cost c and move to a different job, in the spirit of Lucas and Prescott (1974) island model.

The main simplifying assumption is that productivity at the level of the firm switches between two values, $\varepsilon_g > \varepsilon_b$, following a symmetric first order Markov process: $\Pr\{\varepsilon' = \varepsilon_i | \varepsilon = \varepsilon_i\} = p > \frac{1}{2}$, $i = g, b$. A general equilibrium model with non convex adjustment costs of the type we consider cannot generally be solved analytically; moreover, nonconvexities bring about challenging numerical issues, that are particularly relevant in estimation routines, where the model has to be solved repeatedly. With this simplifying assumption we will be able to obtain closed form solutions that, as we will argue, can be seen as approximations of those implied by a more general model, with the additional advantage of a clear and intuitive interpretation.

Consider first the workers' problem. Workers cannot save and consume current income. In each period, a worker is employed in a "good" or a "bad" firm, which pay wages w_g and w_b respectively. In equilibrium the wage fluctuates with the firm productivity, so that with probability $p \geq \frac{1}{2}$, the wage remains constant to its good (w_g) or bad value (w_b), with $w_g \geq w_b$. With probability $1 - p$, a good (bad) wage becomes bad (good). A worker employed in a bad firm can move instantaneously to a good one by paying a moving cost c.

In a stationary environment, the value of working at a good or bad firm are as follows:

$$U_g = u(w_g) + \beta [pU_g + (1-p) U_b] \quad (1)$$

$$U_b = \max \{u(w_b) + \beta [(1-p) U_g + pU_b], u(w_g - c) + \beta [pU_g + (1-p) U_b]\} \quad (2)$$

The first expression shows that people that are in a good job draw utility from their wage $u(w_g)$, do not move, and get continuation utility equal to either U_g or U_b with probability p and $1-p$, respectively. The second expression shows that the mobility decision is taken (and the cost c paid) when expected lifetime utility from moving exceeds that from staying.

In an equilibrium featuring both mobility from bad to good jobs and nonzero employment in bad jobs, it must be that the workers at bad jobs are indifferent between moving or staying, which implies that the two terms in curly brackets are equal. To allow for analytical solutions, take the case of linear utility. Then, after some algebra, we obtain:

$$w_g = w_b + \theta c \quad (3)$$

where $\theta = 1 + \beta(1-2p)$. Thus with serially uncorrelated shocks ($p = \frac{1}{2}$), to attract workers the firm must pay a wage premium that equals the moving cost: $w_g = w_b + c$,: given that next period the state can be good or bad with equal probability, the worker wants to recoup the cost immediately. For a similar reasoning, with full persistence ($p = 1$) the firm only needs to pay the annuity value of the moving cost: $w_g = w_b + (1-\beta)c$. Bertola (2004) formally shows that the wedge between wages implied by equation (3) constitutes a lower bound with respect to the more realistic case in which workers are risk averse.[3]

Firms' productivity shocks are realized at the beginning of the period, before the employment decision is taken. We assume a quadratic production function $F(l, \varepsilon) = el - \frac{\phi}{2} l^2$. Following the literature on firing costs (Bentolilla and Bertola 1990), the adjustment cost

[3] Bertola (2004) shows that consumption falls upon moving: workers are trading off current consumption for expected future consumption. The expected reward must therefore be larger the more concave the utility function, because risk averse individuals suffer more from a given reduction in current consumption.

function is linear in employment changes:

$$g(\Delta l) = f * \Delta l * I_{\{\Delta l<0\}} + h * \Delta l * I_{\{\Delta l>0\}}$$

where $I_{\{.\}}$ is the indicator function, f is the firing and h the hiring cost.[4] Firms decide both whether to adjust when hit by a shock and, in the case they do, by how much.

The state of the firm is described by the couple (l, ε). The general formulation of the firm's problem in recursive terms is

$$V(l, \varepsilon) = Max_{l'}\{F(l', \varepsilon) - w(l' - l)l' - g(l' - l) + \beta EV(l', \varepsilon')\}$$

where, consistently with the workers' problem, we allow for the wage to depend on the labor adjustment. As stated above, this general problem is hard to solve even numerically, due to the non convex nature of the adjustment cost function and to the general equilibrium character of the model. The two shocks assumption greatly simplifies the analysis. If firms adjust when productivity changes, then in equilibrium employment will also take up two values l_g, l_b as productivity, implying four distinct states, (l_i, ε_j), $i, j = g, b$, in which the firm can be. Moreover, from the workers' problem, it follows that firms that want to expand employment from l_b to l_g must increase wages by θc to compensate workers for the mobility cost, while when firing a wage reduction of the same amount will make them indifferent between staying or leaving. The wage rate therefore also switches between two states w_g, w_b, as assumed above. Note that, once the firm pays the wage w_g, labor supply is infinitely elastic for labor increases, and the same holds at w_b for labor decreases. The value of the firm in the four states satisfy:

$$V(l_i, \varepsilon_i) = F(l_i, \varepsilon_i) - w_i l_i + \beta(pV(l_i, \varepsilon_i) + (1 - p)V(l_i, \varepsilon_j)) \quad (4)$$

$$V(l_i, \varepsilon_j) = V(l_j, \varepsilon_j) - g(l_i - l_j) \quad (5)$$

[4]Following most of the literature, we assume that f (h) is a costs and not a transfer from the firm to the worker. We will discuss in the empirical section the consequences of this assumption. See Garibaldi and Violante (2005) for a model that studies how the implications of a firing tax differ from those of a severance payment.

for $i = \{g, b\}, j = \{b, g\}$. The first equation characterizes the value of the firm when productivity does not change, so that no employment adjustment is required; the second in the case that productivity switches, triggering adjustment.

To determine the size of employment adjustment, we use the fact that the marginal value of employment must equalize the hiring cost when hiring and the (negative of) the firing costs when firing (see the appendix for details). Then, the optimal employment change when productivity switches from ε_g and ε_b is:

$$\Delta l = \phi^{-1}(\Delta\varepsilon - \theta(c + f + h)) \tag{6}$$

where we use the notation $\Delta x = x_g - x_b$ throughout. The employment change will be larger the stronger the shock, while it will be reduced by the presence of hiring, firing or mobility costs; moreover, the effects are dampened by the degree of concavity of the production function ϕ.

Consider now the optimality of adjusting. To determine the conditions under which adjustment is the optimal policy, we use the one-step-deviation condition: if adjustment is optimal, it must deliver a higher payoff than not adjusting and resuming the optimal policy from next period onward:[5]

$$V(l_i, \varepsilon_j) > F(l_i, \varepsilon_j) - w_i l_i + \beta(p(V(l_j, \varepsilon_j) - g(l_i - l_j)) + (1-p)V(l_i, \varepsilon_i)) \tag{7}$$

In the appendix we show that, assuming that the labor force is of mass 1 and firms are of total mass 2, the model can be fully characterized and the inequality in (7) directly solved. The optimality of adjusting can be expressed in terms of threshold levels for the changes in

[5]This formulation implies that the firm takes into account the fact that, when changing employment, wages change too: in fact, the payoff from deviating is computed using the wage that results from not changing employment. Alternatively, one could assume that firms take state contingent wages as given, in which case w would follow the same process as ε even in the deviating stage. Given that we are considering firm-level labor supply, is seems more reasonable to assume that a firm knows that it has to increase wages if it wants to increase employment.

productivity above (and below) which adjustment is preferable to inaction.

$$\Delta\varepsilon^{*} = \frac{1}{2}\left(\theta(c + 2(h + f)) + \sqrt{\theta c(\theta c + 4\phi)}\right) \tag{8}$$

when hiring and

$$\Delta\varepsilon^{**} = \frac{1}{2}\left(\theta(c + 2(h + f)) + \sqrt{\theta c(\theta c - 4\phi)}\right) \tag{9}$$

when firing.

Equations (3, 6, 8, 9) for the basis of the moment conditions we will use in the empirical analysis. They supply two extensive conditions for wages and employment conditional on adjusting and two intensive conditions for employment for the optimality of adjusting. Even if derived under the two shocks assumption, they have the same structure that would result from a model with a continuum of shocks. For example, the partial equilibrium investment model with a continuum of shocks of Fuentes, Gilchrist and Rysman (2004) also delivers a binary condition on the optimality of adjustment and a size of adjustment condition.

There are some important aspects to note. First, in all the equilibrium conditions we are going to use for identification purposes, firing and hiring costs enter as sum. This implies that the model can only identify the total amount of internal adjustment costs, $k = h + f$, rather that its two components separately. While this is a shortcoming of the model, our primary interest is in distinguishing internal from external costs, rather than firing from hiring costs. We thus believe this is a relatively unimportant issue. From now on, we will refer to k as total internal costs and neglect the distinction between f and h.

