

Firms' margins of adjustment to wage growth: the case of Italian collective bargaining

Francesco Devicienti¹ | Bernardo Fanfani²

¹Università di Torino, Collegio Carlo Alberto and IZA

²Università di Torino and IZA

Correspondence

Francesco Devicienti, Dipartimento di Scienze Economico-Sociali e, Matematico-Statistiche, corso Unione Sovietica 218b, I-10134 Torino, Italy.

Email: francesco.devicienti@unito.it

Abstract

This study analyses firms' adjustment behaviour when facing higher labour costs. The empirical research design considers several outcomes, and exploits, as a source of variation in labour costs, discontinuities in the growth of contractual wages set by Italian collective bargaining institutions. The results indicate that adjustment channels are highly heterogeneous across the firms' productivity distribution. Employment, revenue, productivity and the profit margin are negatively related to contractual wage growth among relatively less efficient companies. Instead, most efficient firms do not downsize, they substitute high- with low-wage workers while preserving their productivity, and they may even increase (or at least keep constant) their profitability. We conclude that more efficient companies, which adjust through cost-saving and labour-hoarding strategies, may benefit from cleansing effects, as their product market shares increase when costs of more constrained rivals are raised.

KEYWORDS

collective bargaining, employment, matched employer–employee data, minimum wage, productivity

JEL CLASSIFICATION

J00; J23; J24; J31; J38; J58; L13

1 | INTRODUCTION

The problem of predicting firms' adjustment behaviour to labour scarcity has long been debated in economics (see, for example, Acemoglu 2002). Labour scarcity is relevant to firms since it typically takes the form of an increase in labour costs or of a reduction in the relative price of other production factors. It is a quite recurrent phenomenon that can be triggered by macroeconomic fluctuations, migration flows and technological shocks, including the availability of

This is an open access article under the terms of the [Creative Commons Attribution](https://creativecommons.org/licenses/by/4.0/) License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2024 The Author(s). *Economica* published by John Wiley & Sons Ltd on behalf of London School of Economics and Political Science.

cheaper investment goods, but also by institutional factors, such as a growth in labour income taxes or tighter wage regulations.

This paper studies firms' adjustment paths to the growth in relative labour costs induced by an institutional mechanism, namely by changes in the level of contractual minimum wages set within the system of industrial relations (contractual wages, for short). Our analysis evaluates the relative importance of a rich set of adjustment channels potentially used by firms when facing this shock, in particular productivity, revenue, employment, profits, capital intensity, wages and workforce composition.

The contractual wages considered in our analysis operate similarly to a strictly binding statutory minimum wage. Rather than being regulated by the government, these pay levels are collectively bargained by trade unions and employers' associations. This wage-setting mechanism can be found in various countries and tends to be quite common in continental Europe (see OECD 2017). Italy is an illustrative case for such an institutional setup, with pay negotiations that are carried out at the nationwide industry level.¹ Several features make the Italian case particularly interesting to study, and motivate the detailed analysis that we undertake in this paper.

The main purpose of Italian collective contracts is to set wage floors at the level of the national sector. Other dispositions contained in these contracts are typically highly stable across time, and they are often amendable by individual firms. For this reason, the policy context considered in this analysis is reminiscent of stylized models where unions set wage levels as a monopolist through centralized bargaining, while individual firms choose employment levels on their labour-demand schedule (see, for example, MaCurdy and Pencavel 1986). Another peculiarity of Italian contractual wages is that they are specific for each occupation, including relatively high-paying ones. Moreover, they represent both a pay floor and a fixed-pay component. Hence their growth tends to affect the pay levels of virtually all workers covered by the collective agreement, not only those close to the pay floor or at the bottom of the wage distribution.² Given this setting, economic theory provides several predictions on the adjustment channels available to firms when contractual wages grow.

According to the standard Hicks–Marshall theory of labour demand, firms adjust to positive labour cost shocks by capital–labour substitution and price increases, which lead to a reduction in output and employment levels (Hamermesh 1993). Other non-competitive adjustment channels could also be relevant, such as profit reductions when rents are available to firms in the product or labour market, lower managerial slack, or increased worker effort. In these latter cases, output and employment effects could be more ambiguous.³

In this paper, we argue that there are relevant market-level interactions among the alternative adjustment mechanisms available to different types of firms. Specifically, firms that have rents to exploit when facing the wage shock also benefit from reduced competition from less efficient employers, which are typically subject to the negative employment and output effects predicted by the competitive model. In this setting, high-rent firms can even increase their profits when labour costs grow. That is, contractual wage growth can be characterized as an institutional shock that *raises the rivals' costs*. This form of profitable exclusionary strategy has often been considered relevant in the industrial organization literature (see, for example, Salop and Scheffman 1983). Various forms of this argument have also been proposed in the theoretical literature on centralized wage-setting. Indeed, by raising costs in a context where firms respond asymmetrically due to heterogeneity in their resources and productivity, this institutional mechanism could give rise to relevant cleansing effects among firms sharing the same industry or local labour market (e.g. Williamson 1968; Moene and Wallerstein 1997; Haucap *et al.* 2001; Barth *et al.* 2014). Despite this interest, empirical investigations on the salience of this phenomenon have been relatively scarce.

In our empirical analysis, we have adopted a fixed effects estimation strategy that exploits changes in the level of contractual wages in order to analyse employers' adjustment behaviour.

In this model, the parameters of interest are identified by comparing the growth in the outcomes of interest between firms affected and unaffected by changes in contractual wage levels, conditioning on a rich set of non-parametric time effects specific for each sector in each geographic location. Once we remove the variation captured by the fixed effects, we are effectively comparing firms within the same sector, province and year, which experience different changes in bargained wage solely because they are applying different contracts due to their historical labour force composition, to the specificity of their activity, or to their broad size at the time when the bargaining was taking place.⁴

Our study is based on the most comprehensive panel of incorporated businesses' balance sheets available for Italy, as provided by Cerved. We have matched this database to social security records on the population of private-sector employees provided by the Italian Social Security Institute (INPS) and to a comprehensive hand-collected dataset on contractual wages set by the majority of Italian collective bargaining agreements. The final sample of analysis comprises almost 400,000 firms per year over the period 2006–15, virtually covering the universe of Italian incorporated businesses in the private sector.

Our results show that, on average, higher labour costs induce a decline in employment and revenue, an increase in wages, and null effects on productivity, workers' average quality, the profit margin and capital intensity. However, we show that these results are highly heterogeneous across firms, and that there are relevant interactions across adjustment channels exploited by different employers that take place at the market level.

The main dimension that we have considered to characterize the heterogeneity in adjustment behaviour across firms is their underlying productivity. This is an interesting dimension for several reasons. Productivity is closely related to quasi-rents available to firms in the product and labour markets (e.g. Card *et al.* 2018). The empirical literature on centralized wage-setting has often considered productivity as a good proxy to characterize more and less advantaged firms in similar systems of industrial relations (Boeri *et al.* 2021; Bartolucci *et al.* 2018; Manacorda and Petrongolo 2006). Theoretical contributions also emphasize productivity differences across firms when characterizing heterogeneity in the effects of collective bargaining (e.g. Barth *et al.* 2014).⁵

We classified firms in four groups reflecting their relative productivity within collective contracts. We find that higher contractual wages have a relatively similar positive effect on the pay level of incumbent workers for all types of employers. In the group of relatively less efficient firms, contractual wage growth has strongly negative effects on employment, productivity and revenue. The effects of this shock on the profit margin are also negative for less efficient employers. By contrast, the productivity effects associated with contractual wage growth are either not significant or marginally positive in the group of most productive firms. For these firms, contractual wage growth has a similar positive effect on revenue, profitability and employment, although these effects are not always significant, depending on the specification adopted. Instead, capital intensity is not significantly affected by the wage shock at both more and less productive firms, indicating that changes in the adoption of technologies do not play a major role, at least with regard to relatively short-run adjustments.

This rich analysis of the outcomes of several firms pinpoints important mechanisms arising from market-level interactions in the adjustment channels across different types of firms.

First, the labour cost shock induces a redistribution in market shares that benefits the most efficient companies. Sales decline more the lower the firm's productivity, which implies that high value-added firms increase their share of production at the expense of the least productive firms as a result of higher contractual wages.

Second, the least productive firms suffer strong employment losses, particularly among low-wage and low-firing-cost workers. Since productivity decreases for these firms, despite more positive selection in their workforce, this suggests that frictions and firing costs (which are typically higher for more qualified workers with open-ended contracts) may not allow these employers

to adjust employment levels optimally. Through similar mechanisms, the least productive firms could also be pushed away from an optimal scale level. Instead, efficient competitors increase the use of cheaper workers while keeping employment levels constant after the labour cost shock.

After the growth in labour costs, the average age of employees, the share of open-ended contracts and the wage ‘quality’ of workers—as measured by their earning potential through Abowd–Kramarz–Margolis (AKM) workers’ fixed effects—decrease at more efficient companies and increase at less efficient firms. An increased reliance on cost-saving strategies and workers’ reallocation may rationalize the employment composition effects among the most efficient companies. In particular, low-wage workers may become more easily available in the labour market after the labour cost shock, as they are laid off by the least efficient firms. This induces more productive firms to use more extensively low-wage workers, a cost-saving strategy that could become more appealing to these employers, which are precisely those that do not cut production levels after the growth in labour costs.

Given the above mechanisms, the highest quartile of relatively more efficient firms within a collective contract is actually able to increase profitability (or keep profitability constant) after the labour cost shock, while the profit margin tends to decline elsewhere. That is, the increase in the product market share of relatively more efficient companies and the related reallocation of the workforce allow them to compensate for the growth in labour costs derived from centralized collective bargaining. Overall, our analysis shows that the effects of higher labour costs on firms’ behaviour are quite complex and heterogeneous, depending on the relative efficiency of firms that are hit by this shock. Moreover, interactions across adjustment channels exploited by different firms can even lead to profit gains among the most efficient, high-rent employers, as costs of more constrained rivals, which are more subject to competitive pressures, are raised by centralized wage growth (Williamson 1968; Salop and Scheffman 1983).

Our results contribute to the literature on firms’ margin of adjustments in the presence of higher labour costs, on which the evidence is still not abundant. The shock to labour costs from contractual wages that we are studying is expected to be stronger than that from a standard minimum wage change, as contractual wages are typically binding across the entire pay distribution, and not just at the bottom of it. This setting offers a fertile ground to better understand the elusive impacts of the minimum wage (e.g. Manning 2021),⁶ as well as the relevance of various adjustment margins and their interactions in the face of large labour cost shocks.

Finally, this study nicely complements the evidence in Dustmann *et al.* (2022) on the reallocation effects of minimum wages, by investigating the effects of a similar, yet distinct wage policy—collectively bargained floors.⁷ On this topic, we show that a hitherto neglected dimension of heterogeneity in firms’ adjustment behaviour to a market-wide growth in relative input costs is given by product market shares, which tend to increase for firms that are more resourceful and able to cope with this shock. This mechanism could also represent an explanatory element for the secular growth in product market power and concentration, whose sources are highly debated in the recent literature (De Loecker *et al.* 2020).

2 | INSTITUTIONAL CONTEXT

According to the Italian Constitution (article 36), each employee is entitled to a pay level that is commensurate with the tasks that he performs, and sufficient to guarantee an adequate standard of living. Italian labour courts have interpreted this provision as a disposition to apply to each worker the minimum contractual wage that is bargained by the most representative collective agreement relevant to the worker’s occupation. Thus contractual wages set within the Italian system of industrial relations *de facto* represent statutory pay floors that apply to virtually all private-sector employees.

The Italian employers' association and trade unions negotiate contractual wages at a highly centralized level. There are several hundred collective contracts, but the 150 largest contracts cover most private-sector employees, as they apply to more than 15 million workers, representing more than 90% of the workforce. The 2017 classification of the INPS included around 300 collective agreements. There are also several other contracts (typically those with very small coverage and often a dubious legal basis for their applicability) that are not included in this classification, but the proportion of workers falling into this group of unregistered agreements was always below 2% during the years covered by our study.⁸

Italian legislation considers contractual wages to be not just a wage floor, below which an employee in the relevant occupation and sector cannot be paid. They are also, in practice, a fixed component of the wage. This implies that whenever a contractual wage grows by a given amount, all pay levels in the relevant occupation must be increased by the same fixed amount, including those already above the new minimum level. There are clauses called *superminimi assorbibili* according to which employees who are paid above the minimum can agree to give up this fixed pay rise, as long as their wage remains above the relevant contractual wage. Although there is no systematic evidence on the incidence of these clauses, they tend not to be very common. Indirect evidence of this phenomenon is provided by Adamopoulou and Villanueva (2022). Their study shows that Italian wages in the metal-manufacturing sector tend to increase across the entire within-contract earning distribution in response to the growth of negotiated pay levels, while in recent years, the 'wage cushion' (i.e. the difference between actual and contractual pay levels) has always been quite stable across time. Importantly, the same study also documents negligible levels of non-compliance with contractual wage growth.

Many studies have stressed the influence of collective bargaining on wage differentials and inequality, including Belloc *et al.* (2023) and Boeri *et al.* (2021) regarding geographic wage dispersion, Devicienti *et al.* (2008) and Faia and Pezone (2024) regarding wage rigidity, and Erickson and Ichino (1995), Manacorda (2004) and Devicienti *et al.* (2019) regarding wage inequality and its evolution.⁹ Discontinuities in the timing of collective contracts have also been exploited as a source of variation in other policy evaluation exercises (e.g. Cappellari *et al.* (2012) and Daruich *et al.* (2023) on temporary employment). In a complementary study, Fanfani (2023) analyses employment losses and wage effects associated with contractual wage growth. While that study focuses on aggregate employment and wage effects across demographic and industry groups, the present analysis considers instead a large set of firm-level margins of adjustment to contractual wage growth.

3 | DATA

Our empirical analysis is based on three main sources of information. First, we rely on the Cerved database on Italian incorporated businesses' balance sheets and profit and loss accounts for the years 2006–15. These data cover virtually all Italian incorporated companies, and we are able to match each of these firms to the population of its employees registered in the INPS social security records archives, our second source.

The INPS records are based on compulsory information compiled by all employers in the private sector that hire at least one employee, thus they cover the universe of workers to which the dispositions of collective bargaining apply. Finally, our third source of information is a hand-collected database on Italian contractual wages settled in around 160 nationwide agreements renewed periodically between 2006 and 2015. Relying on information on the collective agreement applied to each worker reported in the INPS archives, we have been able to match almost 80% of the private-sector employee population to a contractual wage.¹⁰

3.1 | Matching firms to contractual wage data

To better understand how contractual wages work in Italy, Figure 1 plots the evolution of these wage floors over the period 2006–15 for the two largest collective agreements, for the metal-manufacturing and trade sectors. Each contract sets more than one pay floor for different job titles. Which pay level applies depends on the occupation and sometimes on the seniority levels, but the INPS data do not provide information on the specific job title of each worker within collective contracts. Contractual wages are renewed at different dates, with changes that appear to be more frequent in the metal-manufacturing contract (each dot on the graph represents a contractual wage change). As mentioned, contractual wages represent both a wage floor and a fixed component of the pay, so that their growth typically implies that all wages in the relevant job title have to be adjusted.

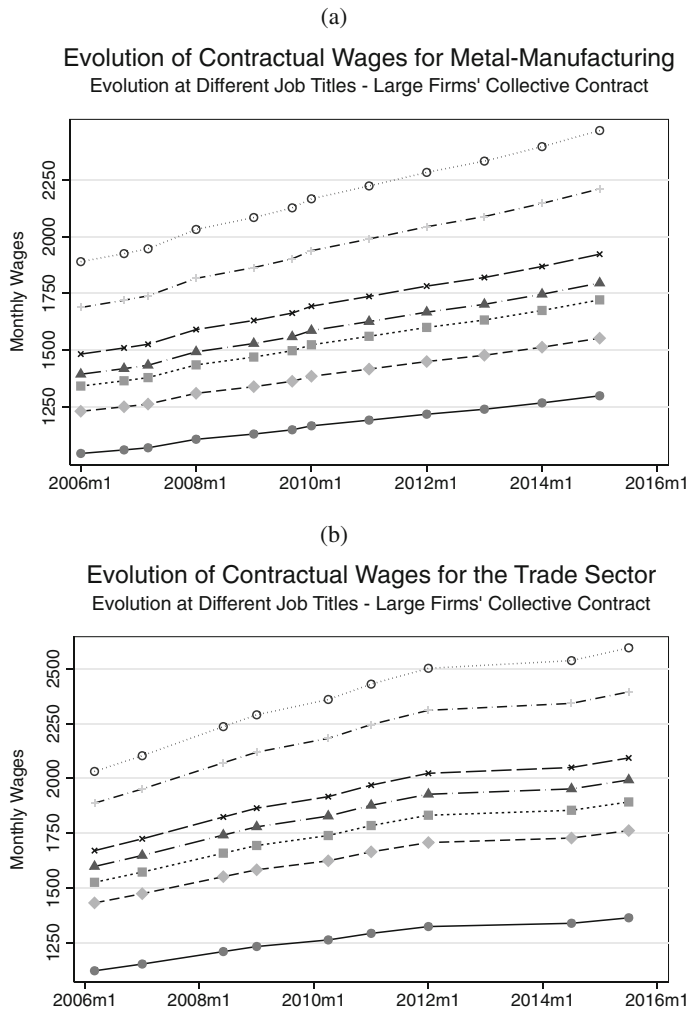


FIGURE 1 Evolution of contractual wages in selected collective agreements. *Notes:* Evolution of contractual wages set by (a) the metal-manufacturing collective contract and (b) the trade collective contract. Each line describes the evolution across time of a job title's pay floor within the same collective contract. Each dot identifies a new level of a job title's pay floor. The slope of the lines is steeper for larger and more frequent contractual wage increases.

