Remote Work and Labor Market Power*

Tito Boeri† Peter John Lambert‡
Bernardo Mottironi§ Paolo Naticchioni¶

March 31, 2026

Preliminary and incomplete

[latest version]

Abstract

This paper examines the causal impact of the post-pandemic transition to remote work on labor market power, connecting to an extensive literature which highlights distaste for commuting as a key source of rent extraction for local employers. Our analysis combines three sources of microdata: (i) millions of job vacancy postings to derive an index of remote work using a frontier large language model; (ii) worker-level administrative records on the workplace and occupational group; and (iii) firm-level production data to estimate wage markdowns. Exploiting the heterogeneity in shifts to remote work across occupations — instrumented with shifts observed abroad — and the heterogeneity in occupation mixes across local labor markets, we implement a shift-share design and identify a substantial reduction in employer rents. Our counterfactual analysis quantifies the impact of the 2019-2023 transition to remote work in one-fifth of the deviation from perfect competition.

Keywords: work from home, monopsony, wage markdowns, job vacancies, shift share
JEL Codes: J31, J42, J63, R23, R32

*We thank John Van Reenen, Attila Lindner, and Chad Syverson for their insightful comments. We are grateful to INPS, and in particular to the staff of its Research Department (DCSR), for invaluable support through the VisitINPS program. This project has received funding from Horizon Europe (ERC grant 101097687) as well as NextGenerationEU (PRIN 2022JYY7SA). The views expressed do not necessarily reflect those of the Bank of Italy or INPS.

†Bocconi University & Centre for Economic Performance
‡University of Warwick & London School of Economics
§Bank of Italy & Centre for Economic Performance
¶Roma Tre University

1 Introduction

Following Manning (2003), an important theoretical and empirical literature has emphasized that labor markets are largely monopsonistic and firms are typically able to set wages below the competitive level.[1] Monopsony power is often explained in terms of competition with limited substitutability between differentiated employers, arising from the idiosyncratic taste for different workplaces (Card et al., 2018). Chief among these sources of horizontal differentiation is location proximity, as workers dislike commuting costs and times (Manning & Petrongolo, 2024). This mechanism shields local employers from competition, allowing them to exert market power, especially in less connected areas. Considering the key role of commuting for monopsony power, a natural question is whether the unprecedented adoption of remote work, following the COVID-19 pandemic, has produced sizable changes in labor market competition.

Remote work arrangements have indeed persisted long after the forcing event subsided (Hansen et al., 2023).[2] The persistence of work-from-home (WFH) reflects both the value workers place on flexibility (Cullen et al., 2025; Mas & Pallais, 2017), as well as the benefits of relaxing geographic constraints and expanding the feasible set of jobs a worker might be willing to accept (Akan et al., 2025; Manning & Petrongolo, 2017). If the diffusion of remote work enlarges the set of appealing matches, then firms must compete for workers in a broader market, which could raise the elasticity of labor supply they face, and consequently, reduce the wedge betweem the marginal revenue product of labor (MRPL) and labor costs per worker.[3] At the same time, WFH represents a job amenity over which workers may have idiosyncratic preferences.[4] Heterogeneity in its adoption may generate further horizontal differentiation across employers, reducing competition. Thus, we need empirical guidance to infer which of the two mechanisms — a declining relevance of the geographical differentiation vs a novel amenity increasing differentiation — dominates.

We look at the relationship between WFH and monopsony power in Italy, combining three complementary data sources. First, we measure the diffusion of WFH using millions of job

[1] Examples include Azar et al. (2020), Dube et al. (2020), Bassier et al. (2022), Berger et al. (2022), Deb et al. (2022), Lamadon et al. (2022), Yeh et al. (2022), Goolsbee and Syverson (2023), Kroft et al. (2025). See Card (2022) for a summary and discussion.

[2] For the United States, Barrero et al. (2023) document that working from home remains a large share of paid workdays in the post-pandemic period and argue that this persistence reflects durable changes in technology, managerial practices, and worker preferences. Cross-country evidence on persistence appears in Zarate et al. (2025).

[3] Under static profit maximization, the percentage difference between MRPL and labor costs per worker equals the elasticity of the inverse labor supply.

[4] Consider, for instance, differences in housing conditions across workers.

vacancy postings, analysing the posting text with frontier large language model (LLM) processing methods and classifying whether jobs offer fully remote or hybrid arrangements. These measures are then aggregated by occupational group and year, to obtain occupation-specific WFH shifts that occurred during the pandemic. Second, we gain access to administrative records on worker-level characteristics to construct the pre-pandemic employment shares of each occupation across local labor markets, defined as interactions of commuting zones and industries. Third, we measure monopsony power at the firm-level by performing a parametric estimation of MRPL from the financial statements of Italian corporations, from which we derive net wage markdowns — henceforth markdowns — defined as percentage differences between MRPL and labor costs per worker. Markdowns are then aggregated by cell and year, to obtain time differences in the departure from perfect competition across local labor markets.

The 2019–2023 variations in aggregate markdowns constitute our dependent variable. Our empirical strategy exploits heterogeneity in the post-pandemic diffusion of remote work across local labor markets from two sources of variation: (i) the diverse capacity to carry out tasks remotely across occupations and (ii) the diverse occupation mix across commuting zones and industries. We construct a shift-share index of exposure to remote work by interacting the local occupational shares in 2019 with the national occupation-specific WFH shifts that occurred between 2019 and 2023.

A legitimate concern is that national shifts in WFH may be correlated at the occupation-level with other contemporaneous shocks. To address this potential source of endogeneity, we leverage the idea that the diffusion of remote work is partly driven by technology and task content transcending national boundaries. We process job postings from the United States with the same LLM applied to Italian data and instrument WFH shifts in Italy with WFH shifts observed in the US. A further advantage in this IV strategy lies in the superior accuracy of measurement based on US data due to sample size, which is larger by an order of magnitude. Additional identification issues are dealt with by specifying control variables and standard errors in accordance with the most recent guidance on the econometrics of shift-shares (Borusyak et al., 2025).

Our core finding is that WFH increases competition in the labor market reducing aggregate markdowns. The inferred impact is substantial in magnitude and robust to several specifications. An increase of one interquantile range in the WFH index lowers the aggregate markdown by nearly 6 basis points. Furthermore, using the same strategy we find a substantial effect on wages, which increase by 4%, consistent with a more competitive labor market. The counterfactual analysis considers scenarios with different diffusion of remote

work. According to our estimates, without the WFH shifts occurring between 2019 and 2023, the country-level markdown in 2023 would be approximately one fifth higher, 0.36 instead of 0.29. Comparing the same industries across different commuting zones, our results reveal that, under the WFH shifts observed in major cities such as Rome or Milan, the aggregate markdown would be 0.26 or 0.25, respectively — a further drop by nearly one tenth.

Market-level results are then corroborated by tests at the firm-level and worker-level, confirming the negative effect on markdowns and the positive effect on wages. Worker-level analysis allows us to uncover an additional element in the mechanism. We document that exposure to WFH leads to a higher probability of mismatch in the commuting zones of the worker's residence and office, supporting the hypothesis of a pro-competitive effect due to an expansion of geographical constraints.

Related literature. Our work contributes to two literatures that have largely developed in parallel: research on labor market power and research on remote work. Growing evidence using a variety of approaches has established the presence of non-trivial employer power in labor markets. For instance, Bassier et al. (2022) use a matched event study to measure the elasticity of employer-employee separations to the wage, from which they obtain a net markdown of 0.24 in their preferred specification. Yeh et al. (2022) estimate plant-level production functions across US manufacturing and derive an average net markdown of 0.53, with an upward trend since 2002. Deb et al. (2022) instrument labor demand shocks with time-series variation in taxes and find a range spanning between 0.32 and 0.53, with a more muted time trend. Berger et al. (2022) develop and structurally estimate a model of oligopsonistic competition, obtaining an aggregate net markdown equal to 0.39, which generates a 21% output loss through resource misallocation.

Furthermore, monopsony power has been shown to explain several features of the labor markets, such as the elusive effect of minimum wages on employment (Dustmann et al., 2021; Manning, 2021), between-firm inequality in wages (Card et al., 2018) falling labor shares (Mertens & Mottironi, 2025) or urban wage premia (Hirsch et al., 2022). Regarding the roots of imperfect competition, recent research has emphasised the role of concentration in hiring flows (Azar et al., 2020, 2022), collusion against workers (Delabastita & Rubens, 2025), non-compete agreements (Alves et al., 2024; Boeri et al., 2025), search frictions (Berger et al., 2024; Burdett & Mortensen, 1998), and commuting distance (Datta, 2024; Manning & Petrongolo, 2024). This paper speaks to the latter factor, providing novel evidence that, by softening the need to commute — along either the extensive or intensive margin, in fully

remote or hybrid arrangements respectively — the diffusion of remote work can increase competition in labor markets.

In parallel, an extensive body of research documents that remote work is persistent, its feasibility varies sharply across tasks, and its adoption differs across places (Althoff et al., 2022; Barrero et al., 2023; Dingel & Neiman, 2020). A central theme is that remote work is a non-wage amenity with heterogeneous valuations (Mas & Pallais, 2017; Song & Gao, 2018). Remote work also has organisational and productivity consequences that vary across settings and interact with selection, coordination, and collaboration mechanisms (Bloom et al., 2015; Emanuel & Harrington, 2024; Gibbs et al., 2023; Yang et al., 2022). Recent work has also considered its impact on worker mobility (Liu & Su, 2025), migration (Bick et al., 2024), and commuting (Monte et al., 2023).

Our analysis extends this literature by uncovering an overlooked channel, contributing to a broader understanding of the labor market effects of flexible work arrangements. The welfare-enhancing effects of remote work may not be limited to the direct effect on workers who experience greater flexibility. By enlarging the size of labor markets, remote work can change the spatial mobility and power dynamics between workers and firms, improving competition. These findings bear significant implications for policy, as the adoption of remote work, possibly combined with the supply of remote working locations to those who have limited work opportunities at home, could be encouraged as a tool to discipline monopsony power, mitigating unequal labor market opportunities across space and improving overall market allocation.

Outline. The remainder of this paper is organised as follows. Section 2 illustrates the diverse data sources employed in this project. Section 3 outlines our measurement methodology, including novel LLM tools for parsing remote work arrangements from job ads and markdown estimation from financial statements. Section 4 provides the details of our research design and how it relates to the recent literature on shift-share instruments. Section 5 reports the results of the 2SLS estimation and the counterfactual exercise. Section 6 concludes.

2 Data

2.1 Job postings

Our analysis employs over 20 million online job vacancy postings from Italy and nearly 200 million from the United States, collected during the period 2018–2024. Italian job postings

cover all NUTS2 regions as well all two-digit ISCO occupations but the military. This database is produced by Lightcast, a company specialised in labor market analytics, successor brand of Burning Glass Technologies. Lightcast collects vacancy advertisements from a large set of online job boards and employer career pages, and then converts the unstructured content into a harmonised statistical dataset through a pipeline that includes de-duplication, entity resolution, and automated coding of occupations, industries, and locations (Lightcast, 2022).

Table 2.1 reports the number of online job vacancy postings for which we could retrieve reliable information on the ISCO classification, before and after the pandemics, i.e. the two periods considered in the empirical strategy discussed in Section 4. The US database shows a substantially larger sample size, which implies more accurate estimates. This advantage leads us instrumenting our index of WFH exposure based on Italian postings with the same index computed using the US database (see Section 4). The yearly count of postings in the full sample is reported in Appendix A.

Table 2.1: Count of LLM-classified job postings

Year	Italy	United States
2019	1,097,174	28,140,958
2023	3,511,278	32,322,852

Note: This table reports the sample size of job postings for which we could retrieve reliable information on the ISCO classification, by year and country. Source of the raw data: Lightcast.

Although online postings are not equivalent to the full universe of vacancies in the economy, they provide uniquely timely and text-rich signals of labor demand and advertised job attributes, which is precisely the dimension we exploit in constructing vacancy-based measures of remote work.

2.2 Social security records

Our analysis requires to combine the occupation-level measure of 2019–2023 variations in remote work adoption with the 2019 employment share of each occupational group (three-digit ISCO) across Italian local labor markets, which we define as interactions of commuting zones ('Sistemi Locali del Lavoro', from ISTAT) and two-digit industries (NACE rev. 2).

