Firm Dynamics, Monopsony, and Aggregate Productivity Differences*

Tristany Armangué-Jubert
UAB, BSE

Tancredi Rapone
UAB, BSE

Alessandro Ruggieri
CUNEF Universidad

December 2, 2025

Abstract

This paper studies how labor market power affects firm dynamics and aggregate productivity. We build a dynamic model of neoclassical monopsony with occupational choice, firm growth, and productivity-enhancing technology adoption. Labor market power lowers efficiency and leads to aggregate output losses by distorting the allocation of labor, entrepreneurship, and innovation decisions. The model is consistent with cross-country evidence of higher life cycle firm growth and higher productivity investment in more competitive labor markets and can explain 25 percent of the differences in aggregate productivity across countries. We find that about 85 percent of the losses are attributable to the lack of technology adoption induced by weaker labor market competition, suggesting that efficiency losses may be greater than those estimated by previous studies.

Keywords: Monopsony, firm dynamics, technology adoption, productivity

JEL Classification: E24, J42, L13

*Armangué-Jubert: Universitat Autonoma de Barcelona, Cerdanyola del Valles, 08193, Spain, tristany.armangue@barcelonagse.eu, Rapone: Universitat Autonoma de Barcelona, Cerdanyola del Valles, 08193, Spain, tancredi.rapone@icloud.com, Ruggieri: Department of Economics, CUNEF Universidad, Calle Leonardo Prieto Castro 2, 28040, Madrid, aruggierimail@gmail.com.

We would like to thank the Editor and the two anonymous referees for their helpful suggestions, which greatly improved the paper. All errors are ours.

1 Introduction

Firms are larger and grow faster over the life cycle in high-income countries. A common explanation is that better-functioning labor markets can improve allocative efficiency, favoring the adoption of better technology (Hsieh and Klenow, 2014). Still, despite being central in the allocation of resources across firms, labor markets around the world are far from competitive and employers operate under imperfect competition almost everywhere (Amodio et al., 2024). Labor market power is a source of inefficiency that can lower returns to adopting more efficient technologies, slow down business dynamism, and explain differences in income across countries.

In this paper, we study how important labor market power is for understanding differences in firm dynamics and aggregate productivity across countries. To this end, we build a general equilibrium model of the labor market featuring an occupational choice between entrepreneurship and wage employment, dynamic investment decisions, and taste for employers à la Card et al. (2018), which limit the firm-level elasticity of labor supply to wages. In the model, agents differ in entrepreneurial productivity and the amenities they could provide as employers. Every period they choose whether to leverage their productivity to open a business or to work for wages. Because workers value employer-specific amenities, entrepreneurs hold wage-setting power and can attract employees despite paying less than their marginal product. Entrepreneurs can also invest to improve their expected future productivity subject to a fixed cost. Through employer turnover, technology adoption, and labor supply decisions, the model generates a host of facts on firm dynamics that can be compared to the data across countries.

We calibrate the model using firm-level micro-data for the Netherlands, a country characterized by a high GDP per capita, i.e., 54,000 USD (in 2017 constant prices) in the year 2020, and an estimate for the average wage markdown as low as 1.301. This value— obtained as the average firm-level ratio between the estimated marginal revenue product of labor and the wage paid (Amodio et al., 2025b)—lies at the lower margin of the estimated spectrum (Sokolova and Sorensen, 2021), and corresponds to a firm-level elasticity of labor supply of

3.318. The benchmark model successfully replicates the observed right-skewed firm size, including the concentration of employers at the top of the distribution. It also matches the average firm size growth over the life cycle, the employer turnover, the firm age distributions, and the share of firms adopting new technologies.

We then move to our counterfactual experiment. Keeping everything else constant, we progressively change the degree of competition between firms by reducing the firms-level labor supply elasticity so to mimic the estimated patterns of wage markdown across middle- and high-income countries. There are several reasons why the labor supply to a firm might be less elastic in poorer countries. Labor markets are often more fragmented due to inadequate transportation and communication infrastructure (Brooks et al., 2021a). Labor mobility costs are higher in poorer countries (Hollweg et al., 2014) and workers are less likely to live and work in urban areas (Ananian and Dellaferrera, 2024), where labor markets tend to be more competitive due to agglomeration effects (Manning, 2010; Luccioletti, 2022). Lastly, imperfect information, social ties, and institutional irregularity could also explain why the labor supply elasticity to firms increases with development (Amodio et al., 2024; Armangué-Jubert et al., 2025; Breza and Kaur, 2025).

The model reproduces the observed cross-country differences in firm growth and technology adoption. Specifically, a counterfactual reduction in the labor supply elasticity to a firm lowers the average firm size and the average firm growth, generates higher labor turnover, and lowers the share of firms investing to improve their productivity. Moreover, we show that differences in wage markdown alone can account for 25 percent of the observed variation in aggregate productivity across middle- and high-income countries, and no less than 11 percent over multiple robustness checks.

In the model, labor market power slows firm dynamics and reduces aggregate income through three channels. The first channel is standard in models of neoclassical monopsony (Card et al., 2018; Dustmann et al., 2022) and it operates through the *static allocation* of workers: lower competition increases the marginal factor cost only for a subset of firms with sufficiently high productivity, spurring employment reallocation towards less-productive, lower-paying employers. The other two mechanisms are relatively novel. Under imperfect competition in the labor market, both *selection into entrepreneurship* and *technology adoption decisions* are altered: lack of competition makes amenities more important determinants of profits, allowing low-productivity agents to reap high benefits from entrepreneurship, and lowering the returns from adopting productivity-enhancing technologies. By penalizing high-productivity employers, labor market power acts as a *skill-biased* force in the labor market, similar to what has been shown for a wide array of size-dependent policies (Guner et al., 2008; Gourio and Roys, 2014; Garicano et al., 2016; Ando, 2021): both mechanisms keep firms inefficiently small and unproductive, reducing firm growth and aggregate output.

Using our calibrated model, we quantitatively decompose each channel. To do so, we compare our benchmark economy with a counterfactual economy featuring the same degree of labor market power observed in Greece. The choice of Greece reflects the following two considerations. First, Greece has approximately one-half the GDP per capita of the Netherlands (i.e., 29,000 USD, in 2017 constant prices, in the year 2020); and second, the degree of labor market competition is much weaker in Greece than in the Netherlands: the estimated wage markdown is equal to 2.623, corresponding to an elasticity of labor supply of 0.616. We find that at least 85 percent of the model-implied difference in aggregate productivity between these two countries is attributable to the lack of technology adoption induced by weaker labor market competition.

This paper relates to recent work on the aggregate costs of labor market power. Berger et al. (2022) find that eliminating labor market power in the US would improve allocative efficiency of labor and rise the average wage by 48 percent, contributing to increase welfare by about 6 percent of lifetime consumption. Armangué-Jubert et al. (2025) find that labor market power can explain 15 percent of the difference in GDP per capita over the development ladder. Deb et al. (2022) show that a less competitive market structure lowered the average wage of low- and high-skilled workers in the US by 12 and 11 percent, respectively. Amodio et al. (2025b) find that fully eliminating labor market power in Peru would increase earnings on average by 26 percent. Bachmann et al. (2022) show that monopsony leads firms to stay inefficiently small and invest less in marketing, and caused a 10 percent loss in aggregate productivity in East Germany. The contribution of this paper is to show that the lack of competition in the labor market also alters selection into entrepreneurship and reduces firms' incentives to innovate, resulting in larger aggregate productivity losses than previously estimated.

Our paper also belongs to the macro literature that focuses on how differences in frictions and distortions could generate the observed cross-country differences in income per capita. (e.g. Bento and Restuccia, 2017; Guner et al., 2018; Guner and Ruggieri, 2022; Da-Rocha et al., 2023; Tamkoç and Ventura, 2024). We contribute to this literature by showing that labor market competition lowers GDP per capita by reducing complementary investments in productivity-enhancing technology.

The remainder of the paper goes as follows. In Section 2 we present cross-country evidence on firm dynamics, technology adoption, and competition in the labor market. We introduce our model in Section 3. In Section 4 we describe the calibration strategy while in Sections 5 and 6 we discuss counterfactual experiments and model mechanisms. We conclude in Section 7.

2 Stylized facts

In this section, we study how firm dynamics and local labor market competition vary across countries with different incomes per capita. For this purpose, we use the World Bank Enterprise Surveys (WBES), conducted by the World Bank. WBES is an establishment-level survey, and it is a representative sample of non-agricultural and non-financial private firms with at least 5 full-time permanent employees, spanning more than 90 countries from 2006 to 2021. The survey follows a stratified sampling methodology along sectors, establishment size, and location, with a common questionnaire for more than 90 countries from 2006 to 2021.

The data covers information on firm-level sales, number of workers, labor cost, the value of machinery, cost of raw materials, and intermediate goods employed in production, together with a large set of additional plant-level demographic characteristics, e.g., age, sector, and location, among others. We complement this data with other aggregate variables, such as real GDP per capita (expressed in USD at 2017 constant prices) from the World Bank's World Development Indicators (WDI).

We restrict our focus to countries that ever had a GDP per capita of above 25,000 USD during the years in which the survey was conducted (Tamkoç and Ventura, 2024) and only consider firms with non-missing observations on annual sales, number of workers, material and capital expenditure.[1] As a result, our sample includes 37,096 firm-year observations spread over 31 countries, consisting of middle- and high-income countries. The poorest country in the sample is Kazakhstan, with a GDP per capita of 19,615 USD in 2009, while the richest one is Ireland, with a GDP per capita of 91,791 USD in 2020. Table A.1 in Appendix A.1 reports the list of countries and years included in the sample.

As it is common in the literature, we conduct our analysis at the local labor market level. In what follows, we define a local labor market as a location-industry pair, where locations are the first administrative level of the country and industries are 2-digit ISIC v.3.1.

Firm dynamics. Figure 1 reports the average firm age (panel A) versus the average annualized firm size growth, residualized from firm age fixed effects (panel B) across countries. Countries are ranked by their real GDP per capita. Each dot refers to the average local labor market in a country. Firm age is computed by subtracting the first year of operations from the year of the survey. The annualized firm size growth is constructed by dividing the observed cumulative firm size growth recorded in the year of the survey by the respective firm age.[2]

[1] Information on production input expenditure is missing for many firms in the construction and service sectors. Therefore, to ensure comparability across countries, we only focus on the manufacturing sector. See also Eslava et al. (2024) for a similar sample selection.

[2] The cumulative firm size growth is computed as a log difference between the current size and the initial size, where the former is defined as the number of employees recorded in the fiscal year before the year of the survey, whereas the latter is defined as the number of employees recorded in the first year of operations.

Figure 1: Firm dynamics across countries

(a) Average firm age **(b)** Average annualized firm size growth

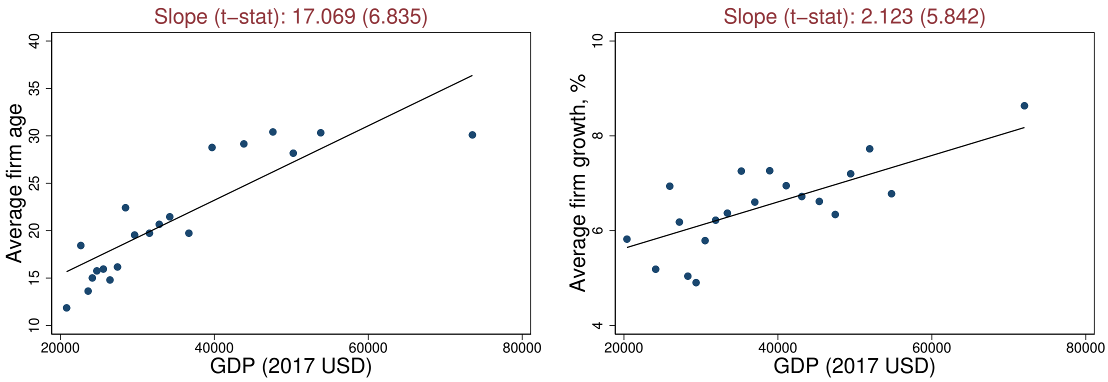

NOTES: Panel A is a binscatter of the average firm age across countries over GDP per capita. Panel B is a binscatter plot of the average annualized firm size growth since birth across countries over GDP per capita, residualized by firm age. Fitted lines are obtained from an auxiliary regression on GDP per capita. In the caption of each panel, we report the estimated slope of the regression and the t-stat (in parentheses). GDP per capita is expressed in 2017 USD constant prices. SOURCE: WBES, WDI, and authors' calculation.

