Labour Market Power and Aggregate Productivity

Bernardo Mottironi[+]

December 1, 2024

[last version here]

Abstract

This paper offers a novel perspective on the general equilibrium effects of labour market power. By developing a tractable model of entrepreneurship with monopsonistic labour markets and endogenous technology adoption, I show that labour market power diminishes aggregate productivity through three distinct channels: (i) misallocation of workers towards small firms, (ii) excess entry of low-ability entrepreneurs, and (iii) limited diffusion of productivity-enhancing technologies. The proposed theory generates testable predictions, which I validate with novel evidence from Italian microdata: at the province-level, weaker competition in labour markets is associated with greater misallocation, more entrepreneurship, reduced firm size, inferior use of intangibles, and lower productivity. To quantify aggregate losses, I calibrate the model with measures of market power from financial statements, information on IT adoption from survey data, and expenditure shares from national accounts. My results reveal a significant impact of labour market power: aggregate productivity in a typical province is 21% lower than in a competitive benchmark, with excess entry and limited technology diffusion being the dominant factors.

Keywords: *Monopsony*; *misallocation*; *firm size*; *technology diffusion*; *productivity*

JEL Codes: *D61; E24; J42; L11; O33*

[+]London School of Economics and Political Science (LSE). Centre for Economic Performance (CEP). Email: b.mottironi@lse.ac.uk. I am profoundly grateful to my supervisors, John Van Reenen, Maarten De Ridder, and Alan Manning, for their guidance. This work benefited from fruitful discussions with David Berger, Giuseppe Berlingieri, Eduard Boehm, Tito Boeri, Filippo Di Mauro, Matthias Doepke, Gaia Dossi, Arnaud Dyèvre, Jan Eeckhout, Luis Garicano, Peter Lambert, Paolo Mengano, Matthias Mertens, Ben Moll, Jane Olmstead-Rumsey, Gianmarco Ottaviano, Giorgio Pietrabissa, Hugo Reichardt, Ricardo Reis, Yannick Schindler, Abhijit Tagade, and Maria Ventura. I acknowledge the generous financial support from the UniCredit Foundation.

1 Introduction

A growing body of evidence documents two key patterns on firm size and performance. At the firm-level, larger companies demonstrate greater capacity for innovation and technology adoption, particularly in relation to intangible assets (Hsieh and Rossi-Hansberg, 2023; De Ridder, 2024). At the industry-level, market concentration is positively associated with aggregate productivity (Autor et al., 2020; Bighelli et al., 2023). These empirical regularities, illustrated in Figures 1 and 2, raise a critical question: is the proliferation of underperforming micro firms the result of specific market frictions? As smaller businesses often struggle to absorb the sunk costs required for technical upgrades, distortions that divert resources away from larger firms may hinder the spread of new technologies, resulting in significant productivity losses.[1]

This paper offers a novel perspective on the issue, starting from a simple observation: the negative impact of monopsony power on wages generates an incentive for entrepreneurship over labour. Firms in several contexts possess substantial pricing power in labour markets, which allows them to set wages below the competitive level, as demonstrated by a solid theoretical and empirical literature (Manning, 2011; Lamadon et al., 2022). This distortion could increase the appeal of entrepreneurship in an occupational choice, inflating the number of active firms beyond optimality, with adverse consequences for technology diffusion and, ultimately, aggregate productivity.

To test and quantify this mechanism, I develop and calibrate a general equilibrium model of entrepreneurship, incorporating monopsony power and endogenous technology adoption in the canonical framework of Lucas (1978). In this setting, agents are endowed with heterogeneous levels of entrepreneurial ability, or 'talent for managing' in Lucas's words. As this ability positively influences firm productivity, only the most talented agents can generate large profits and benefit from entrepreneurship, whereas others are better off earning wages. Technological upgrades, modelled following Bustos (2011), allow firms to increase their productivity by paying a fixed cost. Labour market power arises from job differentiation due to worker-specific taste for job amenities, as in Card et al. (2018).[2] Because of dispersion in taste, workers see employers as poor substitutes.

[1] A similar argument is presented in a recent policy report by former ECB President, Mario Draghi, which identifies the scale of European companies as a major barrier to the diffusion of innovation across firms (Draghi, 2024).

[2] Idiosyncratic taste is an intuitive and tractable source of monopsony power, commonly used in the recent literature. However, the theoretical mechanism and empirical tests do not strictly depend on this specific microfoundation.

Eventually, limited substitutability between jobs reduces labour market competition, with distortionary effects on market outcomes.[3]

The model has closed-form solution, which I use to derive the general equilibrium effects of monopsony power analytically.[4] My theoretical results reveal that imperfect competition in labour markets diminishes aggregate productivity through three distinct channels: (i) misallocation of workers towards small firms, (ii) excess entry of low-ability entrepreneurs, and (iii) limited diffusion of productivity-enhancing technologies. Misallocation and excess entry stem from the way in which monopsony power affects labour costs, whereas the reduced technology diffusion follows from the size distortions induced by the first two.

With limited substitutability between employers, labour supply is not perfectly elastic. As employers set wages along an upward-sloping curve, smaller firms pay lower wages than larger firms. This disproportion prevents large employers from fully expanding their market shares, partially shielding small businesses from competition. I define this channel as misallocation because the employment distribution is distorted towards smaller, less efficient firms.[5] Additionally, the positive slope of the inverse labour supply creates a wedge between the marginal and average cost of labour, pushing wages below the marginal revenue product of labour (MRPL). By lowering salaries, monopsony power induces a larger share of agents to become entrepreneurs, including the less talented. This channel, which I define as excess entry, makes firms smaller and more poorly managed on average.

Reduced technology diffusion is the third channel in the impact of labour market power on productivity. Firms can adopt a productivity-enhancing technology and lower their marginal costs by paying a sunk cost. The profitability of this choice depends on firm size: the fixed cost of upgrading is equal for all companies, but the fall in marginal costs generates more significant benefits for the largest. As a result, only firms above a certain size threshold use the superior technology. By inducing the proliferation of small businesses, labour market power hinders technology diffusion across firms and workers, with detrimental effects on aggregate productivity.

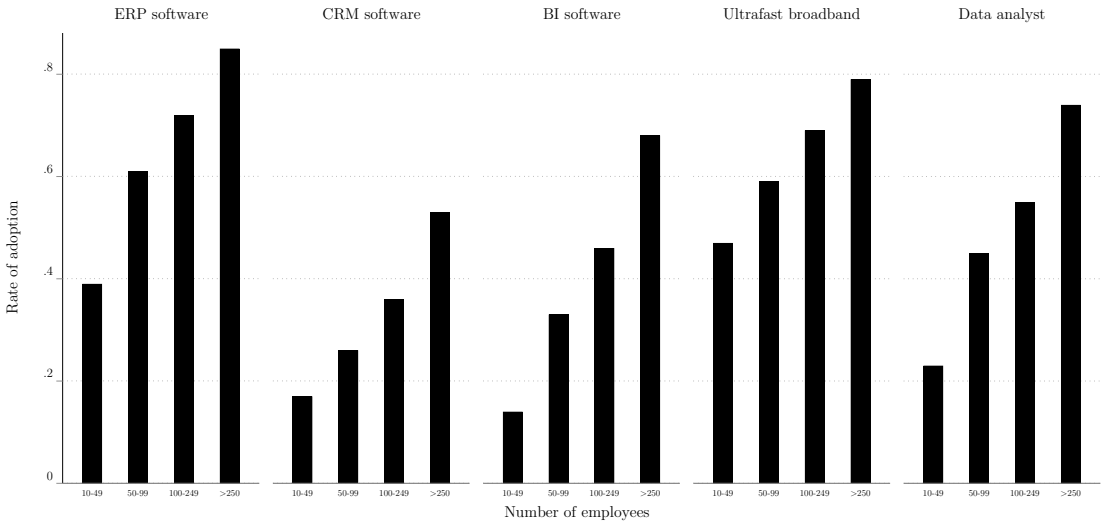
---

[3]Labour market power originates from preferences, analogous to how monopolistic competition emerges from product differentiation. Nevertheless, taste here should not be seen as independent from policy choices. For instance, reducing commuting costs and times could increase job substitutability and competition (Datta, 2024).

[4]Labour market competition is monopsonistic. In Appendix B, I confirm the key theoretical insights in an oligopsonistic framework with variable wage markdowns.

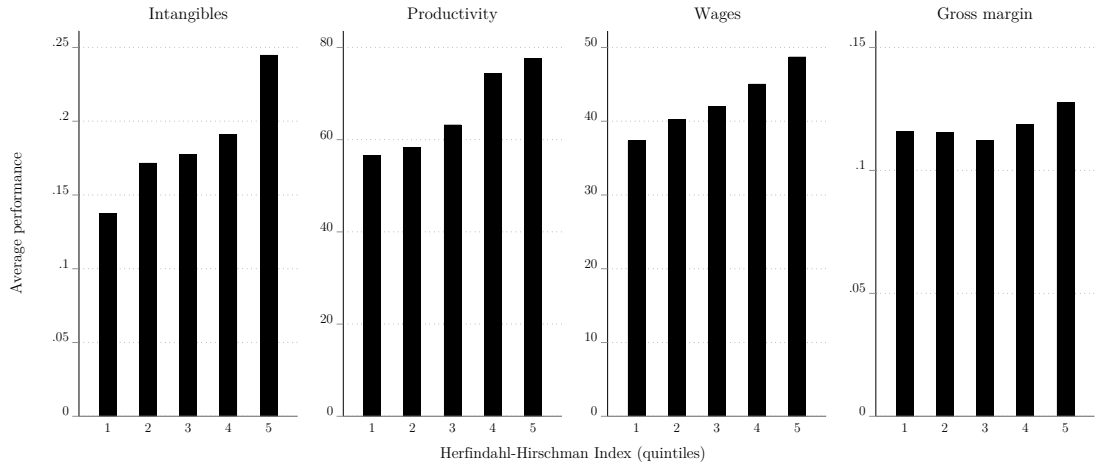
[5]A standard result in the market power literature is that homogenous markups from a CES demand do not cause misallocation, as they do not generate dispersion in marginal revenue products across firms (Edmond et al., 2023). However, this does not apply to labour markets, because homogenous markdowns from a CES labour supply induce dispersion in wages and, consequently, in MRPL and TFPR. I elaborate on this in Section 2.10.

FIGURE 1. TECHNOLOGY ADOPTION BY SIZE CLASS



Note: Columns represent the share of firms adopting the following technologies: (i) Enterprise resource planning (ERP) software; (ii) Customer relationship management (CRM) software; (iii) Business intelligence (BI) software; (iv) Ultrafast broadband (>100 Mbps); (v) Internal or external specialist of data analysis. Size classes are defined by the number of employees. Survey data on Italian firms in 2023 from ISTAT.

FIGURE 2. INDUSTRY PERFORMANCE BY QUINTILE OF MARKET CONCENTRATION



Note: Columns represent the average industry-level ratios of (i) intangible fixed assets to total fixed assets; (ii) value added to employment (thousands of Euros); (iii) labour costs to employment (thousands of Euros); and (iv) variable profits to revenues. 2-digit industries are split in quintiles by levels of sales-based HHI. Microdata on Italian industries in 2019 from CompNet.

The proposed theory generates a set of testable predictions, which I validate using microdata on Italian firms from the CompNet database (Melitz et al., 2018; Mertens et al., 2022).[6] The scope of my reduced-form analysis is to test whether monopsony power is associated with the predicted outcomes, comparing local labour markets in Italy. First, I classify administrative provinces by the level of labour market competition, measured using average wage markdowns (i.e. ratios of MRPL to wage, equal to one in perfect competition and higher under monopsony). In particular, firm-level markdowns are derived from the elasticities of revenue functions, estimated using a control function approach (Petrin and Sivadasan, 2013; Hashemi et al., 2022).[7] The parameters of revenue elasticities are measured by industry with a translog specification, which allows the elasticities to be nonlinear and heterogeneous across firms and periods.

Considering Italian labour markets in 2019, the last pre-pandemic period, 90% of province-level average markdowns fall within the range from 1.28 to 1.73, revealing substantial deviations from the competitive equilibrium, particularly in Southern and rural markets.[8] I then test the relationship between labour market power and several indicators across provinces. My novel evidence supports all theoretical predictions: provinces with higher average markdowns exhibit lower wages, higher entrepreneurship rates, smaller average firm size and lower average productivity, greater statistical dispersion in wages and lower dispersion in firm size, greater dispersion in MRPL and TFPR, and reduced use of intangibles. A set of robustness checks confirms these findings across different specifications, including controls for the aforementioned geographical sources of variation (i.e. North vs South, urban vs rural).

Finally, I quantify the model to estimate aggregate losses. The parameters representing market power in product and labour markets are externally calibrated, using measures of price markups from the literature and wage markdowns from my own estimations. The remaining parameters are identified matching the model to the data. The shape of the ability distribution is derived using the aggregate profit share. The two parameters characterising the superior technology (i.e. the sunk cost and productivity growth upon

[6]I contributed to the last vintage of this data collection project with the CompNet Scientific team. Specifically, I designed the estimation routine to derive the firm-level indicators of productivity and market power used throughout this paper. I describe the CompNet database in Section 3.1.

[7]For robustness, I replicate the analysis estimating output elasticities with a 'production approach' (Section 3.2; Appendix F). There are two key advantages in using revenue functions in my context: (i) unlike quantities, revenues are directly observed in financial statements, so the elasticities are not subject to the 'price bias' (Bond et al., 2021); (ii) the markdown estimator based on revenue elasticities is robust to imperfect competition in intermediate input markets.

[8]Consistently, Hirsch et al. (2022) find that denser areas have more competitive labour markets.

adoption) are pinned down using aggregate data on software expenditure from EUKLEMS and survey data on software adoption from ISTAT.

My results reveal that imperfect competition in labour markets has a significant impact on aggregate outcomes. For the average labour market, monopsony power generates a 34% loss in output from a competitive benchmark, which I define using localities at the bottom of the markdown distribution (such as Milan, the city with the lowest measure of market power in my dataset). This number reflects the combination of a 21% drop in aggregate productivity and a 17% drop in aggregate employment. Decomposing productivity losses into the three key channels, I find that excess entry accounts for 45% of the total and limited technology diffusion accounts for 35%, with misallocation explaining the remaining share. Moving from an equilibrium with MRPL exceeding wages by 25% to an equilibrium with a 75% excess MRPL (i.e. from low to high monopsony) doubles the number of active firms and reduces the share of workers using the superior technology from 45% to less than 10%. Lastly, I use the calibrated model to infer the contribution of labour market power to geographical disparities in Italy. In light of my estimates, the lower level of competition in Southern labour markets appears to be a powerful explanation for its relative underperformance, accounting for 40% of the productivity gap between the North and South.

Related literature. This paper contributes to the growing research on the implications of market power.[9] Within this literature, the closest strand consists of recent studies examining the aggregate costs of imperfect competition in labour markets, typically focused on misallocation. My contribution lies in a novel framework that analytically demonstrates two additional sources of productivity losses: the excess entry of low-ability entrepreneurs and the limited diffusion of productivity-enhancing technologies.

Berger et al. (2022) present a general equilibrium model with a fixed number of firms and Cournot competition in labour markets, quantifying significant output losses due to misallocation stemming from dispersion in wage markdowns. Consistently with my results, they find that concentration is inefficiently low as labour market power distorts the firm size distribution towards small firms. Bachmann et al. (2022) reach a similar conclusion regarding the effect of monopsony power on concentration and allocative efficiency, developing a model of monopsonistic competition with investments in customer acquisition. Armangué-Jubert et al. (2024) perform a quantitative exercise with a numerical model of oligopsonistic competition, showing that labour market power accounts for a substantial component of differences in GDP per capita across countries at different stages of development. Outside the monopsony literature, Edmond et al. (2023)

[9]Syverson (2019, 2024) provides a comprehensive assessment of this extensive literature.

and Rubens (2023) offer insights on the allocative effects of market power considering, respectively, product and intermediate input markets.

My findings on the excess entry of low-ability entrepreneurs relate to earlier work on distortions in entry decisions. A literature dating back to Mankiw and Whinston (1986) suggests that imperfect competition may lead to an excessive number of active firms, resulting in output losses in the presence of entry fees. This effect also occurs in my model, in which the cost of entry is an opportunity cost (i.e. the foregone wage): by encouraging self-selection into entrepreneurship, labour market power reduces aggregate employment, involving a fall in output traditionally defined as deadweight loss. In this paper, however, the reduction in output induced by excess entry goes beyond the deadweight loss, including efficiency losses due to the low productivity of additional entrants.

This type of mechanism, by which a higher number of active firms results in lower average productivity, is typically associated with other classes of frictions, such as trade barriers (Melitz, 2003), limited substitutability between product varieties (Syverson, 2004), credit constraints (Ghatak et al., 2007), policy distortions (Guner et al., 2008), or regulation costs (Gourio and Roys, 2014; Garicano et al., 2016). This paper considers labour market power as an additional source of inefficiencies. Furthermore, as excess entry in the model reduces the average managerial talent across firms, my findings are also relevant to the research on managerial practices (Bloom and Van Reenen, 2007 and 2010; Guner et al., 2018; Bloom et al., 2019; etc.), suggesting that differences in managerial performance across regions and countries can be affected by differences in labour market competition.

Lastly, by showing that monopsony power acts as a potential barrier to the diffusion of productivity-enhancing technologies, this paper contributes to the literature on factors affecting productivity investments (Bento and Restuccia, 2017) and technological upgrades (Bustos, 2011). Bloom et al. (2012) and Schivardi and Schmitz (2017) highlight that the quality of management is a critical factor for fully exploiting the benefits of IT. In my model, the proliferation of small businesses with poor managerial practices hampers technology diffusion, with the root cause being the weaker selection of active entrepreneurs driven by labour market power.

Outline. The remainder of the paper is organised as follows. Section 2 illustrates the model and presents the theoretical results. Section 3 describes the data, the measurement of labour market power, and the theory's empirical validation. Section 4 shows the calibration exercise and the quantitative results. Section 5 concludes.

2 Theory

2.1 Environment

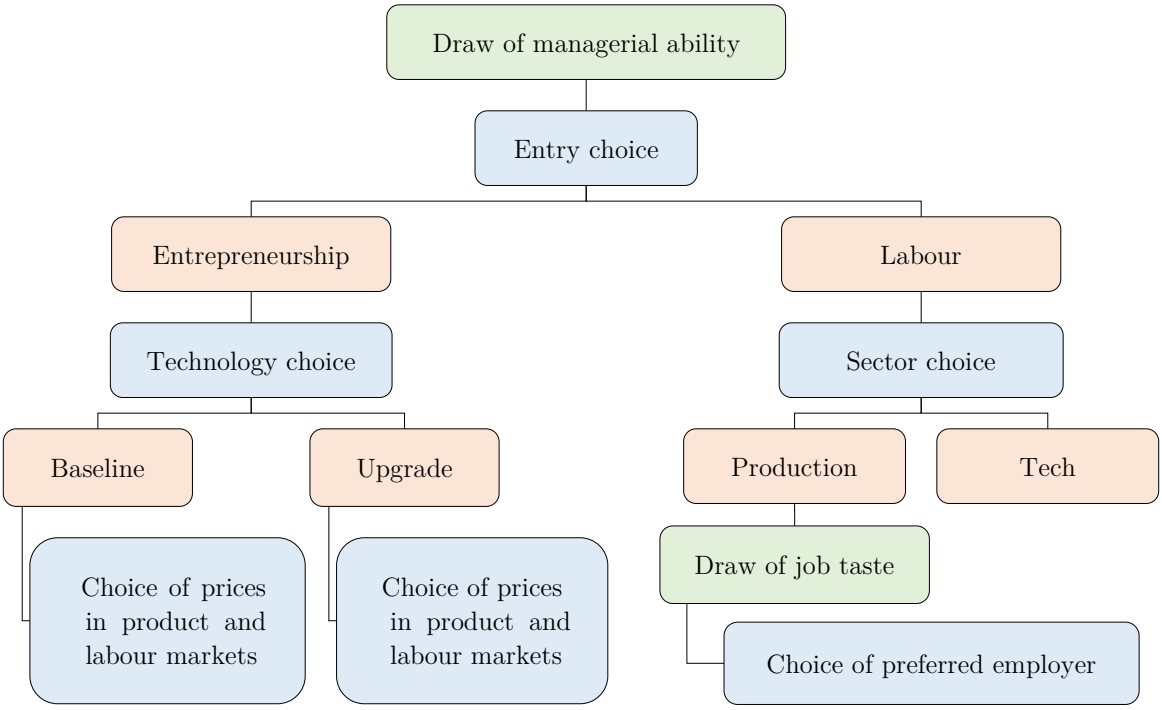
A continuum of agents i with unit mass is endowed with heterogenous levels of managerial ability, μ_i, drawn from an exogenous distribution, $\Phi(\mu_i)$. Depending on their ability, agents self-select into entrepreneurship or labour, maximising expected utility. Agents are risk-neutral and use their income on a consumption bundle, c_i, which is defined as a CES aggregator of product varieties and enters linearly in the utility function. Although leisure is not valued, the aggregate labour supply is elastic as higher wages attract agents from entrepreneurship.

The utility from labour combines the wage, w_k, with a random term, $\varepsilon_{ik} > 0$, which is heterogeneous across workers and reflects the taste of worker i for a job in firm k. This idiosyncratic component captures factors of non-monetary utility (e.g. preferred location, commuting times, working hours, etc.) learnt by workers after they enter the labour market and consider potential employers. The parameter θ represents dispersion in ε_{ik}, which corresponds to job differentiation: with higher θ, workers see employers as poorer substitutes and their choice attributes more importance to non-monetary factors; vice versa, with lower θ, workers care more about wage differentials.

Substitutability and competition are closely related concepts. If substitutability between employers is limited, firms can only increase their size by rising wages, as they need to convince workers who, in terms of taste, prefer other firms. Hence, idiosyncratic job preferences limit the labour supply elasticity, allowing employers to exert monopsony power and extract rents from workers. If θ approaches zero, workers see employers as perfect substitutes and labour markets approach perfect competition, with an equalised market wage.

Firms compete monopolistically in the product market, with the intensity of competition depending on σ, the elasticity of substitution between varieties. Labour is the only factor of production. Output per worker, or productivity, depends on the entrepreneur's managerial ability and the technology used. Firms are endowed with a baseline technology and can increase their productivity by adopting a superior technology. This upgrade, chosen endogenously, involves a fixed cost defined in labour units. For tractability, technology is produced in a competitive sector, in which tech workers are paid the average wage, $\overline{w}$. Entrepreneurship and tech do not provide non-monetary utility.

FIGURE 3. TIMING OF DRAWS AND CHOICES



Note: Green boxes represent random draws; blue boxes represent the agents' choices; orange boxes represent options. The first draw generates firm heterogeneity; the second generates job differentiation and, thus, monopsony power.

