Heterogeneous Labour Market Concentration and Minimum Wage Policy *

Tiago Paúl †
London School of Economics

November 10, 2023
Click here for the latest version

Preliminary and incomplete

Abstract

This paper empirically tests whether minimum wages may have positive employment effects when diverging from a perfect competition framework. I use administrative data from 1986 to 2013 to segment Portugal into several local labour markets and construct Herfindahl–Hirschman Indices to measure concentration levels in each market. I find that increased minimum wage exposure has differential employment effects depending on labour market concentration - estimates indicate employment responds negatively in highly competitive markets and positively in highly concentrated ones. These results are robust to the usage of instrumental variables and to different definitions of local labour markets, although the persistence of these effects varies across specifications. Labour markets are highly heterogeneous in their concentration level, which can help reconcile the conflicting results in the minimum wage literature and provide new policy insights.

Keywords— Local Labour Markets, Labour Market Concentration, Monopsony Power, Minimum Wage

JEL Classification: J2, J3, J38, J42

*For valuable supervision and support at the start of this project, I thank Rui Costa. I am also grateful to Ethan Ilzetzki and Liliana Varela for their insightful comments, and especially for their guidance in my early research career. I am, furthermore, thankful to Tiago Bernardino, Seyed Mahdi Hosseini Maasoum, Ricardo Reis, and participants at the 16$^{\text{th}}$ annual meeting of the Portuguese Economic Journal for their helpful comments and suggestions.

†Contact: t.brinkmann-torres@lse.ac.uk

1 Introduction

After decades of research, employment effects of minimum wage increases remain one of the most contentious debates in labour economics literature. In the US, Card and Krueger's 1994 case-study of minimum wage differences in New Jersey and Pennsylvania, followed by their 1995 book *Myth and Measurement*, challenged the established idea in economic theory that minimum wage increases necessarily lead to disemployment effects. Since then, the debate surrounding this issue has found conflicting empirical evidence. Some local (Dube et al., 2007) and national (Dube et al., 2010) studies find similar results showing no negative employment effects, while others authors find significant disemployment effects following increases in the minimum wage, often affecting particular groups (Neumark and Wascher, 1992, 2007).

When results contradict each other and fail to shed light on a particular question, it is often useful to go back to theory and understand what forces may be behind the empirical conundrum. The prediction that minimum wages will necessarily have a negative impact on employment relies heavily on the assumption that labour markets are perfectly competitive. In such a setting, firms participating in the labour market have no influence over the wage rate, and must take the market wage as given. If this is the case, firms maximise profits by employing an amount of labour that equalises the marginal revenue product of labour (MRPL) to the market wage. Under these circumstances, any economically-binding minimum wage will necessary raise the market wage above MRPL, and hence it becomes optimal for a competitive firm to reduce its level of employment. This simple model was behind the widespread belief that increases in the minimum wage must lead to a fall in employment.

Theoretical predictions are, however, completely different once we relax the assumption that the labour market operates competitively and that firms are wage-takers. If we allow for a setting in which firms exert labour market power, increases in the minimum wage may in fact lead to positive changes in employment. It has become imperative to study this hypothesis, as evidence has been amounting that labour markets are indeed best characterised by an oligopsonistic or monopsonistic competition framework (Manning, 2011). This setting in which the firm has market power and is a wage-setter functions in a mirrored fashion of the more commonly known monopolist profit maximisation problem. The wage-setting firm faces an upward sloping labour supply curve. If it decides to employ more labour, then it must pay a higher wage to the new hires as well to all other already employed workers, implying that the marginal cost of labour is a steeper upward sloping curve than labour supply. The monopsonistic firm maximises profits by equalising marginal revenue product of labour to the marginal cost of labour. However, it then only needs to pay the wage rate required to satisfy labour supply at that level, which will lead to a wage rate which is below MRPL. This wedge between the wage rate and MRPL is what allows for positive employment effects upon an increase in the minimum wage. Because the wage rate is below MRPL, there is room to raise the minimum wage while still making it profitable for the monopsonist to meet the increased labour supply, and so employment levels increase. Of course, if the increase in minimum wage is of such large magnitude that attending to the increased labour supply causes it exceed MRPL, then even in the presence of monopsonistic market power it is still possible to have negative employment effects.

Given that theory predicts different effects of the minimum wage on employment depending on the competitive nature of the labour market, it seems logical to study minimum wage effects in the context of labour market competition. This idea is the conceptual focus of Bhaskar et al. (2002). There does not, however, exist a vast literature trying to relate the two empirically. Recently, Azar et al. (2023) have made groundbreaking progress on shedding light on this issue. They use job postings data to measure labour market concentration, and study the differential employment effects of minimum wages as a function of market concentration. Their results seem to confirm theoretical predictions - minimum wages have negative employment effects in highly competitive markets, but these effects are less severe in more concentrated markets. As far as I am aware, they are to date the only ones to empirically study the interaction between minimum wages and market concentration. From my understanding, the reason why this has not been a topic of wider interest is due to the lack of granular enough data for the US. Measuring concentration of a labour market requires knowing firm-level employment shares. Datasets commonly used by the literature, such as the Quarterly Census of Employment and Wages for the US, contain employment information at the county level, which does not allow to measure market concentration. Azar et al. use data on online job vacancy postings by Burning Glass Technologies to overcome this, which allows them to build concentration indices based on the shares of these postings. Even so, they are restricted in their sample to labour markets that post job vacancies overwhelmingly online and that have high exposure to the minimum wage.

This paper, therefore, takes advantage of granular administrative employment data in Portugal to overcome some of the constraints faced by Azar et al.. In choosing to study Portugal, I use the *Quadros de Pessoal* dataset, an extremely rich administrative worker-level matched employer-employee panel covering the universe of private employment in the country. Portugal is also a country in which the minimum wage has a much higher bite than the US. With these two factors, I am able to construct a sample of local labour markets which covers almost the entirety of Portuguese labour, allowing me to conduct a much more comprehensive inspection of how labour market concentration interacts with minimum wages in their employment effects.

While theory is clear on the effects of minimum wages at the firm level, it is not obvious how to think of aggregate labour market employment effects when firms are heterogeneous in their market power. Therefore, in order to bridge the gap between theory and data, I first present and solve a dominant firm model in which labour markets are composed by one wage-setting firm with market power, and multiple wage-taking competitive firms. This approach allows me to aggregate individual firm-level employment responses, which provides predictions I test out empirically. The model also shows that market concentration can be used as a good proxy for market power.

In the empirical part of the paper, I segment Portugal into several local labour markets based on commuting zones and industries. For each local labour market, I measure its exposure to minimum wage increases and its level of market concentration. I regress changes in employment on market exposure to the minimum wage, and the level of concentration of the labour market. Results show that market concentration is indeed crucial in understanding how minimum wages affect employment. In accordance with theoretical predictions, employment in highly competitive markets responds negatively to higher minimum wage exposure, but highly concentrated markets may actually respond positively. These results are robust to different samples and definitions of local labour markets, as well as when instrumenting for market concentration. I then use a local projections specification to study dynamic responses, and find that employment effects are mostly transitory. This latter result is not as robust, as altering the definition of local labour markets can induce persistence in the responses.

These findings provide new insights into how we interpret the existing evidence. Local labour markets are extremely heteregeneous in their concentration (Azar et al. 2022a show this for the US, I show this for Portugal), which implies that local estimates of employment effects are influenced by the local level of market concentration, while national estimates are affected by cancelling out effects. The results also provide some new policy insights, as they imply minimum wages can play a role in correcting a market failure, while the existing heterogeneity in labour market concentration warrants a decentralisation of minimum wage systems.

Related Literature - This paper contributes mainly to two strands of literature. Firstly, it adds to the literature on the employment effects of minimum wages, and particularly to the recent trend of using micro data to exploit heterogeneities that might affect these effects. The minimum wage literature started out with time-series empirical work, which often found null effects on a national scale (see Brown et al. 1982 for a survey). Card and Krueger (1994) was the foundational paper that kick-started the micro work on minimum wages. This launched a multitude of micro studies for several countries on the employment effects of minimum wages (Dube et al., 2007, 2010; Neumark and Wascher, 1992, 2007; Stewart, 2004), in which results were often at odds with each other. To dive deeper into potential explanations for the ambiguous evidence, a third generation of minimum wage studies ensued, which used richer micro data to explore underlying heterogeneities in treated observations. Giuliano (2013) looked at different sub-groups of low-wage earners, Harasztosi and Lindner (2019) analysed different adjustment mechanisms used by firms, and Cengiz et al. (2019) dissected the underlying wage distributions. My paper therefore contributes to this growing strand of the literature by answering whether heterogeneity in concentration across labour markets is also relevant in the study of minimum wages.

Secondly, this paper contributes to the growing literature on labour market concentration. There is work on how market concentration relates to market power (Azar et al., 2019, 2022a) which shows the two are highly correlated. A lot of studies measure the impact of market concentration on wages (Bassanini et al., 2022, 2023; Benmelech et al., 2022; Félix and Portugal, 2016; Marinescu et al., 2021; Martins and Melo, 2023; Rinz et al., 2018), finding a clear negative causal effect. However, the effect of concentration on employment has more limited empirical backing (Bassanini et al., 2022; Marinescu et al., 2021), but also points towards a negative impact. This paper partly contributes to expanding the evidence on this latter effect.

Finally, this paper combines these two strands of the literature by joining Azar et al. (2023) in their pioneering work of explaining the differential employment effects of minimum wages through the lenses of labour market concentration. I follow their work closely, but expand on their coverage and analysis by exploiting the advantageous Portuguese setting.

Roadmap - The remainder of this paper proceeds as follows. Section 2 presents a dominant firm model to further motivate my research question and to introduce some predictions to take to the data. In Section 3, I describe the Portuguese institutional setting and the dataset used. Section 4 explains how the local labour market observational units were constructed and how I measure their concentration and minimum wage exposure. Section 5 presents the empirical strategy and the main results for both contemporaneous and dynamic employment responses, as well as detailing the robustness checks performed. Section 6 concludes.

2 Model

To further motivate my research question, and to provide theoretical predictions to be tested in the empirical part of the paper, I present and solve a model that has minimum wage having differential employment effects depending on the level of concentration of the labour market.

I develop a dominant firm model, in which labour markets are composed by one wage-setting dominant firm, and a wage-taking competitive fringe. Labour markets differ in the extent to which the dominant firm is more productive than a competitive firm. The model achieves similar results to the Cournot competition model used by Azar et al. (2023). However, contrary to Azar et al., my model allows for heterogeneity across firms, which enables me to achieve variation in market concentration while keeping the number of firms fixed. The source of exogenous variation - productivity of the dominant firm - causes both market concentration and market power to be endogenously determined, this latter result also being achieved by Manning and Petrongolo (2022).

2.1 Model Setup

Consider a setting in which there exists a continuum of self-contained local labour markets indexed by $j \in [0, 1]$. In each labour market j, there exists one wage-setting dominant firm D, and a competitive fringe composed of N wage-taking firms each indexed by i.

Production for each competitive firm is a function of labour input L_{ji}, and a decreasing returns to scale parameter $\alpha < 1$ which is assumed equal across all firms and markets. Prices are normalised to 1 so that revenue equals production, which is formally given by

$$Y_{ji} = L_{ji}^{\alpha}. \tag{1}$$

The dominant firm in each labour market has a similar production function, except that it possesses a technology that makes it A_j times more productive than each competitive firm, where $A_j \sim \mathcal{U}(1, \bar{A})$ is uniformly distributed. Production and revenue for the dominant firm is then formally characterised by

$$Y_{jD} = A_j L_{jD}^{\alpha}. \tag{2}$$

Labour supply in each market is a function of the wage rate in that market, and a labour supply elasticity parameter ε, assumed to be equal across markets.

$$L_j = W_j^{\varepsilon}. \tag{3}$$

The wage-taking competitive firms maximise profits by equalising their marginal revenue product of labour (MRPL) to the market wage rate, which gives a labour demand by each firm of

$$L_{ji} = \left(\frac{\alpha}{W_j}\right)^{\frac{1}{1-\alpha}}. \tag{4}$$

The dominant firm is faced with the residual labour supply not absorbed by the competitive fringe, which implies it faces a labour supply function given by

$$\mathbf{L_{jD}} = W_j^{\varepsilon} - N\left(\frac{\alpha}{W_j}\right)^{\frac{1}{1-\alpha}}. \tag{5}$$

The wage rate plays a dual role in the labour supply of the dominant firm. An increase in the wage rate attracts more workers to the market - increasing aggregate labour supply - but it will also have a reinforcing effect of making labour more unaffordable to the competitive fringe, increasing the residual labour supply available to the dominant firm. Conversely, decreasing the wage rate will decrease the number of workers in the market, and it will also make labour more affordable to the competitive fringe, further reducing the workers available to the dominant firm.

The government intervenes in each labour market by setting a legally binding minimum wage equal to W_j^{min}. Under this setting, the dominant firm then optimally chooses the amount of labour input in order to maximise profits, by solving

$$\underset{L_{jD}}{max} \quad \Pi_{jD} = A_j L_{jD}^{\alpha} - \max\left\{\mathbf{L_{jD}^{-1}}(L_{jD}); W_j^{min}\right\} \times L_{jD}, \tag{6}$$

where $\mathbf{L_{jD}^{-1}}$ is the inverse of the labour supply function of the dominant firm defined in Equation (5). The dominant firm's choice of labour will also pin down the wage rate in market j,

$$W_j = \max\left\{\mathbf{L_{jD}^{-1}}(L_{jD}); W_j^{min}\right\}. \tag{7}$$

Given this wage rate, the competitive fringe chooses its own profit-maximising employment following Equation (4), which clears the labour market at an aggregate employment level determined by Equation (3).

2.2 Free Market Equilibrium

To illustrate the dynamics of the model, I first set the minimum wage to a non-binding value so that $W_j = \mathbf{L_{jD}^{-1}}(L_{jD})$ in all markets. This allows me to initially abstain from minimum wage effects. I set $\alpha = 0.7$, $\varepsilon = 1.6$ (Azar et al., 2022b), $\bar{A} = 25$, and $N = 100$ competitive firms for each market.

As a first step, I relate the productivity level of the dominant firm A_j to the concentration level of the labour market. I measure concentration using the Herfindahl-Hirschman Index $\text{HHI}_j = \sum_i \eta_{ji}^2 + \eta_{jD}^2$, where η_{ji} and η_{jD} represent the employment shares of competitive firm i and the dominant firm, respectively.

Figure 1: Market Concentration as a Function of A_j

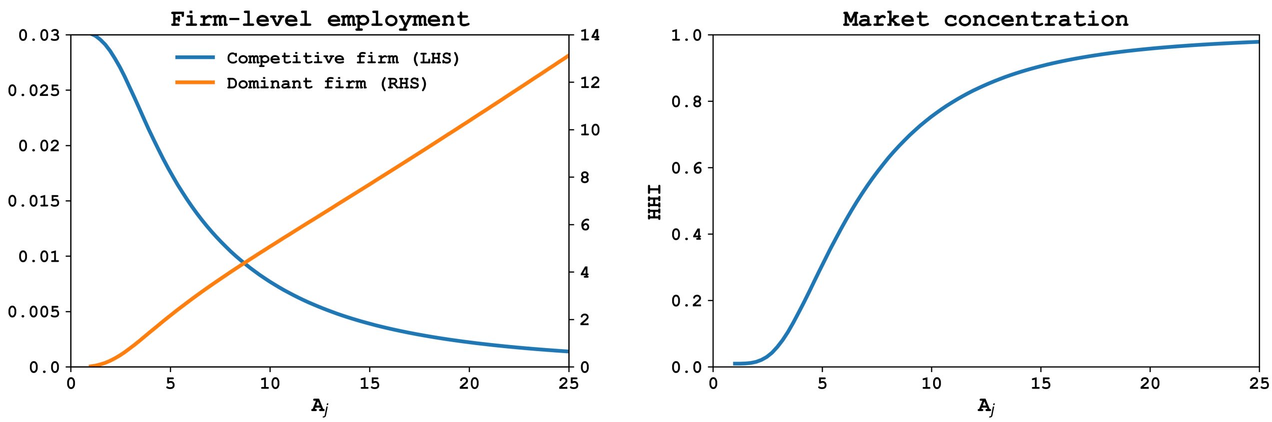

As the figure above shows, market concentration is a monotonically increasing function of, and endogenously determined by, A_j. There are two drivers behind this. Firstly, if the dominant firm is more productive, it will demand more labour, increasing its own employment market share. Secondly, as labour is more productive, the dominant firm pays higher wages, making labour more unaffordable to the competitive firms, reducing their market share. Both these forces contribute towards the labour market becoming more concentrated. As the model aims to provide predictions to later be tested in the data, for the remainder of this section I will characterise labour markets by their HHI rather than their A_j, as this is much more easily measured empirically. Going forward, one should still keep in mind that in this model markets that are more highly concentrated are also those in which productivity is higher.