Firms in countries with high GDP per capita are older on average. As we move from poorer to richer countries, the average firm age triples from 11 to almost 30 years. At the same time, firms in countries with high GDP per capita grow faster in size over their life cycle. The average annualized firm size growth increases steeply with GDP per capita: the estimated slope implies a 1.6 percentage point higher firm growth rate per year as we move to a country with double the GDP per capita. Both relationships are significant at a 95 percent confidence level.[3]

These facts were first uncovered by Hsieh and Klenow (2014), who showed that employment growth from birth at a firm in the US is several-fold the one observed in Mexico or India, and were replicated by Eslava et al. (2019) for the case of Colombia versus the US. Our findings are also consistent with Bartelsman et al. (2004), who document that firm churning is higher in developing countries, resulting in a higher share of new firms created over the total, and

[3]The combination of both evidence results into a higher average cumulative firm size growth over the life-cycle in richer countries, even conditional on survival. See Figures A.1 and A.2 in Appendix A.2.

Figure 2: Technology adoption across countries

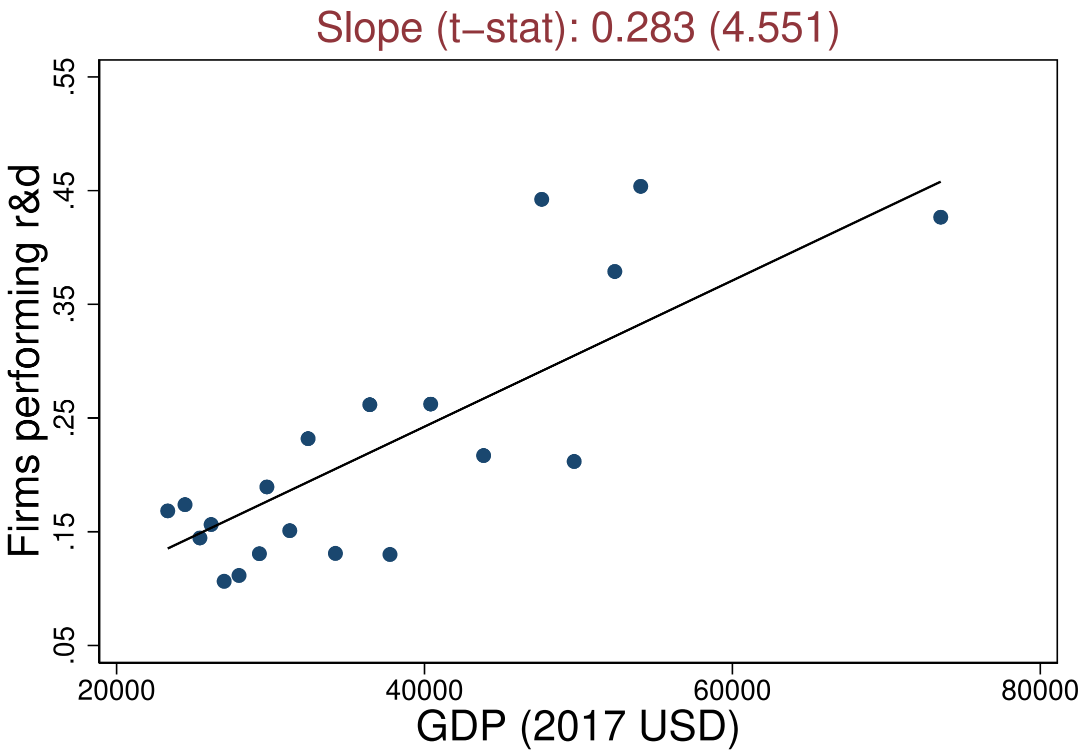

NOTES: This figure shows a binscatter of the share of firms that innovate across countries over their GDP per capita. Fitted lines are obtained from an auxiliary regression on GDP per capita. In the caption, we report the estimated slope of the regression and the t-stat (in parentheses). GDP per capita is expressed in 2017 USD constant prices. SOURCE: WBES, WDI, and authors' calculation.

lower average firm age.[4]

Technology adoption. Figure 2 reports the share of firms that innovate across countries. Like before, each dot refers to the average local labor market in a country. To measure innovation, we use the share of firms that report having conducted formal research and development activities. The WBES reports answers to the question: "During last fiscal year, did this establishment spend on formal research and development activities, either in-house or contracted with

[4] Aga and Francis (2017) document that the probability of firm exit declines with GDP per capita. However, this relation loses significance in more conservative estimations.

Figure 3: Firm innovation over development

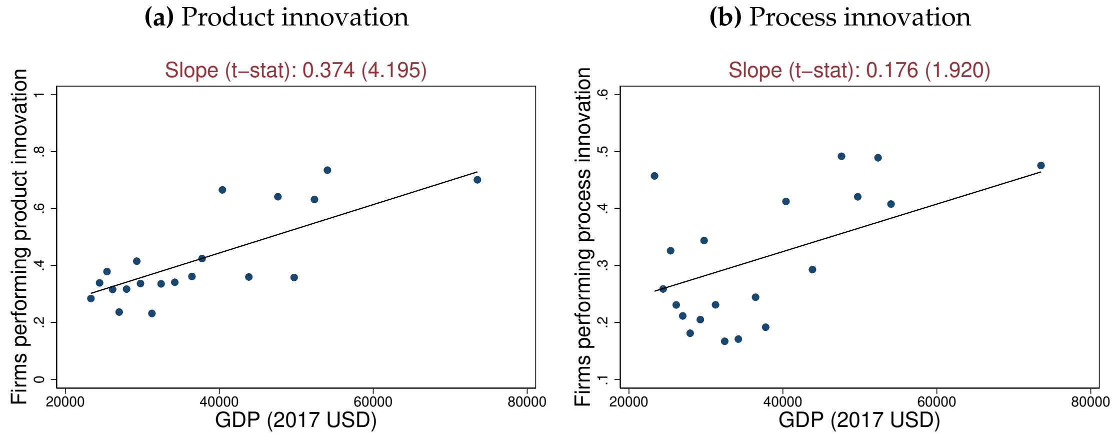

NOTES: Panel A binscatters the share of firms that conduct product innovation across countries over their GDP per capita. Panel B binscatters the share of firms that conduct process innovation across countries over their GDP per capita. Fitted lines are obtained from auxiliary regressions on GDP per capita. GDP per capita is expressed in 2017 USD constant prices. SOURCE: WBES, WDI, and authors' calculation.

other companies, excluding market research surveys?" We construct a binary variable for firm-level innovation by assigning a value of 1 if a firm reports a positive spending on formal research and development activities, 0 otherwise.

Firms in high GDP per capita countries are more likely to perform R&D, whether in-house or contracted. As we move from poorer to richer countries, the share of firms that spend on R&D more than doubles, from around 15 percent to more than 40 percent.

Figure 3 complements Figure 2 by reporting the share of firms performing product innovation (panel A) and process innovation (panel B) across countries with different GDP per capita. These two variables are constructed using answers collected by the WBES to the following two questions: i) "During Last 3 Yrs, Establishment Introduced New/Significantly Improved Process?"; and ii) "New Products/Services Was Introduced Over Last 3 Years?". We construct binary variables for product and process innovation assigning a value 1 to a firm answers affirmatively to the either question, 0 otherwise.

Firms in richer countries are also more likely to conduct both product and process innovation: as we move from middle to high-GDP per capita countries, the share of firms performing product innovation increases from 20 to 80 percent. Similarly, the share of firms performing process innovation increases from 20 to 50 percent.[5]

This evidence is consistent with Lederman and Maloney (2003), who document that R&D expenditure (measured as a share of GDP) is higher in richer countries, and it complements recent findings of Farrokhi et al. (2024), who document relatively low adoption of new technology in low-income countries.

Labor market competition. Finally, we compare the degree of local labor market competition across countries and use firm-level wage markdown as a proxy for labor market power. Specifically, we construct wage markdowns, μ_{it} for firm i at time t as a ratio between the firm-level marginal revenue product of labor and the wage paid (Brooks et al., 2021b; Yeh et al., 2022; Pham, 2023), i.e.

$$\mu_{it} = \frac{\mathrm{MRPL}_{it}}{w_{it}}.$$

Denoting the revenue elasticity of labor as ξ, we can express the wage markdown as a ratio between the revenue elasticity of labor and the wage-bill share of total revenue, i.e.

$$\mu_{it} = \frac{\xi}{\frac{w_{it}\ell_{it}}{y_{it}}}.$$

We measure average revenues, y_{it}/ℓ_{it}, and average wage, w_{it}, using firm annual sales per number of employees and annual payroll per number of employees, respectively, while we estimate ξ separately for each country in the sample using a standard control function approach. We report details of the estimation procedure in Appendix A.4.

Figure 4 scatters the median firm-level markdown in the average local labor market of the sampled countries against GDP per capita. Local labor markets are more competitive in richer countries, and firms charge lower markdowns.[6]

[5] These facts are robust to alternative data sources. See Appendix A.3 for a discussion.

[6] This pattern is robust to estimating firm-level markdown controlling for several firm-level observables. See Figure A.4 in Appendix A.4.

Figure 4: Median markdown across countries

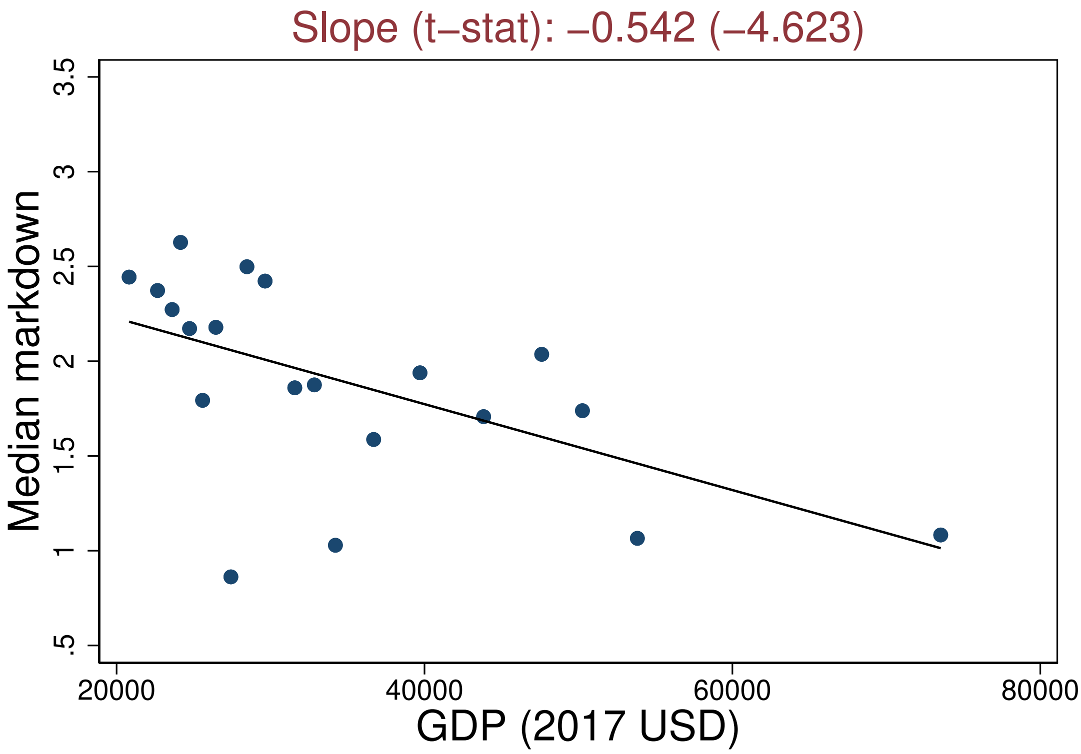

NOTES: This figure shows a binscatter of the median firm-level markdown across countries over GDP per capita. Fitted lines are obtained from an auxiliary regression on log GDP per capita. In the caption of each panel, we report the estimated slope of the regression and the t-stat (in parentheses). GDP per capita is expressed in 2017 USD constant prices. SOURCE: WBES, WDI, and authors' calculation.

