

The distributional effect of employment protection

Emanuele Albarosa

Abstract

In a quasi-experimental framework, I study how firing costs affect the relative outcomes of individuals according to their position on the productivity distribution. I argue that less productive people face a higher risk of dismissal, and are therefore more intensively affected by changes in firing costs. I analyse an Italian employment protection reform in a difference-in-differences setting, using both standard OLS techniques and the Unconditional Quantile Regression framework. My results indicate that the distribution of productivity estimates among new hires shifts substantially towards low values, as predicted. The productivity distribution of dismissed workers, instead, exhibits signs of a mean-preserving spread. The analysis underscores that employment protection has important distributional effects that should be taken into account by policy-makers.

1 Introduction

Labour market institutions are typically designed to protect and sustain jobs that would face the highest risks in an unregulated economic environment. In this light, their impact is inherently distributional: some jobs, individual and firms benefit more intensively than others. These dynamics are clear in the case of employment protection. Firing costs are designed to preserve jobs facing a relatively high risk of dismissal, which should be the most affected by employment protection reform. As long as individual and firms are not randomly selected into high risk jobs, changing firing costs will affect their *relative* outcomes in the labour market.

Using individual productivity estimates, I show that that less productive individuals are associated with higher dismissal rates. For this reason, I argue that reforms that decrease

employment protection should be expected to affect them the most. In particular, the existing literature highlights that lower firing costs increase the average hiring and firing rates. If the impact of firing costs is distributional, workers in the lower part of the productivity distribution should exhibit higher growth rates. Consequently, the distribution of new hires and dismissals should shift towards low-productivity individuals.

To test these predictions, I analyse a quasi experimental framework. I study an Italian employment protection reforms that took place in 2012, affecting small and large firms differently. Firms employing more than 15 employees benefited from a substantial decrease in firing costs associated with contracts signed after the reform. Firms below that threshold, on the other hand, experienced no change at all.

To estimate individual productivity levels, I resort to a methodology firstly introduced by (Abowd, Kramarz, & Margolis, 1999) (henceforth AKM). Using data prior to the main period used in the analysis, I estimate a log-wage equation from which I recover individual fixed-effects. These represent individual specific time-invariant contributions to the wage, once the effect of the other covariates has been netted out.

I analyse the effect of the reform on the distribution of productivity values among new hires and dismissed jobs. To do so, I first run a standard OLS difference-in-differences regression, which captures the shift in the distribution mean. To study the more general distributional impact, I apply the Unconditional Quantile Regression framework proposed by (Firpo, Fortin, & Lemieux, 2009). This allows me to estimate how the unconditional quantiles of the individual productivity distribution changed because of the exogenous decrease in firing costs.

My main datasource is the LoSai dataset, provided by the Italian Labour Ministry. It contains administrative information on the working history of several million workers over the last 30 years. The data it provides is considerably detailed and is reported as declared by firms for administrative matters, making measurement error a minor issue. Moreover, the dataset is of particular interest for this analysis as it links employee, contract and employer information. The length of the period covered and the granularity of the information provided allow me to control for a wide range of potential confounding factors, making the identification strategy more credible.

My analysis shows that employment protection has important distributional consequences. Both the OLS and the unconditional quantile analyses confirm that the distribution of the individual AKM effects among new hires shifts towards low values. This

supports the hypothesis that hiring rates grew with higher intensity among low-productive workers. Conceptually, this outcome supports the hypothesis that firms relax their hiring standards following decreases in firing costs. However, the distribution of productivity estimates among dismissed workers does not exhibit the predicted leftward shift. Instead, a mean-preserving spread of AKM values is observed following the employment protection reform.

These results complement and enrich the existing literature in many ways. Firstly, they expand the scope of the research on employment protection by highlighting its profound distributional implications. Many theoretical contributions, such as (Bentolila & Bertola, 1990), (Mortensen & Pissarides, 1994) and (Blanchard & Katz, 1997), have shown that firing costs are negatively related with hiring and firing rates. These findings have been subsequently corroborated by empirical evidence, as in (Boeri & Jimeno, 2005), (Schivardi & Torrini, 2008), (Kugler & Pica, 2008), (Marinescu, 2009) and (Quintini & Venn, 2013). However, these analyses mostly focus on the average effects of firing costs on the economy. On the contrary, (Butschek & Sauermann, 2019) touch on selection effects. They analyse a Swedish reform that reduced firing costs for some employers through a difference-in-differences setting, and find that it induced a reduction in minimum workers' ability among new hires. Their findings are in line with mine, and underscore that firms decrease their hiring requirements when employment protection laws are relaxed. From a methodological point of view, this paper makes two further contributions to the literature. Firstly, it adds to labour market research based on individual AKM fixed effects estimates. Empirical applications typically concern inter-industry wage differential and assortative matching between highly productive workers and firms - see (Goux & Maurin, 1999), Berth and Dale-Olsen (2003), (Gruetter & Lalive, 2004) and (Card, Heining, & Kline, 2013). Moreover, it also provides an empirical application of the Unconditional Quantile Regression framework, as proposed by Firpo, Fortin and Lemieux (2009). So far, this methodology has been used in several fields, included wage structures (Firpo, Fortin, & Lemieux, 2011), evaluations of child-care systems (Havnes & Mogstad, 2010) and the relation between income inequality and economic growth (Koske, Fournier, & Wanner, 2012).

2 Employment protection in Italy

The Italian labour law has historically discriminated between workers employed in small and large firms, granting the latter higher levels of employment protection. Specifically, the firing costs associated with wrongful dismissals carried out in firms employing more than 15 employees are substantially higher than in other firms. The sanctions depend both on the reasons for the dismissal (economic or disciplinary) and the type of legal violation (formal errors or discriminatory actions). The so-called Fornero reform substantially decreased firing costs for large firms, starting from July 2012. Over the next paragraphs, I first describe the structure of employment protection laws and then how the reform changed firing sanctions for permanent workers.

2.1 Pre 2012

Before July 2012, in most cases, a worker wrongfully dismissed by a large firm could choose between being reinstated and receiving a monetary compensation equal to 15 months of salary. If reinstated, the worker was also due all foregone salary and social security contributions that were not paid during the period in which they were dismissed.

On the contrary, for workers dismissed by small firms reinstatement is rare. Instead, these are typically required to pay a monetary sanction ranging from 2 to 6 months of salary, with higher limits for workers with many years of tenure.

2.2 The Fornero reform

Starting from July 18th, 2012, employment protection rules changed for all existing employment contracts. Firing costs were reduced for firms above the 15 employees threshold. The possibility of reinstatement was limited to dismissals that were discriminatory or whose reason (economic or disciplinary) was non-existent. Moreover, compensation for the lost income was capped at 12 months of salary for most cases. Reinstatement was instead substituted by a monetary compensation between 12 and 24 months of salary. Finally, there were some procedural changes as well. The period in which the worker could take the dismissal in court was reduced by one third, and a mediation phase aiming at resolving the dispute privately was made compulsory.

Permanent and temporary contracts In principle, legal actions can be taken against the termination of both permanent and temporary contracts. In practice, however, the potential rewards for workers who were employed on fixed-term contracts are very low. De facto, the above employment protection laws only apply to workers employed on open-ended contracts.

3 Data

The dataset I use is drawn from the LoSal database, provided by Italian Labour Ministry. It has several interesting features that makes it optimal for the analysis. Firstly, it provides administrative data as reported by firms and workers for social security purposes. Therefore, measurement error should be limited. Moreover, it is based on a sample of several million workers and includes a wide range of contract, worker and firm-related variables allowing to reconstruct in detail their entire career over a considerably long period of time.

The dataset contains administrative information originally collected by the Italian Social Security Institute. It provides the complete working history of around 2.7 million workers, from 1985 to the end of 2014. The data is collected by the Social Security administration for pension and social contributions purposes. It is gathered in a number of different ways, depending on the situation. General employee data is typically collected from the administrative information that firms must report to Social Security. The dataset, however, also contains information on self-employment activities carried out by the sampled workers, as well as any independent contractor relation they might have had with firms. In these cases, the information is reported to the Social Security administration directly by workers.

In my analysis, I focus on dependent employment contracts. Each observation in the dataset represent an employment contract between an employee and a firm in a given calendar year. I observe any employment contract related to the sampled workers that was signed, terminated or ongoing during the period. The data is representative of the entire population of private and public employees.

Contract information The dataset includes information on yearly wages, as reported by firms at the end of each fiscal year. This quantity represents therefore the salary that

was actually paid to a given employee. The yearly wage is rounded to the nearest hundreds, and includes any taxes and social contributions paid by the employee and the employer. The additional information on contracts is quite extensive. It includes the type of contract, the starting and ending date, the occupational category, and the number of paid days and weeks in the year. For part-time workers, I observe the share of part-time. I know whether any active or passive labour market policy was activated, such as wage, employment and apprenticeship subsidies, among others. For contracts ending during the sampled period, I observe whether the employee was dismissed, quit or the agreed termination date was reached (for fixed-term contracts).