Because this disaggregation is not present in publicly available datasets, we get access to Italian social security records through the VisitINPS program. These administrative data include information on worker-level characteristics such as the occupational group, education,

working days, workplace location, and the employer's industry.[5] In particular, we merge two datasets: (i) UniLav, which collects 'compulsory communications' on the beginning, extension, transformation, and conclusion of employment contracts — including, crucially, information on the specific occupation held by the employee as well as the education level — and (ii) Uniemens, which provides yearly information on remuneration, working days, workplace location and industry for all employees.[6]

The resulting database is a cross-section of employer-employee-location matches active in 2019. For job relationships that appear as active in Uniemens but do not show a corresponding communication in UniLav, we impute occupational groups using the most recent compulsory communication made for the same worker by another employer. Our methodology classifies the occupational group of nearly 90% of private sector workers with an active record in 2019. Furthermore, the availability of information on working days allows us to derive the employment shares in full-time equivalent (FTE) units, taking into account potential heterogeneity in job discontinuity and part-time employment across occupations.

2.3 Financial statements

To perform our estimations, we use financial statements of Italian corporations from the Orbis database (AIDA), covering the near-universe of incorporated companies in 2015–2023. Table 2.2 reports summary statistics for the full sample of firms with at least one employee.

Table 2.2: Firm-level summary statistics

	Median	Mean	SD	Obs.
Employment	5	17.9	235.2	4,828,357
Value added per employee	44.17	65.71	323.81	4,828,357
Labor costs per employee	25.56	28.34	92.32	4,828,357
Capital per employee	8.86	107.8	1,211.62	4,828,357

Note: This table reports the median, mean, standard deviation, and number of observations for firms with at least one empoloyee, based on the full 2015-2023 sample, before implementing the data cleaning routine. Monetary values are expressed in thousands of constant Euros (2020). Source of the raw data: AIDA.

Average employment per firm is above the value observed in the universe, partly because we exclude firms with no employee and partly because our firm-level database only includes

[5]Further details are available at https://www.inps.it/it/it/dati-e-bilanci/attivit--di-ricerca/programma-visitinps-scholars.html.

[6]The same data are used by INPS to compute social security contributions.

incorporated companies. As for most papers in the related literature, the reason for choosing this sample is that unincorporated companies are not legally obliged to produce a statutory balance sheet, which is the key source of information for the variables used in our estimations. Nevertheless, our database has a large coverage of micro firms, as shown by a median employment equal to five. Value added is defined by subtracting raw materials (net of changes in the corresponding inventory), purchased goods for resale, and intermediate services from the total value of production. Our sample includes firms with negative value added, which amounts to nearly 5%. Notice that value added is slightly less than twice as large as labor costs per employee, in line with aggregate statistics reported by ISTAT, the source of official statistics in Italy. Capital is defined as tangible fixed assets, net of the stock held for sale.

3 Measurement

3.1 Remote work

Methodology. Our primary use of the job postings is to measure the prevalence of remote and hybrid arrangements by occupational group, based on the textual content of each observation. The Lightcast database does not directly include structured information on the use of remote work, so we extend the methodology of Hansen et al. (2023) and apply an LLM to extract this feature and classify each job posting.

The exercise is carried out in three steps: first, we identify which postings contain any information at all about remote or hybrid work arrangements; second, focusing on the positive cases from the previous step, we identify the actual proposal of remote work, being careful to distinctly tag cases in which remote work is mentioned by being explicitly disallowed; finally, we evaluate the performance of the algorithm using a sample of 2,000 hand-classified postings.[7] Further details are available in Appendix A.

The results of our performance assessment are reported in Table 3.1. The F-score pooling the two years of interest is nearly 96%, showing high consistency between LLM-based and human-based classifications. Auditing points to Type 1 errors largely exceeding Type 2 errors in 2019, based on a subsample that equally represents remote and non-remote arrangements. This difference in the two types of errors implies a vastly superior performance in the population estimates as there are significantly fewer remote arrangements in 2019.

To derive the shift-share index of remote work adoption at the market-level, after classifying job postings by the availability of remote work we compute the occupation-specific shifts in

[7] Our manual classification is available in the replication package.

Table 3.1: LLM PERFORMANCE METRICS

Metric	LLM classification	2019	2023	Pooled
Recall	WFH = 0	100%	99.6%	99.8%
Precision	WFH = 1	88%	96.4%	92.2%
F-score	Harmonic mean	93.6%	98%	95.8%

Note: Table 3.1 reports the confusion matrix assessing the performance of our LLM algorithm, based on human auditing of a stratified subsample, i.e. manual classification of 1,000 postings, equally distributed by LLM classification and year. Recall (Precision) reports the percentage of observations in which the negative (positive) LLM classification is confirmed. The F-score is the harmonic mean of Recall and Precision. Source of the raw data: Lightcast.

the share of WFH jobs occurred during the pandemic. Such 2019-2023 WFH shifts (ΔS_o, with o denoting a three-digit occupational group) are then combined with 2019 shares of workers in the corresponding occupation ($L_{o,c}$, in FTE units) within each labor market c, which we define as the interaction of a commuting zone and a two-digit industry. The resulting shift-share index (ΔWFH_c) constitutes the key regressor in our baseline specification.

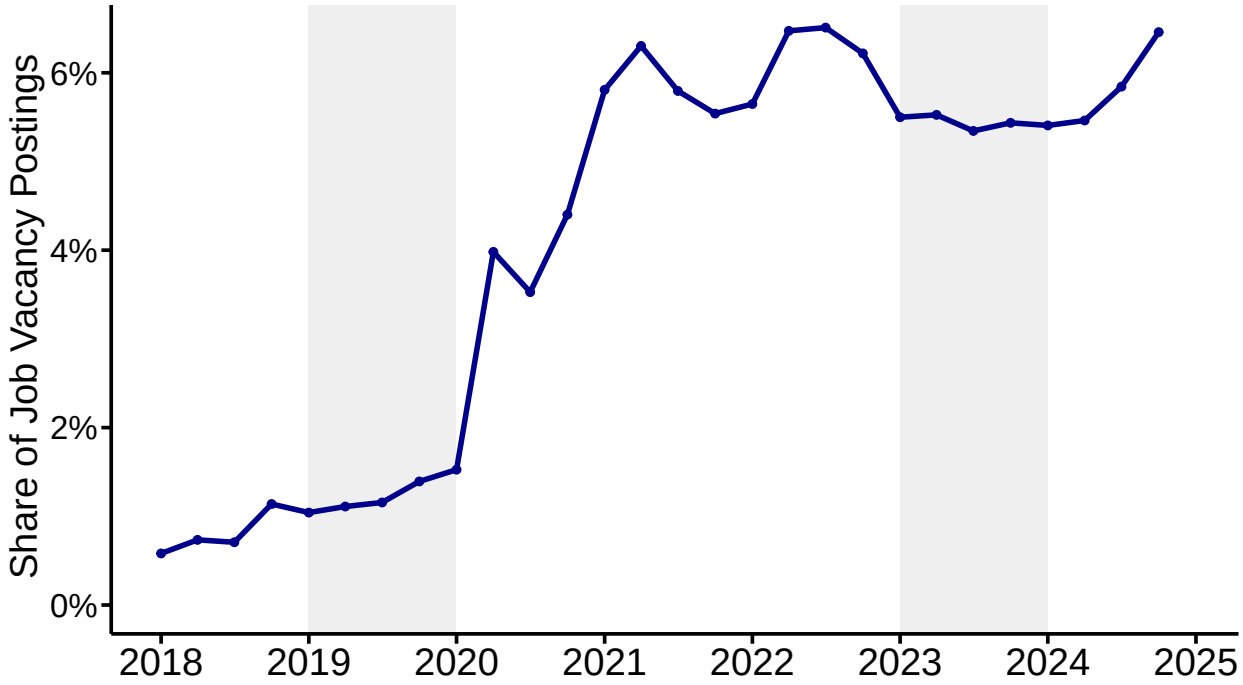
$$\Delta WFH_c \equiv \frac{\sum_o \Delta S_o L_{o,c}}{\sum_o L_{o,c}} \tag{3.1}$$

Descriptives. The rapid rise in remote work practices is now a well-documented fact across the developed world. We provide additional evidence that Italy, one of the earliest countries to be hit by the pandemic, has also experienced a substantial rise in the share of job ads explicitly offering such practices. Figure 3.1 reports the share of job postings offering remote work arrangements by quarter. As our measure is a lower bound for WFH and is confined to job posting on platforms, it should not come as a surprise the fact that the Italian Labor Force Survey points to a larger share of jobs offering WFH arrangements (see Appendix A.1). The trend displayed by our estimates is in line with LFS data: we observe a mild upward trend in 2018–2019, a rapid burst in 2020–2021, and fluctuations around a flat trend afterwards.[8]

Figure 3.2 documents this diffusion across terciles of occupational groups based on the degree of WFH adoption. The 'low WFH' occupations see almost no rise in such arrangements, whereas the 'high WFH' occupations see an almost six-fold increase since the onset of COVID-19. Teaching and ICT professionals are observed at the top of the WFH list. The lowest level of remote work offers is found in food-preparation and construction workers. In our main

[8]This magnitude also aligns with measures from other countries reported by Hansen et al. (2023).

Figure 3.1: WFH DIFFUSION IN ITALY

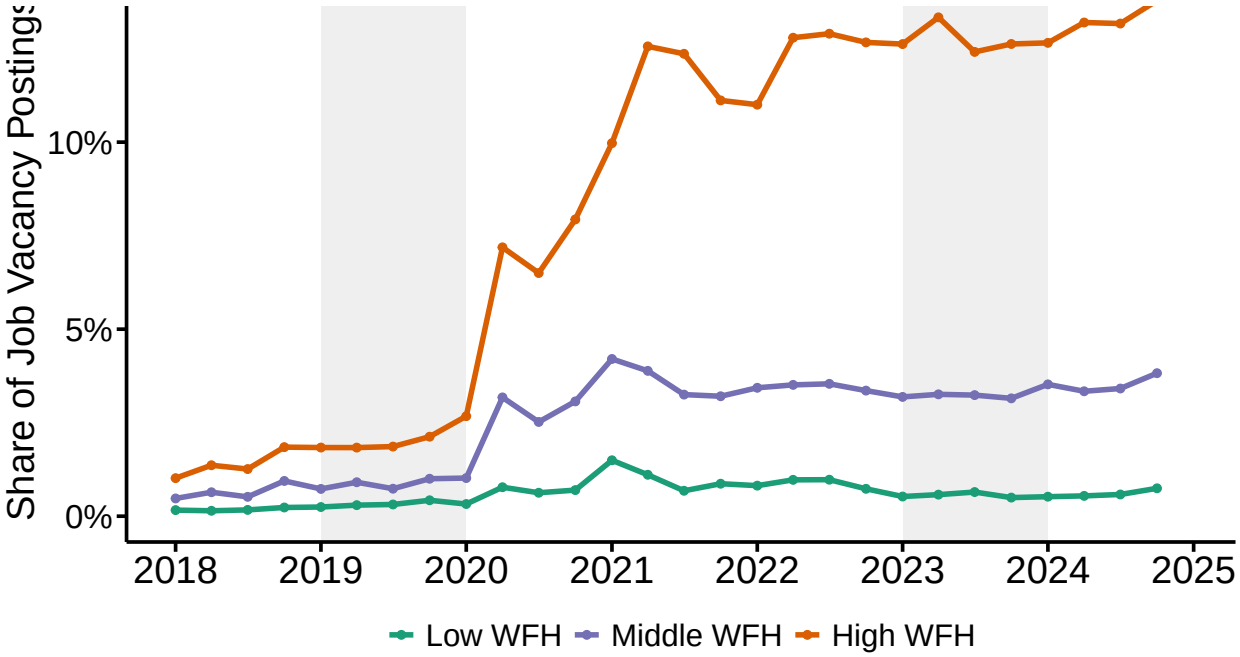


Note: This figure reports the share of online job vacancy postings explicitly offering remote arrangements by quarter. The sample includes 13,616,201 observations in total, where both an occupation code was present and raw textual description allowed for a classification of remote work using our LLM classification algorithm. Quarters in 2019 and 2023 are shaded as these subsets are used in our econometric application. Source of the raw data: Lightcast.

empirical exercise, we exploit this variation between occupational groups to define exposure to the rise of remote working across Italian local labor markets.