In countries with a GDP per capita of 25,000 USD, we estimate a median markdown of 2.25 on average; that is, workers are paid about 55 percent less than their marginal product. These values fall in the range of estimates for middle-income countries obtained from other studies in the literature (Sokolova and Sorensen, 2021). For instance, exploiting variation in firms' exposure to trade and exchange rates, Amodio et al. (2025a) trace out the firm-level labor supply curves of several Latin American countries (GDP per capita ranging from 2,000 USD in Nicaragua to 17,000 USD in Chile, in 2022), and estimate an average labor supply elasticity of 0.82, corresponding to a wage markdown of 2.2. Bassier (2023) uses matched employer-employed data for South Africa (GDP per capita

Figure 5: Mean vs. aggregate markdown across countries

(a) Unweighted average markdown

(b) Aggregate (sales-weighted) markdown

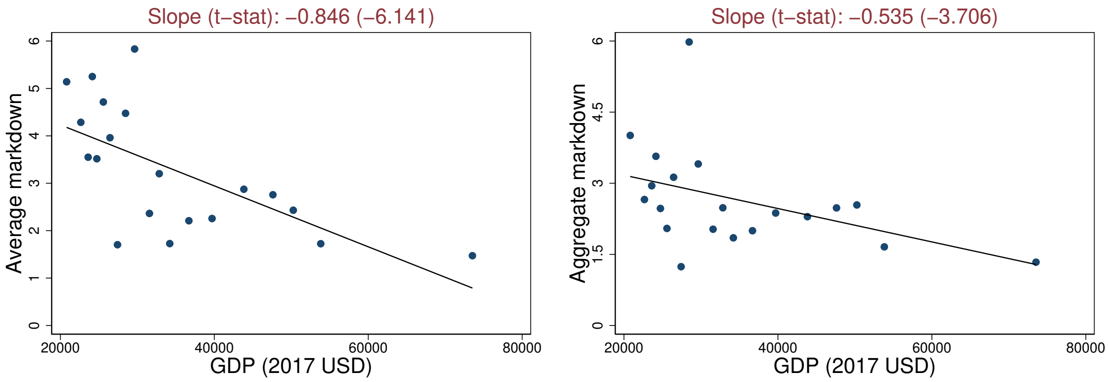

NOTES: Panel A shows a binscatter of the unweighted average markdown across countries over GDP per capita. Panel B shows the aggregate (sales-weighted) markdown across countries over GDP per capita. Fitted lines are obtained from an auxiliary regression on GDP per capita. In the caption of each panel, we report the estimated slope of the regression and the t-stat (in parentheses). GDP per capita is expressed in 2017 USD constant prices. SOURCE: WBES, WDI, and authors' calculation.

of 6,000 USD in 2022) and obtains an estimate for the firm labor supply elasticity of 0.74, corresponding to a wage markdown of 2.35. Garcia-Louzao and Ruggieri (2023) exploits wage variation for workers moving across firms and estimate a labor supply elasticity for Lithuania (GDP per capita of 23,065 USD in 2013) ranging between 0.7 and 0.9, which corresponds to wage markdowns between 2.1 and 2.4. Ogloblin and Brock (2005) document that male workers in Russia (GDP per capita of 25,933 USD in 2012) earn around 35.5 percent less than their potential competitive wage, corresponding to a wage markdown of 1.55.

In high-income countries, the estimated wage markdown is much lower. The median markdown is about 1.25 in countries with a GDP per capita of 60,000 USD, which means workers are paid 20 to 25 percent less than their marginal product. The estimated wage markdowns for high-income countries fall within the range of estimates of 17 and 24 percent provided by Azar et al. (2022) and Berger et al. (2022) for the United States. It also lies between 16 and 25 percent, the estimates obtained by Datta (2022) for the United Kingdom.

Figure 5 complements Figure 4 by reporting the unweighted average markdown (panel A) and the aggregate (sales-weighted) markdown (panel B) across countries. Both measures of markdowns are larger than the median ones but follow a similar pattern over development: labor market power decreases with development regardless of how markdowns are aggregated across firms.

This evidence is consistent with Armangué-Jubert et al. (2025), who show that markdowns decrease with income per capita for countries with GDP per capita over 2,000 USD. This evidence is also consistent with Eslava et al. (2024), who document that markdowns are higher in less developed countries. Using WB-ES data, they show this pattern to be robust across different econometric specifications and alternative sample selections. Their estimated markdown decreases from values around 2 in middle-income economies to 1.2 in high-income countries. Similar to our findings, their estimates of aggregate (sales-weighted) markdowns also decline with GDP per capita, going from about 4 for middle-income countries to 2 for high-income countries. Amodio et al. (2024) document a hump-shaped relationship between GDP per capita and median markdowns for relatively poorer countries than those in our sample, i.e., those with GDP per capita levels below 25,000 USD. Our estimates of median wage markdowns for countries just above this threshold are comparable (in magnitude) to their estimates for countries just below the threshold.

To summarize, firms grow faster and are more likely to innovate in richer countries, where local labor markets are more competitive. In the next section, we build a dynamic model of neoclassical monopsony that fits our stylized facts and use it to study how labor market power affects firm dynamics, productivity-enhancing technology adoption, and aggregate productivity.

3 Model

We extend a standard model of neoclassical monopsony, as discussed in Card et al. (2018) and Dustmann et al. (2022), to a dynamic general equilibrium setting with an entrepreneurial choice and endogenous productivity investment.

Time is discrete. The economy is populated by a unitary measure of agents, each characterized by entrepreneurial productivity z defined over a ladder $\mathcal{Z} = [\underline{z}, ..., z_{-}, z, z_{+}, ..., \overline{z}]$, and amenity a defined over a subset $\mathcal{A}$ of the reals. Agents face a stochastic life cycle, discount the future using a discount factor β, and have a probability of exiting the labor market equal to δ_w. Exiting agents are replaced by an equal share of new entrants who draw a tuple of characteristics (z, a) from two independent distributions, $\Psi_z(z)$, and $\Psi_a(a)$. In what follows we assume both distributions to be log-normal with mean 0 and variances σ_z^2 and σ_a^2, respectively.[7]

Each period following entry, agents decide whether to become wage workers or entrepreneurs. Let L and $E = 1 - L$ denote the (endogenous) aggregate measures of workers and entrepreneurs in the economy, respectively. Entrepreneurial productivity of every agent evolves stochastically over the life cycle, following a discrete time Poisson process which moves it one step up or down the productivity ladder with probability p_n and 1-p_n. Entrepreneurs can invest in innovation, which increases the likelihood of moving up the ladder to $p_i > p_n$, resulting in a higher expected future productivity.

Both workers and entrepreneurs are hand to mouth, they value consumption as well as workplace amenities. Entrepreneurs produce a homogeneous final good, whose price is the numeraire of the economy. Finally, labor markets are assumed to be spot markets that clear every period: entrepreneurs post wages to maximize their profits, with knowledge of workers' labor supply function. Workers observe posted wages and amenities and choose which firms to work for. Job differentiation through amenities endows entrepreneurs with wage-setting power.

3.1 The problem of the workers

Workers derive utility from consumption of a final good and from the work environment that a firm provides. The instantaneous indirect utility for a worker

[7] Assuming independent distributions allows us to restrict the parametric space and simplifies the identification strategy. We will relax this assumption in Section 5.2.

i employed by entrepreneur (firm) j is:

$$u(z_i, a_i, z_j, a_j, \nu_{ij}) = \epsilon^L \ln(w_j) + a_j + \nu_{ij}, \tag{1}$$

with $\epsilon^L > 0$. In equation (1), w_j is the wage paid by firm/entrepreneur j; a_j denotes the amenity provided by firm j, common to all workers within the firm, while ν_{ij} is a match-specific i.i.d. preference shock for working for firm j, assumed to follow a Gumbel distribution with location parameter 0 and scale parameter σ_ν.

The specification of preferences allows for *vertical* and *horizontal* employer differentiation. Both amenities and idiosyncratic preference shocks are meant to capture non-pecuniary job characteristics, such as schedule flexibility, commuting arrangements, workplace safety, provision of parental leave, job autonomy and security, among others (Maestas et al., 2023; Mas and Pallais, 2017; Sorkin, 2018; Sockin, 2024). On the other hand, while differences in a_j imply a ranking across employers, the term ν_{ij} makes workers heterogeneous in their preferences over the same firm.

The indirect utility function in equation (1) is also convenient because it offers a direct map between the parameter ϵ^L and the elasticity of labor supply: lower values of ϵ^L increases the relative utility weight of amenities, making workers' individual labor supply to a firm less responsive to wage changes. It also offers a micro-foundation of the CES preferences for differentiated jobs used, among others, by Berger et al. (2022) and Felix (2022).[8]

Let $\tilde{U}(z_i, a_i)$ be the expected value of being a wage worker for agent i, before a realization of ν_{ij} is observed. Because ν_{ij} is assumed to be independent across

[8] At the aggregate level, CES preferences generate the same labor supply system as in a framework where workers solve a discrete choice problem and firms compete in a monopsonistic labor market. This argument adapts those by Anderson et al. (1988), who establish an equivalence between single-sector CES and single-sector logit, and Verboven (1996), who extends it to nested-logit and nested-CES environments. See Berger et al. (2022), Appendix B for a formal discussion.

alternatives and type-I extreme value distributed, $\tilde{U}(z_i, a_i)$ can be written as

$$\begin{aligned}\tilde{U}(z_i, a_i) &= \mathbb{E}\left[\max_k \{U(z_i, a_i, z_k, a_k) + \nu_{ik}\}\right] \\ &= \sigma_\nu \ln\left(E\int_{\mathcal{Z}\times\mathcal{A}} \exp\left(\frac{U(z_i, a_i, z_k, a_k)}{\sigma_\nu}\right)\psi(z_k, a_k)dz_k da_k\right),\end{aligned}$$

where $U(z_i, a_i, z_j, a_j)$ denotes the worker-i's conditional value of working at firm j, which is equal to

$$\begin{aligned}&U(z_i, a_i, z_j, a_j) = \epsilon^L \ln(w_j) + a_j \\ &+ \beta(1-\delta_w)\big(p_n \max\{\tilde{U}(z_{i+}, a_i), V(z_{i+}, a_i)\} + (1-p_n)\max\{\tilde{U}(z_{i-}, a_i), V(z_{i-}, a_i)\}\big),\end{aligned} \tag{2}$$

while $\psi(z_k, a_k)$ denotes the equilibrium distribution of active entrepreneurs across productivity and amenities, and V is the value of being an entrepreneur, defined below.