In terms of wages, markets with higher concentration levels pay higher wages, but this result is purely driven by these markets also having higher levels of average productivity. When controlling for productivity, more concentrated markets pay lower wages, as shown by the average markdown between the wage rate and marginal revenue product of labour of each firm. This goes in accordance with the vast existing empirical evidence (Bassanini et al., 2022, 2023; Benmelech et al., 2022; Félix and Portugal, 2016; Marinescu et al., 2021; Martins and Melo, 2023; Rinz et al., 2018).

Figure 2: Wages as a Function of Market Concentration

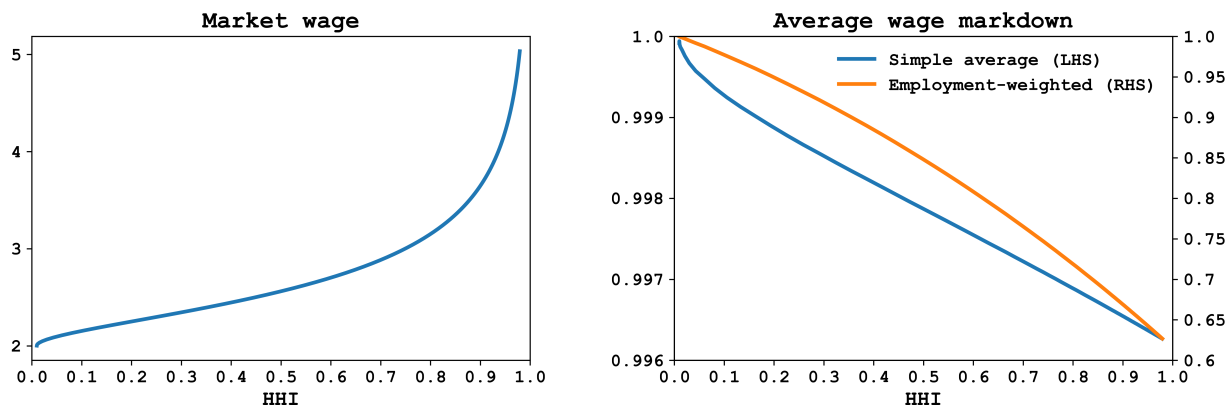

The wedge between the wage rate and the marginal revenue product of labour is a measure of market power, so in this model higher market concentration is also correlated with higher average market power, a result which is also shown by Azar et al. (2019, 2022a). What is more, even though labour supply elasticity is constant, market power is endogenously determined by the productivity of the dominant firm. When the dominant firm is symmetric in terms of productivity to the competitive fringe, then it has no market power - a wage rate below the marginal revenue product of labour will cause labour to be absorbed by the fringe. The dominant firm is only able to exploit its market power when it is productive enough so that it can pay wages below its own MRPL but still above the MRPL of a competitive firm.

Aggregate market employment follows a very similar shape to that of wages, with more concentrated markets having larger levels of employment. Again, this result occurs due to the increased labour demand by the dominant firm, which is willing to pay higher wages so that it can employ more labour when labour is more productive. In truth, employment does not increase as much as would be socially desirable given the higher levels of productivity in the market. As markets become more concentrated, their employment falls increasingly below the socially-optimum level, again as estimated by the literature (Bassanini et al., 2022; Marinescu et al., 2021).

Figure 3: Employment as a Function of Market Concentration

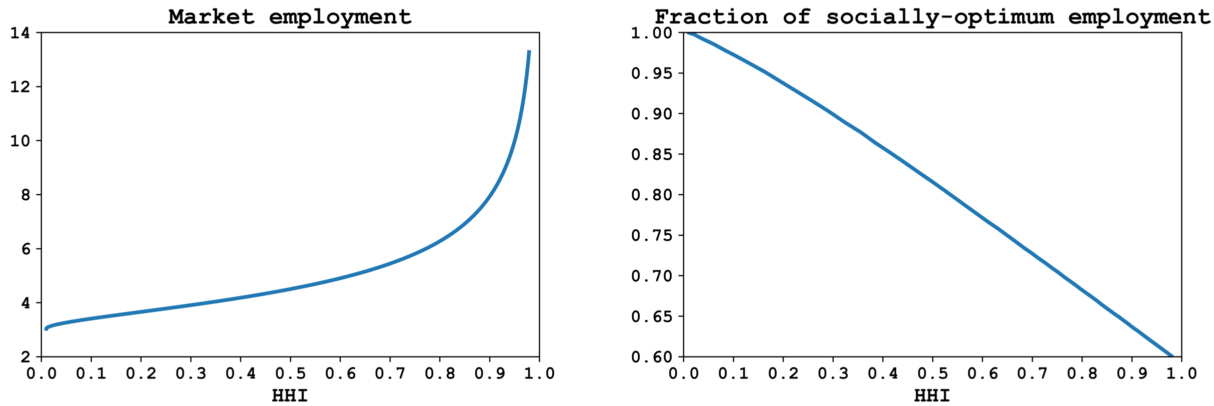

2.3 Minimum Wage Effects

How does the introduction of an economically-binding minimum wage affect the market equilibrium and, more relevant to my research question, how does it impact market employment? To answer this, my starting point is to set the minimum wage equal to $W_j^{min} = \mathbf{L_{jD}^{-1}}(L_{jD})$. This means that the minimum wage is initially non-binding at the level of the free market equilibrium, but any marginal increase will be economically binding. I denote this starting scenario as "0% minimum wage markup". I then analyse how market employment responds to increases of 10%, 20%, and 50% in the minimum wage.

Figure A.1 in the Appendix-A plots the firm-level employment responses to changes in the minimum wage for different levels of A_j. In the figures below, I plot aggregate responses.

Figure 4: Minimum Wage Employment Effects

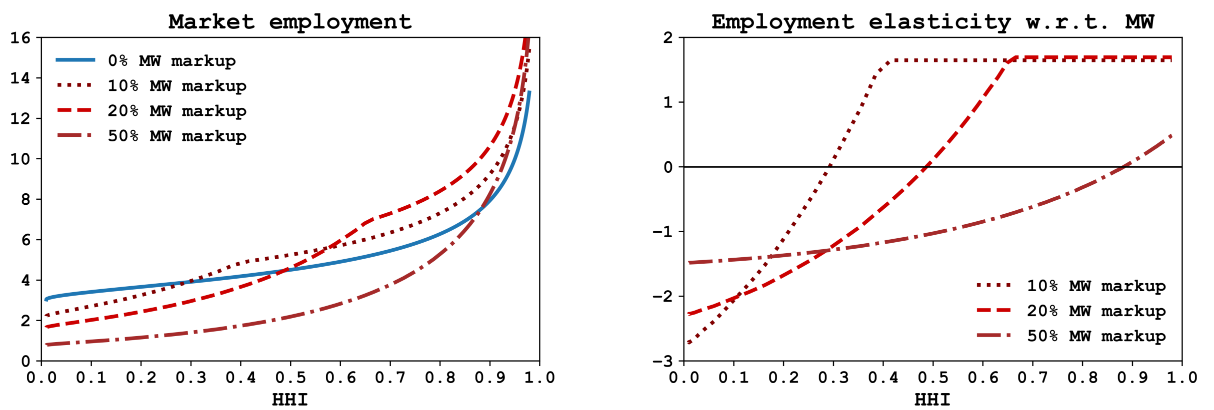

Note: Employment elasticity is defined as average elasticity, $\frac{\Delta \% L_j}{\Delta \% W_j^{min}}$. This is opposed to marginal elasticity $\frac{\partial log(L_j)}{\partial log(W_j^{min})}$ which in this case would be constant for all three markups.
Market concentration is also changing with minimum wage. These figures plot employment responses against initial HHI, i.e. when the minimum wage is not binding.

As the figures suggests, increases in the minimum wage impact employment negatively in markets with low concentration, but this negative effect gradually dissipates as markets become more concentrated. For high enough levels of concentration, the increase in minimum wage may actually lead to an increase in market employment. However, the point at which this occurs will also depend on the magnitude of the minimum wage increase. The intuition for these results is the following. Wage-taking competitive firms maximise profits by choosing labour such that wages and marginal revenue product of labour are equalised. An increase in the minimum wage will for sure raise the wage above MRPL if labour is kept constant, and so a competitive firm will always decrease employment. In contrast, a dominant firm which has market power will have marginal revenue product of labour being above the wage rate. The employment effect will therefore depend on by how much the minimum wage is increased. The dominant firm will be willing to meet the increased labour supply as long as the minimum wage is below MRPL. If instead the government overshoots by setting the minimum wage above MRPL for the corresponding labour supply, then the marginal employment effect at that point will be negative. The overall effect may still be positive if the firm's labour demand for that minimum wage is higher than the labour supply at the original wage rate.

In aggregate terms, positive employment effects are more conspicuous in highly concentrated markets. There are multiple reasons for this. Firstly, higher HHI is correlated with higher market power, and so the starting wedge between the wage rate and MRPL is bigger. There is more room to increase minimum wage and close that wedge while increasing employment. Secondly, higher HHI is also driven by higher A_j, and so again it will take a higher level of minimum wage to cause labour supply to be above MRPL. The government may increase the minimum wage by a larger amount without overshooting and causing negative marginal effects on employment. Finally, as a market becomes more concentrated, the employment response of the dominant firm will weigh increasingly more on the aggregate outcome. For these reasons, for the same minimum wage increase, the change in employment will always be higher in more concentrated markets.

The main takeaway from this model that I wish to test out using the data is whether minimum wage increases will indeed have less severe disemployment effects in highly concentrated markets than in highly competitive market. If this is the case, I also wish to find out whether the difference is big enough so that minimum wage employment effects may in truth become positive if the market is concentrated enough.

3 Institutional Background and Data

In this section, I go deeper into the setting behind my empirical analysis, and how the particularities of Portugal allow me to test the theoretical predictions of the model in a detail that other countries do not allow for.

I first describe the Portuguese national legislation of the minimum wage and the sources behind variation in minimum wage exposure. Here I will highlight the first big advantage of looking at Portugal - the minimum wage has a large bite and directly affects a much bigger share of workers' wages than other countries traditionally studied, such as the US.

I then describe the structure of the dataset I use, its richness and advantages, as well as the caveats that come with it. This is the source of the second large benefit of studying Portugal - unusually rich granular employment data that allows to measure labour market competition in a way that is not possible with the data available for other countries.

3.1 Institutional Setting

The national minimum wage (NMW) was first introduced in Portugal in 1974. It initially discriminated against certain occupations and age groups, but it has been universally applicable to everyone in mainland Portugal since 2004. The national minimum wage is determined by a monthly value, and thus part-time workers are entitled to a minimum salary proportional to their monthly hours of work.

The Portuguese government sets a monthly statutory national minimum wage, which is typically revised and updated every year in January.[1] However, due to extensions of collective bargaining agreements, there exists a very large number of *de facto* minimum wages in Portugal, based on sector and occupation (Martins, 2021). Nevertheless, Card and Cardoso (2022) acknowledge that these negotiated wage floors are likely influenced by the national minimum wage, and so the statutory wage floor impacts wage minima across the entire Portuguese economy.

This paper focuses on the years between 1986 and 2013. During this period, there are four sources of variation in the minimum wage - time, occupation, demography, and geography. Regarding its time variation, there are a few distinct trends in minimum wage evolution. Between 1986 and 2006, the nominal minimum wage increases at a pace which almost exclusively only compensates for inflation, leading its real value to increase at a very slow pace during this period. Following a social concertation agreement, between 2007 and 2011 the nominal minimum wage outpaces inflation and its real value sees a significant increase for the first time in many years (see Carneiro et al. 2011; Centeno et al. 2011 for studies on the ensuing effects). In 2011, this evolution is brought to an abrupt halt. Due to the sovereign debt crisis and under the Economic Adjustment Programme negotiated with the European Commission, the European Central Bank, and the International Monetary Fund, the Portuguese government agrees to keep the nominal value of the minimum wage constant. This would last until October 2014, resulting in its real value being slightly eroded by inflation during these years.[2]

[1] There are a few exceptions to these consistently-timed updates. Between 1974 and 1982, the month in which the national minimum wage was updated varied. In 1983, it began being updated every January more consistently, albeit with some exceptions. In 1989, the NMW was updated twice, once in January, and then again in July. In 2012 and 2013, due to the Economic Adjustment Programme, the minimum wage was not updated. In 2014, the minimum wage was updated in October, which led there to be no update in 2015.

[2] Since 2014, the minimum wage has again been increasing significantly in real terms. See Alexandre et al. (2022) for a study of the impact of these increases.

Figure 5: Evolution of the National Minimum Wage (1986-2013)

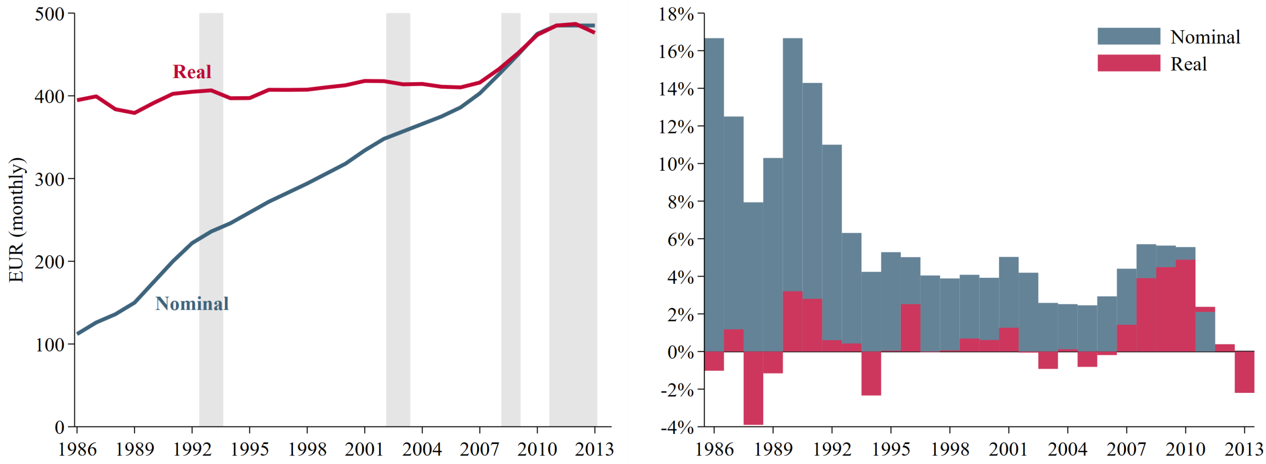

(a) Monthly Value (€) (b) Change y.o.y.

Note: Base prices used is the GDP deflator for 2011Q3. Values show national minimum wage in force in October of each year. In Portugal, wage earners receive a total of 14 months of monthly salary *per annum* (12 regular months plus a Christmas and a holidays bonus). Analogously, even though the national minimum wage is set at a monthly figure, a minimum wage earner will receive 14 monthly minimum wages in a given year. Therefore, for most international comparisons, one would need to multiply the monthly value by a factor of 14/12.
Shaded grey areas represent recession periods, as defined by Reis et al. (2023).

Despite this general trend, other sources of variation foreseen in the law meant that the minimum wage has historically affected workers differently, based on occupation, age, and region. When the national minimum wage was first introduced in 1974, it initially excluded agricultural workers, the armed forces, and domestic workers. In 1977, agricultural workers became entitled to their own occupational-specific minimum wage, which was lower than the national one. The same happened for domestic workers in 1978. Only in 1991 for rural workers, and in 2004 for domestic workers, were these occupational discriminations abolished, and they became entitled to the full value of minimum wage guaranteed to the wider work force.

Young workers have also often been entitled to lower levels of minimum wage. Initially, only those aged 20 and above were entitled to the NMW. This threshold was reduced to 18 years in 1978, and back to 20 years the following year. Between 1979 and 1987, those aged between 18 and 20 were entitled to 75% of the NMW, and workers below the age of 18 could receive 50%. In 1987, the legislation is again altered to entitle those aged 18 above with the full value of the NMW, and minors to receive 75% of this value. In 1998, the law finally enshrines that age-based discrimination is forbidden. These various changes in the age threshold of minimum wage entitlement have been the subject of several natural experiment case studies in the Portuguese minimum wage literature (Cerejeira, 2008; Cerejeira et al., 2012; Pereira, 2003; Portugal and Cardoso, 2006).

In terms of regional variation, the Autonomous Region of Madeira has had since 1987 a minimum wage which is always 2% higher than the value in continental Portugal. Since the year 2000, the Autonomous Region of Azores also has a 5% higher minimum wage than the mainland.

An important point to make of the minimum wage in Portugal is how relevant it is in the Portuguese wage structure. At several points in time, around a fifth of the Portuguese labour force earned the minimum amount they were entitled to, and the minimum wage has more than once been as high as 80% of the median base wage. Figure 6 plots the percentage of minimum wage earners and the Kaitz index (the ratio between the minimum wage and median wage) across time.