Since our unit of analysis is the firm, we have identified for each company the most expensive collective contract, that is, the collective contract covering the largest proportion of the firm's wage bill.¹¹ Using this information, we have assigned each firm to a treatment, defined as the median contractual wage of the most expensive agreement.

The minimum wage assigned to each firm approximates quite well the dynamics of contractual wages observed among all job titles within the same agreement. The correlation coefficient of the contractual wage growth across pay levels within the same contract at renewal dates was around 0.74 during the years covered in our sample. This feature is also visible in Figure 1, where the relative distance across different contractual pay levels is always quite stable for different job titles within the same collective contract. Considering also that the percentage of workers covered by the largest collective contract of the firm is 93% on average, a growth in our treatment variable can be considered a good approximation to a shock in labour costs affecting most workers within the firm, even if the precise magnitude of this shock is potentially measured with error.

3.2 | Descriptive statistics on the sample of analysis

Our final sample of analysis includes only Cerved firms with at least one employee in the INPS archives, and whose most expensive collective agreement was present in our database on contractual wages. The number of firms included in this sample was approximately 600,000, which were observed for a period of nearly 6 years on average.

Figure 2 provides descriptive statistics on the representativeness of the Cerved–INPS–contractual wage matched sample across years, with respect to the universe of private-sector firms with at least one employee, which is observable through the INPS social security archives. The same statistics are also computed among firms with at least 10 employees. Our sample of analysis includes slightly more than 20% of all Italian firms with at least one employee, and this coverage rate grows to around 65% when considering the population of firms with at least 10 employees.¹² The over-sampling of larger firms is due to the fact that the Cerved data include only incorporated businesses, which are mandated to maintain balance sheets and make them publicly available via the Chambers of Commerce. The Cerved data, however, do not include unincorporated enterprises, which are typically small family-run businesses or other partnerships that are not subject to the above dispositions concerning balance sheets. Importantly, Figure 2 shows that

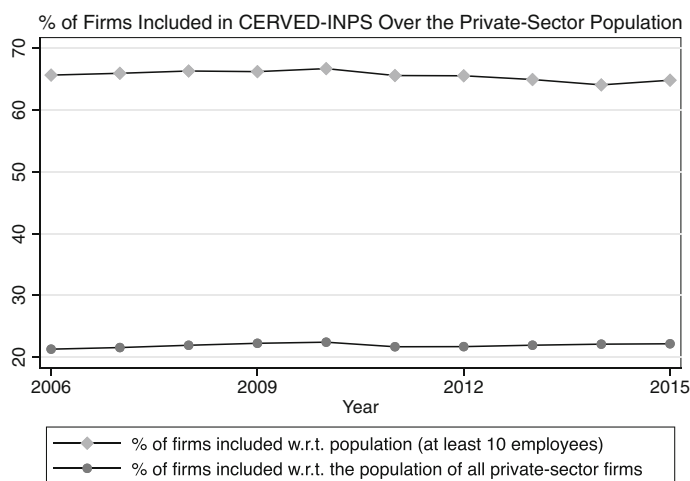


FIGURE 2 Representativeness of the Cerved–INPS matched sample.

TABLE 1 Summary statistics weighted by firms' size (2006–15).

Variables	Mean	S.D.	Available observations
Log median contractual wage	4.099	0.109	3,515,332
Nominal contractual wage growth	0.024	0.014	2,650,312
Main collective contract share in wage bill	0.928	0.138	3,515,299
Firms' closure	1.5%		3,515,332
Log FTE employees	4.367	2.358	3,515,332
Log firms' average wages	4.358	0.315	3,515,277
Incumbents' wage growth	0.027	0.087	2,693,455
Firms' average AKM worker fixed effects	0.000	0.188	3,515,332
Log value-added per worker	3.922	0.575	3,240,727
Log TFP	4.555	0.808	3,150,643
Log revenues	8.239	1.633	3,419,936
Log profit margin	-2.678	0.871	2,736,402
Log physical capital per worker	2.887	1.750	3,267,867
Total number of firms	603,855		
Total number of observations	3,515,332		

Notes: All means and standard deviations are computed weighting by the number of workers observed in the firm each year. Contractual wages refer to the nominal median pay level of the collective contract that covers the greatest proportion of the wage bill. Firms' closure is defined as a permanent exit from the INPS social security archives in the subsequent year. Balance sheet variables are derived from Cerved and are not always available for all firms in every year. AKM worker fixed effects were computed using the Abowd *et al.* (1999) regression model and standardized as the difference from their mean value. TFP was derived from the Levinsohn and Petrin (2003) regression model.

the coverage rate with respect to the underlying population was quite stable across time, which suggests that firms' selection into the sample is relatively homogeneous in all years.

Table 1 provides descriptive statistics for our sample of analysis computed after weighting for the number of workers in each firm. Overall, we have relied on an unbalanced panel of more than 600,000 firms and around 3.5 million firm–year observations. The average size of firms, measured in full-time equivalent (FTE) number of workers is around 77 employees.¹³ Median contractual daily wages are on average 30% lower and three times less dispersed than firms' actual daily average wages. The annual percentage growth of contractual wages is 2.4%, a value that is slightly lower than the average growth of incumbent workers' actual pay levels within firms (2.7%).

For each firm, we have also computed the average of its workers' fixed effects derived from an AKM regression model (Abowd *et al.* 1999), which we have expressed as a difference from their mean value in the sample.¹⁴ Appendix Section A.3 provides more details on the AKM estimation procedure and its results. AKM worker fixed effects allow us to rank employees' time-constant relative earning abilities, conditional on employers' fixed effects and on observable time-varying characteristics. Table 1 shows that the standard deviation of this measure of average workers' quality is 0.188. This dispersion accounts for around 60% of the total dispersion in firms' average wages, whose standard deviation amounts to 0.315.

For each firm, we have estimated a measure of total factor productivity (TFP) using the Levinsohn and Petrin (2003) method and adopting the value-added-based regression approach. This method is based on the use of lagged intermediate goods to instrument for the choice of capital and labour levels, and allows the recovery of a measure of a firm's efficiency conditional on the amount of production factors employed. In general, firms' financial information is not always available due to missing variable problems arising in the Cerved database. The variable that is

most affected by this problem is the profit margin, which is defined as earnings before taxes, interest, and depreciation over revenue.¹⁵

The percentage of firms that change their main collective contract (potentially starting to use a contract not included in our contractual wage sample) is only 3.4%. Appendix Section A.4 provides a regression analysis on whether a firm's selection out of collective contracts is related to contractual wage growth, thereby assessing the relevance of this potentially endogenous sample selection mechanism. The percentage of firms that permanently disappear from the INPS archives in the subsequent year, typically because they go out of business or stop hiring any employees, is only 1.5%. Section A.4 also assesses to what extent such firm exits are related to the wage growth stipulated by collective bargaining.

Appendix Table A1 provides additional statistics on yearly growth rates of the outcome and treatment variables within firms. Since we have adopted a firm fixed effect estimation strategy, this within-firm variation is a close approximation of the variation actually used to identify the treatment effects of interests. As can be noted, the yearly average contractual wage growth was 2.4% on average, with a standard deviation of only 1.4%. The outcome variables instead have a less persistent evolution within firms, as shown by higher standard deviations in their growth rates.

3.3 | Variation in contractual wages

To better characterize the source of variation in contractual wages that is exploited in the empirical analysis, Table 2 provides more detailed descriptives on contractual wages and their yearly growth. As the table shows, the typical growth of contractual wages ranges between 0.8% and 4.1%. Our empirical analysis exploits variations in the timing and size of these shocks across collective agreements in order to identify firms' adjustment behaviour.

TABLE 2 Detailed statistics on contractual wages and their dynamics.

Variables			
<i>Panel A: Log median contractual wage</i>			
	Mean	S.D.	Observations
	4.099	0.109	3,515,332
	10th perc.	50th perc.	90th perc.
	4.007	4.090	4.192
Variance decomposition			
Within years, sectors, regions	55.7%		
Between years, sectors, regions	44.3%		
<i>Panel B: Nominal contractual wage growth</i>			
	Mean	S.D.	Observations
	0.024	0.014	2,650,312
	10th perc.	50th perc.	90th perc.
	0.008	0.024	0.041
Variance decomposition			
Within years, sectors, regions	41.9%		
Between years, sectors, regions	58.1%		
Total number of firms	603,855		
Total number of observations	3,515,332		

Notes: All means and standard deviations are computed weighting by the number of workers observed in the firm each year. Contractual wages refer to the nominal median pay level of the collective contract that covers the greatest proportion of the wage bill. The variance decomposition of each variable is computed within and across groups defined by the interaction of 38 sectors, 107 provinces and 10 years.

Considering that our identification strategy is based on a rich set of non-parametric fixed effects that restrict the identifying variation within sectors and small geographical areas, it is interesting to discuss in more detail the source of variation in contractual wages within sectors and regions, and its magnitude. For each panel of Table 2, the last two rows report the decomposition of the variance into components determined by variation within and between groups. These groups are defined by the interaction of years, a 38-sector ISIC classification, and 107 Italian provinces. The main empirical specification of the paper includes a full set of fixed effects for each of these groups, thus it restricts the variation of contractual wages used for identification purposes within these cells.

Table 2 shows that as much as 55.7% of the variance in contractual wages is driven by differences within the same year, sector and region. This implies that there is a substantial heterogeneity in the collective contracts applied by firms within these cells. Moreover, as much as 41.9% of the variance in the yearly growth of contractual wages is determined by differences among firms within the same year, region and sector. Thus the variance decomposition exercise shows that even within fairly granular industries and geographic locations, firms tend to apply different collective contracts characterized by fairly heterogeneous growth rates of contractual wages.

Firms within the same sector and region may be covered by different contracts for historical and organizational reasons. For example, in most manufacturing sectors, there are different collective contracts for large and mid-sized firms, which are negotiated separately. Moreover, some collective agreements cover heterogeneous activities that may belong to a variety of industries according to a standard sector classification.

To provide a more concrete example of the variation in collective contracts within sectors, Figure 3 shows the evolution of contractual wages bargained by the largest collective agreements in the sector of leather and related products manufacturing, which represent a unique two-digit industry group. The figure shows that this sector is populated by a variety of collective contracts of similar size. Indeed, a separate collective contract exists for workers in firms involved in tanning and leather production, with a coverage of around 22%, in firms involved in production of footwear and leather products (27%), and in firms involved in the production of wearing apparel (18%). Altogether, these three contracts cover around two-thirds of all firms observed in this two-digit industry.¹⁶

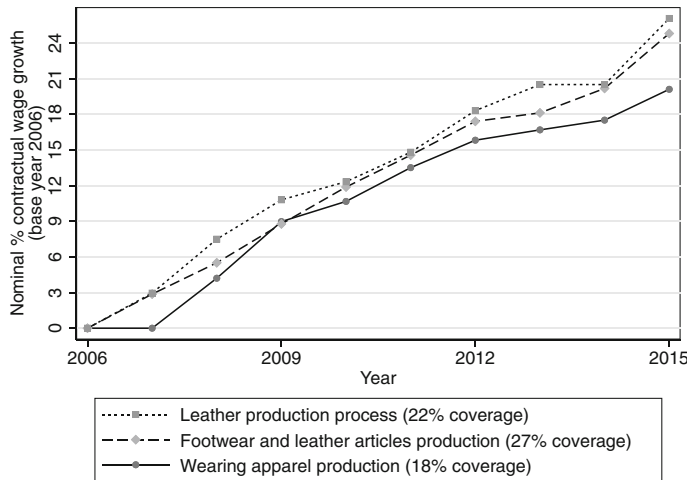


FIGURE 3 Contractual wage growth in the largest collective contracts of the leather and related products manufacturing industry. *Notes:* Evolution of the median contractual wage in the three largest collective agreements observed in the leather and related product manufacturing sector. Each line describes the cumulative growth of a given median contractual pay floor. Each dot identifies the percentage difference of these pay floors in each year with respect to their 2006 level. The slope of the lines is steeper for larger contractual wage increases across years.

Figure 3 shows that the growth in nominal contractual wages bargained by these three collective agreements was heterogeneous. Depending on the agreement, contractual wages increased by between 20% and 25% in nominal terms across the entire period of analysis. The pace of this growth was quite heterogeneous across collective contracts in the short run. For example, the growth in contractual wages in the wearing apparel contract was smaller by 3 percentage points in 2007, but it was larger with respect to the other contracts in 2008 and 2009. Overall, Figure 3 is illustrative of the relatively fragmented structure of collective contracts in Italy, which gives rise to relevant short-run differences in the growth of bargained minimum wages even within the same industry. This type of variation has allowed us to identify the treatment effect of interest relying on a rather saturated fixed effects specification.

4 | REGRESSION APPROACH

In order to study the effects of higher labour costs on various firms' outcomes, we have exploited statutory changes in pay levels induced by collective bargaining. Since these shocks typically imply that firms need to adjust the wages of most of their workforce, contractual wage growth can be considered as a generalized growth in the cost of labour that hit all companies applying the same collective agreement.

We denote by w_{jt}^c the log median contractual wage bargained by a collective agreement c (that is, the median bargained pay level across the job titles defined by contract c). As mentioned, in cases where a firm applies more than one collective agreement to its employees, we have assigned this firm to the most expensive contract, that is, the one that covers the largest proportion of its wage bill. The subscript j is a firm identifier, while t denotes the year. Whenever contractual wages were renewed in the middle of a year, w_{jt}^c was defined as the weighted average of the two (or more) pay levels applied during the year, with weights representing the number of months during which each level was in place.

The baseline specification of our regression model reads as

$$y_{jt} = \beta w_{jt}^c + j * c + s * l * t + e_{jt}, \quad (1)$$

where $j * c$ is a firm by collective contract fixed effect,¹⁷ $s * l * t$ is an ISIC 38-sector (s) fixed effect, specific for 107 administrative provinces (l), interacted by a year fixed effect t .¹⁸ Finally, e_{jt} is the residual. Notice that sectors (s) and collective contracts (c) are different, albeit partially overlapping, categories. Indeed, collective agreements are often specific for given firms' characteristics within sectors, such as their size or corporate structure. Similarly, several contracts cover either heterogeneous activities that can be found in more than one industry, or very specific tasks within a single sector.¹⁹

The main firm-level outcomes considered in our analysis are the following:

$$y_{jt} = \left\{ \begin{array}{l} \log \text{ value-added per worker,} \\ \log \text{ TFP,} \\ \log \text{ FTE number of workers,} \\ \log \text{ average daily wages,} \\ \text{incumbents' average log wage growth,} \\ \text{average AKM worker fixed effects,} \\ \log \text{ revenues,} \\ \log \text{ profit margin,} \\ \log \text{ physical capital per worker.} \end{array} \right.$$

In the case of incumbents' wage growth, we have excluded firm fixed effects from the specification in equation (1), since the outcome is expressed in first differences within firms.²⁰ The TFP was computed using the Levinsohn and Petrin (2003) approach and adopting the value-added-based regression approach. Profit margins were defined as earnings before interests, taxes and depreciations divided by revenue. In Appendix Section A.4, we present complementary evidence on further outcomes, namely firms' closures and firms' switches in the collective contract applied to workers.²¹

Our treatment effect of interest (β) captures the effect of contractual wage growth on the outcomes of the above firms. This effect is estimated by comparing the evolution in the outcomes of firms affected by contractual wage changes with respect to a counterfactual group, represented by other firms in the same sector and geographic location that apply a different collective contract. Since equation (1) contains a firm by contract fixed effect and a year fixed effect that interacts with the fixed effects of 38 sectors (using the ISIC rev. 4 classification) and 107 Italian provinces, the regression model allows us to control for a rich set of unobservable factors, such as local business cycle fluctuations and time-constant differences between firms.

4.1 | Identification concerns and interpretative issues

Changes in contractual wages w_{it}^c could be potentially endogenous to the outcomes of our model, but this concern should not be overstated. Fanfani *et al.* (2024) show that contractual wage growth during the period of our study has been loosely related to productivity dynamics and mostly tied to coordinated inflation targets.²² More generally, contractual wage growth in Italy is bargained at a quite centralized level, and it would be difficult for trade unions and employers' associations to account for business cycle dynamics that are not captured by the non-parametric sector- and geographic-specific time fixed effects included in our regression model. Indeed, contractual wages are uniformly set at the level of the national sector, while the granularity of our data, coupled with the usual presence of several collective contracts within standard industry groups, allows us to control for a rich set of local industry-specific shocks potentially correlated to the wages settled in centralized negotiations.²³

An interpretative issue concerning several outcomes derived from firms' financial statements involves the role of output and input prices. Firms could indeed pass through the higher cost of labour on consumers, or they may cope with this shock by relying on cheaper intermediate inputs.²⁴ In principle, these price dynamics have an influence on most variables derived from balance-sheet and profit and loss accounts, including productivity. Our regression model accounts for differences in input and output prices across firms, as long as such differences remain stable across time among companies belonging to the same industry and province. Any residual variation in prices is nevertheless going to affect our results, even if quantifying the relative importance of such idiosyncratic price dynamics is difficult, given the unavailability of firm-level price data.

In this regard, the joint availability of rich information on the workforce of all employers allows us to provide a more solid interpretative framework to characterize which mechanisms may be driving the treatment effects observed for outcomes derived from firms' financial accounts. For example, a reduction in revenue that coincides in magnitude with a reduction in physical employment can be interpreted more easily as an effect driven by a reduction in production quantities, even if residual output price dynamics may in principle mitigate or strengthen the size of the treatment effect on revenue.