Worker and firm information The data also reports employee and employer information. As for the employee, the information includes the birth (and potentially death) date, sex, age, and region of residence. I also have data on the unemployment benefits received during the period. Concerning the firm, I know how many workers it employs (in categories) and its 2-digit sector code. However, the information on firms is available only from 2005 onward. From the combination of contract, worker and firm data I am also able to recover work experience and job tenure. The former is defined as the number of years preceding each observation in which the employee has at least one active employment contract. The latter is a similar measure which only considers the employment contracts that the worker had within a given firm.

Sample restrictions I restrict the analysed period to the years between 2005 and 2014, as firm-size and sector information is not available for previous years. I only include observations related to employees between 15 and 64 years old. Finally, for some observations the number of paid days during the year, the firm size category and/or the region of residence were not reported. I simply choose to exclude them, as they represent a very small fraction of the data.

3.1 Estimating individual and firm productivity

To obtain measures of individual productivity, I follow the approach initially proposed by Abowd, Kramarz and Margolis (1999; henceforth AKM) and then built upon by Card et al. (2013) and Butshek and Sauermann (2019a). I first divide the sample in two 6-year periods:

2005-2009 and 2010-2014. Using the former subsample, I estimate AKM individual effects from the following wage equation:

$$\log(w_{i,f,t}) = X_{i,f,t}\beta + \alpha_i + \alpha_f + \alpha_t + \eta_{i,f,t} \quad (1)$$

where $w_{i,f,t}$ is the weekly wage. $X_{i,f,t}$ includes controls for the worker's age, sex, work experience and region of residence. I further net out the effects of job tenure, occupation, part-time share, economic sector, firm-size category and type of contract (permanent, temporary or seasonal). α_i , α_f and α_t represent individual, firm and time fixed-effects.

The AKM effects α_i represent the time-invariant individual specific contributions to the wage, once the effect of all other controls has been netted out. They capture productivity levels insofar as these are reflected in higher wages. Clearly, labour market frictions and institutions will prevent wages to reflect exactly the marginal product of labour. However, these two factors should still be positively related, making the AKM effects proxies for productivity.

At this point, I assign to each observation in the 2010-2014 sample its correspondent AKM individual effect, if they exist. I restrict my main analysis to observations in this sample. This allows me to use estimated productivity levels that are pre-determined with respect to my period of interest. However, it also means that only individuals that are observed at least once between 2005 and 2009 will be assigned an AKM effect. Furthermore, in some cases I cannot assign both an individual and firm AKM effect, as some individuals and firms are observed only in a single match. Overall, I am able to assign an estimated productivity value to 70% of the individuals and 72% of the firms observed in the sample.

Empirical performance How well does this measure approximate individual ability? From a theoretical point of view, Card et al.(2013) show that assumptions on which the measure is based are largely supported by the data. In particular, AKM effects assume strong separability between individual and firm effects. The authors show that non-separable models with match fixed effects only slightly improve the explanatory power of the above specification. On the other hand, Butshek and Sauermann (2019b) compare AKM effects and more classical measure of individual ability, such as cognitive test scores and GPAs. They find a correlation varying between 30 and 60%.

A potential issue with this estimation is the absence of education among the controls.

Unfortunately, the data at my disposal does not contain information on education levels, which means that I cannot include them in the estimation. Insofar as education is time-invariant, its effect on wages could be captured by the AKM effect. However, education is thought to be positively related both to innate ability and to work-related productivity itself. As long as this is the case, the relation between the estimated AKM effects and productivity should not be significantly altered.

3.2 The samples

After having estimated the AKM effects is performed on all available observations, I further subset the 2010-2014 sample. Firstly, I exclude any non-permanent contracts. As I argue in the previous section, temporary and seasonal ones were not affected by any of the two reforms under study. Moreover, I only consider jobs in firms between 11 and 20 employees.

To analyse the distributional hiring effect of the two employment protection reforms under study, I only consider new hires. These are defined contracts observed in their first year of existence. As I only consider contracts that are assigned an AKM individual effect, the resulting sample includes around 38.000 observations. Similarly, to examine the firing impact of employment protection, I consider a second sample that only includes dismissals. These are contracts observed in their last year of existence and for which the reason for termination was registered as dismissal. This sample includes around 20.000 observations. Tables 1 and 2 display some descriptive statistics for both samples. The first line provides information on weekly wages. The average weekly wages are not dissimilar in the two samples, although dismissed workers are on average both older and more experienced than new hires. The fact that men account for two thirds of both samples is not surprising. Female employment levels have been historically low in Italy, especially in Southern regions. Part-time jobs represent a non-trivial share of the total, around one fifth of dismissals and one fourth of new hires. The samples exhibit similar shares of blue and white collars, and approximately the same geographical distribution. As for the economic sector, two features seem worthy of mention. Firstly, the share of Industry jobs in both samples is surprisingly high. In fact, although the industrial sectors represent a around 23% of Italian employment, around half of the observations are of industrial jobs. Secondly, industrial jobs are also over-represented among dismissals, relative to new hires. Both features might be explained by the fact that typically permanent jobs are under-represented in the Services

Table 1: Descriptive statistics - New hires

Statistic	N	Mean	St. Dev.
Weekly wage	43,381	221.456	193.547
Age	43,381	35.672	11.121
Female	43,381	0.308	0.461
Part-time workers	43,381	0.257	0.437
Years of work experience	43,381	9.784	7.774
Blue-collar worker	43,381	0.665	0.472
White-collar worker	43,381	0.183	0.387
Manager	43,381	0.009	0.097
Resident in Northern regions	43,381	0.494	0.500
Resident in Central regions	43,381	0.193	0.395
Resident in Southern regions	43,381	0.313	0.464
Sector = Industry	43,381	0.423	0.494
Sector = Services	43,381	0.577	0.494

This table displays some descriptive statistics for new hires. The latter are defined as contracts observed in their first year of existence. Only new hires for which the AKM individual effect is identified are included. The sample includes observations relating to 39573 individuals and 28006 firms.

Table 2: Descriptive statistics - Dismissals

Statistic	N	Mean	St. Dev.
Weekly wage	20,849	229.992	214.365
Age	20,849	40.039	10.807
Female	20,849	0.277	0.448
Part-time workers	20,849	0.185	0.388
Years of work experience	20,849	13.268	8.006
Blue-collar worker	20,849	0.762	0.426
White-collar worker	20,849	0.184	0.387
Manager	20,849	0.007	0.082
Resident in Northern regions	20,849	0.399	0.490
Resident in Central regions	20,849	0.186	0.389
Resident in Southern regions	20,849	0.415	0.493
Sector = Industry	20,849	0.532	0.499
Sector = Services	20,849	0.468	0.499

This table displays some descriptive statistics for dismissals. The latter are defined as contracts observed in their last year of existence, only when the reason for the termination is a dismissal. Only observations for which the AKM individual effect is identified are included. The sample includes observations relating to 19834 individuals and 14719 firms.

sector. On one hand, as I only include permanent contracts, Services jobs should appear less frequently. On the other, dismissals mainly concern permanent contracts, as temporary ones end "naturally" on their agreed termination date.

3.3 AKM productivity estimates

This subsection provides descriptive evidence on the AKM effects estimated as outlined above. As can be seen in Figure 1, individual AKM productivity estimates among new hires and dismissals are approximately normally distributed. Their mean is approximately the same; AKM effects in the dismissal sample, however, exhibit a higher variance.

Figure 2 further shows that the individual AKM effects are negatively correlated with the firm fixed-effects from Equation 1. Although this might seem surprising at first, it is a common feature of virtually all the literature estimating the individual and firm contributions to wage levels and growth. Goux and Maurin (1999), Barth and Dale-Olsen (2003), Abowd et al. (2004), as well as Gruetter and Lalive (2004) all find negative correlation values using linked employer-employee data from different countries. Andrews et al. (2008) show that an important cause of this unintuitive relation lies in the nature of the econometric technique used to estimate the AKM fixed effects. They show that, when the covariates in the wage regression are weakly correlated with the worker and firm dummies, the correlation between the latter two is negative even when in reality there is positive assortative matching between individual and firms. Abowd et al. (2004) also suggest that this correlation-bias is stronger when the number of job movers is limited.