Figure 6: Minimum Wage Coverage (1986-2013)

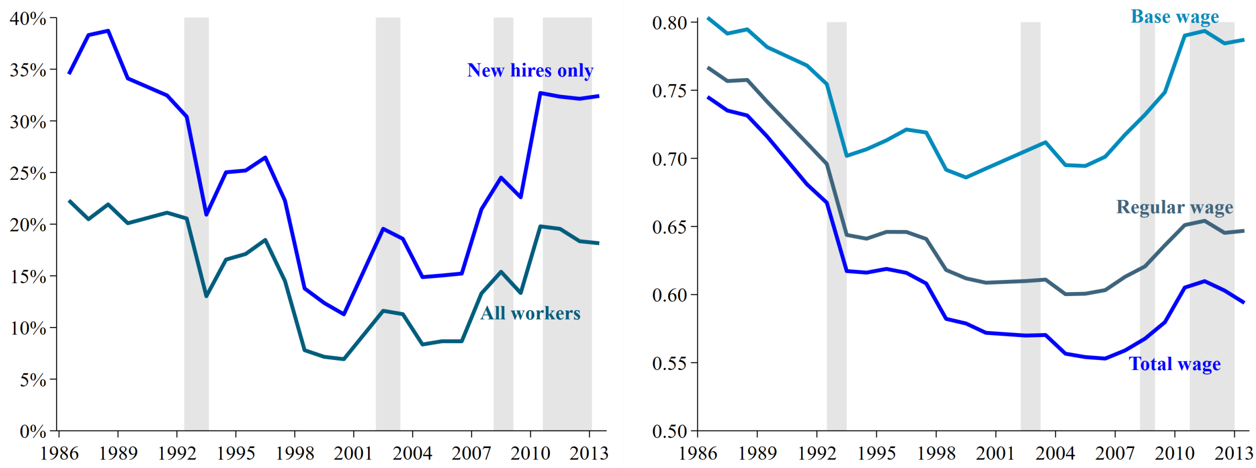

(a) Minimum wage earners (%) **(b)** Kaitz Index

Source: Author's own calculations using *Quadros de Pessoal*.
Note: Three Kaitz Indices are presented: for base wage, since Portuguese minimum wage law establishes a minimum value for base salary, not total; for regular wage, which is the sum of base wage and any other benefits an employee receives regularly, such as meal and commute allowances, risk or nocturnal shift subsidies, and tenure bonuses; for total wage, which includes compensation for extra hours and any other irregular benefits such as productivity bonuses, reparations, or profit participation.
Shaded grey areas represent recession periods, as defined by Reis et al. (2023).

Due to how influential the minimum wage is, any legal changes to its value has a massive impact on the entire wage structure of the Portuguese labour market. Indeed, Oliveira (2023) has estimated that the national minimum wage was responsible for as much as 40% of average wage growth in Portugal between 2006 and 2019.

As a means of comparison, according to the OECD, in 2013 Portugal had a Kaitz index of 0.66 compared to 0.29 in the United States. The fact that the minimum wage is much less impactful in the US has led the literature to often focus on particular industries where the minimum wage is known to be more binding, such as fast-food restaurants (Card and Krueger, 1994; Dube et al., 2007). In Portugal, the minimum wage is impactful on a national scale, allowing for a much wider analysis of the effects of the minimum wage, rather than a more narrow study of a particular local labour market. This is the first big advantage of the setting used in this paper.

3.2 Data

For this paper, I exploit an extremely rich and granular administrative dataset, the *Quadros de Pessoal* (Personnel Records), which consists of a worker-level matched employer-employee panel. Being an annual mandatory census collected by the Portuguese Ministry of Labour for all private sector firms employing at least one worker, it covers the near universe of private sector workers in Portugal since 1986. It was at first collected in January of each year until 1994, and it has been collected in October since then. It contains characteristic and contract information for each worker, such as their age, gender, education, occupation, contract hours, wages, tenure and last promotion. For each worker, the dataset further contains information on the firm plus establishment employing them, including firm age, firm employment, industry, location, sales and composition of equity. Each worker, firm and establishment has time-invariant identifiers. I consider all workers classified as employees, with non-zero wage, aged 15 years and above, with a valid (anonimised) social security number, for the period 1986-2013.[3] Additional data, such as municipality-level population and price levels, was obtained from the *Instituto Nacional de Estatística* (National Statistics Institute).

Measuring concentration of a labour market requires at least firm-level employment data. In the US, commonly used datasets such as the Quarterly Cencus of Employment and Wages contain county-level employment. This has been one of the biggest obstacles in the way of empirically testing the predictions of the monopsony model of minimum wage. Azar et al. (2023) overcome this by using the Burning Glass Technologies dataset, which gives them firm-level online job vacancies, which allowed for the recent breakthrough in the literature. However, the limitations are that they are restricted to studying local labour markets that post most of their job vacancies online, while still being at risk of some measurement error as they miss any other type of vacancy posting, inducing mismeasurement in market concentration.

Quadros de Pessoal contains employment information at the most granular level possible, the worker, while maintaining the link to the employing firm. Moreover, the administrative and mandatory nature of the census means I am not restricted to any particular industry or occupation, and I am able to study the near-universe of the Portuguese labour force. Therefore, the existence of such rich and granular data presents the second major advantage of the Portuguese setting in testing the predictions of the monopsony model.

The dataset is, however, not without its limitations. Unlike US county-level data, it is updated at an annual frequency rather than quarterly. This is mostly not an issue, as historically the national minimum wage is usually updated once a year in January. However, in the rare occasions in which the national minimum wage was updated more than once within the same year,[4] I only capture one of those changes. The dataset also suffers from a discontinuity in 1990 and 2001. One other major drawback of the data is that it does not contain public employees,

[3] More recent years do not contain detailed geographical information, which makes it impossible to accurately define local labour markets.

[4] This only occurs in 1989 for the time period I consider.

which can potentially create measurement errors of concentration for labour markets in which public and private enterprises may coexist, such as education or health services. As a robustness, therefore, I also test specifications which exclude sectors where the government is likely to be a non-negligible competitor in the labour market.

4 Local Labour Markets

The empirical part of this paper uses the local labour market as the unit of observation. In order to empirically test the predictions of the model, I use *Quadros de Pessoal* to cluster workers into local labour markets, and for each local labour market I: i) measure how exposed a labour market is to changes in the minimum wage; ii) calculate the market's concentration level based on firms' shares of yearly new hires.

4.1 Design

I define local labour markets based on the location and industry that workers are employed in. A common challenge in the literature is defining the geographic boundaries of Portuguese local labour markets. Administratively, Portugal is split into 18 districts,[5] which are further sub-sectioned into 308 municipalities, neither of which are adequate definitions of a commuting zone. Most districts are far too large to be considered effective commuting zones, while there exist significant inter-municipal flows (particularly surrounding major urban areas) that make municipalities too granular of a division.

To improve the precision by which local labour markets are geographically defined, I take advantage of the work of Afonso and Venâncio (2016). They use 2011 census data in order to analyse commuting flows across municipalities in mainland Portugal and cluster them into into 52 commuting zones. Although not an official administrative division of Portuguese national territory, these are much likely to be closer representations of local labour markets, as they are directly constructed by analysing the areas in which people are willing to commute to work. Afonso and Venâncio only define commuting zones for mainland Portugal. For completeness, I assume each of the 11 inhabited islands that make up the Azores and Madeira to be a separate commuting zone, and therefore I consider a total of 63 commuting zones in my analysis.[6]

Unlike Azar et al. (2023), who define local labour markets as a combination of location and occupation, I will instead mainly do so on a basis of location and sector of activity. The main reason for this is for labour markets to be more self-contained. Using an occupations-based definition creates higher inter-market mobility than when using sectors, which can mean that labour markets are better defined using sectors rather than occupations. I discuss this issue further in Section 4.2. For now, I use the 88 sectors defined by the Portuguese Classification of

[5]The Autonomous Regions of Azores and Madeira do not have defined districts.

[6]See Figure B.1 in the Appendix-B for a map of the commuting zones.

Economic Activity (CAE-Rev.3) at the two digit level.

Overall, for each year t, I define a local labour market m to be any distinct combination of the 63 commuting zones and 88 sectors of activity. In the data there exist 5,069 local labour markets at some point in time, and I obtain 63,210 observations for the whole period. My analysis, however, is not done on the universe of local labour markets. I make some additional restrictions in order to select my sample of analysis.

Firstly, I exclude pure public sectors, which in this case are sectors classified as *Public Administration and Social Security*, or *International organisations and overseas institutions*. This is done in order to clean the data off markets for which I have little information, since the dataset consists solely of private sector employees. Secondly, I exclude pure monopsony labour markets, i.e. markets for which in a given year there is a single firm hiring. Although this may seem counter intuitive for the purposes of this paper, the vast majority of these markets are very tiny in terms of employment, and not markets in which one firm employs several people. Getting rid of these markets does not substantially decrease the representativeness of the sample in terms of total employment, but it reduces the number of excessively small markets I consider. Thirdly, I exclude labour markets which are not hiring, as for these I will not be able to measure the degree of competition in terms of recruitment. Finally, in order to further ensure markets have a sizeable dimension, I trim the bottom 5% of remaining markets in terms of hires.[7]

It is important that the sample does not include excessively small markets for a couple of reasons. Firstly, the way I defined markets creates some observations which only contain a handful of workers, and conceptually these cannot be realistically considered actual markets. Secondly, and more importantly still, not getting rid of these markets could induce endogeneity in my concentration measure, as small markets mechanically force a high level of concentration. This issue is discussed more in depth when I introduce my instrument in Section 5.2.

These restrictions generate a sample of 3,668 local labour markets, and 38,107 observations over the entire period. Summary statistics for the universe and selected sample of labour markets, are presented in Table 1.

Workers in the selected sample are demographically and economically representative of the universe of local labour markets, being of similar age, education, and tenure. A similar percentage is female and earns the minimum wage. The sample is also representative in terms of hours worked and average wage. Labour markets in the selected sample are larger in terms of average number of firms, average number of workers, and average yearly hires. This is to be expected as many of the restrictions in creating the sample envisioned removing excessively small markets. Nonetheless, despite losing 40% of labour market observations, my sample still covers almost 95% of all workers.

[7]A few markets are also single-firm markets with zero sales. These were also excluded to ensure markets consist of active firms.

Table 1: Summary Statistics for Universe and Selected Sample of Local Labour Markets

Panel A: average across workers		
	All labour markets	**Selected sample**
Age (years)	36.1	36.0
Female workers	42.5%	42.9%
Years of schooling	7.6	7.6
Tenure (months)	88.8	88.1
Regular real wage	€850.19	€847.07
Monthly normal hours	163.5	163.6
Minimum wage earners	14.5%	14.5%
Panel B: average across labour markets		
	All labour markets	**Selected sample**
Number of firms	89	140
Employment (nr. workers)	789	1,240
Total hires (yearly)	162	257
Employment representativeness	100%	94.8%
Observations	63,210	38,107

Note: Observations are pooled across time. Base prices used is the GDP deflator for 2011Q3.

4.2 Market Mobility

In this section, I test how stable and self-contained my defined local labour markets are, in order to validate the way I constructed them. I use the original worker-level data to examine worker mobility across the two dimensions of my local labour markets, location and activity. I also compare this mobility to alternate variables for each dimension.

Firstly, I examine how Afonso and Venâncio's commuting zones compare to the usage of districts in geographically defining local labour markets. Table 2 does this comparison,by showing where each worker was compared to the previous year, using commuting zones and districts as the location variable, respectively. Geographic mobility is strikingly almost exactly equal whether I use commuting zones or districts as workers' location variable. The majority (73-74%) of workers remain in the same location they were in the previous year. Even when a worker changes the firm they work for, it is much more common for them to remain in the same location (5.6-5.7% of all workers) than to move locations (1.3-1.4%). It is also quite rare for a worker to be relocated within the same firm (0.8% of all workers).

Table 2: Geographic Mobility

Geographic mobility of workers	**Location variable**	
Workers in year t compared to $t-1$	Commuting zones	Districts
Same location	**73.4%**	**73.5%**
Same firm	67.8%	67.8%
Different firm	5.6%	5.7%
Different location	**2.2%**	**2.1%**
Same firm	0.8%	0.8%
Different firm	1.4%	1.3%
Entering employment	**24.4%**	**24.4%**
Clusters considered	63	29

Overall, it appears that both commuting zones and districts do well in geographically defining local labour markets, as mobility across location is very low using either variable. What is striking, however, is that commuting zones achieve the same accuracy using much narrower geographical areas. While splitting Portugal into commuting zones creates 63 clusters, doing so by districts creates only 29.[8] At the logical extreme, imagine I considered the whole of Portugal to be one labour market. Then, mobility across "locations" would be zero. The more I can partition Portugal without raising mobility, the more efficiently and accurately I am defining local labour markets. This means that commuting zones are a much more efficient geographic boundary, as for a larger partition of territory I obtain the same mobility.[9] This exercise is reassuring that the usage of commuting zones is the most appropriate variable I have to geographically define local labour markets.

When it comes to the activity dimension of the local labour markets, the choice is not as clear. I define my local labour markets based on the sector of the employing firm, but I could have as well defined the local labour markets based on the occupation classification of the worker. Conceptually, there are arguments for both choices. You could say that an accountant can equally be employed by a firm operating in real estate or by a retail company, and so my labour market should be the labour market for accountants, and not the labour market for the real estate industry, and the labour market for the retail industry. Or, you could consider that there are industry-specific skills and knowledge that are non-transferable, and so an engineer working in energy could not transfer to transportation, implying a sector-based definition of local labour markets would be more appropriate. Table 3 compares the usage of sectors with the usage of occupations in terms of mobility across activities.

[8]Each of the inhabited islands was considered a district, even though officially the Azores and Madeira do not have defined districts.

[9]Note that, although both commuting zones and districts are made up of municipalities, commuting zones are not nested within districts. So a commuting zone may overlap with two or more districts, and vice versa.

Table 3: Activity Mobility

Activity mobility of workers	**Activity variable**	
Workers in year t compared to $t-1$	Sector	Occupation
Same activity	**64.2%**	**62.3%**
Same firm	61.0%	58.6%
Different firm	3.2%	3.7%
Different activity	**11.4%**	**13.3%**
Same firm	7.6%	10.0%
Different firm	3.8%	3.3%
Entering employment	**24.4%**	**24.4%**
Clusters considered	88	98

Note: Sectors are from CAE-Rev.3 at the 2 digit level. Occupations are from CPP-2010 at the 2 digit level.

Overall, the usage of sectors does slightly better than the usage of occupations. A marginally higher percentage of workers (64%) stays in the same activity from one year to the next when activity is defined by sector than when activity is defined by occupation (62%).

This is mostly driven by workers keeping the same activity when staying at the same firm. This is to be expected as sector is a firm-side variable, whereas occupation is a worker-side variable. In contrast, when workers change firm, it is actually more common for them to remain in the same activity when activity is defined by occupation (3.7% vs 3.2%).

What is worrisome about the usage of occupations, is that even when workers stay at the same firm, they change activity classification in terms of occupation more often (10% of all workers), than in terms of sector[10] (7.6%). This can be problematic as it may lead to measurement error - the worker has not changed employment, yet is being classified with a different activity. In my local labour market observations, this will show that employment has fallen in one local labour market and risen in another, when in truth there was no change in that worker's employment situation. For this reason, I prefer to use a sector-based definition of local labour markets. Nonetheless, because in this case the choice is not obvious, later in the paper I present my results also using occupations-based local labour markets as a robustness.

4.3 Minimum Wage Exposure

Now that I have defined my local labour markets, and shown how well they do in terms of inter-market mobility, I explain how I create two additional variables, to measure the impact of the minimum wage and the degree of market power in each local labour market.

My theoretical model assumed that the minimum wage was always economically binding. This, however, might not always be true, as for some of local labour markets it may be the case that all workers earn above the legal minimum, in which case raising the minimum wage should have no effect on employment.

[10]This can happen if the sector classification of the firm changes, or if the worker is relocated to a different establishment within the same firm that has a different sector classification.

Then, to measure by how much labour market m is exposed to the national increase in the minimum wage in year t, I use the following metric:

$$\Delta \tilde{w}_{mt|t-1} \equiv \tilde{w}_{mt|t-1} - w_{m,t-1} \tag{8}$$

in which $\tilde{w}_{mt|t-1}$ is the counterfactual (log) wage bill necessary to comply with the (real) minimum wage if market m had kept all its workers from year $t-1$ working the same number of hours, and $w_{m,t-1}$ is the actual (log real) wage bill of labour market m in the previous year.

In other words, $\Delta \tilde{w}_{mt|t-1}$ is measuring by how much the wage bill of labour market m would have to increase in real terms in year t in order to simultaneously keep the same labour force composition as in year $t-1$ and to comply with the new minimum wage. It is the minimum-wage-induced real increase in the wage bill, *ceteris paribus*.

If all workers in year $t-1$ were earning above the minimum wage for year t, then exposure is zero. The exposure will be maximal when all workers in year $t-1$ were earning the minimum wage. Intermediate levels will depend on the number of previous year workers earning below the new minimum wage.