Another identification concern is related to the potential strategic behaviour of firms, which may decide to apply different collective agreements whenever a given contractual wage is raised. However, this possibility is typically limited by the law, according to which firms must apply the most representative collective contract given their activity. Moreover, the inclusion of firm

by contract fixed effects in the regression model ensures that the parameter of interest is identified only by variations in the outcome of interest within firms whose most expensive collective contract identity did not change across time.

In Appendix Section A.4, we explicitly account for the potential role of firms' self-selection across collective contracts. In particular, we have estimated a model where the outcome of interest is an indicator for firms that change the main collective contract applied to their workforce in the subsequent year.²⁵ Results from this test show that firms' propensity to switch collective contract is not influenced by contractual wage growth.

A related issue concerns the possibility that firms may decide to apply more extensively a less expensive collective contract to part of their workforce, even without changing the main one. In this respect, note that all of our outcomes of interest (including employment) are measured at the firm level and not at the contract–firm level. For example, if a decrease in employment in the main collective contract is compensated by a corresponding growth of workers hired under a secondary collective contract applied within the firm, then this change in workforce composition would have no influence on our firm-level employment measure. This consideration also suggests that in the presence of a similar endogenous reshuffling of workers across collective contracts, our results can be interpreted as a lower bound of the policy effects that would be observed if non-compliance opportunities were completely absent.

Finally, the model of equation (1) includes only a contemporaneous contractual wage term w_{jt}^c , even if adjustments to the wage shock may take time to materialize (Baker *et al.* 1999; Sorkin 2015). A standard 'event-study' analysis, where anticipatory and long-run effects are separately identified, would be unfeasible in the current setting. Contractual wages change quite frequently (on average more than once every 1.5 years), while our database is constructed at the yearly level. Thus it is not possible to identify suitable treated and control groups over a sufficiently long observation window.

Since our treatment variable is continuous and relatively persistent across time, as it is typically characterized by small stepwise increments, the treatment effect estimated in our static specification is also influenced by long-run adjustments to the contractual wage growth. This occurs because the relevant lagged values of w_{jt}^c that are omitted from the model tend to be highly correlated with the included contemporaneous term.²⁶ This bias towards the cumulative effect of the policy is going to be stronger, the stronger the serial correlation among lags and leads of the treatment variable. Section 7 discusses the results obtained using a dynamic version of equation (1) where leads and lags of w_{jt}^c are also included.

5 | BASELINE REGRESSION RESULTS

Table 3 provides the results obtained from the regression model of equation (1), which estimates the effect of the growth in contractual wages on several firms' outcomes. All regressions include year fixed effects interacted by 38 industry and 107 provinces fixed effects. Standard errors are always clustered at the collective contract level,²⁷ and regressions are weighted by the number of workers in the firm.

As Table 3 indicates, the baseline regression results show that, on average, the effect of higher contractual wages on value-added per worker is not significant.²⁸ A similar result is also found when using TFP as the dependent variable, which better accounts for heterogeneity in fixed costs across firms and for endogenous adjustments in the quantity of labour employed. Overall, our results suggest that higher labour costs do not trigger a generalized improvement in efficiency, which is consistent with previous findings in the context of the minimum wage by Draca *et al.* (2011), but which differs from other evidence on the UK (Riley and Bondibene 2017) and China (Mayneris *et al.* 2018).

TABLE 3 Effect of contractual wages on firm's outcomes—baseline regression results.

Regressor Outcomes	Log median contractual wage				Observations (millions)
	Coeff.	S.E.	Adjusted R^2	RMSE	
Log value-added per worker	0.022	0.189	0.769	0.274	2.988
Log TFP	-0.163	0.146	0.829	0.331	2.911
Log FTE employees	-0.785*	0.330	0.977	0.356	3.257
Log firms' average wages	0.262*	0.115	0.909	0.094	3.257
Incumbents' log wage growth	0.259**	0.048	0.062	0.001	2.642
Firms' average AKM worker fixed effects	0.025	0.038	0.948	0.043	3.186
Log revenues	-0.647**	0.248	0.937	0.403	3.167
Log profit margin	-0.272	0.205	0.668	0.499	2.484
Log physical capital per worker	-0.140	0.347	0.903	0.536	3.028

Notes: Results obtained by estimating the regression model of equation (1) on several firms' outcomes. All regressions are weighted by the number of workers in the firm. Standard errors are clustered at the collective contract level. The number of observations in each model is computed excluding singleton groups, i.e. units that are perfectly identified by the fixed effects included in the regression. AKM worker fixed effects were computed using the Abowd *et al.* (1999) regression model. TFP was derived from the Levinsohn and Petrin (2003) regression model.

*, ** indicate significance levels 5%, 1%, respectively.

In principle, higher wage levels could potentially affect productivity through several channels. On the one hand, there could be a reduction in managerial slack, which could be used to align workers' marginal product to the new pay levels, or an improvement in workers' effort (e.g. Coviello *et al.* 2022). On the other hand, there could be indirect effects on productivity triggered by firms' reliance on other adjustment margins, such as selective changes in the employment composition (e.g. Horton 2017; Clemens *et al.* 2021), size reductions, output price increases, or higher investments in capital goods. Thus it is interesting to investigate which other adjustment margins firms rely on when facing higher labour costs.

The third row of Table 3 shows that the average effect of higher labour costs on firms' employment is negative and sizeable. Indeed, a 1% growth in contractual wages is associated with reductions in employment by almost 0.8%. Strong negative employment effects associated with contractual wage growth are also documented by Fanfani (2023) using the entire private-sector workforce and monthly-level administrative data for Italy. Several factors may concur in rationalizing this effect. First, contractual wages affect virtually all workers across the pay distribution, not only those at the bottom, as is the case for minimum wages. Second, economic growth has been always close to zero or negative during the period of our study, and statutory pay growth has been shown to be more detrimental for employment during economic downturns (Clemens and Wither 2019). Finally, the very low inflation levels that have characterized the period of our study could also have contributed to strengthening the size of the employment effects, as low inflation implies that nominal contractual wage increases have not eroded rapidly in real terms.²⁹

The fourth row of Table 3 shows that the elasticity of firms' average wages to the growth in contractual wages is positive and significant, but also smaller in magnitude than the respective employment elasticity. These results suggest that employment losses related to higher labour costs are more proportional than the wage gains. Average firms' wages can be also influenced by employment composition, which changes over time. A more accurate measure of the direct effect of contractual wages on workers' pay is provided in the fifth row of Table 3, which shows that a 1% growth in contractual wages is associated with a 0.26 percentage point increase in the growth rate of incumbents' wages. Incumbents' pay growth is indeed computed on a fixed population of workers within firms, thus it is not affected by differences in composition across time.

Note that contractual wage growth induces a less than proportional effect on the actual pay growth of incumbent workers. The size of this marginal effect can be rationalized by at least two mechanisms. First, contractual wage growth in principle should not directly affect the size of the wage cushion (i.e. the part of the pay representing the positive difference between actual and minimum statutory wage levels).³⁰ Second, the marginal effect of our model represents a relative difference between firms affected by positive contractual wage growth, and a counterfactual group for which pay floors are kept constant. Thus if actual pay growth is still positive (but weaker) even in the absence of contractual wage growth, then the treatment effect for this outcome can be less proportional than the difference in contractual wage growth between the treated and control groups.

In order to shed more light on the adjustment channel of employee selection, we have estimated a measure of workers' quality based on the AKM regression model. This technique, which is presented in Appendix Section A.3, allows us to recover an estimate of worker fixed effects that is conditional on observable characteristics and on firm-specific pay policies. Thus these worker fixed effects can be interpreted as a measure of the employees' time-constant earning abilities. Since they are, by definition, constant across time, a firm can change the average level of its employees fixed effects only through selective hiring and firing.

The sixth row of Table 3 shows that the average quality of the workforce (defined using AKM workers' fixed effects) is actually unaffected by higher labour costs. However, the next section further characterizes and discusses this result, by showing that the treatment effect is instead highly heterogeneous and different from zero across the distribution of firms' productivity, suggesting that companies rely on selective employment adjustments, but use differentiated strategies depending on their efficiency levels.

The seventh row of Table 3 shows that firms' revenues are negatively affected by the growth in contractual wages. This result suggests that the employment losses discussed previously also translate into lower sales. As discussed above, revenue is made up of two components, output prices and quantities. In our context, the presence of negative revenue effects seems likely to be driven mostly by quantity reductions, as the regression model controls for sector-wide price shocks at a quite granular level through sector- and geographic-specific time fixed effects, as well as for time-constant firms' heterogeneity in market power. Moreover, the reductions in physical employment that we have documented appear to be consistent with a drop in physical production levels. In this respect, the presence of potential pass-through mechanisms of higher wage floors to consumers via increases in product market prices would actually characterize the elasticity of revenue to contractual wages as a downward-biased measure of the true effect on output quantities.

Finally, the last two rows of Table 3 show that the profit margin and the intensity in the use of physical capital are not affected by the growth in the cost of labour. The null effect on capital intensity implies an overall reduction in investments, since employment is also negatively affected by the shock. Thus firms seem to have limited possibilities for adopting more capital-intensive (and potentially more productive) production processes (see, for example, Acemoglu (2003) for a theoretical discussion of this point). Instead, the wage shocks to the firm analysed in this application produce mostly scale effects, as on average investments are reduced together with employment and production levels.

Overall, the results from our baseline regression model show that on average, firms responded to higher contractual wages by decreasing production levels and employment. Average firms' wages are instead increased by this shock. Moreover, there are no effects on firms' productivity, on quality of employees, on the profit margin, or on capital intensity. From an aggregate perspective, these results suggest that growing contractual wages contribute to a modest increase in the labour share or, equivalently, to higher unit labour costs (given the positive effect on wages and the null effect on productivity), reducing total output and, more generally, the international competitiveness of Italian companies (see Dustmann *et al.* 2014). The next section further characterizes

these results, by uncovering the heterogeneity in adjustment behaviour along the distribution of firms' productivity.

6 | EFFECTS OF CONTRACTUAL WAGES ACROSS THE PRODUCTIVITY DISTRIBUTION

This section presents the heterogeneity in the effects of contractual wage growth across the distribution of firms' productivity, which we define as value-added per worker. Productivity has often been considered a dimension along which the effects of centralized wage bargaining can be heterogeneous.³¹ According to models of firms' entry and selection in the market, a centralized wage favours relatively more efficient competitors (Moene and Wallerstein 1997; Barth *et al.* 2014). The previous theoretical and empirical literature that has analysed collective bargaining in Italy has also highlighted the fact that productivity differentials do not map very well onto wage differentials (Boeri *et al.* 2021). In this context, it is reasonable to assume that the effects of centralized wage growth could be differentiated across the productivity distribution of firms hit by this shock.

6.1 | Estimation method

To describe heterogeneities in the response to contractual wage growth, we have adopted a regression specification similar to the model provided in equation (1). Using the same notation, the estimated model reads as

$$y_{jt} = \sum_{\theta=1}^4 \beta_{\theta} w_{jt}^c * q(\theta)_j + j * c + s * l * t + e_{jt}. \quad (2)$$

We have interacted contractual wages w_{jt}^c with four indicator variables $q(\theta)_j$, denoting time-constant quartiles of productivity to which each firm j belongs. This approach is similar to running separate regressions on four split samples, but we have used interactions on the full dataset to improve efficiency in the estimation of β_{θ} , which represent the four parameters of interest.

The quartiles $q(\theta)_j$ were time-constant, defined using the contract- and year-specific distribution of value-added per worker. We made these quartiles always constant over time for each firm, by assigning them to their most common quartile across all years. For example, a metal-manufacturing firm was included in the fourth quartile of productivity if it was among the highest value-added per worker companies in the metal-manufacturing collective contract during most of the years of observation.

To further test the robustness of this approach, we also classified firms across the contract-specific productivity distribution adopting an alternative strategy. In particular, we considered value-added per worker in the years 2004–6, and we assigned firms to the most common quartile of the contract-specific distribution of this productivity measure. Then we estimated the heterogeneity effects of model (2) using the years 2007–15.

A more detailed discussion of the advantages and potential weaknesses of each of these two approaches is provided in Appendix Section A.2, along with alternative specifications estimated to check the overall robustness of the results. In general, a drawback of the approach of ranking firms in the pre-estimation period 2004–6 is that only firms surviving in subsequent years can then be used, with a loss of around 60% of the observations. This source of selection tends to over-sample relatively healthier, larger and more productive firms. For this reason, we think that the approach of ranking firms according to their most common quartile over the period 2006–15

allows us to draw conclusions on a sample that is more representative of the underlying firms' population.

With this approach, firms are assigned to quartiles based on their most common level of productivity, rather than based on the unconditional trend of this variable. Thus it seems reasonable to assume that selecting firms with more positive (negative) marginal effects into higher (lower) productivity quartiles is mostly irrelevant.³² Nevertheless, we show in parallel the main results using both classification approaches, one based on the contemporaneous period and one on the pre-estimation period, in order to provide a complete overview of how results are affected depending on the approach adopted.

6.2 | Overview of the main results

Figure 4 presents the main results obtained by estimating the model of equation (2). In each panel, the first quartile refers to the lowest time-invariant productivity quartile of firms, while the fourth refers to the most efficient one. Each panel shows the marginal effect of higher contractual wages for each quartile, together with the 95% confidence interval. Figure 5 shows the heterogeneity results derived from the alternative approach of estimating the model on the years 2007–15, and using the pre-estimation measure of productivity observed in 2004–6 to classify firms into quartiles.

In general, when comparing Figures 4 and 5, results on all outcomes look qualitatively similar. Nevertheless, two main differences emerge. First, estimated coefficients in Figure 5 tend to be smaller in absolute value. Second, some of the results that are only marginally significant in Figure 4 are sometimes not significant when using the alternative classification of firms. This is the case for the positive productivity, revenue and profit effects among the most productive firms, which are not statistically significant in Figure 5. In part, this could be a consequence of the smaller sample size, given that around 60% of the observations are lost with the approach of Figure 5. A more detailed discussion on potential sources of differences between the two approaches, as well as results derived from further alternative specifications, are provided in Appendix Section A.2.³³

That said, differences across the two sets of estimates are quite limited from a qualitative perspective, particularly regarding comparisons between relatively less and more productive firms. For this reason, and to keep the discussion of the results shorter and easier to follow, in the remainder of this section we will refer mostly to the results of Figure 4. Nevertheless, the presence of small differences across estimates between Figures 4 and 5 should be taken into consideration when evaluating the robustness of this evidence.

Figure 4(a) shows that higher contractual wages had a positive effect on pay levels among incumbents for all types of firms.³⁴ Thus contractual pay levels have a relatively uniform and positive effect on wages for all types of firms, and similar evidence emerges from Figure 5. This symmetrical effect can be rationalized considering the Italian legislation. Indeed, in Italy, contractual wages are considered not only a statutory minimum, but also a fixed pay component. This implies that when contractual wages grow, in principle all wages should grow by the same fixed amount, irrespective of whether workers were already paid above the new pay floor.

Figures 4(b) and 4(c) show that the overall null effect of higher labour costs on productivity is instead heterogeneous along its distribution. It is strongly negative for less efficient firms, and small and positive for the two highest quartiles of the productivity distribution. The small positive effect for productive firms is not significant in Figure 5, but a similar heterogeneous gradient between less and more efficient firms emerges still in these latter estimates. The joint analysis on several outcomes allows us to uncover more precisely how the adjustment margins of the most efficient firms differ from those exploited by less productive companies. Two main mechanisms seem to emerge from this analysis. First, efficient firms gain market shares with respect to less

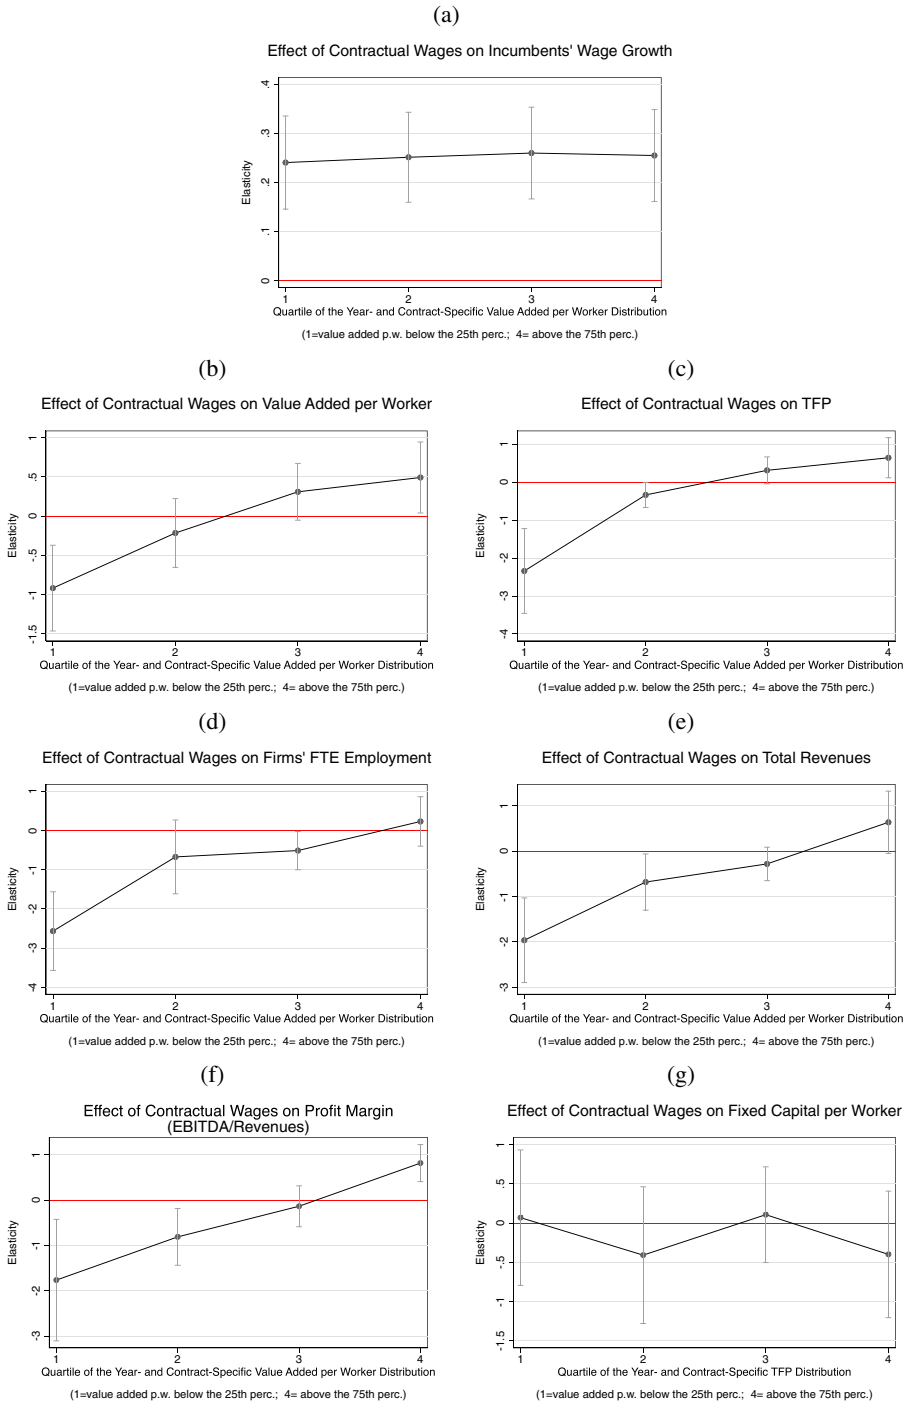


FIGURE 4 Wage growth effects across the contract-specific distribution of value-added per worker—most common quartile classification. *Notes:* Estimates of model (2), where value-added per worker quartiles $q(\theta_j)$ are defined using the most common quartile to which firms belong in the collective contract-specific distribution over the period 2006–15.