AKM effects and firing rates The hypothesis of the distributional impact of employment protection is based on the assumption that firing rates depend on the productivity level of the individual and the firm involved in a job match. Is there any evidence in favour of this assumption, based on the estimated AKM effects? Figures 3 displays the yearly shares of dismissal, by quintile of the individual and firm AKM effect distributions. They clearly show that, based on the AKM estimates, the above hypothesis is clearly supported. The case of firms is evident: throughout the period, jobs in lower quintiles of the firm productivity distribution exhibit higher dismissal shares. The relation between dismissals and individual productivity estimates is not as clearcut, as for instance, jobs in the 3rd quintile of the distribution are more often dismissed than those in the 1st and 2nd ones. However,

the lowest dismissal shares are indeed observed in the upper part of the distribution.

4 Empirical strategy

The Fornero reform decreased firing costs associated with workers employed by firms above the 15 employees threshold. Specifically, it lowered employment protection for all existing jobs in those firms, starting from July 18, 2012. I use this quasi-experimental framework to estimate the effect of firing costs on the distribution of individual and firm productivity levels.

The main hypothesis of this paper holds that lowering firing costs should have a relatively strong positive effect on hiring and firing rates of low-productive workers and by low-productive firms. If that was the case, the distribution of hires and dismissals made after the reform should exhibit a shift towards the low-end of the productivity distribution.

In this section, I outline my difference-in-differences identification strategy. In a standard OLS framework, I test whether the employment protection reform changed the mean of the AKM effect values among workers hired or dismissed. Moreover, I specify the assumption on which the identification strategy relies and describe how I gather evidence on it. Finally, I provide a brief review of the Unconditional Quantile Regression (UQR) method proposed by Firpo et al. (2009). I show how to use a difference-in-differences UQR strategy to evaluate how the unconditional distribution of AKM effects was impacted by the change in firing costs.

4.1 Difference-in-differences strategy

To estimate the impact of the reform, I compare the individual AKM effect distribution of contracts belonging to firms above the 15 employees threshold with those in firms below it, before and after the Fornero reform.

Identification equation Let any new employment contract be associated with an individual i , a firm f and a date t . When estimating the impact of the reform on new hires, t represents the starting date of the contract. Similarly, when studying the effect on dismissals, it indicates the ending date. I model the AKM value of the worker, conditional on the covariates $X_{i,f,t}$, as:

$$AKM_i = \gamma_0 + \gamma_1(\text{Size}_{i,f,t}^{16-20} \times \mathbb{1}\{t \geq \text{July 18, 2012}\}) + \gamma_2\text{Size}_{i,f,t}^{16-20} + \gamma_3 \times \mathbb{1}\{t \geq \text{July 18, 2012}\} + X'_{i,f,t}\beta + \alpha_f + \alpha_t + \epsilon_{i,f,t} \quad (2)$$

$\text{Size}_{i,f,t}^{16-20}$ indicates whether the contract concerns a firm above the 15 employees threshold. The reference category - and the control group - includes firms employing between 11 and 15 employees.

$X_{i,f,t}$ includes worker, firm and contract covariates. I control for years of work experience and job seniority, as well as the occupational category and economic sector. I also include the worker's age, sex, region of residence, the part-time share and time dummies. As the individual AKM values are time-invariant, I can only include firm fixed effects.

The effect of the reform on the mean of individual AKM productivity values is given by $\hat{\gamma}_1$. It represents the post-reform mean variation, in firms above the threshold relative to firms below it. The standard errors are clustered by both individual and firms.

Identification assumption The assumption my analysis rests upon is a standard parallel trends one, widely made in difference-in-differences frameworks. In this context, it translates to assuming that the average productivity values of new and dismissed contracts would have followed the same trend in firms both below and above the 15 employees threshold, had the reform not happened. To gather evidence on the likelihood of this assumption, I estimate the mean AKM effect of contracts started (or dismissed) in a given quarter during the sampled period, in firms above the 15 employees threshold relative to firms below it. This allows me to analyse whether trends in AKM values of treated and control observations significantly diverged before the reforms.

Figure 4 display the results of this analysis for both new hires and dismissals. In both cases, no coefficient is significant before the 3rd quarter of 2012, when the reform period starts. Although this result is not a formal test of the identification assumption, it does bring evidence in favour of its likelihood. Up to the beginning of the employment protection reform period, contracts started or dismissed seem to have followed the same AKM effect trends in firms both above and below the 15 employees threshold.

4.2 Unconditional quantile regression

To evaluate how the reforms changed the distribution of the AKM individual effects beyond its mean, I follow the unconditional quantile regression (UQR) proposed by Firpo et al. (2009). In their paper, they propose a method to estimate the effect of a treatment D on the unconditional distribution of an outcome variable y , $F^{-1}(y | D = 1) - F^{-1}(y | D = 0)$. They show that this effect can be estimated by regressing the Recentered Influence Function (RIF) of a distributional statistics on the treatment and other covariates. The Influence Function $IF(y; v(y))$ measures the effect of removing/adding an observation to a statistics $v(y)$. The RIF, instead, is simply equal to the sum of the statistics $v(y)$ and $IF(y; v(y))$. In particular, the expected RIF for the τ -quantile, q_τ , conditional on the covariates X , is equal to:

$$E[RIF(y; q_\tau) | X = x] = q_\tau - \frac{1 - \tau}{f_Y(q_\tau)} + \frac{\Pr\{y > q_\tau | X = x\}}{f_Y(q_\tau)} \quad (3)$$

where $F(Y)$ is the unconditional cumulative distribution of Y . It turns out that $q_\tau(y | D = 1) - q_\tau(y | D = 0)$ is equal to the marginal effect of D on $E[RIF(y; q_\tau) | X = x]$. This happens because, thanks to one of the properties of RIFs, $E_X[E[RIF(y; q_\tau) | X]] = q_\tau$.

Equation 3 underscores the useful relation between RIF-regression and linear probability models. In fact, the effect of any binary treatment on $E[RIF(y; q_\tau) | X = x]$ is equal to the coefficient of a linear probability model of the probability that $y > q_\tau$, rescaled by the appropriate $f_Y(q_\tau)$. q_τ and $f_Y(q_\tau)$ are estimated through \hat{q}_τ , the sample quantile, and $\hat{f}_Y(\hat{q}_\tau)$, the sample kernel-density at \hat{q}_τ .

In this context, to evaluate the distributional effects of the employment protection reforms under study, I estimate additional UQR regressions based on the same identification equation outlined above. The dependent variable, however, changes. Instead of AKM_i , the outcome variable is $\mathbb{1}\{AKM_i > \hat{q}_\tau(AKM)\}$. As explained above, the unconditional effects of the reform is given by the difference-in-differences coefficient, rescaled by $f_{AKM}(\hat{q}_\tau)$.

Why unconditional quantiles? The UQR framework is not the only possible approach when examining counterfactual distributions. In fact, it is rather recent; traditionally, the preferred methods would rather focus on Conditional Quantile Regressions (CQR). However, I argue that UQR analysis is conceptually more appropriate and more easily interpretable in this context. In fact, this paper focuses on the market-wide distri-

butional consequences of employment protection. That is, its aim is to analyse the effect of firing costs on the relative outcomes of low and highly productive workers in the labour market. From this perspective, I argue that it makes more sense to examine unconditional distributional effects, rather than conditional ones. Moreover, the results are interpretable in a more straightforward way, as they refer to the general and directly observable (unconditional) distribution of productivity levels.

5 Results

In the following paragraphs, I present the results of my analysis. Firstly, I show the graphical analysis of the post reform raw distributional changes. Then, I proceed to discuss the OLS regressions. Finally, I present the outcome of the unconditional quantile regression analysis.

5.1 Raw data analysis

Figure 6 plots the raw difference-in-differences distributional effect of the reform, for both new hires and dismissals. I first compute the AKM density gap between the contracts that are covered by the employment protection reform and those that are not. I compute the gaps separately for contracts signed in firms above or below the threshold. The graphs then display the difference between these two latter groups. The greyish areas indicate the magnitude of the raw density difference-in-differences.

The plot suggests that following the decrease in firing costs the distribution of AKM individual effect shifted towards low values, as expected. The raw difference-in-differences estimator is positive on the lower part of the distributions, and falls to negative values approximately after the mean. These raw data trends are consistent with the research hypothesis, i.e. that lower firing costs increase the hiring rate of low-productive workers relatively more than the average.

The case of dismissals is different. The raw difference-in-differences is positive over both tails of the distribution, while strongly negative around the median. Clearly, this raw picture is not consistent with a lefthand shift in the distribution of AKM individual effects. It rather suggests a sort of mean-preserving spread, which would indicate an increase in the variance of productivity levels among dismissed workers.