The occupational choice splits agents into two sets, entrepreneurship ($\mathcal{E}$) and labour ($\mathcal{L}$), with the latter composed of production and tech workers. Non-monetary utility from production labour, ε_{ik}, plays a fundamental role in job differentiation and competition, whereas its effect on the occupational choice is neutralised by normalising the expected value of non-monetary utility from the preferred employer k^*: $\mathbb{E}[\varepsilon_{ik^*}] = 1$. This is analogous to neutralising love-of-variety in a CES aggregator. It follows that expected utility from production labour equals the average wage, $\mathbb{E}[c_i \varepsilon_{ik^*}] = \overline{w}$, making workers indifferent between production labour and tech in expectation.[10]

$$\mathbb{E}[U_i] = \begin{cases} \overline{w} & i \in \mathcal{L} \\ \pi_i & i \in \mathcal{E} \end{cases} \tag{1}$$

[10]By neutralising the scale effect of non-monetary components on the utility of production workers, I obtain an equivalence between the sum of individual utilities and aggregate output. This assumption aligns a utilitarian welfare analysis to the analysis of production efficiency, coherently with the scope of this paper. Considerations about the welfare value of employer variety can be found in Jha and Rodriguez-Lopez (2023).

The entry condition for firms resembles Lucas (1978): agents self-select into entrepreneurship whenever potential profits, π_i, are higher than the average wage. Overall, this theoretical framework can be understood as the Lucas' model with two key differences: first, the size distribution does not result from the 'span of control' (i.e. decreasing returns to scale) but from product and job differentiation, which make firms price-maker in both markets; second, firms can adopt a productivity-enhancing technology for a fixed cost, which generates economies of scale.

The timing of draws and choices is illustrated in Figure 3.

2.2 Consumer's problem

Consumers spend their income x_i on a consumption bundle c_i, defined as a CES aggregator of varieties c_{ik}. The aggregator has an elasticity of substitution σ and includes a correction term based on the mass of firms, E, to neutralise love of variety.

$$Max_{c_{ik}} \left(\frac{1}{E^{\frac{1}{\sigma}}}\int_{k\in\mathcal{E}} {c_{ik}}^{\frac{\sigma-1}{\sigma}}\,dk\right)^{\frac{\sigma}{\sigma-1}} \; s.t. \int_{k\in\mathcal{E}} p_k c_{ik}\,dk \leq x_i \tag{2}$$

Combining the first-order conditions with $x_i = \int_{k\in\mathcal{E}} p_k c_{ik}\,dk$, I obtain the price index $P = \frac{c_i}{x_i}$, which I normalise to one, and the relationship between c_{ik} and p_k.

$$P = \left(\frac{1}{E}\int_{k\in\mathcal{E}} p_k^{1-\sigma}\,dk\right)^{\frac{1}{\sigma-1}} = 1 \tag{3}$$

$$c_{ik} = \frac{c_i}{E} p_k^{-\sigma} \tag{4}$$

The demand observed by firm k is the sum of the individual consumption of its variety, $y_k = \frac{\int_i c_i di}{E} p_k^{-\sigma}$. I combine the latter with equation (3), to derive aggregate output $Y \equiv \left(\frac{1}{E^{\frac{1}{\sigma}}}\int_{k\in\mathcal{E}} {y_k}^{\frac{\sigma-1}{\sigma}}\,dk\right)^{\frac{\sigma}{\sigma-1}} = \int_i c_i\,di$, which leads to the following equation for the product demand, as well as its elasticity.

$$y_k = \frac{Y}{E} p_k^{-\sigma} \tag{5}$$

$$\left|\frac{\partial y_k}{\partial p_k}\frac{p_k}{y_k}\right| = \sigma \tag{6}$$

2.3 Worker's problem

Production workers choose employers that maximise their utility, considering both the wage offered and an idiosyncratic taste component, $\tilde{\varepsilon}_{ik} > 0$. The latter is independently drawn from a Fréchet distribution with shape parameter $\frac{1}{\theta}$, leading to tractable choice probabilities.

$$Pr(argmax_k\{w_k\tilde{\varepsilon}_{ik}\} = k^*) = \frac{{w_{k^*}}^{\frac{1}{\theta}}}{\int_{k\in\mathcal{E}} w_k^{\frac{1}{\theta}}\, dk} \tag{7}$$

Equating probabilities to employment shares, I derive the labour supply and its elasticity, with L^P defined as the mass of production workers.

$$l_k = L^P \frac{{w_k}^{\frac{1}{\theta}}}{\int_{k\in\mathcal{E}} w_k^{\frac{1}{\theta}}\, dk} \tag{8}$$

$$\frac{\partial l_k}{\partial w_k}\frac{w_k}{l_k} = \frac{1}{\theta} \tag{9}$$

Higher θ implies more job differentiation and lower substitutability between employers, which generates a lower labour supply elasticity and, ultimately, weaker competition in the labour markets.

The taste shock enters in the utility function of production workers rescaled by its expected value: $U_i|_{k=k^*} = w_{k^*}\varepsilon_{ik^*} = w_{k^*}\frac{\tilde{\varepsilon}_{ik^*}}{\mathbb{E}[\tilde{\varepsilon}_{ik^*}]}$. This normalisation means that workers value non-monetary utility from a random firm k in terms of the expected non-monetary utility from their preferred option k^*. The correction term $\mathbb{E}[\tilde{\varepsilon}_{ik^*}]$ does not affect the relative utility between job options, hence it is irrelevant for the worker's problem. However, by affecting the scale of the utility, it matters for the occupational choice. The scope of neutralising non-monetary utility, by setting $\mathbb{E}[\varepsilon_{ik^*}] = 1$, is to make the utility from production labour comparable to the utility from entrepreneurship and tech, as the latter ones do not have any non-monetary component. A related advantage of this normalisation is that, by fixing the expected value of non-monetary utility to a constant, it makes θ a dimensionless measure of dispersion, limiting its effect to job substitutability and monopsony power.

2.4 Firm's problem

Production labour is the only input. Firm productivity, a_k, can be decomposed into the managerial ability of the entrepreneur, μ_k, and the technology employed by the firm, t_k.

$$y_k = a_k l_k = \mu_k t_k l_k \tag{10}$$

Two technologies are available in the economy: a baseline technology, costless and normalised to one, and a superior technology that allows productivity to grow by a factor $t_k = 1 + \gamma$. Upgrading to the productivity-enhancing technology involves a fixed cost, $f_k = \overline{w}\tau$; the constant τ represents the mass of tech workers required to provide a firm with the superior technology.

$$t_k \in \{1,1+\gamma\}, \gamma > 0 \tag{11}$$

$$f_k \in \{0, \overline{w}\tau\}, \tau > 0 \tag{12}$$

First, I solve the firm's problem for a given level of productivity; I then show which firms upgrade their technology. Firms maximise profits internalising the finite elasticities of the product demand and labour supply.

$$Max_{l_k}\ p_k a_k l_k - w_k l_k - f_k \tag{13}$$

Taking the first-order condition and substituting the two elasticities, I obtain the marginal revenue product of labour as the wage times a markdown. The positive slope of the inverse labour supply creates a gap between marginal and average labour costs, inducing firms to set wages below the MRPL. Intuitively, the magnitude of this wedge depends on θ, which denotes limited substitutability between employers and, consequently, monopsony power.

$$MRPL_k = p_k a_k \frac{\sigma - 1}{\sigma} = w_k(1+\theta) \tag{14}$$

Profitability depends on competition in both markets, as shown by the ratio of revenues to variable costs, equal to the product of price markups and wage markdowns (respectively, $\frac{\sigma}{\sigma-1}$ and $1+\theta$).

Substituting the product demand and labour supply, after some algebra I can express firms' employment and output in terms of productivity. Then, I compute aggregate output and substitute it in the equations of wages and profits.

$$l_k = L^P \frac{a_k^{\frac{\sigma-1}{\sigma\theta+1}}}{\int_{k\in\mathcal{E}} a_k^{\frac{\sigma-1}{\sigma\theta+1}}\, dk} \tag{15}$$

$$y_k = L^P \frac{a_k^{\frac{\sigma(1+\theta)}{\sigma\theta+1}}}{\int_{k\in\mathcal{E}} a_k^{\frac{\sigma-1}{\sigma\theta+1}} dk} \tag{16}$$

$$Y = L^P \frac{\left(\frac{1}{E^{\frac{1}{\sigma}}} \int_{k\in\mathcal{E}} a_k^{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}} dk\right)^{\frac{\sigma}{\sigma-1}}}{\int_{k\in\mathcal{E}} a_k^{\frac{\sigma-1}{\sigma\theta+1}} dk} \tag{17}$$

$$w_k = \frac{\sigma-1}{\sigma(1+\theta)} \left(\frac{1}{E} \int_{k\in\mathcal{E}} a_k^{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}} dk\right)^{\frac{1}{\sigma-1}} a_k^{\frac{(\sigma-1)\theta}{\sigma\theta+1}} \tag{18}$$

$$\pi_k = \frac{\sigma\theta+1}{\sigma(1+\theta)} L^P \frac{\left(\frac{1}{E} \int_{k\in\mathcal{E}} a_k^{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}} dk\right)^{\frac{1}{\sigma-1}}}{\int_{k\in\mathcal{E}} a_k^{\frac{\sigma-1}{\sigma\theta+1}} dk} a_k^{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}} - f_k \tag{19}$$

$$\overline{w} = \frac{\sigma-1}{\sigma(1+\theta)} \frac{\left(\frac{1}{E^{\frac{1}{\sigma}}} \int_{k\in\mathcal{E}} a_k^{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}} dk\right)^{\frac{\sigma}{\sigma-1}}}{\int_{k\in\mathcal{E}} a_k^{\frac{\sigma-1}{\sigma\theta+1}} dk} = \frac{\sigma-1}{\sigma(1+\theta)} \frac{Y}{L^P} \tag{20}$$

2.5 Technology adoption

Technological upgrades are convenient to firms if the increase in variable profits (i.e. revenues minus variable costs) is larger than the sunk cost. Substituting the profit function, the following condition characterises firms adopting the superior technology.

$$\frac{\sigma\theta+1}{\sigma(1+\theta)} L^P \frac{\left(\frac{1}{E} \int_{k\in\mathcal{E}} a_k^{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}} dk\right)^{\frac{1}{\sigma-1}}}{\int_{k\in\mathcal{E}} a_k^{\frac{\sigma-1}{\sigma\theta+1}} dk} \mu_k^{\frac{(\sigma-1)(\theta+1)}{\sigma\theta+1}} \left[(1+\gamma)^{\frac{(\sigma-1)(\theta+1)}{\sigma\theta+1}} - 1\right] > \overline{w}\tau \tag{21}$$

Whilst the cost of adoption (on the right-hand side) is the same for all firms, the benefit (on the left-hand side) increases in μ_k. This implies that only firms with managerial ability above a threshold, $\mu_k > \mu_T$, adopt the productivity-enhancing technology. It follows that ranking firms by productivity or managerial ability gives the same order: employment, output, wage, and profits increase not only in productivity, but also in managerial ability.

The technology adoption cutoff, μ_T, is defined by the following condition:

$$\pi(\mu_T|_{t_T=1+\gamma;\, f_T=\overline{w}\tau}) = \pi(\mu_T|_{t_T=1;\, f_T=0}) \tag{22}$$

Substituting the profit function and average wage, and rearranging, I derive the following equation for the adoption cutoff.

$$\mu_T = \left(\frac{\sigma-1}{\sigma\theta+1} \frac{\tau}{(1+\gamma)^{\frac{(\sigma-1)(\theta+1)}{\sigma\theta+1}} - 1} \frac{\int_{k\in\mathcal{E}} a_k^{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}} dk}{L^P} \right)^{\frac{\sigma\theta+1}{(\sigma-1)(\theta+1)}} \tag{23}$$

The adoption cutoff goes to zero (i.e. all firms adopt) when no tech worker is required, $\tau \to 0$, or the implied productivity growth is unbounded, $\gamma \to \infty$. Vice versa, the cutoff diverges (i.e. no firm adopts) when the cost is unbounded, $\tau \to \infty$, or it does not provide gains, $\gamma \to 0$.

The definition of the adoption cutoff allows to derive the mass of tech workers, L^T.

$$L^T = \int_{\mu_T}^{\infty} \tau\, \phi(\mu) d\mu \tag{24}$$

2.6 Firm entry

Agents self-select into entrepreneurship if this choice provides higher expected utility. Because non-monetary utility is neutralised, indifference between entrepreneurship and labour, at the margin, corresponds to the equivalence of potential profits and the average wage, resembling Lucas (1978).

$$\mathbb{E}[U_i|_{i\in\mathcal{E}}] \gtrless \mathbb{E}[U_i|_{i\in\mathcal{L}}] \Leftrightarrow \pi(\mu_i) \gtrless \overline{w} \tag{25}$$

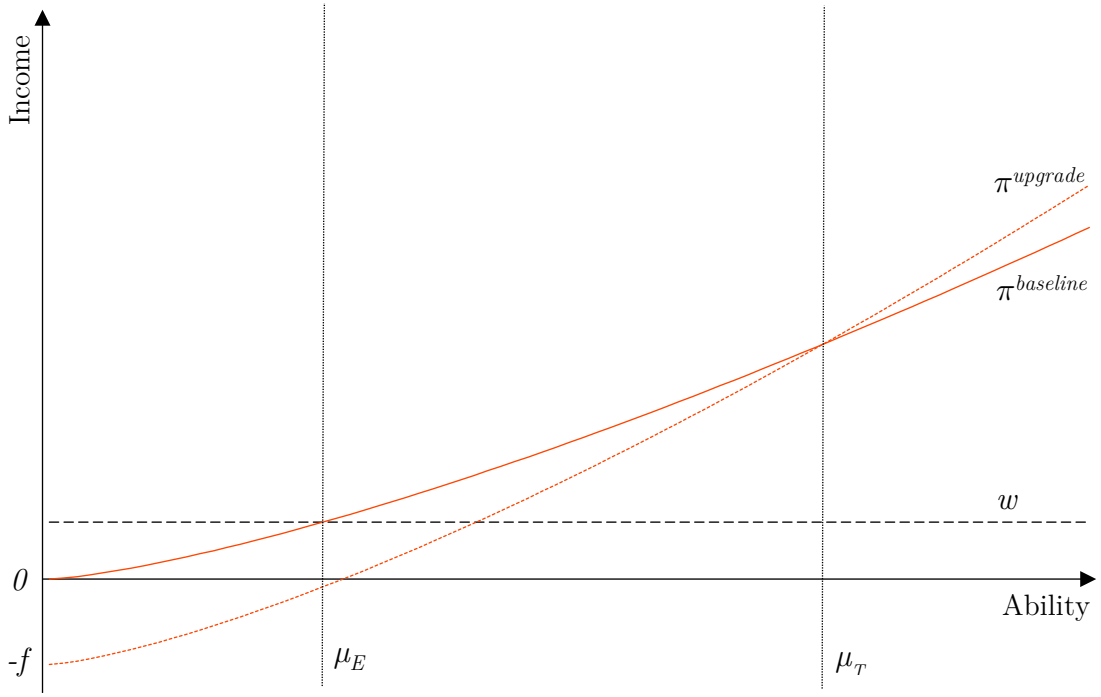
As profits increase in μ_i, only those with managerial ability above a threshold, $\mu_i > \mu_E$, become entrepreneurs.

The entry cutoff is identified by the following condition:

$$\pi(\mu_E) = \overline{w} \tag{26}$$

By excluding the trivial case where all active firms use the same technology, I assume that the marginal entrant employs the baseline technology, i.e. $t_E = 1$. Substituting profits and the average wage, and rearranging, I obtain the following equation for the entry cutoff.

FIGURE 4. EXPECTED INCOME BY LEVELS OF MANAGERIAL ABILITY



Note: The wide-dashed black horizontal line is the expected income of workers (i.e. average wage). Orange curves are the expected income of entrepreneurs: the continuous curve is profits under the baseline technology; the narrow-dashed curve is profits in case of technological upgrade, whose fixed costs corresponds to the negative intercept. Dotted vertical lines mark the entry and adoption cutoffs, respectively μ_E and μ_T.

$$\mu_E = \left(\frac{\sigma - 1}{\sigma\theta + 1} \frac{\int_{k \in \mathcal{E}} a_k^{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}} dk}{L^P} \right)^{\frac{\sigma\theta+1}{(\sigma-1)(\theta+1)}} \tag{27}$$

The definition of the entry cutoff allows to characterise the set of active entrepreneurs (and, residually, the set of workers), as well as its mass and the mass of production labour.

$$\mathcal{E} = \{i|_{\mu_i > \mu_E}\} \tag{28}$$

$$E = \int_{\mu_E}^{\infty} \phi(\mu) d\mu \tag{29}$$

$$L^P = 1 - E - L^T \tag{30}$$

The payoff of occupational choice is represented in Figure 4. The expected wage is the same for any level of managerial ability. Profits with the baseline technology start from zero and increase in the ability, as firm productivity increases in the latter. Technology adoption has two effects on the profit function: (i) the curve shifts down, as adoption involves a sunk cost; (ii) the curve becomes steeper, because adoption makes firms more productive, larger and, ultimately, more profitable. The entry cut-off is identified by the intersection of the average wage and profits with the baseline technology, and marks the threshold between workers and entrepreneurs. The adoption cut-off is identified by the intersection of the two profit functions, and splits entrepreneurs by the type of technology used.

2.7 Scale bias

Adopting the productivity-enhancing technology allows to cut marginal costs, $MC_k = \frac{(1+\theta)w_k}{\mu_k t_k}$, by a factor $(1+\gamma)$ for any level of employment. Following Reichardt (2024), I characterise this superior technology as *scale-biased*, because it affects the cost structure in a way that is more beneficial to large firms: the sunk cost, $\tau\overline{w}$, is the same for all firms, whereas the benefits of lowering marginal costs are scaled up in firm size.[11]

The final effect of technological upgrades on marginal costs, after adjustments in employment, may not be obvious in the presence of monopsony power, because the increase in productivity induces firms to employ more production workers, which require higher wages. Using the wage function derived in equation (18), I can show that technology adoption still generates a reduction in marginal costs in equilibrium by a factor that depends on γ.[12]

$$\frac{(MC_k|_{t_k=1+\gamma})}{(MC_k|_{t_k=1})} = \left(\frac{1}{1+\gamma}\right)^{\frac{1+\theta}{\sigma\theta+1}} \tag{31}$$

For any level of γ, the higher τ the more the technology is scale-biased: the benefit from reducing marginal costs must be scaled up to a larger size to offset the sunk cost. Combining equations (23) and (27), the ratio of the adoption to entry cutoff, defined as $b(\tau,\gamma)$, can be seen as a market-specific measure of scale bias. It represents the combination through which the two parameters characterising the superior technology, γ and τ, affect technology diffusion across firms.

[11]This argument dates back to Sutton (1991) and the related literature on endogenous sunk costs.

[12]Variable costs, $VC_k = \frac{MC_k}{1+\theta}$, fall by the same factor.

$$b(\tau,\gamma)=\frac{\mu_T}{\mu_E}=\left(\frac{\tau}{(1+\gamma)^{\frac{(\sigma-1)(\theta+1)}{\sigma\theta+1}}-1}\right)^{\frac{\sigma\theta+1}{(\sigma-1)(\theta+1)}} \tag{32}$$

The numerator of the ratio within brackets, τ, represents the mass of tech workers required by each firm willing to adopt the productivity-enhancing technology, whereas the denominator is the growth rate in 'gross profits' upon adoption. The latter depends not just on γ, the percentage increase in firm productivity, but also on the slopes of the product demand and labour supply. The intuition is that the benefits from technological upgrades are lower if scaling up requires substantial decreases in prices and increases in wages. Thus, the scale required to benefit from the superior technology becomes larger.

For $b=0$, a technology has no scale bias as no fixed cost is required. For $0<b\leq 1$, all agents with managerial ability above the entry cutoff opt for the superior technology. For $b>1$, the case considered throughout this paper, the adoption choice is non-trivial as the degree of scale bias is sufficiently high that some active entrepreneurs do not adopt. The higher b, the larger the size required to benefit from the superior technology.

2.8 General equilibrium

DEFINITION 1 (Competitive equilibrium). *Given a cumulative distribution of entrepreneurial ability $\Phi(\mu)$, the utility function's parameters σ and θ, the available technology characterised by $t_k\in\{1,1+\gamma\}$ and $f_k\in\{0,\tau\overline{w}\}$, and a production function linear in labour, $y_k=\mu_k t_k l_k$, a* competitive equilibrium *consists of (i) vectors of prices p_k, wages w_k, and fixed costs f_k, (ii) vectors of production workers l_k and technologies t_k, and (iii) an entry cutoff μ_E and a technology adoption cutoff μ_T, such that: in expectation, no agent would be better off changing occupation, no worker would be better off changing employer, no entrepreneur would be better off changing technology or employment, and both the product and labour markets clear. The system of equations (15)-(20), (23), (24), (27), (28-30) satisfies this equilibrium.*

Henceforth, I assume that the probability density function of the ability distribution is Pareto, $\phi(\mu)=\frac{\alpha}{\mu^{\alpha+1}}$, with scale normalised to one and α corresponding to the shape parameter. To obtain a defined mean in the revenue distribution and solve the model analytically, I assume $\alpha>\max\{\frac{(\sigma-1)(1+\theta)}{\sigma\theta+1},1\}$. Under this assumption, the equilibrium in Definition 1 admits a closed-form solution. Derivations are reported in Appendix A.

To show how labour market power affects the size distribution and aggregate productivity, in the following two subsections I consider the general equilibrium effects of

variations in θ, the exogenous parameter that determines job differentiation and, thus, the level of labour market power.

2.9 Theoretical results: size distribution

The distribution of workers across firms is shaped by two factors: (i) the allocation of employment shares across active firms, which depends on the productivity distribution as well as the slopes of the product demand and labour supply, and (ii) the average firm size, determined by the ratio of the mass of production workers to the mass of entrepreneurs. Labour market power distorts both factors, affecting the variance and mean of the size distribution, as shown in Panel A of Figure 5.

PROPOSITION 1 (Reallocation): *Labour market power reallocates inputs from larger to smaller firms, reducing size dispersion.*

$$l_k > l_j \;\Rightarrow \frac{\partial (l_k/l_j)}{\partial \theta} < 0 \tag{33}$$

Proof. See Appendix A.

Although labour market power reduces wages on aggregate, the effect is weaker for large firms, because lower substitutability between employers generates a steeper inverse labour supply at the firm-level. By making the wages of large firms relatively higher compared to their smaller competitors, monopsony reallocates workers towards smaller firms, ultimately compressing the size distribution. Even in settings in which large firms have higher wage markdowns, a larger size is still associated with higher wages as long as employers are imperfect substitutes. Moreover, a positive relationship between markdowns and firm size would be an additional source of reallocation towards small firms. Appendix B provides a framework with variable markdowns, in which increasing the market power of large firms would still compress the size distribution.

PROPOSITION 2 (Excess entry): *Labour market power increases the mass of entrepreneurs and decreases the mass of production workers, reducing average firm size.*

$$\frac{\partial E}{\partial \theta} > 0 \tag{34}$$

$$\frac{\partial L^P}{\partial \theta} < 0 \tag{35}$$

Proof. See Appendix A.