FIGURE 5 Wage growth effects across the contract-specific distribution of value-added per worker—2004–6 quartile classification. *Notes:* Estimates of model (2), where value-added per worker quartiles $q(\theta)$, are defined using the most common quartile to which firms belong in the period 2004–6, and estimating the model on the years 2007–15.

productive competitors. Second, employment levels and composition are adjusted differently, depending on the relative productivity of each firm.

6.3 | Employment, revenue, and cleansing effects

Figure 4(d) shows that the negative employment effects of higher labour costs are mostly borne by low-productivity firms, while they are not significantly different from zero in the highest quartile. Thus firms that are more efficient tend to be more resourceful, as they are able to absorb the cost shock without incurring employment losses. This suggests that productive companies exploit non-competitive adjustment mechanisms that rely on rents that they may possess in the product or labour market. In this subsection, we show that these differences in adjustment mechanisms across firms provide profitable externalities to dominant firms within a collective contract.

The hypothesis that higher labour costs could provide a barrier to entry by dominant companies within a sector has a long tradition in economics (Williamson 1968). Vintage models of firms' survival in the context of centralized wage-setting, as developed by Moene and Wallerstein (1997) and Barth *et al.* (2014), similarly suggest that employment reallocation towards more productive firms could emerge in a context where wage moderation imposes similar pay levels to a heterogeneously efficient population of firms. The recent literature on regional misallocation further suggests that wage compression could produce perverse welfare effects and excess rents among employers in most productive regions (see, in particular, Boeri *et al.* 2021). However, the importance of this mechanism has rarely been investigated in the empirical literature.

Figure 4 shows that while revenue declines as a consequence of contractual wage growth for less efficient firms, it actually slightly increases in the highest quartile of most efficient companies. The positive effect on revenue at the top is only marginally significant in Figure 4, and not significant in Figure 5, but the same gradient between more and less productive firms emerges in both sets of results. As mentioned, revenue is made up of prices and quantities. However, given the inclusion of year fixed effects by industry and geographic location, we can assume that heterogeneity in revenue effects across the contract-specific productivity distribution is mostly driven by relative differences in quantities produced. This implies that contractual wage growth increases the product market shares of more productive companies, as long as firms belonging to the same collective agreement are also likely to share the same product market. Moreover, considering that our evidence shows that there are negative effects on productivity for less productive companies, their reduction in production could constitute a shock away from an optimal scale level.³⁵

An alternative hypothesis could be that efficient companies systematically select into segmented product markets characterized by a rigid demand or by monopolistic power, where price or markup adjustments are possible. However, the drops in revenue and employment are very similar in shape and relative magnitude across the productivity distribution, which suggests that revenue changes largely map onto output quantity changes, rather than to changes in output prices. Thus market shares of productive companies are likely to increase as a result of contractual wage growth, as heterogeneity in revenue effects does not seem to be driven simply by heterogeneity in price adjustments. That said, the presence of output price mechanisms influencing our results cannot be completely ruled out. In particular, reduced production levels among low-productivity firms, which are suggested by the symmetric reduction of revenue and inputs (labour and capital) for these firms, may induce an upwards effect on output prices at the market level, which could explain the small positive revenue effects among relatively more productive firms, even if this effect is statistically different from zero only for high significance levels.³⁶

Another test for the hypothesis that differential cost-price pass-through across firms is not the major mechanism driving our results is provided by Appendix Figure A1. In this robustness test, we interact the treatment effect with a dummy variable indicating industries (at the three-digit level) characterized by a nationwide Herfindahl index above the median.³⁷ The industry-wide

Herfindahl index, computed on the basis of firms' revenue in each industry, can be considered a proxy for the level of product market competition and the price-setting power of firms. As can be noted, even when separately estimating revenue and employment effects within more versus less competitive industries, a similar shape emerges across the productivity distribution for both outcomes. However, this test could be too weak to rule out any potential heterogeneity in price adjustment mechanisms.

Notice that when considering profit margins, the effect of higher labour costs appears to be negative for relatively less productive firms, while the effect is instead slightly positive in the quartile of most efficient companies. The positive effect at the top is not significant in Figure 5, but a similar gradient between more and less productive firms also emerges according to this set of estimates. This heterogeneity in profit effects indicates that companies that are able to deal with higher labour costs without cutting production may actually increase their performance due to a cleansing effect on less productive competitors. The hypothesis that increasing one's own and rivals' costs can be profitable has a long tradition in the industrial organization literature (e.g. Salop and Scheffman 1983), and it seems to be relevant for our results. The recent literature on cleansing effects during recessions (see, in particular, Foster *et al.* 2016; Osotimehin and Pappadá 2016) has shown that positive selection of firms can be limited during severe downturns, due to potential distortions in the credit market. In this respect, our evidence suggests that input cost shocks, as opposed to negative movements in the product market demand, tend to be more cleansing as they provide a competitive advantage to more productive firms.

To further explore this issue, Appendix Section A.4 provides an analysis of the impact of contractual wage growth on firms' exit rates. Overall, we did not find significant changes in firms' closure rates in response to higher labour costs in our sample of analysis, which is made up of incorporated businesses only. However, when extending this analysis to the entire INPS archives covering the private sector, we found significant increases in firms' exit rates among very small companies, that is, those with fewer than five employees. This evidence suggests that relatively larger companies reduce output levels on the intensive margin, but the cost shock related to contractual wages is not strong enough to drive these employers out of the market. However, this can be the case among very small firms, which are less likely to be included in our main sample of analysis, as they tend to be unincorporated, while they also represent the group for which the wage bill tends to be more relevant in proportion to total costs. Overall, even if the more extreme event of a firm's closure was not affected by contractual wages in our sample, this mechanism appears to be relevant when focusing on a subgroup of firms that are smaller and more intensively hit by the labour cost shock. The exit channel thus further contributes to the determination of 'cleansing effects', which refer to a combination of firm closures (in the case of small enterprises) and a redistribution of market shares towards more efficient firms within the sample of surviving incorporated business.

Figure 4 also shows that capital intensity is not affected by wage growth across the entire firms' productivity distribution. This evidence entails that investments in physical capital decrease at less efficient companies, for which we have documented a strong and significant employment reduction as a consequence of the contractual wage shock. It also implies that there is a reallocation of capital away from less productive firms, even if it is not accompanied by detectable improvements in investment levels among more efficient companies.³⁸

6.4 | Employment composition effects

In order to provide evidence on workforce selection, Figure 6 shows the effect of contractual wage growth on several measures of workforce composition. Figure 6(a) shows the effects on employee earning potential (quality for short), as measured by AKM workers' fixed effects. While this measure tends to increase among the least efficient companies, which also reduce employment,

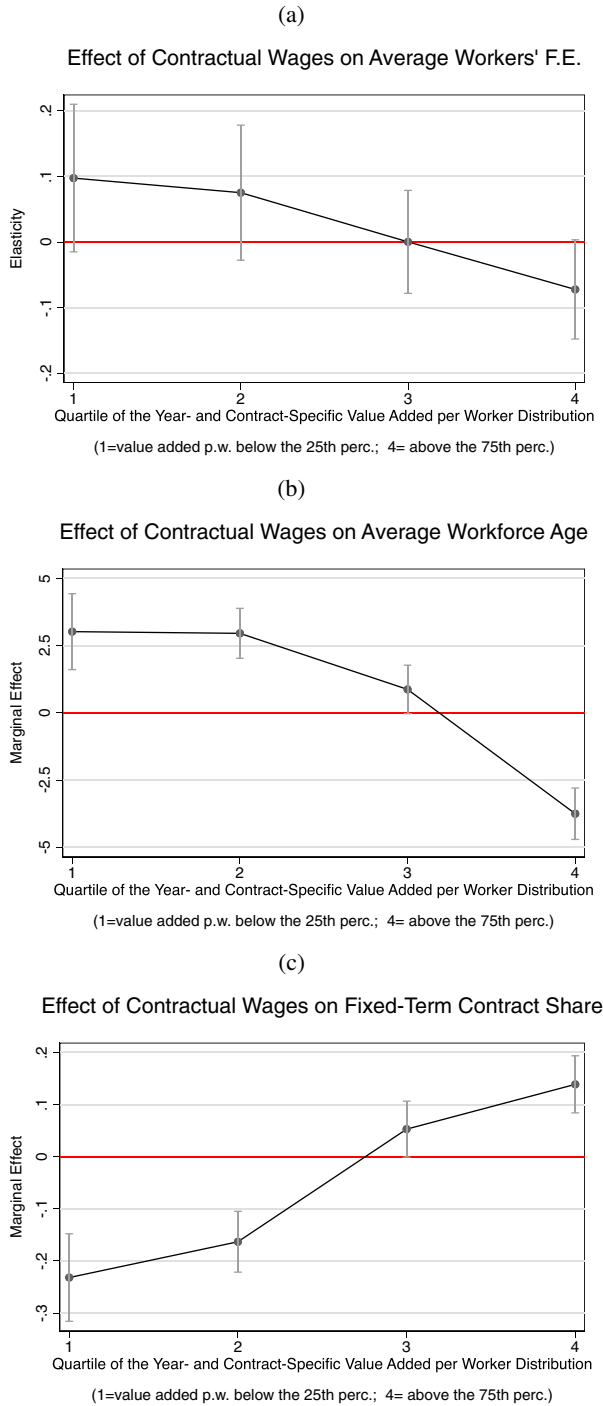


FIGURE 6 Heterogeneity of wage growth effects across the contract-specific distribution of value-added per worker—workforce quality, age and fixed-term contract share.

the quality of workers actually drops among the most productive firms. Further evidence on this pattern is provided by considering as outcomes of interest the average age and the proportion of fixed-term employees within firms.³⁹ Figures 6(b) and 6(c) show that, consistent with the pattern emerging from the analysis of AKM workers' fixed effects, average workers' age drops at more productive firms, and it increases at less productive ones, whereas the share of fixed-term workers increases at efficient firms, while it decreases elsewhere.

The patterns of workforce selection described above are consistent with several mechanisms. First, low-productivity firms, which tend to cut employment, may decide to rely less on workers who are less protected by firing costs, changing their skill mix. Indeed, most of the employment losses observed at these companies hit marginal workers, that is, those who have low earning potential, who are younger, and are on a fixed-term contract. Evidence of similar mechanisms is provided in the context of the minimum wage by Horton (2017) and Clemens *et al.* (2021). However, productivity does not grow among less efficient firms as a result of this process, which suggests that their negative employment adjustments may be suboptimal due to frictions and potentially high firing costs, which are considerably more pronounced for older employees with open-ended contracts.

On the other hand, productive firms hit by the wage shock do not cut overall employment. In this context, their more intensive reliance on low-quality workers could be a cost-saving strategy. Indeed, the reduction in the average age of the workforce among efficient companies seems consistent with a 'young in, old out' strategy, where typically more expensive older workers with open-ended contracts are pushed to retire, while firms start relying more on cheaper young and fixed-term workers.⁴⁰

Finally, when considering general equilibrium effects of the labour cost shock, it is also possible that productive companies are able to absorb part of the job losses observed at their less efficient competitors. In this context, as low-quality workers become more easily available in the labour market, reallocation mechanisms, emphasized also in the context of minimum wage policies by Dustmann *et al.* (2022), could be a relevant channel through which less productive employees end up in more productive firms.

7 | EFFECTS OF CONTRACTUAL WAGES ACROSS TIME

This section discusses the role of dynamics in firms' adjustment behaviour. Moreover, it discusses some interpretative issues on the relationship between cumulative policy effects and our main regression approach.

7.1 | Direct estimates of anticipatory and long-run effects

Several of the outcomes considered in the main analysis, in particular those related to employment, investment decisions and productivity, could take time to adjust to higher labour costs. Moreover, our identifying assumptions may become more credible by testing and rejecting the hypothesis that contractual wage changes had an influence on the dependent variable long before their actual occurrence. Failing to reject this hypothesis would suggest that systematic differences exist between treated and control firms even conditioning on the model fixed effects.

For these purposes, we have directly estimated anticipatory and long-run effects of contractual wages on the main outcomes of interest relying on the following distributed lags regression model:

$$y_{jt} = \sum_{i=-1}^2 \beta_i w_{j(t+i)}^c + j * c + s * l * t + e_{jt}, \quad (3)$$

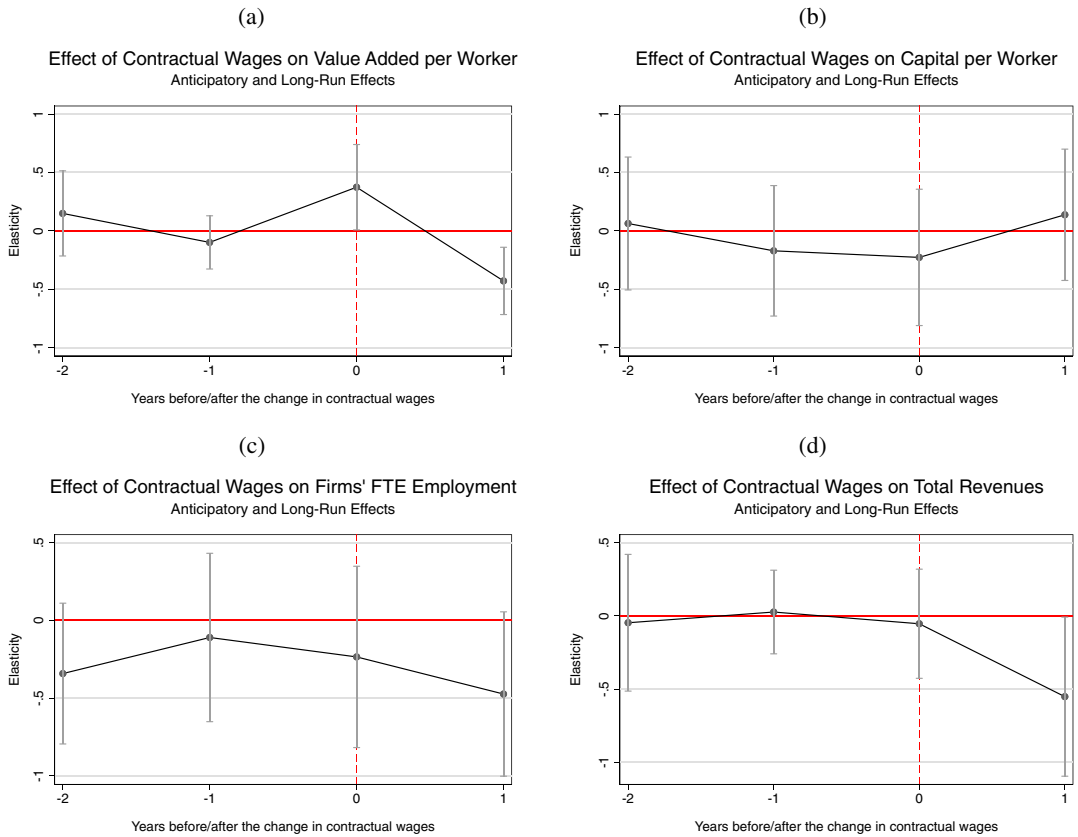


FIGURE 7 Long-run and anticipatory effects of wage growth.

where we include the contemporaneous level of w_{jt}^c together with two leads and one lag. In this model, we have adopted a specification similar to equation (1), including firm by contract and year by sector and province fixed effects. However, the inclusion of two leads of contractual wages allows us to estimate anticipatory policy effects, while the lagged term allows us to study long-run adjustments one or more years after the change in contractual wages has occurred.

