Costs of Adjustment and Firms' Responsiveness: Evidence from a Labour Market Reform

José Pedro Garcia†

European University Institute, Italy

March, 2026

Abstract

Following the European debt crisis various countries added more flexibility to their labour markets in an attempt to boost productivity. However, in Portugal the distribution of labour productivity has remained unchanged even after a comprehensive reform effectively slashed labour adjustment costs. Why did this happen? Is productivity immune to changes in adjustment costs which constrain firms' labour demand decisions? I study this puzzling development with a dataset of Portuguese firms. While relying on a revenue function estimation exercise I identify the unobserved profitability shocks in the data and show that firms actually respond more to those exogenous innovations in the period following the reform: the responsiveness coefficient in the employment policy function nearly doubled. Importantly, I also document a significant fall in the curvature of firms' revenue function in the post-reform period. I trace this back to a more inelastic demand schedule which raises markups. Since lower costs of labour adjustment and increasing product market power have disparate effects on allocative efficiency and productivity, I develop and estimate a structural model of firm dynamics to disentangle those. A counterfactual exercise demonstrates that, although in opposite directions, reduced costs of adjustment and a falling curvature exert quantitatively similar effects on the variance of labour productivity, a measure linked to misallocation. The flat dispersion observed in the data emerges as a tension between these channels: a more flexible labour market enhances firms' employment responsiveness to shocks but the corresponding effects on allocation and productivity are offset by a concurrent increase in market power.

JEL codes: E20, E24, J20, J23, L22, L25

Keywords: Productivity, Responsiveness, Labour adjustment costs, Misallocation

†I am particularly thankful to both my supervisors, Russell Cooper and Giancarlo Corsetti, as well as staff members from the SSL division at the ECB. Helpful comments from Afonso Vaz de Castro, Chengzi Yi, Francesco Saverio, José Alberto Ferreira, Leonardo Indraccolo, Leonardo Melosi, Marius Gruenewald, Raul Santaeulalia-Llopis, Vasco Botelho, Vida Maver, and participants at various EUI working groups, ECB DG-Economics internal seminar, XVII Jornadas de Economía Laboral, and 18th Annual Meeting of the Portuguese Economic Journal are also greatly appreciated. All errors are my own.

Correspondence to: `josepedro.garcia@eui.eu`.

1 Introduction

Slowing growth and flat productivity have been ailing most European economies since the start of the century. Naturally, pressure has mounted on policy-makers to design sound and effective strategies which can reverse these trends. This paper studies the productivity effects of a comprehensive reform which reduced labour adjustment costs borne by Portuguese firms. The context is particularly interesting. In 2008, prior to the enacted changes, the labour market in Portugal featured as the most stringent across all OECD countries. Owing to significant reductions in the amount of severance paid to dismissed employees and the introduction of a new and more flexible reason to justify dismissals, amongst numerous other provisions, costs of firing were slashed. Ideally, these looser constraints would free firms to be more responsive to shocks hitting their profitability. This would then speed up the labour reallocation process from low- to more-profitable firms, and thus contribute to raise aggregate productivity in the economy.

This was arguably the ultimate goal the reform attempted to accomplish. Shortly before the first set of measures was introduced, for example, the Portuguese government itself acknowledged the need to promote *"an accelerated increase of productivity"* in an agreement signed with employers' associations and labour unions. A similar assessment would be found years later in a report commissioned to evaluate the impact of those policies: OECD (2017) posited very clearly that *"cuts in severance pay are expected to result in significant gains in both productivity and growth"*. Yet these gains have been nonexistent. Although new contracts have been under the umbrella of more flexible clauses since late 2013, neither the cross-sectional mean nor the dispersion of log labour productivity have exhibited noticeable swings (Figure 1, panel A). These also seem to have been incapable of reversing the declining pace of job flows across firms (Figure 1, panel B).

Figure 1: Recent trends in the distribution of (log) labour productivity and labour reallocation

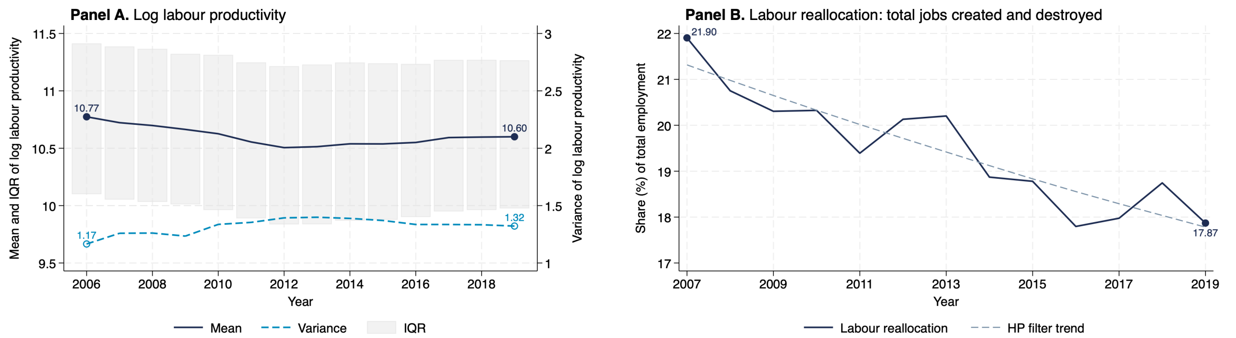

Notes: Labour productivity is measured as the ratio of a firm's yearly revenues to their stock of workers. The cross-sectional mean and interquartile range (IQR) are displayed on the left axis, the variance on the right. *Labour reallocation* is the sum of jobs created by entering and expanding firms, and jobs destroyed by exiting and contracting firms. The fitted trend relies on an HP filter with smoothing parameter 100. Data from Banco de Portugal Microdata Research Laboratory (2025).

Despite the increasing flexibility added to the labour law, the productivity distribution remained largely unchanged. **Why did this happen? Is productivity surprisingly immune to changes in adjustment costs which constrain firms' labour demand choices?** A potential answer to these questions resides in the concept of employment responsiveness: if, contrary to what had been expected, these lower costs of adjustment did not prompt firms to respond more to changes in their profitability circumstances, then the efficient labour reallocation process would be kept impaired and be a drag on productivity. I discard this hypothesis with the aid of three facts taken from a rich dataset of Portuguese firms. First, an estimated log linear approximation of the employment policy function reveals that the responsiveness coefficient nearly doubles in the post-reform period. This is consistent with additional reduced-form evidence: both the covariance and the correlation coefficient between profitability and employment rise in the period following the reform. Second, I show that this enhanced responsiveness can be found at the extensive and intensive margins of labour adjustment. I report the existence of a convex adjustment hazard function and a concave hiring rule, similar to that of Ilut et al. (2018), which becomes more linear in the presence of lower costs of adjustment. And third, I demonstrate that rising responsiveness is not matched by an increasing probability of exit when poor profitability innovations are drawn.

All these findings hinge on a revenue function estimation exercise which retrieves those unobserved profitability innovations from the data. But the usefulness of such a procedure expands beyond the identification of these exogenous shocks. Importantly, my GMM-IV estimates indicate that the period following the reform coincided with a statistically significant fall in the curvature of firms' revenue function. I trace this back to rising product market power: firms face a more inelastic demand schedule, and charge higher markups as a consequence. Since these have a detrimental impact on allocative efficiency, a falling curvature can also bring serious implications for productivity (see Baqaee and Farhi (2020) and Edmond et al. (2023)).

The existence of a nearly constant labour productivity variance merits particular attention here given its connection to the concept of misallocation, *i.e.* the dispersion of TFPR, defined in the seminal work of Hsieh and Klenow (2009). **Does that flat dispersion reflect the interaction between a more flexible labour market and changes to market structures impacting the curvature of the revenue function?** An answer to this question requires that the productivity effects of reduced costs of adjustment and those of a falling curvature be separately identified. I do so in a canonical model of firm dynamics in the tradition of Hopenhayn (1992) and Hopenhayn and Rogerson (1993), and which I estimate through simulated method of moments. In my partial equilibrium framework idiosyncratic shocks to profitability induce firms to alter their employment decisions and, in some cases, to endogenously exit the economy. Albeit simple, the model retains very useful properties. First, it allows to rationalise the magnitude of these responses by appealing to three components which can be disentangled: *i)* a piecewise formulation of convex and non-convex labour adjustment costs, *ii)* the market structure embedded in the curvature of firms' revenue function, and *iii)* the stochastic process which governs their profitability. And second, it also allows to express moments of the labour productivity distribution as (explicit and implicit) functions of these three elements.

I highlight two main results delivered by the estimated model. First, the dispersion of labour productivity declines when adjustment costs are cut but it increases when the curvature of the revenue function falls. This tension is present in the data and is explained by the model. In a counterfactual exercise I show that if the curvature had not fallen, the standard deviation of labour productivity would have dropped 14%. Conversely, if labour adjustment costs had not been slashed, that standard deviation would have actually risen 11%. Such similar magnitudes suggest the two mechanisms exert a quantitatively neutral impact on the dispersion of labour productivity. And second, although the distribution is left unchanged, decomposing its variance does reveal rich dynamics. The parameters governing the marginal revenue of labour and the stochastic profitability process play a pivotal role here: they fully determine the increase in the misallocation measure of Hsieh and Klenow (2009), and they also raise the weighted variance of employment.

Ultimately, the paper argues that lower costs of adjustment in the Portuguese economy did enhance firms' employment responsiveness, but the corresponding impact on productivity was fully offset by a concurrent increase in market power encoded in the curvature of their revenue function. Although I focus on Portugal, these conclusions can be farther extended. For example, following the European debt crisis, larger countries like Italy and Spain also added more flexibility to their national labour laws in an attempt to boost productivity and revamp their economies. Pinelli et al. (2017) and OECD (2014) provide an overview of the Italian Jobs Act of 2014-2015 and of the labour market reform undertaken by the Spanish Labour Ministry in 2012, respectively. But as widespread as these policies might have been, a complete assessment of their productivity effects is still missing. This is particularly important now that muted growth and the lack of competitiveness have retaken centre stage in European policy circles. The paper contributes to that debate by drawing insightful lessons from the Portuguese case. First, because the reform introduced flexibility to the most stringent of labour markets in the OECD. This is a non-trivial development. And second, because that reduced stringency interacted with rising market power, a documented trend in most European and North American economies (De Loecker and Eeckhout (2018)).

The paper also speaks to an ongoing debate in the US on the link between productivity, business dynamism, and firms' responsiveness. Decker et al. (2016) and Decker et al. (2020) have supported the hypothesis that the well-documented decline in the pace of job reallocation is explained by a weaker responsiveness of firms to exogenous shocks. While Cooper et al. (2024a) have subsequently claimed this is essentially the reflex of increased costs of labour adjustment, De Loecker et al. (2020) have instead pointed out that waning business dynamism is justified by changes in market structures and specifically a rise in firms' market power.

I reconcile both views. On the one hand, I show that an enhanced employment responsiveness is mostly the product of reduced labour adjustment costs. On the other, my model implies that the observed changes in market power embedded in the curvature of the revenue function *i)* fully neutralise the positive effect the labour market reform could have on the variance of labour productivity, and *ii)* together with the parameters governing firms' profitability process, explain the dynamism found in the decomposition of that variance.

The paper also contributes to a literature pioneered by Hamermesh (1989), Bentolila and Bertola (1990), and Caballero et al. (1997) which studies firms' labour demand decisions in the presence of adjustment costs, and how these impact the macroeconomy. Hopenhayn and Rogerson (1993) showed that a linear cost of job destruction can reduce average productivity by more than 2%. I rely on a labour adjustment cost function which supplements theirs. My piecewise formulation distinguishes net costs of job creation from job destruction, and includes both fixed and linear forms of adjustment. This specific coexistence of non-convex and convex adjustment costs is intended to capture prominent features of the Portuguese labour market (pervasive inaction, as documented by Varejão and Portugal (2007)) and targets of the reform (cut in severance payments). If all these costs were to be removed from the economy, the model predicts that the covariance between employment and profitability would rise by a factor of four, and the standard deviation of labour productivity would fall 35%.

Lastly, I bridge this literature with another one which assesses the effects of labour market reforms. Autor et al. (2007), Cingano et al. (2016), and Caggese et al. (2022) have all relied on purely empirical exercises to evaluate the impact of policy interventions in the US, Italy, and Belgium, respectively, which strengthen workers' rights at the expense of mobility and flexibility in the labour market. Cooper et al. (2018), on the other hand, also estimated a general equilibrium model of labour demand to explore the effects of a similar reform in China. Eventually, they all agree rising firing costs have a negative impact on labour mobility, aggregate employment, and productivity. This is not entirely the case in Portugal. While studying a reform of the opposite sign in a severely rigid labour market, I show that the distribution of labour productivity actually remained unaltered. Reduced costs of adjustment enhance firms' employment responsiveness, but the fall in the curvature of their revenue function fully neutralised the positive effect of looser constraints on the variance of labour productivity.

Road map

The paper is organised as follows. In Section 2 I describe in detail the labour market reform which was implemented in Portugal between 2011 and 2013. In Section 3 I introduce the dataset and demonstrate how the estimated fall in the curvature of firms' revenue function mirrors their increasing market power. From this estimation procedure I also retrieve the unobserved profitability innovations which are used to explore the effects of the reform in different responsiveness regressions. The partial equilibrium model of labour demand with adjustment costs is presented in Section 4 and estimated in Section 5. In Section 6 I rely on counterfactual scenarios and a decomposition exercise to quantitatively assess the impact of the various channels on the variance of labour productivity. Section 7 concludes.

2 Institutional setting: the labour market reform in Portugal

The Portuguese labour market ranked as the most stringent across OECD countries prior to the GFC and the ensuing Euro Area crisis. Over the subsequent eleven years that strictness would fall and be brought much closer to the OECD unweighted average. This striking evolution is seen in both panels of Figure 2, which plots the standard OECD composite indicator of employment protection laws.

Figure 2: OECD Employment protection legislation index in 2008 (left) and 2019 (right)

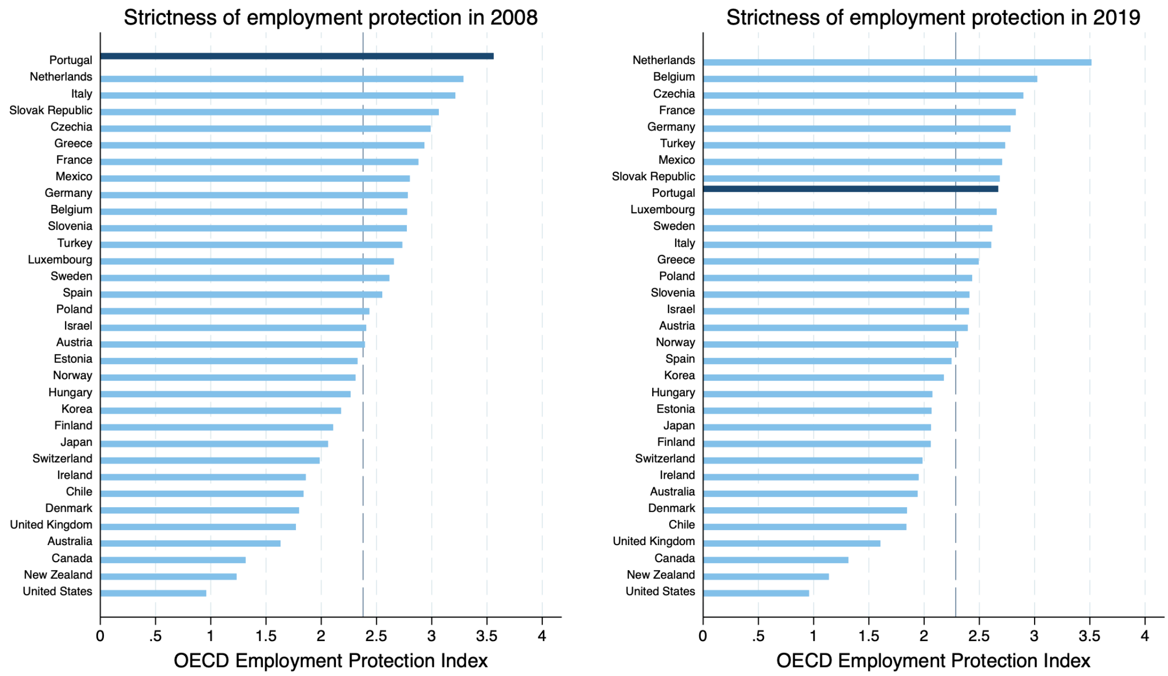

Notes: Data from OECD Employment Protection Legislation Index, version 3. Greater/smaller values of the index are understood as more/less rigid labour markets. Only OECD countries with 2008 data available were considered. Gray dashed line is the cross-section unweighted average.