Another important aspect is that, given the continuum of firms assumption and the symmetric Markov transition matrix, in the aggregate there is a constant and equal share of firms in each state. This is therefore an economy with only idiosyncratic shocks, i.e., shocks that affect the single production units without altering aggregate outcomes. This aspect will have a very strict counterpart in the empirical analysis, where the shocks will be defined at the level of the single firm, after filtering out aggregate and local shocks.

Some further properties of the equilibrium can be more easily described using Figure 1.

In a frictionless world, in which both $k = 0$ and $c = 0$, firms face an infinite elastic labor supply at the prevailing wage $\widetilde{w}$: wages do not respond to the idiosyncratic firm conditions. Moreover, there is no lumpiness in the employment response to shocks. In terms of Figure 1, the wage is fixed at $\widetilde{w}$ and employment fluctuates between $\widetilde{l}_g$ and $\widetilde{l}_b$:

$$\Delta l = \phi^{-1} \Delta \varepsilon \tag{10}$$

The introduction of frictions has several implications. First, the response of employment becomes lumpy: firms only adjust when the shocks are sufficiently large. In Figure 1, for movements of the MPL within the $[w_b, w_g]$ range, employment does not respond. The smallest change in the frictionless employment at which adjustment occurs even with frictions is:

$$\Delta \widetilde{l}^* = \phi^{-1} \Delta \varepsilon^* = \gamma^H \tag{11}$$

and similarly for downward adjustment. One important implication of (11) is that *any* type of friction induces lumpiness: in fact, from (8) and (9) it follows that $\Delta \widetilde{l}^* \neq 0$ if either $c \neq 0$ or $k \neq 0$. This implies that lumpy adjustments will signal the presence of frictions, but cannot be used to determine their nature. This is in contrast with most of the literature on employment adjustment at the level of the firm, where lumpy behavior is usually taken as signaling adjustment costs at the level of the firm only (Hamermesh and Pfann 1996).

The second implication relates to wage changes. Given that firms need to compensate workers from the moving costs they bear upon changing employer, wage changes only occur together with employment changes. Moreover, the wage response to the shocks is due to the cost of moving: if firing costs were the only friction in the market, then we should observe no wage response at the firm level. In terms of the figure, the broken wage schedule only holds with $c \neq 0$. In particular, one can write $\Delta w = \theta c$ for workers employed in firms that adjust employment upward, and $\Delta w = -\theta c$ for firms that adjust downward.

The third implication relates to employment changes. With respect to a frictionless world, frictions not only induce lumpy adjustments, but also dampen employment changes when they take place, as can be seen by expressing actual adjustment in deviation from the frictionless counterpart:

$$\Delta l = \Delta \widetilde{l} - \psi \tag{12}$$

where $\psi = \phi^{-1}\theta(c + k)$. Again in Figure 1, employment expansions are further reduced to l_g^{cf} by the presence of firing costs (and similarly for contractions).

4 Data

We rely on two administrative data sets, one for firms and one for workers. Data for firms are obtained from *Centrale dei Bilanci* (Company Accounts Data Service, or CAD for brevity), while those for workers are supplied by *Istituto Nazionale della Previdenza Sociale* (National Institute for Social Security, or INPS). Since for each worker we can identify the firm he/she works for, we combine the two data sets and use them in a matched employer-employee framework. There is a burgeoning empirical literature on the use of matched employer-employee data sets (see Hamermesh (1999) for an account).

The CAD data span from 1982 to 1994, a period that comprises two complete business cycles. It contains detailed information on a large number of balance sheet items together with a full description of firm characteristics (location, year of foundation, sector of operation, ownership structure), plus other variables of economic interest usually not included in balance sheets, such as employment and flow of funds. Balance sheets are collected for approximately 30,000 firms per year by *Centrale dei Bilanci*, an organization established in the early 1980s jointly by the Bank of Italy, the Italian Banking Association, and a pool of leading banks to gather and share information on borrowers. Since the banks rely heavily on it in granting and pricing loans to firms, the data are subject to extensive quality controls by a pool of professionals, ensuring that measurement error should be negligible.

INPS provides us with data for the entire *population* of workers registered with the social security system whose birthday falls on one of two randomly chosen days of the year. Data are available on a continuous basis from 1974 to 1994. We use the data after

1981 for consistency with the timing of the CAD data. The INPS lacks information on self-employment and on public employment (public firms are also absent in the CAD). The INPS data set derives from forms filled out by the employer that are roughly comparable to those collected by the Internal Revenue Service in the US.[6] Misreporting is prosecuted.

Given that the INPS data set includes a fiscal identifier for the employer which is also present in the CAD data set, linking the employer's records to the employees is relatively straightforward. As in other countries where social security data are available, the Italian INPS data contain some detailed information on worker compensation but information on demographics is scant.

Table 1 reports various descriptive statistics for the firms (Panel A) and workers (Panel B) present in our sample. We report separate statistics for the whole sample and for the sample obtained after matching firm and worker information. From an initial sample of 177,654 firm/year observations, we end up with 116,686, corresponding to 16,037 firms. We exclude firms with intermittent participation (40,225 observations) and those with missing values on the variables used in the empirical analysis (20,620 observations) or extreme employment changes (123 observations). Since the panel is unbalanced, the firms in this sample appear from a minimum of one to a maximum of 13 years.

The whole sample ranges from very small firms to firms with almost 180,000 employees, with an average of 204 and a median of 60. As expected, most of the firms are in the North (75 percent). As for the distribution by industry, manufacturing firms account for about 75 percent of the final sample. Construction firms account for about 15 percent. The remaining 10 percent is scattered in the service and retail sectors. The matched sample includes larger firms, but the distribution by region and industry is similar to that in the whole sample.

Panel B reports sample characteristics for the workers in the 1982-1994 INPS sample. We start with an initial sample of 267,539 worker/year observations (including multiple observations per year for the same worker due to multiple jobs, intra-firm position change,

[6]While the US administrative data are usually provided on a grouped basis, INPS has truly individual records. Moreover, in the US earnings records are censored at the top of the tax bracket, while the Italian data set is not subject to top-coding.

and inter-firm mobility) and end up with 162,184. Sample selection is made with the explicit aim of retaining workers with stable employment and tenure patterns. First we exclude those younger than 18 or older than 65 (2,652 observations), circumventing the problem of modeling human capital accumulation and retirement decisions. To avoid dealing with wage changes that are due to job termination (quits or layoffs) or unstable employment patterns, we exclude workers with part-time employment and those with multiple jobs (81,117 observations). For similar reasons, we drop individuals who worked for less than 12 months (43,750 observations). Moreover, we keep only individuals with non-zero recorded earnings in all years (105 observations lost). Finally, we eliminate those with missing values on the variables used in the empirical analysis (8,627 observations). Since these selections − particularly those that exclude firm-movers − can potentially affect our results, we account for sample selection bias (see Section 4.3).

Our measure of earnings covers remuneration for regular and overtime pay plus non-wage compensation. We deflate earnings using the CPI (1991 prices). For workers with intermittent participation we treat two strings of successive observations separated-in-time as if they pertained to two different individuals.

Workers in the whole sample are on average 39 years old in 1991; production workers account for 62 percent of the sample, 37 percent are clericals, and about 2 percent managers. Males are 73 percent of our sample and those living in the South 14 percent. Finally, net earnings in 1991 are roughly 17,000 euro on average. In the matched sample individual characteristics are fairly similar to the ones in the whole sample.

5 Identification

The identification procedure is based on the equilibrium relations obtained form the model. The most important shortcoming of the simple GE model is that it has only two productivity states, a clearly untenable assumption when bringing it to the data. We depart from the model and generalize this structure by allowing the shock (and the consequent labor adjustment) to take any value. Without the two-shocks assumption, we would not be able to directly calculate the value functions and obtain a closed form solution. Without closed form, the estimation procedure would still take the form of a threshold rule plus an extensive equation but would require the numerical solution of a nested fixed point problem, with the additional complexity of determining equilibrium wages.[7] While doable in principle, this would greatly increase the computational complexity and reduce the transparency of our procedure; moreover, the estimation results would still depend on the functional form and distributional assumptions. On balance, we believe that the increase in complexity is not matched by the increase in explanatory power. The main advantage of the strategy proposed here is that it is simple and transparent without being unrealistic. Indeed, the conditions obtained from the two shocks model can be thought of as approximations to the exact solution with a continuum of productivity shocks.