This measure allows me to fully capture all variations of the minimum wage, including age-based discrimination. The local labour markets were defined based on location and sector, which encompasses two of the historical sources of variation in the minimum wage. But, even though labour markets are not defined on an age dimension, the granularity of my data allows me to assess how changes to minimum wage legislation affects each worker individually (which may be affected by the worker's age), and then aggregate those individual impacts to the whole labour market. Two markets which are identical in terms of wage distribution may still be impacted differently if their demographic composition is different.

Looking at the selected sample, 68% of labour market observations are exposed to the minimum wage, which is an encouraging reaffirmation of how Portugal is a good setting in which to study minimum wage effects. Of those observations that do get impacted, the average minimum-wage-induced increase in the wage bill is 0.4%. As an illustrative example, such number could occur in a market where the minimum wage is raised by 2% in real terms, and 20% of workers were earning the minimum wage.[11]

4.4 Labour Market Concentration

The model in Section 2 showed that the existence of labour market power is key for the positive employment effects of minimum wages. The model then went to show that labour market concentration and market power are highly correlated, and so the latter can be used as a proxy for the former. This result is also shown by Azar et al. (2019) and Manning and Petrongolo

[11] This is an approximation. If 20% of workers earn the minimum wage, and this is increased by 2%, then the 0.4% minimum wage impact is a lower bound. There can be other workers earning above the previous minimum wage but below the new legal minimum which will also be impacted.

(2022).

A popular measure of market concentration in IO literature is the Herfindahl-Hirschman Index (HHI), which was also used in my model. Both Azar et al. (2022a) and Azar et al. (2023) calculate HHI using the shares of online job postings. They do so due to the difficulty of finding granular employment data for the US, and it raises the constraint of being limited to markets that intensively hire online. Since I am not constrained by lack of granularity in *Quadros de Pessoal*, I will diverge from this and instead calculate HHI using the share of new hires, as done by Marinescu et al. (2021) and Martins and Melo (2023)[12]. Labour market concentration is then measured as the sum of squared hires shares across all firms in a given labour market and year. More concretely, and defining $\mathcal{M}$ as the set of all firms i participating in local labour market m in year t, the formula for HHI is given by

$$H_{mt} \equiv \sum_{i \in \mathcal{M}} \left(\frac{\text{New Hires}_{i,mt}}{\sum_{i \in \mathcal{M}} \text{New Hires}_{i,mt}} \right)^2 . \tag{9}$$

HHI, by construction, can vary between 0 and 1. The higher its value, the more concentrated is the labour market. If a labour market is active (i.e. is hiring), then HHI must be strictly greater than 0. If HHI tends to 0, then the labour market is perfectly competitive and has many recruiters. If HHI equals 1, then the labour market is a monopsony, i.e. it has a single recruiter.

Figure 7 shows the distribution of local labour markets in the selected sample by their concentration level, as well as the workers in these markets. Three things are important to note: firstly, the majority of local labours markets have low concentration levels. As a reference point, under US Department of Justice/Federal Trade Commission horizontal merger guidelines, markets with a HHI below 0.25 are considered highly competitive or unconcentrated;[13] secondly, low-concentration markets have larger levels of employment than high-concentration markets, so an even higher proportion of workers belong to highly competitive labour markets; thirdly, labour markets are heterogeneous in their concentration, and there are labour markets across the entire HHI spectrum, providing me with sufficient variation across observations.

In the Appendix-C, I show the heterogeneity in HHI and minimum wage exposure across time, regions, and industries.

[12]The dataset includes information at the firm and establishment level. While local labour markets were defined using establishment-level information, HHI was calculated using firm-level shares. This implicitly assumes that establishments belonging to the same firm in the same labour market do not compete against each other, and that there is some degree of coordination in hiring policy across the multiple establishments of a firm.

[13]These guidelines are mostly applied to product markets, and not labour markets.

Figure 7: Distribution of Selected Sample

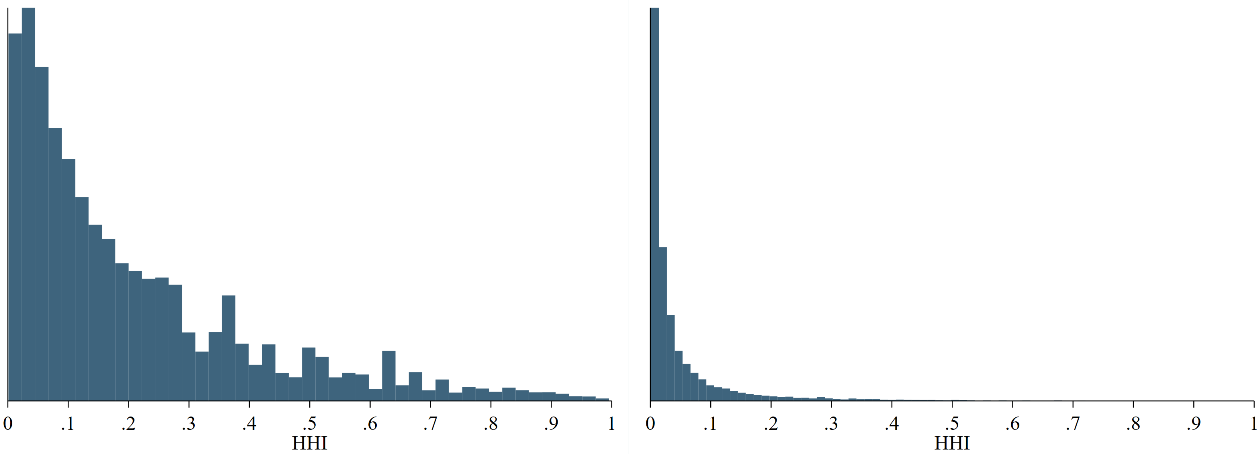

(a) Distribution of Labour Markets

(b) Distribution of Workers

Note: Observations are pooled across time.

5 Empirical Results

5.1 OLS Estimates

I now have all the variables necessary to empirically estimate the effects of minimum wage increases on employment on a local labour market level. For this, I estimate the following panel two-way fixed effects model by OLS,

$$\begin{aligned} \Delta \ell_{mt} &= \beta_1 \cdot \Delta \tilde{w}_{mt|t-1} + \beta_2 \cdot H_{mt} + \beta_3 \cdot \Delta \tilde{w}_{mt|t-1} \times H_{mt} \\ &+ \sum_{j=1}^{2} \gamma_j \cdot \Delta \tilde{w}_{m,t-j|t-j-1} + \mathbf{X}'_{mt} \cdot \mu + \alpha_m + \phi_t + \varepsilon_{mt}, \end{aligned} \tag{10}$$

where $\Delta \ell_{mt}$ is the percentage change in employment in terms of number of workers in labour market m and year t, and $\Delta \tilde{w}_{mt|t-1}$ and H_{mt} are, respectively, market minimum wage exposure and HHI as defined in Equations (8) and (9). I include two lags of the minimum wage (MW) exposure variable in order to control for any delayed responses in employment. $\mathbf{X}'_{mt}$ is a vector of variables controlling for economic and demographic factors that may influence employment equilibrium. In terms of demographics, it includes fraction of female workers, average worker age, average years of schooling, average tenure, number of firms, and median firm age. The growth rate of total commuting zone population aims to control for changes in labour supply. To control for shocks in worker productivity and consequently labour demand, I include as covariates the growth rate of real sales per worker, the growth rate of average sector real wage, and the growth rate of average commuting zone real wage. I further include the growth rate of average real wage in the given labour market, to capture all other shocks that may affect the market equilibrium. The labour market and year fixed effects are respectively represented by α_m and ϕ_t, and ε_{mt} is the error term. Standard errors are clustered at the sector $\times$ NUTS-II region[14] level, therefore assuming no correlation between labour markets of the same sector across different NUTS-II regions.

The β coefficients are my main coefficients of interest. β_1 indicates the percentage change in employment following a 1 p.p. increase in minimum wage exposure in perfectly competitive labour markets, i.e. when $H \to 0$. β_2 indicates by how much employment changes as you increase the level of concentration in the market, absent the existence of a minimum wage. Finally, β_3 represents the additional change in employment induced by the minimum wage (on top of β_1), conditional on the level of labour market concentration $0 < H \leq 1$.

The main threat to identification arises if firms anticipate the increase in minimum wage and adjust employment accordingly *before* the change is enacted. The minimum wage exposure variable I constructed, $\Delta\tilde{w}_{mt|t-1}$, is arguably exogenous as it does not reflect actual minimum-wage-induced growth in wages, but rather the counterfactual change if firms maintained their employment from year $t-1$. However, this variation will only be exogenous if employment in year $t-1$ has not itself been affected by minimum wage changes in year t. This can occur if firms anticipate the change in minimum wage and begin adjusting employment preemptively in order to smooth out potential convex adjustment costs. If this is the case, then firms are manipulating how exposed they are to future changes in the minimum wage, and my minimum wage exposure variable will be endogenous. To dismiss such concerns, in Section 5.3, when I estimate dynamic employment responses, I look at the pre-trends and find that such manipulation is not much of a concern, particularly for lowly concentrated markets. For now, I proceed with my described empirical strategy and show estimates for contemporaneous employment responses to increased minimum wage exposure.

Table D.1 in the Appendix-D presents the OLS estimates for the effect of minimum-wage-induced wage growth and concentration on employment growth regressions. Column (1) presents a base regression using solely MW exposure, HHI, and their interaction, with no additional controls. Column (2) adds the controls contained in vector $\mathbf{X'_{mt}}$. Besides the controls, column (3) adds both year and market fixed effects. Column (4) is the complete specification, with full controls, fixed effects, and also adding the two lags of minimum wage exposure.

The coefficients concerning MW exposure, β_1 and β_3, are only simultaneously statistically significant with the inclusion of fixed effects. The two-way fixed effects gives the specification a triple difference interpretation, as I am comparing different changes to employment across labour markets with different levels of concentration. Interestingly, the estimated coefficient

[14]Portugal is split into seven NUTS-II regions: North, Centre, Lisbon Metropolitan Area, Alentejo, Algarve, Autonomous Region of the Azores, and Autonomous Region of Madeira. Some commuting zones belong to more than one NUTS-II region. These are rare, however, and for those cases I assigned the commuting zone to the NUTS-II region to which most constituent municipalities belong.

on HHI, β_2, is initially positive, but when I add fixed effects and then lags of minimum wage exposure, it becomes sequentially non-significant and negative. This is in accordance with the predictions of the model. As was noted in Figure 3, the model predicted employment to increase with HHI, but this correlation was inverted once market productivity was controlled for, something that would be captured by the fixed effects.

The results from column (4) - the full specification - can be interpreted as follows. In a perfectly competitive labour market, when the minimum wage forces wages to grow by an additional 0.1 p.p.[15], employment is estimated to grow by 1 p.p. less. However, as indicated by the estimate for β_3, the more concentrated a market is, the less severe this disemployment effect will be. These results are in accordance with the first prediction of the model. In perfectly competitive labour markets where labour market power is non-existent or negligible, minimum wage increases will have a negative impact on employment. As market power increases with concentration, this negative impact dissipates.

The question still remains whether the difference in employment effects can be strong enough such that increases in the minimum wage actually have a positive impact on employment in highly concentrated markets. Figure 8 answers this question by plotting the estimated response of employment growth to a 1 standard deviation minimum-wage-induced increase in wage growth for the whole spectrum of HHI values. It also plots the distribution of local labour markets (LLM) and workers across that same spectrum.

Figure 8: OLS Employment Response to a 1 Std. Dev. Minimum-Wage-Induced Increase in Wages

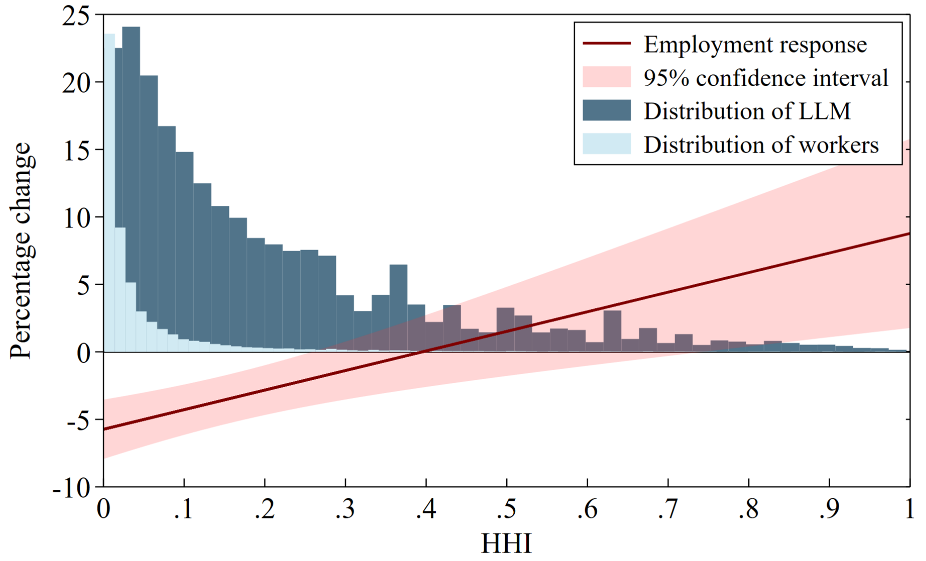

Note: Figure is plotting the estimates for $\beta_1 + \beta_3 \times \text{HHI}$. Standard errors are clustered at the sector (1 letter) by NUTS-II level.

As it shows, at a HHI level of around 0.25, minimum wage disemployment effects become

[15] The average value of the minimum wage exposure variable is 0.3%, so it doesn't make much sense to speak of magnitudes of 1 p.p. increases.

statistically non-significant. Moreover, for markets with concentration levels above 0.72, minimum wage increases actually have a statistically significant positive effect on employment. Nonetheless, as the figure shows, the majority of local labour markets (and an even higher proportion of workers) are located in highly competitive levels on the HHI spectrum. As a result, increasing minimum wages would have larger negative effects than positive effects on a national scale.

5.2 IV Estimates

Despite these results, there may be reasons for concerns regarding the exogeneity of the HHI measure. There are several potential sources of endogeneity in H_{mt} which are worth discussing.

Firstly, there may exist some measurement error in labour market concentration, particularly in sectors with a strong government presence. *Quadros de Pessoal*, consists solely of private sector workers, but in sectors such as education and health the government might be a significant employer and I could therefore severely miscalculate local labour market concentration in these cases.

Secondly, there could exist some reverse causality between employment and concentration given the way the concentration variable is constructed. Imagine this scenario: a labour market has an employment level of only two workers in a given year. At most, the number of hires for that market was two in that year, if both workers are new hires. If both workers were hired by the same firm, then HHI is 1. If they were hired by separate firms, the HHI is 0.5. In any case, a labour market with two workers mechanically creates a lower bound on HHI of 0.5. More generally, a labour market with L workers will impose a lower bound of $\frac{1}{L}$ on HHI for that market, by construction. This can create simultaneity between concentration and employment. This issue is partially addressed when I removed the bottom 5% of markets in terms of hires, but the lower-bound restriction always exist, even if it becomes less binding as labour markets increase in size.

Finally, there is a second plausible source of reverse causality between the regressand and HHI, as firms in rapidly expanding labour markets may not grow symmetrically, and thus affect the levels of concentration.

These concerns together are enough threaten the validity of the previous estimates, and to warrant the use of an instrument for HHI. I follow the instrument commonly used in IO and labour economics literature (Autor et al., 2013; Azar et al., 2022a; Martins and Melo, 2023; Nevo, 2001), which is the average of $log\left(\frac{1}{N}\right)$ in all other commuting zones for the same sector and year, where N corresponds to the number of firms in the labour market. This instrument is less likely to be endogenous as it does not directly depend on market shares. It instead measures concentration as the inverse of the number of firms in a market, and it captures variation in local labour market concentration driven by nation-wide shocks affecting the labour markets of the same sector as the instrumented labour market, excluding itself.

Table D.2 in Appendix-D shows the 2SLS estimated coefficients.[16] Figure 9 plots the estimates across the range of HHI values. As the figure shows, the IV and OLS estimates are very similar, both qualitatively and quantitatively. Standard errors are larger, particularly for highly concentrated markets. This means that only for the cases close to monopsony can I say that minimum wages have positive employment effects. Nonetheless, the main results remain: minimum wage increases have negative employment effects in highly competitive markets, but this negative effect becomes less severe as markets become more concentrated. In extreme cases of concentration, the employment effect may actually be positive.