What can then explain this move from first to ninth? In April 2011 the Portuguese government requested financial assistance to the European Commission, the European Central Bank, and the International Monetary Fund. The memorandum of understanding signed a month later tied a bailout package of 78 billion euros to a series of public spending cuts and significant reforms in the economy. The most important ones were arguably those pertaining to the labour market.

Indeed, over the following years lawmakers introduced multiple changes to the Portuguese labour law. A comprehensive overview of all those, as well as their intended goals, can be found in OECD (2017). Yet my emphasis here is placed on three very specific policies which were meant to reduce the costs borne by firms when letting go employees: *i*) severance payments were slashed, *ii*) firm-worker funds were set up in order to accommodate up to half of severance due in case of dismissal, and *iii*) a new and more flexible justification for fair dismissals was created. Below I go over these in detail.

Start with the severance payments cut. Initially, for every full year of tenure at the firm, the severance pay of workers on open-ended contracts (2/3 of all employment relationships) corresponded to 30 days of base wage, plus tenure-based increments. A minimum of three months' wage would have to be paid and no maximum would apply. All these conditions were eventually scrapped: lawmakers progressively reduced the amount of severance to only 12 days of base wage (from 30 to 20 in November 2011 and ultimately to just 12 in October 2013); the minimum was removed, and a cap of 240 times the national monthly minimum wage was introduced. Similar provisions were also set for fixed-term contracts, albeit with additional conditional clauses on the length of the terminated contract.

This was not the only change made to severance payments. In an attempt to prevent liquidity shortages caused by the costs borne at the time of dismissal, firms were required to constitute worker-specific funds. This provision was instituted in October 2013 and encompasses no risks for employers. Each month they make a mandatory contribution of 0.925% of each employees' wage to their respective fund. Should a worker leave the firm voluntarily, no severance is due and the firm is paid back all the contributions made so far. On the other hand, if the worker is entitled to severance, then the fund covers half of that amount.

Lastly, the reform also introduced a more flexible procedure to justify the fair dismissal of an employee who is revealed to be unsuitable for the job. Before August 2012 a worker could only be let go on the grounds of unsuitability after a lengthy series of steps had been painstakingly followed: first, the employer would have to prove that the worker's inability had been preceded by a change in the nature of the assigned tasks; conditional on those changes, the employer would then have to demonstrate that *i*) the lack of ability had not been remedied by the provision of additional and certified training, *ii*) the worker could not be transferred to a more suitable role within the company, and *iii*) the situation had not been generated by a lack of health and safety conditions for which the employer could be held accountable. The process has been significantly shortened. First, it can now start even when the nature of the tasks remains unchanged. And second, the employer is no longer required to offer an alternative and suitable position within the firm.

Two final remarks about the policies described here. First, they were part of a wider labour market reform which took place in Portugal. For example, the memorandum of understanding also encompassed minimum wage freezes, easier access to less generous unemployment benefits, and changes to the system of collective bargain agreements. The ultimate goal of these policies was clear: to boost job reallocation, and eventually productivity, through a less stringent labour market. The second one concerns their immediate effectiveness. All enacted provisions on severance payments were grandfathered in, *i.e.* they would only apply to new contracts set from then on. In the next section I show this may have mitigated the reduction in employment inaction but may have also produced an immediate effect on the responsiveness to positive profitability shocks.

3 Data facts

Lower costs of labour adjustment prompt firms to respond more to changes in their profitability. By speeding up the reallocation of workers throughout the economy, this enhanced responsiveness can have a sizeable impact on labour productivity. In this section I leverage on a rich Portuguese dataset to study the evolution of the employment responsiveness. This requires that a measure of purely exogenous innovations to firms' profitability be first retrieved from an estimated revenue function. I do so via quasi first-differences, a technique also deployed in Cooper and Haltiwanger (2006).

The annual data are provided by Banco de Portugal Microdata Research Laboratory (2025), an autonomous unit within the Bank of Portugal's Economics and Research Department, and comprise all balance sheet and income statement items of non-financial firms operating in Portugal. The dataset is based on information these firms are required to report by law. As a result, the coverage is extraordinary: the near-universe of non-financial firms is included. Table B1 displays some summary statistics.

Given the staggered implementation of the policies there is not a unique discontinuity which helps to causally disentangle their effects. For the remainder of the paper let the pre-reform period be 2006-2013, and the post-reform period 2014-2019. The data are available since 2006 and choosing 2019 as the final year avoids contamination from the covid shock. The choice of 2014 as a cut-off year is justified by two reasons. First, all policy measures described in the last section were implemented until late 2013. And second, because that coincides with the peak of employment inaction observed in the economy[1]. The 2014-2019 time window also retains a sufficiently large number of observations.

[1] A useful timeline of events is provided in OECD (2017), p.155. For the evolution of employment inaction see Figure 3. I will further elaborate on this in in Section 3.3.

3.1 Revenue function estimation

Consider an economy where heterogeneous firms operate in monopolistic competition. The *real* side of the economy is described by a Cobb-Douglas production function: $y_{i,t} = z_{i,t} \cdot (n_{i,t}h_{i,t})^{\zeta} \cdot k_{i,t}^{1-\zeta}$, where $z_{i,t}$ is the idiosyncratic productivity level of firm i in period t, $n_{i,t}$ is the stock of employees, $h_{i,t}$ is the average number of hours worked, and $k_{i,t}$ is the stock of capital. For a reason I will discuss later, this formulation assumes constant returns to scale. On the *nominal* side, say these firms face a downward sloping demand curve given by $p(y_{i,t}) = (\theta_{i,t}/y_{i,t})^{1/\nu}$, where $\theta_{i,t}$ is a stochastic demand shifter, and ν the price elasticity of demand.

Under fully flexible capital, the revenue function simplifies to $R_{i,t} = A_{i,t} \cdot (n_{i,t}h_{i,t})^{\alpha}$. In Appendix A I show how to derive this closed form function from the primitives indicated above. Two elements are worth noticing. First, the idiosyncratic profitability level, $A_{i,t}$, encompasses shocks to productivity, demand, and the cost of capital. The literature commonly refers to it as a TFPR measure. It evolves according to a stochastic AR(1) process, such that $A_{i,t} = \xi_{i,t} \cdot A_{i,t-1}^{\rho}$ with $|\rho| < 1$ and $\log(\xi_{i,t}) \overset{iid}{\sim} \mathcal{N}(0, \sigma^2)$. And second, the curvature $\alpha \in (0,1)$ embeds both the labour elasticity in the production function, $\zeta \in (0,1)$, and the price elasticity of demand, $\nu \geq 0$.

While capital and hours can be freely adjusted, the number of employees cannot. Regardless of the specification of labour adjustment costs, the employment policy function can be conceived as $n_{i,t} \equiv \varphi(n_{i,t-1}, \xi_{i,t})$. The past stock of workers is relevant to the decision because it is informative about the costs firms bear with potential adjustments. And since they also factor in their contemporaneous technological and demand conditions, the profitability innovation, $\xi_{i,t}$, is included as a second argument.

Despite their importance to firms' employment decisions, these innovations are unobserved to the researcher. But they can be retrieved from the data once the parameters which characterise the revenue function are estimated. To do so I apply logs and quasi first-difference it:

$$\log(R_{i,t}) = \rho \log(R_{i,t-1}) + \alpha\big(\log(n_{i,t}) + \log(h_{i,t})\big) - \alpha\rho\big(\log(n_{i,t-1}) + \log(h_{i,t-1})\big) + \log(\xi_{i,t}) \tag{1}$$

Estimating (1) via OLS would yield a biased estimator of α. Recall that from the policy function laid out before employment decisions are contemporaneously correlated with profitability innovations. The fact that $\log(n_{i,t}) \not\perp \log(\xi_{i,t})$ implies the existence of omitted variable bias. Instead, I rely on a set of instrumental variables to estimate (1) through GMM. The existence of adjustment costs adds some sluggishness to labour demand decisions. One-period lagged wages and effective labour usage must then be correlated with today's stock of workers. And by the definition of an i.i.d. innovation, the former must also be orthogonal to $\log(\xi_{i,t})$. Finally, a full set of year dummies controls for the effects of aggregate shocks. Table 1 shows the outcome of the estimation.

The estimated curvature of the revenue function ranges between 0.44 (post-reform) and 0.56 (pre-reform).

Table 1: GMM-IV revenue function estimation

	Full sample: 2006-2019	Pre-reform: 2006-2013	Post-reform: 2014-2019
Curvature, α	0.526^{***}	0.563^{***}	0.442^{***}
	(0.005)	(0.015)	(0.017)
Serial correlation, ρ	0.880^{***}	0.878^{***}	0.901^{***}
	(0.003)	(0.004)	(0.008)
Std deviation, σ	0.601^{***}	0.601^{***}	0.594^{***}
	(0.005)	(0.015)	(0.017)

Notes: Parameters α and ρ are estimated directly from equation (1). The corresponding standard errors are clustered at the industry level. The series of profitability innovations are then generated with the parameterised revenue function. The estimator of σ is the standard deviation of that data series. The variance of this estimator is obtained by sampling with replacement 100 000 observations 100 times.
$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

This is a statistically significant fall which can only be traced back to *i)* changes in the labour elasticity in the production function, and/or *ii)* changes the price elasticity of demand which eventually mirror varying aggregate markups[2]. Which of the two is the culprit? I answer this question below.

What explains the fall in α?

Under the hypothesis of constant returns to scale, the output elasticity of labour can be read off standard firm-level data as the ratio between the wage bill and revenues, *i.e.* the labour share [3]. Given the GMM-IV estimates of α documented in Table 1 above, and the values of the labour elasticity estimated with this factor share approach, one can ultimately infer the price elasticity of demand. This exercise is carried out in Table 2 and allows to understand which of the two parameters drove the fall in the curvature of the revenue function.

The labour elasticity averaged 0.29 in 2006-2013 and 0.32 in 2014-2019. With such negligible change most of the variation in the estimated curvature must have come from the associated markups: an increase from 22% in the pre-reform period to 40% after the reform. This proves the period following the labour market interventions coincided with rising market power exerted by firms. Just as importantly, this development in the product market can be shown to drive most of the fall in the curvature of the revenue function.

To better understand why consider the following counterfactual scenario: suppose the labour elastic-

[2]A supplementary explanation could potentially rest on a changing composition of the economy: sectors with a lower curvature could have become more prominent. This is ruled out for two reasons. First, employment shares remained broadly unchanged throughout these years. And second, because the curvature of the five largest sectors, *i.e.* those which jointly comprise 75% of all employment, exhibit a similar declining trend. See Appendix A.

[3]In addition to the assumption of constant returns to scale, this factor share approach also requires that labour be chosen frictionlessly by firms. This obviously contradicts one of the main premises of the paper: the presence of labour adjustment costs. Here I follow De Loecker and Syverson (2021) and average the ratio across producers to neutralise the bias they induce.

Table 2: Disentangled changes in the estimated curvature of the revenue function

	Pre-reform: 2006-2013	Post-reform: 2014-2019
Curvature, α	0.563	0.442
Output elasticity of labour, ζ	0.285	0.316
Implied elasticity of demand, $\nu \equiv \lvert(1-\varepsilon)^{-1}\rvert$	5.527	3.511
Implied markup, $\mu \equiv \varepsilon^{-1}$	1.221	1.398
Counterfactual α: constant labour elasticity		-5.734%
Counterfactual α: constant markup		$+33.05\%$

Notes: The curvature of the revenue function corresponds to $\alpha \equiv \frac{\zeta\varepsilon}{1-\varepsilon(1-\zeta)}$. The labour share of a firm is computed as the ratio between the wage bill (including social security contributions) and its revenues. The output elasticity of labour, ζ, is obtained by taking the unweighted mean of (1% winsorized) firm-level values and then the time-series average across the corresponding time window. The number of observations in each time window is the same as reported in Table 1.

ity would have been kept unchanged at $\zeta = 0.285$ and only the implied elasticity of demand would have changed to its observed values in the post-reform period. The resulting curvature would have been $\tilde{\alpha} = 0.413$. This is very close to its true value, $\alpha = 0.442$. On the contrary, if only the labour elasticity had been allowed to vary while the elasticity of demand would have remained the same, the counterfactual curvature would have been $\tilde{\alpha} = 0.558$, which is more than 33% greater than the observed value.

It is worth to single out two conclusions. First, I document a large fall in the curvature of the revenue function in the period following the labour market reform. And second, I show that a growing measure of the aggregate markup is the main responsible for that fall[4]. All else equal, these developments can have sizeable repercussions for productivity. While in a perfectly competitive market profitability shocks are passed one-to-one to consumers owing to a fully elastic demand, that pass-through is lower with increasingly inelastic demand and grater market power. As a consequence, the same profitability shock prompts a smaller labour adjustment, employment responsiveness dwindles, and the reallocation process falters.

Throughout the paper I work under the assumption that the observed increase in market power is orthogonal to the effects produced by the reform which cut labour adjustment costs. This is consistent with the narrative of De Loecker et al. (2021), who have recently attributed the documented rise in market power of US firms to both changes in technology and market structures. Here, I showed that the fall in the estimated curvature of the revenue function mirrors essentially a change in the structure of product markets: as

[4] I further validate these points in Appendix A and Appendix B. First, the decline in the curvature of the revenue function is robust to alternative estimation strategies. Relying on Ackerberg et al. (2015), for example, yields a similar result. And second, while resorting to the method of De Loecker and Warzynski (2012) I show that there is indeed a rightward shift in the underlying distribution of firm-level markups. Interestingly, the median of that distribution is fairly close to the aggregate markup.

demand becomes more inelastic, firms can charge higher markups.

That curvature is jointly estimated with two other parameters which fully characterise the profitability process of firms. The serial correlation of TFPR ranges between 0.88 and 0.9, whereas the standard deviation of profitability innovations hardly moves from 0.6. Although admittedly small, these changes also have serious implications. Notice these two parameters suffice to describe the variance of log TFPR, $\sigma^2/(1-\rho^2)$, which increases from 1.60 to 1.82. This is a measure of worsening allocative efficiency (see Hsieh and Klenow (2009)), concurrent to increased market power. I will later show that, when taken altogether, these three parameters (α, ρ, and σ) produce disproportionately large effects on the decomposition of the labour produtivity variance.

In conclusion, this estimation exercise allows to obtain all three parameters which fully characterise the revenue function. But the usefulness of the exercise also lies on the ability to retrieve the series of exogenous firm-specific profitability innovations, $\{\log(\xi_{i,t})\}_{t=1}^{T_i}$. In the following sub-sections I show how firms' decisions on employment adjustment (both on the extensive and intensive margins) and exit are qualitatively and quantitatively affected by these innovations.

3.2 Log-linear approximation of the employment policy function

I initially claimed that labour demand is a function of past employment and current profitability innovations: $n_{i,t} \equiv \varphi(n_{i,t-1}, \xi_{i,t})$. Consider now a log-linear approximation of this policy function:

$$\log(n_{i,t}) = \lambda_1 \log(n_{i,t-1}) + \lambda_2 \log(\xi_{i,t}) + \epsilon_{i,t} \tag{2}$$

There are two parameters of interest: while λ_1 measures the serial correlation of log employment, λ_2 gauges the responsiveness of labour demand to contemporaneous profitability innovations: $\lambda_2 = \partial \log(n_{i,t})/\partial \log(\xi_{i,t})$. Table 3 shows their estimates.

There are three results which stand out. First, the impressive quality of the log linear approximation. Roughly 96% of the variation in employment is explained with two single variables: lagged employment and current profitability. Second, the estimated value of λ_1 is exceedingly high and close to unity. This indicates employment remains an extremely persistent phenomenon all throughout the sample. Third, and more importantly, the estimated responsiveness is almost twice as large in the post-reform period (0.051) than it was before the intervention (0.028). These values are statistically different and, as expected, greater than zero.

These results are robust to alternative specifications of the DGP — see Appendix B. The third one, however, merits further consideration. Employment reacts positively to profitability innovations and that effect is almost twice as large in the period following the labour market reform. It is a result which builds on a

Table 3: Log linear approximation of employment policy function

	Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Serial correlation, λ_1	0.953^{***}	0.949^{***}	0.960^{***}
	(0.005)	(0.006)	(0.004)
Responsiveness, λ_2	0.035^{***}	0.028^{***}	0.051^{***}
	(0.008)	(0.008)	(0.010)
Industry & Year FE	Yes	Yes	Yes
R^2	0.959	0.958	0.959
Observations	3 308 659	1 709 959	1 350 617

Notes: Standard errors are clustered at the industry level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

very specific and structural view of the economy. Yet it is also mirrored in reduced-form objects. The covariance between log profitability and log employment, for example, rose from 0.64 in 2006-2013 to 0.81 in 2014-2019. For a more intuitive interpretation, the corresponding correlation coefficient had a statistically significant increase from 0.47 to 0.55. In the following sections I will place a greater emphasis on the evolution of that covariance because it will be a particularly relevant element to study the decomposition of the labour productivity variance.