Notice that entrepreneurial productivity evolves exogenously for wage workers, i.e., it increases by one step on the ladder with probability p_n, while it decreases with the opposite probability, $1 - p_n$. Notice also that every periods agents face to option of remaining a wage worker or becoming an entrepreneur. The max operator in equation (2) implies a policy function for entrepreneurial choice, $\rho^e(z_i, a_i)$, defined as

$$\rho^e(z_i, a_i) = \begin{cases} 1 & \text{if} \quad V(z_i, a_i) > \tilde{U}(z_i, a_i), \\ 0 & \text{otherwise} \end{cases}$$

Finally, wage worker i will choose to work for entrepreneur j if

$$U(z_i, a_i, z_j, a_j) + \nu_{ij} \geq U(z_i, a_i, z_k, a_k) + \nu_{ik}, \quad \forall k \neq j.$$

Given the assumption on ν_{ik}, the probability that a wage worker i chooses to work for a firm j is given by the following continuous logit formulation (Mc-

Fadden, 1976; Ben-Akiva et al., 1985):

$$p_{ij} = \frac{\exp\left(\frac{U(z_i,a_i,z_j,a_j)}{\sigma_v}\right)}{\int_L^1 \exp\left(\frac{U(z_i,a_i,z_k,a_k)}{\sigma_v}\right) dk}.$$

The aggregate labor supply to a firm j is then equal to:

$$L(z_j, a_j) = L_j = L \int_{\mathcal{Z}\times\mathcal{A}} p_{ij}\phi(z_i, a_i) dz_i da_i \tag{3}$$

where $\phi(z_i, a_i)$ is the equilibrium distribution of workers across productivity and amenities. Re-arranging terms, equation (3) can be re-written as to:

$$L_j = L\Theta \exp\left(\epsilon^L \ln(w_j) + a_j\right)$$

where Θ is an aggregate market shifter.[9] The labor supply solution resembles the one obtained in Card et al. (2018): because the labor market is a spot market, dynamic forces only affect Θ.

3.2 The problem of the entrepreneurs

Entrepreneurs produce a homogeneous final good using a decreasing return to scale production function,

$$Y_j = z_j L_j^{\xi}, \tag{4}$$

where z_j denotes their entrepreneurial productivity, L_j is the labor supplied to their firm, and $\xi \in (0, 1)$ denotes the revenue elasticity of labor. Every period, entrepreneurs post a wage w_j to maximize profits given knowledge of the labor supply function in equation (3). Since entrepreneurs do not observe the preference shocks of individual workers, they cannot perfectly discriminate and will offer the same wage to all of their workers.

[9] We report the full derivation of the labor supply function in Appendix B.1.

The static problem of an entrepreneur j is then given by

$$\max_{w_j} \pi_j(z_j, a_j) = z_j L_j^{\xi} - w_j L_j \tag{5}$$
$$\textit{subject to} \quad L_j = L\Theta \exp\left(\epsilon^L \ln(w_j) + a_j\right).$$

where a_j denotes the amenity provided by entrepreneur j to their employees. A solution to this problem is an optimal wage schedule, $W(z_j, a_j), \forall j$.

Given the solution to the static profit maximization problem, entrepreneurs choose whether to invest in their productivity. Innovation allows entrepreneurs to increase their expected productivity by raising the likelihood of productivity improvement to $p_i > p_n$, hence, by construction, lowering the likelihood of productivity depreciation. To innovate, entrepreneurs incur a per-period fixed cost c_x, while all entrepreneurs pay a fixed cost of operation c_f. Both costs are defined in terms of final goods.[10]

The value to an agent i of being an entrepreneur is then given by

$$V(z_i, a_i) = \max\{V^I(z_i, a_i), V^N(z_i, a_i)\} \tag{6}$$

where $V^I(z_i, a_i)$ is the value of investing in productivity, equal to

$$V^I(z_i, a_i) = \epsilon^L \ln(\pi(z_i, a_i) - c_f - c_z) + a_i$$
$$+ \beta(1-\delta_w)\big(p_i \max\{V(z_{i+}, a_i), \tilde{U}(z_{i+}, a_i)\} + (1-p_i)\max\{V(z_{i-}, a_i), \tilde{U}(z_{i-}, a_i)\}\big)$$

while $V^N(z_i, a_i)$ is value of not investing,

$$V^N(z_i, a_i) = \epsilon^L \ln(\pi(z_i, a_i) - c_f) + a_i$$
$$+ \beta(1-\delta_w)\big(p_n \max\{V(z_{i+}, a_i), \tilde{U}(z_{i+}, a_i)\} + (1-p_n)\max\{V(z_{i-}, a_i), \tilde{U}(z_{i-}, a_i)\}\big)$$

Notice that entrepreneurs value their own amenities, which they also provide to their employees. We interpret this as the consequence of the entrepreneurs sharing the same work environment of their employees (Stephan et al., 2024).

[10] In Section 5.3, we test the robustness of our finding to defining both costs in terms of labor.

This choice also ensures that the instantaneous indirect utility functions of entrepreneurs and wage workers are consistent.[11]

Finally, the max operator in equation (6) implies a policy function for investment into innovation, $\rho^z(z_i, a_i)$, defined as

$$\rho^z(z_i, a_i) = \begin{cases} 1 & \text{if} \quad V^I(z_i, a_i) > V^N(z_i, a_i), \\ 0 & \text{otherwise.} \end{cases}$$

3.3 Equilibrium

A stationary recursive equilibrium for this economy is a list of value functions $V(z_i, a_i)$, $U(z_i, a_i, z_j, a_j)$ and $\tilde{U}(z_i, a_i)$, an associated entrepreneurship policy function $\rho^e(z_i, a_i)$ and innovation policy function $\rho^z(z_i, a_i)$, a wage schedule $W(z_i, a_i)$, an allocation of labor $L(z_i, a_i)$, an aggregate measure of workers L, a distribution of agents over productivity and amenities, $\Omega(z_i, a_i)$, and distributions of wage workers and entrepreneurs over productivity and amenities, $\phi(z_i, a_i)$ and $\psi(z_i, a_i)$, such that:

- The labor supply to firm i, $L(z_i, a_i)$, satisfies equation (3);
- $\rho^e(z_i, a_i)$ and $\rho^z(z_i, a_i)$ solve the entrepreneurial and the innovation choices, and the value functions $V(z_i, a_i)$, $U(z_i, a_i, z_j, a_j)$ and $\tilde{U}(z_i, a_i)$ attain their maxima;
- The aggregate measure of workers is consistent with the entrepreneurial choices:

$$L = \int_{\mathcal{Z}\times\mathcal{A}} L(z_i, a_i)\psi(z_i, a_i) dz_i da_i = \int_{\mathcal{Z}\times\mathcal{A}} (1 - \rho^e(z_i, a_i))\Omega(z_i, a_i) dz_i da_i;$$

- The distribution of agents over productivity and amenities, $\Omega(z_i, a_i)$ is stationary and replicates itself through entry and exit, and the policy functions, as in equations (14), (15) and (16), defined in Appendix B.2.

[11] We experimented with a version of the model where entrepreneurs are precluded from valuing their own amenities and their indirect utility only depends on profit flows. We found no significant quantitative differences in counterfactual outcomes.

- The distributions of wage workers and entrepreneurs over productivity and amenities are stationary and defined as

$$\phi(z_i, a_i) = \frac{(1-\rho^e(z_i,a_i))\Omega(z_i,a_i)}{\int_{\mathcal{Z}\times\mathcal{A}}(1-\rho^e(z_i,a_i))\Omega(z_i,a_i)dz_i da_i},$$

  and

$$\psi(z_i, a_i) = \frac{\rho^e(z_i,a_i)\Omega(z_i,a_i)}{\int_{\mathcal{Z}\times\mathcal{A}}\rho^e(z_i,a_i)\Omega(z_i,a_i)dz_i da_i},$$

  respectively.

A solution algorithm is presented in Appendix B.3.

3.4 Discussion

In the model, competition in the labor market operates as a “skill-biased" force, in the sense of favoring high-productivity entrepreneurs, and it does so through different channels.

To gain some insights, let us assume that the aggregate labor supply L is fixed and constant. Notice that profit maximization (5) subject to equation (3) yields the following equilibrium employment choice by firm j:

$$\ln(L_j) = \frac{\epsilon^L}{1+(1-\xi)\epsilon^L}\ln(z_j) + \frac{1}{1+(1-\xi)\epsilon^L}a_j + C$$

where $C = \frac{1}{1+(1-\xi)\epsilon^L}\left[\epsilon^L \ln\left(\frac{\epsilon^L}{1+\epsilon^L}\right) + \epsilon^L \ln\xi + \ln(L) + \ln(\Theta)\right]$ is a market-level constant. Rearranging the equation above, we obtain that the relative employment between firms with a low- and high-productivity, $\underline{z}$ and $\overline{z}$, and same amenities a, equals:

$$\frac{L(\overline{z},a)}{L(\underline{z},a)} = \left(\frac{\overline{z}}{\underline{z}}\right)^{\frac{\epsilon^L}{1+(1-\xi)\epsilon^L}} \tag{7}$$

Similarly, the relative employment between firms with low- and high-amenities,

$\underline{a}$ and $\overline{a}$, and the same level of productivity z, is equal to:

$$\frac{L(z,\overline{a})}{L(z,\underline{a})} = \left(\frac{\exp(\overline{a})}{\exp(\underline{a})}\right)^{\frac{1}{1+(1-\xi)\epsilon^L}} \tag{8}$$

Equations (7) and (8) predict that when the labor supply elasticity rises, relative employment falls at the lower-productivity and higher-amenities firms. This effect is standard in static models of classical monopsony (Card et al., 2018; Autor et al., 2023; Armangué-Jubert et al., 2025). With a constant aggregate labor supply L, an equilibrium reduction in relative employment at lower-productivity and higher-amenities firms implies labor reallocation towards high-productivity firms and away from high-amenities firms. We summarize this result in our first proposition.

Proposition 1 *Fix the aggregate labor supply. Everything else equal, labor reallocates towards high-productivity firms and away from high-amenities firms when the labor supply elasticity, ϵ^L, increases.*

Proof 1 *Immediate from equations (7) and (8). Full derivation in Appendix B.4.*

Through reallocation of employment across firm types, higher labor competition increases allocative efficiency. To the extent that resource allocation is more efficient when firms' labor market power is lower, a higher elasticity of labor supply ϵ^L will result in lower differences in the average revenue product of labor (APL_j) across firms.[12] To see this, notice that the elasticities of APL_j with respect to firm-level productivity, z_j, and amenities, $\exp(a_j)$, are equal to

$$\frac{\partial \ln \text{APL}_j}{\partial \ln z_j} = 1 - \frac{(1-\xi)\epsilon^L}{1+(1-\xi)\epsilon^L}, \tag{9}$$

and

$$\frac{\partial \ln \text{APL}_j}{\partial a_j} = -\frac{(1-\xi)}{1+(1-\xi)\epsilon^L}, \tag{10}$$

[12] Labor misallocation generates differences in the marginal revenue products of labor (Hsieh and Klenow, 2009, 2014). Given production technology, dispersion in marginal revenue product result into dispersion of average revenue product, as the latter is proportional to the former.

respectively. When ϵ^L is low, the average revenue product becomes largely dispersed across firms: it increases with firm-level productivity and declines with firm-level amenities. Which suggests high-productivity low-amenity firms face higher barriers in less competitive markets, rendering them smaller than optimal. As ϵ^L increases, both $\partial \ln \text{APL}_j / \partial \ln z_j$ and $\partial \ln \text{APL}_j / \partial a_j$ approach zero, and APL_j equalize across firms. We summarize this result in the following proposition.

Proposition 2 *Fix the aggregate labor supply. The dispersion of the average product of labor across firms reduces when the labor supply elasticity, ϵ^L, increases.*

Proof 2 *Immediate from equations (9) and (10). Full derivation in Appendix B.4.*

Like employment, profits reallocates from high-amenities firms to high-productivity firms when the elasticity ϵ^L increases. Notice that the elasticities of firm-j variable profits with respect to firm-level productivity, z_j, and amenities, $\exp(a_j)$, are equal to

$$\frac{\partial \ln \pi_j}{\partial \ln z_j} = \frac{1+\epsilon^L}{1+(1-\xi)\epsilon^L}, \tag{11}$$

and

$$\frac{\partial \ln \pi_j}{\partial a_j} = \frac{\xi}{1+(1-\xi)\epsilon^L}, \tag{12}$$

respectively. The increase in variable profits from a proportional increase in productivity is higher when the labor supply elasticity ϵ^L is high. On the other hand, variable profits increase less with amenities when the labor supply elasticity ϵ^L is high. We summarize this result in our third proposition.