Equation (10) indicates a linear relation between shocks and frictionless labor changes that we maintain. We include an error term that captures unobserved determinants of frictionless employment growth, i.e.,[8]

$$\Delta\widetilde{l}_{jt} = \phi^{-1}\Delta\varepsilon_{jt} + \zeta_{jt} \tag{13}$$

Our theory delivers the adjustment thresholds:

$$\gamma^H = \phi^{-1}\frac{\theta(c+2k) + \sqrt{\theta c(\theta c + 4\phi)}}{2} \tag{14}$$

$$\gamma^L = -\phi^{-1}\frac{\theta(c+2k) + \sqrt{\theta c(\theta c - 4\phi)}}{2} \tag{15}$$

We estimate our structural parameters using a multi-step strategy. First, we estimate the parameters that affect the probability of adjusting using observations on all firms; then, the size of adjustment equation using only the observations on the firms that adjust; finally, the wage equation. The parameters we estimate at these three stages are non-linear combinations of the structural parameters; the latter are over-identified from these restrictions and therefore we can test the overidentifying restrictions. We use optimal

[7]The only paper we are aware of that estimates a dynamic programming model of labor demand is Cooper, Haltiwanger and Willis (2003). However, they take wages as given at the plant level, and so their setting is partial equilibrium in nature. Other authors (Rota 2004) derive an approximated labor demand Euler equation in the presence of adjustment costs.

[8]For the remainder of the paper i, j, and t index workers, firms, and years, respectively.

minimum distance to map reduced form parameters onto structural parameters.

More specifically, rewrite actual labor adjustment as:

$$\Delta l_{jt} = \begin{cases} \Delta\widetilde{l}_{jt} + \psi & \text{if } \Delta\widetilde{l}_{jt} < \gamma^L \\ 0 & \text{if } \gamma^L \leq \Delta\widetilde{l}_{jt} \leq \gamma^H \\ \Delta\widetilde{l}_{jt} - \psi & \text{if } \Delta\widetilde{l}_{jt} > \gamma^H \end{cases} \tag{16}$$

Firms can be in one of three regimes: hiring, firing, or doing nothing. These regimes are defined by the dummies:

$$\begin{aligned} s^+_{jt} &= \mathbf{1}\left\{\Delta\widetilde{l}_{jt} > \gamma^H\right\} = \mathbf{1}\left\{\zeta_{jt} > \gamma^H - \phi^{-1}\Delta\varepsilon_{jt}\right\} \\ s^-_{jt} &= \mathbf{1}\left\{\Delta\widetilde{l}_{jt} < \gamma^L\right\} = \mathbf{1}\left\{\zeta_{jt} < \gamma^L - \phi^{-1}\Delta\varepsilon_{jt}\right\} \\ s^0_{jt} &= \mathbf{1}\left\{\gamma^L \leq \Delta\widetilde{l}_{jt} \leq \gamma^H\right\} = \mathbf{1}\left\{\gamma^L - \phi^{-1}\Delta\varepsilon_{jt} \leq \zeta_{jt} \leq \gamma^H - \phi^{-1}\Delta\varepsilon_{jt}\right\} \end{aligned}$$

with $s^+_{jt} + s^-_{jt} + s^0_{jt} \equiv 1$. Assume that $\zeta \sim N\left(0, \sigma^2_\zeta\right)$. The likelihood function for the regime a firm happens to be in is:

$$L = \prod_{s^-_{jt}=1} \Phi^L_{jt} \prod_{s^+_{jt}=1} \left(1 - \Phi^H_{jt}\right) \prod_{s^0_{jt}=1} \left(\Phi^H_{jt} - \Phi^L_{jt}\right) \tag{17}$$

with $\Phi(.)$ the c.d.f. of the standard normal, $\Phi^k_{jt} = \Phi\left(\frac{\gamma^k - \phi^{-1}\Delta\varepsilon_{jt}}{\sigma_\zeta}\right)$ $(k = H, L)$, and σ_ζ the scale factor. Thus, γ^H, γ^L, and ϕ^{-1} can only be identified up to scale at this stage.

The next step is to consider the continuous aspect of the labor adjustment process. Given that, by definition, $E\left(\Delta l_{jt} | s^0_{jt} = 1\right) = 0$, we can use the law of iterated expectations to write:

$$E\left(\Delta l_{jt} | \Delta\varepsilon_{jt}\right) = E\left(\Delta l_{jt} \middle| s^+_{jt} + s^-_{jt} = 1, \Delta\varepsilon_{jt}\right) \Pr\left(s^+_{jt} + s^-_{jt} = 1 | \Delta\varepsilon_{jt}\right)$$

where $s^+_{jt} + s^-_{jt} = 1$ identifies an adjusting firm. Since the probability of adjusting is $\Pr\left(s^+_{jt} + s^-_{jt} = 1_{jt} | \Delta\varepsilon_{jt}\right) = 1 - \Phi^H_{jt} + \Phi^L_{jt}$, it follows that -using again the law of iterated expectations-

$$E\left(\Delta l_{jt}\left|s_{jt}^{+}+s_{jt}^{-}=1,\Delta\varepsilon_{jt}\right.\right)=\sum_{s=\{s^{+},s^{-}\}}\frac{E\left(\Delta l_{jt}\left|s_{jt}=1,\Delta\varepsilon_{jt}\right.\right)}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}$$

$$=\frac{\left[\phi^{-1}\Delta\varepsilon_{jt}-\psi+E\left(\zeta_{jt}\left|s_{jt}^{+}=1\right.\right)\right]\left(1-\Phi_{jt}^{H}\right)+\left[\phi^{-1}\Delta\varepsilon_{jt}+\psi+E\left(\zeta_{jt}\left|s_{jt}^{-}=1\right.\right)\right]\Phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}$$

$$=\phi^{-1}\Delta\varepsilon_{jt}-\psi\frac{1-\Phi_{jt}^{H}-\Phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}+\sigma_{\zeta}\frac{\phi_{jt}^{H}-\phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}} \tag{18}$$

where we have used the properties of the truncated normal distribution repeatedly, together with equations (13) and (16). This regression can be run on the subset of firms that adjust their level of employment. The terms $\frac{\phi_{jt}^{H}-\phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}$ and $\frac{1-\Phi_{jt}^{H}-\Phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}$ are generalized selection terms that account for the fact that we are selecting only the firms that are adjusting their level of employment. A two step strategy can be adopted. In the first step, one estimates (17) and construct consistent estimates of $\frac{\phi_{jt}^{H}-\phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}$ and $\frac{1-\Phi_{jt}^{H}-\Phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}$. In the second step, one estimates (18) by OLS using the estimates of $\frac{\phi_{jt}^{H}-\phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}$ and $\frac{1-\Phi_{jt}^{H}-\Phi_{jt}^{L}}{1-\Phi_{jt}^{H}+\Phi_{jt}^{L}}$ in the place of the true ones.[9]

The use of the firms' adjustment policies allows for the identification of the total costs of adjustment. In fact, by combining the parameters of the ordered probit (that identify γ^{H}, γ^{L}, and ϕ up to a scale) and that of the size of adjustment (that identify the scale and supply additional overidentifying restrictions), one can recover the value of $\theta(c+k)$. This is in fact the strategy that, under different forms, has been followed by the literature on factor demand in the presence of adjustment costs (Hamermesh and Pfann 1996). Unfortunately, these estimates do not help to separately identify c and k since these two parameters enter the structural equations jointly[10]. However, the separate identification of c and k can be

[9] Needless to say, one could estimate the two equations in one single step by writing the likelihood function as an ordered Tobit. The two step strategy used here greatly reduces the computational burden. Indeed, it is well known that estimating this type of likelihood function in one step tends to give rise to serious convergence problems (see, for example Nilsen et al. 2003). Efficiency is not an issue here because all the standard errors are computed by the block bootstrap.

[10] In principle, the two parameters could be identified by the non-linearity implied by the the functional form assumptions of the adjustment thresholds. In practice, this identification strategy is very tenuous and indeed it fails due to inability to fit the restrictions on the lower threshold. For this reason, our minimum distance mapping omits the lower threshold and uses the wage restrictions which have the advantage of being economically transparent.

achieved by considering the implications that our model has for the behavior of wages, which offer an equation that uniquely identifies θc.

Wage changes in our model depend on the employment change behavior of the firm, and in particular

$$\Delta w_{ijt} = \begin{cases} \theta c + X'_{ijt}\beta + \omega_{ijt} & \text{if } s^{+}_{jt} = 1 \\ X'_{ijt}\beta + \omega_{ijt} & \text{if } s^{0}_{jt} = 1 \\ -\theta c + X'_{ijt}\beta + \omega_{ijt} & \text{if } s^{-}_{jt} = 1 \end{cases}$$

or

$$\Delta w_{ijt} = X'_{ijt}\beta + \theta c \left(s^{+}_{jt} - s^{-}_{jt} \right) + \omega_{ijt} \tag{19}$$

where X'_{ijt} is a vector of observable individual characteristics affecting wage growth (such as age or tenure). The variable $\left(s^{+}_{jt} - s^{-}_{jt} \right)$ is an adjustment indicator defined for the firm the individual is working for. It equals 1 (-1) if the firm has adjusted its level of employment upward (downward) and zero if it has remained inactive. Note that $\left(s^{+}_{jt} - s^{-}_{jt} \right)$ is *not* an indicator for whether the worker has moved, but one for whether the firm has changed its labor force.

