
Monopsony and Automation

Marina Chugunova (Max Planck Institute for Innovation and Competition)

Klaus Keller (Max Planck Institute for Innovation and Competition)

Jose Azar (University of Navarra, School of Economics and Business and IESE Business School)

Sampsa Samila (IESE Business School, University of Navarra.)

Discussion Paper No. 432

October 17, 2023

Monopsony and Automation

José Azar*, Marina Chugunova[†], Klaus Keller[‡], Sampsa Samila[§]

October 11, 2023[¶]

Abstract

We examine the impact of labor market power on firms' adoption of automation technologies. We develop a model that incorporates labor market power into the task-based theory of automation. We show that, due to higher marginal cost of labor, monopsonistic firms have stronger incentives to automate than wage-taking firms, which could amplify or mitigate the negative employment effects of automation. Using data from US commuting zones, our results show that commuting zones that are more exposed to industrial robots exhibit considerably larger reductions in both employment and wages when their labor markets demonstrate higher levels of concentration.

Keywords: automation, employment, labor market concentration, industrial robots, wage setting.
JEL-classification: J23, J30, J42, L11, O33.

*University of Navarra, School of Economics and Business and IESE Business School. Email: jazar@iese.edu

[†]Max Planck Institute for Innovation and Competition. Email: marina.chugunova@ip.mpg.de

[‡]Ludwig Maximilian University of Munich and Max Planck Institute for Innovation and Competition. Email: klaus.keller@ip.mpg.de

[§]IESE Business School, University of Navarra. Email: ssamila@iese.edu

[¶]Marina Chugunova and Klaus Keller gratefully acknowledge support by DFG CRC TRR 190 (project number 280092119).

1 Introduction

Over the last decades, firms in the United States have steadily increased their use of artificial intelligence and other forms of automation (Acemoglu and Restrepo, 2020; Alekseeva et al., 2021; Zolas et al., 2021). At the same time, US labor markets have become increasingly concentrated, especially in manufacturing (Benmelech et al., 2022; Rinz, 2020). Both automation and labor market power may have contributed to increasing labor market inequality (Acemoglu and Restrepo, 2021; Yeh et al., 2022). Yet, they might be intertwined, and both the level of automation and its effects on employment and wages could depend on the characteristics of the labor market, including the intensity of competition for workers.

In this paper, we study theoretically and empirically the interrelation between automation and labor market power. We provide empirical evidence that labor market power can amplify the negative effects of automation on employment and wages. We also propose a theory that can help make sense of these empirical findings. If a firm is a wage-taker in the labor market – that is, it has no monopsony power – automating its processes would result in a reduction of its wage expenses due to hiring fewer workers. However, if a firm has monopsony power, the impact of automation on the firm's total wage bill consists of two components: the reduction in wages due to fewer workers being hired, *and* the negative impact of automation on the wages of the remaining workers. That is, since the monopsonistic firm faces an upward-sloping labor supply, automating a marginal worker enables the firm to pay lower wages for the infra-marginal workers as well.

We build a model that formalizes these ideas. The model adds labor market power to the task-based theory of automation of Acemoglu and Restrepo (2018c). In this model, a firm must choose which tasks will be performed by humans, and which tasks it will automate with machines. There exists a threshold beyond which tasks cannot be automated with the existing technology. Technological change can be modeled in two ways: either an increase in productivity of labor or capital for their existing tasks, or an increase in the automation threshold, which allows more tasks to be automated. Labor market power arises as the result of jobs being differentiated, as in Berger et al. (2022).

We first show that, if we start from an equilibrium where the competitive level of automation is below the threshold of automatable tasks, an increase in monopsony power leads to an increase in automation. We then examine whether higher labor market power amplifies or mitigates the effect of an increase in the automation threshold on employment. We show that the effect is ambiguous.

On the one hand, if the automation threshold is binding in both the high and low labor market power economies, increasing the threshold would result in stronger reductions in employment in the economy with *lower* labor market power. On the other hand, if the threshold is not binding in the low labor market power economy, but it is binding in the high labor market power one, increasing the threshold would result in stronger reductions in employment in the economy with *higher* labor market power.¹ We show in numerical simulations that this effect tends to make the expected effect of automation on employment stronger when labor market power is high.