Lastly, notice the dynamics of employment is the outcome of two broader decisions made by the firm: one at the *extensive* margin (adjustment/no adjustment), and another at the *intensive* margin (conditional on adjustment, size of job creation/destruction). These can be differently impacted depending on the formulation of labour adjustment costs. Take a reform which would essentially reduce a symmetric linear cost of adjustment. The extensive margin decision would remain largely unchanged as the band of inaction is mostly driven by non-convex costs. But conditional on adjusting both before and after the enacted change, the size of job growth would be impacted. As discussed in the previous section, the Portuguese reform targeted a reduction in fixed and linear costs alike. Given how important these are to decisions on extensive and intensive margins of adjustment, in the next sub-section I take a closer look at these.

3.3 Responsiveness regressions: employment adjustment decisions

Let the employment growth rate of firm i in period t be $\varkappa_{i,t} \equiv (n_{i,t} - n_{i,t-1}) / \big(0.5 \times (n_{i,t} + n_{i,t-1})\big)$. The firm is assumed to be inside the band of inaction whenever the change in the number of workers is very small, *i.e.* smaller than 2.5% in absolute value, $\varkappa_{i,t} \in [-0.025, 0.025]$. The indicator variable $\mathbb{1}_{i,t}$ tracks this: let $\mathbb{1}_{i,t} = 1$ when the firm is outside that band of inaction, and 0 when it is inside.

Figure 3: Evolution of employment adjustments within band of inaction, $\varkappa_{i,t} \in [-0.025, 0.025]$

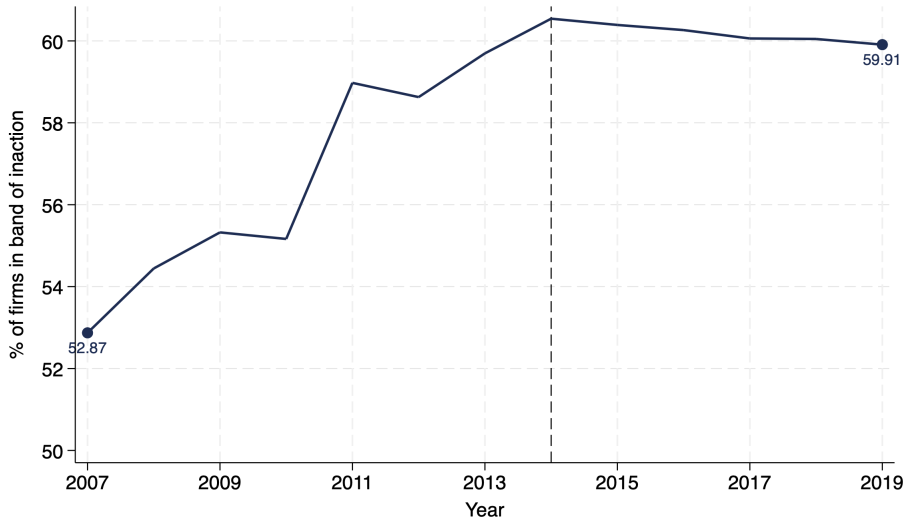

Figure 3 above shows the pervasiveness of employment inaction in the Portuguese economy. More than half of all yearly observations between 2007 and 2019 entailed employment adjustments smaller than 2.5% in absolute value [5]. This is a typical effect exerted by non-convex labour adjustment costs. Yet that inaction has been falling at a very slow pace since 2014: an average short of 0.15 percentage points each year. As mentioned before, this is likely to be caused by the fact that new provisions introduced in the labour law were grandfathered in: although new hirings are covered by more flexible rules, pre-existing contracts are not. This duality can then act as a drag on inaction. I will return to this point below.

The responsiveness regression at the extensive margin corresponds to the following:

$$\mathbb{1}_{i,t} = \delta_1 \log(\xi_{i,t}) + \delta_2 \log(\xi_{i,t})^2 + \delta_3 \log(n_{i,t-1}) + \epsilon_{i,t} \tag{3}$$

Equation (3) is a linear probability model: it posits how current profitability innovations and lagged employment determine the likelihood of an employment adjustment today. I include both linear and quadratic terms of the log profitability innovation to capture the fact that firms may only adjust in the presence of sufficiently large shocks, either positive or negative.

I use an identical formulation for responsiveness at the intensive margin. Conditional on action, the size of the employment adjustment is determined as follows:

$$\varkappa_{i,t\{\mathbb{1}_{i,t}=1\}} = \theta_0 + \theta_1 \log(\xi_{i,t}) + \theta_2 \log(\xi_{i,t})^2 + \theta_3 \log(n_{i,t-1}) + \epsilon_{i,t} \tag{4}$$

The results are reported in tables B6 and B7. I rely on those to project the left-hand side of both (3) and (4)

[5] For comparison, of the four European countries considered by Cooper et al. (2024b), Italy is the one where inaction is the greatest and only amounts to 35%.

against a very fine grid of log profitability innovations[6]. These are shown in Figure 4 below.

Figure 4: Responsiveness regressions at the extensive (left) and intensive (right) margins of employment

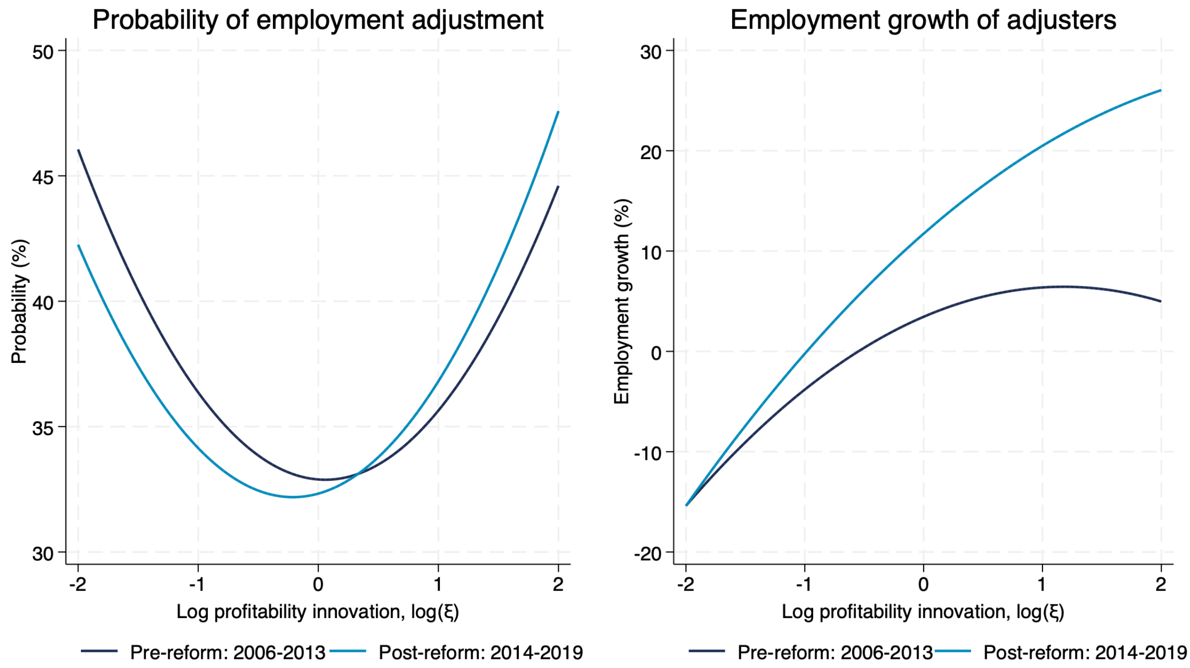

Notes: Extensive and intensive margin decisions of the average-sized firm. For the post-reform period of 2014-2019 size is kept constant at the 2006-2013 level.

The left panel on the probability of employment adjustment yields three major results. The LPM predicts a convex adjustment hazard: firms are more likely to adjust employment when hit by profitability innovations which are increasingly larger in absolute value. This is explained by the persistently positive estimates of δ_2. Second, even though those point estimates remain unchanged throughout the several time windows, those of δ_1 actually increase by 50% in the post-reform period. This joint evolution of δ_1 and δ_2 implies that the probability of adjustment responds significantly more to profitability innovations after the labour market intervention. And finally, there is an interesting asymmetry in the U-shaped curves. Conditional on drawing a *negative* profitability innovation, firms were more likely to respond to it prior to the reform; but when faced with a *positive* innovation, the likelihood of adjustment is greater in the post-reform period. This again appeals to the duality introduced by the reform. Although the new provisions were mostly aimed at firing costs, their immediate effect implied greater incentives for job creation. Quantitatively, a positive one standard deviation profitability innovation would raise the probability of adjustment by 0.9 percentage points prior to the reform; after the set of more flexible rules were put in place, an identical shock would

[6] The log profitability innovations are normally distributed with mean 0 and standard deviation close to 0.6 — see Table 1. The range $[-2, 2]$ in that grid encompasses 99.91% of the possible realisations.

actually raise the same probability by 1.9 percentage points.

The increase in responsiveness can also be found at the intensive margin. The right panel shows that, after the reform, employment adjustments carried out by firms do respond more to profitability innovations. This is mostly driven by the point estimate of θ_1: it was 0.048 in the 2006-2013 period but it more than doubled to 0.104 in 2014-2019. The estimated values of the quadratic component, θ_2, did not change as much. They remain statistically smaller than zero, which implies the existence of a concave hiring rule similar to that documented by Ilut et al. (2018). Yet that concavity is dampened as the estimated value of θ_2 approaches zero. Basically, as costs of adjustment loosen, the hiring rule becomes less concave and increasingly linear. This further amplifies the response to both positive and negative one standard deviation profitability shocks. For firms which do adjust, and conditional on size, one such positive shock would raise employment growth by 2.3 percentage points before the reform, and by 5.6 afterwards. In case of a negative innovation, employment growth would fall by 3.8 percentage points before the reform, and by 6.7 afterwards.

3.4 Responsiveness regression: exit decision

Another extensive margin decision of firms concerns exit. This also relates to the existence of labour adjustment costs. Take a firm which is hit by a very negative profitability innovation. In the absence of labour frictions the firm would be able to costlessly downsize to its new optimal scale. Even accounting for potential fixed costs of operation, some firms would decide stay in the economy. But if the costs of firing are prohibitively high, that labour adjustment may not take place and the firm decides to exit altogether.

To study this decision let $\chi_{i,t}$ be an indicator variable which takes the value 1 when firm i exits the economy in the beginning of period t. The responsiveness to profitability can then be assessed with the following regression:

$$\chi_{i,t} = \psi_1 \log(n_{i,t-1}) + \psi_2 \log(\xi_{i,t-1}) + \epsilon_{i,t} \tag{5}$$

The timing of the independent variables is slightly different from that used in equations (3) and (4). A firm which exits in period t can only be observed in the data up to period $t-1$. This implies that the exit decision responds to the stock of workers and the profitability innovation drawn in that previous period. Table 4 below shows the estimated coefficients.

Two results stand out. First, the probability of exiting depends negatively on size. The more workers a firm has, the less likely it is to exit in the next period. This can be intuitively related to the concept of severance payments: an exit cost which grows linearly with the size of the firm. And second, those same exit decisions also depend negatively on profitability innovations. This is an expected result as firms which draw increasingly positive shocks are less likely to exit. However, that coefficient is not statistically different

Table 4: Log linear approximation of exit policy function

	Full sample: 2006-2019	Pre-reform: 2006-2013	Post-reform: 2014-2019
Lag employment, ψ_1	-0.026^{***}	-0.028^{***}	-0.023^{***}
	(0.004)	(0.004)	(0.004)
Responsiveness, ψ_2	-0.073^{***}	-0.075^{***}	-0.072^{***}
	(0.010)	(0.010)	(0.010)
Industry & Year FE	Yes	Yes	Yes
Observations	3 308 659	1 709 959	1 350 617
R^2	0.116	0.125	0.106

Notes: Standard errors are clustered at the industry level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

across the two time windows. In other words, this means that, conditional on size, the probability of exiting remains unaltered despite the documented changes in adjustment costs and the curvature of the revenue function.

3.5 Summary of collected evidence

The preceding empirical exercises show that employment responsiveness in the 2014-2019 period was greater than before. This happens both at the extensive and intensive margins of labour adjustment. Firms are more likely to adjust, particularly to hire, and when they do, they carry out larger adjustments. There is no noticeable change in the responsiveness of the exit decision. All these features coincide with two important developments in the economy, also documented above: the inclusion of more flexible provisions in the labour law, and a declining curvature of the revenue function mirroring greater product market power.

Despite this reported increase in responsiveness, the Portuguese economy is yet to experience a productivity boost. I showed in Figure 1 that the dispersion of labour productivity remained stubbornly flat after the reform. Why did this happen? Is the variance of productivity immune to changes in the flexibility of the labour market? Or can the interaction between the two mechanisms, *i.e.* reduced labour adjustment costs and a lower curvature of the revenue function, account for this seeming lack of dynamism?

In order to provide an answer to these questions, in the following sections I develop and estimate a partial equilibrium model of labour demand. The goal is twofold. First, the model provides a framework which can be relied upon to rationalise the observed changes in employment responsiveness and labour productivity. And second, it allows to formally express moments of the labour productivity distribution as a function of three pivotal mechanisms: *i)* labour adjustment costs, *ii)* the curvature of the revenue function, and *iii)* the stochastic process of profitability. While Table 1 provides estimates for the latter two, the labour adjustment cost function is estimated via simulated method of moments.

4 Model economy

I build a partial equilibrium model of dynamic labour demand with adjustment costs. The unit of analysis is the firm, which also makes endogenous decisions about entering and exiting the economy. The framework follows closely those of Hopenhayn (1992) and Hopenhayn and Rogerson (1993).

At the beginning of the period surviving firms draw their idiosyncratic profitability level, $A \in \mathcal{A} \subset \mathbb{R}_+$, and know the stock of workers they inherit from the previous period, $n_{-1} \in \mathbb{R}_+$. The stationary AR(1) process which governs A is that specified in Section 3. Firms immediately rely on this duple, (A, n_{-1}), to determine if they want to remain in the economy as incumbents or to exit permanently. Formally:

$$V(A, n_{-1}) = \max \left\{ \underbrace{V^i(A, n_{-1})}_{\text{Incumbents}}, \underbrace{V^o(n_{-1})}_{\text{Exiters}} \right\} \tag{6}$$

Incumbents:

Firms which decide to continue their operations use (A, n_{-1}) to optimally choose employment (number of workers to be employed instantly (n), and hours per worker (h). Formally, their dynamic programming problem is expressed as follows:

$$V^i(A, n_{-1}) = \max_{n \geq 0, h \geq 0} \left\{ R(A, n, h) - \omega(n, h) - c - C(n_{-1}, n) + \beta \int_{A' \in \mathcal{A}} V(A', n) \cdot \mathrm{d}\Pi(A'|A) \right\} \tag{7}$$

The flow of revenue to the firm is determined by the profitability level and the labour demand choices: $R(A, n, h) \equiv A \cdot (nh)^{\alpha}$. The chosen combination of n and h requires, however, that compensation be paid in the amount of $\omega(n, h) \equiv n \cdot (w_0 + w_1 h^{\eta})$, where w_0 is the base wage, w_1 the hourly rate, and $\eta > 1$ the constant elasticity between the hourly wage cost and hours per worker[7]. This formulation assumes that the decision on average hours is static: firms merely trade off the marginal revenue associated with raising average hours of work with the contemporaneous marginal wage cost of doing so.

The choice of the number of workers, on the other hand, is dynamic as it is subject to an adjustment cost function, $C(n_{-1}, n)$. This function has a piecewise formulation to create the distinction between *net* costs of job creation and destruction, and nests both fixed and linear forms of adjustment.

[7] The convexity of the hourly wage schedule is justified with the data. Hours and lagged employment covary negatively: the correlation is -0.119 in the whole sample. The optimal decision on hours implies that $\frac{\partial h}{\partial n_{-1}} = \frac{\alpha-1}{\eta-\alpha} \cdot \varphi(A, n_{-1})^{\frac{\alpha-1}{\eta-\alpha}-1} \cdot \varphi_2(A, n_{-1})$. Under $\varphi_2(A, n_{-1}) > 0$, as estimated in Section 3.2, $\eta > 1 > \alpha$ ensures the existence of that negative relationship.

$$C(n_{-1}, n) \equiv \begin{cases} \Gamma_f + \gamma_f \cdot (n_{-1} - n) & \text{if } n_{-1} > n \\ 0 & \text{if } n_{-1} = n \\ \Gamma_h + \gamma_h \cdot (n - n_{-1}) & \text{if } n_{-1} < n \end{cases}$$

The inclusion of non-convexities, Γ_f and Γ_h, is meant to capture the pervasiveness of inaction first reported by Varejão and Portugal (2007) and also documented in Section 3 before. These can also encompass measures such as the introduction of a more flexible reason to justify fair dismissals. On the other hand, the linear components, γ_f and γ_h, can capture changes to severance payments and searching costs. The subscripts f and h denote *firing* and *hiring*, respectively.