Monopsony power affects the occupational choice by lowering wages below the MRPL, transferring rents from workers to entrepreneurs. This effect increases the appeal of entrepreneurship, generating a fall in the entry cutoff. By increasing the mass of firms and decreasing the mass of production workers, labour market power directly reduces average firm size, $\overline{l} = \frac{L^P}{E}$.

In the presence of costly technology adoption, the effect on the mass of production workers is not trivial, as more entrepreneurs could be compensated by fewer tech workers. However, I show that the effect remains unambiguously negative, and its magnitude is independent of technology specifications: in this model, a scale-biased technology induces agents to move from entrepreneurship to tech, increasing business concentration with no effects on production labour.

Imperfect competition in labour markets affects the occupational choice by acting on both sides of the entry condition represented in equation (25): on the one hand, it reduces the average wage; on the other hand, it increases potential profits. In different settings, in which wage markdowns depend on firm size and the marginal entrant has little market power, higher levels of monopsony would still generate excess entry because of the wage channel induced by a weaker labour demand from large firms. Appendix B provides an example.

Concentration.[13] The effects of labour market power described in Propositions 1 and 2 (i.e. reallocation towards small firms and excess entry) reduce employment concentration. This result may seem to contradict the idea of a positive relationship between market power and concentration, which is based on strategic interactions and proceeds in the opposite causal direction, with an opposite sign. As strategic interactions are absent in this model, monopsony power purely depends on job differentiation and is not influenced by firm size.

In an oligopsonistic framework, the relationship between concentration and market power is not obvious and depends on the source of heterogeneity. For instance, factors that raise the productivity gap of larger firms over laggards (e.g. scale-biased technical change) increase employment concentration and, with strategic interactions, monopsony power. On the contrary, factors that increase monopsony power by lowering the substitutability between employers (e.g. strong distaste for long commuting) generate excess entry and reallocation towards small firms, thus lowering employment

[13]This discussion resembles an old debate in the IO literature on the limits of the Structure-Conduct-Performance paradigm, which assumes a specific causal nexus from market structure to firm conduct (see Clarke et al., 1984, and, more recently, Syverson, 2019).

concentration. The former type of shock produces a positive relationship between concentration and market power, whereas the latter generates a negative relationship.

Ultimately, the sign of the unconditional correlation between monopsony power and employment concentration across labour markets is an empirical question and depends on which causal nexus dominates in the specific context. As this paper does not focus on the sources of monopsony power but on its implications, the shock under consideration (i.e. variations in job differentiation) leads to a negative relationship, with or without strategic interactions. This theoretical result contributes to understanding why, empirically, there is mixed evidence on the relationship between monopsony power and employment concentration across labour markets. An analysis of the theoretically and empirically ambiguous relationship between market power and concentration is available, for product markets, in Syverson (2019).[14]

2.10 Theoretical results: productivity

The dual effect of monopsony power on the size distribution, represented by Propositions 1 and 2, has direct implications for aggregate productivity. First, reallocation towards smaller firms implies reallocation towards less efficient firms, as higher firm productivity determines larger size. Second, expanding the pool of entrepreneurship to agents with poorer ability worsens the quality of managerial practices among active firms. Finally, size distortions generate aggregate losses through an additional channel: the reduced diffusion of productivity-enhancing technologies.

COROLLARY 1 (Inefficient allocation): *Labour market power shifts workers towards firms in which they are less productive, reducing allocative efficiency.*

$$MRPL_k > MRPL_j \Rightarrow \frac{\partial(l_k/l_j)}{\partial\theta} < 0 \tag{36}$$

Proof. See Appendix A.

As firm size positively depends on managerial ability, the distortion in the distribution of workers towards small firms discussed in Proposition 1 implies a reallocation towards less efficient firms, in terms of productivity as well as MRPL. I define this as misallocation because the resulting allocation does not maximise output. Reallocating

[14]In my data, provinces with higher average markdowns have significantly lower employment concentration. Yeh at al. (2022) test the relationship across local labour markets in the US, without finding clear patterns. Azar et al. (2020, 2022) find instead a positive relationship, with two key methodological differences: (i) different occupations are classified as different labour markets; (ii) concentration is defined in terms of flows rather than stocks. Both choices are better suited to capture the role of strategic interactions, in line with the purpose of those studies.

workers towards larger firms, which pay higher wages and have higher MRPL, would increase aggregate production.[15] This observation could also be seen in the canonical model of Hsieh and Klenow (2009) by considering distorted input prices as heterogeneous wedges multiplying an equalised market wage. Labour market power generates dispersion in wages and, as a result, in MRPL and TFPR (i.e. firm productivity times price).[16]

In Panel B of Figure 5, I consider two firms, with one twice as productive as the other. Their size ratio falls in the level of labour market power, consistently with Proposition 1. At the same time, their MRPL ratio increasingly departs from one, denoting misallocation.

COROLLARY 2 (Inefficient selection): *Labour market power induces a shift in the cumulative distribution of managerial ability among active entrepreneurs towards the left.*

$$\theta > \theta' \Rightarrow \Phi(\mu|_{\mu>\mu_E(\theta)}) > \Phi(\mu|_{\mu>\mu_E(\theta')}) \tag{37}$$

Proof. See Appendix A.

A direct consequence of the decrease in the entry cutoff discussed in Proposition 2 is that the ability distribution of active entrepreneurs shifts left, worsening the quality of managerial practices in the economy. Intuitively, weaker competition reduces the effectiveness of market forces in selecting only the most capable firms, allowing agents with lower managerial ability to enter and benefit from entrepreneurship.

Equation (37), represented in Panel C of Figure 5, corresponds to the first-order stochastic dominance of ability distributions generated by lower levels of labour market power. Another way to see the effect of monopsony on the ability distribution is by considering moments and quantiles. Under the Pareto assumption, both the mean and median as well as any other percentile grow linearly in the entry cutoff.

$$\begin{cases} \bar{\mu}_{mean} = \left(\dfrac{\alpha}{\alpha-1}\right)\mu_E \\ \bar{\mu}_{median} = (\sqrt[\alpha]{2})\mu_E \end{cases} \tag{38}$$

Whilst Corollary 1 considers the effect of distorting the allocation of workers within the pool of active firms, Corollary 2 adds the additional productivity loss induced by enlarging

[15] A critical difference between homogeneous markups and homogeneous markdowns is that only the latter generate dispersion in marginal revenue products across firms, inducing losses in aggregate output. In other settings (e.g. Berger et al., 2022), this effect is not considered misallocation because gains in non-monetary utility compensate output losses. Ultimately, whether this loss is classified as misallocation depends on the model's underlying assumptions and scope. See also footnote 10.

[16] The effect of wage dispersion on misallocation can be seen from $TFPR_k = \frac{\sigma}{\sigma-1} MRPL_k = \frac{\sigma(1+\theta)}{\sigma-1} w_k$.

the pool of entrepreneurship to agents with lower managerial ability. The former channel induces misallocation, whereas the latter affects aggregate productivity by decreasing its average. Both channels also hold in a setting with strategic interactions and variable markdowns, as I show in Appendix B.

PROPOSITION 3 (Inefficient technology diffusion): *Labour market power reduces the share of production workers using the superior technology.*

$$\frac{\partial \left(\frac{\int_{\mu_T}^{\infty} l(\mu)\phi(\mu)d\mu}{L^P} \right)}{\partial \theta} < 0 \tag{39}$$

Proof. See Appendix A.

Size distortions induced by monopsony power become increasingly relevant in the presence of scale-biased technology. As discussed in Section 2.7, technologies that allow lower marginal costs by paying a fixed cost are more beneficial to large firms. By distorting the employment distribution, reallocating resources towards smaller firms and reducing firm size, labour market power also reduces the diffusion of such scale-biased technologies across workers, as I represent in Panel D of Figure 5. The presence of a scale-bias technology acts as an amplification factor of the losses associated with Corollaries 1 and 2.

Taken together, Propositions 2 and 3 imply that monopsony power also reduces the *mass* of workers using the superior technology. This negative effect is larger than the effect on the share, as the mass of production workers falls as well.

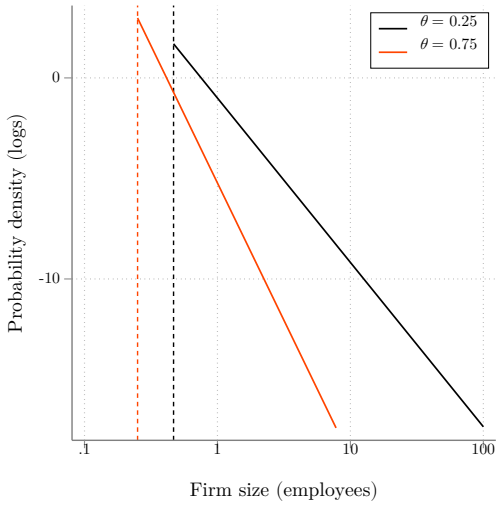
Proposition 3 can also be framed in terms of the share of firms, as labour market power hinders the diffusion of scale-biased technologies not just across workers, but also across companies.[17]

$$\frac{\partial \left(\frac{1-\Phi(\mu_T)}{1-\Phi(\mu_E)} \right)}{\partial \theta} < 0 \tag{40}$$

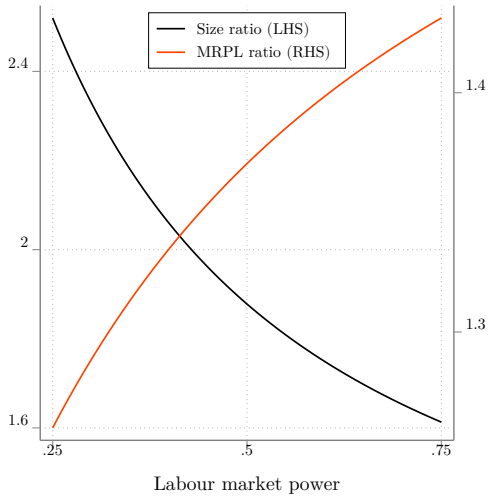
[17]The effect of labour market power on the *mass* of adopting firms depends on the level of scale bias. With the parameters estimated in Section 4, the mass of adopting firms unambiguously falls in the level of labour market power. In Appendix C, I discuss the various channels in the impact of labour market competition on firm-level adoption.

FIGURE 5. THEORETICAL RESULTS

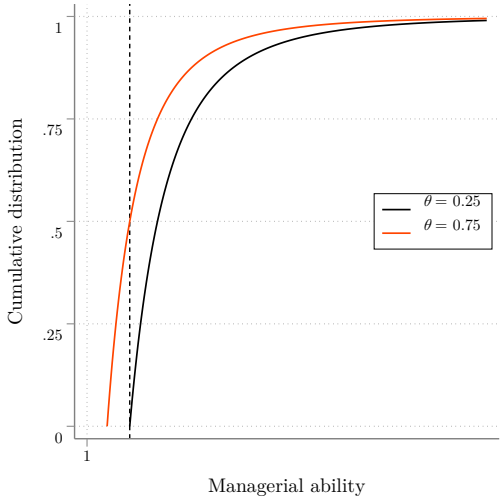
Panel A: Size distortions



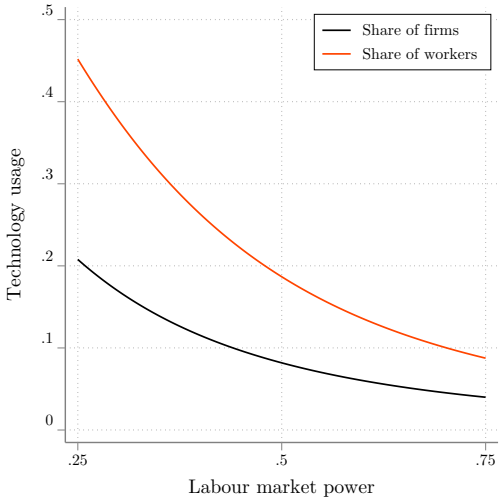
Panel B: Inefficient allocation



Panel C: Inefficient selection



Panel D: Inefficient technology diffusion



Note: *Panel A* represents the dual effect of labour market power on the size distribution: with higher θ, the density gets steeper (i.e. Proposition 1: reallocation from large to small firms) and shifts towards left (i.e. Proposition 2: excess entry and lower average firm size). *Panel B* considers two firms, with one twice as productive as the other, showing that the reduction in their size ratio induced by labour market power is accompanied by an increasing deviation of their MRPL ratio from one (i.e. Corollary 1: inefficient allocation). *Panel C* illustrates how higher θ generates a lower entry cut-off and a shift towards left in the ability distribution of active firms (i.e. Corollary 2: inefficient selection of entrepreneurs). *Panel D* shows the fall in the share of workers using the productivity-enhancing technology induced by labour market power (i.e. Proposition 3: inefficient technology diffusion).

3 Validation

This Section considers the empirical validity of the model's testable predictions. First, I present the data sources (3.1) and measurement strategies (3.2), focusing on the estimation of labour market power. I then discuss the theoretical predictions (3.3) and test them using microdata on Italian firms (3.4).

3.1 Data sources

CompNet data collection. My empirical analysis is based on the ninth vintage of the CompNet micro-aggregated database. I participated in this data collection with the CompNet Scientific team and a network of European institutions. The collection procedure works as follows: data providers implement harmonised estimation routines developed by CompNet on their confidential data, to share indicators of firm behaviour without giving direct access to protected information.[18] The estimation package computes mean, percentiles, and dispersion measures of several variables on production costs and firm performance, including productivity and market power indicators, based on parametric estimates of production and revenue functions.[19] This large information set is then aggregated at different levels (country, industry, province, size class, etc.) and made available to the research community.[20] More information on the CompNet database is available in Melitz et al. (2018) and Mertens et al. (2022).

TABLE 1. SUMMARY STATISTICS

Variable	Mean	SD	P10	P50	P90	Obs.
Revenues	3719	46640	94	583	4913	461384
Intermediates	2651	30615	52	355	3420	486932
Labour costs	632	11681	11	113	862	489144
Capital	853	24831	2	51	940	483338
Employees	16	268	1	5	24	479128

Note: Monetary values in thousands of current Euros. Microdata on Italian firms in 2019 from CompNet.

[18]The dataset on Italian firms used in this paper was produced in 2023 by officers of the European Central Bank. The original data sources are CERVED and Bank of Italy.

[19]I directly designed the parametric estimation of wage markdowns used in this paper. The procedure is explained in Section 3.3.

[20]Researchers can request data access at comp-net.org.

Table 1 presents summary statistics from the database. Empirical tests are performed using the full sample (2007-2019). The quantification exercise in Section 4 calibrates the model to data from 2019.

Other sources. The CompNet database is complemented with other data sources. In particular, I use survey data from ISTAT on IT adoption across Italian firms and ISTAT data on occupations to compute the entrepreneurship rate across Italian provinces. From EUKLEMS-IntanProd (O'Mahoney and Timmer, 2009; Bontadini et al., 2023), I obtain data on national accounts and intangibles used in the calibration.

3.2 Measurement

Most variables employed in my tests are directly taken from financial statements. On the contrary, the index of labour market power requires a parametric estimation which deserves further discussion, presented in this subsection.

Markdowns: assumptions. Whilst wages equal the MRPL in perfect competition, under monopsony they are set below the competitive level. The ratio of the MRPL to the wage, defined as markdown, is a commonly used indicator of imperfect competition in labour markets. This definition mirrors the price markup: product market power increases prices above the marginal cost (mark-*up*) and labour market power reduces wages below the MRPL (mark-*down*). Although these wedges are widely interpreted as measures of competition (Syverson, 2024), they could also reflect other sources of distortion (Hashemi et al., 2022). In this subsection, I derive markdowns in a framework in which labour is a flexible, homogenous input, only used for production. In Appendix D, I show that my results are robust to labour adjustment costs, worker heterogeneity, and labour used as overhead.

Wage markdowns are a component of the labour wedge, defined as the ratio of the marginal product of labour (MPL) to the marginal rate of substitution between consumption and leisure (MRS). In addition to labour market power and the confounding factors considered in Appendix D, several other frictions, extensively discussed in the business cycle literature (Shimer, 2009), affect the labour wedge without creating bias in my markdown estimator. A subset of them, which is quantitatively more important according to Karabarbounis (2014), does not affect the MPL to wage ratio, but the ratio between the wage and the MRS (e.g. income taxes). Among those affecting the MPL to wage ratio, price markups, connected to the labour wedge by Galì et al. (2007), do not enter in the MRPL to wage ratio, as I show below.

Markdowns: derivation. Firm k in period t chooses labour to maximise the difference between revenues, r_{kt}, and variable labour costs, $w_{kt}l_{kt}$.

$$Max_{l_{kt}} \; r_{kt} - w_{kt} l_{kt} \tag{41}$$

From the first-order condition, I can express the markdown, χ_{kt}, as a function of the labour supply elasticity observed by the firm.

$$\chi_{kt} \equiv \frac{MRPL_{kt}}{w_{kt}} = \left(1 + \frac{\partial w_{kt}}{\partial l_{kt}} \frac{l_{kt}}{w_{kt}}\right) \tag{42}$$

Markdowns depart from one whenever the labour supply is not perfectly elastic. The intuition is that firms profit from an additional labour unit as long as its marginal revenue exceeds its marginal cost. Absent monopsony, the marginal cost of labour equals the wage. Hence, firms will employ more workers until the MRPL aligns with the latter. However, if the (inverse) labour supply is upward-sloping, the marginal cost of labour lies above the wage. Thus, it is optimal for firms to reduce employment below the competitive level and keep the MRPL above the wage. The steeper the inverse labour supply, the stronger the degree of monopsony power, and the higher the markdown.

Decomposing the MRPL into the revenue elasticity of labour, $\varepsilon^{r,l}_{kt} \equiv \frac{\partial r_{kt}}{\partial l_{kt}} \frac{l_{kt}}{r_{kt}}$, and revenues per worker, I can express markdowns as the ratio of the revenue elasticity to the labour share, $LS_{kt} \equiv \frac{w_{kt} l_{kt}}{r_{kt}}$. The latter is observed in income statements, whereas the elasticity requires a parametric estimation.

$$\chi_{kt} = \frac{\varepsilon^{r,l}_{kt}}{LS_{kt}} \tag{43}$$

This simple derivation does not require assumptions on competition in product markets. The reason is that the demand elasticity would affect both the numerator and denominator in the same way, leaving the ratio unaffected. Intuitively, a steeper product demand reduces the revenue elasticity because prices fall in input usage and, simultaneously, reduces the labour share because of higher price markups.

A related approach, frequently used in the literature (Dobbelaere and Mairesse, 2013; Yeh et al., 2022), replaces the revenue elasticity of labour with the ratio of its output elasticity to the markup, in a cost-minimisation framework ('production approach'). Markups are then implicitly measured using the 'ratio estimator' (i.e. output elasticity of a flexible input divided by its expenditure share). Under the assumption of perfectly competitive input markets, the markdown equals the ratio estimator of labour divided by the ratio estimator of materials (see Yeh et al., 2022).

$$\chi_{kt}^{prod} = \left(\frac{\varepsilon_{kt}^{y,l}}{LS_{kt}} \Big/ \frac{\varepsilon_{kt}^{y,m}}{MS_{kt}} \right) \tag{44}$$

The limit of the production approach to markdown estimation is that monopsony power is underestimated if firms have some degree of pricing power over their suppliers. For robustness, I replicate the empirical analysis following this alternative approach in Appendix F. Results remain consistent with the theoretical mechanism.

Estimation of revenue functions. The revenue elasticity of labour is the only unobserved component in equation (43). By estimating firm-level revenue functions, I obtain revenue elasticities and, thus, wage markdowns. For this estimation, I generalise the firm's problem of Section 2, in which labour is the only input, by considering a production function in capital, labour, and materials. The key assumption in this approach is that revenues can be approximated by a translog specification, which allows nonlinear relationships between revenues and production factors up to second-order and heterogeneity in revenue elasticities across firms.[21]

By expressing revenue elasticities in terms of output elasticities and demand elasticities, it is possible to see that revenue functions have the same form as the underlying production functions if, as in my theoretical model, the product demand is CES (Petrin and Sivadasan, 2013; De Ridder et al., 2022). For instance, a Cobb-Douglas production function, with homogenous output elasticities across firms, leads to a Cobb-Douglas revenue function if competition is monopolistic; if, instead, demand elasticities are heterogenous, revenue elasticities will be heterogenous too. Thus, a Cobb-Doug specification cannot approximate revenue functions with variable markups.

However, the issue is attenuated with a translog specification. Because translog parameters interact with inputs, revenue elasticities can vary in firm size. This flexibility allows revenue elasticities to account for firm heterogeneity not just in output elasticities but, potentially, also in demand elasticities. Hence, translog production function and CES demand are sufficient but not strictly necessary conditions for revenues to be well approximated by a translog specification.

The functional form of the revenue function is represented in equation (45), with variables' superscripts denoting logarithmic transformations. I omit sectoral subscripts for visual clarity, although the estimation is separately implemented for each 2-digit industry. The first term in square brackets contains (logs of) inputs entering linearly, the second term contains quadratic transformations, the third term contains interactions, and the

[21]For robustness, I replicate the analysis with a Cobb-Douglas specification, which is less flexible but may perform better in small samples thanks to its simpler structure (see Appendix F).

fourth corresponds to the residual. The latter consists of a transitory shock unobserved to the firm, ϵ_{kt}, and an observable (observable to the firm, not to the econometrician) persistent term, ω_{kt}, which includes demand and supply shocks affecting revenues.[22] Capital and materials are denoted by κ_{kt} and m_{kt}, respectively.

$$
\begin{aligned}
& r_{kt}^{ln} \\
(45) \qquad & = [\beta^{\kappa}\kappa_{kt}^{ln} + \beta^{l}l_{kt}^{ln} + \beta^{m}m_{kt}^{ln}] + \left[\beta^{(\kappa^2)}(\kappa_{kt}^{ln})^2 + \beta^{(l^2)}(l_{kt}^{ln})^2 + \beta^{(m^2)}(m_{kt}^{ln})^2\right] \\
& + [\beta^{\kappa l}\kappa_{kt}^{ln}l_{kt}^{ln} + \beta^{km}\kappa_{kt}^{ln}m_{kt}^{ln} + \beta^{lm}l_{kt}^{ln}m_{kt}^{ln}] + [\omega_{kt} + \epsilon_{kt}]
\end{aligned}
$$

As in the case of production function estimation, a key issue in estimating revenue functions is the endogeneity of inputs: the observable revenue shock, ω_{kt}, is a component of the residual and, at the same time, affects the input demand, distorting the estimated elasticities. Following Petrin and Sivadasan (2013), I deal with this problem by estimating revenue elasticities with a 'control function' approach, traditionally used for production function estimation (Olley and Pakes, 1996; Ackerberg et al., 2015).[23] The key difference in the case of revenue functions is that ω_{kt}, a first-order Markov process, does not correspond to productivity as it also captures demand shocks. Appendix E illustrates the key assumptions and operational steps of this methodology.