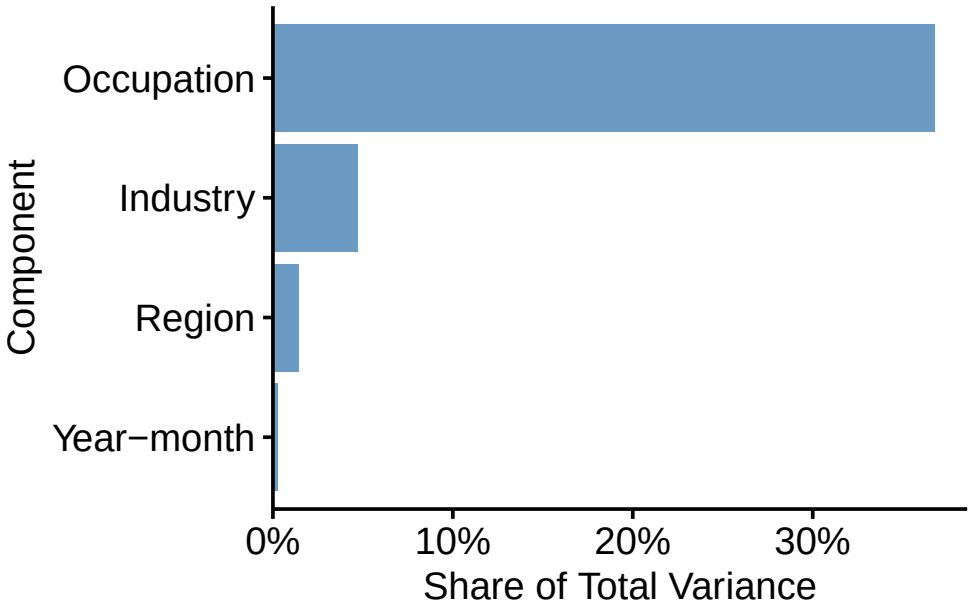
To ensure that our main empirical design exploiting occupational variation is sensible, we conduct a variance decomposition exercise using alternative dimensions of the job vacancy posting, namely regions, industries, and periods. We run a fixed-effects regression with these dimensions along with the occupational groups. Figure 3.3 reports the variance that each dimension explains, showing that occupation is by far the dominant source of variation in remote work opportunities. We see that the job ads occupation explains 38% of total variation, whereas industry and region have less explanatory power after conditioning on the occupation.

Figure 3.2: WFH HETEROGENEITY ACROSS OCCUPATIONS



Note: This figure reports the share of online job vacancy postings explicitly offering remote arrangements, broken into three groups of occupations, by quarter. The low, middle, and high WFH groups are defined using the share of job postings offering remote arrangements in 2023, which we break into terciles of around 14 two-digit ISCO occupations. The sample includes 13,616,201 observations in total, where both an occupation code was present and raw textual description allowed for a classification of remote work using our LLM classification algorithm. Quarters in 2019 and 2023 are shaded as these subsets are used in our econometric application. Source of the raw data: Lightcast.

Figure 3.3: WFH VARIANCE DECOMPOSITION



Note: This figure reports the variance decomposition of remote work calculated using fixed effects for (i) two-digit ISCO occupations, (ii) NUTS-2 regions, (iii) ISIC letters, and (iv) year-months from January 2022 to December 2024. Regression is fit across 281,482 cells defined using the interaction of these four dimensions. Each cell is weighted by the count of postings. Source of the raw data: Lightcast.

3.2 Labor market power

Methodology. We measure the degree of labor market power of firm f in period t by estimating its wage markdown ($\gamma_{f,t}$), defined as the percentage difference between its marginal revenue product of labor ($MRPL_{f,t}$) and labor costs per worker ($w_{f,t}$). Markdowns are commonly used to assess competition in labor markets as their theoretical value is null under perfect competition — one if expressed in gross terms (i.e. $1 + \gamma_{f,t}$) — but can be higher with monopsony power and lower in the presence of strong unions.[9] By simply decomposing $MRPL_{f,t}$ into the revenue elasticity of labor times revenues per worker, markdowns can be expressed in terms of such elasticity ($\varepsilon^{R,L}_{f,t}$), revenues ($R_{f,t}$), and labor costs ($w_{f,t}L_{f,t}$).[10]

$$\gamma_{f,t} \equiv \frac{MRPL_{f,t} - w_{f,t}}{w_{f,t}} = \frac{\varepsilon^{R,L}_{f,t} R_{f,t} - w_{f,t}L_{f,t}}{w_{f,t}L_{f,t}} \tag{3.2}$$

Revenues and labor costs are directly observable in financial statements, whereas the revenue elasticity of labor requires a parametric estimation. We implement the latter by estimating revenue functions with a translog specification in capital, labor, and intermediate inputs, denoted by $F_i(k_{f,t}, l_{f,t}, m_{f,t})$, using small letters for logarithmic transformations and subscript i for a four-digit industry, the level at which the estimation is performed.[11] The form of the firm-level revenue function is shown in Equation 3.3, denoting the parameter interacted with factor x as $\alpha_{i,x}$ and decomposing the residual into a persistent revenue shock ($\omega_{f,t}$) and a measurement error ($\epsilon_{f,t}$).[12]

$$\begin{aligned} r_{f\in i,t} &= F_i(k_{ft}, l_{ft}, m_{ft}) + \omega_{ft} + \epsilon_{ft} \\ &= \alpha_i, 1 + \alpha_{i,k}k_{f,t} + \alpha_{i,l}l_{f,t} + \alpha_{i,m}m_{f,t} + \alpha_{i,k^2}k^2_{f,t} + \alpha_{i,l^2}l^2_{f,t} + \alpha_{i,m^2}m^2_{f,t} \\ &\quad + \alpha_{i,kl}k_{f,t}l_{f,t} + \alpha_{i,km}k_{f,t}m_{f,t} + \alpha_{i,lm}l_{f,t}m_{f,t} + \omega_{f,t} + \epsilon_{f,t} \end{aligned} \tag{3.3}$$

To deal with the endogeneity arising from the correlation between revenue shocks and the input demand, we adapt the 'control function approach' of Ackerberg et al. (2015) — typically employed for production function estimation — to the case of revenue functions.[13] This

[9] As discussed in footnote 3, there is a direct relationship between markdowns and the labor supply elasticity.

[10] By definition, $MRPL_{f,t} \equiv \frac{\partial R_{f,t}}{\partial L_{f,t}} = \varepsilon^{R,L}_{f,t} \frac{R_{f,t}}{L_{f,t}}$.

[11] Translog specifications approximate any twice continuously differentiable function up to second order, allowing nonlinear relationships and heterogeneity in the revenue elasticities across firms and periods.

[12] In standard production function applications, there is physical output — rather than revenues — on the left-hand side, and $\omega_{f,t}$ corresponds to the persistent productivity process (e.g. Olley & Pakes, 1996). Instead, in our framework, $\omega_{f,t}$ is not limited to the supply side but also includes revenue shocks coming from the product demand.

[13] Applications of the 'control function approach' to revenue function estimation can be found in Petrin and Sivadasan (2013), Hashemi et al. (2022), and Mottironi (2024).

methodology requires to purge the measurement error before estimating the parameters. In the first step, the demand for intermediate inputs is inverted to express the persistent shock, $\omega_{f,t}$, as a function of current input usage (hence the term 'control function'). As a result, we can purge $\epsilon_{f,t}$ by simply projecting revenues on a high-order polynomial combination of inputs, which leads to a prediction of $\phi_{f\in i,t} \equiv F_i(k_{f,t}, l_{f,t}, m_{f,t}) + \omega_{f,t}$. In the second step, $\omega_{f,t}$ is assumed to follow a first-order Markov chain, i.e. $\omega_{f,t} = g_i(\omega_{f,t-1}) + \eta_{f,t}$. Under this assumption, $\phi_{f,t}$ can be expressed in the following way.

$$\phi_{f\in i,t} = F_i(k_{f,t}, l_{f,t}, m_{f,t}) + g_i\left(\phi_{f,t-1} - F_i(k_{f,t-1}, l_{f,t-1}, m_{f,t-1})\right) + \eta_{f,t} \tag{3.4}$$

By approximating $g_i(.)$ with a polynomial in its argument, it is possible to estimate the parameters of $F_i(.)$ using the lags of the inputs (as well as their translog interactions) and $\phi_{f,t-1}$ — estimated in the first step — as they are all orthogonal to $\eta_{f,t}$, the innovation term of the persistent process.[14] From the estimated parameters we obtain the revenue elasticity of labor, which can vary across firms and years with input quantities, reflecting both firm size and relative input usage.

$$\varepsilon^{R,L}_{f\in i,t} = \alpha_{i,l} + 2\alpha_{i,l^2} l_{f,t} + \alpha_{i,kl} k_{f,t} + \alpha_{i,lm} m_{f,t} \tag{3.5}$$

Finally, we insert our estimation of Equation 3.5 in the firm-level markdown estimator defined in Equation 3.2.

For robustness, we derive an alternative measure based on output elasticities, following Yeh et al. (2022). Unlike revenue elasticities, output elasticities do not capture the elasticity of product prices to input usage. Therefore, the wedge between the output elasticity and the labor share includes the price markup. To isolate the wage markdown, the approach of YMH derives the markup from the same wedge but using intermediate inputs, under the assumption that the latter are traded in perfectly competitive markets — an unnecessary assumption in our preferred revenue-based methodology.[15] Further details are discussed in Appendix B.

Firm-level markdowns are averaged using labor costs as weights, to obtain aggregate markdowns as ratios of the market-level MRPL to market-level labor costs, in the spirit of Edmond et al. (2023). We denote the aggregate markdown as $\Gamma_{c,t}$ and use this measure to

[14] Further details on the theoretical assumptions and practical implementation of our approach to revenue function estimation are discussed in Appendix B.

[15] An additional difference concerns the market-level aggregation. YMH derive the aggregate MRPL by estimating production functions at a broader level (two-digit instead of four-digit industries) and combining these translog parameters with the market-level input usage.

define the market-specific index capturing variations in labor market power occurred during the pandemic (ΔLMP_c, the dependent variable of our baseline specification).[16]

$$\Gamma_{c,t} \equiv \frac{\sum_{f \in c} \gamma_{f,t} w_{f,t} L_{f,t}}{\sum_{f \in c} w_{f,t} L_{f,t}} = \frac{MRPL_{c,t} - w_{c,t}}{w_{c,t}} \tag{3.6}$$

$$\Delta LMP_c \equiv \Gamma_{c,2023} - \Gamma_{c,2019} \tag{3.7}$$

Descriptives. Table 3.2 reports summary statistics of firm-level and market-level markdowns in 2019 and 2023, the two periods considered in our main research design. Means and medians range between 0.14 and 0.43, an interval consistent with the measures provided by the literature on monopsony power (see Section 1).

At the market-level, markdown has substantially risen from 2019 to 2023, and this increase is not attributable to outliers as medians increase more than averages. One of the most likely drivers of this change lies in the limited capacity of workers, at least in the short run or medium run, to counterbalance inflationary bursts with equal adjustments in nominal wages.[17] This argument has solid empirical and theoretical evidence, from Shiller (1996) to more recent contributions (Afrouzi et al., 2024; Stantcheva, 2024). Italian collective bargaining institutions allow for forward-looking adjustments to inflation and collective agreements are typically signed with long delays, preventing adjustments in the period between two collective agreements. Limited mobility across jobs is another factor as workers locked into old contracts have limited options to renegotiate individually their wage. This drop in compensation may last for months or even years, especially if workers have little power in wage bargaining (Mengano, 2025; Stansbury & Summers, 2020).[18]

It is worth noting that the average time change in markdowns is not per se an object of study for our analysis, as we are only interested in how these changes correlate with WFH across local labor markets. Our results are not necessarily affected by aggregate trends or shocks operating at the national level. A source of concern may instead arise from the uneven relevance of this mechanism across labor markets. We believe that our large set of controls, illustrated in Section 4, can effectively attenuate this issue.

[16]The second equality in Equation 3.6 uses $MRPL_{c,t} \equiv \frac{\sum_{f \in c} MRPL_{f,t} L_{f,t}}{\sum_{f \in c} L_{f,t}}$ and $w_{c,t} \equiv \frac{\sum_{f \in c} w_{f,t} L_{f,t}}{\sum_{f \in c} L_{f,t}}$.

[17]In 2022, the Italian CPI exceeded 8% for the first time since 1986, and Italy is the EU country with the largest fall in real wages since the inflation upsurge of August 2021 (OECD, 2025).

[18]Through the lenses of our hypothesis, the availability of remote work may increase the boundaries of labor markets and facilitate job-to-job transitions, which has been shown to play a key role in helping workers to keep up with inflation bursts (Afrouzi et al., 2024)

Table 3.2: (Net) Markdowns

Aggregation	Year	Median	Mean	SD	Obs.
Firm-level ($\gamma_{f,t}$)	2019	0.14	0.43	1.05	493,437
Firm-level ($\gamma_{f,t}$)	2023	0.19	0.4	0.84	507,468
Market-level ($\Gamma_{f,t}$)	2019	0.17	0.3	0.7	27,179
Market-level ($\Gamma_{f,t}$)	2023	0.25	0.36	0.62	27,478

Note: This table reports the median, mean, standard deviation, and number of observations of our markdown estimates in 2019 and 2023, the two periods considered in the research design. Net markdowns are defined as the percentage difference between MRPL and labor costs per worker. Firm-level estimates are mapped into market-level estimates using labor costs as weights. Labor markets are defined as interactions of commuting zones and two-digit industries. Source of the raw data: AIDA.