Lastly, incumbents must also bear a fixed cost of production, c. The existence of such a cost generates endogenous exit, which is taken to be permanent.

Exit and entry:

Exit is not free. Firms which find it optimal to permanently shut down their operations must pay severance to all workers who are on payroll. Since no other liabilities are due, the value of exiting is the following:

$$V^o(n_{-1}) = -\gamma_f \times n_{-1} \tag{8}$$

The model also includes endogenous entrance in the economy. For each exiting firm there exists a potential entrant which is endowed with a signal about their prospective profitability in the next period. That signal, $\tilde{A} \in \mathcal{A}$, is drawn from the stationary distribution. The firm can then decide to enter the following period and start operating at the lowest level of employment, $\underline{n} \in \mathbb{R}_+$, if and only if the conditional expected value of becoming an incumbent is weakly positive:

$$\int_{A' \in \mathcal{A}} V(A', \underline{n}) \cdot \mathrm{d}\Pi(A'|\tilde{A}) \geq 0 \tag{9}$$

Stationary equilibrium:

Absent aggregate shocks, a stationary equilibrium of the economy corresponds to a steady-state distribution over (A, n). That distribution must exist and is defined by the following elements:

1. The profitability conditional CDF, $\Pi(\cdot|A)$;

2. The employment policy function, $n \equiv \varphi(A, n_{-1})$, which solves the dynamic problem in (7);

3. The exit policy function, $\chi(A, n_{-1})$, which takes the value 1 when $V^o(n_{-1}) > V^i(A, n_{-1})$, and 0 otherwise;

4. The entrance policy function, $\psi(\tilde{A})$, which takes the value 1 when condition (9) holds, and 0 otherwise.

5 Estimation

The model of the previous section includes a revenue function whose parameters α, ρ, and σ were estimated in Section 3. There are seven remaining parameters to identify: $\Theta \equiv (\gamma_f, \Gamma_f, \gamma_h, \Gamma_h, \eta, w_0, c) \in \mathbb{R}^7$. In this section I describe their estimation procedure, the results, and discuss the strength of the identification scheme.

5.1 Estimation approach

I rely on a simulated method of moments (SMM) to estimate Θ. The corresponding estimator, $\hat{\Theta}$, minimises the following objective function, F:

$$\hat{\Theta} \equiv \arg\min_{\Theta \in \mathbb{R}^7} \left\{ \left(\left(m^d - m^s(\Theta) \right) / m^d \right) W \left(\left(m^d - m^s(\Theta) \right) / m^d \right)' \right\}, \tag{10}$$

where $m^s(\Theta)$ is the vector of simulated moments obtained when the model is parameterised by Θ, and m^d the data counterparts of those moments. Below I describe the seven targeted moments and offer some preliminary justification as to why they were picked. Since the model is just-identified, the weighting matrix, W, is the conforming identity matrix.

Identification:

Identification concerns the choice of targeted moments. These are listed in Table 5.

Table 5: Targeted moments in structural estimation

ϱ_{An}	Covariance b/w TFPR and employment	λ_1	Serial correlation log employment
θ_2	Concavity of hiring rule	ψ_2	Responsiveness of exit decision
$\sqrt{\Sigma_{\overline{r}}}$	Std deviation of log labour productivity	$\overline{n}$	Mean of log employment
$\overline{x}$	Time-series average of exit rate		

As mentioned in Section 3 I now focus more on the covariance between the log level of profitability and log employment, ϱ_{An}. This is justified by two reasons. First, because this is an object with an intrinsic connection to concept of responsiveness and adjustment costs. If adjustments are prohibitively expensive, for example, then both the covariance and responsiveness approach zero because firms do not alter their labour force regardless of how profitable they may be. And second, because it will be shown to be important to decompose the variance of labour productivity. Cooper et al. (2024a) also rely on it to study the productivity implications of higher adjustment costs in the US.

The following three moments correspond to coefficients from regressions explored in the previous section. The effect exerted by non-convexities, for instance, is well captured by the serial correlation of log employment in the policy function (2), λ_1. Although those non-convexities also impact the responsiveness of the exit decision, the latter is highly influenced by the size of linear costs of firing. The coefficient ψ_2 in (5) can be quite important to identify those. Reduced severance payments slash the cost of exiting and this may prompt firms to respond differently when hit by a negative realisation of profitability. This notwithstanding, the existence of convex costs is best mirrored in the intensive margin decisions of employment adjustment. As a consequence, I also rely on θ_2, the coefficient which defines the concavity of the hiring rule in (4).

The remaining moments are usual in the literature. Decisions on hours are tamed by two components of the wage function: the hourly rate, w_1, and the elasticity between the hourly wage cost and hours per worker, η. I keep the hourly rate outside the estimation procedure and calibrate it such that the steady-state value of hours per worker matches the yearly average found in the data. The elasticity, on the other hand, is kept within the estimand vector and is mostly targeted by the standard deviation of labour productivity, $\sqrt{\Sigma_{\overline{r}}}$. Finally, the mean of log employment, $\overline{n}$, is highly influenced by the base wage (w_0), while the average exit rate, $\overline{x}$, is exceptionally informative about the fixed cost of operation (c).

5.2 Estimation results

<u>*Parameters:*</u>

I estimate the model parameters for two time windows: one targeting the data moments of 2006-2013, and another one targeting those of 2014-2019. In doing so I can capture the effects the labour market intervention has had on the labour adjustment cost parameters and therefore on the various moments. Table 6 below shows those point estimates in the two periods.

Table 6: Parameter estimates

	Labour adjustment costs				Wage function		Fixed cost
	γ_f	Γ_f	γ_h	Γ_h	η	w_0	c
2006-2013	0.23	0.55	0.18	0.13	4.25	1.25	0.04
2014-2019	0.19	0.35	0.27	0.07	4.16	0.73	0.06

Notes: All labour adjustment costs parameters and the fixed cost of operation are expressed as a share of the average revenue of the simulated economy. The economy is simulated at a quarterly frequency and aggregated to match the annual moments of the data. The quarterly discount factor is set to $\beta = 0.99$, and the hourly rate w_1 is calibrated such that the steady-state value of hours per worker replicates the observed yearly average.

All four labour adjustment cost parameters, as well as the fixed cost of operation, are expressed as shares of the average revenue in the simulated panel. These estimated figures are consistent with a sizeable cut in firing costs. Between 2006 and 2013, a firm was predicted to forego 78% of the economy's average revenue when sacking a single employee. This entails both the fixed cost and a linear component reminiscent of severance payments. That cost was brought down to 54% in the subsequent period. Evidence on hiring costs is not as clear. Taking one additional employee would originally wipe out 31% of the average revenue. Now, it does so by 34%.

I also allow the parameters of the wage function to vary over time. Yet the elasticity between average hours of work and the hourly wage cost, η, remains largely unchanged. For a 1% increase in average hours, the hourly wage bill of firms was estimated to increase by 4.25% in the first time window, and by 4.16% in the second one. Finally, the fixed cost of operation is estimated to increase from 4% to 6% of the average revenue.

Targeted moments:

The fitness of the SMM estimation is reported in Table 7 below. Each of the two blocks shows the data and simulated moments in their corresponding time window.

Table 7: SMM estimation fit — targeted moments

		Responsiveness				Distributions		Exit	Distance
		ϱ_{An}	λ_1	θ_2	ψ_2	$\sqrt{\Sigma_{\overline{r}}}$	$\overline{n}$	$\overline{x}$	$F(\hat{\Theta})$
2006-2013	*Data*	0.64	0.95	-0.02	-0.08	1.14	1.24	8.16	
	Model	0.58	0.83	-0.02	-0.07	1.77	1.17	6.03	0.41
2014-2019	*Data*	0.81	0.96	-0.02	-0.07	1.16	1.14	6.44	
	Model	0.67	0.83	-0.02	-0.08	1.59	1.16	6.19	0.21

Notes: Targeted moments are the covariance between the log level of TFPR and log employment ϱ_{An}, the serial correlation of log employment λ_1 in (2), the concavity of the hiring rule θ_2 in (3), the exit responsiveness coefficient ψ_2 in (5), the standard deviation of log labour productivity $\sqrt{\Sigma_{\overline{r}}}$, the mean of log employment $\overline{n}$, and the mean of the exit rate $\overline{x}$. The measure of distance between data and simulated moments defined in (10) is $F(\hat{\Theta})$.

The sum of squared distances between data and simulated moments, $F(\hat{\Theta})$, is noticeably small in the two periods. This indicates the model does a successful job replicating the targeted moments. This is particularly true for the responsiveness moments, both qualitatively and quantitatively. For example, the model is able to reproduce the observed concavity of the employment hiring rule, $\theta_2 < 0$; it also picks up the significant increase in the covariance between log TFPR and log employment. Although not shown here, this model-based increase in responsiveness can also be found in untargeted moments such as the parameter in the policy function for labour, λ_2, and the correlation coefficient between log profitability and log employment.

Figure 5: Local elasticities of moments wrt parameters (left) and parameters wrt moments (right)

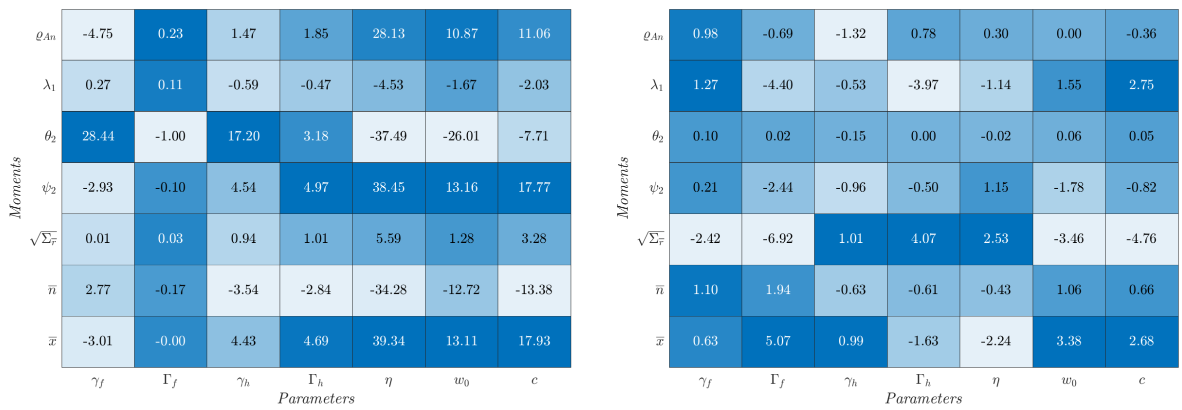

Notes: In the *left panel* the element (m, p) of the matrix gives the local elasticity of moment m following a 1% change in parameter p. In the *right panel* element (m, p) of the matrix gives the local elasticity of parameter p following a 1% change in moment m. All these local elasticities are computed when the parameters are identified for the full period of the sample, 2006-2019.

The predicted values of the *standard deviation* of log labour productivity are not substantially off from their data counterparts: 1.77 vs. 1.14 in the pre-reform period, 1.59 vs. 1.16 afterwards. Yet, however small these differences may be, they compound once I rely on the *variance*. This moment is then projected to be 3.14 in the 2006-2013 period (1.31 in the data), and 2.53 in 2014-2019 (1.35 in the data)[8]. Conversely, both the exit rate and specially the mean of the employment distribution are very well targeted.

Robustness of the identification scheme:

The strength of the identification scheme is assessed in two different ways. In Figure 5 below I uncover the intricate mapping from parameters to moments. More specifically, elements in left panel indicate the percentage change in each targeted moment following a 1% change in each parameter, all else equal. In the right panel I rely on the methodology of Andrews et al. (2017) and implement the reverse exercise: how much do parameters vary following changes in the targeted moments.

The linear cost of firing, reminiscent of severance payments, is particularly well identified. All else equal, raising those linear costs by 1% reduces the covariance between log TFPR and log employment by 4.75%, increases the average size of firms by 2.77% and, as a consequence, the exit rate drops by 3.01%. These are meaningful results: severance payments took centre stage in the reform and they are shown to exert widespread and sizeable effects in the economy.

[8] An alternative option would then be to immediately target the variance instead of the standard deviation. Yet the overall performance of the model is considerably worse in that case. To reduce the dispersion of productivity the minimisation algorithm raises adjustment costs and η to the point where the economy is only populated by very large firms with similar labour productivity levels. Since there is hardly any exit, the new solution gets the variance almost right but at the expense of the two moments connected to that decision, ψ_2 and $\overline{x}$.

The dispersion of labour productivity, on the other hand, is immune to them. Although the standard deviation does respond almost on a one-to-one fashion to linear and fixed costs of hiring, it is mostly sensitive to changes in the elasticity between the hourly wage cost and hours per worker, and the fixed cost of operation. And as argued before, the average exit rate is mostly informed by the latter.

Alternative parameterisation:

The estimation algorithm described in equation (10) aimed to minimise a sum of equally-weighted squared distances between simulated and data moments. The conforming identity matrix provided those equal weights. I now consider an alternative estimation strategy. More specifically, I follow Adda and Cooper (2003) and rely on a two-stage SMM procedure which takes advantage of the optimal weighting matrix. Conditioning on the baseline estimates reported in Table 6, I simulate 1 000 economies and store each of the resulting vectors of simulated moments. The inverted variance-covariance matrix of all those moments is then plugged back in (10) to take the role of weighting matrix in a new iteration of the minimisation procedure.

The point estimates can be found in Table C1, whereas the comparison between data and simulated moments is shown in Table C2. Overall, the results are consistent with those discussed here. If anything, the reduction in labour adjustment costs is made slightly more prominent. Sacking one employee would cost up to 100% of the average revenue between 2006 and 2013, but only 64% in the subsequent period. The evidence on unit costs of hiring, on the other hand, remains less clear. Hiring one more employee would initially cost 35% of the average revenue; it now takes 39%.

6 Quantitative analysis

In this section I rely on the estimated version of the model to shed light on the muted evolution of the labour productivity variance. Although the whole distribution remained largely unaltered (see Figure 1, panel A), I place a greater focus on its dispersion given because it is tightly linked to the concept of misallocation.

I start with two exercises which confirm that, albeit simple, the estimated model encompasses the necessary mechanisms to better understand the documented facts regarding not just labour productivity but also employment responsiveness. First, I conduct a counterfactual analysis. While keeping everything else constant, I separately alter the parameter values of the labour adjustment cost function, and the curvature of the revenue function. The goal is to disentangle the effects each channel produces. Second, to emphasise the impact of the labour market reform alone, I keep all other parameters constant at their 2006-2013 values and just re-estimate the components of $C(n_{-1}, n)$ while targeting responsiveness moments and the standard deviation of labour productivity in 2014-2019. This partial re-estimation exercise aims to show that both responsiveness and labour productivity are sensitive to changes in costs of adjustment.

In the last subsection I demonstrate that the variance of labour productivity can ultimately be written as a closed-form expression of the curvature of the revenue function and the variance-covariance matrix of TFPR, employment, and hours worked. Since the cross-sectional dispersion of TFPR has long been interpreted as a measure of misallocation, studying the decomposition of the labour productivity variance is of unequivocal importance for macroeconomists. I pinpoint which elements are more prominent to explain that dispersion in the data, in the baseline version of the model, and in counterfactual economies where labour adjustment costs and the whole profitability process of firms had not changed.

6.1 Counterfactuals

I consider three types of counterfactual cases. In the first one I ask what would happen to the economy if, all else equal, the labour market reform would be fully reversed. Call this scenario "**Counterfactual,** $C(n_{-1}, n)$". To be clear, this means that while the four elements which describe the labour adjustment cost function $C(n_{-1}, n)$ retake their initial values, all other parameters of the economy are kept at their 2014-2019 levels. This includes the curvature of the revenue function, the serial correlation and variance which describe the profitability process, the components of the wage function, and the fixed cost of operation. It is by doing so that I can precisely identify the effects exerted by labour adjustment costs alone. In the second one I report the effects which would have been produced had policy-makers gone farther and removed all sorts of adjustment costs. Call this case "**Counterfactual**, *free adjustments*". While the first exercise describes a movement towards a more stringent labour market, the second one implies a full liberalisation. Lastly, in the third case I study a fictional economy where all parameters would keep their most recent values with the exception of the curvature of the revenue function. Let this scenario be called "**Counterfactual,** α".

Table 8 shows the outcome of these exercises. I focus on two moments which were targeted in the estimation procedure: *i)* the covariance between profitability and employment, which mirrors the evolution of responsiveness, and *ii)* the standard deviation of log labour productivity. The distance between the whole set of data and simulated moments is reported in the last column.