The resulting revenue elasticities can vary across firms depending on the levels of each input, capturing both the relative input usage and firm size, which means that it can adjust to better fit variable output elasticities or variable demand elasticities. For the case of labour, the revenue elasticity has the following form.

$$
(46) \qquad \varepsilon_{kt}^{r,l} = \frac{\partial r_{kt}^{ln}}{\partial l_{kt}^{ln}} = \beta^{l} + 2\beta^{(l^2)}l_{kt}^{ln} + \beta^{\kappa l}\kappa_{kt}^{ln} + \beta^{lm}m_{kt}^{ln}
$$

A summary of the estimated elasticities is presented in Table 2.

Market-level index. My empirical approach allows the estimated markdowns, $\hat{\chi}_{kt}$, to vary across firms and periods. In a subsequent step, I average the estimates at the market-level, defined in terms of administrative provinces (i.e. NUTS3), to obtain an index of local monopsony power corresponding to $(1+\theta)$ in the theoretical model. The purpose of this aggregation is to validate the general equilibrium mechanism proposed in

[22]Note that the residual of a revenue function, used for the moment conditions, does not correspond to TFPR as revenue elasticities do not correspond to output elasticities (Foster et al., 2016). More details on this methodology are available in Appendix E.

[23]As revenues are directly observed, my estimation does not suffer from the 'price bias' (Bond et al., 2021).

Section 2, by comparing the economic performance of labour markets with different levels of monopsony power.

TABLE 2. FIRM-LEVEL REVENUE ELASTICITIES

Input	Mean	SD	P10	P50	P90	Obs.
Capital elasticity	.03	.02	.01	.02	.05	360515
Labour elasticity	.27	.14	.11	.26	.45	360515
Intermediates elasticity	.76	.13	.58	.77	.91	360515

Note: Summary statistics of revenue elasticities in 2019, computed using a translog specification and a control function approach. Parameters are separately estimated for each 2-digit industry using the full sample: Italy, 2006-2019. Microdata from CompNet.

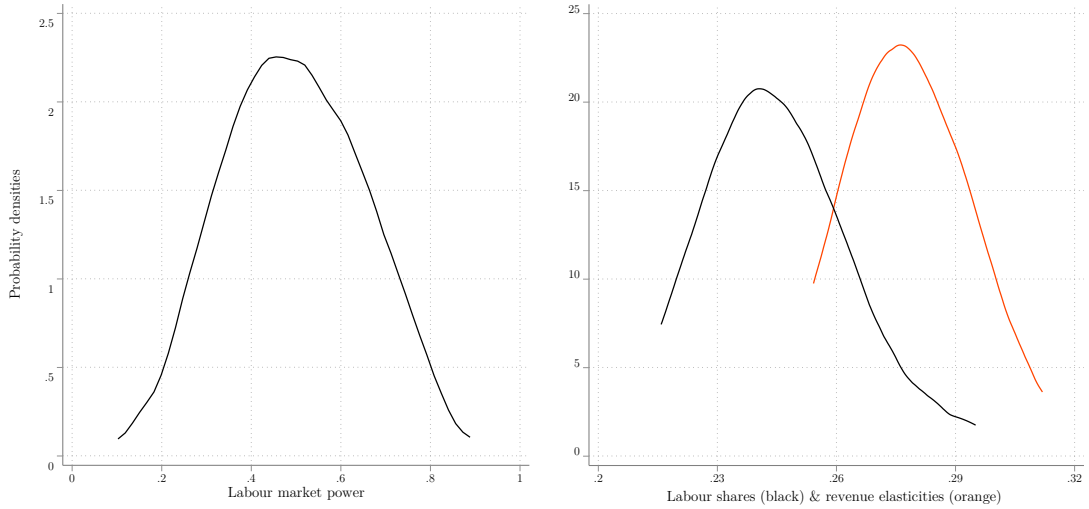
Equation (47) reports the index of labour market power, $\widehat{\theta}_{Kt}$, for locality K in period t. E_{Kt} denotes the number of firms, borrowing the notation of Section 2.

$$1 + \hat{\theta}_{Kt} = \frac{1}{E_{Kt}} \sum\nolimits_{k \in K} \hat{\chi}_{kt} \tag{47}$$

The baseline estimation uses markdown averages as a market-level measure of monopsony power. Additionally, I replicate the empirical analysis using medians and 95$^{\text{th}}$ percentiles, and results remain consistent with the theoretical mechanism (see Appendix F). An alternative aggregation, based on the harmonic means of wage markdowns and price markups, has the advantage of recovering the market-level MRPL to wage ratio, mirroring firm-level markdowns (Yeh et al., 2022). Despite this desirable theoretical property, its application is constrained by the availability of price markups, whose levels cannot be identified without observing firm-level prices (Bond et al., 2021; De Ridder et al., 2022).

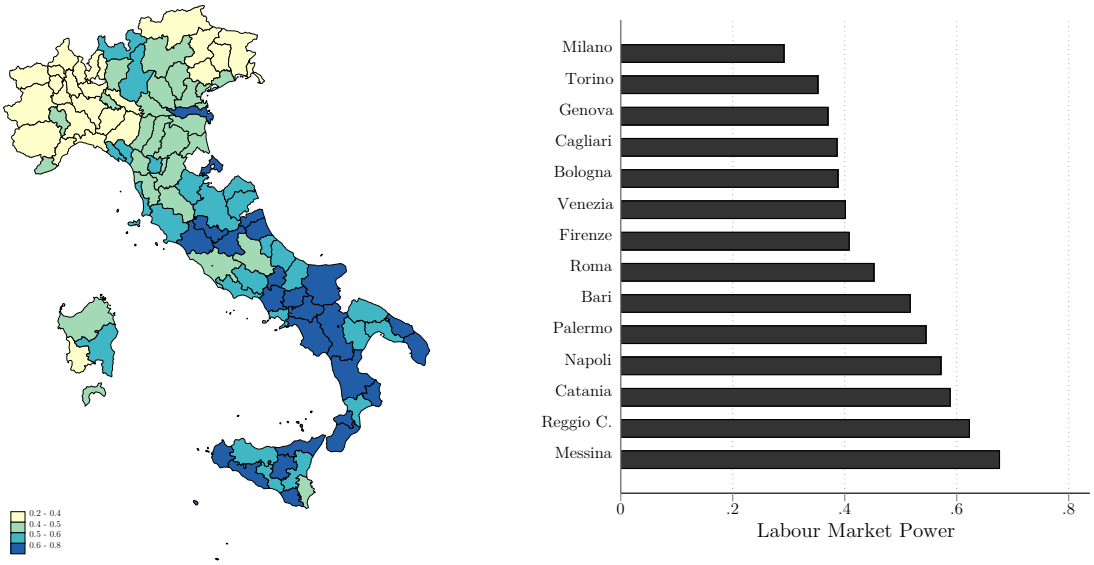
Estimation results, represented in Figures 6 and 7, exhibit a significant departure from the competitive equilibrium. Under perfect competition, labour shares of revenues are equal to revenue elasticities (i.e. markdowns equal to one), and the index of labour market power falls to zero. Instead, in 2019, 90% of the distribution of $\widehat{\theta}_{Kt}$ lies between 0.28 and 0.73, with an average of 0.49. The deviation is particularly noticeable in Southern provinces, where $\widehat{\theta}_{Kt}$ averages 0.61. Furthermore, there is a significant gap between cities, classified using the administrative definition (i.e. 'Metropolitian City'), and other provinces: 0.38 versus 0.44 in the North; 0.56 versus 0.62 in the South.

FIGURE 6. MARKET-LEVEL ESTIMATES OF LABOUR MARKET POWER



Note: On the left: probability density of the province-level index of labour market power, measured as the average markdown (minus one), from the estimation of firm-level revenue functions. On the right: probability densities of province-level averages of revenue elasticities and labour shares of revenues, i.e. the two firm-level components of the markdown estimator. Microdata on Italian firms in 2019 from CompNet.

FIGURE 7. GEOGRAPHIC HETEROGENEITY IN LABOUR MARKET POWER



Note: On the left: labour market power across provinces, measured as the average markdown (minus one), from the estimation of firm-level revenue functions. On the right: labour market power for the subset of provinces classified as 'Metropolitan cities' (administrative definition). Microdata on Italian firms in 2019 from CompNet.

3.3 Testable predictions

The key theoretical propositions and corollaries demonstrated in Section 2 correspond to testable predictions on the role of labour market power, which I empirically use to validate the proposed mechanism. Because my theory is centred on a general equilibrium effect, the relationships to test are defined at the market-level. Thus, the scope of this analysis is to compare a set of indicators of firm behaviour and performance across labour markets, classified by the estimated level of monopsony power. To increase comparability across observations, I focus on markets belonging to a single country, Italy, so that they share a similar context from a social and regulatory perspective.

I use province-year observations to proxy local labour markets (103 provinces in 14 periods, 2006-2019). Provinces are administrative divisions corresponding to NUTS3 regions in the EU framework, each with, on average, 230,000 individuals involved in economic activity (labour or entrepreneurship). I choose this level of aggregation, broader than commuting zones, to capture the choice of workers across a large set of options with different appeals, as my theoretical setting features workers with heterogeneous taste for employers.

Table 3 summarises ten testable theoretical predictions and the indicators used for their empirical validation.

Misallocation. The first set of testable predictions relates to the first channel of my theoretical mechanism: misallocation towards small firms. The model predicts that higher levels of labour market power are associated with more dispersion in wages, MRPL, and TFPR, as well as lower dispersion in firm size. In monopsonistic and oligopsonistic labour markets, small firms pay lower wages than large firms for the same worker type, because the labour supply observed by each firm is upward-sloping. This effect creates wage dispersion and, as a result, reduces dispersion in firm size: more productive firms will still be larger than the less productive ones, but the size gap is lower as the former ones face relatively higher labour costs per worker. Because of this distortion, workers are reallocated inefficiently towards small firms, despite the latter having lower levels of MRPL. Hence, differences in the MRPL across firms increase and dispersion in MRPL translates into dispersion in TFPR.

In the theoretical model, firm-level variables are transformations of the managerial ability, the primitive μ_k, thus they follow a Pareto distribution. Higher dispersion corresponds to a lower value in the underlying shape parameter. It is straightforward to prove analytically that the shape of wages, MRPL, and TFPR falls in θ, the parameter representing labour market power, whereas the shape of firm size increases in θ. As labour

market power also affects the scale of these variables, the appropriate test requires a unitless measure of dispersion, which prevents, for instance, the use of standard deviations. I use the ratio of the 75$^{\text{th}}$ to 25$^{\text{th}}$ percentile, which is robust to outliers and corresponds to a simple transformation of the shape parameter: $\frac{p75}{p25} = 3^{\frac{1}{\tilde{\alpha}}}$, with $\tilde{\alpha}$ being the shape of a generic Pareto distribution.

TABLE 3. TESTABLE THEORETICAL PREDICTIONS

Channel	Prediction	Indicator	Figure	Table
	Higher levels of labour market power are associated with:			
Misallocation	i. More dispersion in wages	Labour costs per worker: p75/p25	8	4
Misallocation	ii. Less dispersion in firm size	Number of employees: p75/p25	8	4
Misallocation	iii. More dispersion in MRPL	MRPL (parametric): p75/p25	9	4
Misallocation	iv. More dispersion in TFPR	TFPR (cost-based): p75/p25	9	4
Excess entry	v. Lower average wages	Labour costs per worker: aggregate	10	5
Excess entry	vi. Higher entrepreneurship rates	Share of self-employed[24]	10	5
Excess entry	vii. Lower average firm size	Number of employees: average	11	5
Excess entry	viii. Lower average productivity	Value added per worker: average	11	5
Hindered technology	ix. Reduced technology diffusion	Intangible assets: average (in logs)	12	6
Hindered technology	*As above*	Intangible assets per worker: aggregate	12	6
Resulting effect	x. Lower aggregate productivity	Value added per worker: aggregate	13	6
Resulting effect	*As above*	TFPR (cost-based): weighted average	13	6

Note: Testable predictions are classified according to the three channels discussed in Section 2 (misallocation, excess entry, hindered technology). Referenced figures and tables report the results of the empirical tests.

I measure firm-level wages in gross terms, as the ratio of labour costs to the number of employees. The latter is then used as a measure of firm size. Whilst the MRPL is measured from revenue function estimation (see Section 3.2), TFPR is estimated non-parametrically, using cost-shares to proxy output elasticities.[25] This choice reflects that

[24]As I discuss in the text, results also hold using the share of employers instead of all self-employed.

[25]The result based on TFPR complements the result based on MRPL, but its measurement is somewhat problematic in my context as input market power distorts cost shares.

the residual of revenue functions does not correspond to TFPR, because revenue elasticities differ from output elasticities (Foster et al., 2016).

Excess entry. The second set of testable predictions concerns the excess entry of low-ability entrepreneurs. The model predicts that higher levels of labour market power are associated with lower wages, higher entrepreneurship rates, smaller firm size, and lower average productivity. Imperfect competition in labour markets allows firms to pay wages below the competitive level, generating a fall in the average wage across workers. The reduction in salaries is an advantage for entrepreneurs, which face lower costs, and a disadvantage for workers, who observe a drop in their income. Such effect distorts the occupational choice, making entrepreneurship more prevalent. With additional firm entry, workers are allocated across a larger number of firms, reducing firm size on average. Furthermore, because the additional mass of entrepreneurs consists of agents whose ability is at the margin of the entry cutoff, average productivity falls.

I measure the average wage across workers in gross terms, as the ratio of aggregate labour costs to aggregate employment. Empirically, entrepreneurship can be defined considering all self-employed or employers exclusively. If only employers are considered, the entrepreneurship rate becomes a simple transformation of the average firm size in labour units: the entrepreneurship rate would correspond to $\frac{E}{E+L} = \frac{1}{1+(L/E)}$, with E denoting employers and L denoting employees. Because I also test the relationship between average firm size in labour units and labour market power, to avoid redundancy I validate the prediction on entrepreneurship considering all self-employed, regardless of whether they employ workers or not. Average productivity is defined as the average value of value added per worker across firms.

Hindered technology. The last set of predictions involves technology diffusion as well as the combined effect of all three channels on aggregate productivity. The model predicts that higher levels of labour market power are associated with lower usage of productivity-enhancing technologies and lower aggregate productivity. Technologies characterised by higher sunk costs and lower marginal costs are more beneficial to large firms (see Section 2.7). By distorting the size distribution, labour market power limits the diffusion of such technologies: first, because workers are reallocated towards smaller firms, which have lower rates of adoption rates; second, because it generates the excess entry of micro firms which, again, have lower adoption rates. Eventually, the three key channels (i.e. misallocation, excess entry, and hindered technology) contribute to the fall in aggregate productivity.

Figure 1, in the Introduction, shows that smaller firms have significantly lower rates of software adoption, usage of ultrafast broadband, and involvement in data analysis; the

same patterns, consistent with my theory, can be replicated for the other intangible technologies available. The underlying data, which come from a large data collection of ISTAT, are not disaggregated at the province-level, preventing the test of their relationship with labour market power. The only relevant indicator available in firm-level financial data is the stock of intangible fixed assets. Although it can include unrelated information (e.g. accounting goodwill), it largely includes intangible technologies characterised by a scale-bias. My tests consider two indicators: the stock of intangibles (in logs) and, in order to also have a measure relative to size, the market-level ratio of intangibles to employment.

Finally, aggregate productivity is measured in terms of value added per worker, aggregated using employment weights, and TFP, measured by aggregating cost-based TFPR values with value added weights.[26] The former indicator corresponds to labour productivity and recreates the market-level ratio of aggregate value added to aggregate employment, coherently with the model's definition of aggregate productivity. The latter, aggregate TFP, is used as a complementary indicator (see footnote 24), to ensure that the results using labour productivity are not critically dependent on differentials in capital intensity.

3.4 Empirical results

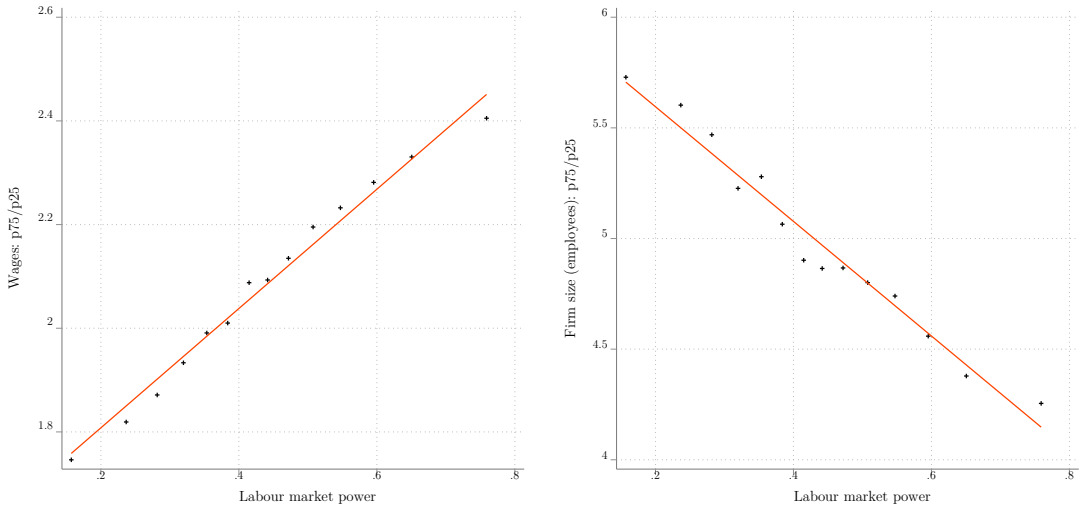
I test the aforementioned predictions in two ways: (i) a visual representation of the unconditional correlations using binned scatter plots, and (ii) OLS regressions of the indicators (in logs) with a set of controls (i.e. year fixed effects and dummy variables for Southern provinces and large cities) and two-way clustered standard errors. Results of the former are presented in Figures 8-13.[27] Tables 4-6 contain the regression results for the main specification, whereas Appendix F has the same tests using alternative estimators of labour market power.

The evidence confirms the relationships predicted by the theory: across local labour markets, higher levels of labour market power are associated with (i) more dispersion in wages, (ii) less dispersion in firm size, (iii) more dispersion in MRPL, (iv) more dispersion in TFPR, (v) lower average wages, (vi) higher entrepreneurship rates, (vii) lower average firm size, (viii) lower average productivity, (ix) reduced technology diffusion, and ultimately, (x) lower aggregate productivity.

[26]Note that TFPR, not TFPQ, is the model-consistent firm-level variable to derive aggregate productivity.

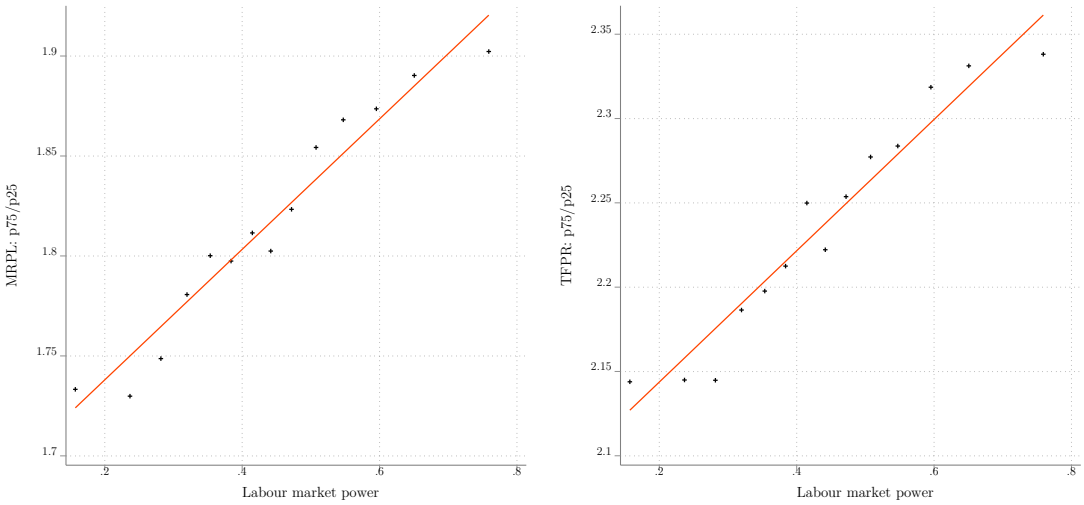
[27]See Cattaneo et al. (2024) for a discussion of the econometric properties of binned scatter plots as a method of data visualisation.

FIGURE 8. DISPERSION IN WAGES AND FIRM SIZE



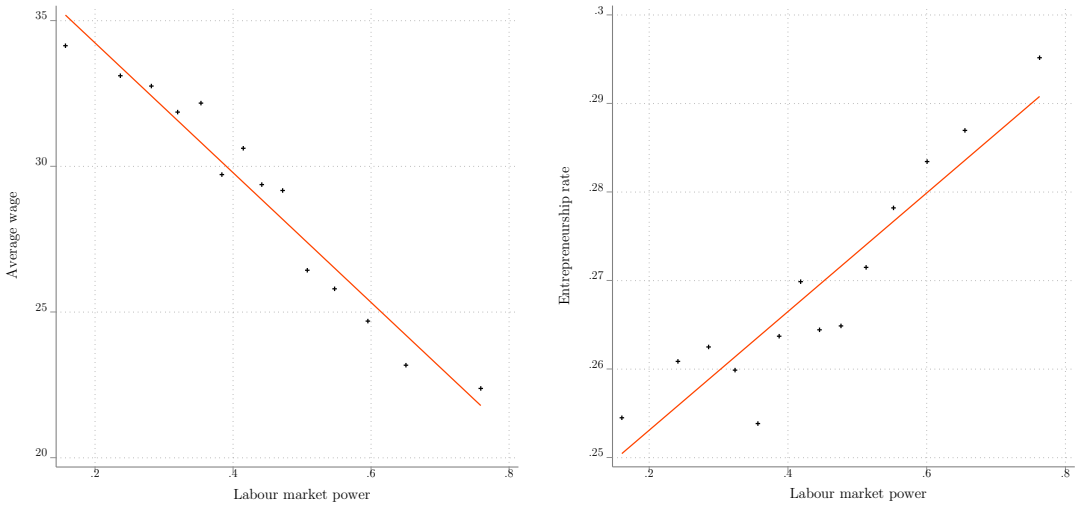
Note: Each cross represents 100 local labour markets in Italy, defined as province-year observations (2006-2019). Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. On the left: ratio of percentiles (75^{th} to 25^{th}) of labour costs per employee. On the right: ratio of percentiles (75^{th} to 25^{th}) of the number of employees. Microdata from CompNet.