Proposition 3 *Fix the aggregate labor supply. Variable profits are more (less) correlated to firm-level productivity (amenities) when the labor supply elasticity, ϵ^L, increases.*

Proof 3 *Immediate from equations (11) and (12). Full derivation in Appendix B.4.*

With a low labor supply elasticity, agents with high amenities have wage-setting power as entrepreneurs. High enough to attract enough workers to produce and make sufficient profits to compete in the market, even with relatively low productivity. When labor supply elasticity increases, the competitive advantage shifts away from high-amenities firms and moves towards high-productivity firms. This alters the decision of becoming entrepreneurs, improving selection in favor of high-productivity agents. It also alters the decision of adopting productivity-enhancing technology, which becomes more profitable as the return from investment is increasing in the labor supply elasticity. This results in higher firm growth, higher productive efficiency, and higher output per capita.[13]

In the next section, we use our model to quantify how much each mechanism, i.e., labor allocation across employers, selection into entrepreneurship, and innovation decision, contributes to fostering firm dynamics and productivity, allowing labor supply to react to changes in labor market power.

4 Estimation

We discipline the model using WBES data for the Netherlands, one of the richest countries in the sample, with an annual GDP per capita of 54,275 USD. We follow Armangué-Jubert et al. (2025) and calibrate the model to replicate the average labor market in the country, as defined by a region-industry pair.

Some parameters are calibrated without solving the model. We chose a model period of a year. We normalize the scale parameter of the Type-I GEV shock, σ_ν, to 1. We set the discount factor, β, to 0.961, consistent with an annual interest rate of 0.04, and choose δ_w to be 0.025 such that agents spend on average 40 years in the labor market. We calibrate the revenue elasticity of labor, ξ, using the corresponding estimate obtained in Section 2, i.e., 0.333. Finally, we use the estimated wage markdown to back out the labor supply elasticity. Given the

[13]While these three mechanisms could be generated, using the language of the misallocation literature, by a reduced-form *wedge* on the marginal product of labor (e.g., Hsieh and Klenow (2014)), in our framework labor-market monopsony alters marginal factor costs asymmetrically over the distribution of firms, and productivity- and amenity-correlated distortions arise *endogenously* as a result.

monopsonistic labor market structure, the elasticity of labor supply is equal to

$$\epsilon^L = \frac{1}{\mu - 1}$$

where μ is the wage markdown for firms in the local labor market. We set μ equal to the median wage markdown in the Netherlands. In Section 2, we estimated this value to be 1.301. Which implies a labor supply elasticity of 3.318. Table C.1 in Appendix C.1 summarizes the value of these parameters and their targets.

The remaining parameters, collected in the following vector,

$$\vartheta := \{c_f, c_x, p_i, p_n, \sigma_z, \sigma_a\},$$

are estimated using the method of simulated moments. The estimator $\hat{\vartheta}$ is the minimizer of the following objective function:

$$\hat{\vartheta} = \arg\min_{\vartheta} d(\vartheta)' W d(\vartheta)$$

where $d(\vartheta)$ denotes the absolute distance between a vector of empirical targets, $\bar{g}$ and their model counterpart, $g(\vartheta)$, while the weighting matrix, W, is chosen to be diagonal with entries equal to the squared inverse of each empirical moment.[14]

Table 1 describes the list of estimated parameters, their point estimates, together with standard errors computed using the standard asymptotic variance expression (Newey and McFadden, 1994).[15] Table C.2 in Appendix C.1 reports the list of targeted moments and their model counterpart.[16]

[14]We opted for this weighting matrix to ensure the stability of our estimator while maintaining consistency and keeping the estimation loss independent of units of measurement.

[15]Specifically, the variance covariance matrix is $(J'WJ')^{-1}J'W\hat{Q}WJ(J'WJ')^{-1}$, where J is the Jacobian matrix, whose (i,j)-entry is equal to $\partial d(\vartheta(i))/\partial\vartheta(j)$, while $\hat{Q} = \hat{\text{cov}}(d(\vartheta))$ is bootstrapped from the sample data with 1000 replications.

[16]While our estimation allows firm size to be arbitrarily small, our database does not cover plants with less than 5 workers. To avoid any inconsistency, we apply the same truncation to our simulated moments (see Coşar et al. (2016) for a similar approach). In Section 5.4, we test the robustness of our results to this approach and re-estimate the model using empirical moments imputed to cover a representative sample of of *existing* producers.

Table 1: Parameters estimates and standard errors

Parameters	Description	Estimates	St.Error
c_f	Operating costs	22.80	0.206
c_x	Innovation costs	122.4	1.343
p_i	Productivity growth of investors	0.649	0.274
p_n	Productivity growth of non-investors	0.499	0.020
σ_z	Productivity dispersion	2.134	0.044
σ_a	Amenities dispersion	0.899	0.069

NOTES: The table shows the list of estimated parameters, their point estimates, and the standard errors computed using the Delta method.

The operating cost, c_f is estimated to match an average firm size of 59 employees while the innovation cost is chosen to match a share of firms investing in R&D of 41 percent. The productivity dynamics of investors, p_i, is informed by the average cumulative employment growth rate since entry among observed incumbents. In the data, the latter amounts to 153 percent. The average firm age, which is approximately 30 y.o., is instead meant to discipline the productivity dynamics of non-investors, p_n through entry and exit into entrepreneurship. Finally, the dispersions in entrepreneurial talents at entry, σ_z, and amenities, σ_a, are disciplined by the standard deviation of (log) firm size (0.988) and (log) wages (0.418), respectively. The fit of the model is very satisfactory: the sum of squared deviations between empirical and simulated moments is equal to 1.7 percent.[17]

The estimates of c_f and c_x are 22.8 and 122.4, respectively. These values correspond to about 8 and 43 percent of the average profits made by incumbent firms in the model. Jump probabilities, p_i and p_n, are estimated to be 0.65 and 0.50, respectively; i.e., firms adopting productivity-enhancing technology are 15 percent more likely to grow compared to those who do not. Finally, initial firm productivity levels and amenities are largely dispersed in the cross-section of firms: the estimates of σ_z and σ_a are 2.13 and 0.90, respectively.

The standard errors on these estimates are small, suggesting the empirical mo-

[17] The estimation loss is minimized across all dimensions in the parametric space. See Figure C.1 in Appendix C.1.

Figure 6: Non targeted moments - Model vs. Data

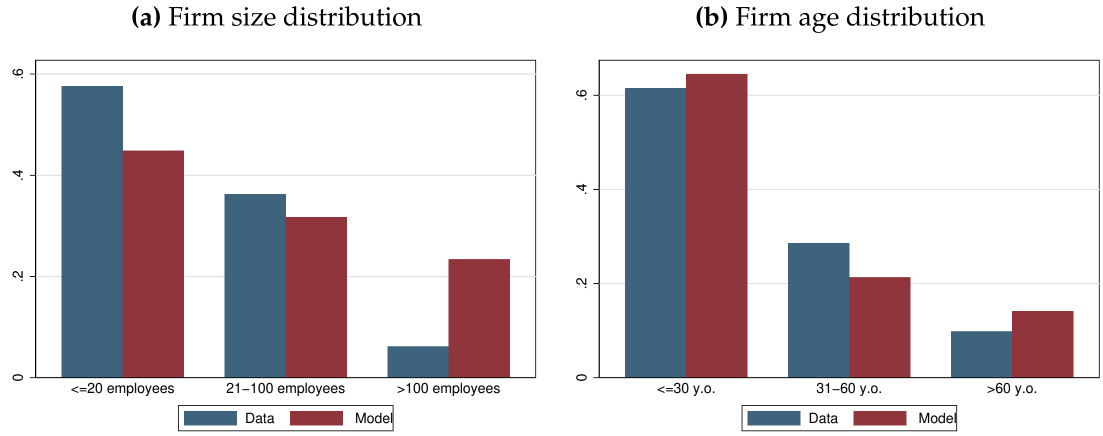

NOTES: Blue bars refer the shares of firms over firm size groups (panel A) and firm age groups (panel B) observed in the data. Red bars refer to the corresponding shares predicted by the model.

ments are informative about the estimated parameters. To check that correct identification is achieved, Figure C.2 in Appendix C.2 shows that each targeted moment moves significantly and monotonically in response to changes in the parameter they are meant to inform.[18]

The model also replicates the empirical firm size and age distributions observed in the Netherlands despite neither being part of the targeted moments. Panel A of Figure 6 reports the percent of firms belonging to different firm size bins, in the data (blue bars) and the model (red bars). In the data, about 57.6 percent of the observed firms have less than 20 employees, while only around 6.2 percent of them employ more than 100 employees. The model reproduces this pattern. Panel B reports the percentage of firms across different firm age groups in the data and the model. In both cases, around 60 percent of firms are under 30 years old, while around 10 percent of them are over 60 years old.

[18] In Appendix C.2, we also show that i) the cross-contamination in the estimates is limited (Figures C.3, C.4, and C.5); and ii) the parameters estimates are not sensitive to the calibrated value of the revenue elasticity of labor, ξ (Figure C.6).

5 Labor market power and firm dynamics

We are now ready to discuss how labor market power affects firm dynamics and aggregate productivity. To this end, we construct counterfactual economies that differ from the benchmark only with respect to their labor supply elasticity while leaving all other parameters unchanged. As a result, the counterfactual economies are replicas of the Netherlands, except for differences in ϵ^L. In the benchmark economy, the labor supply elasticity is equal to 3.318, a value chosen to match a median markdown of 1.301. In the counterfactual economies, we let the elasticity vary between 0.6 and 5.5. These values correspond to wage markdowns ranging from 1.18 to 3, approximately the same values we estimate for our sample of middle- and high-income countries in Section 2.

Figure 7 reports the average cumulative life cycle firm growth (panel A), the average firm size (panel B), the share of firms adopting productivity-enhancing technologies (panel C), and the average firm productivity (panel D), for counterfactual economies with different degrees of labor market competition. The red dot refers to the benchmark economy, the Netherlands. The blue dots refer to counterfactual scenarios.

The model is consistent with the evidence described in Section 2. Labor market power affects firm dynamics over the life cycle. Reducing labor market competition leads to a significant reduction in the cumulative firm size growth. As wage markdown increases from 1.2 to 3, the average firm growth rate shrinks by more than a half, from 150 to about 50 percent. Lower firm growth reduces the average firm size (panel B). As labor supply elasticity reduces, the average firm size drops from approximately 60 to around 45 employees.

Like in the data, firms also are more likely to innovate when the labor market is more competitive (panel C). As we increase the labor supply elasticity, the share of firms investing in productivity doubles from 25 to 40 percent. As a result, the average firm productivity is higher (by about 20 percent) in more competitive labor markets (panel D).