Start with the effects produced by labour adjustment costs only. The first two counterfactual exercises show that, all else equal, a more rigid/flexible labour market is associated with falling/rising responsiveness but rising/falling labour productivity dispersion. If the labour market reform would be entirely reversed, the covariance would be halved (from 0.67 to 0.31) and the dispersion of labour productivity would increase 11% (from 1.59 to 1.76). Conversely, if policy-makers had completely wiped out all adjustment costs, the covariance would have increased by a factor of four, and the standard deviation of labour productivity would have dropped 35%.

Table 8: Counterfactual exercises

		Targeted moments		Distance
		ϱ_{An}	$\sqrt{\Sigma_{\overline{r}}}$	$F(\hat{\Theta})$
2006-2013	*Data*	0.64	1.14	
	Model	0.58	1.77	0.41
2014-2019	*Data*	0.81	1.16	
	Model	0.67	1.59	0.21
	Counterfactual, $C(n_{-1}, n)$	0.31	1.76	1.53
	Counterfactual, *free adjustments*	2.79	1.03	19.08
	Counterfactual, α	0.24	1.36	8.95

Notes: In **Counterfactual,** $C(n_{-1}, n)$ only the four adjustment cost parameters ($\gamma_f, \Gamma_f, \gamma_h$, and Γ_h) are kept at their 2006-2013 levels. All other parameters take their 2014-2019 values.

In **Counterfactual,** *free adjustments* all four adjustment costs parameters ($\gamma_f, \Gamma_f, \gamma_h$, and Γ_h) are set to zero. All other parameters take their 2014-2019 values.

In **Counterfactual,** α only the curvature of the revenue function (α) is kept at its 2006-2013 levels. All other parameters take their 2014-2019 values.

Consider now the counterfactual economy where, all else equal, the curvature of the revenue function would return to its original higher value. The dispersion of labour productivity would be smaller, as the standard deviation would fall from 1.59 to 1.36. This means that had it not been for a lower curvature of the revenue function, the dispersion of labour productivity would have fallen 14%. Recall that reversing the labour market reform would have increased it 11%. In absolute terms, these are remarkably similar effects, which suggests that the two channels do exert a quantitatively neutral impact on the labour productivity dispersion. The standard deviation of labour productivity hardly moved because the reduction which would be implied by the labour market reform was fully offset by the fall in the curvature of the revenue function.

The impact of α on employment responsiveness, on the other hand, is slightly more nuanced. Notice the covariance effectively drops from 0.67 in the baseline model to only 0.24. This is seemingly at odds with what was to be expected: a larger marginal revenue of labour, prompted by a rising curvature, should boost employment responsiveness. I do obtain this result once I account for the unparalleled change in the endogenous distribution of employment[9]. Instead of a simple covariance, consider the correlation coefficient between profitability and employment. This is a more useful metric because it incorporates the varying

[9]The estimated model predicts the mean of log employment in 2014-2019 to be 1.16 and its standard deviation 1.94. In the "**Counterfactual,** α" case, however, this distribution is largely affected: the mean jumps to 3.66 and the dispersion shrinks to 0.45. As the curvature rises and labour adjustment costs remain low, the economy becomes populated by a mass of identically large firms where there is hardly any exit.

cross-sectional dispersion of labour. In that case, responsiveness actually doubles from 0.2 in the estimated version of the model to 0.43 in the counterfactual economy.

Two conclusions should be drawn from these counterfactual exercises. First, the model corroborates a theoretical prediction which had been put forward in advance: while reduced costs of adjustment enhance firms' employment responsiveness, a lower curvature of the revenue function shrink it. Second, and more importantly, these two channels exert a quantitatively similar effect on the labour productivity dispersion, although in opposite directions. This explains why the documented increase in responsiveness was not followed by noticeable changes in the variance of labour productivity: as the curvature of the revenue function fell, it fully offset the positive effects lower costs of adjustment would bring.

Lastly, notice that the fit of the model worsens in all three counterfactual scenarios relative to the baseline. Yet the case where the change is the smallest is the one where I simulate a full reversal of the labour market reform — the distance measure "only" increases by a factor of seven, from 0.21 to 1.53. This could potentially be construed as evidence that the set of targeted moments is not very sensitive to changes in the parameters of the adjustment cost function. To discredit this hypothesis, in the next sub-section I demonstrate that appealing solely to changes in labour adjustment costs is sufficient to target most responsiveness coefficients and the standard deviation of labour productivity.

6.2 Re-estimation of labour adjustment costs

In the previous exercise I separately switched off the two mechanisms deemed responsible for the change in responsiveness and the standard deviation of labour productivity. I now do something slightly different. With the exception of the four parameters which describe the labour adjustment cost function, suppose all others would have been kept constant at their 2006-2013 levels. Again, this includes the curvature of the revenue function, the serial correlation and variance of the profitability process, the components of the wage function, and the fixed cost of operation. In order to focus more explicitly on the effects brought about by a more flexible labour market, I re-estimate the four parameters which fully describe $C(n_{-1}, n)$ while targeting *i)* the covariance between profitability and employment, *ii)* the serial correlation of log employment, *iii)* the concavity of the hiring rule, and *iv)* the standard deviation of log labour productivity. This just-identified estimation exercise aims to reinforce how adjustment costs alone can exert a significant influence in both employment responsiveness and the dispersion of labour productivity. The results are displayed in Table 9.

The distance between data and simulated moments is only 0.03, a remarkably smaller value than in the baseline (0.21). This indicates the model does a great job replicating the targeted moments. Put differently, by appealing solely to changes in labour adjustment costs one can reproduce various patterns observed in the data. The concavity of the hiring rule (-0.07) and the covariance between profitability and employment

Table 9: Partial re-estimation of labour adjustment costs

		Targeted moments				Distance
		ϱ_{An}	λ_1	ψ_2	$\sqrt{\Sigma_{\bar{r}}}$	$F(\hat{\Theta})$
2006-2013	*Data*	0.64	0.95	-0.08	1.14	
	Model	0.58	0.83	-0.07	1.77	0.41
2014-2019	*Data*	0.81	0.96	-0.07	1.16	
	Model	0.67	0.83	-0.08	1.59	0.21
	Partial re-estimation	0.83	0.84	-0.07	1.30	0.03

Notes: In **Partial re-estimation** all but the four adjustment cost parameters are kept at their 2006-2013 levels. These are estimated while targeting the four moments displayed in the table: ϱ_{An}, λ_1, ψ_2, and $\sqrt{\Sigma_{\bar{r}}}$.

(0.83), for example, match almost exactly their data counterparts: -0.07 and 0.81, respectively. The model-predicted value for the standard deviation of labour productivity is also of interest. The closer the model can be to the observed value is actually in this partial re-estimation exercise.

6.3 Variance decomposition of labour productivity

Changes in responsiveness patterns should have macroeconomic implications. This is a conclusion shared by Decker et al. (2020) and Cooper et al. (2024a) alike. In Tables 8 and 9 I assessed those through the dispersion of labour productivity, *i.e.* the standard deviation of log labour productivity, because it is a measure intrinsically connected to the concept of labour misallocation. Surprisingly, despite all the significant transformations in the Portuguese economy, that dispersion remained fairly stable between 2006-2013 and 2014-2019: the standard deviation went from 1.15 to 1.16, and the corresponding variance from 1.31 to 1.35. In fact, I showed in Figure 1 that the whole distribution remained unchanged. But that does not imply that there cannot be substantial dynamism in the individual elements which jointly determine that dispersion. In this subsection I start by deriving a closed form solution to the decomposition of that variance and then assess the relative contribution of all those elements both in the data and the estimated model. The usefulness of this exercise is twofold. First, it allows me to formally establish the connection between the dispersion of labour productivity and the misallocation measure defined in Hsieh and Klenow (2009): the cross-sectionl variance of log TFPR. And second, it offers an adequate framework to study which economic forces are more important to understand the variation in labour productivity.

The natural logarithm of labour productivity corresponds to the difference between log revenues and log employment. Using the revenue function derived initially, and omitting firm- and time-specific subscripts for convenience of notation, I can express it as $\log(R) - \log(n) \equiv \log(A) + (\alpha - 1)\log(n) + \alpha \log(h)$. From the properties of the variance of a sum:

$$\operatorname{Var}\big(\log(R) - \log(n)\big) = \mathbf{tr}\left(\begin{bmatrix} 1 & 2(\alpha-1) & 2\alpha \\ 0 & (\alpha-1)^2 & 2\alpha(\alpha-1) \\ 0 & 0 & \alpha^2 \end{bmatrix} \cdot \begin{bmatrix} \operatorname{Var}\big(\log(A)\big) & \operatorname{Cov}\big(\log(n), \log(A)\big) & \operatorname{Cov}\big(\log(h), \log(A)\big) \\ \operatorname{Cov}\big(\log(A), \log(n)\big) & \operatorname{Var}\big(\log(n)\big) & \operatorname{Cov}\big(\log(h), \log(n)\big) \\ \operatorname{Cov}\big(\log(A), \log(h)\big) & \operatorname{Cov}\big(\log(n), \log(h)\big) & \operatorname{Var}\big(\log(h)\big) \end{bmatrix}\right) \tag{11}$$

Proof. See Appendix C. □

The variance of labour productivity is simplified to an explicit function of two elements: *i)* the curvature of the revenue function, which fully parameterises the first matrix, and *ii)* the variance-covariance matrix of profitability, employment, and average hours of work. Given these two matrices one just needs to compute the trace of the product between them. Notice this does not mean labour adjustment costs and the stochastic process of profitability have a dismissive role. Actually, they influence the dispersion of labour productivity through their implicit effect on the individual variances and covariances.

For what follows I will partition the two matrices and just focus on the first 2×2 block matrices. This allows me to abstract from any variance or covariance elements associated with hours decisions. This is justified by two reasons. First, because the paper focuses on evaluating how labour adjustment costs and the curvature of the revenue function influence the choice on the number of employees. And second, because the discarded elements explain no more than 4% of the observed labour productivity variance. In other words, roughly 96% of the observed dispersion in labour productivity can be explained by the variance-covariance elements between profitability and employment:

$$\operatorname{Var}\big(\log(R) - \log(n)\big)^{-1} \times \underbrace{\mathbf{tr}\left(\begin{bmatrix} 1 & 2(\alpha-1) \\ 0 & (\alpha-1)^2 \end{bmatrix} \cdot \begin{bmatrix} \Sigma_A & \varrho_{An} \\ \varrho_{An} & \Sigma_n \end{bmatrix}\right)}_{=\Sigma_A + (\alpha-1)^2 \Sigma_n + 2(\alpha-1)\varrho_{An}} \approx 0.96, \tag{12}$$

where Σ_A and Σ_n are the respective variances of log profitability and employment, and ϱ_{An} is the covariance between them. This last moment is one of the targets in the structural estimation exercise.

Decomposing the variance

The decomposition derived in (12) is useful to understand which economic forces more prominently explain the variance of labour productivity. To reiterate, here I focus on the three components which jointly explain approximately 96% of the observed dispersion: *i)* the variance of log TFPR, Σ_A, *ii)* the weighted variance of log employment, $(\alpha-1)^2\,\Sigma_n$, and *iii)* the weighted covariance between profitability and employment, $2(\alpha-1)\,\varrho_{An}$. In Table 10 I report the total variance, the level of these three elements which comprise it, and the sum of their corresponding shares.

Table 10: Variance decomposition of log labour productivity

		Variance	Elements of variance decomposition			Share of variance
		$\Sigma_{\bar{r}}$	Σ_A	$(\alpha-1)^2\,\Sigma_n$	$2(\alpha-1)\,\varrho_{An}$	
2006-2013	*Data*	1.31	1.60	0.23	-0.56	0.96
	Model	3.14	1.15	1.37	-0.51	0.64
2014-2019	*Data*	1.35	1.82	0.37	-0.90	0.96
	Model	2.53	1.34	1.17	-0.74	0.70
	Counterfactual, $C(n_{-1}, n)$	3.09	1.35	1.27	-0.34	0.74
	Counterfactual, $R(\cdot)$	1.86	1.58	0.04	-0.21	0.76

Notes: The first column displays the variance of log labour productivity, $\Sigma_{\bar{r}}$. The following three columns show the elements which are part of the variance decomposition: the variance of log profitability, Σ_A, the weighted variance of log employment, $(\alpha-1)^2\Sigma_n$, and the weighted covariance between log profitability and log employment, $2(\alpha-1)\varrho_{An}$. The last column reports the sum of the three preceeding elements divided by the variance of log labour productivity.

In **Counterfactual,** $C(n_{-1}, n)$ only the four adjustment cost parameters (γ_f, Γ_f, γ_h, and Γ_h) are kept at their 2006-2013 levels. All other parameters take their 2014-2019 values. In **Counterfactual,** $R(\cdot)$ all elements of the revenue function (α, ρ, and σ) are kept at their 2006-2013 levels. All other parameters take their 2014-2019 values.

Start with the decomposition found in the data. The variance of TFPR plays a disproportionately large role: it corresponds to $(1.6/1.31) = 122\%$ and $(1.82/1.35) = 135\%$ of the labour productivity variance in the pre- and post-reform periods, respectively. The share of the weighted variance of employment is smaller: only 17% and 28%. These are two non-trivial evolutions which can be fully attributed to the estimated parameters of the revenue function. Since the total variance of labour productivity remained fairly stable across the two periods, these varying shares must essentially mirror changes in the numerators, *i.e.* changes in the variance of TFPR, and changes in the weighted variance of employment. The variance of profitability did increase from 1.60 to 1.82 and, more importantly, that rising dispersion can only be accounted for by the documented changes in ρ and σ because $\Sigma_A = \sigma^2/(1-\rho^2)$. The *unweighted* variance of employment, on the other hand, remained constant: 1.18 in the pre-reform period, 1.20 afterwards. But what matters for the decomposition is the *weighted* variance, where the weight is given by $(\alpha-1)^2$. As the estimated curvature of the revenue function falls, that weight increases and, all else equal, so too does the share of the weighted variance of employment. Lastly, the share of the weighted covariance between profitability and employment is a more complex object to be studied. In absolute terms, the observed increase in that share reflects mostly two channels: a larger weight (again in absolute terms) prompted by the fall in α, and a greater covariance.

The model corroborates the claim that the parameters of the profitability process fully determine the change in the first two decomposition elements. Suppose again the labour market reform would be reversed. I had already shown that this "**Counterfactual,** $C(n_{-1}, n)$" case would imply a greater dispersion of labour productivity. More specifically, the variance would rise from 2.53 in the baseline model to 3.09. However, both the variance of TFPR and the weighted variance of employment would be kept almost unchanged. This contrasts with a counterfactual scenario wherein all estimated parameters of the revenue function would retake their original values. Call this case "**Counterfactual,** $R(\cdot)$". Although the variance would drop and approach its data counterpart, there would be considerably more dynamism in its decomposition than that prompted by a tightening of labour adjustment costs.

7 Conclusion

This paper investigated whether a labour market reform which reduces adjustment costs can be an effective tool to boost productivity. In theory it should: looser constraints allow firms to be more responsive to changes in their profitability circumstances, and this improves the allocation of labour across the economy. Surprisingly, a reform of this sort in Portugal did not deliver on its goal. Although labour adjustment costs were slashed, the declining trend in the pace of job flows across firms remained in place and productivity gains were nonexistent.

I start by addressing this puzzling result with a rich dataset of Portuguese firms. The conclusion drawn from an estimated log linear approximation of the employment policy function is unequivocal: the responsiveness of firms to exogenous profitability shocks almost doubled in the period after the labour market reform. This can be verified both at the extensive (a convex adjustment hazard) and intensive (a concave hiring rule) margins of labour adjustment. The likelihood of exit, however, was not affected.

These findings suggest that a rising employment responsiveness of firms has been incapable of speeding up the labour reallocation process from low- to more-profitable firms. The reason lies in the documented increase of firms' product market power. As they charge higher markups, the curvature of the revenue function dwindles. This not only dampens the response of firms to profitability shocks, but can also be detrimental to the allocative efficiency of labour in the economy.

To corroborate this point I estimate a partial equilibrium model of labour demand with adjustment costs, and endogenous entry and exit decisions. The main usefulness of the model resides in its ability to separately identify the channels deemed responsible for the evolution of the labour productivity variance, an object intrinsically linked to the concept of misallocation. Through counterfactual exercises I demonstrate that looser constraints in the labour law and a falling curvature of the revenue function exert quantitatively similar effects, but in opposite directions. A constant dispersion arises because these two mechanisms cancel each other. Nevertheless, I also show that a flat variance hides substantial dynamism in its decomposition, most of which governed by the change in the curvature of the revenue function and the parameters which describe firms' stochastic profitability process.

This paper takes use of a very particular context to contribute to the debate on the sources of firms' employment responsiveness, and the implications for labour productivity. A promising line of research which can be pursued from here on concerns the potentially endogenous link between smaller costs of adjustment and the rising market power which accounts for the fall in the curvature of the revenue function. Throughout the paper I took these to be orthogonal: Portuguese lawmakers cut labour adjustment costs and this *coincided* with a fall in the estimated curvature of the revenue function. What if this was not a coincidence? With cheaper exit and looser constraints some firms could have taken advantage of exceedingly positive profitability shocks to grow disproportionately more and thus solidify a more dominant position. Although this possibility does not taint my identification strategy or my results, it may cast serious doubt on the ability to boost labour productivity and improve allocation through a more flexible labour market with reduced costs of adjustment.