FIGURE 9. DISPERSION IN MRPL AND TFPR



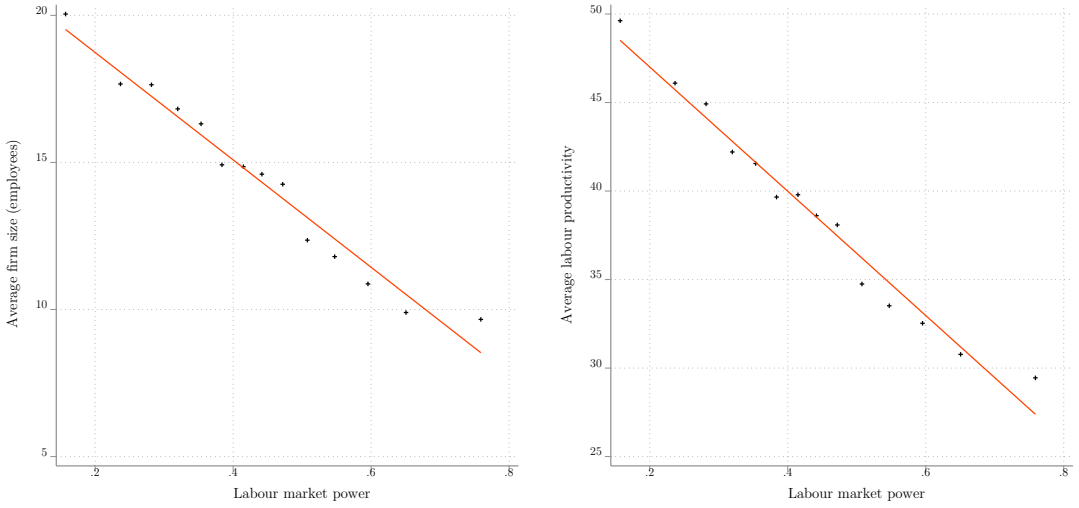
Note: Each cross represents 100 local labour markets in Italy, defined as province-year observations (2006-2019). Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. On the left: ratio of percentiles (75^{th} to 25^{th}) of the MRPL, measured from the estimation of firm-level revenue functions. On the right: ratio of percentiles (75^{th} to 25^{th}) of TFPR, measured with a cost-share approach. Microdata from CompNet.

FIGURE 10. WAGES AND ENTREPRENEURSHIP



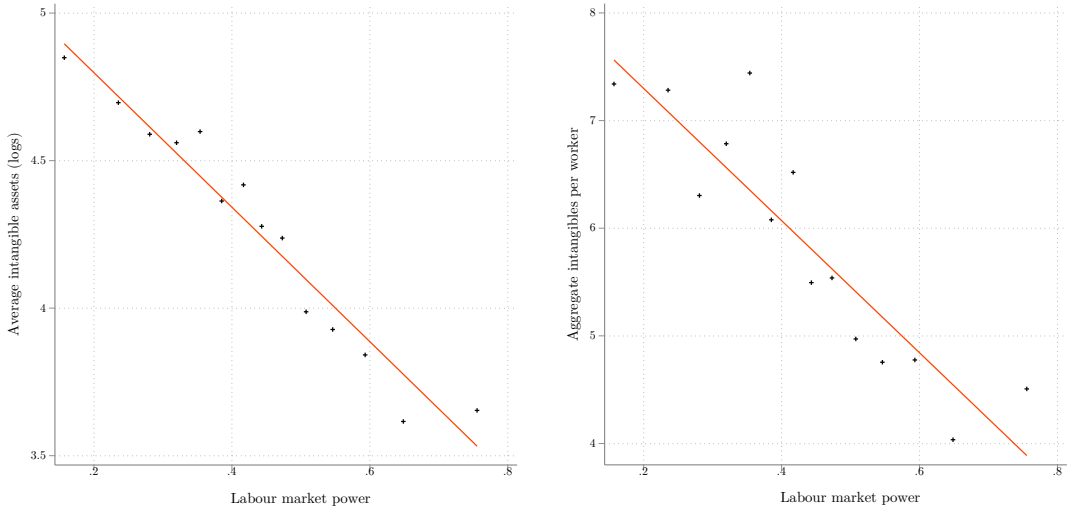
Note: Each cross represents 100 local labour markets in Italy, defined as province-year observations (2006-2019). Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. On the left: ratio of aggregate labour costs to aggregate employment (thousands of 2005 Euros). On the right: share of self-employment. Microdata from CompNet. Self-employment data from ISTAT.

FIGURE 11. AVERAGE FIRM SIZE AND PRODUCTIVITY



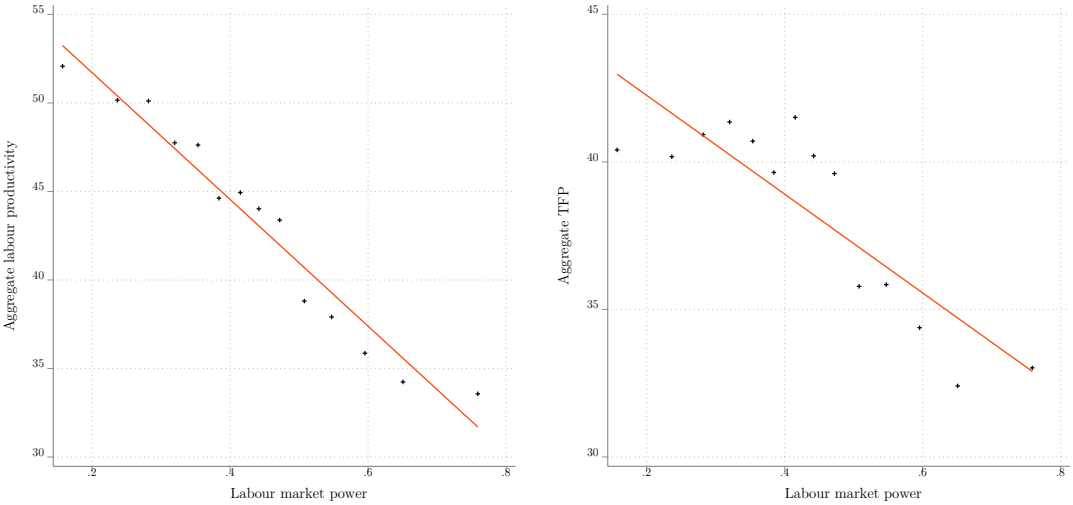
Note: Each cross represents 100 local labour markets in Italy, defined as province-year observations (2006-2019). Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. On the left: average number of employees. On the right: average value added per employee (thousands of 2005 Euros). Microdata from CompNet.

FIGURE 12. INTANGIBLE ASSETS



Note: Each cross represents 100 local labour markets in Italy, defined as province-year observations (2006-2019). Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. On the left: average stock of intangible assets (thousands of 2005 Euros, in logs). On the right: ratio of aggregate intangible assets to aggregate employment. Microdata from CompNet.

FIGURE 13. AGGREGATE PRODUCTIVITY



Note: Each cross represents 100 local labour markets in Italy, defined as province-year observations (2006-2019). Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. On the left: value added per employee, aggregated using employment weights (thousands of 2005 Euros). On the right: TFPR, measured with a cost-share approach, aggregated using value added weights (thousands of 2005 Euros). Microdata from CompNet.

TABLE 4. FIRST SET OF PREDICTIONS

	(1) Wage dispersion (logs)	(2) Firm size dispersion (logs)	(3) MRPL dispersion (logs)	(4) TFPR dispersion (logs)
Labour market power	0.355**	-0.257*	0.113**	0.113*
	(0.161)	(0.143)	(0.0401)	(0.0554)
South (dummy)	0.00378	-0.0902**	0.0313***	0.0174
	(0.0292)	(0.0301)	(0.00849)	(0.0120)
City (dummy)	0.0128	0.0226	0.0198*	0.0299**
	(0.0159)	(0.0292)	(0.00948)	(0.0134)
Observations	1,436	1,436	1,436	1,436
R-squared	0.777	0.314	0.373	0.291

Note: *** p<0.01, ** p<0.05, * p<0.1. OLS regressions using province-year observations from Italy, 2006-2019. Year fixed effects included. Two-way clustered standard errors in parentheses. Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. South defined using the NUTS1 classification (i.e. 'ITF'). City defined using the administrative definition (i.e. 'Metropolitan City'). Dispersion defined as ratio of percentiles (75^{th} to 25^{th}). Dependent variables defined as follows: (1) labour costs per employee; (2) number of employees; (3) MRPL, measured from the estimation of firm-level revenue functions; (4) TFPR, measured with a cost-share approach. Microdata from CompNet.

TABLE 5. SECOND SET OF PREDICTIONS

	(5) Average wage (logs)	(6) Entrepreneurship rate (logs)	(7) Average firm size (logs)	(8) Average labour productivity (logs)
Labour market power	-0.548***	0.394**	-0.948***	-0.497***
	(0.0911)	(0.164)	(0.215)	(0.162)
South (dummy)	-0.235***	-0.00458	-0.317***	-0.205***
	(0.0248)	(0.0358)	(0.0559)	(0.0295)
City (dummy)	0.0596*	-0.103***	0.129**	0.0462**
	(0.0300)	(0.0234)	(0.0480)	(0.0195)
Observations	1,436	1,370	1,436	1,436
R-squared	0.732	0.289	0.542	0.772

Note: *** p<0.01, ** p<0.05, * p<0.1. OLS regressions using province-year observations from Italy, 2006-2019. Year fixed effects included. Two-way clustered standard errors in parentheses. Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. South defined using the NUTS1 classification (i.e. 'ITF'). City defined using the administrative definition (i.e. 'Metropolitan City'). Dependent variables defined as follows: (5) ratio of aggregate labour costs to aggregate employment; (6) share of self-employment; (7) average number of employees; (8) average value added per employee. Microdata from CompNet. Self-employment data from ISTAT.

TABLE 6. THIRD SET OF PREDICTIONS

	(9) Average intangible assets (logs)	(10) Intangibles per worker (logs)	(11) Aggregate labour productivity (logs)	(12) Aggregate TFP (logs)
Labour market power	-0.232***	-0.286	-0.537***	-0.304
	(0.0737)	(0.219)	(0.110)	(0.218)
South (dummy)	-0.202***	-0.149**	-0.254***	-0.259***
	(0.0228)	(0.071)	(0.0335)	(0.0479)
City (dummy)	0.141***	-0.027	0.0222	0.177***
	(0.0226)	(0.0719)	(0.0210)	(0.0354)
Observations	1,376	1,436	1,436	1,433
R-squared	0.621	0.079	0.667	0.593

Note: *** p<0.01, ** p<0.05, * p<0.1. OLS regressions using province-year observations from Italy, 2006-2019. Year fixed effects included. Two-way clustered standard errors in parentheses. Labour market power measured as the average markdown (minus one), from the estimation of firm-level revenue functions. South defined using the NUTS1 classification (i.e. 'ITF'). City defined using the administrative definition (i.e. 'Metropolitan City'). Dependent variables defined as follows: (9) aggregate stock of intangible assets per worker; (10) aggregate intangible assets to aggregate fixed assets; (11) value added per employee aggregated using employment weights; (12) TFPR, measured with a cost-share approach, aggregated using value added weights. Microdata from CompNet.

4 Quantification

The model proposed in Section 2 can be used to quantify aggregate losses induced by monopsony power and decompose them into the three key channels: misallocation, excess entry, and limited technology diffusion. For this purpose, I calibrate the exogenous parameters (σ, θ, α, γ, τ,) using a combination of moments from financial statements, surveys, and national accounts.

4.1 Calibration

The two parameters generating market power in product and labour markets are calibrated externally. The parameter of job differentiation, generating monopsony power, is set as $\theta = 0.61$ to match the average wage markdown, estimated with CompNet microdata following the procedure described in Section 3.2. The elasticity of substitution across product varieties is set as a midpoint in a range of estimates from the literature, $\sigma = 3.5$.[28] Under profit maximisation, this value corresponds to a markup of 1.4 in levels

[28]Other commonly used calibration values for the elasticity of substitution come from Bernard et al. (2003), Broda and Weinstein (2006), and Hsieh and Klenow (2009), $\sigma = 3.79$, $\sigma = 4.25$, and $\sigma = 3$ respectively.

and 0.34 in logs, in line with values found by De Ridder et al. (2022) across French firms (i.e. 0.32 using the arithmetic average of firm-level markups, 0.37 using the sales-weighted harmonic mean). To the best of my knowledge, no contribution has yet provided estimates of markups for the Italian economy observing firm-level prices, which is a critical condition to avoid the 'price bias' in the production approach and identify the markup level (Bond et al., 2021; De Ridder et al., 2022).

The remaining parameters are estimated matching the model to the data. The shape of the ability distribution, α, is identified using a definition of aggregate profit share that excludes the imputed labour income of entrepreneurs, assumed to be equal to the average value of labour compensation per employee. For this specific definition, the model gives a simple formula, reported in equation (48), which grows in σ, falls in α, and is independent of job differentiation (θ) and the technology specifications (γ and τ). The intuition is that both greater product substitutability (higher σ) and greater firm heterogeneity (lower α) lead to a more concentrated market with few firms earning high profits; vice versa, in the equilibrium with lower σ or higher α, firms are smaller and large parts of their profits are the entrepreneur's labour income. Having set σ from the literature, I derive $\alpha = 3.4$ analytically using the aggregate profit share from EUKLEMS (i.e. 21.3%).

$$\frac{Y - \overline{w}(L^P + L^T + E)}{Y} = \frac{\sigma - 1}{\alpha\sigma} = 0.213 \tag{48}$$

Unlike σ, the job differentiation parameter, θ, does not enter in the aggregate profit equation. The reason is that product market power has a larger effect on output than on wages, whereas labour market power affects both in the same way, leaving the profit share unaffected. For the same reason, the aggregate profit share in my model is also independent of the technology specifications. A related implication is that I can quantify the losses due to misallocation and excess entry without estimating the parameters of monopsony power and the superior technology, by just comparing aggregate output and productivity for different levels of θ, given α and σ.

Two firm-level parameters characterise the scale-biased technology, τ and γ, the sunk cost of upgrading in labour units and the productivity growth for adopters, respectively. To estimate their value, I consider how they affect aggregate outcomes. Following a growing literature that links the shift towards sunk costs to digital technologies (Hsieh and Rossi-Hansberg, 2023; De Ridder, 2024), I use data on IT adoption and expenditure. Specifically, I target two moment conditions: (i) the aggregate expenditure on technology adoption in the model is set equal to the aggregate software expenditure in the data, both expressed in proportion to labour costs, and (ii) the share of firms adopting the superior technology in the model is set equal to the rate of software adoption in the data.

TABLE 7. CALIBRATION RESULTS

Calibration	Parameter	Description	Value	Target	Source
External	θ	Job differentiation	0.61	Average wage markdown	CompNet
External	σ	Product substitutability	3.5	Average price markup	*Literature*
Internal	α	Pareto shape parameter of the ability distribution	3.4	Aggregate profit share	EU Klems
Internal	τ	Cost of technology adoption in labour units	0.89	Aggregate ratio of software to labour expenditure	IntanProd
Internal	γ	Growth in firm productivity upon technology adoption	0.23	Share of firms adopting software	ISTAT

Note: The external calibration of θ and σ is based on microdata on markdowns (CompNet) and markups (De Ridder et al., 2022). Other parameters are identified matching the model to the data. α is set to replicate the aggregate profit share (0.213). τ and γ are estimated targeting the ratio of aggregate software to labour expenditure and the adoption rate of software (0.038 and 0.057, respectively).

The aggregate ratio of software expenditure to labour expenditure in Italy is directly observable in EUKLEMS (3.8%). For the adoption rate, I consider the share of firms that use at least one enterprise software, such as Enterprise Resource Planning (ERP), Customer Relationship Management (CRM), or Business Intelligence (BI). From ISTAT, I can find survey data on the adoption rates of such application software by size class in Italy. However, as firms with less than ten employees are not surveyed, I cannot observe the share of adopting firms in the universe of companies. In my theoretical setting, a size cutoff splits firms between the baseline and the superior technology. The share of adopting firms in ISTAT data is indeed steeply growing in firm size (see Figure 1), although with type-1 and type-2 errors that the model does not rationalise. To proxy the share of adopting firms across all size classes, I make an educated guess on the size class in which the adoption cutoff lies, compute the share of adopting firms in such class, and then assume no adoption in inferior classes and full adoption in superior classes. Considering the data on adoption rates (45% among firms with 10-49 employees, 67% among firms with 50-99 employees, etc.), I can reasonably guess that the adoption cutoff lies in the size class corresponding to firms with 10-49 employees.[29] I then compute the mass of adopters as 45% of firms with 10-49 employees plus all firms with at least 50

[29]I formalise this guess in Appendix G by assuming that the probability of adoption depends on the percentage deviation of firm size from the cutoff through a logistic function. Simulating the adoption rates in each size class, I find that the observed data can only be explained by a cutoff between 10 and 49 employees.

employees. This amount corresponds to 5.7% of all employers, which is my estimate of the adoption rate across all size classes.

As σ, θ, and α are already estimated, equations (49) and (50) form a system in two unknowns (γ and τ), which is solved numerically.

$$\frac{\overline{w}L^T}{\overline{w}(L^P + L^T)} = 0.038 \tag{49}$$

$$\frac{1 - \Phi(\mu_H)}{1 - \Phi(\mu_E)} = 0.057 \tag{50}$$

The results of the model's calibration are summarised in Table 2. $\alpha = 3.4$ is consistent with Ghironi and Melitz (2005). The estimated parameters characterising the superior technology are $\tau = 0.89$ and $\gamma = 0.23$, which imply that firms can increase their productivity by 23% paying slightly less than the average annual cost of a worker. This number can be prohibitive for most firms in an economy in which the median number of employees is five (see Table 1). Considering labour costs in Italy, the sunk cost of adoption corresponds to nearly €38,000 in local currency, or $40,000 in US Dollars (current prices, 2024). Using equation (19) and the estimated parameters, the growth in productivity can be expressed as an approximate 30% growth in variable profits ($\Delta\pi_k\% \approx \frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}\Delta a_k\%$).

4.2 Definitions

Decompositions. Mirroring the firm-level production function in equation (10), I define aggregate productivity as the ratio of aggregate output to the mass of production workers, $A \equiv \frac{Y}{L^P}$. Aggregate productivity has the same two components of firm productivity: managerial ability and technology. The aggregate value of managerial ability is defined as the level of aggregate productivity that occurs if all firms use the baseline technology, $M \equiv A|_{\gamma=0}$; the technology component captures the additional productivity induced by the diffusion of the superior technology, $T \equiv \frac{A}{A|_{\gamma=0}}$.

$$Y = AL^P = MTL^P \tag{51}$$

I decompose the management factor in two terms: (i) the *average* productivity that occurs if all firms use the baseline technology, $\overline{M} = (\frac{\alpha}{\alpha-1})(\mu_E|_{\gamma=0})$, which depends on the entry cutoff, μ_E, and captures the toughness of the market selection; and (ii) a residual component, $\Omega \equiv \frac{M}{\overline{M}}$, which corresponds to the market's allocative efficiency. If all active firms have the same level of managerial ability, $\overline{M}$, market shares are equalised and the

average managerial ability represents the aggregate component, i.e. $M = \overline{M}$. However, with firm heterogeneity, market forces raise M above $\overline{M}$ by reallocating workers towards better managed firms.

The three key channels theorised in Section 2 can be understood by considering the effect of variations in θ, the parameter of job differentiation that generates monopsony power, on the three components of aggregate productivity: (i) *allocation*, Ω, which falls with the reallocation of employment shares towards less productive firms, (ii) *selection*, $\overline{M}$, which falls with the entry of low-ability entrepreneurs, and (iii) *technology*, T, which falls when technology diffusion across workers is hindered.

$$A = \Omega \overline{M} T \tag{52}$$

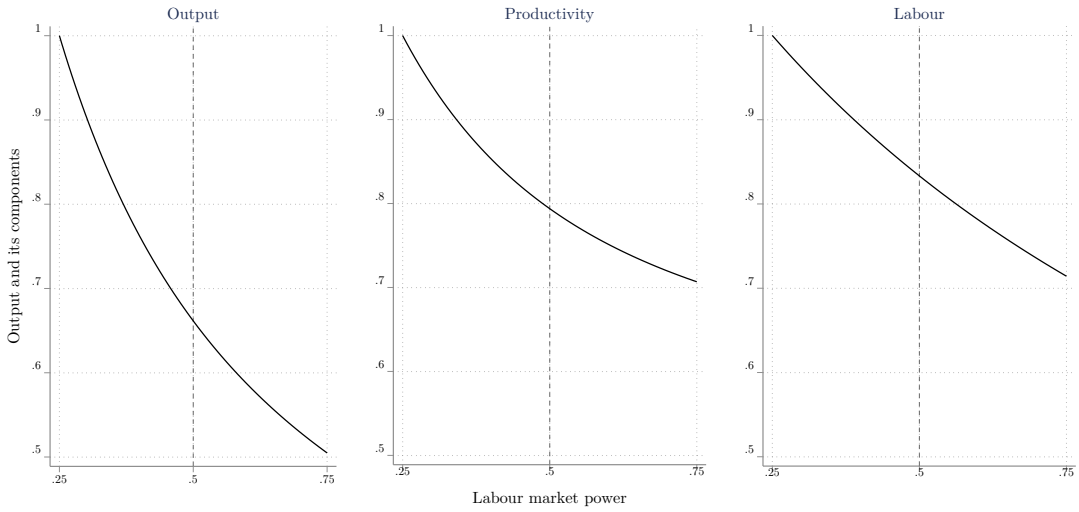
Competitive benchmark. Given the widespread prevalence of monopsony power across labour markets, perfect competition ($\theta = 0$) is a hardly attainable counterfactual. For a more meaningful comparison, I consider a range of values consistent with my empirical analysis. I define $\theta^* = 0.25$ as a *competitive benchmark*, and show how aggregate output and its components fall moving from θ^* along the range $[0.25, 0.75]$. In my data, this interval corresponds to approximately 93% of the distribution of local monopsony power in 2019, derived from province-level average markdowns. The typical labour market corresponds to $\theta = 0.5$.

Although my interest focuses on aggregate productivity, the output losses induced by labour market power are not exclusively caused by a drop in productivity. An additional effect is due to the fall in employment, which corresponds to the traditional deadweight loss associated with market power. To show the distinct effect of labour market power on aggregate output, productivity, and labour, I represent how $(\frac{Y}{Y^*})$, $(\frac{A}{A^*})$, and $\left(\frac{L^P}{L^{P^*}}\right)$ fall in θ, with star superscripts corresponding to the competitive benchmark (i.e. $\theta^* = 0.25$). Similarly, I show the role of each of the three channels discussed in Section 2 by representing the fall in $(\frac{\Omega}{\Omega^*})$, $\left(\frac{\overline{M}}{\overline{M}^*}\right)$, and $(\frac{T}{T^*})$, respectively the allocation, selection, and technology components.