We test the predictions of our model in an empirical setting using data on robot adoption, US local labor market employment, wages, and concentration (Acemoglu and Restrepo, 2020). We first replicate the main results of Acemoglu and Restrepo (*ibid.*), and then explore the heterogeneity with respect to labor market concentration, measured by the Herfindahl-Hirschman index of employer concentration for each industry by commuting zone. This index serves as a proxy for the degree of labor market power within the local labor market. Our results show that commuting zones that are more exposed to industrial robots exhibit considerably larger reductions in both employment and wages when their labor markets demonstrate higher levels of concentration. This is consistent with the model predictions for the case when the automation threshold is binding in the high labour market power economy but not in the low one.

Our research contributes to several distinct strands of literature. First, we extend the recent work on the labor market effects of robots by introducing labor market power in the task-based framework of automation by Acemoglu and Restrepo (2018b) and by testing the implications empirically. We build on the canonical model allowing firms to possess wage-setting power due to upward-sloping labor supply curves and demonstrate that such wage-setting power can affect firms' equilibrium level of automation. We identify conditions under which firms could engage in excessive automation (Acemoglu and Restrepo, 2018a,c) from the stand point of a social planner. Thereby, we can revisit the assumption that firms take wages as given when deciding to automate (Acemoglu and Restrepo, 2020; Adachi et al., 2020; Bessen et al., 2022; Koch et al., 2021). Empirically, we find that the negative impacts of industrial robots on employment and wage growth in US commuting zones are amplified in more concentrated local labor markets where firms hold more wage-setting power. This finding is important, considering that US labor markets show considerable variation in monopsony power (Azar

¹As we will show, it is not possible for the threshold to be binding in the low labor market power case, but not binding in the high labor market power case, because an increase in labor market power shifts the MPL/MPK curve to the right.

et al., 2022; Berger et al., 2022; Yeh et al., 2022).

Second, we contribute to a recent research on the determinants of automation (Acemoglu and Restrepo, 2022; Danzer et al., 2020; Dechezleprêtre et al., 2021; Fan et al., 2021) by showing that idiosyncratic differences in firms labor market power can affect incentives to automate. The existing literature has predominantly focused on regional differences in labor markets like demographic change, migration, or minimum wage policies. We offer a novel explanation for the heterogeneity in robot adoption that has consistently been documented in recent micro-data studies (Brynjolfsson et al., 2023; Deng et al., 2021).

Third, we further the literature on the wage and employment implications of monopsony (Manning, 2021; Robinson, 1933; Sokolova and Sorensen, 2021). While previous research has primarily concentrated on the direct effects of labor market power on wages and employment, our study extends this perspective. We show that labor market power can also indirectly influence labor markets through its impact on firms adoption of automation technologies, thus influencing wage and employment growth.

Fourth, there has been a growing literature on the connection between automation and market concentration recently. Firooz et al. (2022) find empirical evidence suggesting that automation plays a role in augmenting sales concentration within US industries, without notable repercussions on employment concentration. Concurrently, Leduc and Liu (2022) posits that the prospect of workforce displacement due to automation can bolster employers bargaining power, subsequently dampening real wage growth in a business cycle boom. We extend this nascent literature on automation and market concentration by shedding light on the reverse relationship: specifically, how initial differences in labor market concentration may impact the optimal level of automation, and its effects on labor market outcomes.

The paper is structured as follows: Section 2 presents the model and core findings, Section 3 discusses the empirical approach, Section 4 presents empirical results, and Section 5 concludes.