References

Ackerberg, D. A., Caves, K., and Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica*, 83(6):2411–2451.

Adda, J. and Cooper, R. W. (2003). *Dynamic economics: quantitative methods and applications*. MIT press.

Addison, J. T., Portugal, P., and Vilares, H. (2017). Unions and collective bargaining in the wake of the great recession: evidence from portugal. *British Journal of Industrial Relations*, 55(3):551–576.

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the sensitivity of parameter estimates to estimation moments. *The Quarterly Journal of Economics*, 132(4):1553–1592.

Autor, D. H., Kerr, W. R., and Kugler, A. D. (2007). Does employment protection reduce productivity? evidence from us states. *The Economic Journal*, 117(521):F189–F217.

Banco de Portugal Microdata Research Laboratory (2025). Central Balance Sheet Harmonized Panel.

Baqaee, D. R. and Farhi, E. (2020). Productivity and misallocation in general equilibrium. *The quarterly journal of economics*, 135(1):105–163.

Bentolila, S. and Bertola, G. (1990). Firing costs and labour demand: how bad is eurosclerosis? *The Review of Economic Studies*, 57(3):381–402.

Bond, S., Hashemi, A., Kaplan, G., and Zoch, P. (2021). Some unpleasant markup arithmetic: Production function elasticities and their estimation from production data. *Journal of Monetary Economics*, 121:1–14.

Caballero, R. J., Engel, E. M. R. A., and Haltiwanger, J. (1997). Aggregate employment dynamics: Building from microeconomic evidence. *American Economic Review*, 87(1):115–137.

Caggese, A., Guler, O., Mariathasan, M., and Mulier, K. (2022). Firing costs and productivity: Evidence from a natural experiment. *Available at SSRN*.

Card, D. and Cardoso, A. R. (2022). Wage flexibility under sectoral bargaining. *Journal of the European Economic Association*, 20(5):2013–2061.

Cardoso, A. R. and Portugal, P. (2005). Contractual wages and the wage cushion under different bargaining settings. *Journal of Labor economics*, 23(4):875–902.

Cingano, F., Leonardi, M., Messina, J., and Pica, G. (2016). Employment protection legislation, capital investment and access to credit: evidence from italy. *The Economic Journal*, 126(595):1798–1822.

Cooper, R., Gong, G., and Yan, P. (2018). Costly labour adjustment: general equilibrium effects of china's employment regulations and financial reforms. *The Economic Journal*, 128(613):1879–1922.

Cooper, R., Haltiwanger, J. C., and Willis, J. (2024a). Declining responsiveness at the establishment level: Sources and productivity implications. Technical Report 32130, National Bureau of Economic Research.

Cooper, R., Horn, C.-W., and Indraccolo, L. (2024b). Covid and productivity in europe: A responsiveness perspective. *European Economic Review*, 163:104655.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633.

De Loecker, J. and Eeckhout, J. (2018). Global market power. Technical Report 24768, National Bureau of Economic Research.

De Loecker, J., Eeckhout, J., and Mongey, S. (2021). Quantifying market power and business dynamism in the macroeconomy. Technical Report 28761, National Bureau of Economic Research.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644.

De Loecker, J. and Syverson, C. (2021). An industrial organization perspective on productivity. In *Handbook of industrial organization*, volume 4, pages 141–223.

De Loecker, J. and Warzynski, F. (2012). Markups and firm-level export status. *American economic review*, 102(6):2437–2471.

De Ridder, M., Grassi, B., and Morzenti, G. (2024). The hitchhiker's guide to markup estimation: assessing estimates from financial data. *Work. Pap., London Sch. Econ./Bocconi Univ./Analysis Group, UK/Milan, Italy/Paris, France*.

Decker, R. A., Haltiwanger, J., Jarmin, R. S., and Miranda, J. (2016). Declining business dynamism: What we know and the way forward. *American Economic Review*, 106(5):203–207.

Decker, R. A., Haltiwanger, J., Jarmin, R. S., and Miranda, J. (2020). Changing business dynamism and productivity: Shocks versus responsiveness. *American Economic Review*, 110(12):3952–3990.

Edmond, C., Midrigan, V., and Xu, D. Y. (2023). How costly are markups? *Journal of Political Economy*, 131(7):1619–1675.

Hamermesh, D. S. (1989). Labor demand and the structure of adjustment costs. *American Economic Review*, 79(4):674–689.

Hopenhayn, H. and Rogerson, R. (1993). Job turnover and policy evaluation: A general equilibrium analysis. *Journal of political Economy*, 101(5):915–938.

Hopenhayn, H. A. (1992). Entry, exit, and firm dynamics in long run equilibrium. *Econometrica*, pages 1127–1150.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, 124(4):1403–1448.

Ilut, C., Kehrig, M., and Schneider, M. (2018). Slow to hire, quick to fire: Employment dynamics with asymmetric responses to news. *Journal of Political Economy*, 126(5):2011–2071.

OECD (2014). The 2012 labour market reform in spain, a preliminary assessment.

OECD (2017). Labour market reforms in portugal 2011–2015, a preliminary assessment.

Pinelli, D., Torre, R., Pace, L., Cassio, L., and Arpaia, A. (2017). The recent reform of the labour market in italy: A review. *European Commission Discussion Paper 072*.

Santos, C. D., Costa, L. F., and Brito, P. B. (2022). Demand, supply and markup fluctuations. *The Economic Journal*, 132(644):1620–1645.

Varejão, J. and Portugal, P. (2007). Employment dynamics and the structure of labor adjustment costs. *Journal of Labor Economics*, 25(1):137–165.

Yeh, C., Macaluso, C., and Hershbein, B. (2022). Monopsony in the us labor market. *American Economic Review*, 112(7):2099–2138.

A Revenue function

A.1 Derivation

Firms produce according to $y = z \cdot (nh)^{\zeta} \cdot k^{1-\zeta}$, and face a demand price schedule given by $p(y) = (\theta/y)^{1/\nu}$. A detailed explanation of these two functions can be found in the main text. For the sake of clarity, I am omitting time- and firm-specific subscripts.

The revenue flow net of capital payments can be defined as follows:

$$R(\cdot) \equiv \max_k \left\{ p(y) \cdot y - rk \right\} = \max_k \left\{ \theta^{\frac{1}{\nu}} \cdot \left(z \cdot (nh)^{\zeta} \cdot k^{1-\zeta} \right)^{\frac{\nu-1}{\nu}} - rk \right\},$$

where r corresponds to exogenous price of capital, *i.e.* the real interest rate. For convenience of notation let $\varepsilon \equiv (\nu - 1)/\nu$. Under the assumption of fully flexible capital taking and developing the associated FOC is necessary and sufficient to yield the following demand function:

$$k = \left(\frac{\varepsilon(1-\zeta) \cdot \theta^{1-\varepsilon} \cdot \left(z(nh)^{\zeta} \right)^{\varepsilon}}{r} \right)^{\frac{1}{1-\varepsilon(1-\zeta)}}$$

Plugging this solution back in the objective function allows to retrieve a revenue function which is entirely expressed in terms of labour demand, and the stochastic evolution of technological and demand shocks:

$$\begin{aligned} R(\cdot) &= \theta^{1-\varepsilon} \left(z(nh)^{\zeta} \right)^{\varepsilon} \cdot \left(\frac{\varepsilon(1-\zeta) \cdot \theta^{1-\varepsilon} \cdot \left(z(nh)^{\zeta} \right)^{\varepsilon}}{r} \right)^{\frac{\varepsilon(1-\zeta)}{1-\varepsilon(1-\zeta)}} - r \left(\frac{\varepsilon(1-\zeta) \cdot \theta^{1-\varepsilon} \cdot \left(z(nh)^{\zeta} \right)^{\varepsilon}}{r} \right)^{\frac{1}{1-\varepsilon(1-\zeta)}} \\ &= \theta^{\frac{1-\varepsilon}{1-\varepsilon(1-\zeta)}} \left(z(nh)^{\zeta} \right)^{\frac{\varepsilon}{1-\varepsilon(1-\zeta)}} \cdot \left(\left(\frac{\varepsilon(1-\zeta)}{r} \right)^{\frac{\varepsilon(1-\zeta)}{1-\varepsilon(1-\zeta)}} - \left(\varepsilon(1-\zeta) \right)^{\frac{1}{1-\varepsilon(1-\zeta)}} \cdot \left(\frac{1}{r} \right)^{\frac{\varepsilon(1-\zeta)}{1-\varepsilon(1-\zeta)}} \right) \\ &= \underbrace{\left(1 - \varepsilon(1-\zeta) \right) \cdot \left(\frac{\varepsilon(1-\zeta) \cdot z^{\frac{1}{1-\zeta}}}{r} \right)^{\frac{\varepsilon(1-\zeta)}{1-\varepsilon(1-\zeta)}} \cdot \theta^{\frac{1-\varepsilon}{1-\varepsilon(1-\zeta)}}}_{\equiv A(z,\theta,r)} \cdot \underbrace{(nh)^{\frac{\zeta\varepsilon}{1-\varepsilon(1-\zeta)}}}_{\equiv (nh)^{\alpha}} \end{aligned} \tag{A1}$$

The reduced form of the revenue function is simply given by $R(A, n, h) \equiv A \cdot (nh)^{\alpha}$. From equation (A1) it is clear the profitability level, A, encompasses changes to productivity (z), demand (θ), and the cost of capital (r). It also includes various parameters of the production and demand functions. The curvature, on the other hand, corresponds to $\alpha \equiv \zeta\varepsilon / \left(1 - \varepsilon(1-\zeta)\right)$. As pointed out in the main text, it embeds both the labour elasticity in the production function (ζ), and the price elasticity of demand ($\varepsilon \equiv (\nu - 1)/\nu$).

A.2 Aggregate measure of markups

Let $C(y)$ be the total cost function of the firm, such that profits can be written as follows:

$$\pi(y) \equiv y \cdot p(y) - C(y) = y \cdot \left(\frac{\theta}{y} \right)^{1/\nu} - C(y)$$

A simple FOC with respect to y yields the following:

$$\left(\frac{\theta}{y}\right)^{1/\nu} \times \left(\frac{\nu-1}{\nu}\right) = \frac{\partial C(y)}{\partial y} \tag{A2}$$

All three components of condition (A2) above are known. The one on the right-hand side is the marginal cost, $mc(y) \equiv \partial C(y)/\partial y$. The product of the left-hand side comprises the price schedule faced by firms, $p(y)$, and a composite parameter I had already labelled $\varepsilon \equiv (\nu-1)/\nu$. The associated markup, μ, is the ratio between price and marginal cost, which then corresponds to the following:

$$\mu \equiv \frac{p(y)}{mc(y)} = \frac{\nu}{\nu-1} = \varepsilon^{-1} \tag{A3}$$

A.3 Disentangling the fall in α

I showed in equation (A1) that the curvature of the revenue function is pinned down by two primitive parameters: the output elasticity of labour (ζ), and the price elasticity of demand ($\varepsilon \equiv (\nu-1)/\nu$).

$$\alpha \equiv \frac{\zeta\,\varepsilon}{1-\varepsilon(1-\zeta)} = \frac{\zeta}{\mu+\zeta-1}, \tag{A4}$$

where the equality comes from the definition of the aggregate markup displayed in (A3). Given there is a one-to-one mapping between the price elasticity of demand and the aggregate markup, changes in the curvature α ultimately reflect changes in ζ and/or μ. Using simple derivatives, one can easily show that the curvature is a decreasing function of the aggregate markup: $\partial\alpha/\partial\mu < 0$.

In Table 2 I showed the estimated fall in the curvature was mostly driven by firms' growing market power, measured by an increase in the aggregate markup from 22% to almost 40%. But this conclusion builds upon a very specific set of assumptions: with constant returns to scale technology and labour adjustment costs, averaging the labour share of revenue across producers and time recovers the labour elasticity in the production function. A detailed explanation on the topic can be found in De Loecker and Syverson (2021).

An immediate concern is whether the assumption on constant returns to scale is an adequate one. I resort to the literature to show that indeed it is. And to corroborate the conclusion that firms do exhibit greater product market power, I carry out two additional exercises. First, I show that using an employment-weighted mean to average out firm-level labour shares yields the same conclusion. And second, I document a rightward shift in the distribution of firm-specific markups, which are uncovered from the data via the method laid out in De Loecker and Warzynski (2012).

Constant returns to scale

The assumption about constant returns to scale technology requires that a production function be estimated. This can be found in Santos et al. (2022). Although their main goal is to study the cyclicality of markups in the

Portuguese economy, in the process they estimate (industry-wide) Cobb-Douglas and translog production functions using a large firm-product panel of firms between 2004 and 2014. Some of their results are reported in Table A1 below. More details can be found in the original paper.

Table A1: Estimates and test for constant returns to scale in Santos et al. (2022)

2 digit industry	Cobb-Douglas specification		Translog specification	
	Degree of RtS	CRS (p-value)	Degree of RtS	CRS (p-value)
Food	0.97	0.006	1.02	0.023
Beverages	0.80	0.001	0.84	0.892
Textiles	0.99	0.739	1.02	0.367
Apparel	1.10	0.001	1.06	0.695
Leather products	0.97	0.017	1.00	0.119
Wood products	1.04	0.018	1.07	0.211
Paper and pulp	1.00	0.988	1.07	0.784
Printing	1.02	0.311	1.04	0.007
Chemicals	1.03	0.166	1.05	0.591
Rubber and plastics	1.03	0.023	1.04	0.124
Other non-metal minerals	0.95	0.262	1.00	0.829
Basic metals	0.97	0.136	1.04	0.271
Manufactured metal products	1.00	0.913	1.03	0.000
Electrical equipment	1.01	0.738	0.99	0.630
Machinery	1.02	0.830	1.23	0.071
Motor vehicles	0.94	0.000	0.98	0.120
Furniture	1.05	0.000	1.06	0.373
Other manufacturing activities	0.68	0.000	0.96	0.527

Notes: *"Degree of RtS"* corresponds to the sum of all elasticity coefficients in the production function. Inputs in the production function are capital, labour, and materials. *"CRS (p-value)"* is the p-value of a test whose null hypothesis posits that sum equals the unity. All values are taken from Tables 3, C.1, and C.2 in Santos et al. (2022): respectively, pp. 1633 in the main text, and 24 and 25 in the appendix.

The existence of constant returns to scale technology is not a far-fetched premise. Under Cobb-Douglas, the sum of coefficients in the production functions ranges between 0.68 (*Other manufacturing activities*) and 1.10 (*Apparel*). The null hypothesis that the sum equals 1 is not rejected at 10% in half of the selected industries. This figure is even greater when a translog specification is assumed. In that case, the authors do not reject the existence of constant returns to scale in fourteen out of the eighteen industries.

Employment-weighted mean of the labour share

With constant returns to scale and *absent labour adjustment costs*, the output elasticity of labour matches the ratio between the wage bill and revenues. This is the labour share, *i.e.* $s_{i,t} \equiv W_{i,t}/R_{i,t}$, and the equivalence can be proved with the FOC of a static cost minimisation problem. But costs of adjusting employment are very much present in the Portuguese economy. To circumvent this issue, De Loecker and Syverson (2021) suggest averaging the labour share across producers and time:

$$\zeta = T^{-1} \sum_{t=1}^{T} \left(I_t^{-1} \sum_{i=1}^{I_t} s_{i,t} \right)$$

This is the approach followed in Table 2 of the main text. In Table A2 I follow a similar approach but with a caveat: I take a *weighted* mean of firms' labour shares. More formally:

$$\zeta = T^{-1} \sum_{t=1}^{T} \left(I_t^{-1} \sum_{i=1}^{I_t} s_{i,t} \cdot \frac{n_{i,t}}{N_t} \right)$$

The results are shown in Table A2 below and are consistent with those reported in the main text. This alternative approach places more weight on larger firms and, as a result, the average level of the output elasticity of labour increases: 0.65 in 2006-2013 instead of 0.29, and 0.74 instead of 0.32 in 2014-2019. More importantly, rising market power exerted by firms remains the culprit behind the fall in the curvature of the revenue function. First, because the estimated change in the aggregate markup is far larger than the one observed in the output elasticity of labour. As demand becomes more inelastic, the associated markup increases from 50% to 93%. And second, because in a counterfactual economy where the labour elasticity ζ would have been kept unchanged and only the aggregate markup would have been allowed to vary, the resulting curvature would have been quite similar to the actual one, *i.e.* just 7.34% smaller.