This model has some limitations, as it tends to suffer from severe multicollinearity due to the autocorrelation in contractual wages, so that its results can be quite volatile. The model's volatility is a consequence of the fact that the same variation that was used to estimate one treatment effect in the specification of equation (1) is now used to estimate four treatment effects. Moreover, the strong persistence in contractual wages, which tend to be adjusted by small increments rather than drastically reduced or increased across time, makes the correlation among the terms $w_{j(t+i)}^c$ quite strong.⁴¹

The distributed lag model has often been estimated in the minimum wage literature (see, for example, Meer and West 2016; Cengiz *et al.* 2019). However, as noted by Cengiz *et al.* (2019), this model is more demanding than standard falsification tests in event study analyses, as distributed lags also measure the presence of differences in outcome trends in periods long before the policy change. Indeed, in this model the first $t + i$ lead (last $t + i$ lag) is typically interpreted respectively as the effect of the policy i years *or more* before (after) its level change.

Figure 7 presents the results obtained by estimating the dynamic model of equation (3) for value-added per worker, capital per worker, employment and revenue. In all panels, the confidence intervals are computed at the 5% significance level using standard error clustered at the

collective contract level. These outcomes were selected for several reasons. First, as mentioned, investments, productivity and employment could take time to adjust to higher labour costs. Thus it is interesting to study potential differences between short- and long-run adjustments. Second, the presence of significant differences in revenue, employment or productivity trends long before contractual wage shock occurrence could be indicative of endogeneity problems in our estimation approach.

Figure 7 shows that the effects of higher contractual wages on productivity were not significant before the policy change; they were slightly positive in the short run, and negative afterwards. Overall, this result is consistent with a cumulative effect that is close to zero, but there are small short-run productivity gains lost one year after the shock in labour costs. In the short run, revenue and employment levels were not affected by the policy change, while this effect was negative in the longer run for both outcomes. This suggests that firms take time to adjust production and employment levels when facing higher costs, which is consistent with the presence of frictions that slow down adjustment effects.⁴² Instead, no significant effects are detected regarding capital intensity even one or more years after the policy change. However, it is possible that the relatively short estimation period of our analysis is not well suited to capture potential technological shifts induced by higher labour costs that may gradually take place over long periods of time.

Results show that contractual wage growth had no influence on any of the outcomes considered two or more years before its occurrence. As mentioned, this marginal effect can be interpreted as a placebo test on the parallel trend hypothesis that should hold in the presence of a correct identification strategy. Indeed, the assumptions of our model could be questionable in the presence of significant differences in conditional employment or productivity growth between treated and control firms before the labour cost shock took place. Thus the fact that none of the coefficients associated with the two-year lead was statistically significant can be interpreted as evidence supporting the validity of our main identifying assumptions.

It is interesting to note that anticipation effects are not relevant, even if firms may be aware of future contractual wage growth. This growth is the result of centralized negotiations, and the typical Italian firm (much smaller compared to its German or French counterpart) may have limited influence in this bargaining process. For this reason, the outcome of contract renewals can be largely considered an unexpected shock for most individual employers. However, in many cases, contract renewals establish a sequence of wage adjustments that will take place during the period of validity of the contract (usually two years). Thus part of the contractual wage growth can be predicted by firms as it is announced before its implementation. However, the relevance of anticipation effects seems quite limited according to the evidence in Figure 7.

In order to further test for the robustness of the results, we have also estimated the model of equation (3) separately for firms in the two top and two bottom quartiles of productivity, defined as their most common in the years 2006–15.⁴³ This specification is even more demanding than the one estimated on the full sample, given the reduction in the available identifying variation when using the split sample approach. Appendix Table A2 provides the results separately for firms below (plots (a), (c), (e), (g)) and above (plots (b), (d), (f), (h)) the median contract-specific productivity level.

As can be noticed, for all outcomes and both firms below and above median productivity there are no significant pre-trends in the first period, evidence that supports the validity of our identifying assumptions. Revenue and employment effects are negative and significant only among low-productivity firms, and only in the period after the occurrence of contractual wage growth. Productivity effects are significant and negative, with some anticipation close to the contractual wage growth event in the case of less productive firms. They are positive and significant in the year of the contractual wage growth event, then negative in the subsequent period in the case of most productive firms. Capital intensity effects are never statistically significant.

Overall, albeit partly noisy, the evidence of Appendix Figure A2 provides results that are broadly consistent with our main findings of null effects on revenue and employment among

TABLE 4 Employment effects of contractual wage growth within restricted time windows.

Outcome Regressors	Log FTE employees		
	Coeff.	Coeff.	Coeff.
Log median contractual wage	-0.785* (0.330)	-0.512* (0.227)	-0.467* (0.197)
<i>Fixed effects</i>			
Contract by firm	✓	✓	✓
Time by sector and location	✓	✓	✓
Contract by 5 years window		✓	
Contract by 3 years window			✓
Adjusted R^2	0.977	0.977	0.977
RMSE	0.331	0.355	0.355
Observations (millions)	3.257	3.257	3.257

Notes: Results obtained by estimating the regression model of equation (4) on employment. All regressions are weighted by the number of workers in the firm. Standard errors (in parentheses) are clustered at the collective contract level. The number of observations in each model is computed excluding singleton groups, i.e. units that are perfectly identified by the fixed effects included in the regression.

*, ** indicate significance levels 5%, 1%, respectively.

productive firms, and negative ones on less productive firms. Productivity effects are also broadly consistent, even if the negative productivity effects among less productive firms seem partly anticipated, while they oscillate between positive and negative effects in the case of the most efficient companies.

7.2 | Influence of long-run effects on the static specification

The main results discussed in this paper were based on static fixed effects specifications. Given the typical evolution of contractual wages (see, for example, Figures 1 and 3), the independent variable of interest in our model follows a persistent dynamic trend, with leads and lags that tend to be highly correlated. This autocorrelation implies that the static specification adopted for the main results should also be influenced by long-run adjustments, as it is likely biased towards the cumulative effect of contractual wage growth. Indeed, long-run adjustments would be captured by the coefficient associated with a lagged level of contractual wages. Since this term is omitted in the static specification, and it is positively correlated with the included contemporaneous contractual wage level, long-run effects should influence the static coefficient according to the standard omitted variable bias formula.⁴⁴

To test this hypothesis indirectly, we have estimated the static regression model adopting an alternative specification. In particular, we have transformed the model of equation (1) as follows:

$$y_{jt} = \beta w_{jt}^c + j * c + s * l * t + c * p + e_{jt}, \quad (4)$$

where $c * p$ is an interaction between collective contracts fixed effects and a time-window fixed effect p . In particular, we have defined p as either a 5-year or 3-year time window. With this specification, the coefficient of interest is estimated using only variation across time within a restricted sub-period.⁴⁵ Thus in the specification of equation (4), long-run effects should contribute less to β , and this coefficient should be closer to zero in absolute value if long-run adjustments had the same sign of the contemporaneous effect.

Table 4 shows the results of this specification estimated on employment, which is one of the main outcomes of interest. This is also a variable that, according to Figure 7, adjusts

relatively slowly across time, possibly due to frictions and rigidity in the labour market. As can be noted, the estimated coefficient reduces in absolute value as the period over which it is estimated gets shorter. However, the result is still negative and statistically significant when using a 3-year window fixed effect interacted with collective contract fixed effects. Overall, Table 4 is consistent with the hypothesis that the main results presented in the paper partly reflect the influence of long-run adjustments, and that they tend to be biased towards the cumulative policy effect.⁴⁶

The above considerations suggest that our main static regression approach is well suited to study policy effects that develop across time, rather than through sharp adjustments. This seems to be the case for the present application, where the presence of frictions, such as those arising from the labour or capital market, could be relevant. In such contexts, adjustments to higher pay floors could be attenuated in the short run and carried out over longer time periods (see, for example, Pinoli 2010).⁴⁷ Thus the bias towards the cumulative policy effect of the static estimator can be considered a desirable property in our context.

8 | CONCLUSIONS

Our analysis shows how higher contractual wages set by Italian collective bargaining affect the behaviour of firms. On average, the growth in contractual wages induced firms to cut employment and revenue. Instead, companies' average wages increased, while workers' quality, productivity, capital intensity and the profit margin were not affected by the shock.

When looking at the heterogeneity in adjustment behaviour across the productivity distribution, higher labour costs induced a small or null growth in efficiency for more productive firms, and a strong decline for the least efficient ones. We have argued that this heterogeneity in efficiency effects is driven partly by cleansing mechanisms that increase the product market share of relatively more productive firms. Consistent with this hypothesis, we have found that relatively more efficient companies within a sectoral collective agreement increase (or keep constant) their revenue, they do not cut employment and investments, and they slightly increase (or keep constant) their profit margin in response to higher contractual wages. We did not find differences in exit rates among incorporated companies after a growth in labour costs, but we have found evidence of increased firms' closures when extending the sample of analysis to the entire private sector and focusing on very small firms.

Our results have more general implications, as they show that increases in relative labour costs can have nuanced effects on the economy, decreasing production levels and employment on average, but providing the most productive establishments with a competitive advantage. Profitability of such companies can even increase due to greater product market shares, as rival costs are raised (e.g. Salop and Scheffman 1983; Williamson 1968). Cleansing mechanisms have been discussed in the literature with reference to other kinds of firm-level shocks. They have been linked to the presence of credit market imperfections, which tend to hit firms relying more on external finance (see, for example, Pagano and Pica 2012; Giroud and Mueller 2017), and which induce companies to increase workers' quality (e.g. Berton *et al.* 2018). Instead, cleansing mechanisms have been found to be potentially weaker in the context of strong negative demand shocks (e.g. Foster *et al.* 2016; Osotimehin and Pappadá 2016). In this respect, the statutory increases in wage levels considered in our analysis seem to generate a more pronounced positive selection in the underlying composition of firms.

Our results are consistent with hypotheses linking average productivity to wage-setting structures, as developed by Acemoglu (2003) to explain cross-country heterogeneities in inequality and productivity, or the vintage approach theories that explain differential survival rates of firms across the efficiency distribution in the context of collective bargaining (see, in particular, Moene and Wallerstein 1997; Barth *et al.* 2014). However, our results show that productivity gains related

to higher wage floors are not significant on average. Moreover, reallocation effects towards more productive firms are accompanied by overall reductions in employment levels. Less skilled workers tend to suffer most of the employment losses at less efficient companies, but, consistent with workers' reallocation evidence documented by Dustmann *et al.* (2022) in the context of the German minimum wage, we also find that their employment share increases at the most efficient firms.

This evidence contributes to the literature on the southern European productivity puzzle (see Calligaris *et al.* 2018; Schivardi and Schmitz 2020), as it shows that labour market institutions have a relevant effect on the allocation of resources and market shares, which could potentially influence even management practices. The mechanism emphasized in this study also shows how institutional factors related to the system of industrial relations may contribute to concentration in the product market, for which the recent literature has documented a positive secular trend (De Loecker *et al.* 2020). Moreover, our results further support the conclusions of studies on the relationship between centralized wage-setting and regional misallocation (see, in particular, Manacorda and Petrongolo 2006; Boeri *et al.* 2021), as we have shown the relevance of several mechanisms emphasized by this literature using granular data and a causal research design. Finally, by documenting the presence of rather sizeable adjustments while analysing wage shocks that were more extensive than those typically arising in the context of minimum wage hikes, where instead pay rises tend to affect only the bottom of the income distribution, we have provided novel evidence that could help to rationalize the elusive effects of the minimum wage discussed by Manning (2021).

ACKNOWLEDGMENTS

This paper was made possible thanks to the sponsorships and donations in favour of the 'VisitINPS Scholars' Type B programme promoted by the Istituto Nazionale per la Previdenza Sociale (Italian Social Security Institute, INPS). We gratefully acknowledge financial support of MIUR-Ministry of Education, University and Research (Financing Fund for Departments of Excellence), Italy. We would like to thank John Addison, Alex Bryson, Pierre Cahuc, Lorenzo Cappellari, Pietro Garibaldi, Claudio Lucifora, the VisitINPS staff, members of CRILDA and LABORatorio Revelli, and participants at various conferences and seminars for their useful comments.

The data used in this paper are the property of the INPS, and are accessible by researchers at the INPS premises through the VisitINPS programme. To access data for replication purposes, researchers should contact INPS DC Research (dcstudiericerche@inps.it). The findings and conclusions expressed are solely those of the authors and do not represent the views of the INPS.

ENDNOTES

- ¹ In 2016, the 150 largest collective contracts alone covered around 92% of all private-sector workers in the country, affecting the pay levels of more than 15 million employees.
- ² The growth of a contractual pay floor implies that virtually all wages in the relevant occupation have to be increased by the same fixed amount, even those that already complied with the new minimum. Thus, different from the case of a minimum wage, all workers (including those at the top of the pay distribution) tend to be affected by this policy.
- ³ These mechanisms, and the limited available evidence on them, have recently been discussed by Clemens (2021) with reference to the minimum wage literature. Other studies analysing the effects of direct and indirect labour costs on firms' performance and employment have focused on the role of union density (e.g. Addison and Hirsch 1989; Barth *et al.* 2020) and the tax burden (e.g. Cahuc *et al.* 2018). Card and Cardoso (2022) provide a recent contribution to this literature focusing on the effects of collective bargaining in Portugal.
- ⁴ For example, we may be comparing two firms, one applying the leather production collective contract, the other the footwear production collective contract, both belonging to the leather and related products manufacturing two-digit sector, and both located in the same Italian province.
- ⁵ Similar arguments have also been proposed in the minimum wage literature. For example, Dustmann *et al.* (2022) show that the German minimum wage has induced a shift of employment towards more efficient firms.
- ⁶ Several adjustment margins considered in this paper have been analysed in the context of standard minimum wage policies. See, in particular, Aaronson and French (2007) and MaCurdy (2015) for a discussion of product market price effects, Draca *et al.* (2011) for a discussion of the effects on profits, Riley and Bondibene (2017) for a discussion of the

- effects on productivity, and Harasztsi and Lindner (2019) for a joint analysis of several firms' adjustment margins. A detailed literature review is provided by Clemens (2021).
- ⁷ The recent literature analysing cleansing effects, reallocation and labour-hoarding hypotheses in the presence of adverse shocks to the firm includes Foster *et al.* (2016), Giroud and Mueller (2017), Berton *et al.* (2018) and Faia and Pezone (2024).
 - ⁸ See Lucifora and Vignani (2021) and Garnero and Lucifora (2022) for a discussion on these less representative agreements.
 - ⁹ Italian wage inequality has also been analysed recently by Franzini and Raitano (2019) and by Hoffmann *et al.* (2020), while historical evidence on the effects of wage compression induced by collective bargaining has been re-evaluated recently by Leonardi *et al.* (2019).
 - ¹⁰ Contractual wages observable in our sample are the same as are available to INPS labour inspectors (*applicativo 'Vela'*) and to bookkeepers (*Il Sole 24 Ore* archive). We have hand-collected and reorganized the data from these archives, where they are available, only at the disaggregated collective contract-period of validity level. Contracts for which information on wages was unavailable tend to be the less representative ones, which often have a dubious legal validity regarding wage-setting dispositions.
 - ¹¹ The average share of the wage bill covered by the largest contract within firms was 0.93 in our sample. While cases where a firm applies only one collective contract tend to be the vast majority (the 75th percentile of the main collective contract share is 1, while the 90th percentile is 0.94), in several instances a company may also apply different contracts to some of its employees. For example, managers' wages are sometimes negotiated in separate nationwide collective agreements, while large firms may apply different contracts depending on the activities of their production units.
 - ¹² The proportion of workers employed by firms included in the sample over the population of private-sector employees is instead close to 60% in all years.
 - ¹³ FTE workers are obtained as the total days worked in a year at the firm divided by 312, the standard length of full-time contracts.
 - ¹⁴ AKM worker fixed effects represent the difference in conditional wages with respect to an arbitrary reference worker, thus expressing them as a deviation from their sample average does not represent a loss of information.
 - ¹⁵ This measure of profitability is quite standard in the literature; see, for example, Draca *et al.* (2011).
 - ¹⁶ Other firms are typically covered by collective contracts for small employers operating in the same sector, or by multi-sector collective agreements.
 - ¹⁷ In order to account for cases where the number of wage levels set within a collective contract had changed across time, we have also included in all specifications a collective contract by number of wage levels fixed effect. Results were qualitatively similar when excluding this further set of fixed effects.
 - ¹⁸ In equation (1), the symbol $*$ denotes an interaction operator, so that all time effects included in the model are specific for each sector in each geographic location.
 - ¹⁹ For example, professional counsellors are typically hired under the trade collective agreement, but they are classified in the support service activities and not in the trade sector. Similarly, airline employees are covered by different agreements depending on whether a carrier is Italian or foreign.
 - ²⁰ Given that the identity of incumbent workers changes in each period, adopting a standard fixed effects approach is not feasible for this outcome.
 - ²¹ The share of fixed-term employees and the average age of the workforce are two further outcomes that we have considered in complementary analyses.
 - ²² Matano *et al.* (2023) document that Italian contractual wages were slightly negatively affected by sector-wide import penetration during the late 1990s and early 2000s, but the size of these adjustments was quantitatively very small.
 - ²³ In a robustness exercise, results were qualitatively similar when excluding the largest firm of each collective contract. This suggests that the potential endogenous influence of the most powerful firm of the sector on wage negotiations is not particularly relevant in driving the results.
 - ²⁴ See MaCurdy (2015) for a discussion on output price adjustments in the context of minimum wage hikes, and on their welfare effects.
 - ²⁵ As Table 1 notes, around 3.4% of the firms switched collective contract in the subsequent year when considering our sample of analysis.
 - ²⁶ Lags of w_{jt}^c would capture long-run adjustments, since they measure the contemporaneous effect of a wage level applied in the past. In our static specification, these lags can be conceptualized as relevant omitted variables that are positively correlated with the treatment, given that w_{jt}^c is a persistent time series.
 - ²⁷ We have alternatively produced standard errors clustered at the firm level, and the statistical significance of the results was largely unaffected.
 - ²⁸ Since, according to Table 3, there are also negative employment effects, an implication of this result is that higher contractual wages have a negative influence on value-added.
 - ²⁹ Sorkin (2015) emphasizes the importance of inflation in affecting the size of employment adjustments in the context of the minimum wage.
 - ³⁰ In some instances, the wage cushion could even be reduced if it is made up of pay premiums that are flexible enough.