4.3 Quantitative results

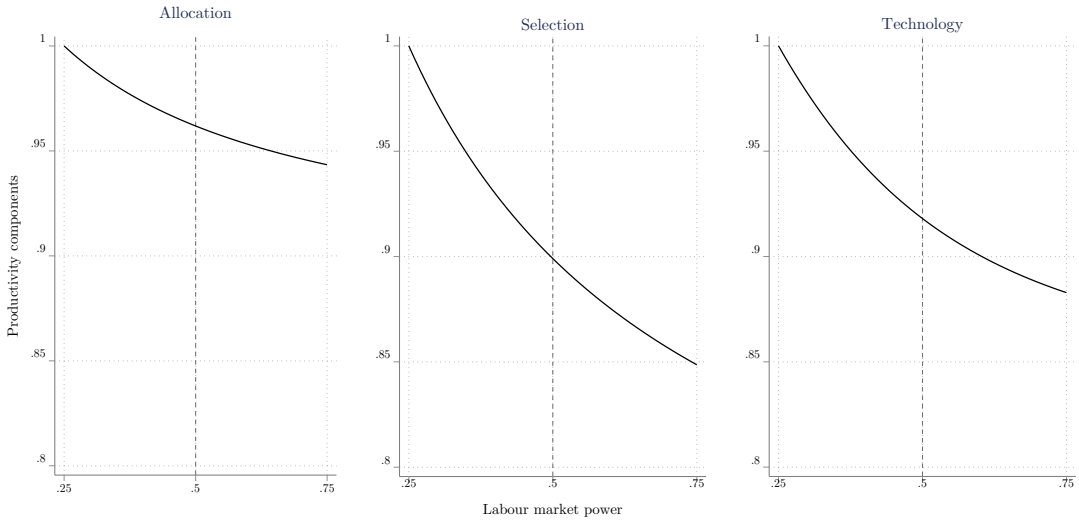
The aggregate losses induced by labour market power are substantial. Moving from the values of θ corresponding to the competitive benchmark to the values of the typical labour market (i.e. from $\theta^* = 0.25$ to $\theta = 0.5$), output falls by 34%. Using equation (51), I attribute a larger share of this loss to aggregate productivity, which has a 21% drop, and the rest to employment, falling by 17%. The evolution in aggregate output, productivity, and labour along the $(0.25, 0.75)$ range of labour market power is reported in Figure 14.

FIGURE 14. AGGREGATE LOSSES



Note: Evolution in aggregate output and its two components (see equation (51)), by levels of labour market power. Variables are expressed in proportion to the competitive benchmark, corresponding to $\theta^* = 0.25$. The interval of labour market power represents approximately 93% of Italian labour markets in 2019. $\theta = 0.5$ denotes the typical labour market. Parameters are set in accordance with the calibration exercise (Table 7).

FIGURE 15. PRODUCTIVITY LOSS: CHANNELS



Note: Evolution in aggregate output and its two components (see equation (52)), by levels of labour market power. Variables are expressed in proportion to the competitive benchmark, corresponding to $\theta^* = 0.25$. The interval of labour market power represents approximately 93% of Italian labour markets in 2019. $\theta = 0.5$ denotes the typical labour market. Parameters are set in accordance with the calibration exercise (Table 7).

TABLE 8. QUANTIFICATION RESULTS

Variable	Definition	Competitive benchmark	Typical labour market	Highest monopsony
θ	Labour market power	0.25	0.5	0.75
Y	Aggregate output	1	0.66	0.5
A	Aggregate productivity	1	0.79	0.71
L^P	Aggregate production labour	1	0.83	0.72
Ω	Allocation term	1	0.96	0.94
$\overline{M}$	Selection term	1	0.9	0.85
T	Technology term	1	0.92	0.88

Note: Aggregate variables (Y, A, L^P) and productivity components (Ω, $\overline{M}$, T) normalised by the value corresponding to the competitive benchmark ($\theta^* = 0.25$), to illustrate the aggregate losses occurring in the typical labour market ($\theta = 0.25$) and in most monopsonistic labour markets ($\theta = 0.75$).

Decomposing aggregate productivity in the three components introduced in equation (52), I find that losses in the selection term and technology term are dominant factors: moving from $\theta^* = 0.25$ to $\theta = 0.5$, the selection term falls by 10% because of the excess entry of low-ability entrepreneurs; the technology term falls by 8% because of the reduced diffusion of the productivity-enhancing technology across firms and workers; the last term, allocation, falls by 4%, reflecting the distortion in the size distribution towards small firms. Figure 15 shows the evolution in the three components of productivity by values of labour market power. Table 8 summarises the quantification of aggregate losses.

Another way to see the effect of labour market power through the three distinct channels is represented by Figure 5, in Section 2. Each panel is produced using parameters from the calibration exercise summarised in Table 5, allowing a direct interpretation of the graphical results.[30] Panel B considers the shift in the cumulative distribution function (CDF) of entrepreneurial ability, moving from the top to the bottom of labour market competition in the usual interval $(0.25, 0.75)$. The entry cutoff in the low-monopsony case, marked by the vertical dashed line, intersects the CDF of the high-monopsony case at approximately 50% of the distribution. This result implies that the

[30]For visual clarity, Panel A considers the case with one technology only (i.e. $\gamma = 0$), as technology adoption creates a discontinuity in the size distribution. The two firms considered in Panel C share the same technology (baseline or superior).

least efficient firm in the competitive benchmark has the same productivity as the median firm in most monopsonistic environments: differential levels of competition across labour markets could generate entirely different levels of managerial quality, even if the underlying distribution of talent is the same.

Panel D (Figure 5) represents the fall in superior technology usage across firms and workers. The share of workers using the superior technology falls from 45% in the competitive benchmark to slightly less than 20% in the typical labour market and to 9% in the least competitive ones. The share of adopting firms falls from 21% to 8% in the typical labour market and 4% in the least competitive. The opportunities offered by new technologies are not fully exploited if high levels of labour market power prevent their firms from reaching the sufficient scale for technological upgrades.

Panel C (Figure 5) considers the size ratio and the MRPL ratio of two firms, with one twice as productive as the other, showing that labour market power lowers the gap in size and increases the gap in MRPL. Quantitatively, the change in the size ratio (on the left scale of Panel C) suggests a large distortion, as the size ratio falls from 2.5 in the competitive benchmark to 1.9 in the typical labour market. However, changes in the MRPL are more muted (on right scale of panel C), consistently with the relatively small misallocation effect (see Figure 15, Allocation). This apparent contradiction is explained by product differentiation. Because small firms can pay higher prices in monopolistic product markets, large productivity differences translate into smaller gaps in MRPL and TFPR.

As a sanity check, in Figure 16 I compare the evolution in aggregate productivity predicted by the calibrated model to the evolution in aggregate productivity observed in the data. Results are comparable for most of the interval of labour market power considered $(0.25, 0.75)$, corresponding to approximately 93% of the estimated distribution; for the highest values of labour market power, the observed fall in aggregate productivity is even larger than what the model predicts, possibly reflecting additional factors.

Quantification: robustness.[31] The fall in aggregate productivity induced by labour market power is the main object of analysis of this paper. To assess the robustness of my estimates, I compute productivity losses for symmetrical deviations in each estimated parameter. Losses are defined as a percentage drop in aggregate productivity moving from the competitive benchmark ($\theta^* = 0.25$) to the typical labour market ($\theta = 0.5$). Under the

[31]Robustness for the empirical validation of the mechanism is provided in Appendix F.

deviations considered, productivity losses range from 14% to 26%, not critically far from the 21% baseline estimation.

The drop in aggregate productivity increases in the elasticity of substitution, σ, which implies that labour market power is more detrimental if product markets are more competitive. In the baseline estimation, I set $\sigma = 3.5$. For $\sigma = 3$, productivity losses are 15%; for $\sigma = 4$, the loss is 26%. The shape of the ability distribution, which is set as $\alpha = 3.4$ in the baseline, has an inverse relationship with the estimated losses. Moving to $\alpha = 2.9$, with greater firm heterogeneity, losses increase to 26%. For $\alpha = 3.9$, the loss is 17%. The availability of better technologies, denoted by a higher growth in productivity upon adoption from the baseline $\gamma = 0.23$ to $\gamma = 0.46$, would increase losses to 26%. Excluding the endogenous technology choice, setting $\gamma = 0$, implies a 14% drop.

The relationship between the sunk cost required for technological upgrades, τ, and the loss in aggregate productivity is non-monotonic: initially, higher τ implies larger losses because the technology has a greater degree of scale-bias, which amplifies the efficiency costs of the size distortions induced by labour market power; after a peak, as adoption becomes increasingly rare even in the competitive benchmark, the loss shrinks and converges to the case of $\gamma = 0$. I find that the loss decreases with both lower and higher values of θ, although the difference from the baseline estimation is small. For both the case of $\tau = 0.3$, in which all entrepreneurs choose to adopt the technology, and $\tau = 1.48$, the loss is 19%.

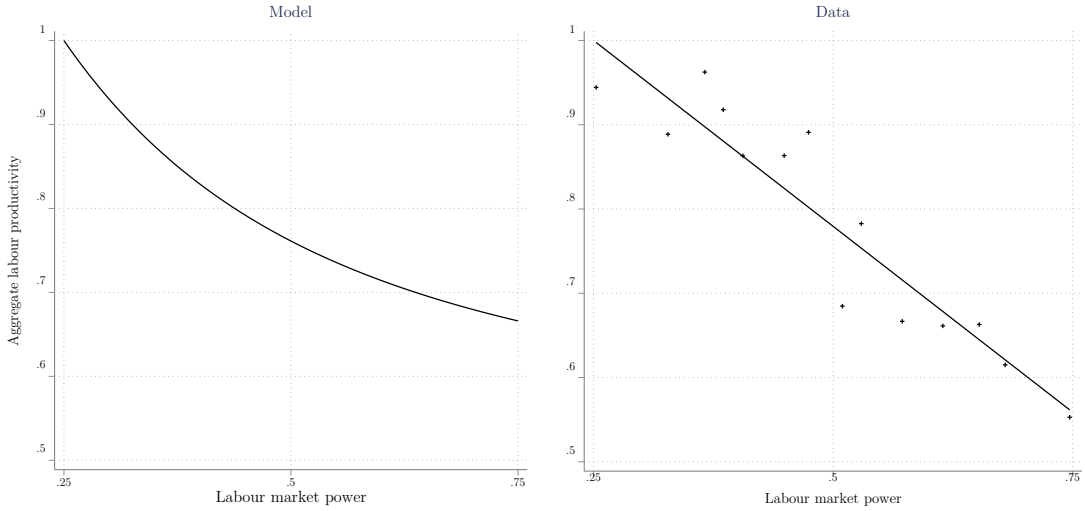
Table 9 summarises the results of these robustness checks.

TABLE 9. PRODUCTIVITY LOSSES: ROBUSTNESS

Baseline	Loss	Lower value	Loss	Higher value	Loss
$\sigma = 3.5$	21%	$\sigma = 3$	15%	$\sigma = 4$	26%
$\alpha = 3.4$	21%	$\alpha = 2.9$	26%	$\alpha = 3.9$	17%
$\tau = 0.89$	21%	$\tau = 0.3$	14%	$\tau = 1.48$	26%
$\gamma = 0.23$	21%	$\gamma = 0$	19%	$\gamma = 0.46$	19%

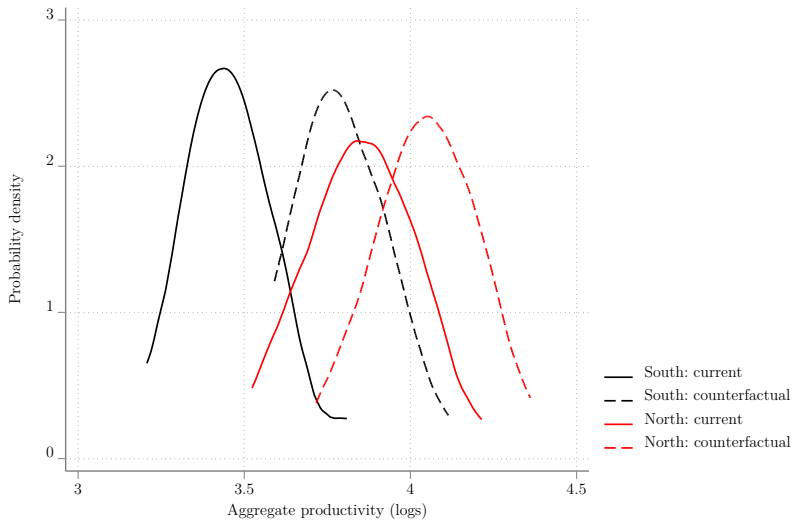
Note: Losses in aggregate productivity induced by labour market power, defined comparing the typical labour market ($\theta = 0.5$) to the competitive benchmark ($\theta^* = 0.25$). Estimates reassessed for symmetrical deviations in each estimated parameter.

FIGURE 16. AGGREGATE PRODUCTIVITY: MODEL VERSUS DATA



Note: On the left: evolution in aggregate productivity predicted by the calibrated model, from Figure 14. On the right: evolution in aggregate productivity observed in the data (data from 2019 for consistency with the model, Figure 13 instead uses 2006-2019). Variables are expressed in proportion to the competitive benchmark, corresponding to $\theta^* = 0.25$. The interval of labour market power represents approximately 93% of Italian labour markets in 2019.

FIGURE 17. GEOGRAPHICAL DISPARITIES



Note: Density of aggregate productivity (in logs) across local labour markets in Northern and Southern Italy. Continuous lines represent real data. Dashed lines represent the model's prediction for the competitive benchmark, corresponding to $\theta^* = 0.25$. Parameters set in accordance with the calibration exercise (Table 7). Microdata on Italian firms in 2019 from CompNet.

Geographical disparities. Lastly, I explore the contribution of labour market power to the productivity gap between Northern and Southern Italy. Using data from 2019, my index of labour market power, $\hat{\theta}_{Kt}$, averages 0.43 across Northern provinces and 0.61 in the South. This large difference could capture several factors: the lower population density in Southern labour markets, which makes job opportunities spatially more dispersed, potentially reducing labour market competition; the poorer quality of local public transports, which further limits employer substitutability; a stronger taste against geographical mobility and commuting, possibly due to cultural reasons; etc. My approach studies the productivity implications of monopsony power without taking a stance on its source, as the estimation is robust to several microfoundations.

The calibrated model allows me to predict the level of aggregate productivity that each province would experience under $\theta^* = 0.25$, the degree of labour market power defined as the competitive benchmark. I use this estimation to measure the evolution of the gap between North and South, moving from the competitive benchmark to the observed values of labour market power. Results are represented in Figure 17. The average level of aggregate productivity across Northern provinces is 18% lower than in the competitive benchmark, whereas the fall across Southern provinces corresponds to 28%. In the benchmark, which I use as a counterfactual, the productivity advantage of Northern provinces is 29%, substantially lower than the 49% advantage observable in real data: differential levels of labour market power can explain 40% of the productivity gap between Northern and Southern Italy.

5 Conclusion

This paper proposes a theory of labour market power as a critical source of efficiency losses. My theoretical and empirical results show that competition in labour markets affects not only wage setting and rent extraction, but also the equilibrium level of aggregate productivity. The mechanism, proved analytically and validated using microdata from Italy, is based on a general equilibrium effect: by lowering wages, labour market power prompts the entry of less talented entrepreneurs, who manage small businesses lacking the necessary scale for technological upgrades.

To test the model, I measure firm-level wage markdowns through revenue function estimation and classify Italian local labour markets by their degree of competition. The empirical evidence aligns with the theoretical predictions: areas characterised by higher markdowns exhibit lower wages, higher rates of entrepreneurship, a lower average firm

size, reduced use of intangibles and, crucially, inferior levels of aggregate productivity. Calibrating the model to match data on product and labour market power, as well as software adoption rates and expenditure, I quantify a substantial effect of labour market power: aggregate productivity in the typical labour market is 21% lower than under a competitive benchmark. Considering aggregate losses in both productivity and employment, output falls by 34%.

These findings can also offer insights on the roots of persistent disparities across regions. Using my measure of competition, I estimate that the higher levels of monopsony power in Southern Italy account for 40% of the productivity gap with the North. These quantitative results highlight that labour market power acts as a significant barrier to productivity, limiting the ability of the most performing firms to scale up and fully exploit the opportunities presented by new technologies.

References

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

Armangué-Jubert, T., Guner, N., & Ruggieri, A. (2024). Labor market power and development. *American Economic Review: Insights*, forthcoming.

Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, *135*(2), 645-709.

Azar, J., Marinescu, I., Steinbaum, M., & Taska, B. (2020). Concentration in US labor markets: Evidence from online vacancy data. *Labour Economics*, *66*, 101886.

Azar, J., Marinescu, I., & Steinbaum, M. (2022). Labor market concentration. *Journal of Human Resources*, *57*(S), S167-S199.

Bachmann, R., Bayer, C., Stüber, H., & Wellschmied, F. (2022). Monopsony Makes Firms Not Only Small but Also Unproductive: Why East-Germany Has Not Converged. *CESifo Working Paper*, No. 9751.

Berger, D., Herkenhoff, K., & Mongey, S. (2022). Labor market power. *American Economic Review*, *112*(4), 1147-1193.

Bento, P., & Restuccia, D. (2017). Misallocation, establishment size, and productivity. *American Economic Journal: Macroeconomics*, *9*(3), 267-303.

Bernard, A. B., Eaton, J., Jensen, J. B., & Kortum, S. (2003). Plants and productivity in international trade. *American Economic Review*, *93*(4), 1268-1290.

Bighelli, T., Di Mauro, F., Melitz, M. J., & Mertens, M. (2023). European firm concentration and aggregate productivity. *Journal of the European Economic Association*, *21*(2), 455-483.

Bloom, N., & Van Reenen, J. (2007). Measuring and explaining management practices across firms and countries. *The Quarterly Journal of Economics*, *122*(4), 1351-1408.

Bloom, N., & Van Reenen, J. (2010). Why do management practices differ across firms and countries?. *Journal of Economic Perspectives*, *24*(1), 203-224.

Bloom, N., Sadun, R., & Van Reenen, J. (2012). Americans do IT better: US multinationals and the productivity miracle. *American Economic Review*, *102*(1), 167-201.

Bloom, N., Brynjolfsson, E., Foster, L., Jarmin, R., Patnaik, M., Saporta-Eksten, I., & Van Reenen, J. (2019). What drives differences in management practices?. *American Economic Review*, *109*(5), 1648-1683.

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

Bontadini, F., Corrado, C., Haskel, J., Iommi, M., & Jona-Lasinio, C. (2023). EUKLEMS & INTANProd: industry productivity accounts with intangibles.

Bustos, P. (2011). Trade liberalization, exports, and technology upgrading: Evidence on the impact of MERCOSUR on Argentinian firms. *American Economic Review*, *101*(1), 304-340.

Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, *36*(S1), S13-S70.

Cattaneo, M. D., Crump, R. K., Farrell, M. H., & Feng, Y. (2024). On binscatter. *American Economic Review*, *114*(5), 1488-1514.

Clarke, R., Davies, S., & Waterson, M. (1984). The profitability-concentration relation: market power or efficiency?. *The Journal of Industrial Economics*, *32*(4), 435-450.

Datta, N. (2024). Local Monopsony Power. *CEP Discussion Paper*, No 2012.

De Ridder, M, B Grassi and G Morzenti (2022), The Hitchhiker's Guide to Markup Estimation: Assessing Estimates from Financial Data, *CEPR Discussion Paper*, No. 17532.

De Ridder, M. (2024). Market power and innovation in the intangible economy. *American Economic Review*, *114*(1), 199-251.

Dobbelaere, S., & Mairesse, J. (2013). Panel data estimates of the production function and product and labor market imperfections. *Journal of Applied Econometrics*, *28*(1), 1-46.

Draghi, M. (2024). The Future of European Competitiveness Part A: A competitiveness strategy for Europe.

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

Foster, L., Grim, C., Haltiwanger, J., & Wolf, Z. (2016). Firm-level dispersion in productivity: is the devil in the details?. *American Economic Review*, *106*(5), 95-98.

Gali, J., Gertler, M., & Lopez-Salido, J. D. (2007). Markups, gaps, and the welfare costs of business fluctuations. *The Review of Economics and Statistics*, *89*(1), 44-59.

Garicano, L., Lelarge, C., & Van Reenen, J. (2016). Firm size distortions and the productivity distribution: Evidence from France. *American Economic Review*, *106*(11), 3439-3479.

Ghatak, M., Morelli, M., & Sjöström, T. (2007). Entrepreneurial talent, occupational choice, and trickle up policies. *Journal of Economic Theory*, *137*(1), 27-48.

Ghironi, F., & Melits, M. J. (2005). International trade and macroeconomic dynamics with heterogeneous firms. *The Quarterly Journal of Economics*, *120*(3), 865-915.

Gourio, F., & Roys, N. (2014). Size-dependent regulations, firm size distribution, and reallocation. *Quantitative Economics*, *5*(2), 377-416.

Guner, N., Ventura, G., & Xu, Y. (2008). Macroeconomic implications of size-dependent policies. *Review of Economic Dynamics*, *11*(4), 721-744.

Guner, N., Parkhomenko, A., & Ventura, G. (2018). Managers and productivity differences. *Review of Economic Dynamics*, *29*, 256-282.

Hashemi, A., Kirov, I., & Traina, J. (2022). The production approach to markup estimation often measures input distortions. *Economics Letters*, *217*, 110673.

Hirsch, B., Jahn, E. J., Manning, A., & Oberfichtner, M. (2022). The urban wage premium in imperfect labor markets. *Journal of Human Resources*, *57*(S), S111-S136.

Hsieh, C. T., & Klenow, P. J. (2009). Misallocation and manufacturing TFP in China and India. *The Quarterly Journal of Economics*, *124*(4), 1403-1448.

Hsieh, C. T., & Rossi-Hansberg, E. (2023). The industrial revolution in services. *Journal of Political Economy*, *1*(1), 3-42.

Jha, P., & Rodriguez-Lopez, A (2023). Minimum Wage and Employer Variety. *CESifo Working Paper*, No. 9312.

Karabarbounis, L. (2014). The labor wedge: MRS vs. MPN. *Review of Economic Dynamics*, *17*(2), 206-223.

Lamadon, T., Mogstad, M., & Setzler, B. (2022). Imperfect competition, compensating differentials, and rent sharing in the US labor market. *American Economic Review*, *112*(1), 169-212.

Lucas Jr, R. E. (1978). On the size distribution of business firms. *The Bell Journal of Economics*, 508-523.

Mankiw, N. G., & Whinston, M. D. (1986). Free entry and social inefficiency. *The RAND Journal of Economics*, 48-58.

Manning, A. (2011). Imperfect competition in the labor market. *Handbook of Labor Economics* (Vol. 4, pp. 973-1041). Elsevier.

Melitz, M. J. (2003). The impact of trade on intra-industry reallocations and aggregate industry productivity. *Econometrica*, *71*(6), 1695-1725.

Melitz, M. J., Altomonte, C., Bartelsman, E., van de Kerke, J.-P., Lopez-Garcia, P., di Mauro, F., Polder, M., & Roux, S. (2018). Assessing the reliability of the CompNet micro-aggregated dataset for policy analysis and research: Coverage, representativeness and cross-EU comparability.