A firm-level measure of markups

According to Table 2 the aggregate measure of the markup increased from 22% in 2006-2013 to nearly 40% in 2014-2019. I read this result as proof of the increasing product market power exerted by firms. However, this particular measure is admittedly difficult to interpret as it encodes a whole distribution of firm-level markups.

I now address this issue. More specifically, I follow the seminal work of De Loecker and Warzynski (2012) to identify those firm-level markups and track how the underlying distribution evolved throughout time. At its core, this method posits that markups can be retrieved as the wedge between an input's output elasticity and its revenue share. But not any input. Consider a firm which engages in cost minimisation and uses an input m which satisfies the following assumptions:

Assumption I: There are no adjustment costs for input m.

Table A2: Disentangled changes in the estimated curvature of the revenue function (weighted mean)

	Pre-reform: 2006-2013	**Post-reform: 2014-2019**
Curvature, α	0.563	0.442
Output elasticity of labour, ζ	0.649	0.741
Implied elasticity of demand, $\nu \equiv \lvert(1-\varepsilon)^{-1}\rvert$	2.988	2.071
Implied markup, $\mu \equiv \varepsilon^{-1}$	1.503	1.934
Counterfactual α: constant labour elasticity		-7.339%
Counterfactual α: constant markup		$+34.65\%$

Notes: The curvature of the revenue function corresponds to $\alpha \equiv \frac{\zeta\varepsilon}{1-\varepsilon(1-\zeta)}$. The labour share of a firm is computed as the ratio between the wage bill and its revenues. The aggregate labour elasticity ζ is obtained by taking the employment-weighted mean of (1% winsorized) firm-level values and then the time-series average across the corresponding time window. The number of observations in each time window is the same as reported in Table 1.

Assumption II: Input m is chosen statically.

Assumption III: Input m is not subject to monopsony forces.

Assumption IV: Input m is used for the production of output only.

If assumptions I-IV do hold, then the firm-level markup can be computed as follows:

$$\mu_{i,x,t} = \frac{\beta^m_{x,t}}{s^m_{i,x,t}}, \tag{A5}$$

where $\beta^m_{x,t}$ is the output elasticity of input m in sector x period t, and $s^m_{i,x,t}$ is the revenue share of firm i in sector x in period t. For what follows, I take intermediate materials to take on the role of input m. Two key questions are immediately raised concerning identification. The first one is whether those intermediates do lack adjustment costs and monopsony power. Yeh et al. (2022) provide such discussion in the context of the US. The second one relates to the estimation of $\beta^m_{x,t}$. Unlike $s^m_{i,x,t}$, this is an object which cannot be read straight off the data. To do so I rely on Ackerberg et al. (2015) to estimate revenue functions at the 2-digit sectoral level, and which take labour, capital, and intermediate inputs as arguments. Although a more detailed explanation can be found in Appendix B, this procedure comes at the expenses of the following assumption:

Assumption V: The demand function for input m is strictly monotone in unobserved productivity.

Given the estimated values of $\beta^m_{x,t}$ and the observed values of $s^m_{i,x,t}$ one can compute the product markup and follow its distribution over time. Some descriptive statistics are described in Table A3 below.

The distribution of firm-level markups is consistent with the point made in the main text: the post-reform period of 2014-2019 saw an increase in the market power exerted by firms relative to the pre-reform years of

Table A3: Distribution of firm-level markups

	Mean	Std. deviation	p25	p50	p75	Observations
2006-2013	2.58	3.89	0.87	1.21	2.05	1 627 782
2014-2019	2.73	4.17	0.98	1.28	2.04	1 255 966

Notes: The firm-level markup is computed as $\mu_{i,x,t} = \beta^m_{x,t}/s^m_{i,x,t}$, where $\beta^m_{x,t}$ is the revenue elasticity of materials in a 2-digit sector x, and $s^m_{i,x,t}$ is the corresponding revenue share of a firm i in that sector. Both the revenue share of materials and the markup have been winsorized at the 5% tails.

2006-2013.The median markup shifted from 21% to 28%. Interestingly, these values are in the neighbourhood of the estimated aggregate markup.

A final remark on this exercise is in order. The markup estimator defined in equation (A5) requires that a measure of the *output* elasticity of input m be used. However, physical output is not observed in the dataset. Instead of a production function, I am using Ackerberg et al. (2015) to actually identify a *revenue* elasticity of input m. Bond et al. (2021) claim that estimated markups built on deflated revenues, as is the case here, must inevitably equal unity. If they do not, as is also the case, then one of assumptions I-IV is not held. Their critique implies that my markup estimates above are biased. The point is taken but does not alter the main message I want to convey. This procedure is not meant to accurately identify the values of markups in the Portuguese economy but simply to document a trend: market power exerted by firms has risen in the second half of the sample. According to De Ridder et al. (2024), revenue-based estimates can still capture that dynamics.

A.4 Sectoral decomposition of the Portuguese economy

In the preceding subsection I corroborated the conclusion that the fall in α could be essentially attributed to an increasing market power of firms. There is, however, an immediate objection to be made. Suppose there is a sector x where one single firms exerts monopoly power. If that sector has become more prominent in the Portuguese economy, the estimated fall in the curvature of the revenue function will actually be the outcome of a composition effect. According to this reasoning there would be no change in market power, but actually an endogenous reshuffle of employment towards firms which already had higher markups in the first place. I now prove this is not the case.

An initial piece of evidence is the evolution of employment shares. These are shown in Table A4 below. Between 2008 and 2019 the ranking of the five largest sectors has remained unchanged: Manufacturing, Wholesale and retail trade, Construction, Administrative and supportive services, and Accommodation and food services jointly account for roughly three quarters of all employment in non-financial firms. The only noticeable change concerns the construction sector. Before the detrimental effects of the GFC kicked in, approximately 13.7% of all employees worked in construction. Eleven years later, that figure only amounted to 9.16%. Could this particular evolution explain the fall in α?

Table A4: Sectoral employment shares (%)

2-digit industry	**2008**	**2019**
Manufacturing	25.37	22.39
Wholesale and retail trade	21.93	20.89
Construction	13.73	9.16
Administrative and supportive services	9.69	11.23
Accommodation and food services	7.23	9.23
Transportation and warehousing	5.59	5.48
Consulting, scientific and technical services	4.33	5.79
Information and communication services	2.34	3.57
Health care and social assistance	2.30	3.56
Agriculture, forestry, hunting and fishing	1.54	2.43
Educational services	1.35	1.35
Real estate	1.06	1.39
Other services	1.02	0.99
Water, sewage and waste management	0.89	0.94
Arts, entertainment and recreation	0.62	0.97
Extractive industries	0.45	0.28
Electricity	0.35	0.33
Finance and insurance	0.16	0.02
Public administration	0.04	0.00
Production activities for self-consumption	0.00	0.00
Observations	274 485	327 680

To address this question I estimate sectoral revenue functions. More specifically, I repeat the estimation exercise of equation (1), but this time I do it for each of the five largest sectors in the economy. The goal is to understand whether these sector-specific curvatures mimic the behaviour of the aggregate one.The results are displayed in Table A5 below. The fall is observed in all five sectors. And just as interestingly, all those drops have similar orders of magnitude. According to Table 1 in the main text, the curvature of the revenue function for the whole economy dropped 21%, from 0.56 to 0.44. For the sectoral ones that value ranges between 22% (Administrative and supportive services) and 30% (Accommodation and food services).

Table A5: Sectoral estimation of the curvature

2-digit industry	**Employment share (%)**	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Manufacturing	24.30	0.653***	0.495***
		(0.017)	(0.020)
Wholesale and retail trade	21.93	0.648***	0.467***
		(0.012)	(0.012)
Construction	10.90	0.969***	0.706***
		(0.018)	(0.026)
Administrative and supportive services	10.40	0.611***	0.474***
		(0.028)	(0.028)
Accommodation and food services	7.97	1.090***	0.767***
		(0.022)	(0.021)

Notes: Parameter α is estimated directly from equation (1). The corresponding standard errors are robust. The number of observations for the full period, pre-reform, and post-reform are respectively as follows: *Manufacturing* has 480 369, 204 430, and 141 962; *Wholesale and retail trade* has 1 124 121, 404 030, and 288 223; *Construction* has 444 120, 178 376, and 110 635; *Administrative and supportive services* has 135 198, 41 684, and 35 069; and *Accommodation and food services* has 386 467, 136 391, and 113 758. The employment share corresponds to the average throughout the 2006-2019 period.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

B Additional empirical facts

B.1 Dataset

The dataset is the Central Balance Sheet Harmonized Panel (CBHP), maintained by the Banco de Portugal Microdata Research Laboratory (2025). It contains yearly data since 2006 on the numerous elements of non-financial firms' balance sheets and income statements. The data are collected through *Informação Empresarial Simplificada*, a declaration firms are required by law to submit to the national tax authority until July 15th every year. Given its mandatory nature, the coverage of the data is impressive: more than 600 000 unique firms across the entire sample. Table B1 shows this and other relevant statistics.

Table B1: Summary of statistics

		Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Observations		3 983 356	2 182 179	1 801 177
Nr. unique firms		616 441	445 447	435 628
Employees	Mean	9.41	9.56	9.22
	p25	1.00	2.00	1.00
	p50	3.00	3.00	3.00
	p75	6.00	6.00	6.00
Log revenue pw	Mean	10.60	10.63	10.57
	Std. deviation	1.15	1.14	1.16
Exit rate (%)	Mean	7.42	8.16	6.44

The table reiterates some of the points which had already been made. First, the distribution of labour productivity, as measured by the natural logarithm of revenue per worker, remained largely unchanged. Both the mean and the standard deviation do not exhibit any noticeable swings. Second, the same can be said about the employment-size distribution. Notwithstanding the changes in the labour law, the average sized Portuguese firm employs 9 workers, while the one at the median has 6.

B.2 Robustness checks on employment policy function

The excessive inaction

The pervasiveness of employment inaction had already been documented in Varejão and Portugal (2007). An additional and very important feature of the Portuguese labour market concerns the widespread usage of collective bargain agreements. According to Card and Cardoso (2022), 87% of all private sector full-time workers were covered by one of those — see also Cardoso and Portugal (2005) and Addison et al. (2017). These agreements can regulate various aspects of the employment relationship, chiefly amongst them the minimum wage level for the different job titles.

This constraint on wage adjustments naturally compounds the effect exerted by the non-convexities studied in the main text. As such, the exceedingly large values of inaction reported in Figure 3 may actually indicate that more than one lag should be included in the employment policy function. I now consider the alternative specification:

$$\log(n_{i,t}) = \tilde{\lambda}_1 \log(n_{i,t-1}) + \tilde{\lambda}_2 \log(n_{i,t-2}) + \tilde{\lambda}_3 \log(\xi_{i,t}) + \epsilon_{i,t} \tag{B1}$$

Under this formulation, firms include in their state space not just the stock of workers from the preceding period but also from the one before. Responsiveness is now measured by $\partial \log(n_{i,t}) / \partial \log(\xi_{i,t}) = \tilde{\lambda}_3$. The estimates are shown in Table B2 below.

Table B2: Log linear approximation of employment policy function (two lags)

	Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Serial correlation order 1, $\tilde{\lambda}_1$	0.913^{***}	0.911^{***}	0.913^{***}
	(0.012)	(0.012)	(0.012)
Serial correlation order 2, $\tilde{\lambda}_2$	0.056^{***}	0.053^{***}	0.062^{***}
	(0.011)	(0.011)	(0.011)
Responsiveness, $\tilde{\lambda}_3$	0.017^{**}	0.008	0.034^{***}
	(0.008)	(0.007)	(0.008)
Industry & Year FE	Yes	Yes	Yes
R^2	0.966	0.966	0.967
Observations	2 761 799	1 325 349	985 817

Notes: Standard errors are clustered at the industry level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

The main conclusions remain largely unchanged. First, although the introduction of a second lag of employment naturally lowers the serial correlation of order 1, its estimated values are still exceedingly high: above 0.91 in the two time windows and not statistically different from each other. And second, there is an unequivocal increase in the employment responsiveness of firms. In the main formulation explored in Table 3 I showed that responsiveness nearly doubled after the reform. When accounting for additional lags of employment, that jump is actually fourfold.

Non-flexible capital

A contentious assumption concerning the identification of the unobserved profitability innovations in (1) is the one about the absence of frictions in capital adjustments. Since that identification scheme ultimately impacts how the responsiveness coefficient λ_2 is estimated, I now turn to this issue.

In Appendix A I showed how the existence of fully flexible capital allows to express the revenue function solely in terms of labour demand. Back then I argued that as long as the assumption of full flexibility remains, one can develop a FOC with respect to capital and plug it back in the revenue function. Let me now relax this assumption. Basically, say capital is not a static choice, *i.e.* it cannot be freely adjusted. This is consistent with the existence of time to build and adjustment costs. I remain agnostic about the potential causes; the main point is that the revenue function can no longer be simplified to the same expression. As a consequence, the profitability innovations must be identified with an alternative identification scheme which can account for the dynamic nature embedded in capital choices.

I do so via the procedure in Ackerberg et al. (2015). Suppose the revenue function of firm i in period t reads as follows:

$$\log(R_{i,t}) = \log(\tilde{A}_{i,t}) + \alpha_n \log(n_{i,t}) + \alpha_k \log(k_{i,t}) + \alpha_m \log(m_{i,t}) + \epsilon_{i,t} \tag{B2}$$

I addition to labour ($n_{i,t}$) and capital ($k_{i,t}$), I also consider intermediate inputs ($m_{i,t}$) as an argument of the revenue function. The profitability level $\log(\tilde{A}_{i,t})$ is assumed to follow a similar AR(1) process to the one described in the main text: the serial correlation is $|\tilde{\rho}| < 1$ and the underlying i.i.d. gaussian innovations $\log(\tilde{\xi}_{i,t})$. The usage of the tilde in these three objects is meant to reinforce the idea that this profitability process is identified with a different procedure of that of equation (1). There are three additional assumptions which must be emphasised. First, the choice of hours was dropped for convenience. Second, identifying the various elasticities in (B2) requires that **Assumption V** defined in Appendix A be held: the demand function for input intermediate inputs must be strictly monotone in the unobserved profitability level. And third, the estimation is ultimately carried out through GMM-IV. The vector of instruments is $\left[\log(k_{i,t})\ \log(n_{i,t-1})\ \log(m_{i,t-1})\ \hat{\Phi}_{t-1}(k_{i,t-1}, n_{i,t-1}, m_{i,t-1})\right]'$, where $\hat{\Phi}_{t-1}(\cdot)$ corresponds to the lag of the fitted values of revenue on a polynomial of order 3 in capital, employment, and intermediate inputs. A full set of year dummies also controls for the existence of aggregate shocks.

This estimation procedure allows to obtain estimates of all elasticities in (B2), as well as a non-parametric estimate of $\tilde{\rho}$. Ultimately, all these objects are sufficient to back out the series of unobserved profitability innovations without the need to appeal to full flexibility of capital demand. A slightly modified version of the employment policy function can then be set up as follows:

$$\log(n_{i,t}) = \hat{\lambda}_1 \log(n_{i,t-1}) + \hat{\lambda}_2 \log(k_{i,t-1}) + \hat{\lambda}_3 \log(\tilde{\xi}_{i,t}) + \epsilon_{i,t} \tag{B3}$$

Table B3: Log linear approximation of modified employment policy function

	Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Lag employment, $\hat{\lambda}_1$	0.942***	0.935***	0.951***
	(0.005)	(0.006)	(0.004)
Lag capital, $\hat{\lambda}_2$	0.013***	0.015***	0.011***
	(0.002)	(0.002)	(0.002)
Responsiveness, $\hat{\lambda}_3$	0.050	0.074	0.126***
	(0.061)	(0.053)	(0.031)
Industry & Year FE	Yes	Yes	Yes
R^2	0.971	0.969	0.973
Observations	2 082 050	1 148 154	781 895

Notes: Standard errors are clustered at the industry level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

Table B3 above displays its estimates. The results are again aligned with those reported in Table 3. First, the exceedingly high values of $\hat{\lambda}_1$ show that employment is indeed a very persistent phenomenon. Second, the inclusion of the lagged stock of capital[10] does not significantly improve an already large value of R^2. And lastly, despite the larger confidence bands, this alternative specification still predicts a sizeable increase in employment responsiveness.

Non-linearities

As a final robustness exercise I also add non-linear components to the employment policy function. The general specification reads as follows:

$$\log(n_{i,t}) = \check{\lambda}_1 \log(n_{i,t-1}) + \check{\lambda}_2 \log(\xi_{i,t}) + \check{\lambda}_3 \cdot \pi_{i,t} + \epsilon_{i,t}, \tag{B4}$$

The focus lies particularly on the third regressor, $\pi_{i,t}$, which is considered to take two different forms:

$$\pi_{i,t} = \begin{cases} \log(\xi^2_{i,t}) & \textbf{specification I} \\ \mathbb{I} \cdot \big(\log(\xi_{i,t}) < 0\big) & \textbf{specification II} \end{cases}$$

[10]If I had instead considered that the state-space of the firm includes the *contemporaneous* stock of capital, and not its *lagged* one, the results would have been very similar.