2 A Model of Task-Based Production and Monopsony

2.1 Households

Consider a representative household which derives utility from aggregate consumption C , which is a bundle of the different consumption goods produced by firms and disutility from labor. Its preferences over consumption and labor are represented by the following utility function:

$$U(C, L) = C \frac{1 - L^{1+\frac{1}{\varphi}}}{-\frac{1}{\varphi} \frac{1}{1+\frac{1}{\varphi}}} \quad (2.1)$$

where aggregate consumption C and labor L are bundles of firm-level consumption and labor, given by

$$C = \int_0^1 c_j dj \quad (2.2)$$

$$L = \int_0^1 l_j^{\frac{\theta+1}{\theta}} dj^{\frac{\theta}{\theta+1}} \quad (2.3)$$

Thus, all firms produce a homogeneous consumption good and there is perfect competition in the product market. However, jobs that firms offer are differentiated (with a constant-elasticity of substitution across jobs), which gives them some degree of monopsony power in the labor market.

The household has a fixed capital endowment K , which it rents out to the firms at an endogenous rate R . Capital is undifferentiated, and therefore firms are perfectly competitive in the capital market. The household supplies labor and capital l_j and k_j to firm j , and obtains profits π_j from firm j . The price of the consumption good is the numeraire. The household's budget constraint is

$$C = \int_0^1 w_j l_j dj + RK + \int_0^1 \pi_j dj \quad (2.4)$$

The first-order condition with respect to l_j yields the inverse labor supply function to firm j which takes aggregate labor supply L as given:

$$w_j = \frac{1}{-\frac{1}{\varphi}} L^{\frac{1}{\varphi}} \frac{l_j}{L}^{\frac{1}{\theta}} \quad (2.5)$$

We can rewrite this in terms of the wage index $W = \int_0^1 w_j^{1+\theta} dj$ as follows:

$$w_j = \frac{l_j}{L} W^{\frac{1}{\theta}} \quad (2.6)$$

This is the inverse labor supply curve faced by firm j given aggregate labor L and the aggregate wage level W .

2.2 Firms

There is a continuum of firms with measure one, indexed by j . Firm j produces good c_j . Aggregate output for firm j is produced through a continuum of tasks, indexed by i :

$$y_j = \int_0^1 y_j(i)^{\frac{\sigma-1}{\sigma}} di \quad (2.7)$$

where σ is the elasticity of substitution between tasks. There is a threshold of automatable tasks I . Tasks $i > I$ can be produced with labor according to $y_j(i) = l_j(i)$. Tasks $i < I$ can be produced with labor or capital: $y_j(i) = k_j(i) + l_j(i)$. We assume that $k_j(i)$ is strictly increasing in i .

Wages are endogenous, and the firm is a monopsonist with respect to the wage paid to its own workers, and does not discriminate between workers. At the same time, it is a wage-taker with respect to the aggregate wage index W . The rate of return on capital R is also endogenous, and the firms are price-takers in the capital market.

The firm faces an upward-sloping labor supply curve $w_j(l_j)$ with constant elasticity θ . The overall amount of labor that the firm demands is the integral of the labor that it demands across tasks, that is $l_j = \int_0^1 l(i) di$. We also define $k_j = \int_0^1 k_j(i) di$.

2.3 Equilibrium and Comparative Statics

The profit maximization problem of firm j is

$$\max_{l(i), k(i)} y_j - Rk_j - w_j(l_j)l_j \quad (2.8)$$

There is a unique level of automation I at which the firm would be indifferent between automating a task or have it done by humans. In the competitive case ($\theta = 1$), that level of automation is given by

$$\frac{w_j}{R} = \frac{I}{I} \quad (2.9)$$

In the monopsonistic case ($\theta < 1$), the automation level takes into account the extra cost due to the fact that hiring an additional worker increases wages for all existing workers as well

$$\frac{w_j(1 + \frac{1}{\theta})}{R} = \frac{I}{I} \quad (2.10)$$

If $I > I$, the firm cannot produce all tasks up to I with capital because some of the tasks below I are not yet automatable. Thus, the equilibrium level of automation is $I = \min I, I$