Figure 7: Firm Dynamics, Technology Adoption, and Labor Market Power

(a) Cumulative firm growth

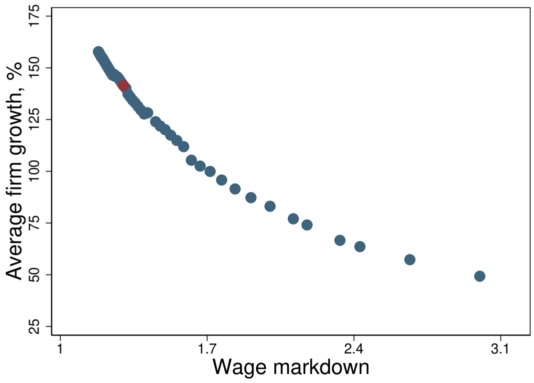

(b) Firm size

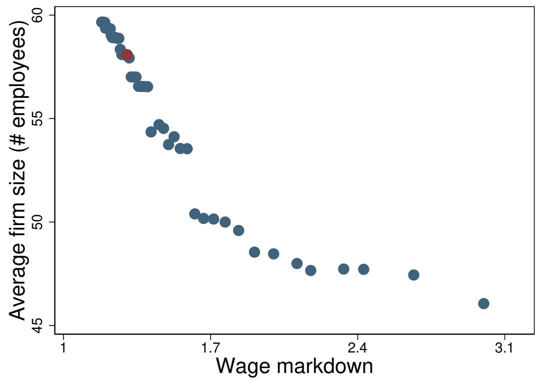

(c) Innovating firms

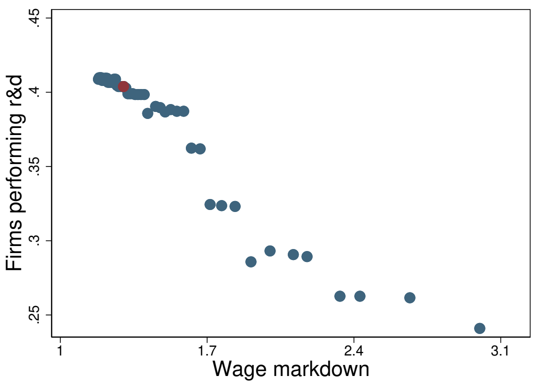

(d) Firm-level productivity

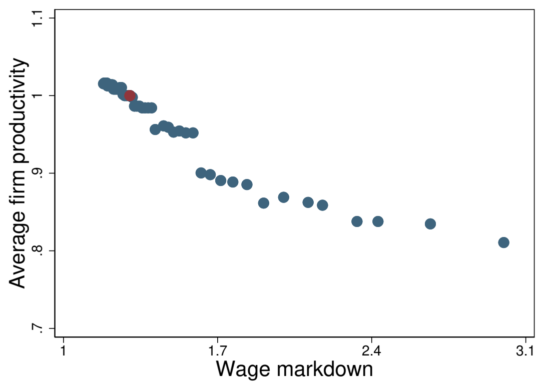

NOTES: The red circle refers to the Netherlands. Blue circles refer to counterfactual economies differing in their labor supply elasticity. Panel A reports the average cumulative firm size growth. Panel B reports the average firm size. Panel C reports the share of firms adopting productivity-enhancing technology. Panel D reports the average firm productivity, relative to the baseline (normalized to 1).

5.1 Cross-country productivity differences

How important is labor market power in generating aggregate productivity dispersion in our sample of countries? We can answer this question through the lens of our model.

Figure 8 scatters the total factor productivity (TFP) observed of each country in our sample against the estimated wage markdown (blue dots) and compares

Figure 8: Cross-Country TFP Differences: Model vs. Data

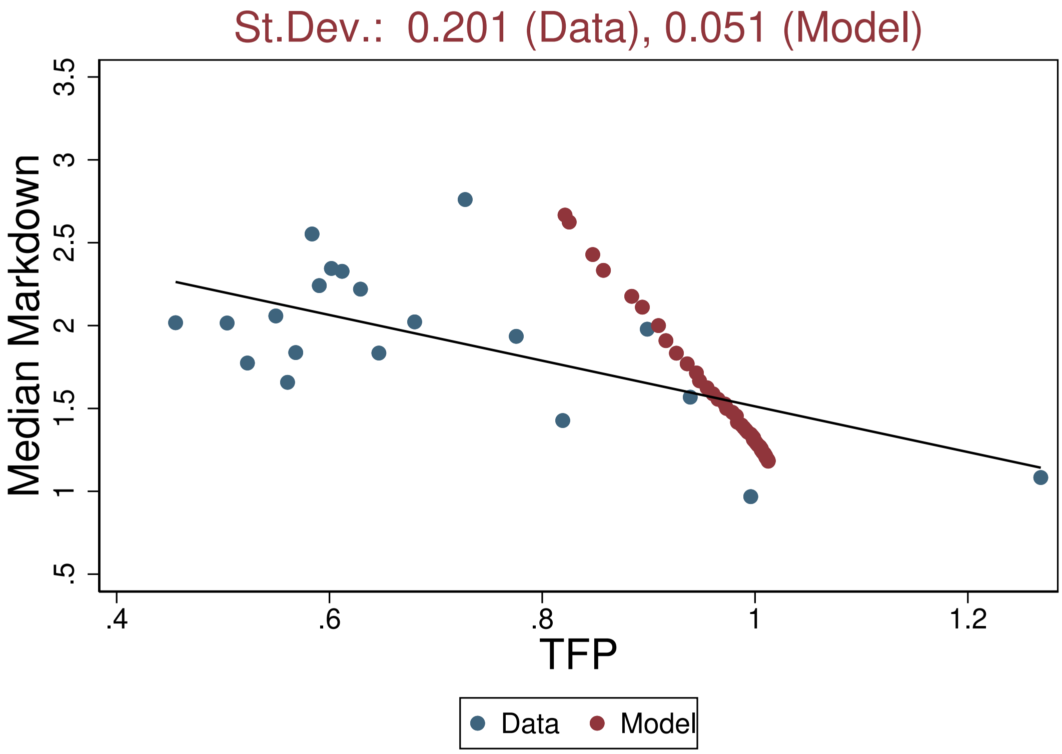

NOTES: Blue dots refer to measured TFP, constructed using data from the Penn Word Table, v. 10.01; that is, TFP$\equiv \left(\frac{rgdpna}{(emp \times avh)^{2/3} rnna^{1/3}}\right)$, where *rgdpna* denotes real GDP at constant 2017 USD, *emp* is the number of persons engaged, *avh* is the average annual hours worked by persons engaged, and *rnna* is the capital stock at constant 2017 USD. Red dots refer to model-based predicted TFP (i.e., GDP per worker), obtained by changing the value of ϵ^L to match the corresponding markdown in the data. Both variables are reported as fraction of the value measured in the Netherlands. SOURCE: PWT, WDI and authors' calculation.

that to simulated TFP obtained in counterfactual economies that feature the labor supply elasticity in line with our estimates of wage markdown reported in Section 2 (red dots). As before, all other parameters are kept fixed at their benchmark values. Both measured and simulated TFP are reported as fraction of the value measured for the Netherlands. Notice that, despite considering only middle and high-income countries, the observed differences in TFP are significant, and amount up to a factor of 2.5.

A few comments are in order. First, there is a positive correlation between measured and simulated TFP across countries. This is because a lower labor supply elasticity generates a high wage markdown, which slows down firm dynamics, reduces efficiency, and lowers aggregate output. On the other hand, the model generates less variation in TFP than is observed in the data: through the lens of our model, the observed cross-country variation in the estimated markdown only explains a portion of TFP dispersion observed in our sample.

To quantify the contribution of labor market power, we compute the standard deviation in model-based and observed TFP. We find values of 0.05 and 0.20, respectively. We interpret it as the ability of the model to account for about 25 percent (=0.051/0.201*100) of the observed variation in TFP across countries.[19] This value is similar to the estimates of GDP losses caused by size-dependent distortions (Restuccia and Rogerson, 2008; Bento and Restuccia, 2017; Tamkoç and Ventura, 2024), and is larger than gains from reducing firms' labor market power obtained using static models of imperfect competition (Berger et al., 2022; Amodio et al., 2025b; Armangué-Jubert et al., 2025).

5.2 The role of productivity-amenity correlation.

Our quantitative results rely on the assumption that initial entrepreneurial productivity is independent of amenities. In Appendix C.5, we assess the robustness of our findings by relaxing this assumption. To do so, we solve a version of the model where new entrants draw their initial productivity and amenities from a bi-variate zero-mean log-normal distribution, $\Psi(z, a)$, with variances equal to σ_z^2 and σ_a^2, and correlation σ_{za}.

We estimate σ_{za} through indirect inference, i.e., we choose σ_{za} to match the correlation between firms' (log) wage and job amenity documented in Sockin (2024). They estimate firm-specific (log) wage and satisfaction premia using matched employer-employee data on wages and employer ratings in the United States, and find a positive (conditional) correlation of 0.622 (Table 2, column 3): satisfaction improves significantly at higher-paying firms. We extend our set of target to include this correlation, and estimate σ_{za} to be 0.296. We report further

[19]Similar results are obtained when we compare observed versus model-based predicted GDP per capita. See Figure C.7 in Appendix C.3.

estimation details in Appendix C.5.

When initial entrepreneurial productivity and amenities are positively correlated, the counterfactual experiments generate a standard deviation in model-based TFP of 0.03. That is, the model is able to account for about 15 percent (=0.030/0.201*100) of the observed variation in GDP per capita across countries—see Figure C.10 in Appendix C.5. This value is slightly lower than the one obtained in main analysis (i.e., 25 percent): this is because, when firm-level productivity and amenities are positively correlated, the scope for reallocating labor and innovation from high-amenities firms to high-productivity firms is smaller, hence reducing the gains of having a more competitive labor market.

5.3 The role of entry and investment costs.

In the baseline version of the model, operating and investment costs are both fixed output costs. Klenow and Li (2025) argue against this assumption. They document that average employment per firm increases with the level of overall labor productivity, both over time and across space, and argue that specifying costs in terms of final goods in models of firm dynamics might produce implications that are at odds with this evidence. By contrast, they suggest denominating costs in terms of labor.

We test the robustness of our results to expressing costs in terms of labor. To do so, we assume that c_f and c_x are both overhead labor costs entrepreneurs need to incur to operate in the industry and to innovate. In Appendix C.6, we report value functions, definition of equilibrium, and estimation details.

We then use this model to conduct the same analysis as in Section 5.1. That is, we generate a series of counterfactual economies that feature the labor supply elasticity in line with our estimates of wage markdown reported in Section 2, and compared the simulated TFP of these economies to the observed ones of each countries in our sample—see Figure C.11 in Appendix C.6.

Expressing operating and innovation costs in terms of labor does not alter the importance of labor market power. The standard deviation in model-based aggregate productivity is 0.04. That is, the model can account for about 22 percent

(=0.044/0.201*100) of the observed variation in TFP across countries. This value aligns to the one obtained in the main analysis (i.e., 25 percent), suggesting our findings extends to an alternative scenario where fixed costs are expressed in terms of labor.

5.4 Robustness

Imputed firm-level moments. The WB-ES data only provides information for firms with at least 5 employees, making the sample biased towards relatively larger firms. In Appendix C.7, we re-estimate our baseline model using firm-level moments that are imputed to cover the entire span of *existing firms* in the Netherlands.

Our findings are robust. We find that the standard deviation in model-based TFP is 0.028— see Figure C.12 in Appendix C.5. This version of the model can explain about 14 percent (=0.028/0.201*100) of the observed variation in TFP across countries, a value slightly smaller than the one obtained in the main analysis.

Alternative identification. Finally, we test how robust our findings are to an alternative identification strategy and re-estimate the baseline model using a different set of moments. Specifically, we make the model over-identified by targeting the shares of firms with different size (≤ 20, 21-100, and > 100 employees), the shares of firms of different ages (≤ 30 y.o., between 30-60 y.o., and > 60 y.o.), and the shares of firms investing in R&D by number of employees. Moreover, we replace the average cumulative firm growth rate with the average annualized growth rate. Details of the estimation are reported in Appendix C.8.

Expanding the set of targets only slightly alter our counterfactual results. With a standard deviation in model-based TFP of 0.023, the over-identified model can still explain about 11 percent (=0.023/0.201*100) of the observed cross-country variation in TFP—see Figure C.13 in Appendix C.8.