- ³¹ A second, potentially related dimension to explore could be the degree of monopsony power held by firms. However, a similar analysis would be out of the scope of the present analysis, and is better suited for future research that may employ more specific estimation methods.
- ³² Furthermore, the size of the estimated marginal effects of contractual wage growth on productivity suggests that this form of mechanical selection is not likely to be an issue. The elasticity of contractual wages to productivity is generally estimated at a level lower than 1 in absolute value, and the average growth of contractual wages per year is slightly above 2%. This suggests that, for example, a positive productivity effect potentially induced by a typical wage shock faced by firms should rarely exceed 1–2% per year. On the other hand, Appendix Table A1 shows that the within-firm growth in log value-added per worker has standard deviation 0.3, suggesting that productivity shocks of an order of magnitude of 30% per year (which can depend on a myriad of firm-specific factors unrelated to contractual wages) are not uncommon in the sample. Thus our classification is more likely to depend on underlying characteristics of firms, such as their efficiency and positioning in the product market, rather than simply on the size of their marginal effect to contractual wage growth.
- ³³ In Appendix Section A.2, we show that the effect on productivity for firms above the median in the pre-estimation period 2004–6 is positive and significant. This seems consistent with the hypothesis that the more demanding approach of Figure 5 could entail a loss in efficiency, with estimates that are generally closer to zero.
- ³⁴ Focusing on incumbents' pay growth allows us to better capture wage effects net of the firms' overall employment composition, which may change from year to year.
- ³⁵ A second potential mechanism driving the negative productivity effects, which we discuss in more detail below, could be a suboptimal choice of the inputs that are reduced due to frictions and firing costs.
- ³⁶ On the other hand, positive revenue effects in the absence of higher employment or capital could also be rationalized as a result of efficiency–wage effects or managerial slack reduction mechanisms.
- ³⁷ Lacking firm-level data on output prices, this can be considered a more direct feasible test for the role of firm-level heterogeneity in price-setting power.
- ³⁸ From an aggregate perspective, capital misallocation tends to be particularly relevant in the Italian case (Cardullo *et al.* 2015; Calligaris *et al.* 2018).
- ³⁹ Appendix Table A1 shows that on average, there was significant growth in the average age of the workers, and a negative but statistically insignificant effect on the share of fixed-term contracts in response to higher contractual wages. Note that growth in the average share of fixed-term contracts does not necessarily imply growth in fixed-term employment, as this also depends on whether fixed-term intensive firms adjust their employment differently from other companies. In this respect, in a complementary analysis of the employment effects of collective bargaining, Fanfani (2023) shows that employment losses associated with this policy are stronger among young and fixed-term workers.
- ⁴⁰ The relationship between employment selection across age groups and institutional mechanisms has often been emphasized in the Italian context with reference to tax credits and firing costs (e.g. Ardito *et al.* 2023) or pension rules (e.g. Bianchi *et al.* 2023). Evidence on the influence of collective bargaining is much less abundant (Fanfani 2023).
- ⁴¹ The consequences of near perfect multicollinearity are quite difficult to predict *ex ante*; see Spanos and McGuirk (2002) and Hill and Adkins (2003).
- ⁴² See Pinoli (2010) and Sorkin (2015) for a theoretical discussion on this hypothesis.
- ⁴³ See Section 6 and Appendix Section A.2 for more details on the assignment of firms into time-constant productivity groups.
- ⁴⁴ A similar argument has been discussed by Baker *et al.* (1999), among others.
- ⁴⁵ Intuitively, the marginal effect of an omitted 4-year lag of the treatment variable would always be absorbed by the contract by period fixed effect when using a 3-year time window p .
- ⁴⁶ Notice that the difference in the parameter β estimated using 5-year and 3-year time window fixed effects cannot be interpreted as a policy effect taking place between 3 and 5 years after contractual wage growth. Indeed, when using, for example, a 5-year time window fixed effect, an omitted 5-year lag may influence the static coefficient only when the contractual wage growth events occur in the first year of this time window. Subsequent contractual wage growth events can be affected only by closer omitted lags conditional on the other controls of the model.
- ⁴⁷ In a frictionless environment, employers can instead wait until the new wage schedule comes into effect and immediately adjust to the new optimal levels of inputs and production.
- ⁴⁸ The recent methodological literature (see Kline *et al.* 2020) has shown that the sorting component of this decomposition may be downward biased. Since the focus of this study is on first moments of the AKM parameters, which are used to derive a measure of employees' earning ability, we have not implemented a correction procedure for this problem.
- ⁴⁹ See Card *et al.* (2016) for a discussion on normalization issues concerning firm and worker fixed effects in the context of AKM regression models.
- ⁵⁰ The relevance of a firm's closure has been considered by several studies analysing the impact of minimum wages, e.g. Draca *et al.* (2011), Luca and Luca (2019), and Alexandre *et al.* (2022).

REFERENCES

- Aaronson, D. and French, E. (2007). Product market evidence on the employment effects of the minimum wage. *Journal of Labor Economics*, **25**, 167–200.
- Abowd, J., Kramarz, F. and Margolis, D. (1999). High wage workers and high wage firms. *Econometrica*, **67**, 251–333.
- Acemoglu, D. (2002). Directed technical change. *Review of Economic Studies*, **69**, 781–809.
- (2003). Cross-country inequality trends. *Economic Journal*, **113**, F121–F149.
- Adamopoulou, E. and Villanueva, E. (2022). Wage determination and the bite of collective contracts in Italy and Spain. *Labour Economics*, **76**, 102147.
- Addison, J. T. and Hirsch, B. T. (1989). Union effects on productivity, profits, and growth: has the long run arrived? *Journal of Labor Economics*, **7**, 72–105.
- Alexandre, F., Bação, P., Cerejeira, J., Costa, H. and Portela, M. (2022). Minimum wage and financially distressed firms: another one bites the dust. *Labour Economics*, **74**, 102088.
- Ardito, C., Berton, F. and Pacelli, L. (2023). Combined and distributional effects of EPL reduction and hiring incentives: an assessment using non-linear DiD. *Journal of Economic Inequality*, **21**, 925–54.
- Baker, M., Benjamin, D. and Stanger, S. (1999). The highs and lows of the minimum wage effect: a time-series cross-section study of the Canadian law. *Journal of Labor Economics*, **17**, 318–50.
- Barth, E., Bryson, A. and Dale-Olsen, H. (2020). Union density effects on productivity and wages. *Economic Journal*, **130**, 1898–936.
- , Moene, K. O. and Willumsen, F. (2014). The Scandinavian model: an interpretation. *Journal of Public Economics*, **117**, 60–72.
- Bartolucci, C., Devicienti, F. and Monzón, I. (2018). Identifying sorting in practice. *American Economic Journal: Applied Economics*, **10**, 408–38.
- Belloc, M., Naticchioni, P. and Vittori, C. (2023). Urban wage premia, cost of living, and collective bargaining. *Journal of Economic Geography*, **23**, 25–50.
- Berton, F., Mocetti, S., Presbitero, A. F. and Richiardi, M. (2018). Banks, firms, and jobs. *Review of Financial Studies*, **31**, 2113–56.
- Bianchi, N., Bovini, G., Li, J., Paradisi, M. and Powell, M. (2023). Career spillovers in internal labor markets. *Review of Economic Studies*, **90**, 1800–31.
- Boeri, T., Ichino, A., Moretti, E. and Posch, J. (2021). Wage equalization and regional misallocation: evidence from Italian and German provinces. *Journal of the European Economic Association*, **19**, 3249–92.
- Cahuc, P., Carcillo, S. and Le Barbanchon, T. (2018). The effectiveness of hiring credits. *Review of Economic Studies*, **86**, 593–626.
- Calligaris, S., Del Gatto, M., Hassan, F., Ottaviano, G. I. and Schivardi, F. (2018). The productivity puzzle and misallocation: an Italian perspective. *Economic Policy*, **33**, 635–84.
- Cappellari, L., Dell’Arlinga, C. and Leonardi, M. (2012). Temporary employment, job flows and productivity: a tale of two reforms. *Economic Journal*, **122**, F188–F215.
- Card, D. and Cardoso, A. R. (2022). Wage flexibility under sectoral bargaining. *Journal of the European Economic Association*, **20**, 2013–61.
- , ———, Heining, J. and Kline, P. (2018). Firms and labor market inequality: evidence and some theory. *Journal of Labor Economics*, **36**, S13–S70.
- , ——— and Kline, P. (2016). Bargaining, sorting, and the gender wage gap: quantifying the impact of firms on the relative pay of women. *Quarterly Journal of Economics*, **131**(2), 633–86.
- Cardullo, G., Conti, M. and Sulis, G. (2015). Sunk capital, unions and the hold-up problem: theory and evidence from cross-country sectoral data. *European Economic Review*, **76**, 253–74.
- Cengiz, D., Dube, A., Lindner, A. and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs: evidence from the United States using a bunching estimator. *Quarterly Journal of Economics*, **134**, 1405–54.
- Clemens, J. (2021). How do firms respond to minimum wage increases? Understanding the relevance of non-employment margins. *Journal of Economic Perspectives*, **35**, 51–72.
- , Kahn, L. B. and Meer, J. (2021). Dropouts need not apply? The minimum wage and skill upgrading. *Journal of Labor Economics*, **39**, S107–S149.
- and Wither, M. (2019). The minimum wage and the Great Recession: evidence of effects on the employment and income trajectories of low-skilled workers. *Journal of Public Economics*, **170**, 53–67.
- Coviello, D., Deserranno, E. and Persico, N. (2022). Minimum wage and individual worker productivity: evidence from a large US retailer. *Journal of Political Economy*, **130**, 2315–60.
- Daruich, D., Di Addario, S. and Saggio, R. (2023). The effects of partial employment protection reforms: evidence from Italy. *Review of Economic Studies*, **90**, 2880–942.
- De Loecker, J., Eeckhout, J. and Unger, G. (2020). The rise of market power and the macroeconomic implications. *Quarterly Journal of Economics*, **135**, 561–644.

- Devicienti, F., Fanfani, B. and Maida, A. (2019). Collective bargaining and the evolution of wage inequality in Italy. *British Journal of Industrial Relations*, **57**, 377–407.
- , Maida, A. and Pacelli, L. (2008). The resurrection of the Italian wage curve. *Economics Letters*, **98**, 335–41.
- Draca, M., Machin, S. and Van Reenen, J. (2011). Minimum wages and firm profitability. *American Economic Journal: Applied Economics*, **3**, 129–59.
- Dustmann, C., Fitzenberger, B., Schonberg, U. and Spitz-Oener, A. (2014). From sick man of Europe to economic superstar: Germany's resurgent economy. *Journal of Economic Perspectives*, **28**, 167–88.
- , Lindner, A., Schönberg, U., Umkehrer, M. and Vom Berge, P. (2022). Reallocation effects of the minimum wage. *Quarterly Journal of Economics*, **137**, 267–328.
- Erickson, C. and Ichino, A. (1995). Wage differentials in Italy: market forces, institutions, and inflation. In R. B. Freeman and L. F. Katz (eds), *Differences and Changes in Wage Structures*. Chicago, IL: University of Chicago Press, pp. 265–306.
- Faia, E. and Pezone, V. (2024). The cost of wage rigidity. *Review of Economic Studies*, **91**, 301–39.
- Fanfani, B. (2023). The employment effects of collective wage bargaining. *Journal of Public Economics*, **227**, 105006.
- , Lucifora, C. and Vigani, D. (2024). Employer associations in Italy: trends and economic outcomes. *British Journal of Industrial Relations*, **62**, 206–32.
- Foster, L., Grim, C. and Haltiwanger, J. (2016). Reallocation in the Great Recession: cleansing or not? *Journal of Labor Economics*, **34**, S293–S331.
- Franzini, M. and Raitano, M. (2019). Earnings inequality and workers' skills in Italy. *Structural Change and Economic Dynamics*, **51**, 215–24.
- Garnero, A. and Lucifora, C. (2022). Turning a 'blind eye'? Compliance with minimum wage standards and employment. *Economica*, **89**, 884–907.
- Giroud, X. and Mueller, H. M. (2017). Firm leverage, consumer demand and employment losses during the Great Recession. *Quarterly Journal of Economics*, **132**, 271–316.
- Hamermesh, D. (1993). *Labor Demand*. Princeton, NJ: Princeton University Press.
- Haraszosi, P. and Lindner, A. (2019). Who pays for the minimum wage? *American Economic Review*, **109**, 2693–727.
- Haucap, J., Pauly, U. and Wey, C. (2001). Collective wage setting when wages are generally binding: an antitrust perspective. *International Review of Law and Economics*, **21**, 287–307.
- Hill, R. C. and Adkins, L. C. (2003). Collinearity. In B. H. Baltagi (ed.), *A Companion to Theoretical Econometrics*. Oxford: Blackwell, pp. 256–78.
- Hoffmann, F., Lee, D. S. and Lemieux, T. (2020). Growing income inequality in the United States and other advanced economies. *Journal of Economic Perspectives*, **34**, 52–78.
- Horton, J. J. (2017). Price floors and employer preferences: evidence from a minimum wage experiment. CESifo Working Paper no. 6548.
- Kline, P., Saggio, R. and Solvsten, M. (2020). Leave-out estimation of variance components. *Econometrica*, **88**(5), 1859–98.
- Leonardi, M., Pellizzari, M. and Tabasso, D. (2019). Wage compression within the firm: evidence from an indexation scheme. *Economic Journal*, **129**, 3256–91.
- Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *Review of Economic Studies*, **70**, 317–41.
- Luca, D. L. and Luca, M. (2019). Survival of the fittest: the impact of the minimum wage on firm exit. NBER Working Paper no. 25806.
- Lucifora, C. and Vigani, D. (2021). Losing control? Unions' representativeness, 'pirate' collective agreements and wages. *Industrial Relations: A Journal of Economy and Society*, **60**, 188–218.
- MaCurdy, T. (2015). How effective is the minimum wage at supporting the poor? *Journal of Political Economy*, **123**, 497–545.
- and Pencavel, J. H. (1986). Testing between competing models of wage and employment determination in unionized markets. *Journal of Political Economy*, **94**, S3–S39.
- Manacorda, M. (2004). Can the Scala Mobile explain the fall and rise of earnings inequality in Italy? A semiparametric analysis, 1977–1993. *Journal of Labor Economics*, **22**, 585–614.
- and Petrongolo, B. (2006). Regional mismatch and unemployment: theory and evidence from Italy, 1977–1998. *Journal of Population Economics*, **19**, 137–62.
- Manning, A. (2021). The elusive employment effect of the minimum wage. *Journal of Economic Perspective*, **35**, 3–26.
- Matano, A., Naticchioni, P. and Vona, F. (2023). The institutional adjustment margin to import competition: evidence from the Italian collective bargaining system. *Oxford Economic Papers*, **75**, 631–51.
- Mayneris, F., Poncet, S. and Zhang, T. (2018). Improving or disappearing: firm-level adjustments to minimum wages in China. *Journal of Development Economics*, **135**, 20–42.

- Meer, J. and West, J. (2016). Effects of the minimum wage on employment dynamics. *Journal of Human Resources*, **51**, 500–22.
- Moene, K. O. and Wallerstein, M. (1997). Pay inequality. *Journal of Labor Economics*, **15**, 403–30.
- OECD (2017). Collective bargaining in a changing world of work. In *OECD Employment Outlook 2017*. Paris: OECD, pp. 125–86.
- Osootimehin, S. and Pappadá, F. (2016). Credit frictions and the cleansing effect of recessions. *Economic Journal*, **127**, 1153–87.
- Pagano, M. and Pica, G. (2012). Finance and employment. *Economic Policy*, **27**, 5–55.
- Pinoli, S. (2010). Rational expectations and the puzzling no-effect of the minimum wage. IZA Discussion Paper no. 4933.
- Riley, R. and Bondibene, C. R. (2017). Raising the standard: minimum wages and firm productivity. *Labour Economics*, **44**, 27–50.
- Salop, S. C. and Scheffman, D. T. (1983). Raising rivals' costs. *American Economic Review*, **73**, 267–71.
- Schivardi, F. and Schmitz, T. (2020). The IT revolution and southern Europe's two lost decades. *Journal of the European Economic Association*, **18**, 2441–86.
- Sorkin, I. (2015). Are there long-run effects of the minimum wage? *Review of Economic Dynamics*, **18**, 306–33.
- Spanos, A. and McGuirk, A. (2002). The problem of near-multicollinearity revisited: erratic vs systematic volatility. *Journal of Econometrics*, **108**, 365–93.
- Williamson, O. E. (1968). Wage rates as a barrier to entry: the Pennington case in perspective. *Quarterly Journal of Economics*, **82**, 85–116.