Figure 9: IV Employment Response to a 1 Std. Dev. Minimum-Wage-Induced Increase in Wages

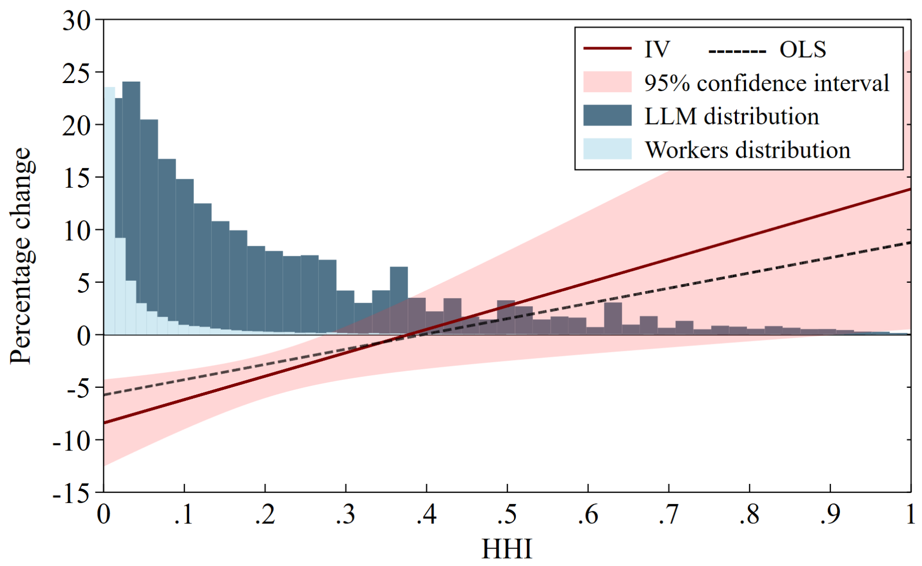

Note: Figure is plotting the estimates for $\beta_1 + \beta_3 \times \text{HHI}$. Standard errors are clustered at the sector (1 letter) by NUTS-II level.

5.3 Dynamic Responses

While so far I have been able to verify the main predictions of the model, what I have shown are only contemporaneous responses of employment to changes in the minimum wage. The usage of the panel dataset allows me to answer more questions that go beyond that of the static model. Namely, are there delayed responses of employment to changes in the minimum wage? Are the employment effects persistent or do they eventually die out?

To answer these questions, I adapt Equation (10) to a Local Projections (Jordà, 2005) specification. I estimate the cumulative employment response from period $t-1$ to period $t+h$ to an increase in minimum wage exposure at time t.[17]

[16] See Table D.3 for the first-stage regression.

[17] Regarding the controls present in $\mathbf{X}'_{m,t+h}$, static controls (fraction of female workers, average worker age, average years of schooling, average tenure, number of firms, and median firm age) are included at their level in period $t+h$. Dynamic controls (growth rate of commuting zone population, growth rate of real sales per worker, growth rate of average sector real wage, growth rate of average commuting zone real wage, and growth rate of average real wage) were included as the cumulative growth rate between period $t-1$ and $t+h$.

$$\Delta_{h+1}\ell_{m,t+h} = \beta_{1,h} \cdot \Delta\tilde{w}_{mt|t-1} + \beta_{2,h} \cdot H_{mt} + \beta_{3,h} \cdot \Delta\tilde{w}_{mt|t-1} \times H_{mt} \\ + \sum_{j=1}^{2} \gamma_{j,h} \cdot \Delta\tilde{w}_{m,t-j|t-j-1} + \mathbf{X}'_{m,t+h} \cdot \mu_h + \alpha_m + \phi_t + \varepsilon_{mt}. \tag{11}$$

Below, I plot the dynamic responses of employment for the two extreme cases of concentration - perfect competition (corresponding to $\beta_{1,h}$), and monopsony (corresponding to $\beta_{1,h} + \beta_{3,h}$) - for a horizon of 5 years. I also plot this for the OLS regressions and for when instrumenting HHI. While the OLS estimates show that for perfect competition the negative employment effects are highly persistent, the IV responses tell us otherwise. Both for perfect competition and monopsony, the respective negative and positive employment effects are quite transitory. One year after the increase in minimum wage, there is no statistically discernible change in employment.

Figure 10: Cumulative Employment Response to a 1 Std. Dev. Minimum-Wage-Induced Increase in Wages

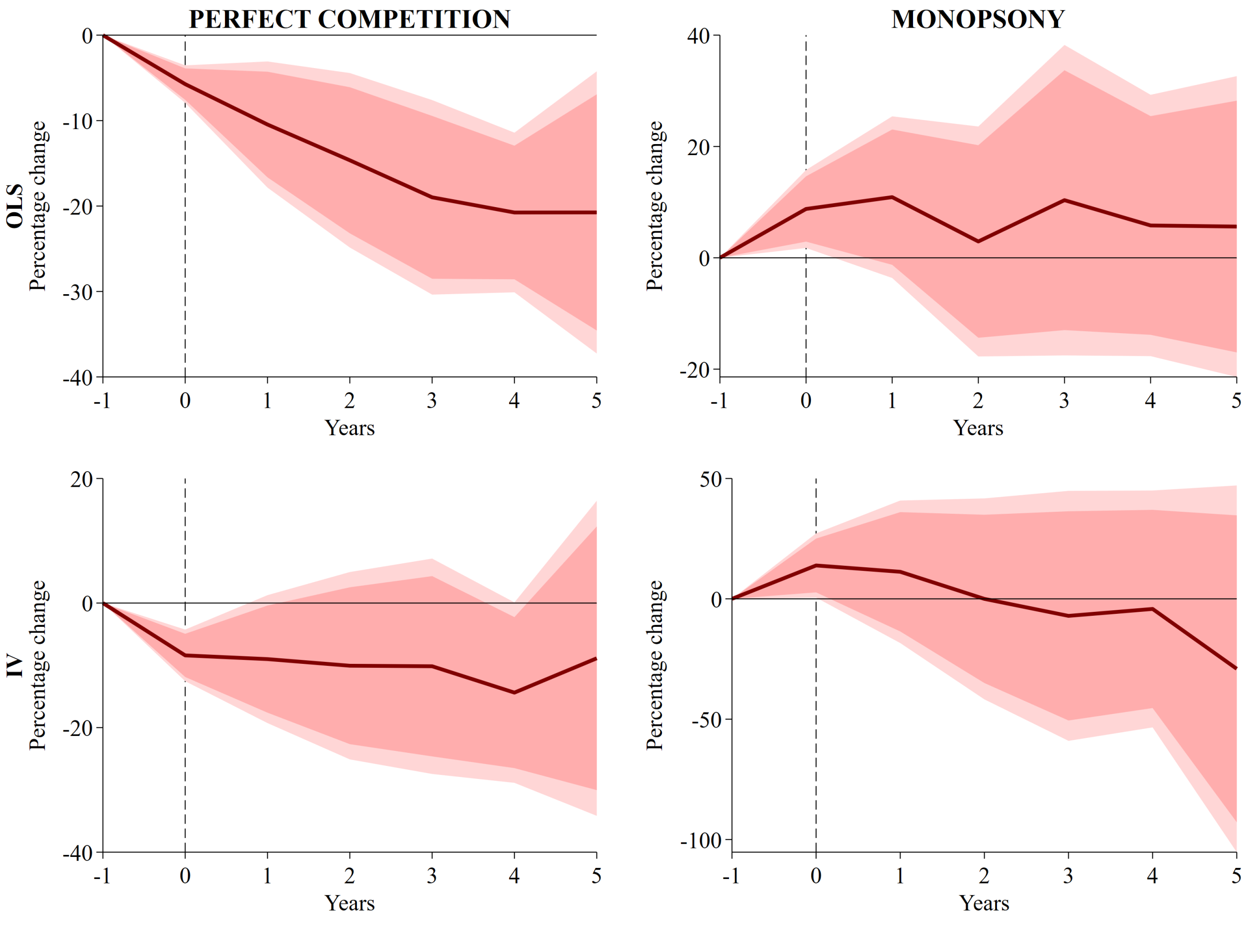

Note: Dark pink and light pink shaded areas show 90% and 95% confidence bands, respectively. In the case of perfect competition, impulse response at horizon h corresponds to $\beta_{1,h}$. In the case of monopsony, the impulse response is plotting $\beta_{1,h} + \beta_{3,h}$. Standard errors are clustered at the sector (1 letter) by NUTS-II level.

In the Appendix-D Figure D.1, I show a visualisation of the pre-trends, in order to dismiss concerns of the previously discussed potential manipulation in the minimum wage exposure variable, originating from anticipated minimum wage changes. For perfectly competitive markets, there is no statistically significant pre-trend to be concerned with. There is more of a pre-trend in monopsonistic labour markets, which may indicate a slight bias in my estimates for these markets.

5.4 Robustness Checks

I now turn to perform a series of robustness checks. For conciseness, I include detailed tables and plots in the Appendix-E. Here, I merely describe the checks I perform.

Firstly, while the selection of labour markets included in my sample resulted in a 95% employment representativeness, it still caused me to lose 40% of local labour market observations. To ensure my results are not driven by this selection, I run the same regressions whilst keeping all observations. Secondly, as noted earlier, the dataset does not contain information regarding public sector employment. For this reason, I might be mismeasuring the level of concentration in labour markets where both private firms and the government compete, such as education or health. For this reason, I perform a second robustness that excludes local labour markets belonging to sectors in which the government is likely to take part in.[18] Thirdly, I run the same regressions but measuring concentration based on employment shares rather than new hires shares, to ensure results are not driven by the particular measure of HHI used. Finally, and perhaps most importantly, I redesign my local labour markets to be defined as a commuting zone $\times$ occupation pair, rather than commuting zone $\times$ sector. I perform all the same sample restrictions as I did for my sector-based observations.

Results for including all observations, excluding public sectors, and an employment-based HHI are very similar to the benchmark specification, both qualitatively and quantitatively.[19] It is in the occupation-based sample that results differ. As shown in Figure E.1, while there is evidence for negative employment effects in highly competitive markets, nothing can be said regarding what happens at higher concentration levels. Employment effects are statistically non-significant, but the confidence bands are so large that I cannot even rule out the possibility that employment effects are even more negative in highly concentrated markets. However, this is only for the contemporaneous employment response, as the dynamic responses in Figures E.2 and E.3 paint a different picture. Contrary to a sector-based definition of local labour markets, when using an occupations-based sample the employment effects have a delayed response and are highly persistent. While it is true that not much can be said at a horizon of $h = 0$, after one year the effects are clear - minimum wage increases have negative employment effects in perfectly competitive markets, and positive effects in monopsonistic markets. Moreover, these

[18] The excluded sectors correspond to CAE-Rev.3 2-digit codes 35, 36, 37, 38, 49, 51, 53, 59, 60, 61, 64, 72, 81, 84, 85, 86, 87, 88, and 99.

[19] See Tables E.1 and E.2 in the Appendix-E for coefficient estimates of contemporaneous response.

respective effects persist for at least five years after the exposure to a minimum wage increase. The magnitude of the employment responses are also much larger than the analogous sector-based results, particularly in IV estimates.

6 Conclusion

This paper set out to test whether labour market power is a relevant factor when studying minimum wage employment effects, and whether this could help reconcile the diverging empirical evidence. Using market concentration as a proxy for market power, I indeed find that whether minimum wages impact employment positively or negatively is dependent on the level of concentration in each market. A common result across my specifications is that higher minimum wage exposure leads to a drop in employment in highly competitive markets, and an employment increase in highly concentrated markets, despite most workers being located in markets with low concentration. For intermediate levels, employment effects are mostly statistically non-significant. Whether or not these effects are persistent depends on the definition of local labour market, although my benchmark definition suggests the effects are purely transitory.

Why does this matter? The differential effects of minimum wages based on market concentration would not matter if labour markets were homogeneous. However, local labour markets tend to be very heterogeneous in their level of concentration. This had been shown for the US by Azar et al. (2022a), and now I show this result also applies to another country with drastically different labour laws, Portugal. This heterogeneity is crucial and has large implications to how we understand the literature and consequential policy implications.

Minimum wage studies differ on whether they estimate nation-wide employment responses to the minimum wage (Dube et al., 2010), or whether they look at a particular geographical area or industry (Card and Krueger, 1994; Dube et al., 2007). The fact that concentration matters, and labour markets are heterogeneous in that respect, sheds new light into how we interpret these studies. Estimates of employment responses in local case studies will be affected by the concentration level in the particular labour market being looked at, which can explain why some find negative responses while others find positive. National-level studies will be affected by the underlying distribution of market concentration. If some local labour markets are impacted negatively by the minimum wage and others positively, when these effects are aggregated at a national scale, they will be at least partially cancelled out, which may explain the null employment effects often found in national estimates.

In terms of policy, my results may also have some implications regarding the social function of the minimum wage. Currently, minimum wages are seen mostly as an instrument that guarantees a base standard of living guaranteed to all workers. My results suggest that it may gain a new function, as with the existence of market power it can also be used as a tool to correct a market failure and consequently improve labour market efficiency, by bringing employment closer to its socially-optimum level. This is particularly useful given that competition authorities are mostly focused on product markets rather than factor markets.

To make the best use of this policy tool to fight market power, my results also suggest that a minimum wage applied uniformly at a national level - as is the case in Portugal - might not be the best system. Arguments for more decentralised wage minima varying across sector and/or region often invoke reasons such as heterogeneity in productivity and costs of living across different sectors and regions. My findings provide an additional argument in favour of decentralising minimum wages - heterogeneity in labour market concentration, which causes differential employment effects of minimum wages. Figures C.2 and C.3 in the Appendix-C showed that, in Portugal, regions and industries both vary quite a lot in terms of concentration, so these could be used as dimensions upon which to base different minimum wages. Policy-makers could take into account that in regions/industries with higher concentration, the trade-off between raising the minimum wage and employment is less of a concern and they can increase the minimum wage in those labour markets by more. They can also be aware that despite regions such as Lisbon having a higher cost of living, it also has very low labour market concentration, and so the employment costs of increasing the minimum wage may be steeper. Policy-makers may need to prioritise their conflicting social welfare goals.

Azar et al. (2023) began the groundbreaking path of understanding minimum wage employment effects through the lenses of labour market concentration. I expanded on their work by choosing a setting which allows for a more comprehensive coverage of the national labour force. Going forward, it would be beneficial for our comprehension of this issue if more evidence could be gathered from other countries and settings, and possibly using other methods to estimate labour market power.

References

Afonso, A. and Venâncio, A. (2016). The relevance of commuting zones for regional spending efficiency. *Applied Economics*, 48(10):865–877.

Alexandre, F., Bação, P., Cerejeira, J., Costa, H., and Portela, M. (2022). Minimum wage and financially distressed firms: another one bites the dust. *Labour Economics*, 74:102088.

Autor, D. H., Dorn, D., and Hanson, G. H. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American economic review*, 103(6):2121–2168.

Azar, J., Huet-Vaughn, E., Marinescu, I., Taska, B., and von Wachter, T. (2023). Minimum Wage Employment Effects and Labour Market Concentration. *The Review of Economic Studies*, page rdad091.

Azar, J., Marinescu, I., and Steinbaum, M. (2019). Measuring labor market power two ways. In *AEA Papers and Proceedings*, volume 109, pages 317–21.

Azar, J., Marinescu, I., and Steinbaum, M. (2022a). Labor market concentration. *Journal of Human Resources*, 57(S):S167–S199.

Azar, J. A., Berry, S. T., and Marinescu, I. (2022b). Estimating labor market power. Technical report, National Bureau of Economic Research.

Bassanini, A., Batut, C., and Caroli, E. (2023). Labor Market Concentration and Wages: Incumbents versus New Hires. *Labour Economics*, 81:102338.

Bassanini, A., Bovini, G., Caroli, E., Fernando, J., Cingano, F., Falco, P., Felgueroso, F., Jansen, M., Martins, P. S., Melo, A., et al. (2022). Labour Market Concentration, Wages and Job Security in Europe.

Benmelech, E., Bergman, N. K., and Kim, H. (2022). Strong employers and weak employees: How does employer concentration affect wages? *Journal of Human Resources*, 57(S):S200–S250.

Bhaskar, V., Manning, A., and To, T. (2002). Oligopsony and monopsonistic competition in labor markets. *Journal of Economic Perspectives*, 16(2):155–174.

Brown, C., Gilroy, C., and Kohen, A. (1982). The effect of the minimum wage on employment and unemployment: a survey. Technical report, National Bureau of Economic Research.

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

Card, D. and Krueger, A. B. (1994). Minimum Wages and Employment: A Case Study of the Fast-Food Industry in New Jersey and Pennsylvania. *The American Economic Review*, 84(4):772–793.

Card, D. and Krueger, A. B. (1995). *Myth and Measurement: The New Economics of the Minimum Wage - Twentieth-Anniversary Edition*. Princeton University Press, REV - Revised edition.

Carneiro, A., Sá, C. A. d. S. P., Cerejeira, J., Varejão, J., Portela, M., et al. (2011). Estudo sobre a retribuição mínima mensal garantida em Portugal. Technical report, Ministério da Economia e do Emprego.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *The Quarterly Journal of Economics*, 134(3):1405–1454.

Centeno, M., Duarte, C., and Novo, A. (2011). The impact of the minimum wage on low-wage earners. *Banco de Portugal Economic Bulletin Autumn*, pages 107–121.