6 Mechanisms

In this section, we shed light on the model mechanisms behind the outcomes presented in the previous sections. To keep the discussion compact, we compare the benchmark economy (Netherlands) with a single counterfactual economy, featuring the same degree of labor market power observed in Greece. This choice is motivated by two reasons, i.e. i) Greece has one of the lowest GDP per capita in the sample, approximately one-half of that of the Netherlands (29,000 vs. 54,000 USD); and ii) the degree of labor market competition is much weaker in Greece than the Netherlands: the estimated wage markdown is equal to 1.301 (vs. 2.623), corresponding to an elasticity of labor supply of 3.318 (vs. 0.616). To implement this experiment, we keep all parameters fixed at their benchmark values (Netherlands), and change the elasticity of labor supply to reproduce the median wage markdown estimated for Greece.

Compared to the Netherlands, the average size of firms is smaller in Greece (26 vs. 59 employees), firms grow less over the life cycle (84.5 vs. 153 percent), they are on average younger (the average age is 22.5 years vs. 30), and are less likely to invest in productivity innovation (18 vs. 41 percent). Differences in labor market competition can explain 29 percent of the differences in firm size, account for differences in average firm growth, and explain 27 and 74 percent of the differences in average firm age and share of firms investing in R&D, respectively (see Table C.3 in Appendix C.4).

Why does firm dynamics slow down when labor markets are less competitive? Figure 9 shows how firm-level employment (Panel A) and profits (Panel B) change when we move from a counterfactual economy with low labor market competition (Greece) to the benchmark economy with higher labor competition (Netherlands). Warmer-colored areas refer to firms with states (z_j, a_j) expanding in size and making higher profits when we increase the labor supply elasticity, ϵ^L. Cooler-colored areas refer to firms shrinking and losing profits.

As discussed in Section 3.4, labor reallocates to high-productive, low-amenity firms and away high-amenities firms when labor market competition is stronger. As the elasticity of labor supply reduces, the relative importance of amenities

Figure 9: Employment, Firms, and Investment Reallocation

(a) Profits **(b)** Employment

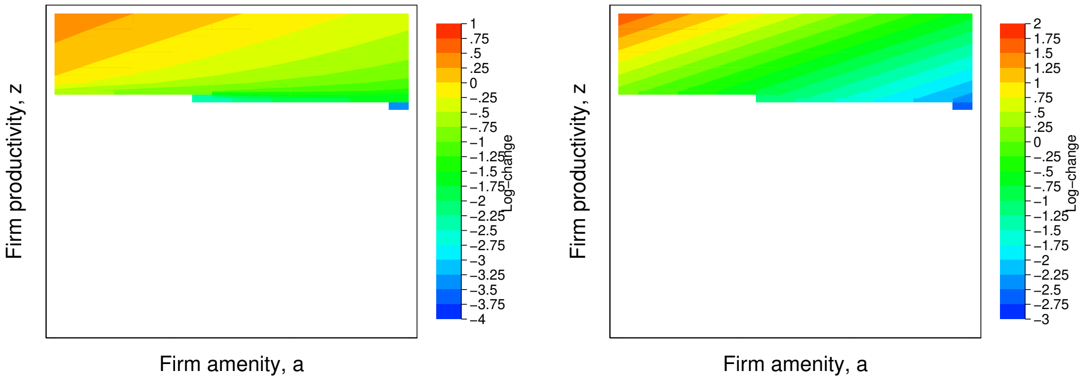

NOTES: Panel A shows the log-difference in firm-level profits for firms with different levels of firm productivity and amenities, between the benchmark economy, i.e. the Netherlands, and counterfactual economy, i.e. Greece. Panel B shows the log difference in firm-level employment between benchmark and counterfactual.

increases, workers become less responsive to wage differences, and lack of competition pushes firms to post wages farther from the marginal revenue product of labor. As a result, profits increase for low-productivity high-amenity firms, allowing them to survive in the market, distorting allocative efficiency.

Changes in employment and profits also alter entrepreneurial and innovation decisions. Panels A and B of Figure 10 show the share of agents that choose to become entrepreneurs in the benchmark (blue line) and the counterfactual economy (red line), by levels of productivity and amenities. Similarly, Figure 11 compares the share of entrepreneurs that adopt productivity-enhancing technologies by levels of productivity (panel A) and amenities (panel B) in the benchmark (blue line) and the counterfactual economy (red line). In Appendix C.4 we report the associated policy functions for entrepreneurship and technology adoption.

In the model, agents choose whether to become entrepreneurs or wage workers aware of the equilibrium labor supply and demand curves. When ϵ^L is low, and the relative importance of amenities is higher, agents with low entrepreneurial productivity but high amenities can attract workers despite pay-

Figure 10: Entrepreneurs

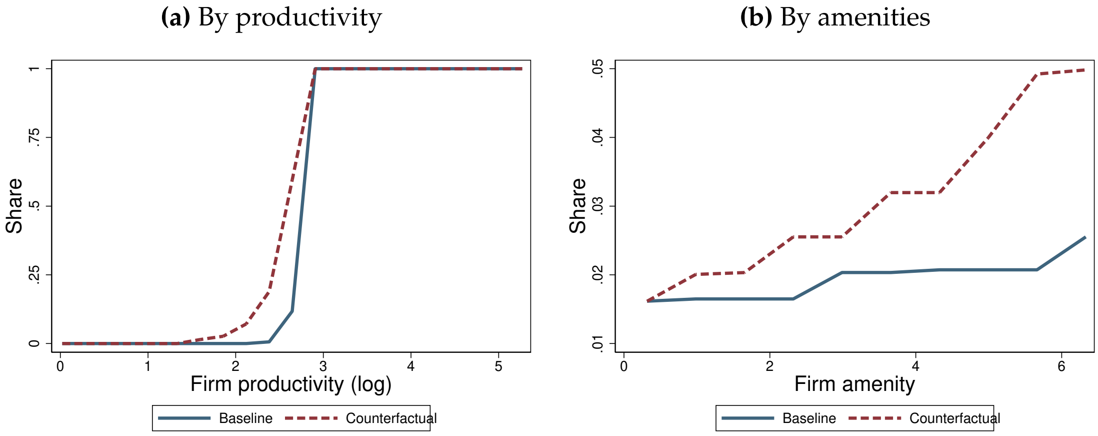

NOTES: This figure shows the share of agents that become entrepreneurs by level of (log) productivity (Panel A) and level of amenities (Panel B) in the benchmark (blue line) and counterfactual (red dashed line) economy. The benchmark economy refers to the Netherlands. The counterfactual economy refers to an economy where ϵ^L is chosen to match the median markdown observed in Greece (leaving other parameters to their benchmark values).

Figure 11: Firms Adopting Productivity-Enhancing Technology

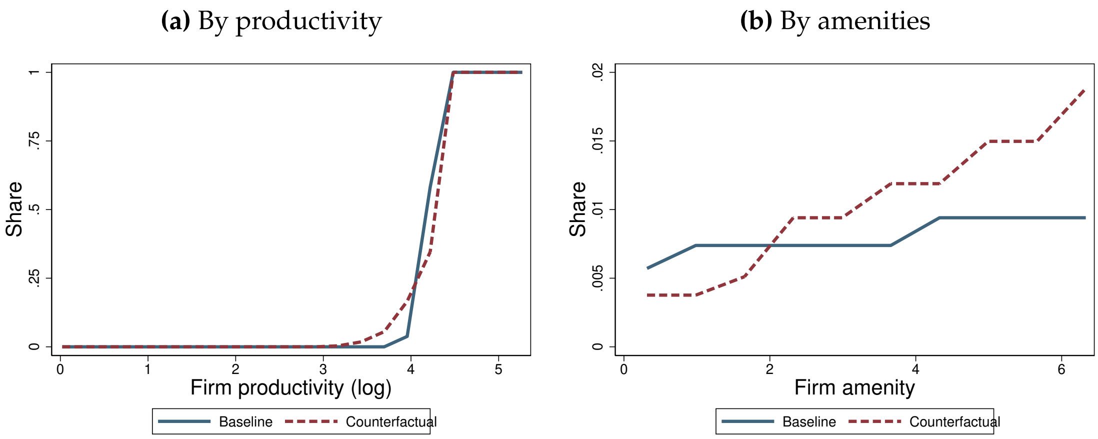

NOTES: This figure shows the share of entrepreneurs that adopt productivity-enhancing technologies by levels of (log) productivity (panel A) and amenities (panel B) in the benchmark (blue line) and counterfactual (red dashed line) economy. The benchmark economy refers to the Netherlands. The counterfactual economy refers to an economy where ϵ^L is chosen to match the median markdown observed in Greece (leaving other parameters to their benchmark values).

ing lower wages. Hence they anticipate being able to make net profits beyond what they could earn as wage workers, and choose to do so by self-selecting into entrepreneurship. Similarly, when ϵ^L is low and the role of productivity in shaping profits is weaker, returns to innovation diminish. This, together with the reallocation of profits towards high amenities and low-productivity firms, reduces dynamic investment, particularly for high-productivity firms. By favoring high-productivity firms, labor market competition operates a skill-biased force.

Figure 12 describes the overall distortionary effects of labor market power. It reports the simulated (log) average revenue product of labor (APL) for firms with different levels of productivity (panel A) and amenities (panel B) in the baseline (blue line) and counterfactual (red dashed line) economies. In a fully undistorted economy, the average revenue product of labor would not vary across firms. Instead, in both scenarios it increases with productivity and declines with amenities. At the same time, it changes much less steeply with productivity and amenity when labor market is more competitive: the elasticities of the APL with respect to productivity and amenities are 0.31 and -0.28 in the baseline economy, respectively, versus 0.79 and -0.81 in the counterfactual.

Figure 12: Revenue product versus productivity and amenities

(a) By productivity **(b)** By amenities

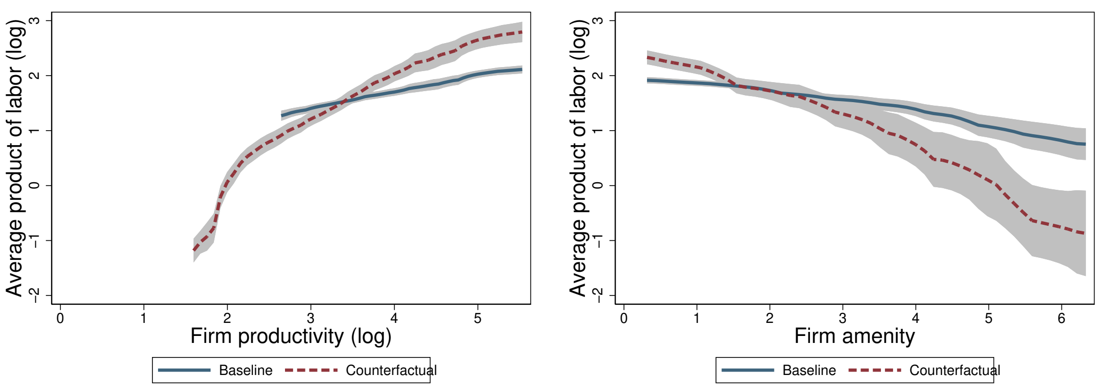

NOTES: This figure reports the simulated average revenue product of labor (in logs) for firms with different levels of productivity (panel A) and amenities (panel B) in the baseline economy (blue line) and counterfactual scenario (red dashed line). The shaded areas refer to 95 percent confidence interval constructed using simulated data.

Finally, we decompose how much each margin, meaning, allocation of labor, selection into entrepreneurship, and technology adoption, contribute to differences in aggregate productivity between benchmark and counterfactual. We do it using two alternative experiments. In the first one, we change ϵ^L to the level observed in Greece while fixing entry and investment policy functions to the ones in the benchmark economy, i.e. we construct a counterfactual economy keeping selection into entrepreneurship and technology adoption fixed. In the second alternative scenario, we perform the same exercise while fixing only the entry policy function from the benchmark, i.e. we construct a counterfactual economy keeping only selection into entrepreneurship at the benchmark level. Table C.4 in Appendix C.4 reports the full outcomes of this decomposition.