For tasks below the automation threshold I , the first-order condition with respect to $k_j(i)$ is

$$y_j(i) = \frac{(i)}{R} y_j \quad (2.11)$$

For tasks above I , the first-order condition with respect to $l_j(i)$ is

$$y_j(i) = \frac{(i)}{w_j(l_j) (1 + \frac{1}{\theta})} y_j \quad (2.12)$$

We can solve each of these for y_j and replace it in the production function of firm j to obtain a production function in terms of aggregate labor and capital for the firm:

$$y_j = F(k_j, l_j) = \int_0^{I^*} (i)^{\frac{1}{\sigma}} k_j^{\frac{\sigma-1}{\sigma}} di + \int_{I^*}^1 (i)^{\frac{1}{\sigma}} l_j^{\frac{\sigma-1}{\sigma}} di \quad (2.13)$$

To simplify the expression, we define the productivity factors of capital and labor as $A_K = \int_0^{I^*} (i)^{\frac{1}{\sigma}} di^{\frac{1}{\sigma-1}}$ and $A_L = \int_{I^*}^1 (i)^{\frac{1}{\sigma}} di^{\frac{1}{\sigma-1}}$ so that we can rewrite the production function as $y_j = F(k_j, l_j) = (A_K k_j)^{\frac{\sigma-1}{\sigma}} + (A_L l_j)^{\frac{\sigma-1}{\sigma}}$

The marginal product of firm j 's capital bundle is equal to the cost of capital R :

$$F_K(k_j, l_j) = A_K \frac{y_j}{A_K k_j}^{\frac{1}{\sigma}} = R \quad (2.14)$$

The marginal product of firm j 's labor bundle is equal to the marginal cost of labor to the firm:

$$F_L(k_j, l_j) = A_L \frac{y_j}{A_L l_j}^{\frac{1}{\sigma}} = w_j(l_j) + w(l_j)l_j = w_j(l_j) \left(1 + \frac{1}{\sigma}\right) \quad (2.15)$$

Imposing symmetry in the first-order condition for labor and combining it with the aggregate inverse labor supply $W = (L^{-1})^{\frac{1}{\varphi}}$ yields a nonlinear equation in aggregate labor (conditional on a level of automation). Although the equation does not have a closed-form solution, we can use it to characterize the equilibrium conditional on the level of automation. Equilibrium employment conditional on the level of automation implies an MPL/MPK curve as a function of the level of automation i . The equilibrium level of automation is given by the intersection of this curve and the $(i) = (i)$ curve, or the automation threshold if the latter is lower.

We can also show that automation increases (though not strictly) when labor market power is higher. This is because the schedule given by the equilibrium ratio of MPL and MPK conditional on a level of automation shifts to the right when labor market power increases. A shift to the right in the MPL/MPK schedule implies a higher intersection between this curve and the $(i) = (i)$ curve, and implies a higher level of desired automation I . This is summarized by the following Proposition, and illustrated in Figure 1. If we start at point a in the figure, and shift the MPL/MPK curve to the right, equilibrium automation increases until I hits the maximum possible level of automation, given by the threshold I . Further increases in market power would not increase automation, because I would be above the threshold.


Proposition 2.1. *An equilibrium of the model exists and is characterized by the intersection between the $(i) = (i)$ schedule and the $\frac{MPL}{MPK}(i)$ schedule which indicates the ratio of MPL and MPK conditional on the level of automation (ignoring the automation threshold). If the intersection of the two curves is below the automation threshold I , then the equilibrium level of automation I^* is given by I , i.e., the level of the intersection. If the intersection of the two curves is above the automation threshold, then the equilibrium level I^* is equal to the threshold I .*

The equilibrium level of employment is characterized by the solution to the following equation in L at I :

$$A_L \left(1 + \frac{A_K K}{A_L L}\right)^{\frac{\sigma-1}{\sigma}} = \frac{L}{L}^{\frac{1}{\varphi}} \left(1 + \frac{1}{\sigma}\right) \quad (2.16)$$

If the equilibrium level of automation is below the threshold I , then an increase in labor market

Figure 1: Equilibrium Automation under Low and High Levels of Labor Market Competition