5.2 OLS regressions

Tables 3 and 4 display the results of OLS analysis. They report the coefficients of standard difference-in-differences regressions, in which the outcome variable is the AKM individual effect. The coefficients indicates the estimated shift in the mean of the AKM distribution. The different columns show how the coefficients vary when including controls for worker related characteristics, job and firm covariates, time dummies and firm-fixed effects.

The results concerning new hires seem to support the research hypothesis. The difference-in-differences coefficients are always negative and statistically significant across specifications, except when I include firm-fixed effects. This, however, might be because only very few firms in the sample appear more than once, thus eliminating most of the variation in the data. The magnitude of the coefficients, however, is small. The estimated coefficients equal around 2% of the overall sample AKM-mean.

On the contrary, the results of the regressions on the dismissal sample clearly indicate no shift in the mean of the distribution. In fact, the coefficients are never significant. Depending on the specification, they are either positive or negative. In terms of magnitude, however, they are very close to 0. Only when I introduce firm-fixed effect does the size of the coefficient increase. However, it still is not significant.

These results provide mixed evidence on the distributional impact of employment protection. Among new hires, the mean of the productivity distribution shifted towards low-value following the reform. However, no such shift can be detected among dismissals.

5.3 Unconditional quantile regressions

Figures 9a to 15b display the results of the unconditional quantile regressions on the AKM individual effect distribution among new hires and dismissals. They plot the average coefficients for the nine deciles of the AKM individual effect distribution, together their associated 95% confidence intervals. These coefficient measure the variation in the unconditional deciles of the AKM effect distribution due to the reform, once the effect of the covariates have been netted out. The standard errors are bootstrapped by clustering at the intersection of individual and firm.

New hires The results of the unconditional quantile regression are sizeable and rather stable across specifications. The coefficients are negative for the entire distribution; the

only exception of the 9th decile, which becomes statistically non significant once the worker covariates are accounted for. The magnitude of the significant coefficients ranges between 5 and 10% of a standard deviation. In particular, the coefficients are highest for the first decile of the distribution, and decrease for subsequent ones. These estimates are stable across all specifications: the inclusion of worker, firm and job characteristics, as well as time dummies, does not influence neither the significance nor the size of the estimates.

The fact that the coefficients are negative and significant across all quantiles, together with the shift in the mean estimated through the OLS regression, strongly supports the idea of a leftward shift in the productivity distribution among new hires.

Dismissals The outcomes of the UQR on the dismissal sample are more nuanced. The unconditional quantile coefficients are once again significant across specifications. Moreover, their magnitude is not much affected by the inclusion of any covariates. However, it is smaller than in the case of new hires, as it ranges between 2 and 4% of a standard deviation.

The interesting feature of the results is that, while the unconditional quantile coefficients are negative up to the median, they become positive in the upper part of the distribution. The coefficients for the 6th to 8th deciles even stay significant in all specifications. This seems to be consistent with what was suggested by the analysis of raw difference-in-differences, which indicated a mean preserving spread of the distribution. Therefore, it appears that individuals in both tails of the productivity distribution experienced a relatively strong increase in firing rates. Despite the fact that the unconditional quantile effect seems to be stronger in the lefthand part of the distribution, these dynamics cannot be reconciled with the hypothesis of a shift towards low productivity values.

6 Conclusions

This analysis undertakes to study the distributional implications of employment protection. While the latter is a much discussed and polarising topic, the existing research rarely focuses on its effect on relative outcomes in the labour market. Firing costs affect especially jobs that suffer from a high risk of dismissal. As individual in the lower part of the productivity distribution are likely to be selected in those jobs, employment protection reforms might change substantially their relative standing in the labour market.

In this paper, I take advantage of a rich Italian linked employer-employee dataset to estimate individual productivity values, using a log wage regression in the spirit of (Abowd et al., 1999). This allows me to change how the distribution of productivity-proxies changes following an employment protection reform that affected firms differently based on their size. Through a difference-in-differences setting, I estimate the shift in the mean of the productivity distribution among new hires and dismissed workers. Moreover, I use Unconditional Quantile Regression techniques to estimate the changes at other points of the distribution.

The results highlight two facts. First, the productivity distribution among new hires does indeed shift towards low values. This supports the view that hiring rates grow more intensively for less productive workers, following a decrease in firing costs. In turn, this effect is consistent with a decrease in firms' hiring standards. Secondly, firing rates seem to increase at both tails of the productivity distribution. While this second dynamic is not consistent with the research hypothesis of this paper, it does highlight that the effect of firing costs is essentially distributional.

The implications of these findings are many. On one hand, they provide an original approach to the study of labour market reforms, which takes into account the individual's estimated productivity level. Secondly, they highlight that labour market institutions should be studied also through the lenses of distributional changes. Finally, they are of particular importance for policy making. Whether policy-maker focus more on the potential gains or losses, they should take into account that the strongest repercussions of employment protection reforms are likely to fall on the least productive individuals.

References

- Abowd, J. M., Kramarz, F., & Margolis, D. N. (1999). High wage workers and high wage firms. *Econometrica*, *67*(2), 251–333.
- Bentolila, S., & Bertola, G. (1990). Firing costs and labour demand: how bad is eurosclerosis? *The Review of Economic Studies*, *57*(3), 381–402.
- Blanchard, O., & Katz, L. F. (1997). What we know and do not know about the natural rate of unemployment. *Journal of Economic Perspectives*, *11*(1), 51–72.
- Boeri, T., & Jimeno, J. F. (2005). The effects of employment protection: Learning from variable enforcement. *European Economic Review*, *49*(8), 2057–2077.

- Butschek, S., & Sauermann, J. (2019). The effect of employment protection on firms' worker selection.
- Card, D., Heining, J., & Kline, P. (2013). Workplace heterogeneity and the rise of west german wage inequality. *The Quarterly journal of economics*, *128*(3), 967–1015.
- Firpo, S., Fortin, N. M., & Lemieux, T. (2009). Unconditional quantile regressions. *Econometrica*, *77*(3), 953–973.
- Firpo, S., Fortin, N. M., & Lemieux, T. (2011). Occupational tasks and changes in the wage structure.
- Goux, D., & Maurin, E. (1999). Persistence of interindustry wage differentials: a reexamination using matched worker-firm panel data. *Journal of Labor Economics*, *17*(3), 492–533.
- Gruetter, M., & Lalive, R. (2004). The importance of firms in wage determination.
- Havnes, T., & Mogstad, M. (2010). Is universal child care leveling the playing field? evidence from non-linear difference-in-differences. iza discussion papers 4978. *Institute for the Study of Labor (IZA)*.
- Koske, I., Fournier, J.-M., & Wanner, I. (2012). Less income inequality and more growth—are they compatible? part 2. the distribution of labour income.
- Kugler, A., & Pica, G. (2008). Effects of employment protection on worker and job flows: Evidence from the 1990 italian reform. *Labour Economics*, *15*(1), 78–95.
- Marinescu, I. (2009). Job security legislation and job duration: Evidence from the united kingdom. *Journal of Labor Economics*, *27*(3), 465–486.
- Mortensen, D. T., & Pissarides, C. A. (1994). Job creation and job destruction in the theory of unemployment. *The review of economic studies*, *61*(3), 397–415.
- Quintini, G., & Venn, D. (2013). Back to work: re-employment, earnings and skill use after job displacement. *Final report, OECD, October*.
- Schivardi, F., & Torrini, R. (2008). Identifying the effects of firing restrictions through size-contingent differences in regulation. *Labour Economics*, *15*(3), 482–511.