How to cite this article: Devicienti, F. and Fanfani, B. (2025). Firms' margins of adjustment to wage growth: the case of Italian collective bargaining. *Economica*, **92**(365), 107–149. <https://doi.org/10.1111/ecca.12558>

APPENDIX

A.1 Other figures and tables

In Tables A1 and A2, and Figures A1 and A2, this section provides further evidence on the sample of analysis and on supplementary regression models that have been estimated to check the overall robustness of our findings. A more detailed description of all evidence reported here is provided in the main text.

A.2 Further evidence on the heterogeneity analysis across the productivity distribution

In this subsection, we discuss further evidence related to the heterogeneity analysis presented in Figures 4 and 5. As discussed in Section 6, Figure 4 is based on an estimation approach where a firm is assigned to the most common quartile of its year-specific and collective contract-specific productivity distribution over the period 2006–15. For brevity, we refer to this approach as the *most common quartile classification*. Instead, in the approach of Figure 5, a firm is assigned to its most common quartile of the year-specific and collective contract-specific productivity distribution in the years 2004–6. The model is then estimated in the years 2007–15 using only firms for which a classification based on the years 2004–6 could be recovered. For brevity, we refer to this approach as the *pre-estimation classification*.

Both approaches could be vulnerable to potential drawbacks. However, in Section 6 we considered the *most common quartile classification* as our baseline reference for a series of reasons that are worth discussing in more detail. First, in our setting there is no clearly defined 'pre-treatment' or 'pre-reform' period, as in the collective bargaining system that we are analysing, industry-wide wage negotiations have been in place even in the years prior to our analysis. Thus classifying firms in some pre-estimation years does not provide a condition where firms are not subject to the potential effects of contractual wage growth. Second, our baseline approach is based on

TABLE A1 Growth rates of treatment and outcome variables (2007–15).

Variables' growth rates	Mean	S.D.	Available observations
Log median contractual wage	0.024	0.014	2,650,312
Log FTE employees	0.065	0.462	2,737,630
Log firms' average wages	0.025	0.121	2,737,559
Log incumbents' wages	0.027	0.087	2,693,455
Firms' average AKM worker fixed effects	-0.004	0.079	2,737,630
Log value-added per worker	-0.018	0.324	2,475,742
Log TFP	-0.007	0.391	2,421,887
Log revenues	0.054	0.463	2,654,520
Log profit margin	-0.031	0.603	1,939,900
Log physical capital per worker	-0.026	0.543	2,538,694

Notes: Growth rates are computed as log differences of each variable between consecutive years within firms. Averages and standard deviations are computed weighting by firm size. The first year of the sample (2006) is omitted.

TABLE A2 Effect of contractual wages on average workforce age and fixed-term share.

Outcomes	Dependent variable: Log median contractual wage				Observations (millions)
	Coeff.	S.E.	Adjusted R^2	RMSE	
Workers' average age	0.827*	0.370	0.916	1.473	3.257
Fixed-term share	-0.030	0.023	0.830	0.107	3.257

Notes: Results obtained by estimating the regression model of equation (1) on several firms' outcomes. All regressions are weighted by the number of workers in the firm. Standard errors are clustered at the collective contract level. The number of observations in each model is computed excluding singleton groups, i.e. units that are perfectly identified by the fixed effects included in the regression.

*, ** indicate significance levels 5%, 1%, respectively.

grouping firms using their most common quartile. In principle, this approach does not depend on whether the unconditional trend of firms' productivity is increasing or decreasing, but rather on which is their most common level of productivity across years. Moreover, in the heterogeneity analysis, the productive firms' counterfactual group is composed of other firms belonging to the same productivity quartile, for which the same sample selection mechanisms, if any, should apply.

Regarding the risk of selection depending on the size of marginal effects to contractual wage growth, it should be noted that firms are clearly hit by a myriad of productivity shocks. Indeed, value-added per worker at the firm level is quite volatile in most firm-level data (also reflecting measurement error specific to accounting data, and partly related to accounting practices). In our data, the within-firm log value-added per worker growth per year is close to zero on average, but its standard deviation is as high as 0.3 according to Table A1. This suggests that productivity 'shocks' of an order of magnitude of around 30% from year to year are not uncommon in the sample. With a recorded increase in collectively bargained wages of around 2% per year, and an estimated elasticity of productivity to bargained wages of around 0.5 for most productive firms in our baseline estimates, the (positive) wage-bargaining induced effect on productivity would be of only around 1% per year with respect to a counterfactual with zero contractual wage growth. This

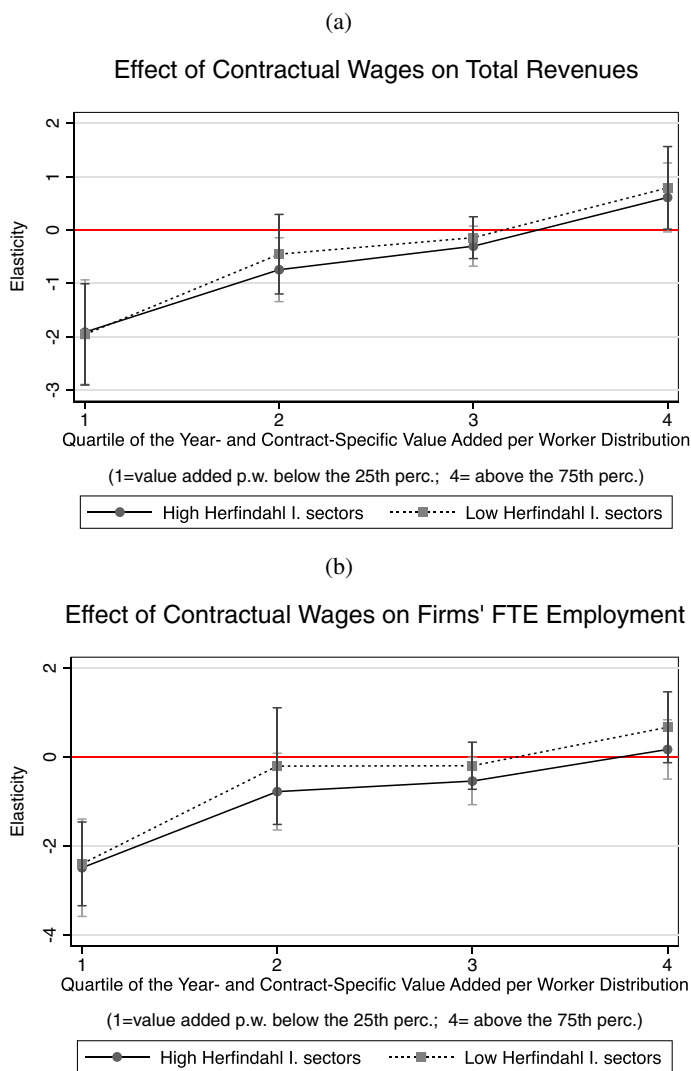


FIGURE A1 Heterogeneity of wage growth effects across the contract-specific distribution of value-added per worker, split between more and less concentrated three-digit industries based on the Herfindahl index. *Notes:* Revenue and employment effects of contractual wage growth. The treatment effect is interacted by quartiles of the contract-specific productivity distribution and by a dummy variable for high Herfindahl index three-digit industries (those above the median Herfindahl index level).

may cast doubts on the hypothesis that sorting of firms into each productivity quartile induced by positive (or negative) productivity effects of wage bargaining is a major concern in practice for our classification.

There are also some drawbacks that are specific to the *pre-estimation classification* only. In particular, the *pre-estimation classification* entails a loss of observations of around 60%, which is non-random but rather skewed towards the largest and most productive firms. Indeed, Table A3 shows that around 60% of the observations are lost due to sample attrition, as the *pre-estimation classification* cannot be defined for firms that are not observed in the years 2004–6. Depending on the outcome considered, the loss of observations for each regression model given by equation (2) is between 50% and 60%.

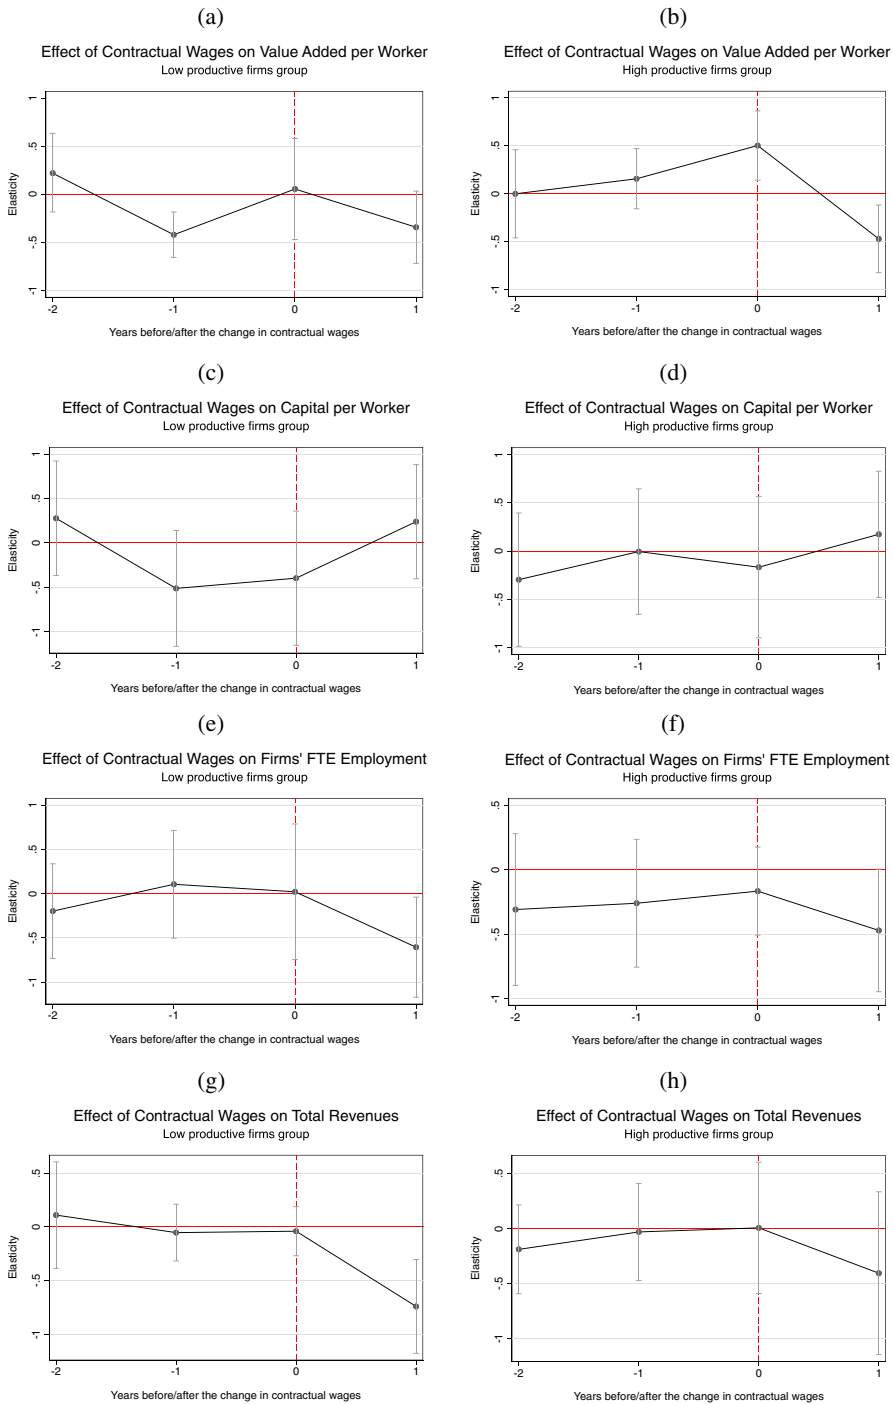


FIGURE A2 Long-run and anticipatory effects of wage growth by above and below the contract-specific median of productivity. *Notes:* Estimates of model (3) by firms classified as above or below the median of the year- and contract-specific productivity distribution. A firm is classified below the median if this is its most common condition across all years of observation.

TABLE A3 Summary of sample size differences between the pre-estimation and most common quartile approaches.

Total observations	3,515,332 (100%)
Observations in overlapping sample	1,331,830 (37.9%)
<i>Loss of observations by outcome</i>	
Incumbents' wage growth	49.1%
Value-added per worker	59.0%
TFP	57.8%
Employment	58.8%
Revenue	59.5%
Profit margin	58.5%
Capital per worker	58.3%

Notes: Overlapping sample refers to the sample where both the most common quartile and the pre-estimation quartile classification could be defined. The percentage loss of observations refers to the difference between the non-singleton sample sizes using the pre-estimation and most common quartile for each outcome in the regression equation (2).

TABLE A4 Summary of differences between the pre-estimation and most common quartile classifications in the overlapping sample.

	Observations	
	Number	%
Observations in overlapping sample	1,331,830	100%
Observations belonging to same quartile	799,379	60.02%
Observations classified differently	532,451	39.98%
Pre-estimation quartile > Most common quartile	151,311	11.36%
Pre-estimation quartile < Most common quartile	381,140	28.62%
Observations in first most common quartile	172,304	12.94%
Observations in first pre-estimation quartile	294,366	22.10%
Correlation pre-estimation and most common quartiles	0.74	

Notes: Overlapping sample refers to the sample where both the most common quartile and the pre-estimation quartile classification could be defined. All percentages are expressed with respect to the size of the overlapping sample.

In Table A4, we further characterize differences across the classifications based on the two approaches, considering only the overlapping sample where both can be defined. As can be noticed, there is a strong positive correlation of 0.74 between the two quartile definitions. However, there are also some relevant differences. In general, almost 40% of the observations are classified in a different quartile, depending on the approach taken. Moreover, it is more likely that a firm is assigned to a lower quartile according to the *pre-estimation* approach with respect to the *most common quartile* approach. This is also reflected in the size of the first quartile, which is much smaller when defined using the *most common quartile* approach.

This evidence suggests that a disproportionate amount of firms that would be typically classified as low-productivity ones in the full sample are actually missing when considering the overlapping sample. Thus low quartiles tend to be underpopulated in the overlapping sample when using the *most common quartile* approach. On the other hand, this latter approach allows us to estimate the model on the full sample of firms, which is more representative of the underlying population, particularly regarding relatively less productive companies that tend to have a higher attrition rate.

In the remainder of this subsection, we present a robustness analysis based on alternative specifications of the *pre-estimation classification* approach. In a first robustness exercise, we classify firms into two groups based only on their most common position with respect to the year-specific and collective contract-specific median during the years 2004–6. We then estimate the regression specification of equation (2) using these two groups instead of the usual quartiles. This model allows us to test for the presence of differences in adjustment behaviour between relatively more and less productive firms using a less demanding and more parsimonious specification.

Results derived from this approach are provided by Figure A3. As can be seen, results show a qualitatively similar pattern to the evidence provided by Figures 4 and 5. There are negative effects of contractual wage growth on productivity, employment, revenue and profits that are concentrated among less productive firms. Wage effects on incumbent workers are positive and similar in size for both more and less productive firms. Capital intensity, however, is not affected by contractual wage growth.

It is interesting to notice that there seem to be positive and significant productivity effects on more efficient firms, a result that is more consistent with Figure 4 than with the pre-estimation classification approach of Figure 5. This may be the result of the improved efficiency in the estimation provided by the more parsimonious approach adopted in Figure A3. On the other hand, employment and revenue effects are more negative for productive firms in Figure A3, if compared to the result for the fourth quartile of most productive firms of Figures 4 and 5. However, the point estimate for the third quartile was always negative in both of these figures, thus a negative point estimate in Figure A3 could be the result of having merged the third and fourth quartiles into a single group.

In Figure A4, we present a further heterogeneity analysis specification. In particular, we adopt an approach similar to that for Figure A3, by defining only two productivity groups based on a firm's most common position with respect to the median in the years 2004–6. However, this time we define the median using the year-, contract- and industry-specific productivity distribution, considering a 38-sector classification. The main approach adopted in the paper was to divide firms into quartiles within collective contracts, which tend to define relatively uniform industry groups. Thus in part, our approach already considers relative differences in productivity among firms that are likely to operate in the same product market. Moreover, the empirical specification includes time by sector and location fixed effects, restricting the identifying variation to comparisons across firms that share the same industry and geographic area. That said, it is still possible that an analysis based on a classification into productivity groups defined within both sectors and collective contracts could more carefully test whether the heterogeneity among even more similar firms in terms of product market position is still relevant and consistent with our main approach.

As can be noted, the evidence of Figure A4, based on a sector and collective contract specific classification of firms, is very similar to the corresponding approach derived from a classification based on collective contracts only, provided by Figure A3. The only relevant difference that emerges concerns the positive productivity effect of contractual wage growth among more productive firms, which is not significant at the 95% confidence level in Figure A4. Nevertheless, results are quantitatively very similar overall. Thus the choice of classifying firms considering only the contract-specific productivity distribution, instead of the contract and industry specific one, is not particularly sensitive in affecting our results.