Cerejeira, J. (2008). Young Employment, Job-Skill Composition and Minimum Wages: Evidence from a 'Natural Experiment'. NIPE Working Papers 04/2008, NIPE - Universidade do Minho.

Cerejeira, J., Kizilca, K., Portela, M., and Sá, C. (2012). Minimum Wage, Fringe Benefits, Overtime Payments and the Gender Wage Gap. IZA Discussion Papers 6370, Institute of Labor Economics (IZA).

Dube, A., Lester, T. W., and Reich, M. (2010). Minimum wage effects across state borders: Estimates using contiguous counties. *The review of economics and statistics*, 92(4):945–964.

Dube, A., Naidu, S., and Reich, M. (2007). The economic effects of a citywide minimum wage. *ILR Review*, 60(4):522–543.

Félix, S. and Portugal, P. (2016). Labor Market Imperfections and the Firm's Wage Setting Policy. IZA Discussion Papers 10241, Institute of Labor Economics (IZA).

Giuliano, L. (2013). Minimum wage effects on employment, substitution, and the teenage labor supply: Evidence from personnel data. *Journal of Labor Economics*, 31(1):155–194.

Harasztosi, P. and Lindner, A. (2019). Who pays for the minimum wage? *American Economic Review*, 109(8):2693–2727.

Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American economic review*, 95(1):161–182.

Manning, A. (2011). Imperfect competition in the labor market. In *Handbook of labor economics*, volume 4, pages 973–1041. Elsevier.

Manning, A. and Petrongolo, B. (2022). Monopsony in Local Labour Markets. *IFS Deaton Review of Inequalities, Institute for Fiscal Studies*.

Marinescu, I., Ouss, I., and Pape, L.-D. (2021). Wages, hires, and labor market concentration. *Journal of Economic Behavior & Organization*, 184:506–605.

Martins, P. S. (2021). 30,000 minimum wages: The economic effects of collective bargaining extensions. *British Journal of Industrial Relations*, 59(2):335–369.

Martins, P. S. and Melo, A. P. (2023). Making their own weather? Estimating employer labour-market power and its wage effects. *Nova SBE Working Paper Series*, (659).

Neumark, D. and Wascher, W. (1992). Employment effects of minimum and subminimum wages: panel data on state minimum wage laws. *ILR Review*, 46(1):55–81.

Neumark, D. and Wascher, W. (2007). Minimum wages, the earned income tax credit, and employment: evidence from the post-welfare reform era.

Nevo, A. (2001). Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2):307–342.

Oliveira, C. (2023). The Minimum Wage and the Wage Distribution in Portugal. Technical report.

Pereira, S. C. (2003). The impact of minimum wages on youth employment in Portugal. *European Economic Review*, 47(2):229–244.

Portugal, P. and Cardoso, A. R. (2006). Disentangling the minimum wage puzzle: an analysis of worker accessions and separations. *Journal of the European Economic Association*, 4(5):988–1013.

Reis, R., Bação, P., Horta Correia, I., Valério, N., Varejão, J., Tavares, J., Aguiar-Conraria, L., and Ferreira, J. A. (2023). Crises na economia portuguesa. Estudo FFMS, Fundação Francisco Manuel dos Santos.

Rinz, K. et al. (2018). Labor market concentration, earnings inequality, and earnings mobility. *Center for Administrative Records Research and Applications Working Paper*, 10.

Stewart, M. B. (2004). The impact of the introduction of the UK minimum wage on the employment probabilities of low-wage workers. *Journal of the European Economic Association*, 2(1):67–97.

A Model Figures

Figure A.1: Firm Choice of Employment as a Function of Minimum Wage

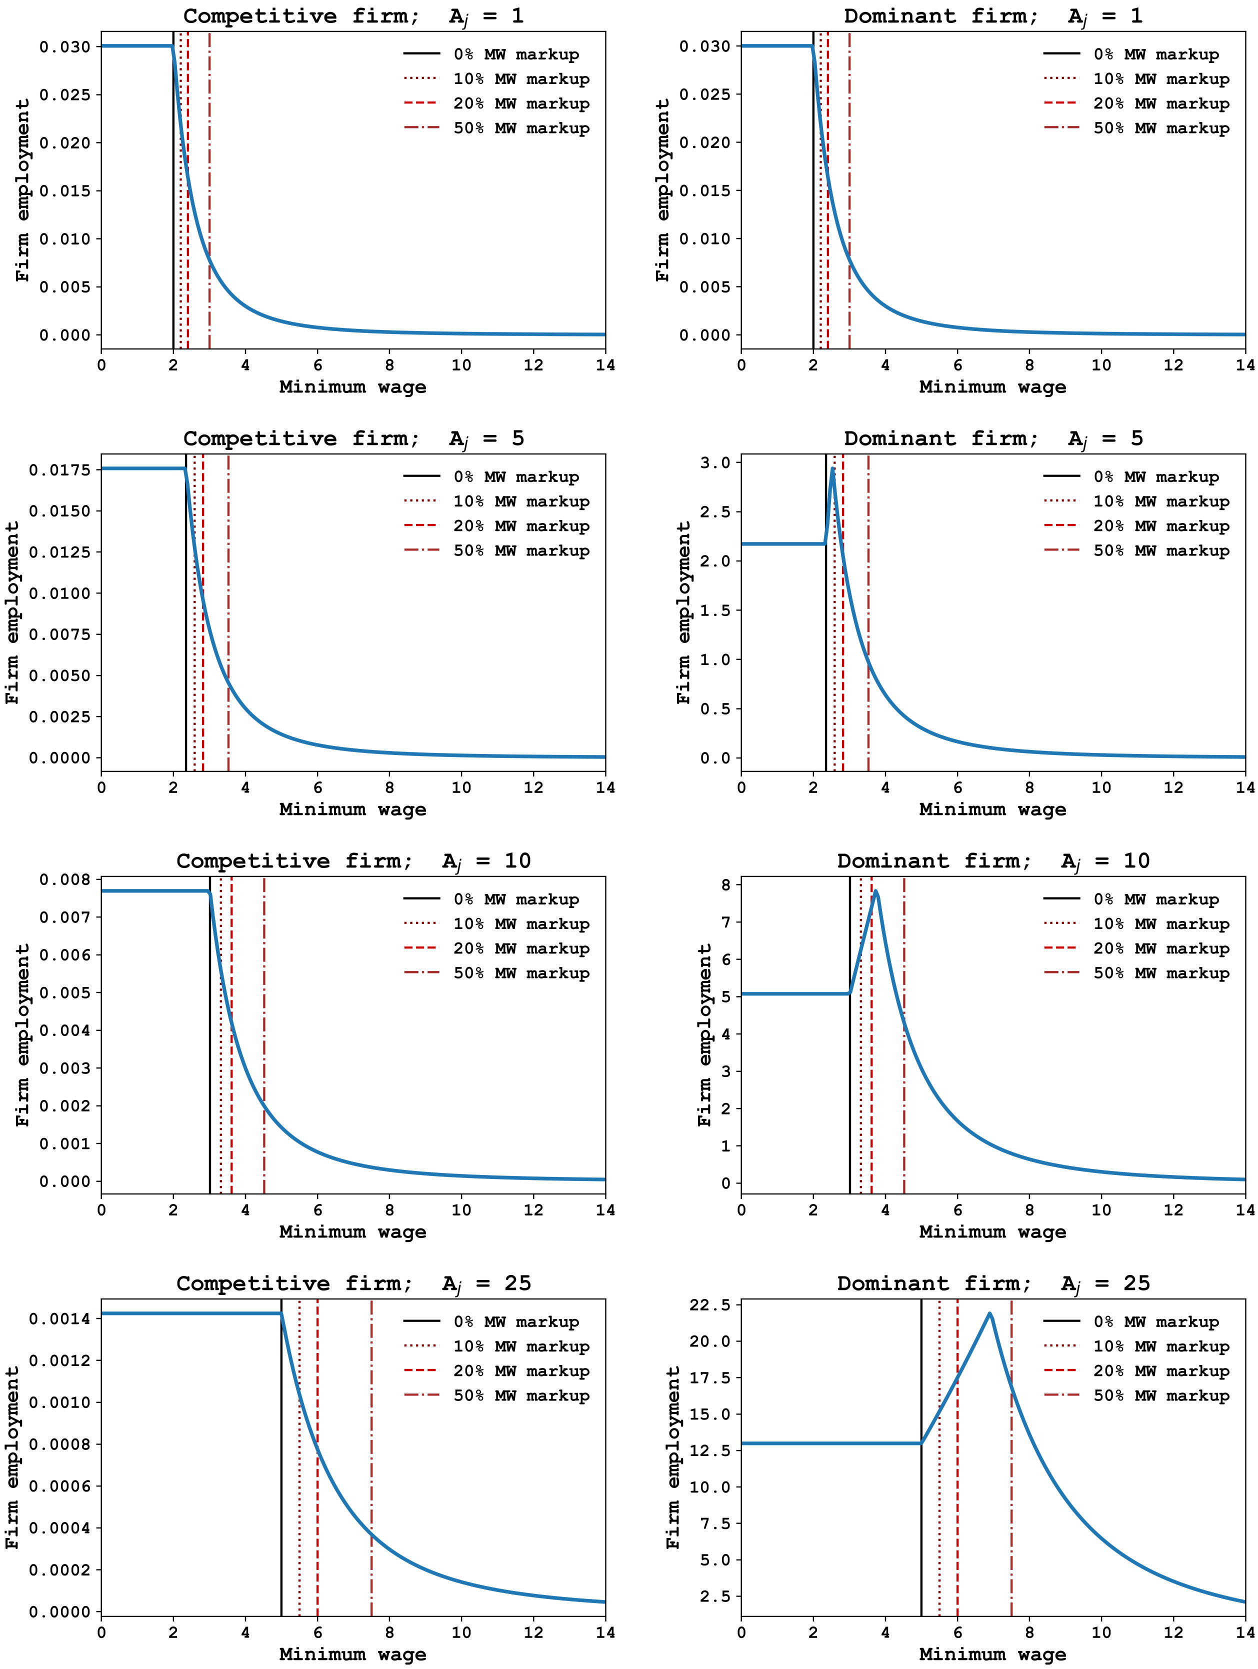

B Local Labour Markets

Figure B.1: The 63 Commuting Zones of Portugal

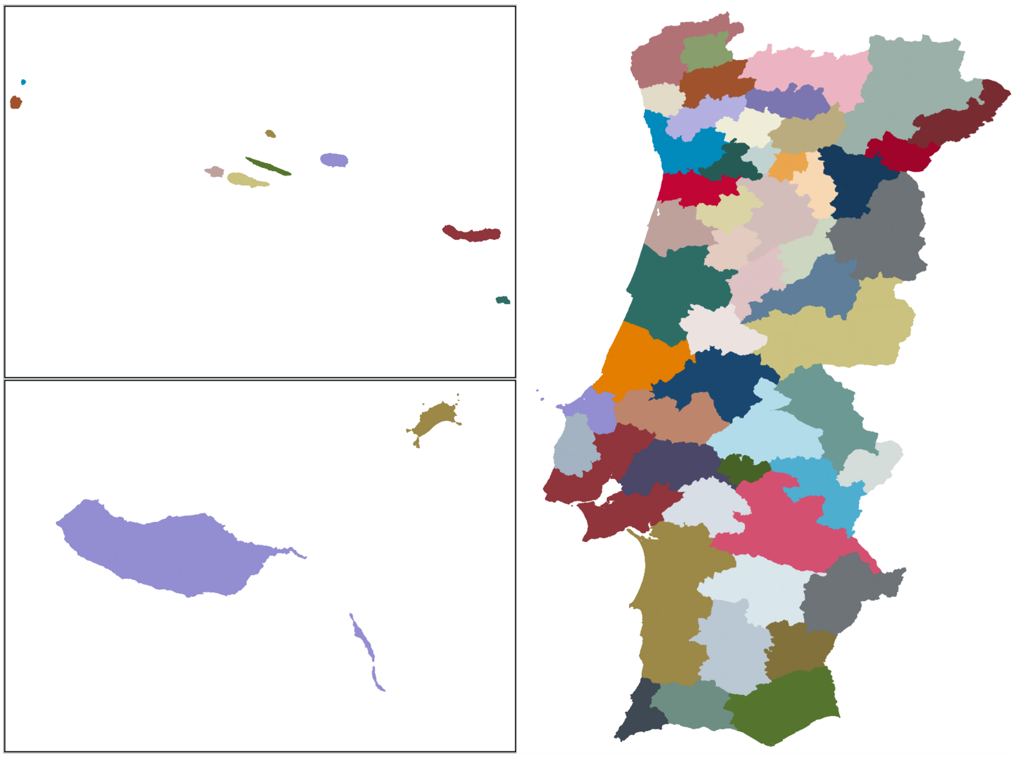

Source: Author's own map, based on Afonso and Venâncio (2016).
Note: Islands are not to scale.

C HHI and MW Exposure

Figure C.1: Temporal Evolution

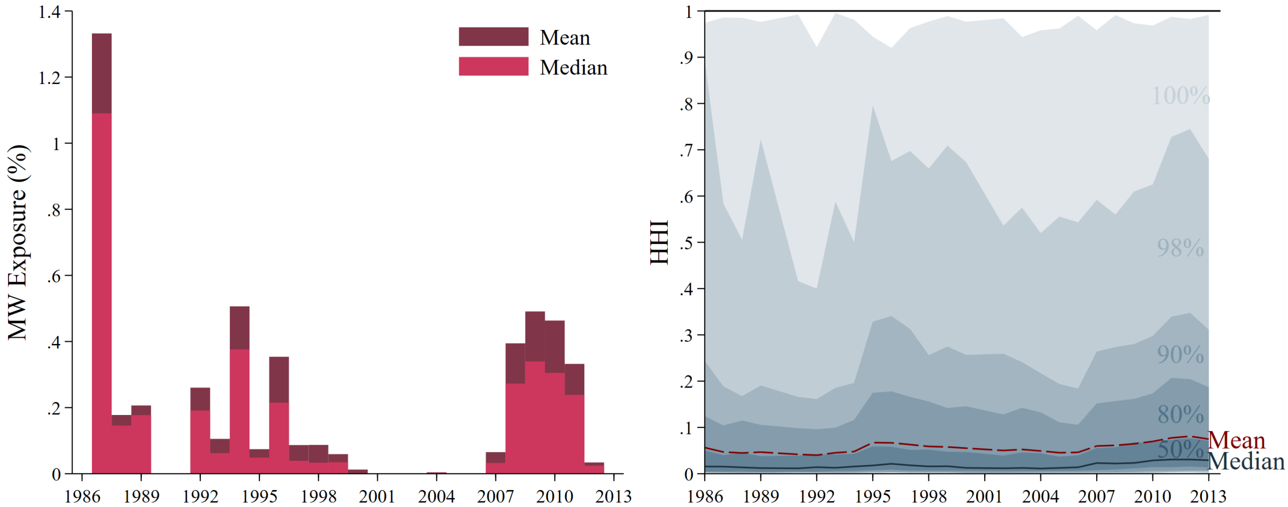

(a) MW Exposure (b) HHI

Note: Employment-weighted average.

Figure C.2: Regional Dispersion

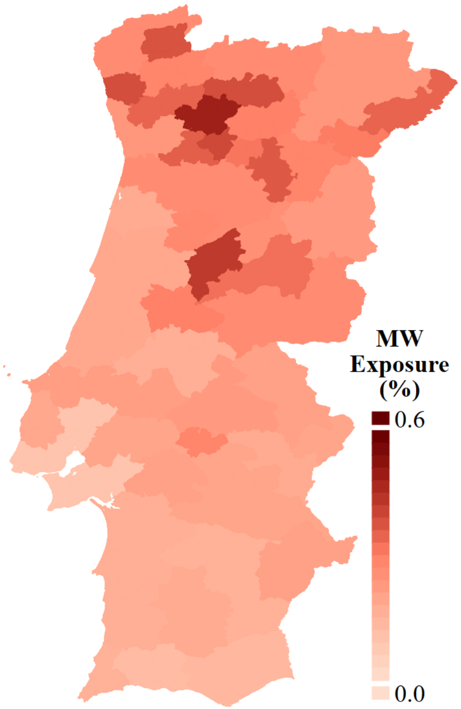

(a) MW Exposure

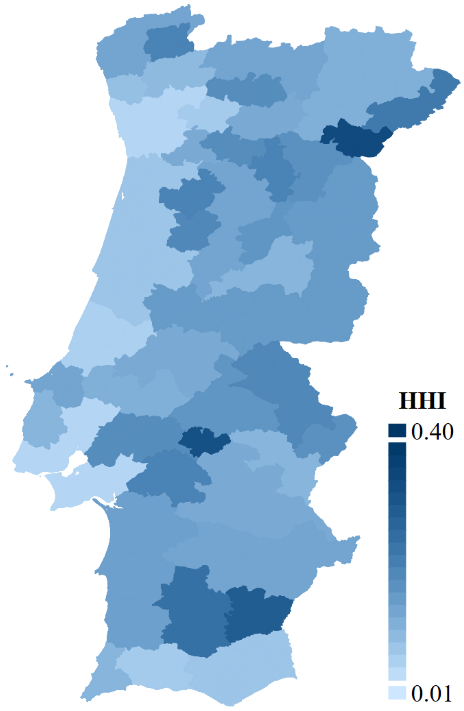

(b) HHI

Note: Employment-weighted average.