Note: This figure illustrates the determination of the equilibrium level of automation when the labor market power parameter is θ_H (less labor market power), and when it is θ_L (more labor market power). If labor market power is low, at $\theta = \theta_H$, the equilibrium level of automation is at point a , at the intersection of the MPL/MPK curve and the $\frac{\gamma(i)}{\eta(i)}$ curve, which is below the automation threshold. If labor market power increases, so that $\theta = \theta_L$, the MPL/MPK curve shifts to the right, and the intersection is at point b . However, because tasks above I are not automatable, the equilibrium is in point c , at the intersection of the MPL/MPK curve and the vertical line that indicates automation level I . Thus, the equilibrium level of automation I is higher than in the low labor market power case.

power ($\theta_L < \theta_H$) increases automation. If the equilibrium level of automation is at the threshold I , then an increase in labor market power leaves automation unchanged.

We are also interested in the effect of technological progress, modeled as an increase in the set of automatable tasks, on employment and wages, and in particular in the heterogeneity of this effect by the level of labor market power. Consider two economies that are identical, except that in one of them labor market power is low (i.e., the elasticity of substitution parameter is high, $\theta = \theta_H$) and in the other labor market power is high (i.e., the elasticity of substitution parameter is low $\theta = \theta_L < \theta_H$). We can compare the response of employment and wages to an increase in the automation threshold I in these two economies.

There are three possible cases. (1) If the automation threshold is not binding for both the θ_L and

H economies, an increase in the threshold has no effect in either economy. (2) If the automation threshold is binding in both, a marginal increase in the automation threshold increases automation by the same amount in both. When this happens, it is possible for the effect of an increase in the automation threshold on employment and wages to be stronger in the more competitive economy. (3) It is possible that the automation threshold is not binding for the low labor market power economy, but is binding in the low labor market competition economy (see Figure 1). In this case, an increase in the automation threshold has no effect on the equilibrium with low labor market power, that is, at point a , because the equilibrium is below the automation threshold. However, it does affect the equilibrium for the high labor market power economy, at point c , because it is at the threshold.


In the Appendix, we show that the effect of an increase in the automation threshold I on employment and wages can be negative or positive, depending on whether the (negative) displacement effect dominates the (positive) productivity effect (as in Acemoglu and Restrepo, 2018b). Regardless of the sign of the effect on employment and wages, if we are in case (2), the effect can be stronger (i.e., higher in absolute value) in the more competitive labor market, while in case (3) the effect can be stronger in the more monopsonistic labor market. The overall effect is therefore ambiguous. If the automation threshold has high variance, it is more likely that it will bind when labor market power is higher, and the mechanism in case (3) is more important. However, if the automation threshold does not vary across labor markets, the mechanism in case (2) is more important. To illustrate this, we provide a Monte-Carlo simulation of the expected marginal effect with a stochastic automation threshold in the following subsection.

2.4 Simulation

We solve the model numerically for the following parameter values: $\beta = 0.7$, $\alpha = 0.5$, $\tau = 1$, $K = 1$, $\phi(i) = 1$, $\psi(i) = e^{Ai}$, with $A = 1$, and for a range of values for μ between 1 and 5. For each value of μ , we take 10,000 random draws for the automation threshold I . We do it first for a “low I dispersion case”, with I uniformly distributed in the interval $[0.33, 0.331]$, in which case the threshold is always binding for this range of μ . We then run the simulation for a high I dispersion case, uniform over $[0.33, 0.45]$, which approximately covers the range of I for our range of μ parameters.

For each draw, we calculate the derivative with respect to I (which can be zero if the threshold is non-binding), and take expectation across draws for each value of μ . The results are plotted in Figure

Figure 2: Simulated Average Effect of Automation Threshold on Log Employment as a Function of Labor Market Power



Note: This figure shows the average marginal effect on log employment with respect to the automation threshold I . The model is solved numerically for the following parameter values: $\sigma = 0.7$, $\varphi = 0.5$, $\bar{\varphi} = 1$, $K = 1$, $\eta(i) = 1$, $\gamma(i) = e^{Ai}$, with $A = 1$, and for a range of values for θ between 1 and 5. For each value of θ , we take 10,000 random draws for the automation threshold I . In the low dispersion case, I is uniform in the interval $[0.33, 0.331]$, and in the high dispersion case it is uniform in the interval $[0.33, 0.45]$ (about the range of \bar{I} corresponding to our range of θ s).