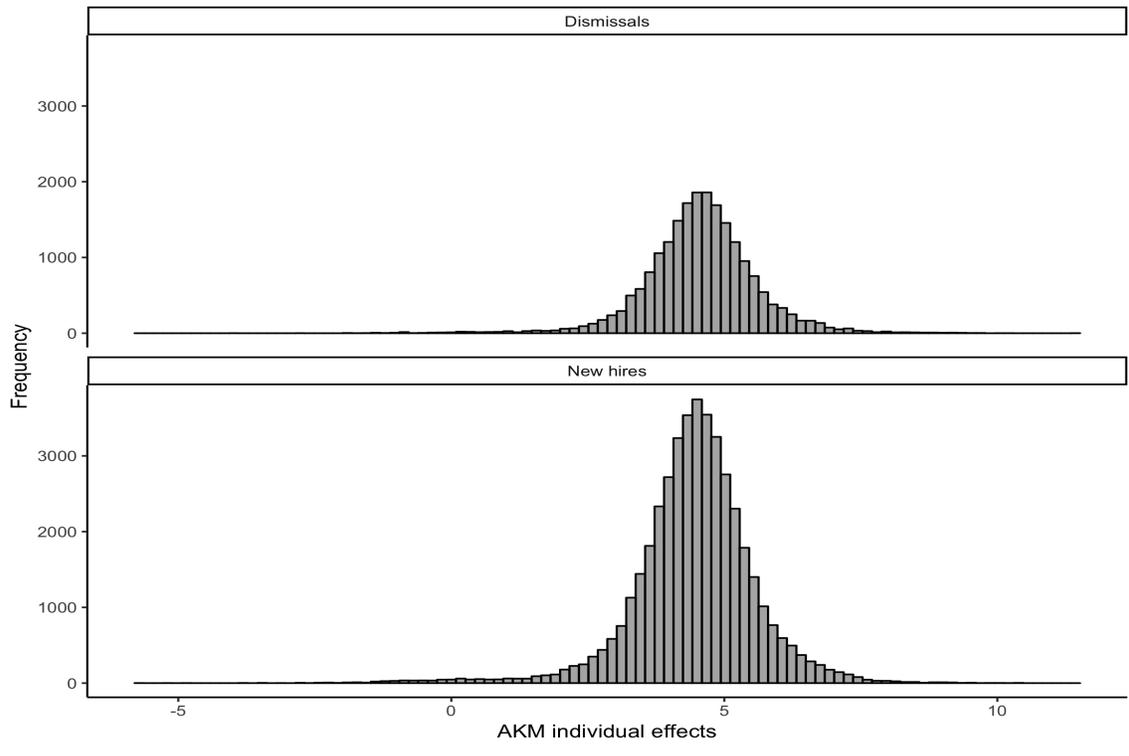


Figure 1: Distributions of AKM individual effects

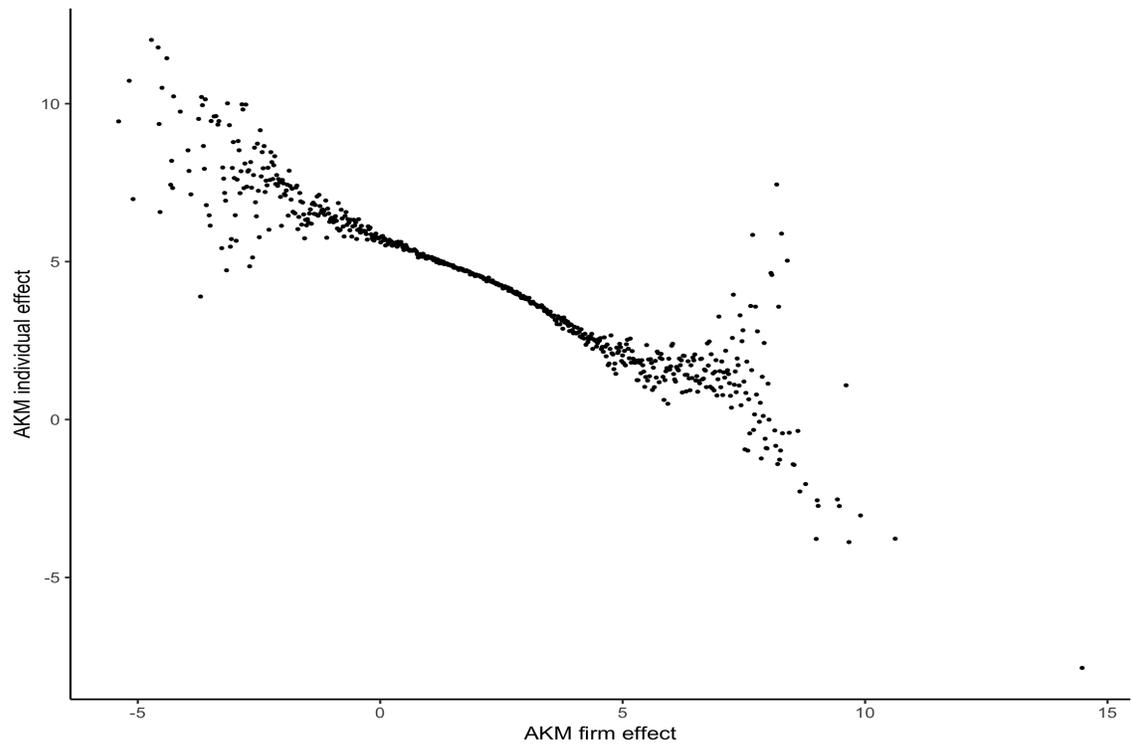


Figure 2: Binned scatterplot of AKM individual and firm fixed-effects.

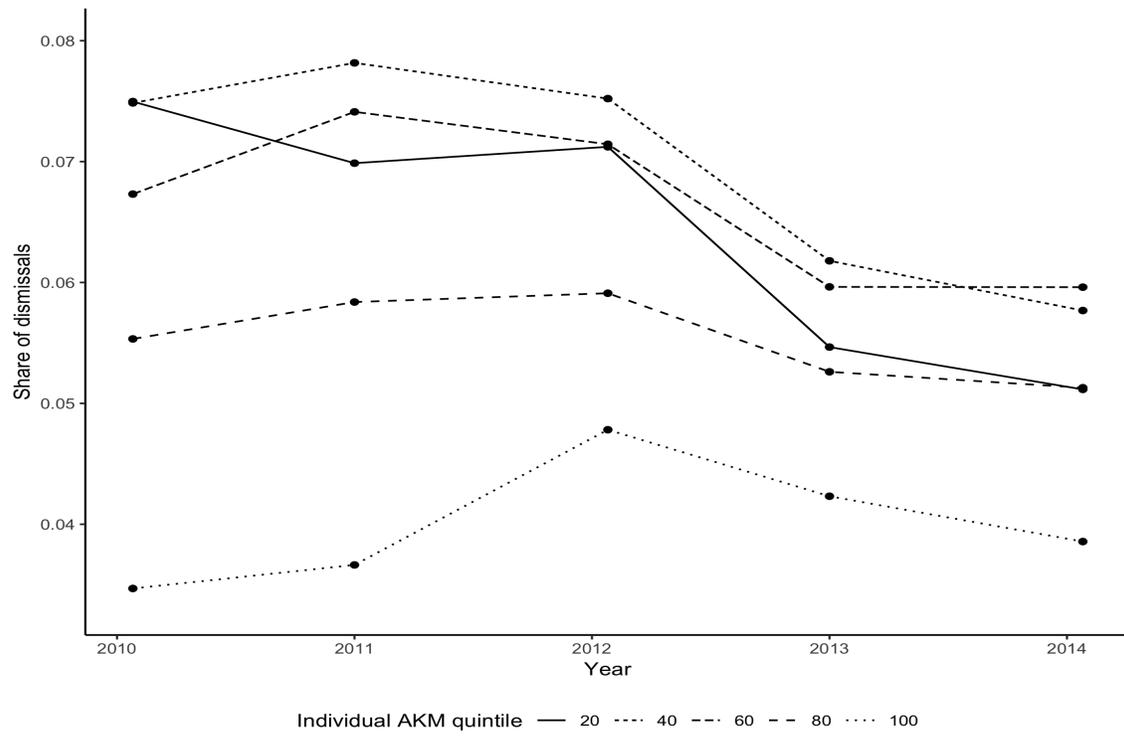
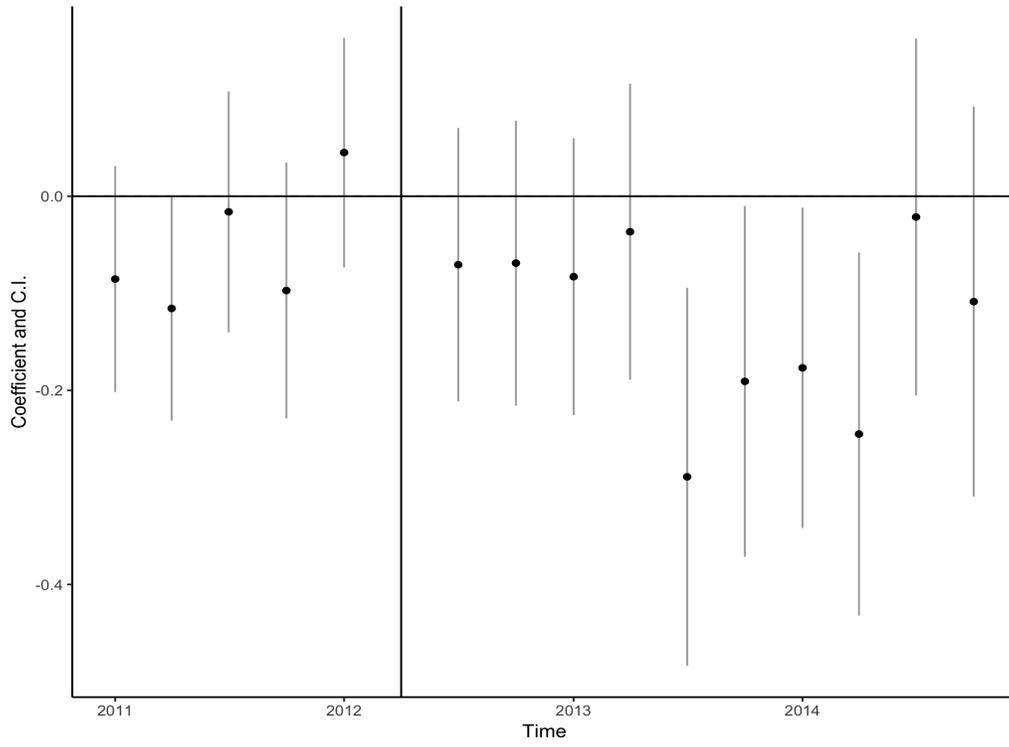
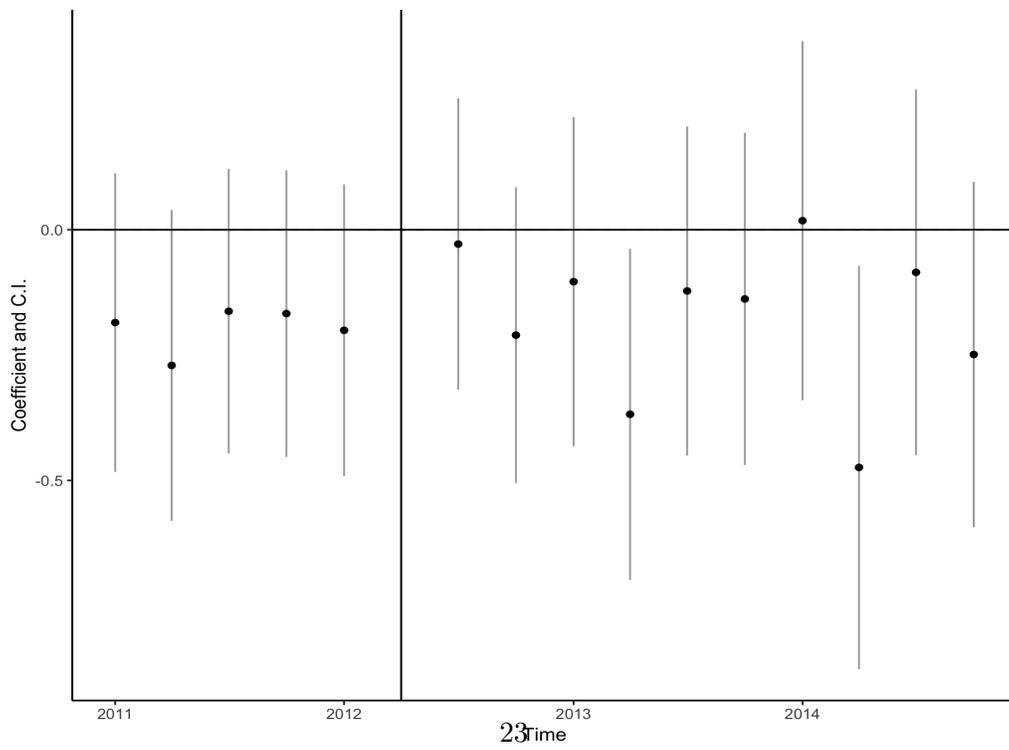