Figure C.3: Industry Dispersion

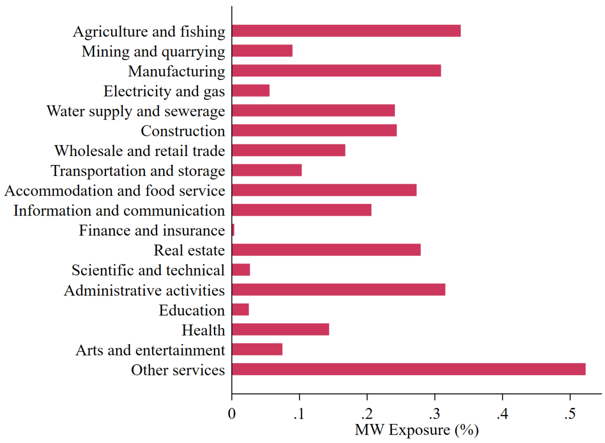

(a) MW Exposure

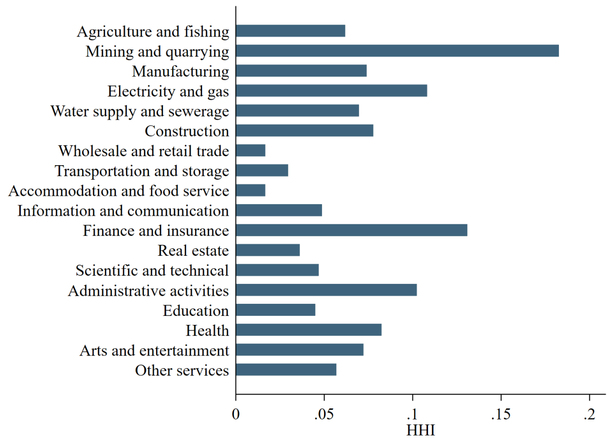

(b) HHI

Note: Employment-weighted average.

D Additional Regression Tables and Figures

Table D.1: OLS Estimates

Δ Log(Employment)	(1)	(2)	(3)	(4)
MW Exposure	-1.434	-1.437	-3.870***	-10.224***
	(0.897)	(0.909)	(1.014)	(1.978)
HHI	0.136***	0.182***	0.003	-0.126***
	(0.030)	(0.028)	(0.039)	(0.044)
MW Exposure × HHI	8.074**	4.507	8.758**	25.871***
	(3.477)	(3.627)	(3.900)	(7.217)
Additional Controls		✓	✓	✓
Fixed Effects			✓	✓
Lags MW Exposure				✓
Observations	31,640	31,611	31,465	20,560
Adjusted R^2	0.004	0.103	0.176	0.183

Standard errors in parentheses are clustered at the sector (1 letter) × NUTS-II level
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table D.2: IV Estimates Compared

Δ Log(Employment)	OLS	IV
MW Exposure	-10.224***	-14.989***
	(1.978)	(3.730)
HHI	-0.126***	-0.611***
	(0.044)	(0.123)
MW Exposure × HHI	25.871***	39.711***
	(7.217)	(15.104)
Additional Controls	✓	✓
Fixed Effects	✓	✓
Lags MW Exposure	✓	✓
Observations	20,560	20,560
Adjusted R^2	0.183	0.077

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table D.3: IV First-Stage

	HHI	MW Exposure × HHI
Average Log(1/N)	0.138***	0.000***
	(0.006)	(0.000)
MW Exposure × Average Log(1/N)	-0.761***	0.096***
	(0.220)	(0.008)
Observations	20,560	20,560
Adjusted R^2	0.655	0.723
H0 = Under-identified		
Sanderson-Windmeijer χ^2	631.63	115.33
P-value	0.000	0.000
H0 = Weakly identified		
Sanderson-Windmeijer F statistic	625.48	114.20

Standard errors are clustered at the sector (1 letter) × NUTS-II level

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure D.1: Pre-Trends for Dynamic Responses

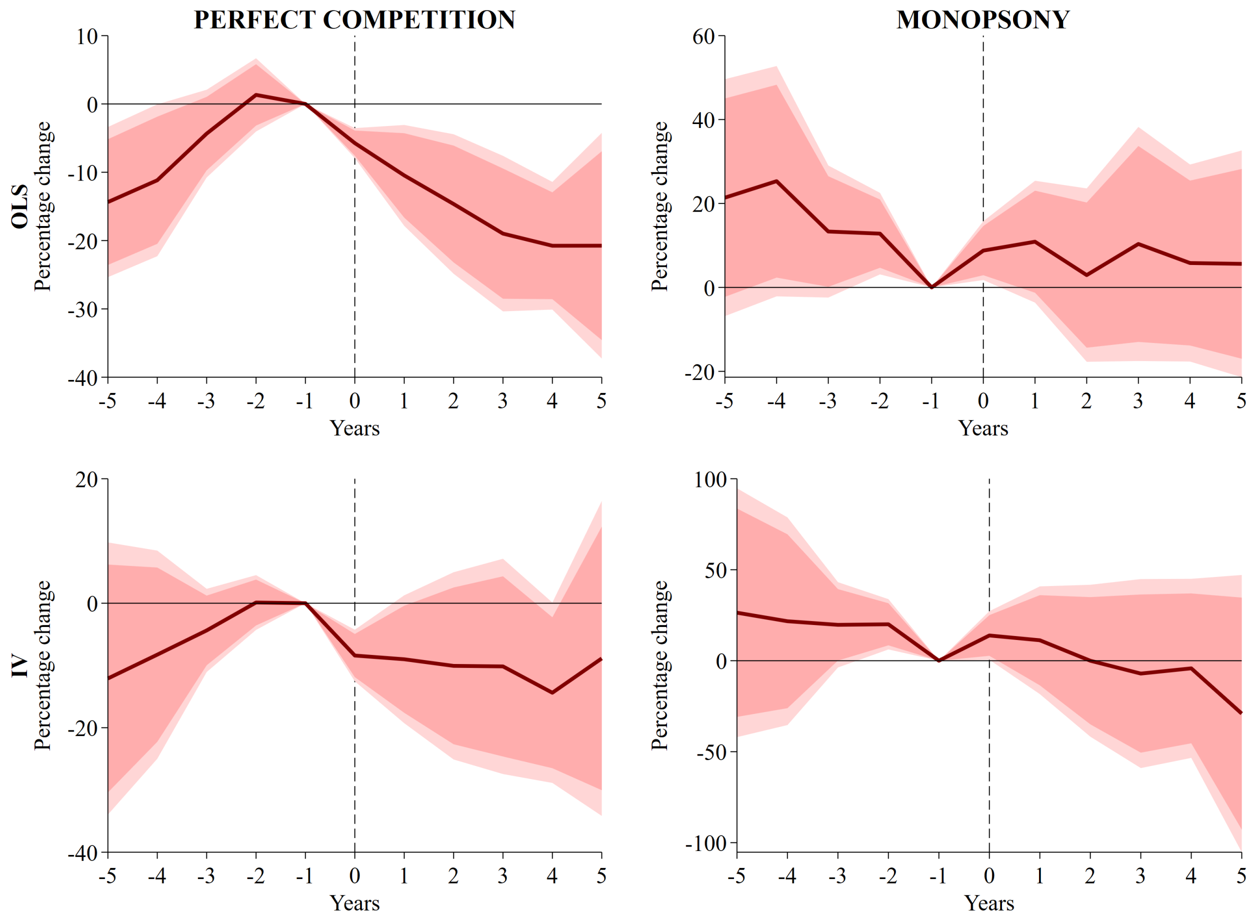

Note: Dark pink and light pink shaded areas show 90% and 95% confidence bands, respectively. In the case of perfect competition, impulse response at horizon h corresponds to $\beta_{1,h}$. In the case of monopsony, the impulse response is plotting $\beta_{1,h} + \beta_{3,h}$. Standard errors are clustered at the sector (1 letter) by NUTS-II level.

E Robustness Checks

Table E.1: Robustness Checks - OLS Estimates

Δ Log(Employment)	Benchmark	All LLM	No Pub. Sect.	Empl. HHI	Occupations
MW Exposure	-10.224***	-8.624***	-8.855***	-9.532***	-5.041**
	(1.978)	(2.693)	(2.239)	(1.995)	(2.160)
HHI	-0.126***	-0.375***	-0.123**	-0.017	-0.281***
	(0.044)	(0.036)	(0.052)	(0.059)	(0.053)
MW Exposure × HHI	25.871***	16.504***	23.254***	33.844***	4.882
	(7.217)	(4.102)	(7.961)	(10.216)	(7.230)
Additional Controls	✓	✓	✓	✓	✓
Fixed Effects	✓	✓	✓	✓	✓
Lags MW Exposure	✓	✓	✓	✓	✓
Observations	20,560	28,514	16,496	20,560	24,941
Adjusted R^2	0.183	0.204	0.193	0.183	0.079

Standard errors in parentheses are clustered at the sector/occupation (1 letter/digit) × NUTS-II level
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table E.2: Robustness Checks - IV Estimates

Δ Log(Employment)	Benchmark	All LLM	No Pub. Sect.	Empl. HHI	Occupations
MW Exposure	-14.989***	-17.304***	-15.181***	-13.283***	-7.388*
	(3.730)	(3.688)	(4.123)	(3.089)	(3.835)
HHI	-0.611***	-0.519***	-0.747***	-0.644***	-1.895***
	(0.123)	(0.077)	(0.159)	(0.137)	(0.174)
MW Exposure × HHI	39.711***	33.473***	42.335**	43.326**	11.804
	(15.104)	(7.945)	(17.428)	(16.826)	(20.984)
Additional Controls	✓	✓	✓	✓	✓
Fixed Effects	✓	✓	✓	✓	✓
Lags MW Exposure	✓	✓	✓	✓	✓
Observations	20,560	28,514	16,496	20,560	24,941
Adjusted R^2	0.077	0.093	0.079	0.073	-0.117

Standard errors in parentheses are clustered at the sector/occupation (1 letter/digit) × NUTS-II level
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure E.1: Robustness Checks - Contemporaneous Response

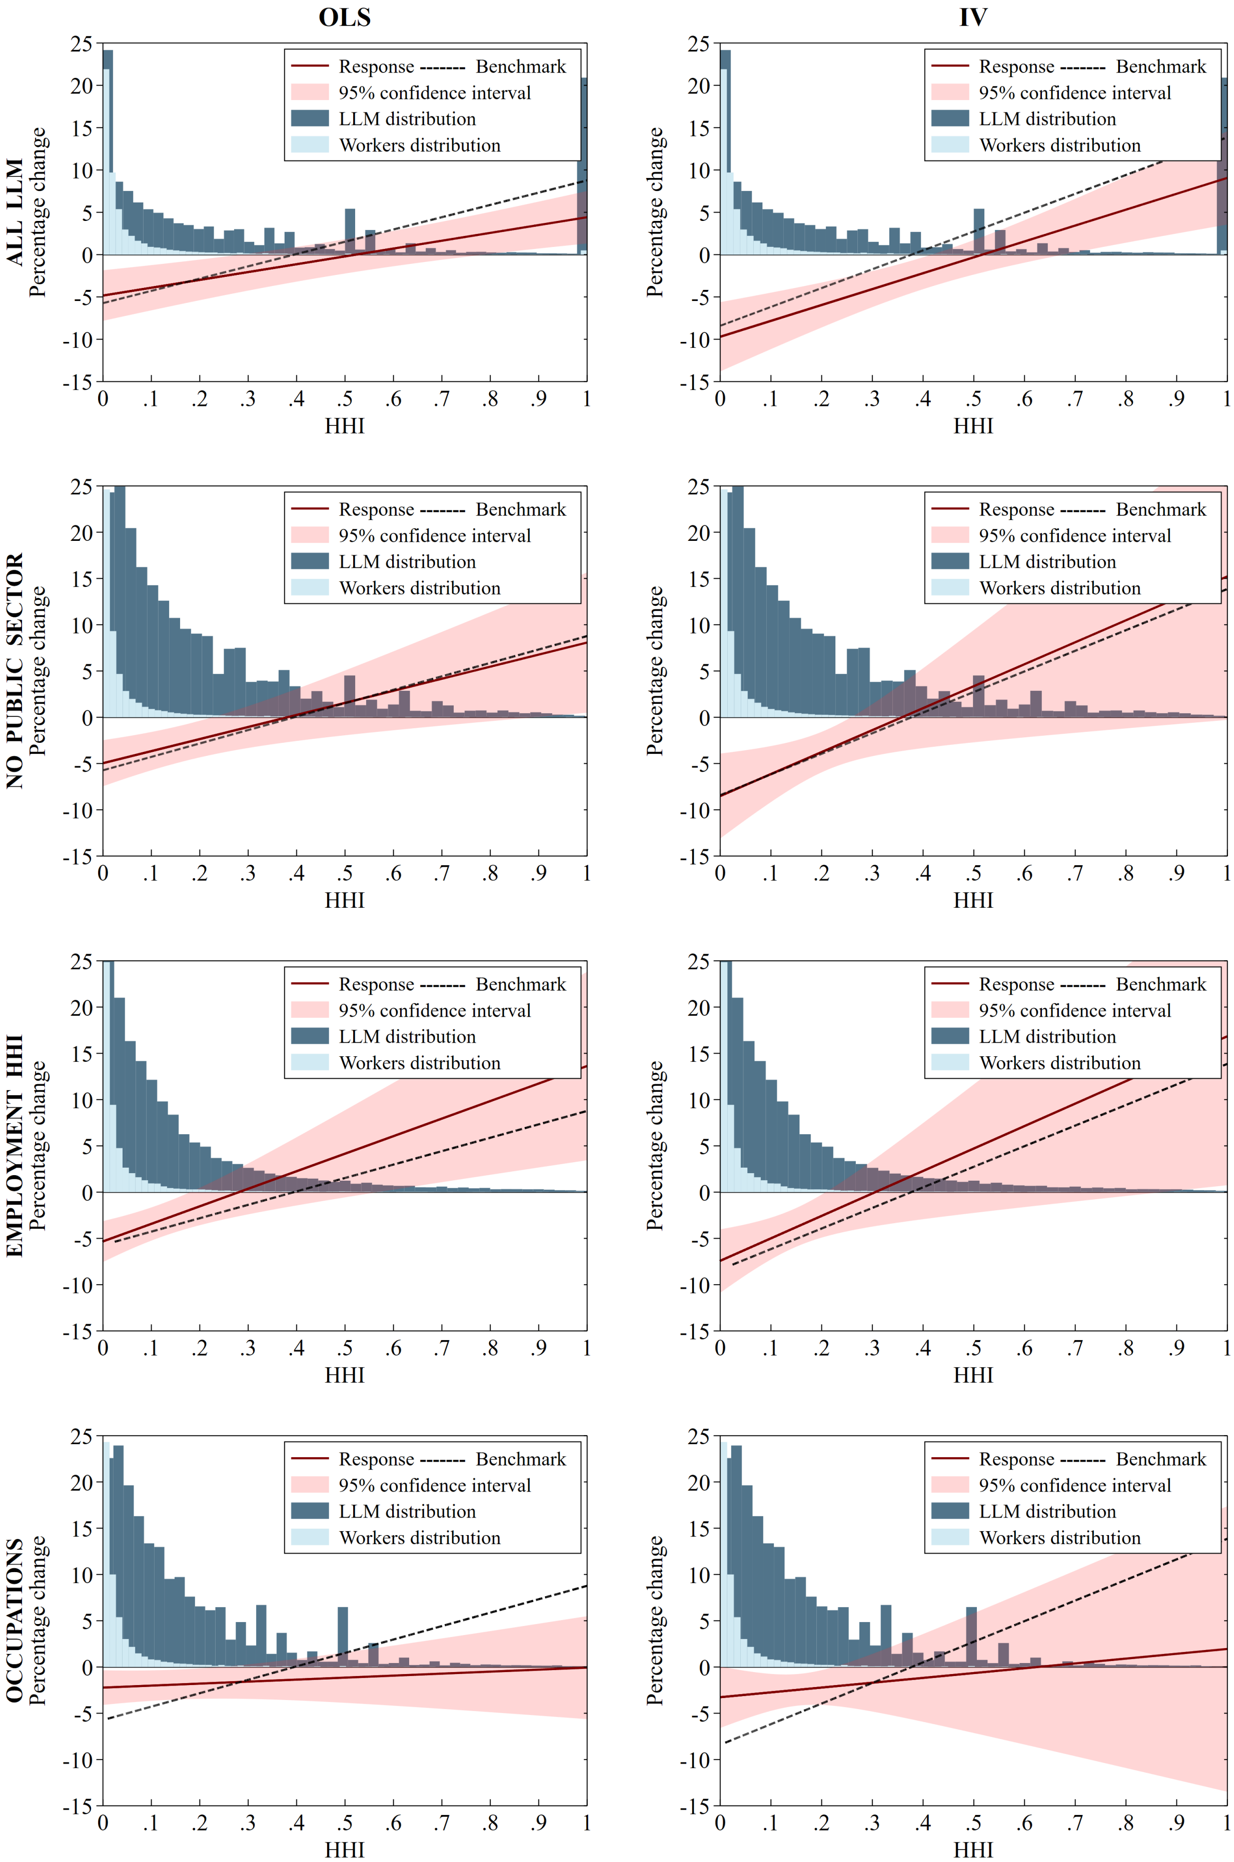

Note: Figures are plotting the estimates for $\beta_1 + \beta_3 \times \text{HHI}$. Dashed lines show the analogous benchmark response for an impulse of the same magnitude. Standard errors are clustered at the sector/occupation (1 letter/digit) by NUTS-II level.