4 Identification

The scope of our research design is to estimate the general equilibrium effect of market-level variations in WFH adoption on market-level variations in labor market power. In our analysis, we define labor markets c as interactions of commuting zones and two-digit industries, for a total of nearly 23,000 cells.

The baseline specification is as follows:

$$\Delta LMP_c = \beta_1 + \beta_{\Delta WFH}\, \Delta WFH_c + \beta_X X_c + u_c \tag{4.1}$$

ΔLMP_c and ΔWFH_c denote the 2019-2023 variations in labor market power and WFH adoption, respectively. As discussed in Section 3, labor market power is measured by aggregating firm-level wage markdowns, whereas the WFH index is derived as a shift share index, combining occupation-specific WFH shifts with the predetermined occupation shares.

The recent econometric literature highlights that identification in shift-share research designs requires an assumption of exogeneity either in the shares or the shifts (Borusyak et al., 2025), conditional on control variables. Accordingly, we select controls in X_c that allow us to claim the exogeneity of the shifts, i.e. conditional on controls, shifters ΔS_o can be considered as good as randomly assigned. To net out potential confounding factors, the set of controls include the 2019 values of following market-level variables: (i) share of high-skilled workers; (ii) share of college-educated workers; (iii) average age of workers; (iv) aggregate value added per worker (in logs); (v) aggregate labor costs per worker (in logs); (vi) number

of workers per firm (in logs); (vii) number of workers in FTE units (in logs); (viii) lag of the dependent variable (i.e. the 2015–2019 variation in the aggregate markdown).

The share of high-skilled workers is defined using the first three categories of the one-digit ISCO classification (i.e. 1. Managers; 2. Professionals; and 3. Technicians and associate professionals). As our claim on the exogeneity of the shifts is conditional on controls, WFH shifters need to be as good as randomly assigned only within the same skill level. By including this control variable, we ensure that the identification exploits within-skill heterogeneity in WFH rather than the heterogeneity between high-skilled and low-skilled workers, potentially correlated with unobserved shocks.

The other six control variables do not have a direct correspondence with the occupation classification, but are included to make local labor markets comparable and net out confounding factors correlated with other market characteristics. We control for between-market heterogeneity in education, age, productivity, wages, firm size, and market size. Finally, the lag of the dependent variable accounts for time persistence or regression to mean.

Even after conditioning on this large set of controls, there might be unobserved occupation-specific shocks at the national level contaminating our identification. To prevent these unobservables from distorting our estimates, we follow Autor et al. (2013) and instrument our shift-share regressor with a related shift-share index, in which WFH shifters are measured using job vacancy postings from the United States rather than Italy.

An additional issue highlighted by the econometric literature on shift-share designs concerns the estimation of standard errors. Adão et al. (2019) and Borusyak et al. (2025) argue that, when standard errors have a shift-share structure, typical tests involving heteroskedasticity-robust or clustered standard errors may lead to severe over-rejection of the null hypothesis. Intuitively, local labor markets with similar occupational shares are exposed to similar WFH shifts and will likely have similar residuals. Thus, multiple observations may bear little additional information on the underlying shocks although the 'traditional' standard errors mechanically fall in sample size. To account for this issue, we follow the methodology of Adão et al. (2019) — henceforth AKM — which derive an estimator of standard errors that is asymptotically valid under any correlation structure of the residuals across markets, provided that occupation-level shifters are uncorrelated.

Finally, we face a practical issue in dealing with observations lacking a critical mass to be considered as labor markets in a meaningful way. Especially in rural areas, the interaction of a commuting zone and a two-digit industry may consist in a negligible number of workers. While acknowledging that any artificial threshold is hard to justify per se, we decide to set 15 workers as the the minimum threshold for our baseline estimation, and then test the

robustness of our results under two extreme alternatives: (i) no threshold, which implies including *markets* with less than a single FTE worker; and (ii) a 1,500 threshold, larger than the baseline by two orders of magnitude. The baseline threshold corresponds to a sample size of 16,993 observations, roughly three-fourth of the 22,833 observations considered under no threshold, whereas the 1,500 threshold leads us to only consider the 835 largest markets.

5 Results

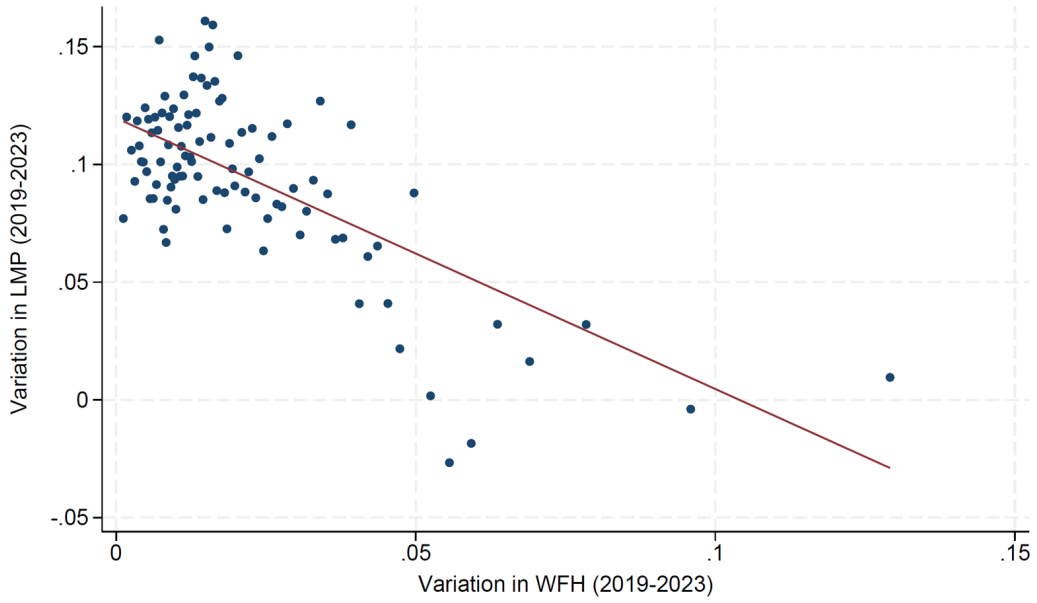
This Section presents the main empirical findings on the relationship between the diffusion of remote work and labor market power. We begin by illustrating the empirical association on a visual level. Next, we report the regression results of the baseline specification using OLS and 2SLS estimators, with heteroskedasticity-robust as well as AKM shift-share standard errors. We test pre-trends by switching the dependent variable of the baseline specification (ΔLMP in 2019-2023) with its lag (ΔLMP in 2015-2019). These results are corroborated by testing the mechanism with 2SLS estimators and AKM shift-share standard errors under a variety of alternative specifications. We then conduct a counterfactual analysis to better interpret the estimated magnitude of the effect. Finally, we deepen the analysis by running additional tests at the firm-level and worker-level.

5.1 Market-level

Graphical results. Figure 5.1 displays the unconditional correlation between time differences in labor market power (LMP) and time differences in work-from-home (WFH), revealing a clear negative association between the two variables. This binned scatter plot represents the raw relationship, equivalent to an OLS regression without controls. Each of the 100 dots represents nearly 200 local labor markets.

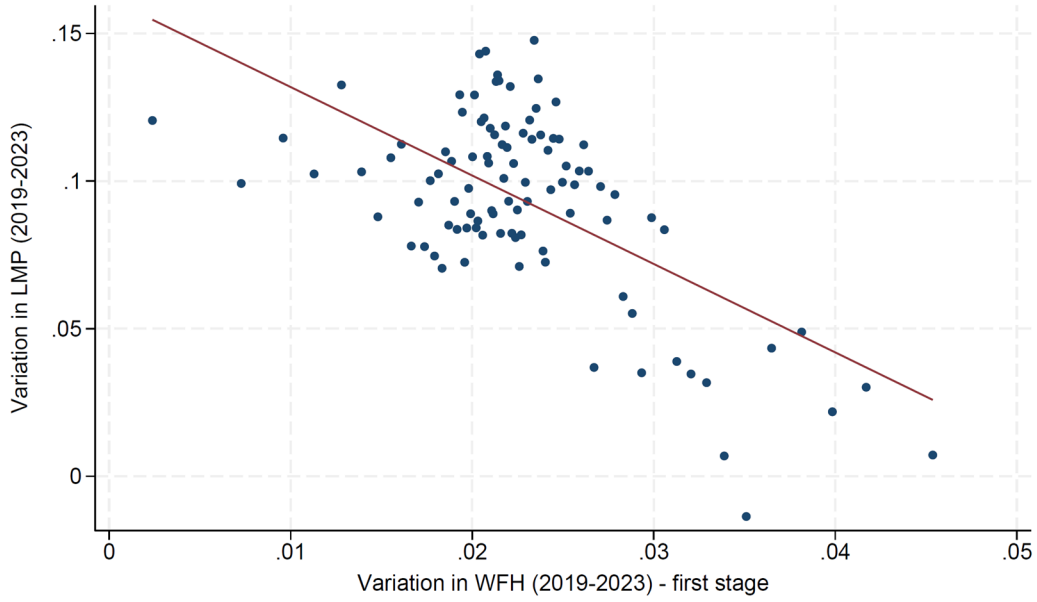
To account for endogeneity, we include a set of controls and instrument WFH shifts in Italy with WFH shifts in the US. Figure 5.2 visualizes our baseline specification. It plots the variation in LMP against the variation in WFH instrumented by the first stage, residualized on the full set of controls discussed in Section 4. The negative slope confirms that the relationship holds when we isolate the exogenous variation in the adoption of remote work.

Figure 5.1: Labor market power and WFH — unconditional correlation



Note: This figure shows a binned scatter plot with variations in aggregate markdowns against variations in remote work at the market-level. Sources of the raw data: Lightcast, VisitINPS, and AIDA.

Figure 5.2: Labor market power and WFH — baseline



Note: This figure shows a binned scatter plot with variations in aggregate markdowns against (instrumented) variations in remote work at the market-level. Both variables are residualised against the full set of controls (see Section 4). Sources of the raw data: Lightcast, VisitINPS, and AIDA.

Regression tables. Table 5.1 presents the OLS as well as both stages of the 2SLS for our baseline specification, using heteroskedasticity-robust standard errors and using AKM shift-share standard errors. OLS coefficients (Columns 1 and 4) are substantially lower than 2SLS coefficient (Columns 3 and 6). A possible reason lies in the superior signal-to-noise ratio in the instrument, due to the vastly larger sample of US job postings compared to Italian job postings, offsetting the attenuation bias induced by measurement error in the regressor. Alternatively, dependent and independent variables may co-move with occupation-specific shocks at the country-level, which would be net out using WFH shifts from the US.[19] As expected, heteroskedastic standard errors (Columns 1-3) are notably lower than AKM standard errors (Columns 4-6). Nevertheless, coefficients remain highly significant in the more conservative specification. Column 6 represents our baseline specification. We find that an increase in WFH exposure causes a significant reduction in labor market power. The coefficient is -3.01, which implies a -0.06 drop in the aggregate markdown for one interquantile range in WFH variations.

We then test the role of pre-trends by replacing the 2019-2023 variation in labor market power with its lag (2015-2019) as the dependent variable. For an accurate comparison, we replicate the baseline removing the lag of the dependant from the set of control variables, and then substitute the dependent with its lag. Column 1 shows that removing the lag from the set of controls has negligible effects. Crucially, switching the dependent variable with its lag (Column 2) reduces the point estimate by nearly 95% and increases the p-value from $<1\%$ to $>80\%$.

We conduct extensive robustness checks to validate our findings. Table 5.3 provides the results — first-stage F-Stat and second-stage key coefficient — across six alternative specifications, using 2SLS and AKM standard errors. Column (1) replaces our revenue-based measure of labor market power with the approach of Yeh et al. (2022), showing a very similar point estimate. Similarily, results do not change substantially in Column (2), in which the set of controls is reduced from eight variables to three (i.e. share of high-skilled workers, aggregate labor costs per worker in logs, and number of workers per firm in logs). Column (3) adds the geometric mean of FTE labor in 2019 and 2023 as weights in the estimation, slightly reducing the magnitude.[20] In column (5) and (6), we make *extreme* manipulations to the threshold of minimum FTE to be considered as a labor market. In column (4), we

[19] For instance, a fall in real wages due to outdated collective contracts, accompanied by higher wage markdowns, may induce employers to offer remote arrangements in order to attract more workers.