Figure 3: Yearly share of dismissed workers, by individual AKM effect quintiles

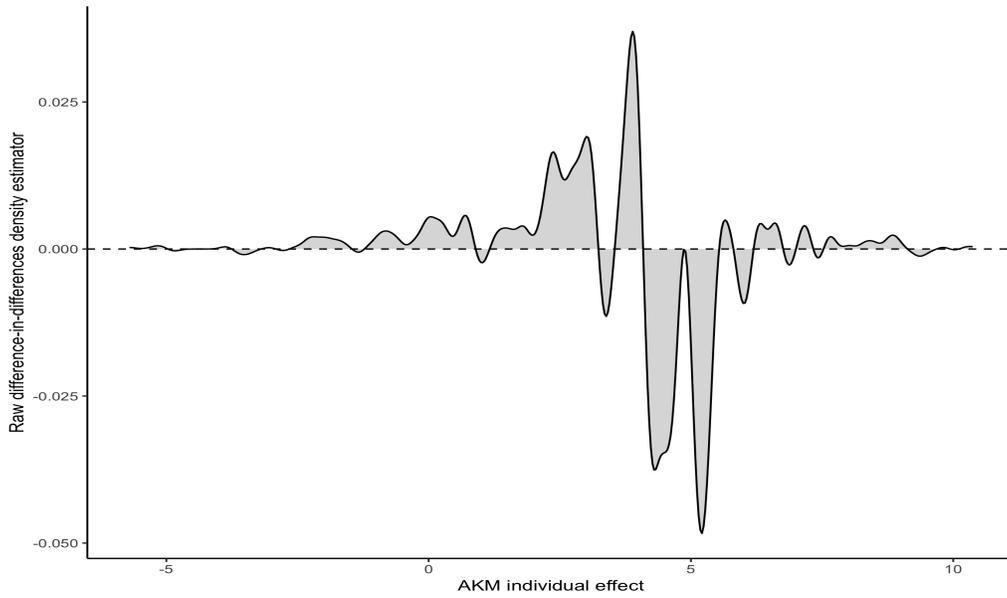


(a) New hires

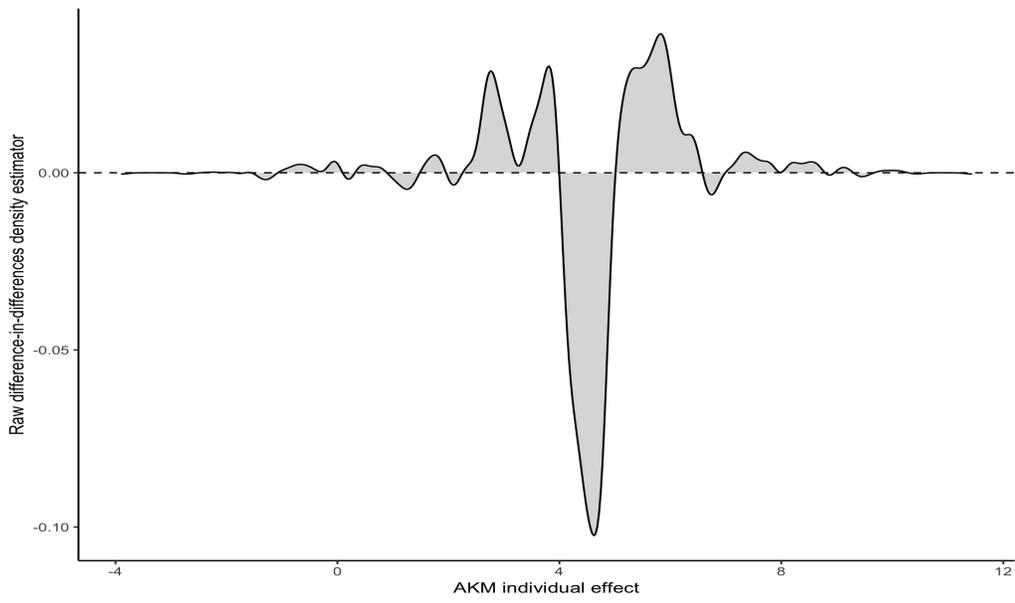


(b) Dismissals

Figure 4: Pre-trends analysis



(a) New hires



(b) Dismissals

Figure 6: Density difference-in-differences raw estimator

Table 3: OLS Difference-in-differences - New hires

<i>Dependent variable:</i>					
AKM individual effect					
	(1)	(2)	(3)	(4)	(5)
Diff-in-diff. Fornero	-0.066** (0.028)	-0.057** (0.027)	-0.052** (0.026)	-0.052** (0.026)	-0.062 (0.074)
Observations	43,381	43,381	43,381	43,381	43,381
Adjusted R ²	0.001	0.095	0.109	0.109	0.633

Note:

*p<0.1; **p<0.05; ***p<0.01

This table displays the difference-in-differences coefficients from an OLS regression in which the dependent variable is the AKM individual effect associated with a given contract. Only new hires are included in the analysis. Column 1 displays the raw data estimates. Column 2 includes controls for work experience, age, sex and region of residence. Column 3 also includes the economic sector, occupation, part-time share and job tenure. Column 4 accounts for time dummies as well. Finally, Column 5 further adds firm fixed effects. All standard errors are clustered by individual and firm.