A.3 AKM regression results

In order to build a time-constant measure of workers' quality, we have estimated an AKM regression model of the form

$$w_{ijt} = x_{ijt} \gamma + \eta_i + \psi_{j=t(i,t)} + \varepsilon_{ijt}$$

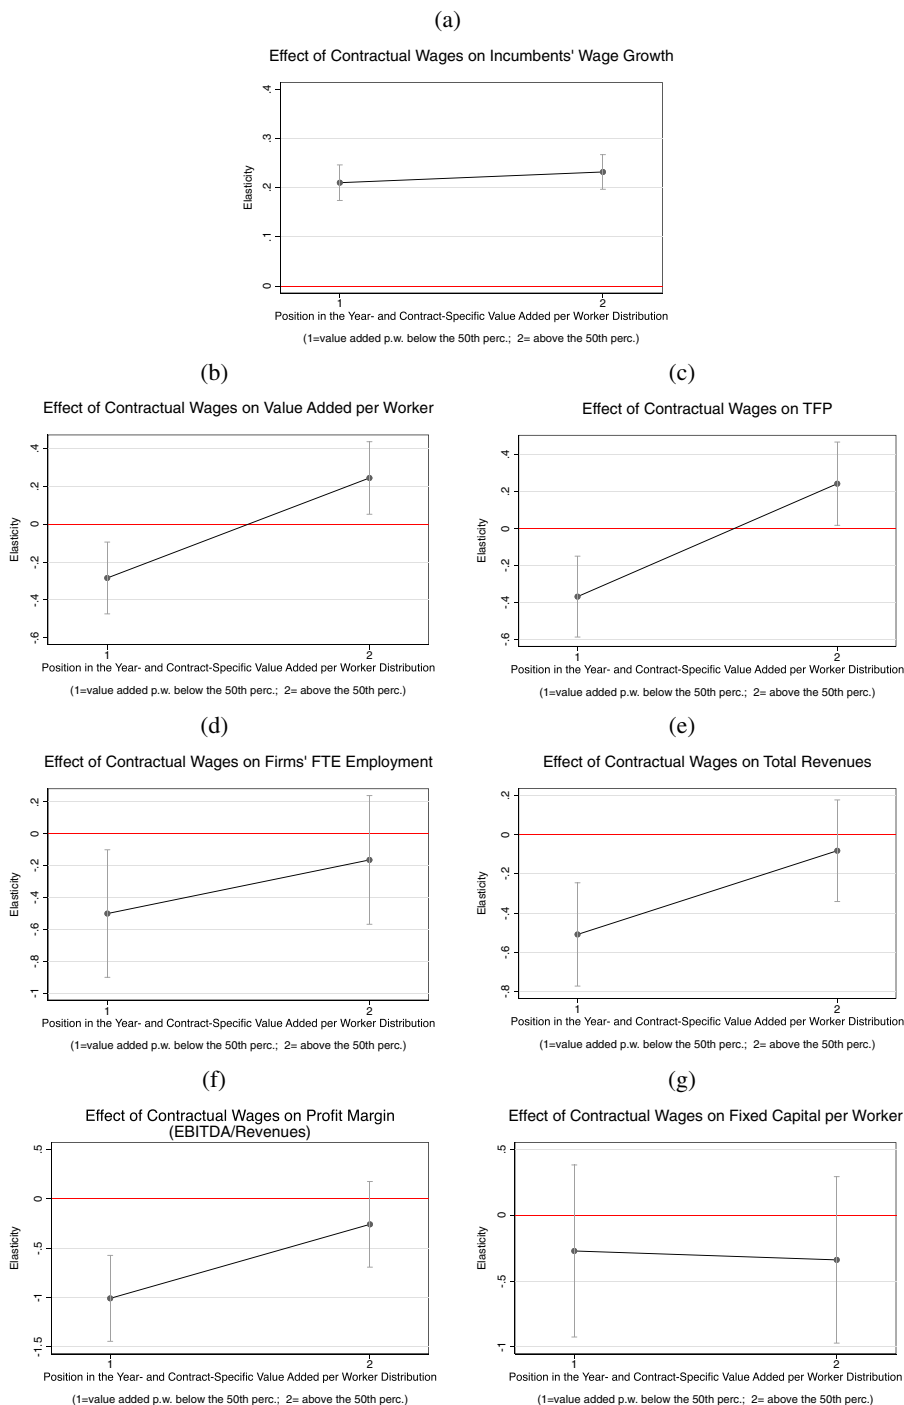


FIGURE A3 Wage growth effects across the contract-specific distribution of value-added per worker—2004–6 above/below median classification. *Notes:* Estimates of model (2), where two value-added per worker groups $q(\theta)_j$ are defined using the most common position with respect to the median to which firms belong in the period 2004–6, and estimating the model on the years 2007–15.

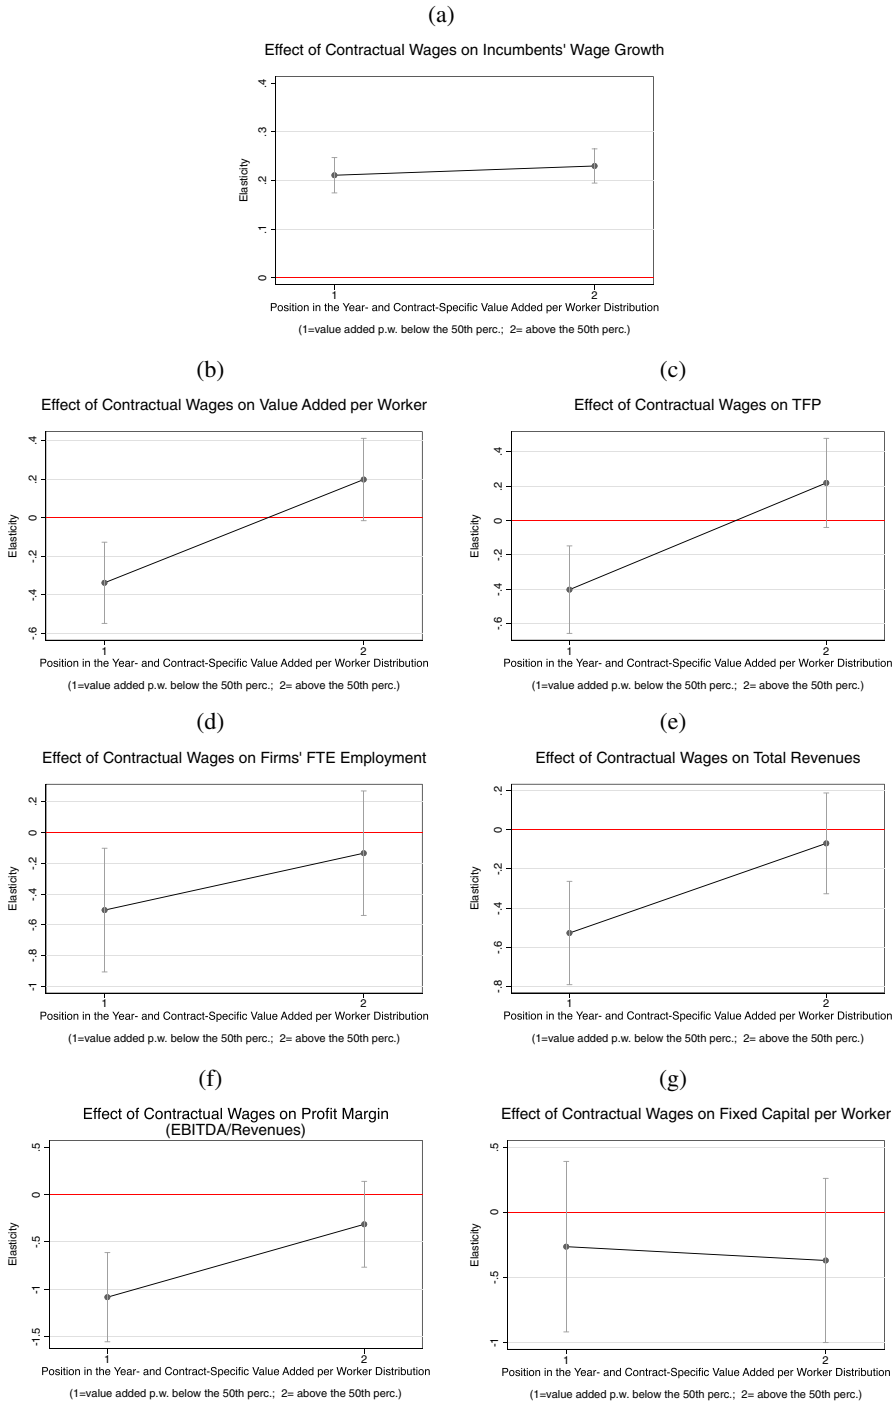


FIGURE A4 Wage growth effects across the contract- and industry-specific distribution of value-added per worker—2004–6 above/below median classification. *Notes:* Estimates of model (2), where two value-added per worker groups $q(\theta)$ are defined using the most common position with respect to the median to which firms belong in the period 2004–6 considering the year-, contract- and industry-specific distribution, and estimating the model on the years 2007–15.

TABLE A5 AKM decomposition of the wage variance.

	$\text{Var}(\phi_j)$	$\text{Var}(\eta_i)$	$\text{Var}(x_{ijt}\beta)$	$\text{Var}(\varepsilon_{ijt})$	$2C(\phi_j, \eta_i)$	$2C(\phi_j, x_{ijt}\beta)$	$2C(\eta_i, x_{ijt}\beta)$	$\text{Var}(w_{ijt})$
<i>All sample</i>								
2006–10	0.041	0.101	0.007	0.017	0.006	0.002	0.011	0.186
% of total	22.2	54.6	3.8	9.2	3.2	1.1	5.9	100
2011–15	0.048	0.104	0.007	0.016	0.001	0.002	0.008	0.187
% of total	25.8	55.9	3.8	8.6	0.5	1.1	4.3	100
<i>INPS–Cerved sample</i>								
2006–10	0.027	0.097	0.007	0.017	0.009	0.001	0.011	0.169
% of total	16	57.4	4.1	10.1	5.3	0.5	6.5	100
2011–15	0.030	0.100	0.007	0.016	0.004	0.001	0.007	0.165
% of total	18.2	60.6	4.2	9.7	2.4	0.6	4.2	100

Notes: Percentage changes for a given quantity z from z_{t-1} to z_t are computed as $100(z_t - z_{t-1})/z_t$, where $z_t = (|z_t| + |z_{t-1}|)/2$.

on two panels created using the universe of social security records (including both men and women). The two datasets cover the years 2006–10 and 2011–15, respectively. The set of controls in x_{ijt} consisted of three occupation dummies, a cubic polynomial in age interacted by sex and occupation, a part-time dummy interacted by sex, a fixed-term contract dummy, and year fixed effects. Table A5 summarizes the AKM wage variance decomposition computed in the two panels, considering both the full sample and the matched Cerved sample of incorporated businesses.

In the AKM variance decomposition results provided by Table A5, all covariances are positive. This indicates that better-paid workers are positively sorted into better-paying firms, and more endowed workers in terms of observables also tend to have higher workers' fixed effects, as one would expect. The relative contributions of firm fixed effects and of worker fixed effects to the total wage variance are always in a reasonable range, which is consistent with previous results on Italy (see, for example, Devicienti *et al.* 2019).⁴⁸

In this regression model, workers' fixed effects η_i measure an employee's earning ability controlling for non-random selection of workers across firms and on time-varying characteristics. Thus it can be considered a time-constant, comprehensive measure of workers' quality. In order to include the average level of workers' fixed effects as an outcome of our main regression model, we have first normalized these parameters across the 2006–10 and 2011–15 panels. In particular, we have defined $\bar{\eta}_i$ as the difference from the panel-specific mean of η_i , and considered for each worker the average of these normalized fixed effects $\bar{\eta}_i$ over the period 2006–15, in order to make them time-constant throughout these years.⁴⁹

A.4 Analysis on firms' exit and contract-switching behaviour

In this subsection, we analyse two outcomes that could potentially be relevant in the context of contractual wage growth. First, we consider a firm's exit from the labour market, defined as a permanent loss of all employees registered in the INPS archives. Second, we consider a firm's change in the main collective contract applied to the workforce. Both outcomes allow us to quantify the importance of alternative adjustment mechanisms available to firms. On one hand, firms could shut down production completely, outsource production, or rely on the black market when facing higher labour costs, and as a consequence they could disappear from the archives covering formal employment relationships.⁵⁰ On the other hand, they could decide to not comply with contractual wage standards, by self-selecting into less expensive collective agreements after a growth in labour costs.

TABLE A6 Effect of contractual wages on firms' closure and contract-switching—Cerved–INPS sample.

Outcomes	Treatment variable: Log median contractual wage				Observations (millions)
	Coeff.	S.E.	Adjusted R^2	RMSE	
Firm exit	0.013	0.013	0.019	0.075	2.635
Change of contract	−0.060	0.059	0.052	0.149	2.179

Notes: Results obtained by estimating the effect of contractual wages on an indicator of firms' permanent exit from INPS archives and on an indicator for firms changing the main collective contract applied to its workers. The estimation method is OLS controlling for a cubic polynomial in firms' age, log of firms' size, collective contract, and two-digit sector fixed effects, 38 industry by 107 province fixed effects interacted with year fixed effects. All regressions are weighted by the number of workers in the firm. Standard errors are clustered at the collective contract. The number of observations in each model is computed excluding singleton groups, i.e. units that are perfectly identified by the fixed effects included in the regression.

*, ** indicate significance levels 5%, 1%, respectively.

With regard to the closure of firms, the outcome variable was defined as equal to 1 if a firm had zero employees registered in the INPS archives during the following three consecutive years. Contract-switching was defined as an indicator variable for firms whose main collective contract applied to its workforce was different during the following year (including changes to collective contracts whose contractual wage was unavailable in our hand-collected database on minimum wages).

We have adopted a different specification of equation (1) when studying these two outcomes, omitting firm fixed effects and replacing them with two-digit sector fixed effects. Indeed, only exiting or contract-switching companies would otherwise contribute to the identification of the parameter of interest if we were exploiting only within-firm variation in the outcomes. The regression model that we adopted also included a cubic polynomial in firms' age in order to control for differences in the likelihood of closing down or switching contract along this dimension. Finally, year by 38 industries and 107 provinces fixed effects were also included, in order to account for general shocks in the probabilities of closing down or switching contract. We have estimated the regression model using ordinary least squares (OLS), so that the treatment effect associated to log contractual wages can be interpreted as an additive effect on the probability of closing down or switching contract.

Table A6 summarizes the results for the two regression models described above. The sample of analysis was composed of all firms included in the matched INPS–Cerved–contractual wage database. As can be seen, contractual wage growth had no significant effects on the probability of shutting down employment, or on the probability of changing the main collective contract applied to the workforce. This last result is reassuring when interpreted as a robustness test on sample selection into our database of analysis. Indeed, endogenous changes in the collective contract applied to workers do not seem to play a relevant role. Instead, the result on firms' mortality suggests that this 'hard' outcome is not relevant, at least for what concerns our sample of analysis, which consists of incorporated businesses only. Thus the employment losses associated with contractual wage growth, which we have documented, were driven mostly by generalized adjustments in the intensive margin of production, rather than by complete shutdowns of selected companies. Results were not significantly different from zero, either, when estimated across the distribution of value-added per worker, although this output has been omitted for brevity.

We have replicated the same analysis considering the full sample in the INPS archives, in order to gain a better understanding on whether focusing on the entire population of private-sector firms leads to the same conclusions. In particular, we have estimated the same regression model on the full sample of INPS records over the years 2006–15, matched with the contractual wage database used for the main analyses of the paper. The results are summarized in Table A7. As can be noted, contract-switching is not significantly affected by the growth in contractual wages.

TABLE A7 Effect of contractual wages on firms' closure and contract-switching—full INPS sample.

Outcomes	Treatment variable: Log median contractual wage				Observations (millions)
	Coeff.	S.E.	Adjusted R^2	RMSE	
Firm exit	0.073*	0.037	0.019	0.068	8.887
Change of contract	-0.016	0.041	0.065	0.143	6.910

Notes: See Table A6.

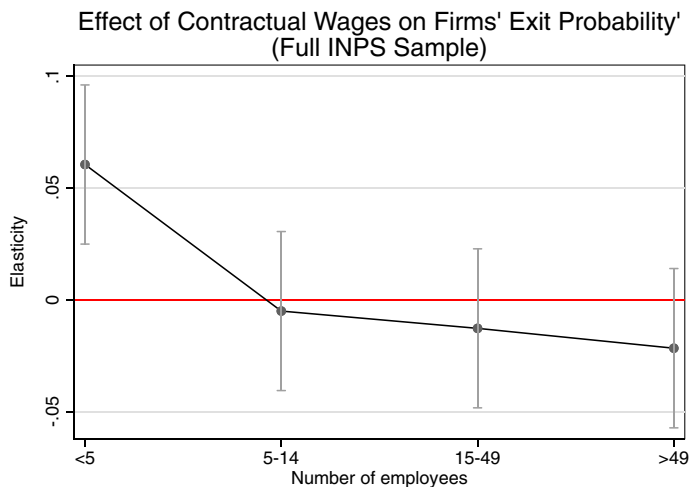


FIGURE A5 Heterogeneity of firms' closure effects across the firm size distribution—full INPS sample.

Instead, firms' exit is positively affected by contractual wage growth, with a significance level that is close to 0.05. However, the size of the coefficient is not particularly strong, as a 10% growth in contractual wages increases the probability of firms' closure by 0.7 percentage points only.

One reason why the effect of contractual wage growth on firms' mortality is marginally significant in the full INPS sample could be the inclusion of very small firms with potentially one or few employees, which were much more likely to be excluded from the matched INPS–Cerved data. For these firms, labour costs are more likely to represent a larger share of total costs, so that changes in contractual wages may trigger their closure. To test this hypothesis, Figure A5 shows the heterogeneity in the effects of contractual wage growth across the firms' size distribution. As can be noted, the positive effect on firms' mortality is driven by very small firms, those with fewer than five employees. For all other groups, the growth in contractual wages has no significant effects on their probability of exiting from the market. The size of the coefficient associated with the smallest group of firms is similar in size to the one estimated in the full sample, but the parameter is now estimated more precisely, and it is significantly different from zero.