For **specification I** I introduce the quadratic component of the profitability innovation. This formulation is identical to that employed in the responsiveness regressions of Section 3. The results are shown in Table B4 below. The policy function is concave in the profitability innovations, $\check{\lambda}_3 < 0$, both in the full sample and in the pre-reform period. Yet that concavity vanishes in the post-reform period. This increasing linearity, which had already been found at the intensive margin of employment adjustment, is also consistent with growing responsiveness. First, because the estimated value of $\check{\lambda}_3$ was initially negative and is now statistically insignificant. And second, because that of $\check{\lambda}_2$ more than doubles: it rises from 0.023 to 0.048.

Table B4: Log linear approximation of employment policy function: **specification I**

	Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Serial correlation, $\check{\lambda}_1$	0.984^{***}	0.975^{***}	0.995^{***}
	(0.003)	(0.004)	(0.002)
Responsiveness, $\check{\lambda}_2$	0.030^{***}	0.023^{***}	0.048^{***}
	(0.008)	(0.008)	(0.010)
Quadratic, $\check{\lambda}_3$	-0.003^{**}	-0.006^{***}	0.002
	(0.001)	(0.001)	(0.001)
Industry & Year FE	Yes	Yes	Yes
R^2	0.958	0.958	0.959
Observations	3 308 659	1 709 959	1 350 617

Notes: Standard errors are clustered at the industry level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

For **Specification II**, on the other hand, I attempt to capture potentially different slopes for positive and negative profitability innovations. This replicates one of the exercises reported in Ilut et al. (2018) and aims to study the existence of asymmetries in the employment policy function. Do firms react more to bad shocks than to positive ones? According to Table B5 below they do. This is wholly encapsulated in the persistently positive values of $\check{\lambda}_3$. Notice, however, that asymmetry is attenuated in the post-reform period. First, because the increase in the estimated value of $\check{\lambda}_2$ implies that firms respond more aggressively to whatever type of shocks hit them. And second, because the additional component of responsiveness which is found in the presence of negative shocks dwindles — there is a 20% fall in the estimated value of $\check{\lambda}_3$.

The inclusion of these non-linearities do not alter the main results conveyed in Section 3. In both cases, the coefficient attached to lagged employment remains high and very close to unity, while that associated with responsiveness invariably increases in the post-reform period. It is also convenient to point out that neither specification improves upon the baseline one expressed in equation (2). This happens mostly because that

Table B5: Log linear approximation of employment policy function: **specification II**

	Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Serial correlation, $\check{\lambda}_1$	0.955^{***}	0.951^{***}	0.961^{***}
	(0.005)	(0.006)	(0.005)
Responsiveness, $\check{\lambda}_2$	0.048^{***}	0.045^{***}	0.061^{***}
	(0.007)	(0.007)	(0.010)
Bad news, $\check{\lambda}_3$	0.028^{***}	0.035^{***}	0.021^{**}
	(0.009)	(0.011)	(0.009)
Industry & Year FE	Yes	Yes	Yes
R^2	0.959	0.958	0.959
Observations	3 308 659	1 709 959	1 350 617

Notes: Standard errors are clustered at the industry level. $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

very simple and fully linear specification already encompassed an exceedingly large value of R^2.

B.3 Responsiveness regressions

Equations (3) and (4) set up the extensive and intensive margin decisions of employment adjustment. I estimated both of them and plotted their projected values against a very fine grid of log profitability innovations — see Figure 4. I now explicitly report the values of the estimated coefficients. The results for the extensive margin are shown in Table B6, while Table B7 displays those of the intensive margin. To explicitly assess the existence of potential asymmetries, in each of them I evaluate the firms' response to a ± 1 standard deviation profitability innovation. Basically, I evaluate how the probability of employment adjustment (Table B6) and the actual adjustment (Table B7) change following both a positive and a negative one-standard deviation shock to profitability. The row labeled "*Different responses?*" investigates whether those responses are statistically different from each other.

Start with the extensive margin decision. The convexity of the employment adjustment hazard is justified by the constantly positive values of δ_2. Although these remain largely unchanged throughout the various time periods, those of δ_1 do not. Its estimated value actually increases by 50% following the labour market interventions, from 0.014 to 0.021. The joint evolution of these two parameters is indicative of a rising, yet asymmetric, responsiveness to profitability innovations. To better understand that asymmetry, consider a firm which is hit by a positive one-standard deviation shock. In the pre-reform period this would raise the probability of adjustment by approximately 0.91 percentage points; post-reform, that likelihood would

Table B6: Extensive margin - probability of employment adjustment

	Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Linear, δ_1	0.015^{***}	0.014^{***}	0.021^{***}
	(0.003)	(0.003)	(0.003)
Quadratic, δ_2	0.012^{***}	0.012^{***}	0.012^{***}
	(0.003)	(0.003)	(0.003)
Lag employment, δ_3	0.176^{***}	0.175^{***}	0.178^{***}
	(0.009)	(0.010)	(0.009)
Positive response	1.143 p.p.	0.907 p.p.	1.903 p.p.
Negative response	1.074 p.p.	1.342 p.p.	0.320 p.p.
Different responses?	Yes***	Yes***	Yes***
Industry & Year FE	Yes	Yes	Yes
Observations	3 308 659	1 709 959	1 350 617
R^2	0.168	0.165	0.172

Notes: Standard errors are clustered at the industry level. *Positive response* and *Negative response* evaluate the percentage point change in probability of adjustment following a positive and negative one-standard deviation shock to profitability, respectively. The row entitled *Different responses?* indicates if those responses are statistically different at various levels of significance.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

increase by 1.9 percentage points. This stands in clear contrast with the probability of adjustment in the presence of a similar shock but of the opposite sign: after 2013, firms are less likely to adjust to a negative shock to their profitability.

These results can be interpreted to a great extent in light of the duality introduced by the reform. As the more flexible provisions introduced in the labour law are grandfathered in, *i.e.* they only apply to new contracts set from that moment on, firms internalise the fact that new hirings will be bound by looser constraints. This has an immediate effect on job creation but not necessarily on job destruction. Upon the realisation of a sufficiently positive profitability shock, firms are more likely to raise their labour force because those new workers will be protected by increasingly flexible rules.

I now turn to the intensive margin. The concavity of the hiring rule is mirrored in the negative values of θ_2. Much like the quadratic component in the formulation of the extensive margin, its coefficient also remains fairly stable throughout the two periods: the slight increase from -0.02 to -0.017 can only go so far to dampen that concavity. In fact, the increasing linearity of the hiring rule is mostly driven by the coefficient θ_1, which more than doubles in the post-reform period.

Table B7: Intensive margin - employment growth of adjusters

	Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Intercept, θ_0	0.250^{***}	0.268^{***}	0.235^{***}
	(0.024)	(0.026)	(0.012)
Linear, θ_1	0.065^{***}	0.048^{***}	0.104^{***}
	(0.013)	(0.012)	(0.015)
Quadratic, θ_2	-0.020^{***}	-0.020^{***}	-0.017^{***}
	(0.001)	(0.001)	(0.001)
Lag employment, θ_3	-0.104^{***}	-0.097^{***}	-0.111^{***}
	(0.011)	(0.011)	(0.011)
Positive response	3.249 p.p.	2.278 p.p.	5.592 p.p.
Negative response	-4.710 p.p.	-3.844 p.p.	-6.725 p.p.
Different responses?	Yes***	Yes***	Yes***
Industry & Year FE	Yes	Yes	Yes
Observations	1 380 478	744 067	538 526
R^2	0.075	0.065	0.080

Notes: Standard errors are clustered at the industry level. *Positive response* and *Negative response* evaluate the percentage point change in employment growth following a positive and negative one-standard deviation shock to profitability, respectively. The row entitled *Different responses?* indicates if those responses are statistically different at various levels of significance.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

The rising responsiveness can be attested by the size of the employment adjustments when firms are hit by one-standard deviation shocks. Whether their profitability improves or worsens, in absolute value their response is always greater in the post-reform period: a positive shock can now boost employment growth by 5.6 percentage points (only 2.3 in the pre-reform period), whereas a negative one reduces growth by 6.7 percentage points (only 3.8 before). The asymmetry is not as striking as that associated with the extensive margin decision. To formally address this issue, I run the following specification:

$$\varkappa_{i,t\{\mathbb{1}_{i,t}=1\}} = \tilde{\theta}_0 + \tilde{\theta}_1 \log(\xi_{i,t}) + \tilde{\theta}_2 \cdot \Big(\mathbb{I} \cdot \big(\log(\xi_{i,t}) < 0\big)\Big) + \tilde{\theta}_3 \log(n_{i,t-1}) + \epsilon_{i,t} \tag{B5}$$

This is one of hiring rules specified in Ilut et al. (2018). If the parameter $\tilde{\theta}_2$ is estimated to be statistically insignificant, then this corroborates the absence of an asymmetry in employment growth amongst adjusting firms. The results are displayed in Table B8 below. The asymmetry does exit and was particularly sizeable in the pre-reform period.

Table B8: Intensive margin of adjustment - employment growth of adjusters conditional on sign

	Full sample: 2006-2019	**Pre-reform: 2006-2013**	**Post-reform: 2014-2019**
Intercept, $\tilde{\theta}_0$	0.206^{***}	0.218^{***}	0.201^{***}
	(0.026)	(0.029)	(0.015)
Linear responsiveness, $\tilde{\theta}_1$	0.093^{***}	0.081^{***}	0.123^{***}
	(0.012)	(0.012)	(0.014)
Bad news, $\tilde{\theta}_2$	0.056^{**}	0.068^{***}	0.039^{**}
	(0.021)	(0.023)	(0.017)
Lag employment, $\tilde{\theta}_3$	-0.098^{***}	-0.089^{***}	-0.106^{***}
	(0.010)	(0.010)	(0.011)
Positive response	5.593 p.p.	4.892 p.p.	7.329 p.p.
Negative response	-8.982 p.p.	-8.981 p.p.	-9.631 p.p.
Different responses?	Yes**	Yes***	Yes**
Industry & Year FE	Yes	Yes	Yes
Observations	1 380 478	744 067	538 526
R^2	0.073	0.064	0.078

Notes: Standard errors are clustered at the industry level. *Positive response* and *Negative response* evaluate the percentage point change in employment growth following a positive and negative one-standard deviation shock to profitability, respectively. The row entitled *Different responses?* indicates if those responses are statistically different at various levels of significance.

$^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$

C Additional quantitative results

C.1 Alternative parameterisation

In the baseline estimation procedure described in Section 5 I used the conforming identity matrix to weight the distances between data and simulated moments. I now report something different. Following Adda and Cooper (2003) the optimal weights are provided by the elements of the inverted variance-covariance matrix of the simulated moments. This requires that a two-stage SMM procedure be followed.

Stage 1: I solve the problem in (10) using the conforming identity matrix as weights. Let the resulting vectors of parameter estimates be labelled $\hat{\Theta}_1$.

Stage 2: I simulate 1000 economies parameterised by $\hat{\Theta}_1$. For each of those simulated economies I record the vector of targeted moments. Eventually, I obtain the variance-covariance matrix of all those

1000 vectors. I invert that matrix and use it as a weight while solving (10) a second time. The vector of estimated parameters is then $\hat{\Theta}_2$.

Although the identification scheme which underlies this two-stage procedure is the same as that employed in the main text, the usage of the optimal weighting matrix allows to downweight moments which are generated less precisely in the simulation. Naturally, this has an impact on the vector of estimated parameters. For example, in Figure 5 I showed that changes in the average exit rate of the economy exert a sizeable influence on the cost of operation, c. Suppose that exit would exhibit a significant variance across the 1000 simulations. The optimal weighting matrix would attach a smaller weight to deviations between the simulated exit rate and its data counterpart, and the estimation of parameter c would inevitably be affected.

Table C1 below show that, albeit different, the vector of estimated parameters remains largely unaltered relative to the baseline case. Importantly, this alternative parameterisation is still consistent with a significant reduction in firing costs. Prior to the reform, sacking one single employee would amount to the entire average revenue in the economy. This again entails both the fixed and linear component of the firing cost. In the period after the reform, that same cost would be equivalent to 64% of the average revenue.

Table C1: Parameter estimates in a 2-stage SMM procedure

	Labour adjustment costs				Wage function		Fixed cost
	γ_f	Γ_f	γ_h	Γ_h	η	w_0	c
2006-2013	0.27	0.73	0.22	0.13	4.19	1.36	0.04
2014-2019	0.20	0.44	0.29	0.10	4.16	0.73	0.06

Notes: All labour adjustment costs parameters and the fixed cost of operation are expressed as a share of the average revenue of the simulated economy. The point estimates are obtained via a two-stage SMM procedure. The optimal weighting matrix is employed after the economy is simulated 1000 times. All simulations are run at a quarterly frequency and then aggregated to match the annual moments of the data. The quarterly discount factor is set to $\beta = 0.99$, and the hourly rate w_1 is calibrated such that the steady-state value of hours per worker replicates the observed yearly average.

Lastly, Table C2 reports the fitness of the two-stage SMM estimation. Unlike the baseline case, the measure of distance now takes exceedingly larger values. This is inevitable, given that $F(\cdot)$ is now a *weighted* sum of distances, and those weights can indeed take very large values. In fact, the greater they are, the more accurate the associated moment is simulated. This notwithstanding, the estimated model is still capable of getting the vector of selected moments just right. The concavity of the hiring rule and the exit responsiveness, for example, are particularly well targeted, just like the average employment in the economy.

Table C2: SMM estimation fit — targeted moments in a 2-stage SMM procedure

		Responsiveness				Distributions		Exit	Distance
		ϱ_{An}	λ_1	θ_2	ψ_2	$\sqrt{\Sigma_{\bar{r}}}$	$\overline{n}$	$\overline{x}$	$F(\hat{\Theta}_2)$
2006-2013	*Data*	0.64	0.95	-0.02	-0.08	1.14	1.24	8.16	
	Model	0.59	0.81	-0.02	-0.08	1.80	0.98	6.31	13736.03
2014-2019	*Data*	0.81	0.96	-0.02	-0.07	1.16	1.14	6.44	
	Model	0.57	0.85	-0.01	-0.07	1.59	1.21	5.44	3153.2

Notes: Targeted moments are the covariance between the log level of TFPR and log employment ϱ_{An}, the serial correlation of log employment λ_1 in (2), the concavity of the hiring rule θ_2 in (3), the exit responsiveness coefficient ψ_2 in (5), the standard deviation of log labour productivity $\sqrt{\Sigma_{\bar{r}}}$, the mean of log employment $\overline{n}$, and the mean of the exit rate $\overline{x}$. The measure of distance between data and simulated moments defined in (10) is $F(\hat{\Theta}_2)$.

C.2 Proof of variance decomposition

From the reduced form revenue function derived initially, log labour productivity can be expressed as $\log(R) - \log(n) \equiv \log(A) + (\alpha - 1)\log(n) + \alpha \log(h)$. The variance of the sum is a sum of variances and covariances, such that:

$$\begin{aligned}\operatorname{Var}\big(\log(R) - \log(n)\big) &= \operatorname{Var}\big(\log(A)\big) + (\alpha-1)^2 \operatorname{Var}\big(\log(n)\big) + \alpha^2 \operatorname{Var}\big(\log(h)\big) \\ &+ 2\Big((\alpha-1)\operatorname{Cov}\big(\log(A), \log(n)\big) + \alpha \operatorname{Cov}\big(\log(A), \log(h)\big) + \alpha(\alpha-1)\operatorname{Cov}\big(\log(n), \log(h)\big)\Big)\end{aligned} \tag{C1}$$

According to equation (11) in the main text this expression above should correspond to the trace of a particular matrix, *i.e.* the product between a matrix parameterised by the curvature of the revenue function and the variance-covariance matrix between profitability, employment, and average hours. As the focus lies on the trace of that resulting matrix, I can abstract from all elements outside that main diagonal. For convenience of notation, say that logged variables are expressed with tildes, such that $\tilde{x} \equiv \log(x)$:

$$\begin{bmatrix} \operatorname{Var}(\tilde{A}) + 2(\alpha-1)\operatorname{Cov}(\tilde{A},\tilde{n}) + 2\alpha \operatorname{Cov}(\tilde{A},\tilde{h}) & \cdot & \cdot \\ \cdot & (\alpha-1)^2 \operatorname{Var}(\tilde{n}) + 2\alpha(\alpha-1)\operatorname{Cov}(\tilde{n},\tilde{h}) & \cdot \\ \cdot & \cdot & \alpha^2 \operatorname{Var}(\tilde{h}) \end{bmatrix} \tag{C2}$$

It is immediate to realise that the expression shown in equation (C1) matches the sum of all three elements of the matrix in (C2).