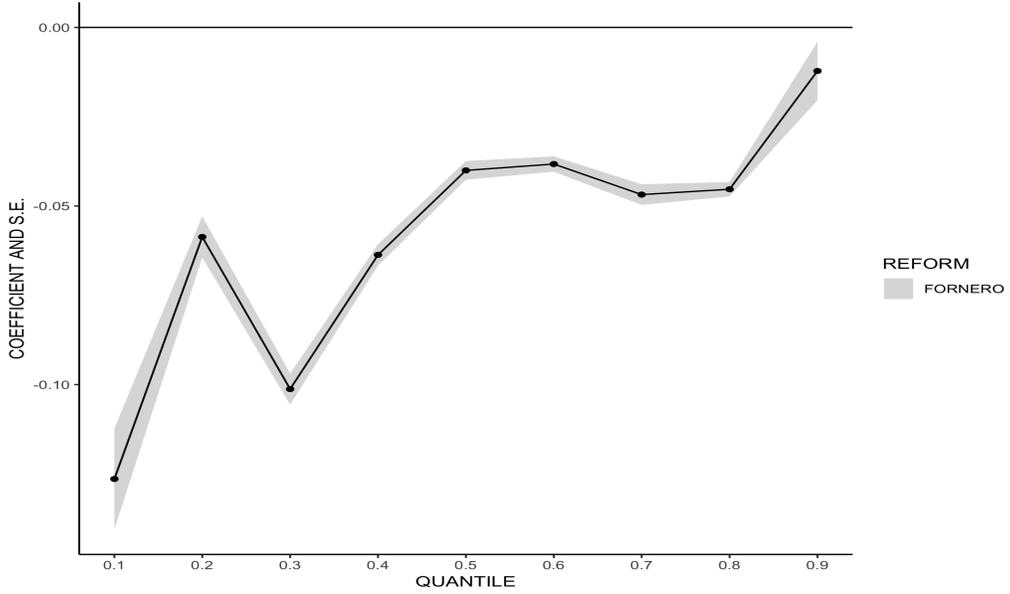
Table 4: OLS Difference-in-differences - Dismissals

<i>Dependent variable:</i>					
AKM individual effect					
	(1)	(2)	(3)	(4)	(5)
Diff-in-diff. Fornero	0.003 (0.035)	-0.004 (0.033)	0.004 (0.033)	0.004 (0.033)	-0.072 (0.088)
Observations	20,849	20,849	20,849	20,849	20,849
Adjusted R ²	0.005	0.093	0.114	0.114	0.780

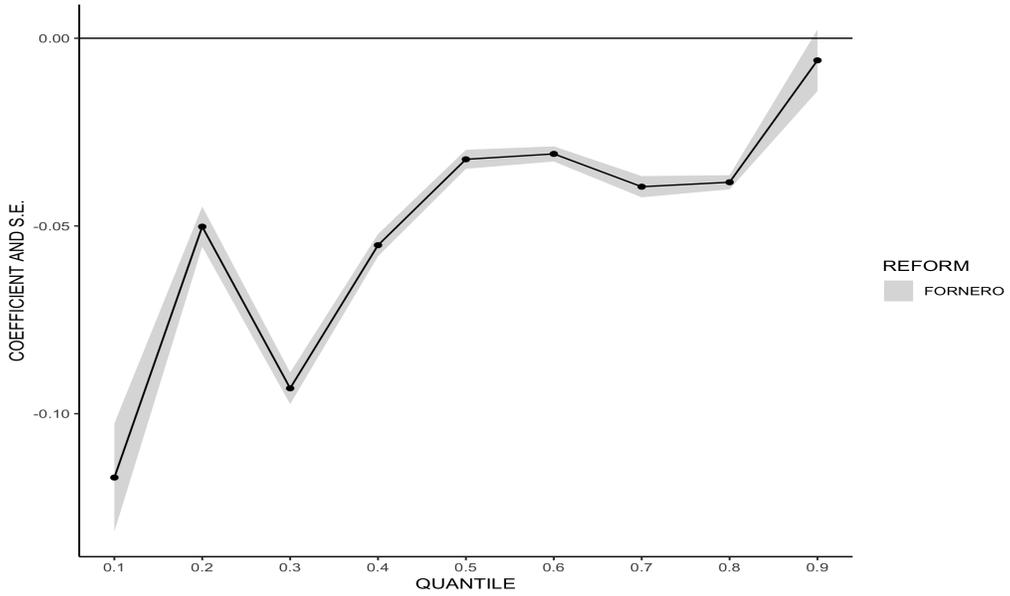
Note:

*p<0.1; **p<0.05; ***p<0.01

This table displays the difference-in-differences coefficients from an OLS regression in which the dependent variable is the AKM individual effect associated with a given contract. Only dismissals are included in the analysis. Column 1 displays the raw data estimates. Column 2 includes controls for work experience, age, sex and region of residence. Column 3 also includes the economic sector, occupation, part-time share and job tenure. Column 4 accounts for time dummies as well. Finally, Column 5 further adds firm fixed effects. All standard errors are clustered by individual and firm.