² In the low dispersion case, the mechanism described in case (2) dominates, and the expected effect of automation on employment is less negative when labor market power is higher. In the high dispersion case, the mechanism in case (3) dominates, and the expected effect of automation on employment (and therefore on wages as well) is more negative when labor market power is higher. As we argue above, the automation threshold is more likely to be binding when labor market power is higher.

This theoretical ambiguity implies that it is an empirical question whether labor market power amplifies or mitigates the labor market effects of automation. The next section describes our methodology to examine this question empirically.

3 Empirical Methodology

To test whether the effect of automation on employment and wages is stronger when labor market power is higher, we make use of the empirical framework presented by Acemoglu and Restrepo (2020). This framework allows to study the long-run equilibrium adjustments of local labor markets

²Note that in all cases, the wage effect goes in the same direction as the employment effect, due to increasing aggregate labor supply.

in the United States in response to changes in labor demand driven by advancements in industrial robot technology.

3.1 Measuring local labor market exposure to robots

Based on the study by Acemoglu and Restrepo (2020), we measure the exposure to industrial robots in 722 continental US commuting zones over the period 1990 to 2015.³ To approximate global advancements in industrial robot technology, we examine changes in the number of industrial robots per worker across 19 different industries in five European countries (Denmark, Finland, France, Italy, Sweden) that have been ahead of the US in adopting robot technology. For each industry i , we compute a measure of average robot penetration over the period from t_0 to t_1 as the average change in the stock of industrial robots relative to the total number of workers in that industry in 1990 subtracting the growth of robot stocks that is due to real output growth. This measure is given by

$$\text{APR}_{i(t_0 t_1)}^{EU5} = \frac{1}{5} \sum_{j \in EU5} \frac{R_{j i t_1} - R_{j i t_0}}{L_{j i 1990}} - g_{j i(t_0 t_1)} \frac{R_{j i t_0}}{L_{j i 1990}} \quad (3.1)$$

where $R_{j i t}$ is the number of robots in industry i in country j at time t , $g_{j i(t_0 t_1)}$ is the output growth rate of industry in country j between t_0 and t_1 and $L_{j i 1990}$ is the total number of workers in industry i in country j in 1990.⁴

To obtain a measure of commuting zone exposure to robots, we finally multiply the industry-specific changes in average robot penetration in the five European countries by $l_{c i 1970}$, the share of industry i in the total employment of commuting zone c in 1970.

$$\text{Robots}_c(t_0 t_1) = \sum_{i \in I} l_{c i 1970} \text{APR}_i^{EU5}(t_0 t_1) \quad (3.2)$$

To compute the shares of industries in commuting zone employment, we make use of micro-data from the decennial census of 1970 as in Acemoglu and Restrepo (*ibid.*). We also use the micro-data from the decennial censuses of 1990 and 2000 combined with micro-data from the American Community Survey to compute outcomes variables in terms of employment, unemployment, non-

³Commuting zones are clusters of counties in which the majority of workers both live and work. This geography is typically used to delineate local labor markets (see Tolbert and Sizer, 1996).

⁴We use data on the industry-specific stocks of industrial robots from the International Federation of Robotics, while both output and employment data for industries in Europe comes from the EUKLEMS database.

participation as well as average wages for our main analysis.