A first approach to estimation is simply to estimate the coefficients in (19) by OLS. This gives unbiased and consistent estimates provided $E\left(\omega_{ijt} | X'_{ijt}, \left(s^{+}_{jt} - s^{-}_{jt} \right)\right) = 0$. However, it is clear that wages and employment are jointly determined at the firm level, so that this assumption is not likely to hold. Consider an exogenous increase in labor costs, due for example to tax incentives offered to firms hiring (or starting up plants) in high unemployment areas. The firm might respond to this by increasing employment, against the exogeneity assumption. Moreover, while $\left(s^{+}_{jt} - s^{-}_{jt} \right)$ is supposed to be related to gross labor adjustment, we only observe net labor adjustment. This is a standard measurement error argument, which again means that $E\left(\omega_{ijt} | X'_{ijt}, \left(s^{+}_{jt} - s^{-}_{jt} \right)\right) \neq 0$. Both problems can be dealt with by using Instrumental Variables (IV).

Note however that the IV estimator is valid only if the instrument is orthogonal to ω_{ijt} *conditioning* on adjustment. The model suggests an obvious instrument for $\left(s^{+}_{jt} - s^{-}_{jt} \right)$, namely $\Delta\varepsilon_{jt}$. However, while $\Delta\varepsilon_{jt}$ and ω_{ijt} are unlikely to be correlated from an unconditional point of view, they may be correlated *conditioning* on adjustment. In fact, if moving costs are heterogeneous, then movers will tend to be those with the lower costs. These more mobile workers will be the ones determining the equilibrium wage differential, which therefore will reflect a level of moving costs below the population average. Thus, one might expect $\Delta\varepsilon_{jt}$ and ω_{ijt} to be negatively correlated conditioning on adjustment. Our solution to this problem is (a) to implement a simple IV strategy, in which $\left(s^+_{jt} - s^-_{jt}\right)$ is instrumented using $\Delta\varepsilon_{jt}$, and (b) a generalized IV strategy which makes distributional assumptions about ω_{ijt} – along with knowledge of what drives the employment decision of the firm – to construct correction (selection) terms in (19) that allow us to estimate the parameters of interest consistently.

To see this point more clearly, assume ω_{ijt} is not orthogonal to the employment regime the firm is in and that it is $\omega_{it} \sim N\left(0, \sigma^2_\omega\right)$. Moreover, assume that its correlation coefficient with ζ_{jt} is $\rho_{\omega,\zeta}$. Then, by the law of iterated expectations:

$$\begin{aligned}
E\left(\Delta w_{ijt}|X_{ijt}, \Delta\varepsilon_{jt}\right) &= E\left(\Delta w_{ijt}|s^+_{jt}=1, X'_{ijt}, \Delta\varepsilon_{jt}\right)\Pr\left(s^+_{jt}=1|X'_{ijt}, \Delta\varepsilon_{jt}\right) \\
&+ E\left(\Delta w_{ijt}|s^-_{jt}=1, X'_{ijt}, \Delta\varepsilon_{jt}\right)\Pr\left(s^-_{jt}=1|X'_{ijt}, \Delta\varepsilon_{jt}\right) \\
&+ E\left(\Delta w_{ijt}|s^0_{jt}=1, X'_{ijt}, \Delta\varepsilon_{jt}\right)\Pr\left(s^0_{jt}=1|X'_{ijt}, \Delta\varepsilon_{jt}\right) \\
&= X'_{ijt}\beta + \theta c\left(1 - \Phi^H_{jt} - \Phi^L_{jt}\right) + \theta\rho_{\omega,\zeta}\sigma_\zeta\left(\phi^H_{jt} + \phi^L_{jt}\right) \qquad (20)
\end{aligned}$$

The strategy is thus the following. Estimation of (17) allows us to construct consistent estimates of the selection terms $\left(1 - \Phi^H_{jt} - \Phi^L_{jt}\right)$ and $\left(\phi^H_{jt} + \phi^L_{jt}\right)$. We can then run a regression of Δw_{ijt} onto X'_{ijt}, $\left(1 - \Phi^H_{jt} - \Phi^L_{jt}\right)$, and $\left(\phi^H_{jt} + \phi^L_{jt}\right)$. The coefficient on $\left(1 - \Phi^H_{jt} - \Phi^L_{jt}\right)$ is an estimate of the mobility costs θc. The term $\left(\phi^H_{jt} + \phi^L_{jt}\right)$ controls for the "sorting" effect mentioned above. Note that without $\left(\phi^H_{jt} + \phi^L_{jt}\right)$, (20) would be, effectively, just a conventional 2SLS regression, since $\left(1 - \Phi^H_{jt} - \Phi^L_{jt}\right) = E\left(\left(s^+_{jt} - s^-_{jt}\right)|X_{ijt}, \Delta\varepsilon_{jt}\right)$ where $\Delta\varepsilon_{jt}$ is the "instrument" from the first stage and so this is the linear projection of $\left(s^+_{jt} - s^-_{jt}\right)$ onto the (excluded and included) instruments. The presence of the term $\left(\phi^H_{jt} + \phi^L_{jt}\right)$ implies that we will be estimating a "generalized" 2SLS specification, where

$\rho_{\omega,\zeta}\sigma_\zeta\left(\phi_{jt}^H + \phi_{jt}^L\right)$ controls for the correlation between the endogenous variable and the error term.

To summarize, the identification procedure thus entails the following steps:

1. Obtain a measure of idiosyncratic shocks to the marginal product of labor ($\Delta\varepsilon_{jt}$).

2. Estimate (17), i.e., an ordered probit for negative, zero, and positive adjustments using the shock and possibly other covariates as explanatory variables; recover estimates of Φ_{jt}^H, Φ_{jt}^L, ϕ_{jt}^H, and ϕ_{jt}^L;

3. Estimate the size of adjustment equation (18) using data on adjusting firms.

4. Estimate the (worker level) wage change equation (20) including an indicator of the firms' adjusting policy, accounting for endogeneity and selection;

5. Recover the main structural parameters of interest θc, θk, ϕ from the reduced form estimates of the previous four steps using optimal minimum distance.

If one uses the theoretical restrictions imposed on γ^H and γ^L, the thresholds of the adjustment decision, the model is over-identified with three overidentifying restrictions. In what follows we use optimal minimum distance on the reduced form estimates to recover estimates of the structural parameters. We use the block bootstrap-generated covariance matrix (based on 200 replications) as the weighting matrix. We do not use the theoretical restriction on γ^L because, as (15) shows, it depends on a square root term that is not defined for some values of the parameters, and so we have two overidentifying restrictions.

Note that we can separately identify θc and θk, but θ cannot be identified. Given that $\theta = 1 + \beta(1 - 2p) \leq 1$, our estimates will provide lower bounds for the true costs of adjustments. We will return to these point when discussing the results.

6 Results

6.1 Employment adjustment

We start by documenting the lumpiness of employment adjustment. Figure 2 plots the distribution of employment changes pooling all years together, excluding for readability the first and last percentile of the distribution (approximately + and -100). The first thing to note is that the amount of adjustment is fairly modest: about 95% of the observations lie between -28 and $+25$. The median employment change is exactly zero (the mean is similar), and about 17% of the firms in our sample do not change their employment from one year to the next; 40% adjust downward, and 42% adjust upward. Not surprisingly, this indicates that lumpiness is an important component of the employment choice.

The model predicts that firms will respond to changes in the marginal product of labor. Given that we do not have a direct measure of it, we follow the previous literature on q-models of adjustment (Abel and Eberly 1994) and use an average measure, defined as value added (in thousand of 1991 euros) at time t divided by employment at time $t-1$, $y_t = \frac{VA_t}{l_{t-1}}$.[11] To obtain a measure of idiosyncratic shocks, we first regress y_t on a full set of year dummies, industry dummies, and area dummies. We then take the first difference of the residual of this regression as our measure of idiosyncratic shocks.[12] This is in fact in line with the theoretical counterpart, where labor adjustment is prompted by *changes* in productivity.[13]

As argued above, the costs estimated using idiosyncratic shocks will be related to those entailed by an employer change rather than to long distance geographical mobility. In fact, by netting out time and industry effects, we are excluding shocks that change aggregate labor demand, and concentrate on those that modify the relative demand of a firm with respect to the others. Most of these changes will be resolved by workers job changes that do not entail geographical mobility. This is indeed the main form of job mobility in the

[11] As usual in this type of regressions, we use the lagged value of employment to avoid simultaneity biases with the left-hand side variable. The use of current value added is justified by the idea that, due to "time to build", it might take some time before new workers are fully operative. We also experimented with lagged value added, obtaining very similar results.