Around 6 percent of aggregate productivity losses induced by higher labor market power can be attributed to changes in employment allocation, keeping innovation policy and selection into entrepreneurship unchanged. About 85 percent of losses in TFP are explained by the distortion to innovation policy, while the remaining 9 percent can be attributed to distorted selection into entrepreneurship. These findings bridge the gap between the estimates of the cost of labor market power (Berger et al., 2022; Amodio et al., 2025b; Armangué-Jubert et al., 2025) and the dynamic inefficiency cost studied in models of misallocation (Restuccia and Rogerson, 2017; Guner et al., 2018; Guner and Ruggieri, 2022). By altering investment and firm growth, our results suggest that losses from labor market power may be greater than those estimated by previous studies that focus solely on the static labor allocation effect.

7 Conclusion

This paper studies how labor market power affects firm dynamics and aggregate efficiency across countries. We do it using a general equilibrium model of labor market monopsony featuring endogenous entrepreneurial choice, dynamic firms' technology adoption, and taste shocks for employers that endow them with wage-setting power. Calibrated to the Netherlands, the model reproduces cross-country differences in firm growth, technology adoption, and firm age structure through changes in labor supply elasticity.

Through the lens of the model, we find that the differences in labor market power can explain 25 percent of differences aggregate productivity across middle- and high-income countries, and no less than 11 percent over multiple robustness checks. Moreover, 85 of the explained dispersion can be attributed to lower innovation rates. Losses from labor market power may be greater than those estimated by previous studies that focus solely on static labor allocation effects.

References

Aga, Gemechu and David Francis, "As the market churns: productivity and firm exit in developing countries," *Small Business Economics*, 2017, *49*, 379–403.

Amodio, Francesco, Emanuele Brancati, Nicolás De Roux Uribe, and Michele Di Maio, "Labor Market Institutions and Wage-Setting Power: Evidence from Latin America and the Caribbean," *Unpublished Manuscript*, 2025. Retrieved from https://repositorio.uniandes.edu.co/entities/publication/1539a754-cc3b-43d9-b83d-1b754a02ad75.

_ , _ , Peter Brummund, Nicolás de Roux, and Michele Di Maio, "Global Labor Market Power," *Unpublished Manuscript*, 2024. Retrieved from https://www.econstor.eu/bitstream/10419/282656/1/dp16529.pdf.

_ , Pamela Medina, and Monica Morlacco, "Labor market power, self-employment, and development," *American Economic Review*, 2025, *115* (9), 3014–3057.

Ananian, Sévane and Giulia Dellaferrera, "Employment and wage disparities between rural and urban areas," *ILO Working Paper*, 2024, (107).

Anderson, Simon P, André De Palma, and J-F Thisse, "A representative consumer theory of the logit model," *International Economic Review*, 1988, pp. 461–466.

Ando, Sakai, "Size-dependent policies and risky firm creation," *Journal of Public Economics*, 2021, *197*, 104404.

Armangué-Jubert, Tristany, Nezih Guner, and Alessandro Ruggieri, "Labor market power and development," *American Economic Review: Insights*, 2025, *7* (2), 177–195.

Autor, David, Arindrajit Dube, and Annie McGrew, "The unexpected compression: Competition at work in the low wage labor market," Technical Report, National Bureau of Economic Research 2023.

Azar, José A, Steven T Berry, and Ioana Marinescu, "Estimating Labor Market Power," Technical Report, National Bureau of Economic Research 2022.

Bachmann, Rüdiger, Christian Bayer, Heiko Stüber, and Felix Wellschmied, "Monopsony makes firms not only small but also unproductive: Why East Germany has not converged," *CEPR Discussion Paper No. DP17302*, 2022.

Bartelsman, Eric J, John Haltiwanger, and Stefano Scarpetta, "Microeconomic evidence of creative destruction in industrial and developing countries," Technical Report, IZA Discussion Papers 2004.

Bassier, Ihsaan, "Firms and inequality when unemployment is high," *Journal of Development Economics*, 2023, *161*, 103029.

Ben-Akiva, Moshe, Nicolaos Litinas, and Koji Tsunokawa, "Continuous spatial choice: the continuous logit model and distributions of trips and urban densities," *Transportation Research Part A: General*, 1985, *19* (2), 119–154.

Bento, Pedro and Diego Restuccia, "Misallocation, establishment size, and productivity," *American Economic Journal: Macroeconomics*, 2017, *9* (3), 267–303.

Berger, David, Kyle Herkenhoff, and Simon Mongey, "Labor Market Power," *American Economic Review*, 2022, *112* (4), 1147–93.

Breza, Emily and Supreet Kaur, "Labor Markets in Developing Countries," *Annual Review of Economics*, 2025, *17.*

Brooks, Wyatt J, Joseph P Kaboski, Illenin O Kondo, Yao Amber Li, and Wei Qian, "Infrastructure Investment and Labor Monopsony Power," *IMF Economic Review*, 2021, *69*, 470–504.

—, —, Yao Amber Li, and Wei Qian, "Exploitation of Labor? Classical Monopsony Power and Labor's Share," *Journal of Development Economics*, 2021, *150*, 102627.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline, "Firms and Labor Market Inequality: Evidence and Some Theory," *Journal of Labor Economics*, 2018, *36* (S1), S13–S70.

Coşar, A Kerem, Nezih Guner, and James Tybout, "Firm dynamics, job turnover, and wage distributions in an open economy," *American Economic Review*, 2016, *106* (3), 625–663.

Da-Rocha, José-María, Diego Restuccia, and Marina M Tavares, "Policy distortions and aggregate productivity with endogenous establishment-level productivity," *European Economic Review*, 2023, *155*, 104444.

Datta, Nikhil, "Local Monopsony Power," *Unpublished Manuscript*, 2022. Retrieved from https://eprints.lse.ac.uk/126772/1/dp2012.pdf.

Deb, Shubhdeep, Jan Eeckhout, Aseem Patel, and Lawrence Warren, "What Drives Wage Stagnation: Monopsony or Monopoly?," *Journal of the European Economic Association*, 2022, *20* (6), 2181–2225.

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp Vom Berge, "Reallocation Effects of the Minimum Wage," *The Quarterly Journal of Economics*, 2022, *137* (1), 267–328.

Eslava, Marcela, A García-Marín, and Julian Messina, "Market Power and Development," *Unpublished Manuscript, Colombia: Universidad de Los Andes, Bogotá*, 2024.

—, John C Haltiwanger, and Alvaro Pinzón, "Job creation in Colombia vs the US:"up or out dynamics" meets "the life cycle of plants".," Technical Report, National Bureau of Economic Research 2019.

Farrokhi, Farid, Ahmad Lashkaripour, and Heitor S. Pellegrina, "Trade and technology adoption in distorted economies," *Journal of International Economics*, 2024, *150*, 103922.

Felix, Mayara, "Trade, labor market concentration, and wages," *Unpublished Manuscript*, 2022. Retrieved from https://www.mayarafelix.com/papers/Felix$_JMP.pdf$.

Garcia-Louzao, Jose and Alessandro Ruggieri, "Labor Market Competition and Inequality," *Unpublished Manuscript.*, 2023. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract$_id = 4557628$.

Garicano, Luis, Claire Lelarge, and John Van Reenen, "Firm size distortions and the productivity distribution: Evidence from France," *American Economic Review*, 2016, *106* (11), 3439–3479.

Gourio, François and Nicolas Roys, "Size-dependent regulations, firm size distribution, and reallocation," *Quantitative Economics*, 2014, *5* (2), 377–416.

Guner, Nezih and Alessandro Ruggieri, "Misallocation and Inequality," *CEPR Discussion Paper*, 2022, (DP17113).

_ , Andrii Parkhomenko, and Gustavo Ventura, "Managers and productivity differences," *Review of Economic Dynamics*, 2018, *29*, 256–282.

_ , Gustavo Ventura, and Yi Xu, "Macroeconomic Implications of Size-Dependent Policies," *Review of Economic Dynamics*, 2008, *11* (4), 721–744.

Hollweg, Claire H, Daniel Lederman, Diego Rojas, and Elizabeth Ruppert Bulmer, *Sticky feet: how labor market frictions shape the impact of international trade on jobs and wages*, World Bank Publications, 2014.

Hsieh, Chang-Tai and Peter J. Klenow, "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 2009, *124* (4), 1403–1448.

_ and _ , " The Life Cycle of Plants in India and Mexico *," *The Quarterly Journal of Economics*, 05 2014, *129* (3), 1035–1084.

Klenow, Peter J and Huiyu Li, "Entry costs rise with growth," *Journal of Political Economy Macroeconomics*, 2025, *3* (1), 43–74.

Lederman, Daniel and William F Maloney, "R&D and development," *Unpublished Manuscript.*, 2003. Retrieved from https://papers.ssrn.com/sol3/papers.cfm?abstract$_i d = 402480$.

Luccioletti, Claudio, "Labor Market Power Across Cities," Technical Report, CEMFI Working Paper 2214. 2022.

Maestas, Nicole, Kathleen J Mullen, David Powell, Till Von Wachter, and Jeffrey B Wenger, "The value of working conditions in the United States and implications for the structure of wages," *American Economic Review*, 2023, *113* (7), 2007–2047.

Manning, Alan, "The plant size-place effect: agglomeration and monopsony in labour markets," *Journal of economic geography*, 2010, *10* (5), 717–744.

Mas, Alexandre and Amanda Pallais, "Valuing Alternative Work Arrangements," *American Economic Review*, December 2017, *107* (12), 3722–59.

McFadden, Daniel, "The revealed preferences of a government bureaucracy: Empirical evidence," *The Bell Journal of Economics*, 1976, pp. 55–72.

Newey, Whitney K and Daniel McFadden, "Large sample estimation and hypothesis testing," *Handbook of econometrics*, 1994, *4*, 2111–2245.

Ogloblin, Constantin and Gregory Brock, "Wage determination in urban Russia: Underpayment and the gender differential," *Economic Systems*, 2005, *29* (3), 325–343.

Pham, Hoang, "Trade reform, oligopsony, and labor market distortion: Theory and evidence," *Journal of International Economics*, 2023, *144*, 103787.

Restuccia, Diego and Richard Rogerson, "Policy Distortions and Aggregate Productivity with Heterogeneous Establishments," *Review of Economic Dynamics*, 2008, *11* (4), 707–720.

_ **and** _ , "The causes and costs of misallocation," *Journal of Economic Perspectives*, 2017, *31* (3), 151–174.

Sockin, Jason, "Show Me the Amenity: Are Higher-Paying Firms Better All Around?," *Unpublished Manuscript*, 2024. Retrieved from https://www.researchgate.net/profile/Jason-Sockin-2/publication.

Sokolova, Anna and Todd Sorensen, "Monopsony in Labor Markets: A Meta-analysis," *ILR Review*, 2021, *74* (1), 27–55.

Sorkin, Isaac, "Ranking Firms Using Revealed Preference*," *The Quarterly Journal of Economics*, 01 2018, *133* (3), 1331–1393.

Stephan, Ute, Karoline Strauss, Marjan J Gorgievski, and Dominika Wach, "How entrepreneurs influence their employees' job satisfaction: The double-edged sword of proactive personality," *Journal of Business Research*, 2024, *174*, 114492.

Tamkoç, M Nazım and Gustavo Ventura, "Rules and Regulations, Managerial Time and Economic Development," Technical Report, The World Bank 2024.

Verboven, Frank, "The nested logit model and representative consumer theory," *Economics Letters*, 1996, *50* (1), 57–63.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein, "Monopsony in the US labor market," *American Economic Review*, 2022, *112* (7), 2099–2138.