Mertens, M., Haug, P., Bighelli, T., Deist, J., & Elhefnawy, A. (2022). User Guide for the 9th Vintage of the CompNet Dataset.

Olley, G. S., & Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, *64*(6), 1263-1297.

O'Mahony, M., & Timmer, M. P. (2009). Output, input and productivity measures at the industry level: the EU KLEMS database. *The Economic Journal*, *119*(538), F374-F403.

Petrin, A., & Sivadasan, J. (2013). Estimating lost output from allocative inefficiency, with an application to Chile and firing costs. *Review of Economics and Statistics*, *95*(1), 286-301.

Reichardt, H. (2024). Scale-Biased Technical Change and Inequality.

Rubens, M. (2023). Market structure, oligopsony power, and productivity. *American Economic Review*, *113*(9), 2382-2410.

Schivardi, F., & Schmitz, T. (2020). The IT revolution and southern Europe's two lost decades. *Journal of the European Economic Association*, *18*(5), 2441-2486.

Shimer, R. (2009). Convergence in macroeconomics: The labor wedge. *American Economic Journal: Macroeconomics*, *1*(1), 280-297.

Sutton, J. (1991). Sunk costs and market structure: Price competition, advertising, and the evolution of concentration. *MIT press*.

Syverson, C. (2004). Product substitutability and productivity dispersion. *Review of Economics and Statistics*, *86*(2), 534-550.

Syverson, C. (2019). Macroeconomics and market power: Context, implications, and open questions. *Journal of Economic Perspectives*, *33*(3), 23-43.

Syverson, C. (2024). Markups and Markdowns. *NBER Working Papers*, (32871).

Yeh, C., Macaluso, C., & Hershbein, B. (2022). Monopsony in the US labor market. *American Economic Review*, *112*(7), 2099-2138.

Appendices

A Proofs

A.1 Closed-form solution

The first step to analytically solve the system described in Definition 1 is finding closed-form solution for the entry and technology adoption cutoffs.

Assuming a Pareto distribution with scale normalised to one, $\phi(\mu) = \frac{\alpha}{\mu^{\alpha+1}}$, and a constraint on the shape $\alpha > \frac{(\sigma-1)(1+\theta)}{\sigma\theta+1}$, I rewrite equation (30) substituting equations (24), (29), and (32).

$$L^P = 1 - \frac{b^\alpha + \tau}{(b\mu_E)^\alpha} \tag{A1}$$

Then, I substitute equation (A1) in equation (27). Solving the integral and rearranging, I obtain a closed-form solution for the entry cutoff and, using equation (32), the technology adoption cutoff.

$$\mu_E = \left[\frac{\alpha(\sigma\theta+\sigma) - (\sigma-1)(1+\theta)}{\alpha(\sigma\theta+1) - (\sigma-1)(1+\theta)}\left(\frac{\tau}{b^\alpha} + 1\right)\right]^{\frac{1}{\alpha}} \tag{A2}$$

$$\mu_T = \left[\frac{\alpha(\sigma\theta+\sigma) - (\sigma-1)(1+\theta)}{\alpha(\sigma\theta+1) - (\sigma-1)(1+\theta)}(\tau + b^\alpha)\right]^{\frac{1}{\alpha}} \tag{A3}$$

Combining equations (29) and (A2), I can solve for E. Combining equation (24) and (A3), I can solve for L^T, and thus for $L^P = 1 - E - L^T$.

Once I have identified the two cutoffs and the mass of the three sets of agents, I can solve for all the remaining equations, (16)-(21).

A.2 Proposition 1 (Reallocation)

Trivial, from $\frac{l_k}{l_j} = \left(\frac{a_k}{a_j}\right)^{\frac{\sigma-1}{\sigma\theta+1}}$.

A.3 Proposition 2 (Excess entry)

Labour market power lowers the mass of entrepreneurs, E, as long as the entry cutoff, μ_E, falls in θ. Considering equation (A2) and assuming $\sigma > 1$: (i) the first ratio falls in θ; and (ii) the variable b grows in θ because the exponent in its denominator falls in θ, whereas the exponent outside square brackets grows in θ with a base above one for any $\mu_E < \mu_T$,

which is the case considered in this paper (see Section 2.7). Note that this result does not require any assumption on technology, which can also be constant across firms (e.g. $\gamma = 0$).

To prove the negative effect of labour market power on the mass of production workers, L^P, I substitute equation (A2) in equation (A1). Rearranging, I obtain the following closed-form solution for L^P, which falls in θ and is independent from the technology's parameters.

$$L^P = \frac{\alpha(\sigma-1)}{(1+\theta)[\sigma(\alpha-1)+1]} \tag{A4}$$

As I proved the positive effect on E and the negative effect on L^P, the effect on average firm size $\bar{l} = \frac{L^P}{E}$ is trivial.

A.5 Corollary 1 (Inefficient allocation)

In equilibrium, labour and its marginal revenue product produce the same rank of firms, as both increase monotonically in productivity: $\frac{MRPL_k}{MRPL_j} = \left(\frac{a_k}{a_j}\right)^{\frac{(\sigma-1)\theta}{\sigma\theta+1}} = \left(\frac{l_k}{l_j}\right)^{\theta}$. This implies that firms with higher MRPL are larger. The rest follows from Proposition 1.

A.4 Corollary 2 (Inefficient selection)

Trivial, from Proposition 2 and the definition of CDF.

A.6 Proposition 3 (Inefficient technology diffusion)

Using the Pareto assumption and equations (15) and (32), I can solve for the share of production labour allocated to firms adopting the superior technology.

$$\frac{\int_{\mu_h}^{\infty} l(\mu)\phi(\mu)d\mu}{L^P} = 1 - \frac{1 - b^{-\frac{\alpha(\sigma\theta+1)-(\sigma-1)}{\sigma\theta+1}}}{1 + \left[(1+\gamma)^{\frac{\sigma-1}{\sigma\theta+1}} - 1\right] b^{-\frac{\alpha(\sigma\theta+1)-(\sigma-1)}{\sigma\theta+1}}} \tag{A5}$$

The term in square brackets obviously falls in θ. A sufficient condition for labour market power to reduce the diffusion of the productivity-enhancing technology across workers is that $b^{-\frac{\alpha(\sigma\theta+1)-(\sigma-1)}{\sigma\theta+1}}$ falls in θ as well. This condition is easily verified using the generic formula $\frac{\partial f(x)^{g(x)}}{\partial x} = f(x)^{g(x)}\left[g'(x)\ln f(x) + g(x)\frac{f'(x)}{f(x)}\right]$, as the base of $b^{-\frac{\alpha(\sigma\theta+1)-(\sigma-1)}{\sigma\theta+1}}$ is positive and grows in θ (see Section 2.7), whereas the exponent is negative and falls in θ.

B Strategic interactions

The model in Section 2 abstracts from the role of strategic interactions and variable supply elasticity: wage markdowns are constant along the size distribution and any change in monopsony power uniformly applies to all firms. A natural question is whether my theory on the effect of labour market power also holds in a different setting, in which small firms have little power and the monopsony shock does not equally affect all firms, but mostly the largest ones.

To address this concern, I explore a simple framework with such characteristics, largely based on the model presented in Section 2. Although I cannot fully solve this version in closed form, I can analytically prove that the core theory, expressed in Propositions 1 and 2, still holds: (i) an increase in monopsony power based on strategic interactions induces excess entry, with negative effects on firm size and managerial ability; and (ii) even if wage are equalised along the size distribution, greater monopsony power of large firms still reallocates resources towards less efficient firms, reducing the market's allocative efficiency.

To model a variable labour supply elasticity, I consider a continuum of sectors k with unit measure. A sector can be thought as a commuting zone or an occupation. Workers see employers as perfect substitutes within sector and imperfect substitutes between sectors. Perfect substitutability implies that, within sector, wages are equalised and all employers provide the same utility.

The worker's problem is similar to the problem in Section 2.3, but workers choose a preferred sector s^*. The taste shock, $\tilde{\varepsilon}_{is} > 0$, is distributed across sectors according to a Fréchet distribution with shape parameter $\frac{1}{\theta}$. As the shape parameter of the taste distribution is negatively related to its dispersion, θ can be regarded as an index of job differentiation between sectors, which leads to limited substitutability and imperfect competition.

As in Section 2, equating choice probabilities to employment shares generates the sector-level labour supply, with L defined as the aggregate mass of workers.

$$l_s = L \frac{w_s^{\frac{1}{\theta}}}{\int_{s \in S} w_s^{\frac{1}{\theta}} \, ds} \tag{B1}$$

I assume that firms internalise the effect of their individual behaviour on sector-level prices and quantities, whereas aggregate values are independent. The resulting inverse

labour supply observed by firm k in sector s has a variable slope, which increases in its employment share.

$$\frac{\partial w_s}{\partial l_{ks}}\frac{l_{ks}}{w_s} = \frac{\partial w_s}{\partial l_s}\frac{l_{ks}}{w_s} = \theta\frac{l_{ks}}{l_s} \tag{B2}$$

The largest firm size, the more a firm's demand of labour can affect sector-level prices and quantities. As substitutability between sectors is limited, the slope of the inverse labour supply grows in firm size. Intuitively, small firms mostly compete within sector, whereas large firms also compete with firms outside their sector, hence they face lower substitutability, which leads to rent-extraction.

To keep the model tractable, I here consider the case with one technology and assume no product differentiation: productivity equals managerial ability and firm-level prices are normalised to one. The first-order condition for profit maximization leads to a variable wage markdown, which ranges from one (for an infinitesimal firm) to $(1+\theta)$ (for a pure (monopsonist).

$$MRPL_{ks} = \mu_{ks} = \left(1 + \theta\frac{l_{ks}}{l_s}\right) w_s \tag{B3}$$

Whilst firms within sectors have heterogenous productivity, for analytical convenience I assume that each sector has an identical ability distribution, which leads to symmetrical equilibria. This implies that, in equilibrium, $w_s = w$ and $l_s = l$ for each sector. As the continuum of sectors has unit measure, equilibrium in each sector mirrors the aggregate equilibrium and the sector-level labour equals the aggregate mass of workers, i.e. $l = L$. With this simplifying assumption, I rewrite firm's employment, production, and profits in terms of their productivity.

$$l_{ks} = \frac{\frac{\mu_{ks}}{w} - 1}{\theta} L \tag{B4}$$

$$y_{ks} = \mu_{ks}\frac{\frac{\mu_{ks}}{w} - 1}{\theta} L \tag{B5}$$

$$\pi_{ks} = y_{ks} - w l_{ks} \ = \frac{(\mu_{ks} - w)^2}{\theta w} L \tag{B6}$$

As in Section 2, I use the entry condition, $\pi(\mu_E) = \overline{w}$, to identify the entry cutoff. This allows characterising the set of entrepreneurs, $\mathcal{E} = \{i|_{\mu_i > \mu_E}\}$, with mass $E = 1 - \Phi(\mu_E)$, as well as its residual, the set of workers $\mathcal{L}$, with mass $L = \Phi(\mu_E)$.

$$\mu_E = w\left[\left(\frac{\theta}{L}\right)^{0.5} + 1\right] \tag{B7}$$

Aggregating the employment shares from equation (B4) and substituting the Pareto distribution, I can express the market wage in terms of the entry cutoff.

$$w = \frac{\int_{\mu_E}^{\infty} \mu\phi(\mu)d\mu}{\theta + 1 - \Phi(\mu_E)} = \frac{\alpha}{\alpha - 1}\frac{\mu_E}{\theta\mu_E^{\alpha} + 1} \tag{B8}$$

Combining (B7) and (B8), and rearranging, I rewrite the cutoff as follows.

$$\mu_E = \left[\frac{1}{\alpha - 1}\left(\frac{\alpha}{\sqrt{\theta L}} + \frac{1}{\theta}\right)\right]^{\frac{1}{\alpha}} \tag{B9}$$

From equation (B9), I can prove that the entry cutoff falls in θ, which is a sufficient condition for excess entry.

PROPOSITION B1 (Excess entry): $\frac{\partial \mu_E}{\partial \theta} < 0$

Proof. Suppose, to the contrary, that μ_E increases in θ. For higher θ, both sides of equation (B9) have to rise, which, for a constant α, also implies lower L. This leads to a contradiction, because $L = \Phi(\mu_E)$ necessarily grows in μ_E. Hence, μ_E falls in θ, QED.

Intuitively, even with variable markdowns, an increase of labour market power decreases the aggregate labour demand and thus market wages, inducing excess entry and, consequently, lower average firm size and poorer managerial practices in the economy.

From equation (B4), I can take the ratio of employment in two firms, k and j, with $\mu_{ks} > \mu_{js}$, and prove that their employment ratio (l_k/l_j) falls in θ (Misallocation).

$$\frac{l_{ks}}{l_{js}} = 1 + \frac{\mu_{ks} - \mu_{js}}{\mu_{js} - w} \tag{B10}$$

PROPOSITION B2 (Reallocation): $l_{ks} > l_{js} \;\Rightarrow \frac{\partial(l_{ks}/l_{js})}{\partial \theta} < 0$

Proof. Equation (B10) shows that a sufficient and necessary condition for the gap to fall in θ is that wages fall in θ. Consider equation (B7) and suppose that θ increases. As I prove above that μ_E falls in θ, both sides of the equation need to fall. As $(\frac{\theta}{L})$ necessarily increases, w needs to fall, QED.

The employment gaps of more productive firms shrink with monopsony power, inducing an inefficient reallocation of employment shares towards smaller, less efficient firms. Whilst in Section 2 the reallocation towards small firms comes from an upward-sloping inverse labour supply, here the same effect depends on the positive markdown-size relationship.

C. Labour market power and the adoption cutoff

Equation (40) shows that labour market power prevents the diffusion of scale-biased technologies across firms. However, a negative effect on the share of adopting firms does not imply that the *mass* of adopting firms falls too, as labour market power simultaneously increases the mass of active firm. In fact, the firm-level relationship between monopsony power and technology adoption is potentially ambiguous. One can identify four separate channels: on the one hand, (i) monopsony reduces firm size, which is a critical driver of adoption for scale-biased technologies, and (ii) it involves an upward-sloping inverse labour supply, which kills parts of the benefits from scaling up through technology adoption; on the other hand, (iii) it reduces the cost of adoption by lowering the wage of tech workers, and (iv) it increases profit margins, making scaling up potentially more appealing.

Eventually, whether the former or latter effects dominate depends on the level of scale bias. If the scale bias is high, monopsony power unequivocally reduces the mass of firms adopting the superior technology. More formally, there exists a $b^*(\tau, \gamma)$ such that labour market power reduces the mass of firms adopting technologies characterised by $b > b^*$, and vice versa.

(C1) $$b \gtrless b^* \Rightarrow \frac{\partial \mu_h}{\partial \theta} \gtrless 0$$

As discussed in Section 2.7, b itself grows in θ, which implies that labour market power expands the set of technologies that satisfies equation (41).

The ambiguity on the effect of monopsony power on firm-level technology adoption does not imply that, on aggregate, labour market power can increase the aggregate benefits of technologies with lower levels of scale bias. In fact, what matters for productivity and output is the share and mass of *workers* using the superior technology, which fall unambiguously.

By reducing average firm size and reallocating resources to smaller firms, monopsony reduces the diffusion of productivity-enhancing technologies across workers, regardless of

the degree of scale bias and the mass of adopting firms. For a technology with a scale bias $b < b^*$, labour market power may increase the mass of adopting firms, but the aggregate benefits from technology use would still fall, because monopsony substantially reduces the market shares of adopting firms. If, instead, the technology introduced has a scale bias $b > b^*$, the negative effect on technology usage appears not only at the market-level, but at the firm-level too.

D. Markdown estimation: identification threats

D.1 Adjustments costs

Yeh et al (2022, Appendix O.4) shows that adjustment costs can lead to mismeasurement of markdown estimators, although the magnitude of the bias is found to be negligible. In this Appendix, I consider whether the bias affects markdowns in a way that could generate the empirical patterns shown in Section 3.4.

I follow Yeh et al (2022) closely in building a simple setup of dynamic optimization, with $\Phi(l, l_{-1})$ denoting convex adjustment cost. I omit subscripts for visual clarity: variables represent a given firm in a given period.

$$v = max\{r - wl - w\Phi(l, l_{-1}) + \beta\mathbb{E}[v']\} \tag{D1}$$

Keeping the notation from the main text, in this case the ratio of MRPL to the wage does not capture only the inverse labour supply elasticity, θ, but also the adjustment costs.[32]

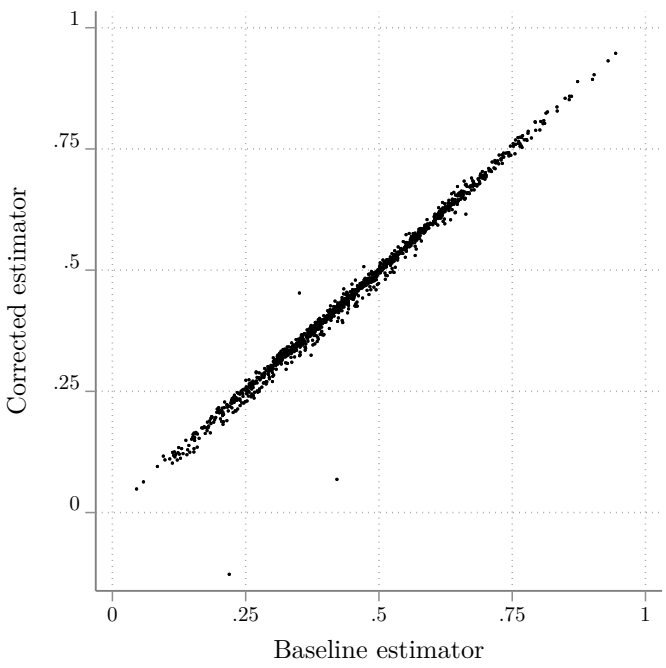
$$\begin{aligned} \frac{MRPL}{w} &= (1+\theta) + \frac{\partial\Phi(l, l_{-1})}{\partial l} + \frac{\partial w}{\partial l}\frac{\Phi(l, l_{-1})}{w} - \beta\mathbb{E}\left[\frac{1}{w}\frac{\partial v'}{\partial l}\right] \\ &= (1+\theta) + \frac{\partial\Phi(l, l_{-1})}{\partial l} + \theta\frac{\Phi(l, l_{-1})}{l} + \beta\mathbb{E}\left[\frac{w'}{w}\frac{\partial^2\Phi(l, l_{-1})}{\partial l^2}\right] \end{aligned} \tag{D2}$$

As in Hall (2004) and Yeh et al. (2022), I consider a quadratic form for the adjustment costs, $\Phi(l, l_{-1}) = \frac{\psi}{2} l \left[\left(\frac{l'-1}{l}\right)^2\right]$, and denote the labour growth rate as $g_l = \frac{l-l_{-1}}{l_{-1}}$.

$$\frac{MRPL}{w} = (1+\theta)\left(1 + \frac{\psi}{2}g_l^2\right) + \psi g_l(g_l+1) - \beta\psi\mathbb{E}[g_l'(1+g_l')(1+g_w')] \tag{D3}$$

[32]The second equality in equation (D2) follows from the envelope theorem.

FIGURE D1. THE ADJUSTMENT COSTS BIAS



Note: Each dot represents a province-year observation in 2006-18 (1,226 observations). Data in 2019 miss the future realised labour growth. For the baseline estimator, see equations (43) and (47). The corrected estimator subtracts the bias generated by adjustment costs, $\psi\left(\overline{g_l} - \beta\overline{\mathbb{E}[g_l']}\right)$. Microdata on Italian firms from CompNet.

In order to isolate dominant contributions, I ignore the squares of growth rates. The MRPL to wage ratio includes, besides a monopsony power term, the difference between the current labour growth and the discounted future labour growth.

$$\frac{MRPL}{w} \approx (1+\theta) + \psi(g_l - \beta\mathbb{E}[g_l']) \tag{D4}$$

Taking the average at the market-level and subtracting one, I approximate my index of labour market power, $\hat{\theta}$, with overlines representing averages of firm-level variables.

$$\hat{\theta} - \theta \approx \psi\left(\overline{g_l} - \beta\overline{\mathbb{E}[g_l']}\right) \tag{D5}$$

I assume that firms do not make systematic forecast errors so market-level averages of the expected growth rates can be proxied by the averages of realised growth rates. Using

these values and setting $\psi = 0.187$ from Hall (2004) and $\beta = 0.96$ (standard value for annual discounting), I measure the deviation and correct my estimator.

This bias has an interquartile range of $(0.003, 0.005)$, which is two orders of magnitude smaller than the interquartile range of $\hat{\theta}$. Figure D1 shows a scatter plot of $\hat{\theta}$ against the corrected estimator, plotting all observations in my database. Correcting the bias does not induce meaningful variations in my measure of labour market power across provinces. Finally, repeating the empirical analysis of Section 3 with the corrected estimator does not generate visible changes in the results.

D.2 Worker heterogeneity

Whilst workers are assumed to be homogenous in Section 3, in this Appendix I show that my estimator is consistent with worker heterogeneity. I consider two worker types, denoted by L and H (low and high skills), subject to heterogeneous levels of monopsony power (i.e. their labour supply elasticities differ: $\varepsilon_k^{w,l^L} \neq \varepsilon_k^{w,l^H}$) and combined in the production process with firm-specific fixed proportions ($l_k^L = \xi_k l_k^H$): firm-level markdowns recover the weighted average of the worker-specific wage markdown.

Firms maximise variable profits choosing the total amount of labour, $l_k \equiv l_k^L + l_k^H$.

$$Max_{l_k}\ r_k - w_k^L l_k^L - w_k^H l_k^H \tag{D6}$$

Considering the relationship between total and type-specific labour, $l_k = (1+\xi_k) l_k^H = \frac{1+\xi_k}{\xi_k} l_k^L$, I derive the following first-order condition.

$$\varepsilon_k^{r,l} \frac{r_k}{l_k} = \frac{\xi_k w_k^L}{1+\xi_k}\left(1+\varepsilon_k^{w,l^L}\right) + \frac{w_k^H}{1+\xi_k}\left(1+\varepsilon_k^{w,l^H}\right) \tag{D7}$$

If markdowns are measured pooling together the labour expenditure across all types of workers, the resulting estimator is a weighted average of type-specific labour supply elasticities, with weights defined as shares of labour costs.

$$\begin{aligned} 1+\hat{\theta}_k &= \varepsilon_k^{r,l} \frac{r_k}{w_k^L l_k^L + w_k^H l_k^H} \\ &= \frac{w_k^L l_k^L}{w_k^L l_k^L + w_k^H l_k^H}\left(1+\varepsilon_k^{w,l^L}\right) + \frac{w_k^H l_k^H}{w_k^L l_k^L + w_k^H l_k^H}\left(1+\varepsilon_k^{w,l^H}\right) \end{aligned} \tag{D8}$$

D.3 Overhead labour

The markdown estimator, based on equations (43) and (47) in the main text, requires *variable* labour costs. If the estimator is measured including overhead labour costs, o_k, markdowns are underestimated.

$$\varepsilon_k^{R,L} \frac{r_k}{w_k l_k + o_k} < \varepsilon_k^{R,L} \frac{r_k}{w_k l_k} = 1 + \varepsilon_k^{w,l} \tag{D9}$$

The evolution of the bias along the size distribution depends on the share of overhead in total labour. Abraham et al. (2024) shows that although fixed costs are higher in large firms their share is substantially lower. This insight is consistent with my model considering as overhead the entrepreneur's imputed labour income proxied by the average wage (i.e. the opportunity cost of entrepreneurship).