[12] Given that, due to accounting rules or special events, such as acquisitions, mergers or breaking-ups, balance sheet data might record extreme values related to events beyond our interests, we run a procedure to exclude outliers. In particular, we exclude the first and last percentile of the resulting shock distribution. The distribution is in fact characterized by extreme values. The median value of the shock is -239 euros, the first and the extreme percentiles are -45,865 and 42,787 euros. We also exclude firms whose employment increases more than 20-folds and those that have negative growth greater than 90% in absolute value and an initial size of more than 100.

[13] We have also experimented with a shock obtained as the residual of a firm fixed effects regression with year dummies. In this case, a shock is a measured as the deviation of value added per worker from a firm-specific average. This definition is less in line with the model; at the same time, the level of productivity has a more natural interpretation than changes in productivity in terms of state variable. Results are similar to those reported in the paper.

Italian labor market, traditionally characterized by high geographical mobility costs. This is confirmed by our data: by considering the location of the employer, we find that 33% of workers that change job remain within the same municipality, 63% within the same province, and 75% within the same region (see Table 7).

We then run an ordered probit for the choice of employment change regime (positive, zero, or negative change in employment). Table 2 report the results for the regression with the shock as sole regressor in column 1. We find that the effect of the shock is positive and statistically significant, as expected: higher shocks imply a higher likelihood of moving from negative to zero to positive adjustments.[14] The two adjustment thresholds are also precisely estimated with signs in line with theoretical predictions. Note that such estimate are identified up to scale, so that their size cannot be interpreted directly. However, we can already infer from these estimates that the (total) costs of adjustments are nonzero.

One potential critique to this regression, especially when considered in conjunction with the size of adjustment one that follows, is the lack of any exclusion restriction, so that identification of the effects only comes from functional form assumptions (Heckman 1990). We have experimented using the number of periods since last adjustment as an exclusion restriction.[15] In fact, in general models where productivity follows a random walk as in Bentolilla and Bertola (1990), the martingale property implies that the expected value of the shock (and therefore the size of the adjustment) is not dependent on the number of periods elapsed since last adjustment, while the variance (and therefore the likelihood of adjusting) increases with them. This makes the number of periods since last adjustment a natural candidate for an exclusion restriction relative to the size of adjustment equation. Results from adding this variable are reported in column 2. Adding the exclusion restriction has no effect on the estimates. As expected, firms that have been inactive for longer are, ceteris paribus, more likely to adjust in the current period.

Using the estimates of the ordered probit, we construct the variables included in the size of adjustment equation (18), that we run on firms that do adjust We find that the shock

[14]In all cases, do account for the genereted regressor bias, we calculate the standard errors by the block bootstrap, based on 200 replications. The block bootstrap accounts for the unbalanced nature of the sample.

[15]This variable is, of course, subject to left censoring.

has a positive and significant impact on the size of adjustment (Table 3). As in Table 2, column 1 (2) again refers to the case without (with) the exclusion restriction. The results are similar across the two specifications, but predictably more precise when an exclusion restriction is used. The estimate of ψ (the coefficient of $\frac{1-\Phi_H-\Phi_L}{1-\Phi_H+\Phi_L}$, which measures how much adjustment is dampened with respect to a frictionless world) is negative, as theory predicts, but not statistically significant. In fact, one typically obtains much more precise estimates on the discrete margin (adjust/don't), while, conditioning on this, the additional information obtained from the continuous one is rather limited. The relative precisions of the estimates will be taken into account by our minimum distance procedure to recover the structural parameters.

6.2 Wage adjustment

To disentangle the external and internal components of total adjustment costs, we now turn to the wage equation. We construct wages as the sum of annual normal compensation and fringe benefits. We include in the wage growth regression a variable, $\left(s_{jt}^{+} - s_{jt}^{-}\right)$ which equals -1 if the firm is reducing employment, 1 if it is expanding it and zero otherwise. As shown in equation 19, the coefficient on this variable is crucial for the identification of the extent of external costs, θc.

In Table 4 we report the results of the wage growth equation. We include the usual regressors of wage equations suggested by the literature, i.e. age, tenure, sector, year, location and job title dummies. The variable of interest is the employment policy of the firm at which the worker is currently attached. The first specification is an OLS regression with $\left(s_{jt}^{+} - s_{jt}^{-}\right)$ assumed exogenous. The results, reported in column 1, show that adjustment does entail external costs: the estimate of θc is positive and statistically significant, with a (bootstrap) t-statistic of 7. In absolute terms, the value is rather modest: it implies that expanding firms pay a yearly premium of 108 euros to their workers. Note that this is an estimate of θc, not of the mobility cost c alone. Given that $\theta < 1$, it represents a lower bound for the cost of adjusting, a point to which we will come back later.

All other regressors in our wage change equation have estimated effects that are in line with expectations: wage growth decreases with age and tenure (reflecting concavity of the wage *level* functions with respect to such variables). The wage of male workers increases on average by 200 euros more that than of females, and blue and white collar are characterized by lower wage growth than managers (the excluded category). Wages also grow less in the South (results not reported). Given that these estimates are very stable throughout the specifications, we will not comment on them any more in what follows.

Even keeping in mind the lower bound argument, the wage premium attached to mobility seems surprisingly low. In fact, the conventional wisdom is that the Italian labor market is characterized by a low willingness of workers to move in the face of better job opportunities. Part of the explanation can be traced back to the fact that we are considering idiosyncratic firm shocks, so that mobility tends to resolve mostly locally. Even so, the value we estimate would imply that mobility costs represents a very small fraction of total adjustment costs.

A possible explanation for the low estimate of mobility costs is that it is downward biased due to the fact that wage and employment adjustment are determined simultaneously, as discussed above. To account for this, in column (2) we use a 2SLS procedure - using the shock to value added as an exclusion restriction, or instrument, for $\left(s_{jt}^{+}-s_{jt}^{-}\right)$. The results change quite dramatically, confirming the importance of the endogeneity issue. The estimate of θc increases by one order of magnitude to 1.2, or 1,200 euros. In column (3) we propose a generalized 2SLS procedure in which we also allow for the possibility that the instrument is correlated with the error term of the wage equation *conditional* on adjusting. The estimate is barely affected, and the extra regressor statistically insignificant. This suggests that such correlation is not likely to be present in our data. In terms of the interpretation given above, this result indicates that heterogeneity in mobility costs is not an important source of bias for our estimates.

In column 4-6 we report the results for the three regressions when the exclusion restriction is used to estimate the employment adjustment equation, which affects the estimated value of $1-\Phi_H-\Phi_L$ and of $\phi_H+\phi_L$. The results change very little.

We performed a number of experiments to assess the robustness of the results. One criticism to our identification strategy is that our measure of workers' earnings include both normal wages and overtime premiums. One thing to notice from the outset is that the use of overtime, *per se*, reflects the existence of frictions; in the absence of frictions, the firm would hire workers in the spot market when facing a positive shock and fire them in the alternative scenario. However, firms that face adjustment costs, internal and external, may use a mixed policy of hiring some workers from outside and also use overtime for their existing workers. Thus, earnings changes may reflect, at least in part, overtime premiums rather than mobility cost differentials.

On the other side, a survey of the Bank of Italy on industrial enterprises shows that firms tend to use overtime and hirings as substitute: output demand changes are satisfied up to a point with changes in the hours worked by current employees, until when the firm decides to hire or fire and brings back hours to normal levels. To check if this is an important issue, we focus on a sample of firms that are adjusting, if at all, at the margin, i.e. by one worker up or down. These firms are the least likely to be using a mixed policy, as long as overtime hours worked by the existing workers are perfect substitutable with normal hours worked by new hired workers: once they hire, overtime should go to zero, so that the effect of adjustment on wage should only come from mobility costs. The results, reported in column 2 of Table 5, show that the estimated coefficient drops from 1.175 to 0.964, pointing to some weak evidence for our coefficient reflecting overtime premiums; however, the drop is so small to suggest that overtime is at most a minor concern.