[20] We prefer the unweighted specification because our hypothesis concerns a general equilibrium effect taking place at the market-level rather than a worker-level mechanism; thus, we view each labor market as an equally relevant unit of observation for our tests.

Table 5.1: Baseline specification

	(1) ΔLMP	(2) ΔWFH	(3) ΔLMP	(4) ΔLMP	(5) ΔWFH	(6) ΔLMP
ΔWFH	-1.22***		-3.01***	-1.22**		-3.01***
	(0.12)		(0.27)	(0.49)		(0.99)
ΔWFH_{US} (IV)		0.35***			0.35***	
		(0.01)			(0.11)	
I stage F-Stat		941.81			10.79	
Method	OLS	2SLS (I)	2SLS (II)	OLS	2SLS (I)	2SLS (II)
Standard errors	Heterosk.	Heterosk.	Heterosk.	AKM	AKM	AKM
LMP measure	Revenues	Revenues	Revenues	Revenues	Revenues	Revenues
Controls	Full set	Full set	Full set	Full set	Full set	Full set
Weighting	None	None	None	None	None	None
Minimum FTE	15	15	15	15	15	15
Observations	16,993	16,993	16,993	22,833	835	16,993
Unit	CZ-industry	CZ-industry	CZ-industry	CZ-industry	CZ-industry	CZ-industry

Note: ***p<0.01, **p<0.05, *p<0.1. This table reports the OLS (Columns 1 and 4), as well as the first stage (Columns 2 and 5) and second stage (Column 3 and 6), of the estimation of 4.1. Heteroskedasticity-robust standard errors (Columns 1-3) and AKM shift-share standard errors (Columns 4-6) are reported in parentheses. Each observation is the interaction of a commuting zone and a 2-digit industry. Sample restricted to a minimum of 15 FTE workers. The full set of controls includes: (i) share of high-skilled workers; (ii) share of college-educated workers; (iii) average age of workers; (iv) aggregate value added per worker (in logs); (v) aggregate labor costs per worker (in logs); (vi) number of workers per firm (in logs); (vii) number of workers in FTE units (in logs); and (viii) lag of the dependent variable (i.e. the 2015–2019 variation in the aggregate markdown). All variables are winsorized at the top and bottom 1%. Sources of the raw data: Lightcast, VisitINPS, AIDA.

remove the threshold including every cell, even those made of one or less FTE workers, and thus increasing the number of observations from 16,993 to 22,833. In column (5), we increase the threshold by two orders of magnitude, from 15 to 1,500 FTE units, reducing the sample to the 835 largest labor markets. The magnitude increases when we include tiny markets and decreases when we restrict to the largest ones, whilst remaining in a relatively confined range. Finally, in column (6), we directly test the effect on wages, specifically on the 2019-2023 growth in labor costs per worker. Although labor market competition is only one of several potential channels through which WFH may affect wages, our hypothesis is consistent with the positive effect on wages that emerges from the data.

Table 5.2: Test of pre-trends

	(1) ΔLMP	(2) $\Delta LMP_{2015-2019}$
ΔWFH	-2.98***	-0.16
	(0.96)	(0.67)
I stage F-Stat	10.79	10.79
Method	2SLS	2SLS
Standard errors	AKM	AKM
LMP measure	Revenues	Revenues
Controls	Full set \lag	Full set \lag
Weighting	None	None
Minimum FTE	15	15
Observations	16,993	16,993
Unit	CZ-industry	CZ-industry

Note: ***p<0.01, **p<0.05, *p<0.1. This table reports the second stage, as well as the F-stat of the first stage, of the 2SLS estimation of 4.1, using the contemporaneous variation of the dependent variable (2019-2023 in Column 1) and using its lag (2015-2019 in Column 2). AKM shift-share standard errors are reported in parentheses. Each observation is the interaction of a commuting zone and a 2-digit industry. Sample restricted to a minimum of 15 FTE workers. The full set of controls includes: (i) share of high-skilled workers; (ii) share of college-educated workers; (iii) average age of workers; (iv) aggregate value added per worker (in logs); (v) aggregate labor costs per worker (in logs); (vi) number of workers per firm (in logs); and (vii) number of workers in FTE units (in logs). All variables are winsorized at the top and bottom 1%. Sources of the raw data: Lightcast, VisitINPS, AIDA.

Table 5.3: Alternative specifications

	(1) ΔLMP	(2) ΔLMP	(3) ΔLMP	(4) ΔLMP	(5) ΔLMP	(6) $\Delta \%W$
ΔWFH	-3.03*	-3.06***	-2.71**	-3.47***	-2.11**	2.07**
	(1.76)	(0.99)	(1.12)	(1.09)	(0.85)	(0.93)
I stage F-Stat	10.81	10.12	10.53	10.64	10.79	11.12
Method	2SLS	2SLS	2SLS	2SLS	2SLS	2SLS
Standard errors	AKM	AKM	AKM	AKM	AKM	AKM
LMP measure	Production	Revenues	Revenues	Revenues	Revenues	Revenues
Controls	Full set	Small set	Full set	Full set	Full set	Full set
Weighting	None	None	FTE (gm)	None	None	None
Minimum FTE	15	15	15	0	1,500	15
Observations	16,993	16,993	16,993	22,833	835	16,993
Unit	CZ-industry	CZ-industry	CZ-industry	CZ-industry	CZ-industry	CZ-industry

Note: ***p<0.01, **p<0.05, *p<0.1. This table reports the second stage of the 2SLS estimation of 4.1, as well as the F-stat of the first stage, under a set of alternative specifications. AKM shift-share standard errors are reported in parentheses. Each observation is the interaction of a commuting zone and a 2-digit industry. Sample restricted to a minimum of 15 FTE workers, except fro Columns (4) and (5). The full set of controls includes: (i) share of high-skilled workers; (ii) share of college-educated workers; (iii) average age of workers; (iv) aggregate value added per worker (in logs); (v) aggregate labor costs per worker (in logs); (vi) number of workers per firm (in logs); (vii) number of workers in FTE units (in logs); and (viii) lag of the dependent variable (i.e. the 2015–2019 variation in the aggregate markdown). All variables are winsorized at the top and bottom 1%. Sources of the raw data: Lightcast, VisitINPS, AIDA.

Counterfactual analysis. To provide a more intuitive interpretation of the magnitude implied by our estimates, Table 5.4 compares the 2023 aggregate markdown against three counterfactual scenarios.

First, using the point estimate of our baseline specification, for each local labor market we derive the aggregate markdown that would have emerged if the 2019–2023 WFH shifts never occurred. According to our estimates, the average value of aggregate markdowns would be 0.36 instead of the realised 0.29, a rise by nearly one fifth from the realised value.

Finally, we consider the greater transition to remote work that took place in large cities as a benchmark. To simulate a realistic counterfactual, we net out differences between commuting zones stemming from differences in the industry composition.[21] For this purpose, we project the observed WFH shifts of local labor market c — defined as the interaction of commuting zone j and industry i — on commuting zone and industry fixed effects.

$$\Delta WFH_{i,j} = \theta_1 + \theta_i + \theta_j + v_{i,j} \tag{5.1}$$

Large part of the heterogeneity in WFH shifts come from heterogeneity between industries, which we take as given in this counterfactual exercise. Instead, we assume that — conditioning on the industry fixed effects — the fixed effect of large cities can be used as a reasonable benchmark for other commuting zones. We compute the counterfactual WFH shifts by substituting θ_j of each commuting zone with the corresponding fixed effect of Rome and Milan. Our results reveal that the country-level aggregate markdown would be 0.26 or 0.25, with Rome or Milan fixed effects respectively, instead of the realised 0.29, a drop by approximately one tenth.

[21] In the thought experiment of this exercise, we consider how labor market power would have evolved if industries across all commuting zones transitioned to remote work by the amount observed in the same industries in large cities, i.e. we control for the fact that large cities are specialised in industries that are more compatible with working from home.

Table 5.4: Counterfactual Analysis

Scenario	WFH shifts	Aggregate markdown
Realized	$\Delta WFH_{i,j}$	*0.29*
No transition	0	0.36
Transition as in Rome	$\Delta WFH_{i,j} - \theta_j + \theta_{Rome}$	0.26
Transition as in Milan	$\Delta WFH_{i,j} - \theta_j + \theta_{Milan}$	0.25

Note: This table reports the estimated value of the country-level aggregate markdown in the realised case and in three counterfactuals: in the first row, we consider the case in which no WFH shift occurred between 2019 and 2023; in the third and fourth rows, we considered the case in which all local labor markets experienced the WFH shifts that occurred in the same industries in Rome and Milan, respectively. Counterfactuals are derived using the coefficient estimated in our baseline specification. Sources of the raw data: Lightcast, VisitINPS, AIDA.

5.2 Firm-level

Although a market-level specification is the most natural setting to test the general equilibrium effect considered in this paper, we complement our findings with further evidence derived at a narrower level — firm-level in this Subsection and worker-level in the next Subsection — taking full advantage of the richness of our data.

Our dependent variable is the 2019-2023 variation in the firm's markdown (or labor costs per worker, in a robustness exercise). The key regressor and the instrument are built analogously to our market-level procedure: occupation-specific WFH shifts (from Italian data in the regressor and from US data in the IV) are combined with the 2019 employment shares of each occupational group. The fundamental distinction is that we are now using the firm-level — rather than market-level — employment shares. The set of controls includes: (i) share of high-skilled workers; (ii) share of college-educated workers; (iii) average age of workers; (iv) value added per worker (in logs); (v) labor costs per worker (in logs); (vi) number of workers in FTE units (in logs); (vii) the lag of dependent variable, i.e. the 2015-2019 variation in the firm's markdown (or labor costs per worker).

A key difference, and perhaps limitation, in the firm-level setting is that we are forced to restrict the sample to firms that have been active between 2015 and 2023. An additional difference is that observations cannot be considered as equally representative of their own general equilibrium effect, thus in this case we opt for a weighted specification, using the geometric mean of FTE workers in 2019 and 2023 as weights, with no minimum FTE threshold.

Table 5.6 provides the first-stage F-Stat and second-stage key coefficient of such specifications. Columns 1 and 2 replicate the firm-level equivalent the market-level baseline, with heteroskedasticity-robust and AKM shift-share standard errors, respectively. Column 3 and 4 test alternative dependent variables, with AKM standard errors: in Column 3 we consider the alternative measure of wage markdowns (Yeh et al. (2022)), and in Column 4 we test the effect on labor costs per worker.

Table 5.5: Firm-level tests

	(1) ΔLMP	(2) ΔLMP	(3) ΔLMP	(4) $\Delta \%W$
ΔWFH	-1.23***	-1.23***	-1.23***	0.94*
	(0.14)	(0.29)	(0.38)	(0.54)
I stage F-Stat	1462.88	43.27	43.28	43.50
Method	2SLS	2SLS	2SLS	2SLS
Standard errors	Heterosk.	AKM	AKM	AKM
LMP measure	Revenues	Revenues	Production	Revenues
Controls	Full set	Full set	Full set	Full set
Weighting	FTE (gm)	FTE (gm)	FTE (gm)	FTE (gm)
Minimum FTE	0	0	0	0
Observations	234,127	234,127	234,127	234,127
Unit	Firm	Firm	Firm	Firm

Note: ***p<0.01, **p<0.05, *p<0.1. This table reports the second stage of the 2SLS estimation of the impact of WFH on labor market power, as well as the F-stat of the first stage, under a set of alternative specifications. Heteroskedasticity-robust standard errors (Column 1) and AKM shift-share standard errors (Columns 2-4) are reported in parentheses. Each observation is a firm. The full set of controls includes: (i) share of high-skilled workers; (ii) share of college-educated workers; (iii) average age of workers; (iv) aggregate value added per worker (in logs); (v) aggregate labor costs per worker (in logs); (vi) number of workers in FTE units (in logs); and (vii) lag of the dependent variable (i.e. the 2015–2019 variation in the firm's markdown). All variables are winsorized at the top and bottom 1%. Sources of the raw data: Lightcast, VisitINPS, AIDA.

In all instances, results are significant and go in the same direction as in the market-level analysis. We do not interpret these findings as a proof that market-level and firm-level *treatments* generate comparable mechanisms. In fact, the shift-share structure makes the regressor an imperfect proxy of the firm's decision in isolation: despite being defined at the firm-level, by attributing an occupation-based treatment to each firm, the regressor still recovers spillovers induced by WFH adoption in competitors with a comparable occupation mix. Firm-level regressions may then be better interpreted as tests of the same market-level

mechanism with more specific controls and an alternative market definition, which in this case depends on the firm's occupation mix rather than commuting zone and industry.