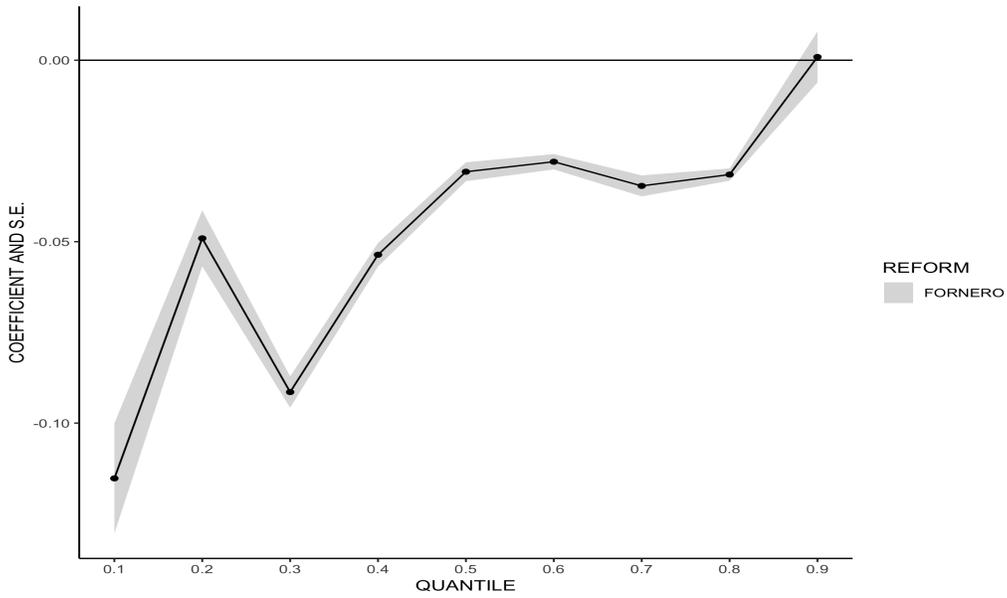


(a) Raw data

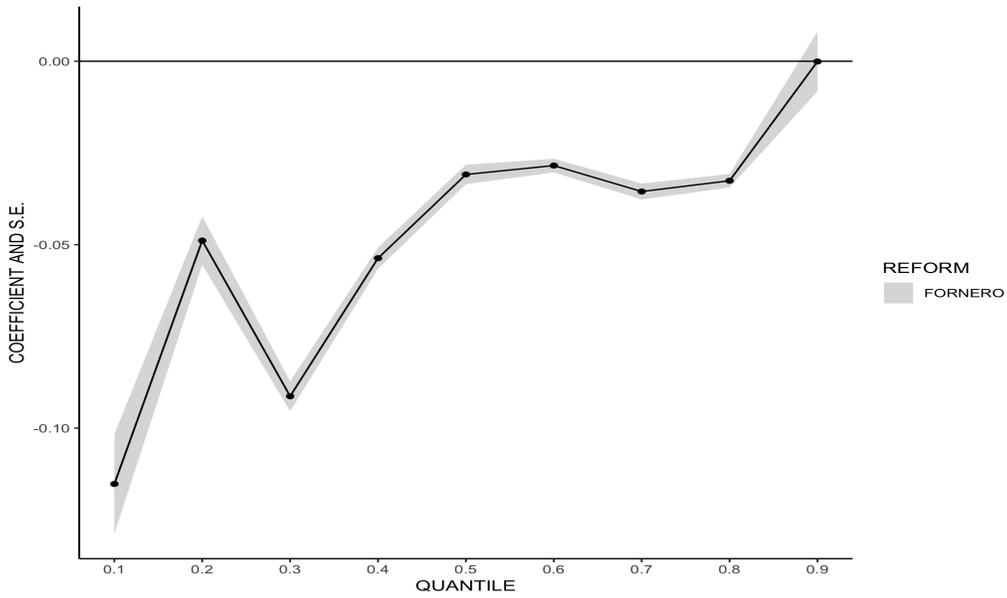


(b) Worker controls

Figure 8: UQR coefficients and S.E. - New hires A

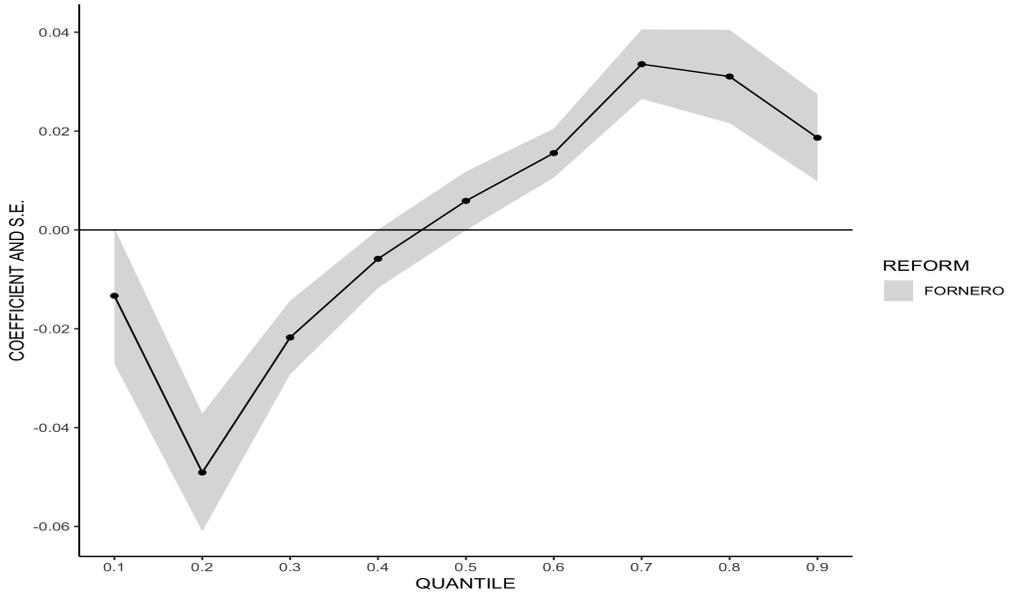


(a) Worker and firm controls

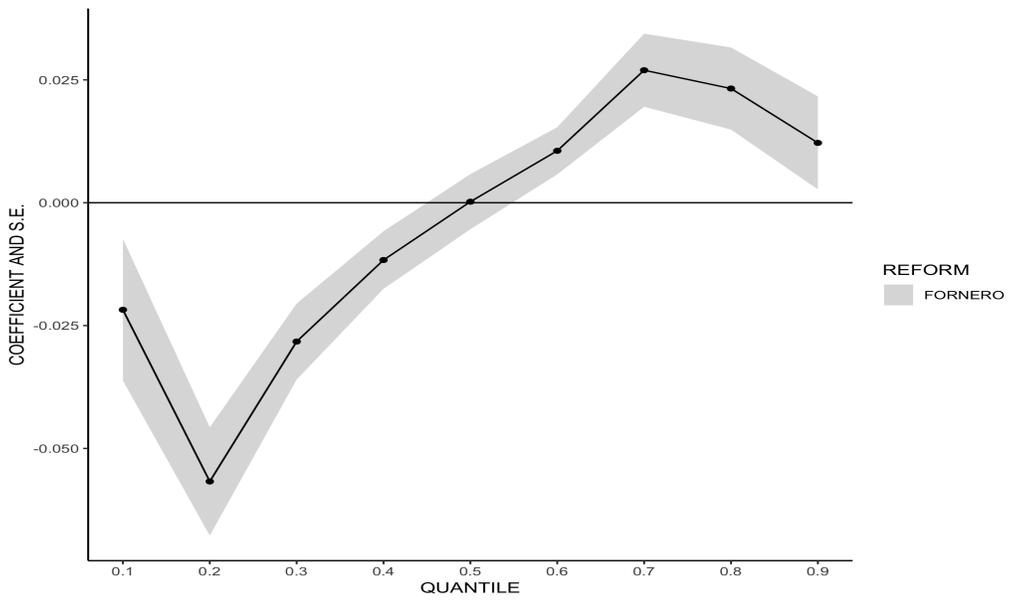


(b) Worker and firm controls, time dummies

Figure 10: UQR coefficients and S.E. - New hires B

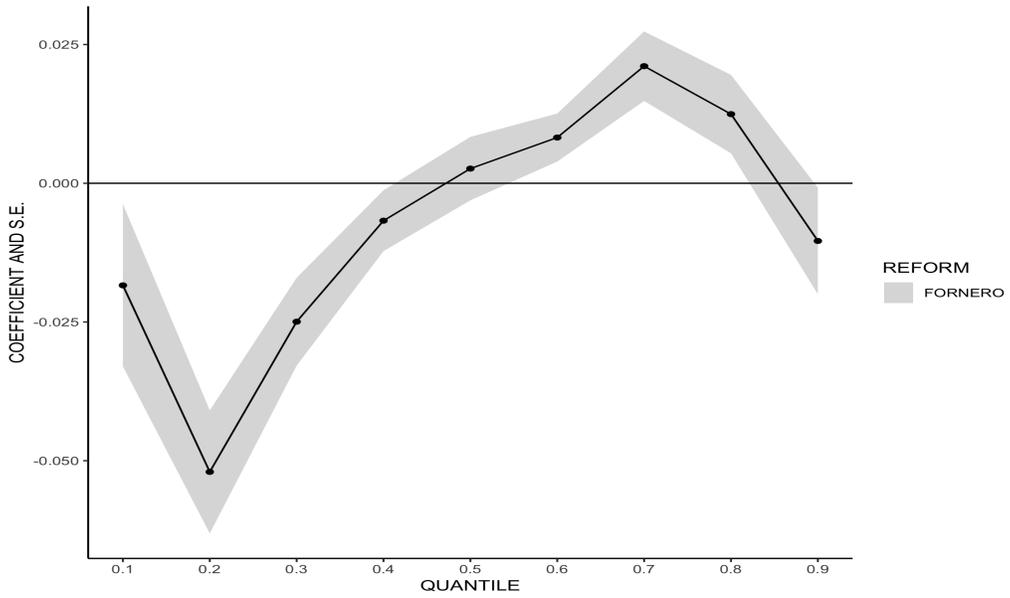


(a) Raw data

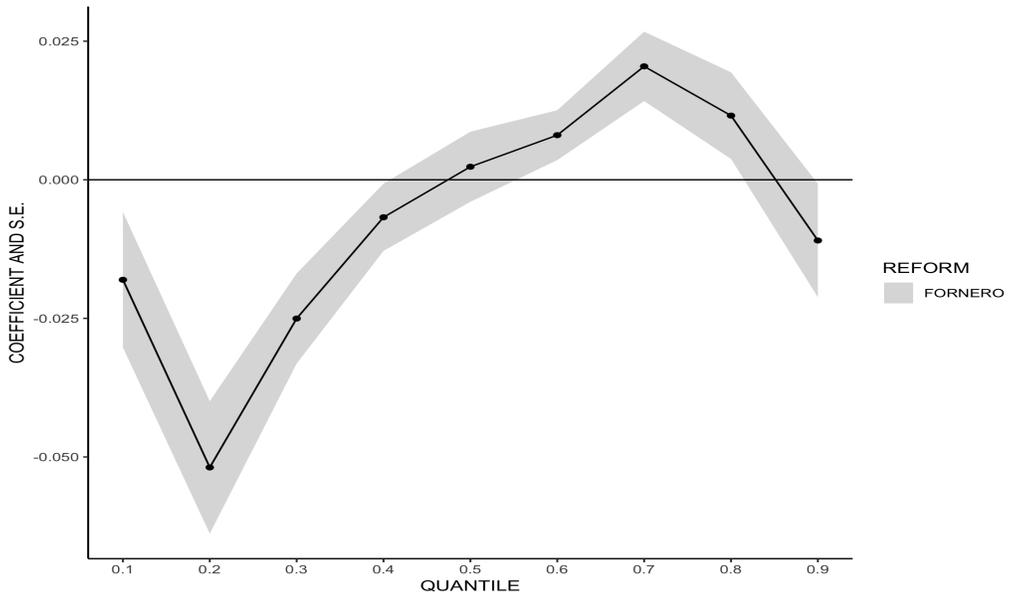


(b) Worker controls

Figure 12: UQR coefficients and S.E. - Dismissals A



(a) Worker and firm controls



(b) Worker and firm controls, time fixed dummies

Figure 14: UQR coefficients and S.E. - Dismissals B