Another issue is that of wage rigidity. In our model, wages adjust upward as well as downward to reflect the employment adjustment policy of the firm. It can be argued that the firm may have leverage to adjust wages upward, but not downward due to contractual obligations and to nominal rigidities. In Column (3) we allow for (instrumented) s^- and s^+ to enter separately in our regression. Given the large standard errors, we fail to reject symmetry. We reach the same conclusions when using different specifications, such as entering s^- and s^+ in two separate regressions, in which case standard errors are small but the point estimates very similar, or running the regressions without instruments. One possible explanation is that, in a period of moderately sustained inflation (average inflation rate in our sample period was 7.7% per year), downward rigidity of nominal wages is not a binding constraint for flexibility of the real wage.

6.3 Joint estimates

As discussed above, our model is over-identified. Table 6 uses optimal minimum distance on the reduced form parameters to back up the structural parameters. As in previous tables, specifications 1 and 2 refer to a model without and with an exclusion restriction in the employment adjustment equation. Given that the results are fairly similar, our comment here will be limited to the simplest case without exclusion restriction. The weighting matrix of OMD is obtained from the block bootstrap.[16] The coefficients are all very precisely estimated, and the test of overidentifying restrictions signals a good fit of the model -despite its simplicity. The estimate of ϕ (the curvature of the production function) is about 1.3. Firms adjust upward if the change in frictionless employment is above 7.6, and downward if frictionless employment goes down by 9.2 or more. The estimate of $\theta(c+k)$ is 8,836 euro. Due to the presence of the scaling parameter $\theta < 1$, this only represents a lower bound for the absolute level of total adjustment costs. However, the relative contribution of internal and external costs is identified from the ratio of the estimate of θc, equal to 1,172, and θk, equal to 7,664. Our results imply that the latter clearly dominates, accounting for around 87% of total costs. Still, the share attributable to moving costs is non trivial and shows that, by disregarding this component, one would overestimate the internal costs of adjustment.

While we have no direct measure of θ that may be used to pin down the absolute level of total adjustment costs, some inference can be drawn for illustrative purposes. Assume that $\beta = .96$, in line with the fact that our data are annualized. Recalling that $\theta = 1 - \beta(2p-1)$ and that $1/2 \leq p \leq 1$, it follows that θ varies between 0.04 ($p = 1/2$, or no persistence) and 1 (full persistence). This implies that the costs of adjusting are included in the rage $8,834 \leq (c+k) \leq 220,855$. While indicative of the bounds, the range is too wide to provide an idea of the size of the costs. We use our data to get an empirical counterpart to p, the measure of productivity persistence. To map actual productivity changes into the two state space of the

[16] In practice, the procedure entails the minimization with respect to the structural parameter of the weighted squared sums of the difference between the theoretical restrictions underlying equations 16, 18 and 19 and the estimated parameters. The weighting matrix is obtained from the bootstrapped variance/covariance matrix.

model, we adopt the following strategy. We first estimate an AR(1) regression on the level of productivity shocks, finding a coefficient of .33. We then use the fact that the transition probability for shocks can be written as $\Pr\{\varepsilon_t|\varepsilon_{t-1}\} = (1-p)(\varepsilon_g+\varepsilon_b)+(2p-1)\varepsilon_{t-1}$. Then, in a regression of ε_t on ε_{t-1} the AR(1) coefficient ρ can be used to obtain the corresponding value of $p = \frac{1}{2}(1+\rho) = .68$. Using this and $\beta = .96$, we obtain a value for $\theta = .68$, which implies a total cost of adjusting employment of approximately 13,000 euros.[17] Internal costs are $11,270$ euros, equal to 66% of the yearly compensation or to 8 months of salary. If we relate the internal costs to legal firing costs, our estimates seems fairly reasonable. As seen in the section on the institutional aspects, costs for a firing ruled as unfair by the judge vary between 2.5 and 6 months of salary for small to up to 15 months in addition to the forgone compensation between firing and the court's ruling for large firms. Our value of 8 months lies in this range.

Using the same calculations, moving costs are around 1,720 euros, equal to 1.2 months of salary. It is harder to assess how plausible this value is, because there is not even indirect evidence on mobility costs. Our estimates imply that the cost of changing a job are a little more than one month of salary. This value seems rather modest. For example, using data from the CPS, Lee and Wolpin (2004) estimates a value of switching *sector* that can be as high as 75% of annual salary. Differently from them, we are not restricting the analysis to workers that change sector. Moreover, most of our job changes take place locally. This is documented in Table 7, that reports the share of workers moving within a given geographical area when changing job. Approximately one third of job changes are confined within the same municipality, more half within the same Local Labor System (LLS),[18] two thirds within the same province, and less that 20 percent entails a change of macro area (North-East, North-West, Center, South and Islands). If we exclude neighboring macro

[17] Previous studies have also find significant costs of adjusting employment. For example, using direct measures the costs of termination from survey data for France in 1992, Abowd and Kramarz (2003) find values in the range of 17,000 and 40,000 euros. Differently from their study, our measure also incorporates any indirect cost of adjusting, such as that coming from productive disruption; moreover, it represents the sum of both the internal and the external adjustment cost.

[18] LLS are defined as groups of municipalities characterized by a self-contained labor market, as determined by the National Statistical Institute on the basis of the degree of work-day commuting by the resident population. Using 1991 census data, the NSI procedure identified 784 LLSs covering the whole national territory.

areas, where mobility might still be local for workers located close to the boundaries, then long distance mobility is even lower. In Panel B of the table we report mobility flows across macro areas. The only substantial flows that surely entail long distance mobility are those from the poor and high unemployment regions of the South and Islands to the rich regions of the North. These flows are part of a secular migration movement South to North that characterizes the Italian labor market. If we exclude this flow, there is little evidence that firms satisfy their employment needs on the whole national territory, as seems more common in the US labor market: most of the job changes occur locally. The picture that emerges is therefore one of a fairly high segmentation is across local labor markets, while mobility within each market is rather inexpensive.[19]

All in all, our estimates suggest that adjusting employment entails non trivial both internal and external costs, with the former playing a more prominent role.

7 Conclusions

In this paper, we have proposed and implemented a method for distinguishing between the internal and the external component of the costs of adjusting employment at the level of the firm. We find that the external costs, ignored by the previous literature, are non trivial, but that the internal one account for a larger proportion of total costs.

These results have important policy implications. Adjustment costs imply that labor might not be allocated to its most productive utilization. Reducing them would imply a more efficient allocation of resources. From this perspective, our results indicate that, while mobility is an issue, the larger gains would occur from reducing adjustment costs internal to the firm, such as firing restrictions or other impediments to employment changes.

This does not imply that mobility costs are unimportant. Indeed, as already stressed above, our estimates of the moving costs should be interpreted as related to costs of changing employer, rather than location or sector. There is a widespread evidence that geographical mobility is low in Italy (Faini et al. 1996). Indeed, an important extension to this paper

[19] This feature is consistent with the traditional view of industrial clusters, where workers move fairly easily within the local market. This is one of the main features of industrial district, that are an important component of the Italian economy (see, for example, Guiso and Schivardi 1999).

would be to apply our method to measure the cost of long-range mobility. This will require to extend the analysis using local labor markets as the unit of observation rather than the firm, a task that we plan to undertake in future work.

A Appendix: Model solution

Consider first the optimal adjustment level. It is easier to work with the marginal shadow value of labor, that in equilibrium also follows the two state structure:

$$\begin{aligned} V_g &= F_l(l_g, \varepsilon_g) - w_g + \beta[pV_g - (1-p)V_b] \\ V_b &= F_l(l_b, \varepsilon_b) - w_b + \beta[pV_b - (1-p)V_g] \end{aligned}$$

Optimality requires that $V_b = -k$ and $V_g = h$: by firing an additional worker a firm pays k, so it will fire workers until the marginal product of labor is $-k$; similarly, when hiring it pays h, so the marginal worker must be worth exactly h. Substituting we obtain:

$$h = F_l(l_g, \varepsilon_g) - w_g + \beta[ph - (1-p)k] \tag{21}$$

$$-k = F_l(l_b, \varepsilon_b) - w_b + \beta[(1-p)h - pk] \tag{22}$$

By using $F_l = \varepsilon - \phi l$ and after some algebra (6) follows.