5.3 Worker-level

In the worker-level setting, we directly combine the worker's occupation with the occupation-specific WFH shift, based on Italian job postings for the regressor and on US job postings for the IV. As there is no need to make use of a shift-share index, our worker-level tests are implemented under occupation-clustered standard errors, unlike the AKM shift-share standard errors used in the market-level and firm-level specifications. The effect on wage markdowns is not tested, as the latter are defined at the firm-level (and through aggregation, at the market-level). We instead evaluate the impact on wage growth and on a separate channel, which could not be tested at a more aggregate level: the probability of a commuting zone mismatch between the worker's residence and workplace.

Table 5.6: Worker-level tests

	(1) $\Delta\%W$	(2) CZ mismatch
ΔWFH	0.48***	0.17***
	(0.14)	(0.06)
I stage F-Stat	25.83	26.58
Method	2SLS	2SLS
Standard errors	Clustered	Clustered
Fixed effects	CZ # industry	CZ # industry
Controls	Full set	Full set
Weighting	FTE (gm)	FTE (gm)
Observations	$\approx$ 7mln	$\approx$ 10mln
Unit	Worker	Worker

Note: ***p<0.01, **p<0.05, *p<0.1. This table reports the second stage of the 2SLS estimation, as well as the F-stat of the first stage, of the impact of WFH on wage growth (Column 1) and on the commuting zone mismatch between the worker's residence and workplace (Column 2). Occupation-clustered standard errors (ISCO three-digit) are reported in parentheses. Fixed effects based on the interaction of a commuting zone and two-digit industry are included. Each observation is a worker. The full set of controls includes: (i) skill level (dummy); (ii) college degree (dummy); (iii) female sex (dummy); (iv) foreign citizenship (dummy); (v) age; (vi) age squared; and (vii) lag of the dependent variable (i.e. the 2015–2019 variation in wage growth for Column 1 and a dummy for the commuting zone mismatch in 2019 in Column 2). Weights are derived as the geometric mean of the individual worker's FTE labor in 2019 and 2023. All variables are winsorized at the top and bottom 1%. Sources of the raw data: Lightcast, VisitINPS.

6 Conclusion

This paper provides novel evidence on an overlooked consequence of the transition to remote work, connecting to an extensive literature that identifies distaste for commuting as a primary source of rent extraction by local employers.

Our analysis combines three sources of microdata: (i) millions of job vacancy postings to measure the WFH transition at the occupation-level using a frontier LLM; (ii) worker-level administrative records on the workplace location and occupational group, which we use to map occupation-level WFH shifts into a market-level index; and (iii) firm-level production data to estimate wage markdowns and infer time changes in the degree of competition across local labor markets.

Exploiting heterogeneity in shifts to remote work across occupations — instrumented with contemporaneous shifts observed abroad — and heterogeneity in occupational shares across commuting zones and industries, we construct a shift-share IV to infer causal effects. Our evidence shows that work from home has substantially attenuated the 2019–2023 increase in wage markdowns, which we observe in the data and attribute to the 2022–2023 inflationary surge. Absent the transition to remote work over this period, local labor markets would have experienced, on average, a one-fifth larger departure from perfect competition.

References

Ackerberg, D. A., Caves, K., & Frazer, G. (2015). Identification properties of recent production function estimators. *Econometrica: Journal of the Econometric Society*, *83*, 2411–2451.

Adão, R., Kolesár, M., & Morales, E. (2019). Shift-share designs: Theory and inference. *The Quarterly Journal of Economics*, *134*, 1949–2010. https://doi.org/10.1093/qje/qjz025

Afrouzi, H., Blanco, A., Drenik, A., & Hurst, E. (2024, December). *A theory of how workers keep up with inflation* (research rep.). Cambridge, MA, National Bureau of Economic Research. https://doi.org/10.3386/w33233

Akan, M., Barrero, J. M., Bloom, N., Bowen, T., Buckman, S. R., et al. (2025). The new geography of labor markets.

Althoff, L., Eckert, F., Ganapati, S., & Walsh, C. (2022). The geography of remote work. *Regional Science and Urban Economics*, *93*, 103770.

Alves, J., Serra Lorenzo, B., Greenberg, J., Guo, Y., Harjai, R., & Van Reenen, J. (2024). Labour market power: New evidence on non-compete agreements and the effects of M&A in the UK. *Centre for Economic Performance Discussion Papers*, 58.

Autor, D. H., Dorn, D., & Hanson, G. H. (2013). The china syndrome: Local labor market effects of import competition in the united states. *American Economic Review*, *103*, 2121–2168. https://doi.org/10.1257/aer.103.6.2121

Azar, J., Marinescu, I., & Steinbaum, M. (2022). Labor market concentration. *The Journal of Human Resources*, *57*, S167–S199. https://doi.org/10.3368/jhr.monopsony.1218-9914r1

Azar, J., Marinescu, I., Steinbaum, M., & Taska, B. (2020). Concentration in US labor markets: Evidence from online vacancy data. *Labour Economics*, *66*, 101886. https://doi.org/10.1016/j.labeco.2020.101886

Barrero, J. M., Bloom, N., & Davis, S. J. (2021). Why working from home will stick.

Barrero, J. M., Bloom, N., & Davis, S. J. (2023). The evolution of work from home. *The Journal of Economic Perspectives: A Journal of the American Economic Association*, *37*, 23–49. https://doi.org/10.1257/jep.37.4.23

Bassier, I., Dube, A., & Naidu, S. (2022). Monopsony in movers: The elasticity of labor supply to firm wage policies. *The Journal of Human Resources*, *57*, S50–s86. https://doi.org/10.3368/jhr.monopsony.0319-10111r1

Basso, G., Boeri, T., Caiumi, A., & Paccagnella, M. (2020). The new hazardous jobs and worker reallocation. *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.3660279

Berger, D., Herkenhoff, K., Kostøl, A. R., & Mongey, S. (2024). An anatomy of monopsony: Search frictions, amenities, and bargaining in concentrated markets. *NBER Macroeconomics Annual*, *38*, 1–47. https://doi.org/10.1086/729194

Berger, D., Herkenhoff, K., & Mongey, S. (2022). Labor market power. *American Economic Review*, *112*, 1147–1193. https://doi.org/10.1257/aer.20191521

Bick, A., Rubinton, H., Blandin, A., & Mertens, K. (2024). *Work from home and interstate migration* (research rep.). Federal Reserve Bank of St. Louis. https://doi.org/10.20955/wp.2024.012

Bloom, N., Liang, J., Roberts, J., & Ying, Z. J. (2015). Does working from home work? evidence from a chinese experiment. *The Quarterly Journal of Economics*, *130*, 165–218.

Boeri, T., Garnero, A., & Luisetto, L. G. (2025). Noncompete agreements in a rigid labor market: The case of italy. *Journal of Law, Economics, & Organization*, *41*, 930–958.

Borusyak, K., Hull, P., & Jaravel, X. (2025). A practical guide to shift-share instruments. *The Journal of Economic Perspectives: A Journal of the American Economic Association*, *39*, 181–204. https://doi.org/10.1257/jep.20231370

Burdett, K., & Mortensen, D. T. (1998). Wage differentials, employer size, and unemployment. *International Economic Review*, *39*, 257. https://doi.org/10.2307/2527292

Cammeraat, E., & Squicciarini, M. (2021). *Burning glass technologies' data use in policy-relevant analysis: An occupation-level assessment* (research rep.). OECD Science, Technology and Industry Working Papers.

Card, D. (2022). Who set your wage? *American Economic Review*, *112*, 1075–1090. https://doi.org/10.1257/aer.112.4.1075

Card, D., Cardoso, A. R., Heining, J., & Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, *36*, S13–S70. https://doi.org/10.1086/694153

Carnevale, A. P., Jayasundera, T., & Repnikov, D. (2014). *Understanding online job ads data* (research rep.). A technical report. MS o. PP Center on Education and the Workforce. https ... McCourt School of Public Policy.

Cullen, Z., Pakzad-Hurson, B., & Perez-Truglia, R. (2025, January). *Home sweet home: How much do employees value remote work?* (Research rep.). American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203. Cambridge, MA, National Bureau of Economic Research. https://doi.org/10.3386/w33383

Datta, N. (2024). *Local monopsony power* (research rep.). London, Centre for Economic Performance.

Deb, S., Eeckhout, J., Patel, A., & Warren, L. (2022). *FBBVA lecture 2022:* what drives wage stagnation: Monopsony or monopoly? *Journal of the European Economic Association*, *20*, 2181–2225. https://doi.org/10.1093/jeea/jvac060

Delabastita, V., & Rubens, M. (2025). Colluding against workers. *The Journal of Political Economy*, *133*, 1796–1839. https://doi.org/10.1086/734780

Dingel, J. I., & Neiman, B. (2020). How many jobs can be done at home? *Journal of Public Economics*, *189*, 104235. https://doi.org/10.1016/j.jpubeco.2020.104235

Dube, A., Jacobs, J., Naidu, S., & Suri, S. (2020). Monopsony in online labor markets. *American Economic Review. Insights*, *2*, 33–46. https://doi.org/10.1257/aeri.20180150

Dustmann, C., Lindner, A., Schönberg, U., Umkehrer, M., & vom Berge, P. (2021). Reallocation effects of the minimum wage. *The Quarterly Journal of Economics*, *137*, 267–328. https://doi.org/10.1093/qje/qjab028

Edmond, C., Midrigan, V., & Xu, D. Y. (2023). How costly are markups? *The Journal of Political Economy*, *131*, 000–000. https://doi.org/10.1086/722986

Emanuel, N., & Harrington, E. (2024). Working remotely? selection, treatment, and the market for remote work. *American Economic Journal. Applied Economics*, *16*, 528–559. https://doi.org/10.1257/app.20230376

Gibbs, M., Mengel, F., & Siemroth, C. (2023). Work from home and productivity: Evidence from personnel and analytics data on information technology professionals. *Journal of Political Economy Microeconomics*, *1*, 000–000. https://doi.org/10.1086/721803

Goolsbee, A., & Syverson, C. (2023, January). *The strange and awful path of productivity in the U.s. construction sector* (research rep.). Cambridge, MA, National Bureau of Economic Research. https://doi.org/10.3386/w30845

Hansen, S., Lambert, P. J., Bloom, N., Davis, S., Sadun, R., & Taska, B. (2023, March). *Remote work across jobs, companies, and space* (research rep. No. 31007). National Bureau of Economic Research. Cambridge, MA, National Bureau of Economic Research. https://doi.org/10.3386/w31007

Hashemi, A., Kirov, I., & Traina, J. (2022). The production approach to markup estimation often measures input distortions. *Economics Letters*, *217*, 110673. https://doi.org/10.1016/j.econlet.2022.110673

Hirsch, B., Jahn, E. J., Manning, A., & Oberfichtner, M. (2022). The urban wage premium in imperfect labor markets. *The Journal of Human Resources*, *57*, S111–S136. https://doi.org/10.3368/jhr.monopsony.0119-9960r1

Kroft, K., Luo, Y., Mogstad, M., & Setzler, B. (2025). Imperfect competition and rents in labor and product markets: The case of the construction industry. *American Economic Review*, *115*, 2926–2969. https://doi.org/10.1257/aer.20220577

Lamadon, T., Mogstad, M., & Setzler, B. (2022). Imperfect competition, compensating differentials, and rent sharing in the US labor market. *American Economic Review*, *112*, 169–212. https://doi.org/10.1257/aer.20190790

Lightcast. (2022). *European data methodology* (research rep.).

Liu, S., & Su, Y. (2025). The effect of working from home on the agglomeration economies of cities: Evidence from advertised wages. *The Review of Economics and Statistics*, 1–45. https://doi.org/10.1162/rest.a.1681

Manning, A. (2003). *Monopsony in motion: Imperfect competition in labor markets*. Princeton University Press.

Manning, A. (2021). The elusive employment effect of the minimum wage. *The Journal of Economic Perspectives: A Journal of the American Economic Association*, *35*, 3–26. https://doi.org/10.1257/jep.35.1.3

Manning, A., & Petrongolo, B. (2017). How local are labor markets? evidence from a spatial job search model. *American Economic Review*, *107*, 2877–2907.

Manning, A., & Petrongolo, B. (2024). Monopsony in local labour markets. *Oxford Open Economics*, *3*, i951–i958. https://doi.org/10.1093/ooec/odad061

Mas, A., & Pallais, A. (2017). Valuing alternative work arrangements. *American Economic Review*, *107*, 3722–3759. https://doi.org/10.1257/aer.20161500

Mengano, P. (2025). *Trends in worker bargaining power* (research rep.). Esade Business School. https://doi.org/10.34724/casd

Mertens, M., & Mottironi, B. (2025). Do larger firms exert more market power? markups and markdowns along the size distribution. *Economics Letters*, *255*, 112500. https://doi.org/10.1016/j.econlet.2025.112500

Monte, F., Porcher, C., & Rossi-Hansberg, E. (2023, July). *Remote work and city structure* (research rep.). National Bureau of Economic Research. Cambridge, MA, National Bureau of Economic Research. https://doi.org/10.3386/w31494

Mottironi, B. (2024). Labour market power and aggregate productivity.