Figure E.2: Robustness Checks - OLS Dynamic Responses

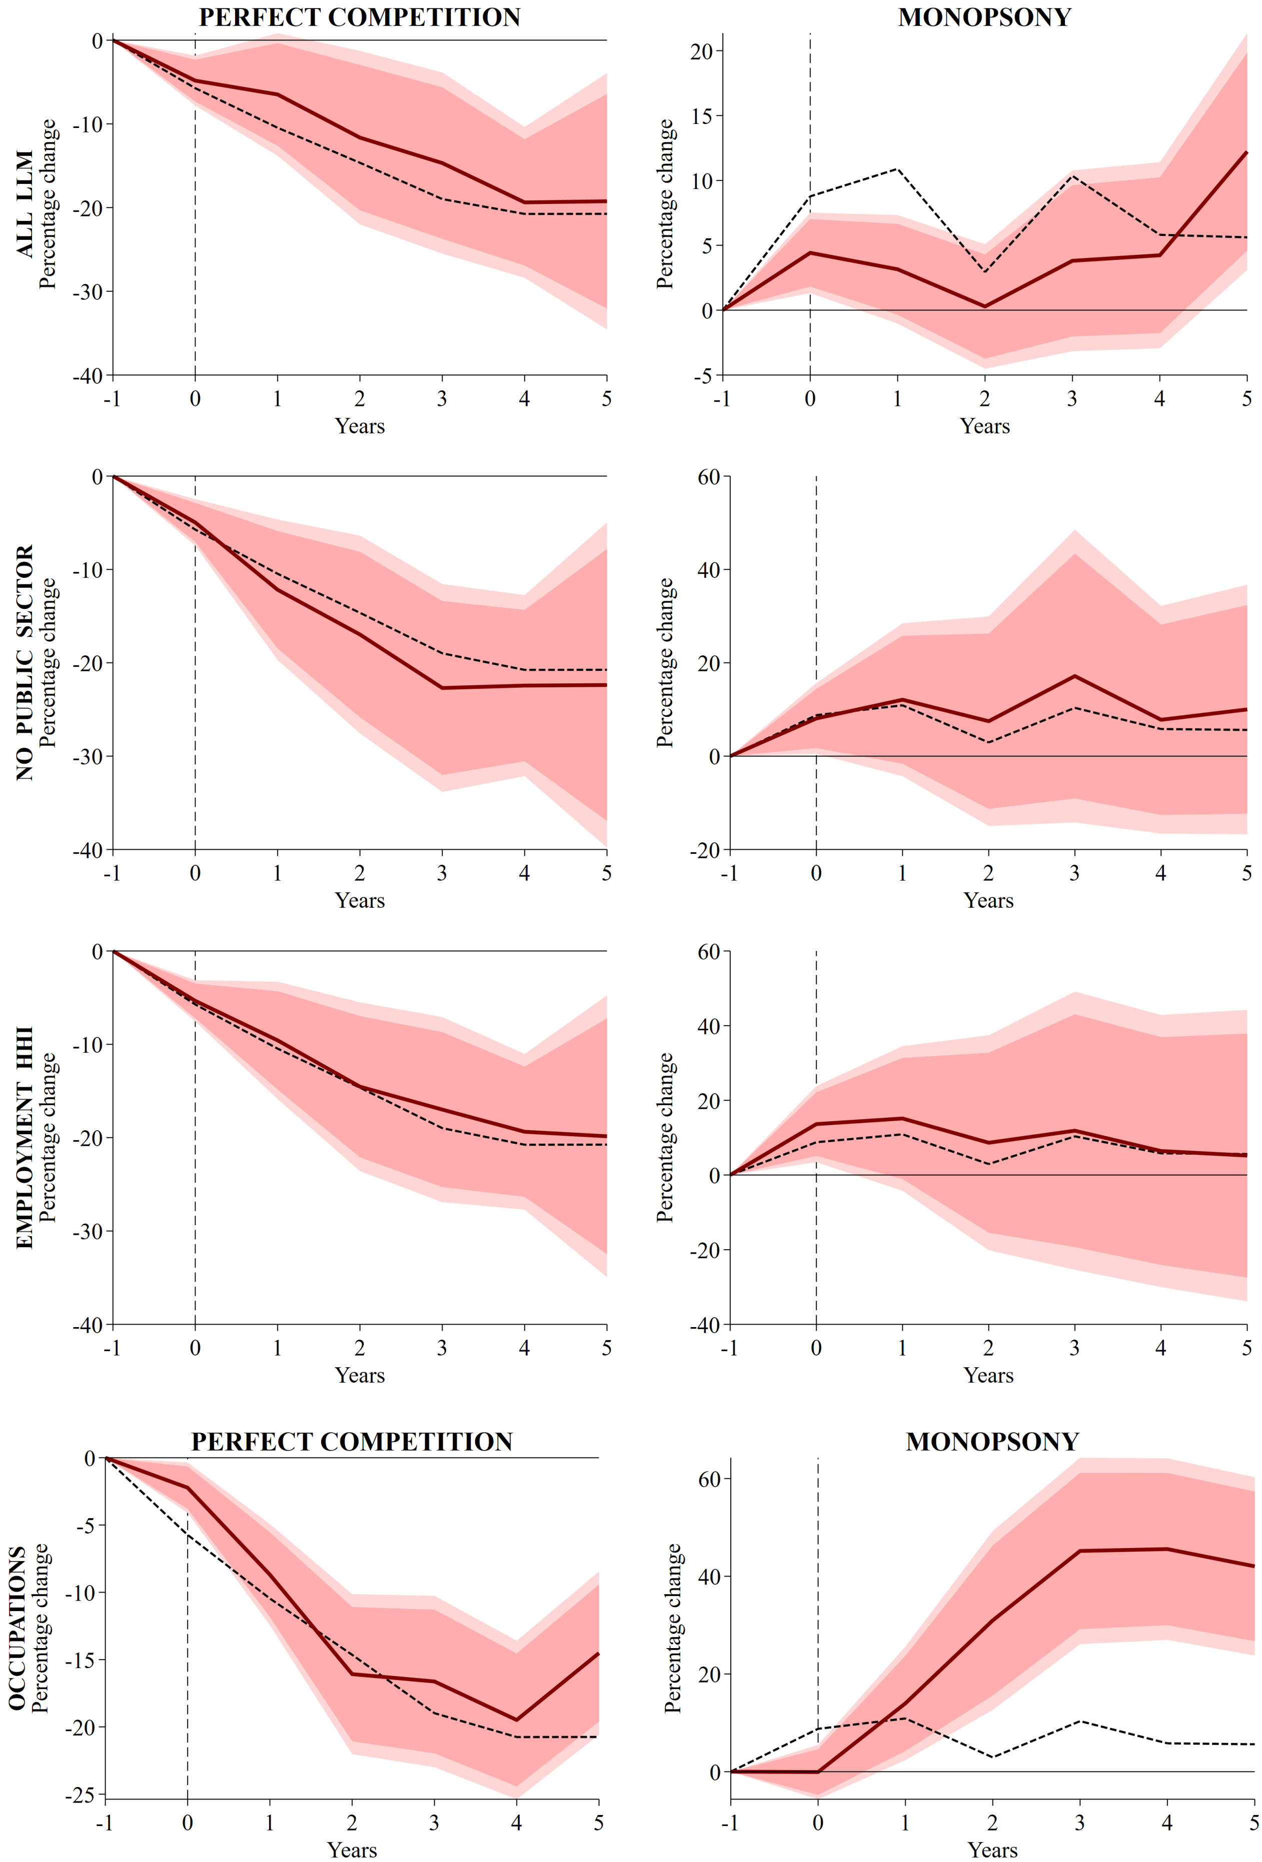

Note: Dark pink and light pink shaded areas show 90% and 95% confidence bands, respectively. Dashed lines show the analogous benchmark response for an impulse of the same magnitude. In the case of perfect competition, impulse response at horizon h corresponds to $\beta_{1,h}$. In the case of monopsony, the impulse response is plotting $\beta_{1,h} + \beta_{3,h}$. Standard errors are clustered at the sector/occupation (1 letter/digit) by NUTS-II level.

Figure E.3: Robustness Checks - IV Dynamic Responses

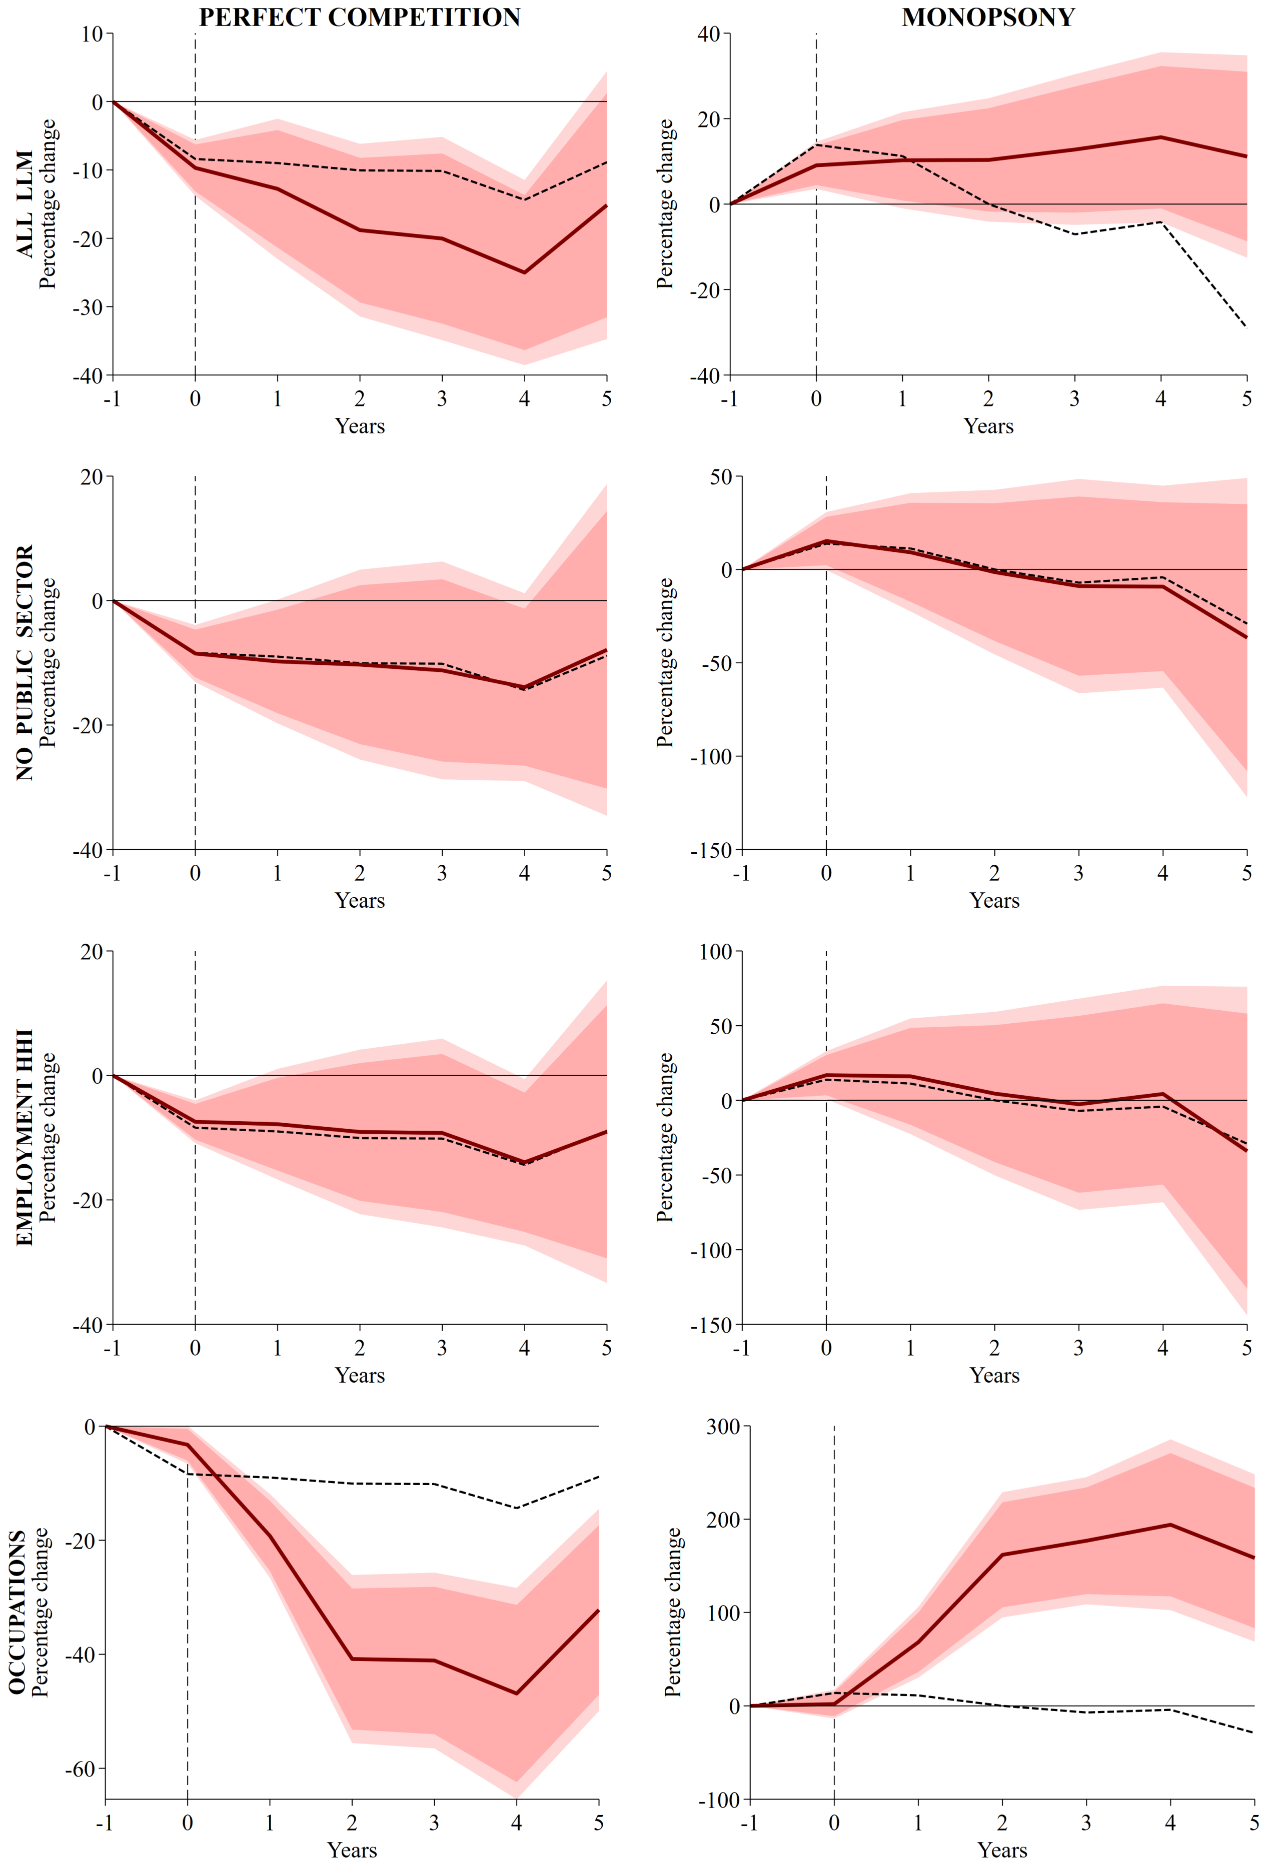

Note: Dark pink and light pink shaded areas show 90% and 95% confidence bands, respectively. Dashed lines show the analogous benchmark response for an impulse of the same magnitude. In the case of perfect competition, impulse response at horizon h corresponds to $\beta_{1,h}$. In the case of monopsony, the impulse response is plotting $\beta_{1,h} + \beta_{3,h}$. Standard errors are clustered at the sector/occupation (1 letter/digit) by NUTS-II level.