To obtain equilibrium levels, note that (3), (21), (22) and the condition $l_g + l_b = 1$ are four equations in four unknown that can be solved out to yield:

$$\begin{aligned} l_g &= \frac{\Delta\varepsilon - \theta(c+k+h) + \phi}{2\phi} \\ l_b &= \frac{-\Delta\varepsilon + \theta(c+k+h) + \phi}{2\phi} \\ w_g &= \frac{1}{2}(\varepsilon_g + \varepsilon_b + \theta c - (1-\beta)(h-k) - \phi) \\ w_g &= \frac{1}{2}(\varepsilon_g + \varepsilon_b - \theta c - (1-\beta)(h-k) - \phi) \end{aligned}$$

Consider now the optimality of adjusting. To save on notation, define $V_{ji} = V(l_i, \varepsilon_j)$. Then, the recursive equations for the value of the firm in the four states, conditional on adjustment being optimal are:

$$V_{gg} = F_{gg} - w_g l_g + \beta(pV_{gg} + (1-p)V_{gb}) \tag{23}$$

$$V_{bb} = F_{bb} - w_b l_b + \beta(pV_{bb} + (1-p)V_{bg}) \tag{24}$$

$$V_{gb} = V_{bb} - k\Delta l \tag{25}$$

$$V_{bg} = V_{gg} - h\Delta l \tag{26}$$

Then, substitute to reduce the system to two equations in two unknowns:

$$V_{gg} = F_{gg} - w_g l_g + \beta(pV_{gg} + (1-p)V_{bb} - (1-p)k\Delta l) \tag{27}$$

$$V_{bb} = F_{bb} - w_b l_b + \beta(pV_{bb} + (1-p)V_{gg} - (1-p)h\Delta l) \tag{28}$$

To determine the optimality of increasing employment when productivity switches from ε_b to ε_g we use Bellman's optimality principle, and check if deviating from one period delivers a higher payoff with respect to following the optimal policy. For adjustment to be optimal, it must be that

$$V_{bg} = F_{gg} - w_g l_g + \beta(pV_{gg} + (1-p)V_{bb} - (1-p)k\Delta l) - h\Delta l \geq$$
$$> F_{bg} - w_b l_b + \beta\{pV_{gg} - ph\Delta l + (1-p)V_{bb})\} \tag{29}$$

Simplifying,

$$F_{gg} - w_g l_g - (\beta(1-p)k + (1-\beta p)h)\Delta l \geq F_{bg} - w_b l_b \tag{30}$$

Using the quadratic production function and the relations $(l_g^2 - l_b^2) = (l_g - l_b)(l_g + l_b) = \Delta l$, $w_b = w_g - \theta c$ and substituting for w_g the condition simplifies to

$$\Delta\varepsilon^2 - \theta(c + 2(h+k))\Delta\varepsilon + \theta(\theta(c+k+h)(h+k) - c\phi) > 0$$

This is a quadratic in $\Delta\varepsilon$ whose solution is (8). Similar calculations yield the condition (9).

References

Abel, A. and Eberly, J. (1994), 'A unified model of investment under uncertainty', *American Economic Review* **84**, 1369–1384.

Abowd, J. M. and Kramarz, F. (2003), 'The costs of hiring and separations', *Labour Economics* **10**, 499–530.

Alonso-Borrego, C. (1998), 'Demand for labour inputs and adjustment costs: Evidence from spanish manufacturing firms', *Labour Economics* **5**, 475–497.

Belzil, C. (2000), 'Job creation and job destruction, worker reallocation and wages', *Journal of Labor Economics* **18**, 183–203.

Bentolilla, S. and Bertola, G. (1990), 'Firing costs and labor demand', *Review of Economic Studies* **57**, 381–401.

Bertola, G. (2004), 'A pure theory of job security and labor income risk', *Review of Economic Studies* pp. 43–61.

Borgarello, A., Garibaldi, P. and Pacelli, L. (2003), Employment protection legislation and the size of firms, LABORatorio Revelli Working papers no. 24.

Bronars, S. and Famulari, M. (2001), 'Shareholder wealth and wages: Evidence for white-collar workers', *Journal of Political Economy* **109**, 328–354.

Buchinsky, M., Fougère, D., Kramarz, F. and Tchernis, R. (2003), Interfirm mobility, wages, and the returns to seniority and experience in the u.s., mimeo, UCLA.

Cooper, R., Haltiwanger, J. and Willis, J. (2003), Dynamics of labor demand: Evidence from plant-level observations and aggregate implications, mimeo, Universtiy of Maryland.

Faini, R., Galli, G. and Rossi, F. (1996), Mobilità e disoccupazione in Italia: un'analisi dell'offereta di lavoro, *in* G. Galli, ed., 'La mobilit'a della società italiana', SIPI.

Fuentes, O., Gilchrist, S. and Rysman, M. (2004), Discrete adjustment costs, investment dynamics and productivity growth: Evidence from Chilean manufacturing plants, mimeo, Boston University.

Garibaldi, P. and Violante, G. (2005), 'The employment effects of severance payments with wage rigidity', *Economic Journal*. Forthcoming.

Guiso, L. and Schivardi, F. (1999), Information spillovers and factor adjustment, CEPR discussion paper n. 2289.

Guiso, L., Pistaferri, L. and Schivardi, F. (2005), 'Insurance within the firm', *Journal of Political Economy*. Forthcoming.

Hamermesh, D. (1999), 'LEEping into the future of labor economics: The research potential of linking employer and employee data', *Labour Economics* **6**, 24–41.

Hamermesh, D. S. and Pfann, G. A. (1996), 'Adjustment costs in factor demand', *Journal of Economic Literature* **34**, 1264–1292.

Heckman, J. (1990), 'Varieties of selection bias', *American Economic Review* **80**, 313–318.

Lee, D. and Wolpin, K. (2004), Intersectoral labor mobility and the growth of the service sector, mimeo, NYU.

Lucas, R. and Prescott, E. (1974), 'Equilibrium search and unemployment', *Journal of Economic Theory* **7**, 188–209.

Mortensen, D. and Pissarides, C. A. (1999), Job reallocation, employment fluctuations and unemployment, *in* M. Woodford and J. Taylor, eds, 'Handbook of Macroeconomics', North-Holland, pp. 1171–1228.

Nilsen, O. A., Salvanes, K. G. and Schiantarelli, F. (2003), Employment changes, the structure of adjustment costs, and plant size, B.C. Working Paper n. 586.

Robinson, J. (1933), *The Economics of Imperfect Competition , 1933*, Cambridge University Press.

Rota, P. (2004), 'Estimating labor demand with fixed costs', *International Economic Review* **45**, 25–48.

Schivardi, F. and Torrini, R. (2004), Threshold effects and firm size: the case of firing costs, Temi di discussione della Banca d'Italia n. 504.

Topel, R. (1986), 'Local labor markets', *Journal of Political Economy* **94**, S111–S143.

Table 1
Firms' and workers' characteristics

Panel A reports summary statistics for the firms in our data set. Panel B shows descriptive statistics for the sample of workers. All statistics refer to 1991. The matched firm sample includes firms that area matched at least once with a worker in the workers' data set.

Panel A: Firm characteristics

	Mean		Stand. dev.	
	Whole sample	Matched sample	Whole sample	Matched sample
Value added (thousand euros)	8712	15485	127028	116589
Number of employees	203	370	2355	2642
South	0.0884	0.0892	0.2839	0.2851
Center	0.1627	0.1672	0.3691	0.3731
North	0.7489	0.7436	0.4337	0.4367
Manufacturing	0.7750	0.7964	0.4176	0.4027
Construction	0.1549	0.1317	0.3619	0.3382
Retail	0.0253	0.0278	0.1571	0.1644
Services	0.0447	0.0441	0.2067	0.2052

Panel B: Workers' characteristics

	Mean		Stand. dev.	
	Whole sample	Matched sample	Whole sample	Matched sample
Earnings (thousand euros)	16.94	17.25	9.39	9.02
Age	38.93	39.15	10.43	10.40
Male	0.7284	0.7423	0.4448	0.4374
Productions	0.6164	0.6188	0.4863	0.4857
Clericals	0.3662	0.3655	0.4818	0.4816
Managers	0.0173	0.0157	0.1305	0.1242
South	0.1427	0.1244	0.3498	0.3301
Center	0.1880	0.1859	0.3907	0.3890
North	0.6693	0.6897	0.4705	0.4626

Table 2
Employment Adjustment: Ordered probit estimates

Dependent variable: a discrete variable taking the value -1 for negative employment changes, 0 for no changes, and 1 for positive changes. Firm shock is the residual in first differences of a regression of value added per (lagged) worker on year, sector, and regional dummies. Bootstrap standard errors in parenthesis.

Regressor	(1)	(2)
Lower threshold	-0.2357	-0.2037
	(0.0058)	(0.0110)
Higher threshold	0.1940	0.2148
	(0.0061)	(0.0109)
Firm shock	0.0195	0.0193
	(0.0005)	(0.0005)
N. periods since last adj.		0.0211
		(0.0073)
# observations	84,771	82,030

Table 3
Employment Adjustment: size of the adjustment

The dependent variable is employment change from one year to the next. The regression only include adjusting firms. See Table 3 for the definition of the firm shock. The other two variables are functionals of the normal p.d.f. and c.d.f. obtained from the estimates of the ordered probit of Table 3, column 1. Bootstrap standard errors are in parenthesis.