OECD. (2025). *Employment outlook* (research rep.). Paris. https://doi.org/0.1787/194a947b-en

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

Shiller, R. (1996, April). *Why do people dislike inflation?* (C. D. Romer & D. H. Romer, Eds.; research rep.). National Bureau of Economic Research. Cambridge, MA, National Bureau of Economic Research. https://doi.org/10.3386/w5539

Song, Y., & Gao, J. (2018). *Does telework stress employees out? a study on working at home and subjective well-being for wage/salary workers* (research rep.). IZA Discussion Papers.

Stansbury, A., & Summers, L. H. (2020). The declining worker power hypothesis: An explanation for the recent evolution of the american economy.

Stantcheva, S. (2024). Why do we dislike inflation? *SSRN Electronic Journal*. https://doi.org/10.2139/ssrn.4787391

Yang, L., Holtz, D., Jaffe, S., Suri, S., Sinha, S., Weston, J., Joyce, C., Shah, N., Sherman, K., Hecht, B., & Teevan, J. (2022). The effects of remote work on collaboration among information workers [Publisher: Nature Publishing Group]. *Nature Human Behaviour*, *6*, 43–54. https://doi.org/10.1038/s41562-021-01196-4

Yeh, C., Macaluso, C., & Hershbein, B. (2022). Monopsony in the US labor market. *American Economic Review*, *112*, 2099–2138. https://doi.org/10.1257/aer.20200025

Zarate, P., Barrero, J. M., Bloom, N., Davis, S. J., Dolls, M., & Aksoy, C. G. (2025). Why working from home varies across countries and people. *Proceedings of the National Academy of Sciences of the United States of America*, *122*, e2529036122. https://doi.org/10.1073/pnas.2529036122

Online Appendix

A Further remarks on measuring remote work

The unit of observation in the Lightcast data is an online job posting (or, more precisely, a de-duplicated advertisement record). Each record includes the job title and the full job description text as scraped from the source website, alongside a set of structured fields produced by Lightcast. In our Italian extract, coverage begins in 2018 and runs through 2024, spanning the pre-pandemic period, the rapid adoption of emergency remote work during 2020, and the subsequent stabilisation of hybrid arrangements. The dataset includes postings from thousands of distinct online sources — job boards as well as employer websites — which is important in Italy where recruiting channels are fragmented and differ across sectors and regions (Lightcast, 2022). Table A.1 reports the number of online job vacancy postings by year.

Table A.1: Count of Online Job Vacancy Postings

Year	Postings	Occupations	Regions	Firms	Sources
2018	1,218,007	40	22	40,266	198
2019	1,724,453	40	22	69,880	253
2020	1,863,369	40	22	102,310	327
2021	2,593,521	40	22	135,553	381
2022	3,388,474	41	22	172,220	1,890
2023	5,026,268	41	22	174,605	2,300
2024	4,427,934	41	22	168,821	2,255
Total	20,242,026	41	22	575,988	3,407

Note: This table reports the number of online job vacancy postings by year. Occupations are defined using the two-digit ISCO classification. Regions correspond to the NUTS-2 regional classification. Firms are unique employer name strings, constructed from cleaned identifiers. Job-board sources correspond to distinct platforms or career sites. Source of the raw data: Lightcast.

A central challenge with online vacancy data is that the same job opportunity can appear multiple times and on multiple platforms, either because employers cross-post an ad across job boards, or because intermediaries syndicate ads. Lightcast addresses through a proprietary de-duplication procedure that attempts to identify and collapse duplicate ads into a single record, using information in job titles, descriptions, employers, and locations (Lightcast, 2022). In addition, Lightcast performs entity resolution on employer names, which is crucial in practice because the same firm may appear under different spellings, legal forms, or language variants across sources. These steps are a key reason that job-posting datasets such as Lightcast/Burning Glass have become widely used in empirical research on labor demand,

where scale and standardisation are essential (Cammeraat & Squicciarini, 2021; Carnevale et al., 2014).

Beyond cleaning and de-duplication, Lightcast enriches postings with harmonised classifications. The pipeline assigns occupation codes and industry codes to postings using the text and any available structured information, enabling aggregation at standard taxonomies — for Europe, commonly ISCO for occupations and NACE for industries (Lightcast, 2022). Locations are geocoded so that postings can be aggregated to municipalities, provinces, regions, or other functional geographies. This geographic structure is central for our application, because we map postings into local labor markets and study the uneven diffusion of remote work across space. Finally, Lightcast extracts additional features from job description text, including skill requirements and credentials, producing structured indicators that can be used either directly as outcomes or as controls in empirical designs (Lightcast, 2022). In our application, we primarily use the raw text and the harmonised occupation field.

The Lightcast data do not provide a consistently populated, structured indicator of remote or hybrid status for Italy over our full sample window, which is why we classify remote/hybrid arrangements directly from the posting text. This approach follows a broader methodological trend in job-posting research, which increasingly relies on modern NLP methods to extract job attributes from text while preserving transparency and validating performance against human labels. Table A.2 shows the prompt template used for the classification remote arrangements, whereas Table A.3 reports examples of classifications.

Table A.2: Prompt Template for the LLM Classification

Component	Description
Overall role	System message: the model is instructed to act as an expert assistant analysing Italian job postings to identify remote or hybrid work options.
Step 1: Key words mentioned	Determine whether the posting explicitly mentions remote work (e.g. working from home or a mix of home and office). The variable remote_hybrid_mentioned is coded 'Y' if any remote term appears anywhere in the text (including negative contexts such as 'no smart working') and 'N' otherwise.
Step 2: Offering type	Classify whether the posting explicitly offers or rules out remote work. The variable offering_type takes value 'P' if the text explicitly offers remote work, 'N' if it explicitly disallows remote work, and 'U' if there is no explicit offer or denial.
Step 3: Exact arrangement	Classify the work arrangement based only on explicit phrasing in the text. The variable work_arrangement_exact takes value 'Hybrid' for mixed office/home arrangements, 'Fully Remote' for entirely remote arrangements, and 'U' if the arrangement is not explicitly stated.
Step 4: Contextual arrangement	For postings with offering_type == 'P', classify the arrangement based on the broader context of the ad. The variable work_arrangement_contextual takes value 'Hybrid' or 'Fully Remote' based on the model's best guess. For postings with offering_type not equal to 'P', the model may respond 'U'.
Step 5: Supporting phrase	Extract the exact phrase or short sentence (approximately 3–10 words) from the original posting that most directly supports the classification of the remote offering. This text is stored in extracted_supporting_phrase and preserves the original language and spelling. It is set to 'null' if no relevant phrase exists.
Step 6: Confidence score	Return a numeric confidence_score between 0.0 (no confidence) and 1.0 (fully confident) capturing how certain the model is about the classification for a given posting.

Continued on next page

Table A.2 continued from previous page

Component	Description
Example of keywords	The prompt lists Italian and English keywords that should trigger remote_hybrid_mentioned, including for example: 'lavoro da remoto', 'smart working', 'telelavoro', 'lavoro a distanza', 'lavoro agile', 'ibrido', 'da casa', 'home office', 'remote work', and 'WFH'. The prompt is clear that this list is not exhaustive.
Disambiguation notes	The prompt instructs the model to ignore uses of 'ibrido' that clearly refer to schedules or shifts (e.g. 'orario ibrido') rather than work location, to treat even negative references such as 'no smart working' as remote_hybrid_mentioned = 'Y', and to base classifications strictly on the text of the job title and description.

Note: This table documents the LLM prompt template used to classify remote and hybrid work options in Italian vacancy postings. The system message specifies that the model acts as an expert assistant and the user message contains the job title and full description. The model returns a structured JSON object with fields remote_hybrid_mentioned, offering_type, work_arrangement_exact, work_arrangement_contextual, extracted_supporting_phrase, and confidence_score, as described in the table. The full Python implementation, including the exact system prompt and JSON schema, is provided in the replication code.

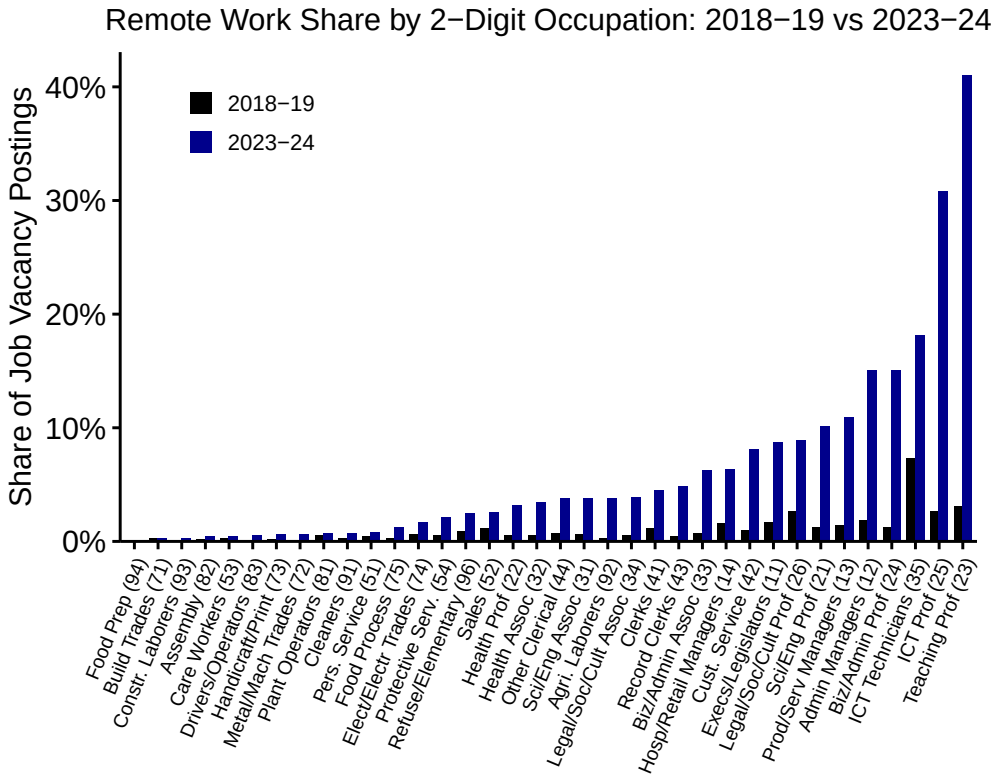
Finally, Figure A.1 documents the relative intensity of postings offering remote work across each sub-major occupation from the ISCO classification. In our main empirical exercise, we exploit this variation across occupations to define exposure to the rise of remote working across Italian local labor markets, aggregated using employment share across each occupation. We see teaching professionals, such as university teaching professionals, at the top of the list, along with ICT professionals. The lowest level of new job opportunities offering remote work are found in food-preparation and building and construction occupations.

Table A.3: Examples of Remote Classifications

Supporting phrase from the Job Ad	Remote type
Lavorare in Full – Remote	Fully Remote
Lavoro agile al 100% or Posizione completely da remoto	Fully Remote
Remote role in Southern Europe	Fully Remote
Work from anywhere	Fully Remote
Flexible work from home policy (3 days in office, 2 days at home)	Hybrid
modalità di lavoro ibrida (3 giorni in sede e 2 giorni in smart working)	Hybrid
Lavora da casa a Firenze e riporta due giorni al mese presso i nostri uffici	Hybrid
Sono previsti 2 giorni a settimana di smart-working	Hybrid
Contratto di lavoro full time in sede (no smart working)	Disallowed
Full remote is not allowed	Disallowed
Modalita' di lavoro: ON SITE 5/5	Disallowed
Smart-Working: NO	Disallowed
Non saranno prese in consider candidature da remoto	Disallowed

Note: This table reports examples of classifications drawn from Italian job postings, labelled by a LLM, with the supporting phrase representing the text that most directly supports the assigned assignment. 'Remote Type' indicates the type of arrangement mentioned in the phrase: Fully Remote, Hybrid, or Disallowed.