Therefore, the index of labour market power should have a more pronounced underestimation in provinces with smaller firms. Considering the robust evidence of a negative relationship between my index and firm size, the bias arising from overhead labour should only attenuate differences in labour market power between provinces, without affecting the rank and, most importantly, with no ability to generate the observed results. If the bias induced by mismeasuring variable labour costs played a major role, correlations would go in the opposite direction.

E. Revenue function estimation

Assumptions. This Appendix illustrates the estimation of revenue functions discussed in Section 3.2. It replicates the methodology of Ackerberg et al. (2015), with revenues instead of production as the dependent variable. This difference implies that the object of estimation is revenue elasticities, which depend not only on the output elasticities of the underlying production function, but also on the elasticity of the product demand. Additionally, the persistent revenue shock, ω_{kt}, differs from firm productivity as it also includes a demand component.

This control function approach to revenue function estimation requires the following five assumptions.

ASSUMPTION 1 (Information Set): *The firm's current information set, I_{kt}, includes the persistent revenue shock, ω_{kt}, but does not include the transitory shock, ϵ_{kt}.*

ASSUMPTION 2 (First Order Markov): *The persistent revenue shock evolves according to a first-order Markov process:* $p(\omega_{kt+1}|I_{kt}) = p(\omega_{kt+1}|\omega_{kt})$.

ASSUMPTION 3 (Timing of Input Choices): *Capital is determined by investments in past periods,* $\kappa_{kt} = \kappa(\kappa_{kt-1}, i_{kt-1})$. *Labour can be chosen before or after observing* ω_{kt+1}, *before or simultaneously with the choice of materials.*

ASSUMPTION 4 (Scalar Unobservable): *The choice of materials is determined by the level of other inputs and, crucially, by the persistent revenue shock ('unobservable' to the econometrician),* ω_{kt}*:* $m_{kt} = m(\kappa_{kt}, l_{kt}, \omega_{kt})$.

ASSUMPTION 5 (Strict Monotonicity): *The relationship between materials and the persistent revenue shock is strictly monotonic:* $\frac{\partial m(\kappa_{kt}, l_{kt}, \omega_{kt})}{\partial \omega_{kt}} > 0$.

Strategy. Whilst in practice my estimation employs the translog specification shown in equation (45), to simplify the notation, in this Appendix I consider the case of a Cobb-Douglas specification (i.e. $\beta^{\kappa^2} = \beta^{l^2} = \beta^{m^2} = \beta^{\kappa l} = \beta^{\kappa m} = \beta^{lm} = 0$). For the same reason, I approximate the Markov process using an AR(1) specification, whereas my estimation includes a quadratic term.

$$r_{kt}^{ln} = \beta^{\kappa}\kappa_{kt}^{ln} + \beta^{l} l_{kt}^{ln} + \beta^{m} m_{kt}^{ln} + \omega_{kt} + \epsilon_{kt} \tag{E1}$$

$$\omega_{kt} = \rho\omega_{kt-1} + \eta_{kt} \tag{E2}$$

Consider equations (E1) and (E2). Because of persistency in the observable revenue shock, lags of inputs are endogenous and cannot be used as valid instruments. The solution proposed by the control function approach is based on controlling past shocks using past revenues and past input usage. However, as revenues also include a transitory shock, this methodology is implemented in two steps: the first one purges ϵ_{kt}, the second one identifies the parameters of interest.

By using the inverted material equation from Assumption 4 as a 'control function' for the persistent shock, $\omega_{kt} = m^{-1}(k_{kt}, l_{kt}, \omega_{kt})$, it is possible to express the difference between revenues and the transitory shock as a function of current inputs, $\Phi_{kt} = \Phi(\kappa_{kt}, l_{kt}, m_{kt})$.

$$r_{kt}^{ln} = \beta^{\kappa}\kappa_{kt}^{ln} + \beta^{l} l_{kt}^{ln} + \beta^{m} m_{kt}^{ln} + m^{-1}(k_{kt}, l_{kt}, \omega_{kt}) + \epsilon_{kt} = \Phi_{kt} + \epsilon_{kt} \tag{E3}$$

The first step in this methodology consists of estimating Φ_{kt}, by projecting revenues on a third-order polynomial in current inputs. In the second step, parameters are estimated controlling the lag of the persistent shock with the difference between Φ_{kt} and a combination of input lags.

$$
\begin{aligned}
& r_{kt}^{ln} \\
& = \beta^{\kappa}\kappa_{kt}^{ln} + \beta^{l}l_{kt}^{ln} + \beta^{m}m_{kt}^{ln} + \rho(\Phi_{kt-1} - \beta^{\kappa}\kappa_{kt-1}^{ln} - \beta^{l}l_{kt-1}^{ln} - \beta^{m}m_{kt-1}^{ln}) \\
& + \eta_{kt} + \epsilon_{kt}
\end{aligned} \tag{E4}
$$

Under this specification, current values of (pre-determined) capital and lags of labour and materials are valid instruments, and can identify the three revenue elasticities. The estimation requires a fourth parameter, the autoregressive coefficient, which is pinned down by the lag of Φ_{kt}, estimated in the first step. The second step of the estimation allows the identification of all parameters using the following moment condition.

$$
\begin{aligned}
& \mathbb{E}\Bigg[\Big(r_{kt}^{ln} - \beta^{\kappa}\kappa_{kt}^{ln} - \beta^{l}l_{jt}^{ln} - \beta^{m}m_{jt-1}^{ln} \\
& - \rho\big(\Phi_{jt-1} - \beta^{\kappa}\kappa_{jt-1}^{ln} - \beta^{l}l_{jt-1}^{ln} - \beta^{m}m_{jt-1}^{ln}\big)\Big) \otimes \begin{pmatrix} \kappa_{kt}^{ln} \\ l_{kt-1}^{ln} \\ m_{kt-1}^{ln} \\ \Phi_{kt-1} \end{pmatrix}\Bigg] = 0
\end{aligned} \tag{E5}
$$

Discussion. The methodology illustrated in this Appendix considers a Cobb-Douglas revenue function (translog in the main text and in my actual estimation), without making explicit assumptions on the functional forms of the underlying production function and product demand. However, it may be difficult to assess the reasonableness of Assumptions 1-5 without clarifying the nature of the persistent shock, which is based on such functional forms. Therefore, I here assume a simple system with a CES demand augmented by a time-variant firm-specific shock, d_{kt}.

$$
p_{kt} = d_{kt}y_{kt}^{-\frac{1}{\sigma}} \tag{E6}
$$

To derive the persistent revenue shock in this framework, consider that, with a CES demand, revenue elasticities equal output elasticities multiplied by $\frac{\sigma-1}{\sigma}$, the inverse of the markup. Therefore, production and revenue functions have the same functional form, whether Cobb Douglas or translog, with each parameter in the revenue function corresponding to the same parameter in production function divided by the markup: $\beta^{x} =$

$\frac{\sigma-1}{\sigma}\tilde{\beta}^x$, with the tilde denoting output elasticities and x denoting any combination of inputs included in the production function ($x = \{\kappa, l, m\}$ in the Cobb Douglas case).

Under this demand assumption, the persistent revenue shock has the following form, depending on both the demand shock and firm productivity. [33]

$$\omega_{kt} = d_{kt}^{ln} + \frac{\sigma - 1}{\sigma} a_{kt}^{ln} \tag{E7}$$

In the standard control function approach to production function estimation, Assumptions 1-5 refer to a persistent shock in production (i.e. firm productivity) instead of a persistent shock in revenues. Assumptions 1-3 do not appear to be directly affected by the difference in my methodology, which includes demand shocks in the Markov process. Assumption 4, in my case, allows materials to also respond to demand shocks and not just to productivity shocks. As for Assumption 5, the additional requirement is that materials strictly increase not only in productivity, but also the in the demand shock. Neither of these two requirements seems particularly problematic.

F. Alternative estimators of labour market power

In the baseline estimation, I estimate the province-level index of labour market power as the average of firm-level wage markdowns, derived from the estimation of translog revenue functions (see equations (43), (46) and (47)).

This Appendix shows the results of the same set of regressions reported in Tables 4-6, but implemented using four alternative measurement choices:

i. Markdowns aggregated using medians instead of arithmetic averages;
ii. Markdowns aggregated using 95$^{\text{th}}$ percentiles instead of arithmetic averages;
iii. Revenue elasticities estimated using a Cobb-Douglas specification;
iv. Markdowns estimated using a 'production approach' (see equation (44)).

[33] The result follows from $\omega_{kt} = p_{kt}^{ln} + y_{kt}^{ln} - \sum_x \beta^x x_{kt}^{ln} = d_{kt}^{ln} + \frac{\sigma-1}{\sigma}\left(y_{kt}^{ln} - \sum_x \tilde{\beta}^x x_{kt}^{ln}\right)$. I omit transitory shocks although they can be easily included as a component of production, prices, or a combination of both, without affecting the derivation of equation (E7).

F.1 Median markdowns (translog revenue functions)

TABLE F1. TABLE 4 REPLICATED WITH MEDIAN MARKDOWNS

	(1) Wage dispersion (logs)	(2) Firm size dispersion (logs)	(3) MRPL dispersion (logs)	(4) TFPR dispersion (logs)
LMP (median)	0.237	-0.339**	0.0456	-0.0978
	(0.134)	(0.141)	(0.0605)	(0.0770)
Observations	1,436	1,436	1,436	1,436
R-squared	0.686	0.302	0.322	0.261

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 4 replicated with a different estimator of labour market power (LMP): median instead of average markdown. Same controls (South, city, year) as in Table 4.

TABLE F2. TABLE 5 REPLICATED WITH MEDIAN MARKDOWNS

	(5) Average wage (logs)	(6) Entrepreneurship rate (logs)	(7) Average firm size (logs)	(8) Average labour productivity (logs)
LMP (median)	-0.913***	0.507**	-0.788**	-0.758***
	(0.103)	(0.194)	(0.317)	(0.179)
Observations	1,436	1,370	1,436	1,436
R-squared	0.723	0.240	0.466	0.752

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 5 replicated with a different estimator of labour market power (LMP): median instead of average markdown. Same controls (South, city, year) as in Table 5.

TABLE F3. TABLE 6 REPLICATED WITH MEDIAN MARKDOWNS

	(9) Average intangible assets (logs)	(10) Intangibles per worker (logs)	(11) Aggregate labour productivity (logs)	(12) Aggregate TFP (logs)
LMP (median)	-0.218*	-0.698*	-0.799***	-0.690***
	(0.119)	(0.392)	(0.141)	(0.208)
Observations	1,376	1,436	1,436	1,433
R-squared	0.601	0.086	0.645	0.605

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 6 replicated with a different estimator of labour market power (LMP): median instead of average markdown. Same controls (South, city, year) as in Table 6.

F.2 95th percentile of markdowns (translog revenue functions).

TABLE F4. TABLE 4 REPLICATED WITH 95TH PERCENTILES OF MARKDOWNS

	(1) Wage dispersion (logs)	(2) Firm size dispersion (logs)	(3) MRPL dispersion (logs)	(4) TFPR dispersion (logs)
LMP (p95)	0.0959**	-0.0594	0.0363***	0.0458***
	(0.0386)	(0.0360)	(0.00865)	(0.0111)
Observations	1,436	1,436	1,436	1,436
R-squared	0.812	0.313	0.420	0.367

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 4 replicated with a different estimator of labour market power (LMP): 95th percentile instead of average markdown. Same controls (South, city, year) as in Table 4.

TABLE F5. TABLE 5 REPLICATED WITH 95TH PERCENTILES OF MARKDOWNS

	(5) Average wage (logs)	(6) Entrepreneurship rate (logs)	(7) Average firm size (logs)	(8) Average labour productivity (logs)
LMP (p95)	-0.110***	0.0818**	-0.246***	-0.108**
	(0.0188)	(0.0307)	(0.0459)	(0.0365)
Observations	1,436	1,370	1,436	1,436
R-squared	0.705	0.272	0.563	0.761

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 5 replicated with a different estimator of labour market power (LMP): 95th percentile instead of average markdown. Same controls (South, city, year) as in Table 5.

TABLE F6. TABLE 6 REPLICATED WITH 95TH PERCENTILES OF MARKDOWNS

	(9) Average intangible assets (logs)	(10) Intangibles per worker (logs)	(11) Aggregate labour productivity (logs)	(12) Aggregate TFP (logs)
LMP (p95)	-0.0571***	-0.0129	-0.116***	-0.0483
	(0.0158)	(0.0471)	(0.0231)	(0.0486)
Observations	1,376	1,436	1,436	1,433
R-squared	0.624	0.072	0.655	0.582

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 6 replicated with a different estimator of labour market power (LMP): 95th percentile instead of average markdown. Same controls (South, city, year) as in Table 6.

F.3 Average markdowns (Cobb-Douglas revenue functions).

TABLE F7. TABLE 4 REPLICATED WITH COBB-DOUGLAS REVENUE FUNCTIONS

	(1) Wage dispersion (logs)	(2) Firm size dispersion (logs)	(3) MRPL dispersion (logs)	(4) TFPR dispersion (logs)
LMP (Cobb-Douglas)	0.333*	-0.263*	0.0802	0.0716
	(0.167)	(0.145)	(0.0479)	(0.0630)
Observations	1,436	1,436	1,436	1,436
R-squared	0.754	0.312	0.343	0.265

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 4 replicated with a different estimator of labour market power (LMP): revenue elasticities based on a Cobb-Douglas specification instead of translog. Same controls (South, city, year) as in Table 4.

TABLE F8. TABLE 5 REPLICATED WITH COBB-DOUGLAS REVENUE FUNCTIONS

	(5) Average wage (logs)	(6) Entrepreneurship rate (logs)	(7) Average firm size (logs)	(8) Average labour productivity (logs)
LMP (Cobb-Douglas)	-0.563***	0.438**	-0.907***	-0.484**
	(0.0849)	(0.164)	(0.196)	(0.170)
Observations	1,436	1,370	1,436	1,436
R-squared	0.726	0.297	0.523	0.759

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 5 replicated with a different estimator of labour market power (LMP): revenue elasticities based on a Cobb-Douglas specification instead of translog. Same controls (South, city, year) as in Table 5.

TABLE F9. TABLE 6 REPLICATED WITH COBB-DOUGLAS REVENUE FUNCTIONS

	(9) Average intangible assets (logs)	(10) Intangibles per worker (logs)	(11) Aggregate labour productivity (logs)	(12) Aggregate TFP (logs)
LMP (Cobb-Douglas)	-0.234***	-0.336	-0.530***	-0.342
	(0.0770)	(0.229)	(0.116)	(0.202)
Observations	1,376	1,436	1,436	1,433
R-squared	0.618	0.081	0.656	0.595

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 6 replicated with a different estimator of labour market power (LMP): revenue elasticities based on a Cobb-Douglas specification instead of translog. Same controls (South, city, year) as in Table 6.

F.4 Average markdowns (translog production functions).

TABLE F10. TABLE 4 REPLICATED WITH A PRODUCTION APPROACH

	(1) Wage dispersion (logs)	(2) Firm size dispersion (logs)	(3) MRPL dispersion (logs)	(4) TFPR dispersion (logs)
LMP (production)	0.370*	-0.261	0.0990*	0.0950
	(0.195)	(0.171)	(0.0516)	(0.0723)
Observations	1,436	1,436	1,436	1,436
R-squared	0.758	0.308	0.350	0.272

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 4 replicated with a different estimator of labour market power (LMP): production approach instead of revenue approach (see Section 3.2). Same controls (South, city, year) as in Table 4.

TABLE F11. TABLE 5 REPLICATED WITH A PRODUCTION APPROACH

	(5) Average wage (logs)	(6) Entrepreneurship rate (logs)	(7) Average firm size (logs)	(8) Average labour productivity (logs)
LMP (production)	-0.591***	0.471**	-1.020***	-0.487**
	(0.0992)	(0.195)	(0.231)	(0.196)
Observations	1,436	1,370	1,436	1,436
R-squared	0.720	0.296	0.529	0.749

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 5 replicated with a different estimator of labour market power (LMP): production approach instead of revenue approach (see Section 3.2). Same controls (South, city, year) as in Table 5.

TABLE F12. TABLE 6 REPLICATED WITH A PRODUCTION APPROACH

	(9) Average intangible assets (logs)	(10) Intangibles per worker (logs)	(11) Aggregate labour productivity (logs)	(12) Aggregate TFP (logs)
LMP (production)	-0.241**	-0.417	-0.569***	-0.329
	(0.0825)	(0.244)	(0.120)	(0.227)
Observations	1,376	1,436	1,436	1,433
R-squared	0.616	0.084	0.655	0.590

Note: *** p<0.01, ** p<0.05, * p<0.1. Tests in Table 6 replicated with a different estimator of labour market power (LMP): production approach instead of revenue approach (see Section 3.2). Same controls (South, city, year) as in Table 6.

G. On the size class of the adoption cutoff

The purpose of this Appendix is to explain how I form the guess that the technology adoption cutoff lies in the size class with firms between 10 and 49 workers (see Section 4.1). Strong evidence shows a positive relationship between firm size and the adoption of digital technologies. Whilst a single size cutoff perfectly predicts technology adoption in my theoretical setting, in real data the share of firms adopting software is only gradually increasing along the size distribution. To make sense of the observed adoption rates, I assume that the probability of adoption is a continuous function increasing in firm size, with the cutoff corresponding to a 50% probability.

First, if the adoption cutoff lies in the smallest size class (i.e. 0-9 employees), all firms in the second size class (i.e. 10-49 employees) would have a probability of adoption above 50%, which is inconsistent with the observed rate, that is 45%. Similarly, if the adoption cutoff lies in the fourth size class (i.e. 100-249 employees), firms in the third size class (i.e. 50-99 employees) would have a probability of adoption below 50%, which is inconsistent with the 67% observed rate. I can then conjecture that the adoption cutoff lies in the two adjacent classes between which the adoption rate crosses 50%.

Because the size distribution is Pareto, firms within each size class concentrate near the lower bound. This suggests that the probability of adoption of most firms in the size class of the cutoff are either very close to 50% or largely below that. The second size has a 45% adoption rate whereas the third has a 67% adoption rate. Because the adoption rate is below but close to 50% in the second size class and is above and far from 50% in the third one, I expect the cutoff to lie in the former.

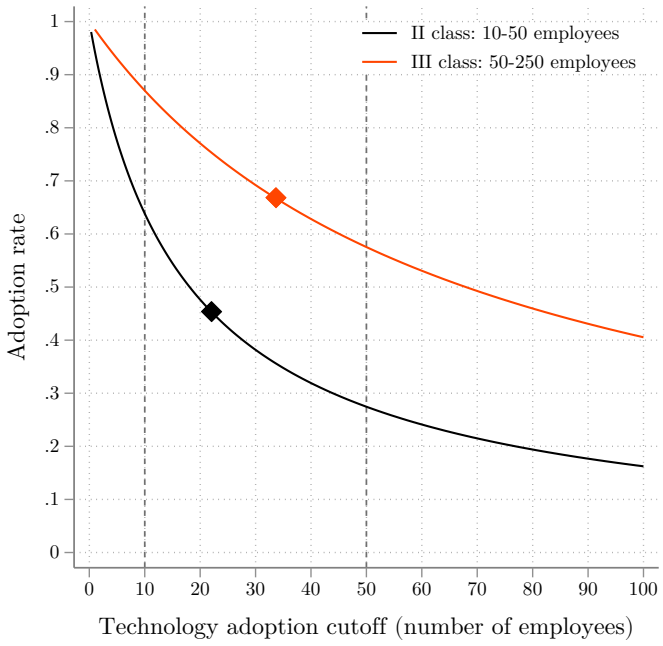
To verify this guess, I formalise the problem. Assume that, for each firm, the probability of adoption, $Pr(t_k = 1+\gamma)$, depends on the log difference between its size, l_k, and the cutoff, l_H, through a standard logistic function (superscripts denote logarithmic transformations).

$$Pr(t_k = 1+\gamma|_{l_k}) = \frac{1}{1+e^{-(l_k^{ln}-l_H^{ln})}} = \frac{l_k}{l_k+l_H} \tag{G1}$$

Denoting $\phi^L(l)$ and $\Phi^L(l)$ as, respectively, the probability density (PDF) and cumulative distribution (CDF) of the size distribution, the adoption rate in the class of firms with size $l_k \in [a,b)$ is given by the following probability:

$$Pr(t_k = 1+\gamma|_{a\leq l_k<b}) = \frac{\int_a^b \frac{l}{l+l_H}\phi^L(l)dl}{\Phi^L(b)-\Phi^L(a)} \tag{G2}$$

FIGURE G1. SHARE OF ADOPTING FIRMS AND THE SIZE CUTOFF



Note: The black curve represents equation (G3) for the boundaries of the second size class. The orange curve represents the same equation for the third size class. The boundaries of the former are highlighted by grey vertical dashed lines. Diamonds denote the observed share of firms adopting software in the corresponding size class.

Assume that firm size follows a Pareto distribution with shape equal to one (Axtell, 2001; Luttmer, 2010). The scale, α_0^L, is irrelevant for my purpose. Substituting the probability density and the cumulative distribution, after some algebra I can express the adoption rate as a function of the cutoff and the lower and upper bounds of the corresponding size class.

$$Pr(t_k = 1 + \gamma|_{a \leq l_k < b}) = \frac{ab}{b-a} log\left(\frac{b(a+l_H)}{a(b+l_H)}\right)\frac{1}{l_H} \tag{G3}$$

Figure G1 represents, for each value of l_H, the corresponding adoption rate in the two size classes in question (boundaries are 10-50 instead of 10-49 as the size distribution in this simulation is continuous). In line with the intuition illustrated above, both observed rates (i.e. 45% in the second size class and 67% in the third one) are only consistent with an adoption cutoff lying within the boundaries of the second size class.

References (Appendix)

Abraham, F., Bormans, Y., Konings, J., & Roeger, W. (2024). Price-cost margins and fixed costs. *The Economic Journal*, *134*(October), 2655–2684

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

Axtell, R. L. (2001). Zipf distribution of US firm sizes. *Science*, *293*(5536), 1818-1820.

Hall, R. E. (2004). Measuring factor adjustment costs. *The Quarterly Journal of Economics*, *119*(3), 899-927.

Luttmer, E. G. (2010). Models of growth and firm heterogeneity. *Annu. Rev. Econ.*, *2*(1), 547-576.

Yeh, C., Macaluso, C., & Hershbein, B. (2022). Monopsony in the US labor market. *American Economic Review*, *112*(7), 2099-2138.