Regressor	(1)	(2)
Firm shock	2.0124	1.6643
	(0.8512)	(0.5549)
$\frac{\phi_H - \phi_L}{1 - \Phi_H + \Phi_L}$	352.2022	322.6161
	(329.8407)	(328.4444)
$\frac{1 - \Phi_H - \Phi_L}{1 - \Phi_H + \Phi_L}$	-144.4637	-116.9288
	(95.9984)	(74.6059)
# observations	70,611	68,678

Table 4
Wage adjustment

The dependent variable is yearly wage change (in thousand euros). For stayers, the change is computed as the difference in the wage from year to year; for movers, it is the change in the annualized wage following the job move. All regressions include sector (1 digit), year, and location (3 macro-areas) dummies. $s^+$is a dummy equal to 1 if the firm the worker is employed at has increased its workforce in the current year; $s^-$is one if it has decreased it. The variables $1-\Phi_{jt}^H+\Phi_{jt}^L$ and $\phi_{jt}^H+\phi_{jt}^L$ are computed from the ordered probit regression of Table 3. Bootstrap standard errors are in parenthesis.

	Specification (1)			Specification (2)		
Regressor	OLS (1)	2SLS (2)	Gen. 2SLS (3)	OLS (4)	2SLS (5)	Gen. 2SLS (6)
$s^+ - s^-$	0.1081 (0.0129)			0.1075 (0.0128)		
$1-\Phi_{jt}^H+\Phi_{jt}^L$		1.1750 (0.1056)	1.1751 (0.1061)		1.1585 (0.1095)	1.1584 (0.1085)
$\phi_{jt}^H+\phi_{jt}^L$			0.0808 (0.7901)			0.0401 (0.8124)
Age	−0.0043 (0.0011)	−0.0027 (0.0012)	−0.0027 (0.0012)	−0.0044 (0.0011)	−0.0023 (0.0012)	−0.0023 (0.0012)
Tenure	−0.0008 (0.0002)	−0.0013 (0.0002)	−0.0013 (0.0002)	−0.0008 (0.0002)	−0.0013 (0.0002)	−0.0013 (0.0002)
Male	0.2679 (0.0176)	0.2761 (0.0176)	0.2762 (0.0173)	0.2715 (0.0188)	0.2847 (0.0188)	0.2847 (0.0186)
Blue collar	−2.6912 (0.1741)	−2.8249 (0.1735)	−2.8251 (0.1736)	−2.6915 (0.1763)	−2.8362 (0.1759)	−2.8363 (0.1760)
White collar	−2.2323 (0.1750)	−2.3459 (0.1746)	−2.3460 (0.1747)	−2.2302 (0.1768)	−2.3492 (0.1764)	−2.3492 (0.1765)
N. obs.	104, 807	85, 689	85, 689	104, 807	84, 218	84, 218

Table 5
Wage Adjustment, Sensitivity analysis

The dependent variable is yearly wage change (in thousand euros). See Table 4 for details.

Regressor	Baseline	$-1 \leq \Delta l \leq 1$	Asymmetry
$1 - \Phi_{jt}^{H} - \Phi_{jt}^{L}$	1.1750 (0.1056)	0.9637 (0.1952)	
$1 - \Phi_{jt}^{H}$			0.8228 (2.0354)
$-\Phi_{jt}^{L}$			1.5273 (2.0360)
Age	−0.0027 (0.0012)	−0.0081 (0.0028)	−0.0027 (0.0013)
Tenure	−0.0013 (0.0002)	−0.0006 (0.0005)	−0.0013 (0.0002)
Male	0.2761 (0.0176)	0.1155 (0.0611)	0.2762 (0.0276)
Blue collar	−2.8249 (0.1735)	−2.6676 (0.1933)	−2.8251 (0.0898)
White collar	−2.3459 (0.1746)	−2.3183 (0.1976)	−2.3460 (0.0911)
N. obs.	85, 689	14,103	85,689

Table 6
Optimal Minimum Distance results

The table reports structural estimates of the parameters, obtained by applying OMD to the reduced form regression coefficients. The variance-covariance weighting matrix is obtained from the block bootstrap.

Structural estimates

	Specification (1)	Specification (2)
γ_H	7.6082 (0.2411)	8.9801 (0.4573)
γ_L	-9.2436 (0.2264)	-8.5160 (0.4588)
σ_ζ	39.2178 (3.9845)	41.8067 (3.5818)
ϕ	1.3044 (0.1499)	1.2363 (0.1186)
θc	1.1719 (0.1060)	1.1727 (0.1083)
θk	7.6642 (0.3166)	9.3340 (0.6039)
OID test	2.3923 (2 d.f.; p-value 30.23%)	2.9845 (2 d.f.; p-value 22.49%)

Table 7
Workers' geographical mobility

The first panel reports the share of workers that move within a given geographical unit. LLS are local labor systems (see footnote 18); Macro-areas are the ones reported in the second panel of the table. The second panel reports the matrix of mobility flows across macro areas.

Panel A: Share of mobility within:

Municipality	LLS	Province	Region	Macro-Area
.33	.54	.63	.74	.81

Panel B: Mobility Across Macro Areas

	To					
	N-W	N-E	Center	South	Islands	N. Obs.
From						
N-W	0.86	0.07	0.04	0.02	0.01	3627
N-E	0.11	0.84	0.03	0.01	0.01	2276
Center	0.18	0.06	0.68	0.07	0.01	1030
South	0.21	0.07	0.13	0.56	0.03	382
Islands	0.16	0.13	0.07	0.03	0.62	259

Figure 1: Equilibrium with no frictions, moving costs and firing costs

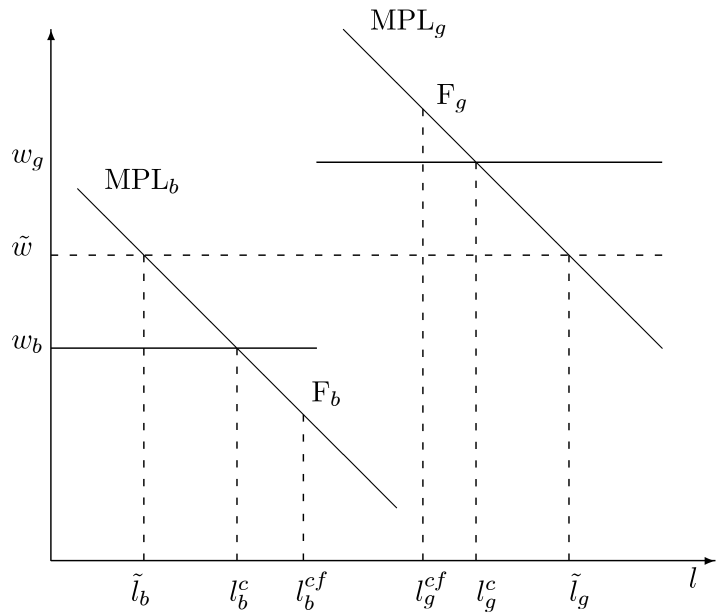

Note: $\tilde{w}$ is the wage that would prevail in a frictionless labor market, l^* is employment in the frictionless case, l^c with moving costs, l^{cf} with both firing and moving costs.

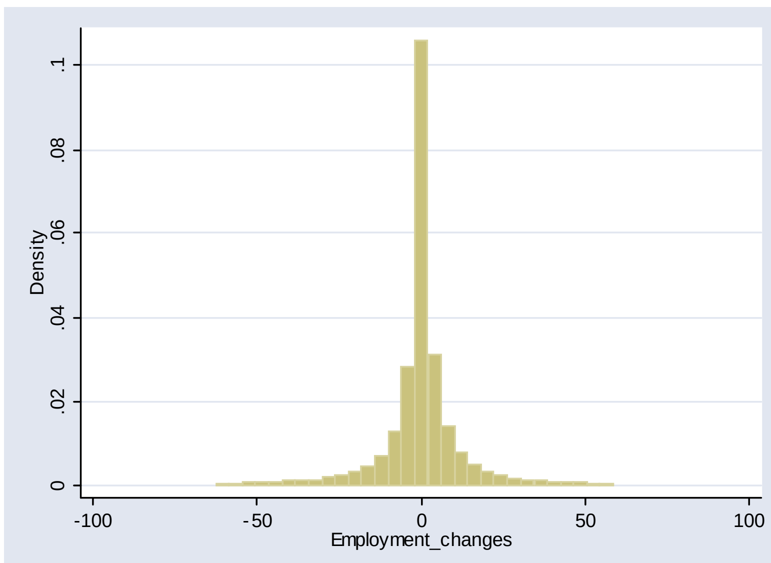

Figure 2: Histogram of employment changes