Figure A.1: Share of Job Postings Offering Remote Work by Occupation

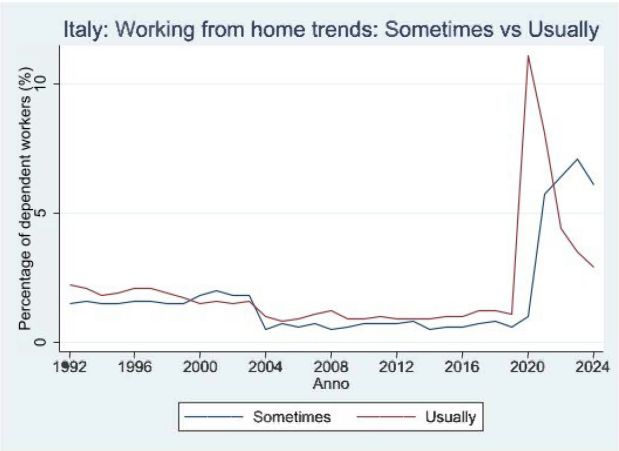


Note: This figure shows the share of job vacancy postings explicitly offering remote work arrangements by two-digit ISCO occupation. The occupation code is in brackets and the label for each is our own short descriptive title. Solid black lines report the share for 2018-19 and solid blue bars report the share for 2023-24. Remote offers are identified using the textual content of each job posting, by applying our LLM classification algorithm. Source of the raw data: Lightcast.

A.1 Comparison of Labour Force Survey (LFS) responses to text-based measures of remote work arrangements

Official survey-based measures of remote work provide a statistically representative measure of the stock of workers engaging in remote work, which is quite complimentary to our vacancy based measure of the flow of new hires. Each measure have advantages and draw-backs. The main advantage of our approach is the sheer scale of job vacancy postings from which to extract our WFH signal, especially for granular occuaption groups which may not be well covered in survey samples. The vacancy-based approach also emphasises forward-looking commitments by employers to offer the amenity in perpetuirty, whereas the stock of workers engaged in remote work is likely to be more volatile, for example during periods of closures and 'lockdowns'. Indeed, we see in Figure A.2 that ISTAT's Labor Force Survey (LFS) which collects information on remote work practices. They allow for respondents to select 'no WFH', 'sometimes WFH', and 'usually WFH'. We see that 'sometimes' work from home has a similar dynamics to our vacancy-based measure, rising during the pandemic and sustaining at it's peak there-after. The pattern for the share of employees who 'usually' work from home has a much higher peak during the COVID-19 lockdows, declines thereafter, and reaches a much lower level in the years since the pandemic. This evidence is largely aligned with our interpretation of our vacancy-based measure as predominantly capturing hybrid work arrangements i.e. those who work 1-3 days a week at home, a finding supported by other survey based measures Barrero et al. (2021).

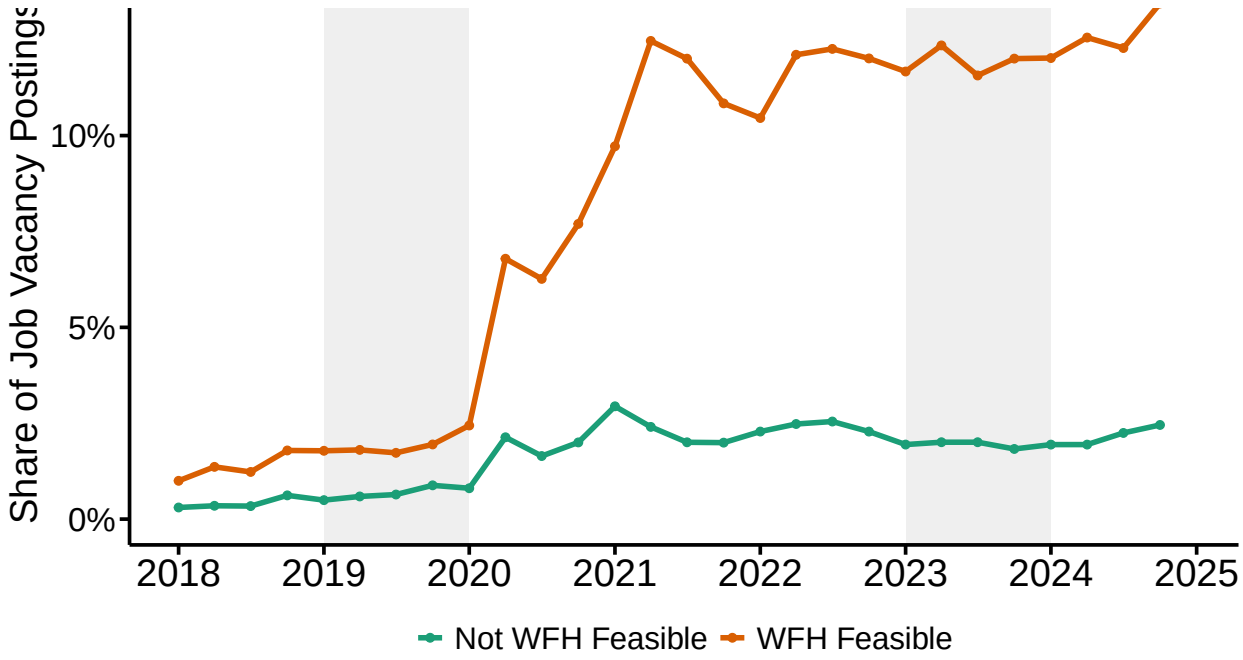
Figure A.2: Share of Dependent Employees with WFH or Hybrid jobs



A.2 Comparison of feasibility classifications to text-based measures of remote work arrangements

A number of studies have sought to measure remote work based on the task-content of occupations (Basso et al., 2020; Dingel & Neiman, 2020). These expert assessments of 'WFH feasibility' provide a binary classification of occupations for which remote work is feasible, based on , Second, they pertain to the ability, but not the obligation, to remote work, meaning market forces governing the provision of the amenity to work from home may be missing. Despite these large methodological differences, we show in Figure A.3 that the share of vacancies offering remote work are very well explained by the classification of Italian occupation-categories by **??**. The set of occupations these authors classify as note feasible for remote work shows almost no increase in the share of WFH postings post-pandemic, whereas the share of job ads offering remote work for those occupations grouped as 'WFH feasible' grew by a factor of five.

Figure A.3: Rising Share of Remote-work Postings since Pandemic is well predicted by 'WFH Feasibility' classification



Note: This figure reports the share of online job vacancy postings explicitly offering remote arrangements, broken into two groups of occupations, by quarter. The groupings come from Basso et al (2022), who reviewed the task-content of ISCO 3-digit occupations to assess whether work is feasible for remote work or not feasible. The sample includes 13,616,201 observations in total, where both an occupation code was present and raw textual description allowed for a classification of remote work using our LLM classification algorithm. Quarters in 2019 and 2023 are shaded as these subsets are used in our econometric application. Source of the raw data: Lightcast.

B Further remarks on measuring labor market power

Practical notes. Our data cleaning routine for the financial variables consists of trimming observations in the top and bottom 1% of the labor share of revenues and the ratios of capital to labor and intermediates to labor. Percentiles used for this procedure are computed separately for different groups of firms based on their broad industry (NACE letters) and size class (<10 employees; 10–249 employees; 250+ employees).

For revenue function estimation, revenues are deflated using the Harmonized Index of Consumer Prices (HICP), from Eurostat. For production function estimation, we obtain quasi-physical quantities of output we use industry-specific gross output deflators from ISTAT. To obtain quasi-physical quantities of capital and intermediate inputs for revenue function estimation, we use input-specific industry-specific deflators from ISTAT.

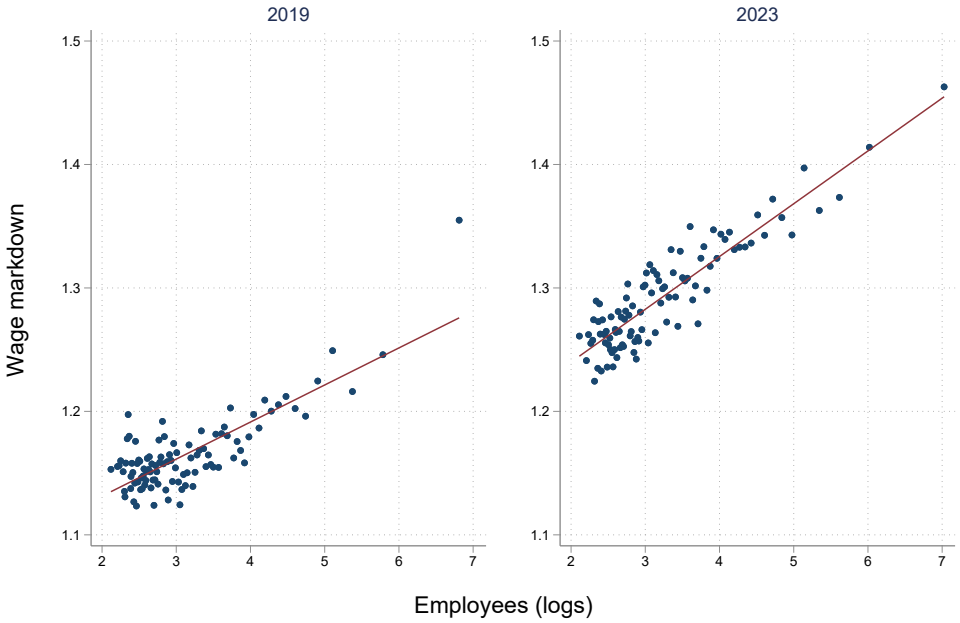
Revenue functions and production functions are estimated using firm-level data from 2015–19. Data are available up to 2023, but information in 2020–22 may not provide an accurate representation of the production process due to the COVID-19 shock as well as the related government interventions. Time-invatiant industry-specific translog parameters (see subsection 3.2 for the details) are thus estimated from 2015–19 data and then used to derive firm-level wage markdowns in the two periods of interest: 2019 and 2023.

For our preferred approach (revenue-based), translog estimations are performed at the four-digit industry level. The aggregation approach of YMH requires the estimation to be run at the same level as the aggregation. Thus, translog elasticities are estimated at the two-digit industry level. This constraint constitutes an additional reason to opt for weight-based aggregation, as the two-digit industry level might be too broad considering the large technology differences between industries.

Top and bottom 2.5% of firm-level markdowns are winsorized before market-level aggregation.

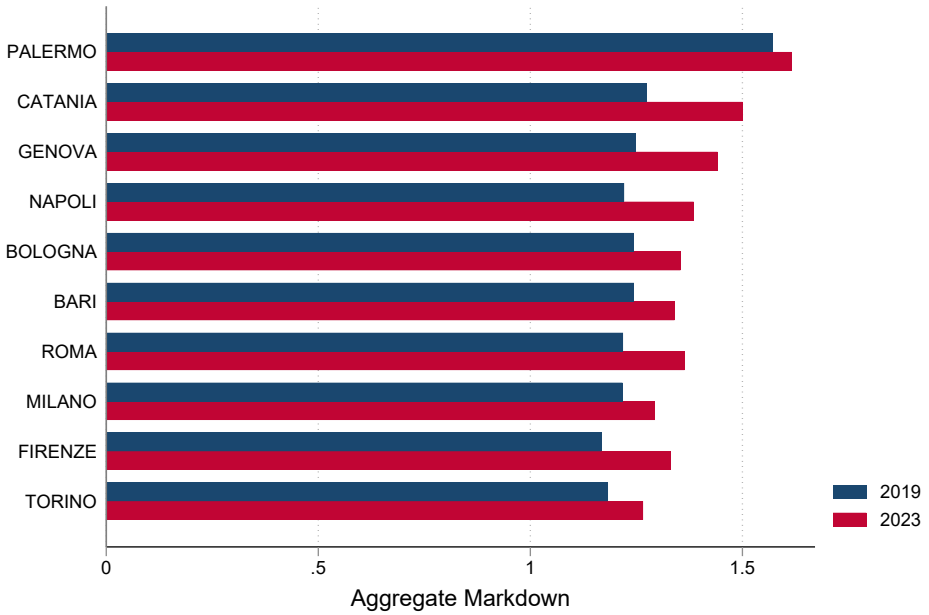
Descriptive results. Figure B.1 shows the positive firm-level relationship between wage markdowns and firm size, conditioning on a commuting zone (local labor market) fixed effect. Figure B.2 reports aggregate markdowns in the commuting zones containing the ten Italian cities with largest population.

Figure B.1: Markdowns and Firm Size



Note: This figure shows binned scatter plots with markdowns against employment at the firm-level for firms with at least 10 employees. The two panels report separate results for the two years considered in the research design. Each dot represents approximately 4,498 firms in the left panel and 3,349 firms in the panel on the right. Source of the raw data: AIDA.

Figure B.2: Markdowns Across Cities



Note: This figure reports the aggregate markdown in 2019 and 2023, the two periods considered in our research design, for a selection of commuting zones including the ten cities with largest population in Italy. Source of the raw data: AIDA.