The identifying assumption of the empirical strategy is that there are no differential shocks or trends affecting labor market outcomes in commuting zones with greater exposure to robots relative to those with less exposure (such as differential trends of industries or regions or negative shocks that are correlated with the increasing exposure to robots and local labor market conditions).⁵

3.2 Measuring local labor market concentration

We extend the empirical framework by Acemoglu and Restrepo (2020) taking into account the competitiveness of local labor markets at the beginning of the observation period. We proxy the competitiveness of local labor markets with the degree of employer concentration within a commuting zone. A recent body of literature shows that measures of employer concentration reflect the extent to which firms face a more or less elastic labor supply curve in the local labor market (e.g. Berger et al., 2022). We therefore utilize the local employer concentration as a proxy for the labor supply elasticity (ϵ) in our model that describes the competitiveness of the local labor market. This allows us to explore how initial differences in local labor market competitiveness impact the effects of improvements in robot technology across different commuting zones.

We compute a measure of local labor market concentration for all 722 continental commuting zones in 1990 using data on county-by-industry establishment counts from the US Census *County Business Patterns*, county industry employment counts from Eckert et al. (2021) and a county-to-commuting zone crosswalk provided by David and Dorn (2013). In each commuting zone, we observe the number of establishments n in a 3-digit SIC industry i by employment bracket s in commuting zone c in 1990.⁶ We take the mid-point m of an employment bracket s as a proxy for the actual employment size of establishments assigned to employment bracket s . We then compute employment shares and the Herfindahl-Hirschman index of employer concentration for each industry i in a commuting zone c as

$$HHI_{ci} = \sum_{s=1}^{12} n_{cis} \left(\frac{m_s}{L_{ci}} \right)^2 \quad (3.3)$$

where L stands for the total employment of industry i in commuting zone c . As the level of analysis will eventually be at the commuting-zone level, we further aggregate the industry-by-commuting zone

⁵See Acemoglu and Restrepo, 2020 for a comprehensive check of the validity of the proposed measure of commuting zone exposure to robots.

⁶County Business Patterns reports 12 employment brackets which are described in detail in Table 3 in the Appendix.

level HHIs to the commuting zone level. To calculate the average level of employer concentration for each commuting zone, we compute a weighted mean of all industry employer Herfindahl indices as

$$\overline{HHI}_c = \sum_{i=1}^{395} l_{ci} HHI_{ci} \quad (3.4)$$

where l is the share of industry i in total employment of the commuting zone c . As in Benmelech et al. (2022), this average HHI at the level of the commuting zone represents the degree of employer concentration that the *average* worker faces in a given local labor market.⁷

3.3 Empirical Specifications

We explore the heterogeneous effect of local labor market exposure to robots on employment, unemployment and non-participation rates across labor markets with different initial employer concentration. We estimate the following model in three stacked differences over three periods from 1990 to 2000, 2000 to 2007 and 2007 to 2015:

$$y_{c(t_0 t_1)} = \beta_1 \text{Robots}_{c(t_0 t_1)} + \beta_2 \text{Robots}_{c(t_0 t_1)} \overline{HHI}_{c 1990} + \beta_3 \overline{HHI}_{c 1990} + \mathbf{X}_{c 1990} + \alpha_t + \gamma_j + \delta_{c(t_0 t_1)} \quad (3.5)$$

In our main specification, $y_{c t}$ stands for the log number of private sector employees in commuting zone c in year t and $\overline{HHI}_{c 1990}$ is the continuous measure of local labor market concentration of commuting zone c in 1990. The coefficient of interest is β_2 which captures the heterogeneous effect of robots on employment across commuting zones with different initial levels of local labor market concentration $\overline{HHI}_{c 1990}$. The sign of the coefficient β_2 allows us to infer whether the automation threshold is binding for monopsonists but not for competitive firms. If this was indeed the case, improvements in automation technology would lead to more negative employment effects in more concentrated labor markets. We keep $\overline{HHI}_{c 1990}$ fixed to initial levels in 1990 to avoid any endogeneity between increasing automation and *contemporaneous* changes in labor market concentration.

Following Acemoglu and Restrepo (2020), we control for unobserved period-specific regional trends by including dummies for census divisions γ_j and period indicators α_t . Hence, our regression identifies the coefficients β_2 from variation in exposure to labor market shocks between commuting

⁷Figure 4 in the Appendix displays the regional variation in labor market concentration across the 722 US commuting zones in 1990, attributing the value for each commuting zone to one out of seven equal-sized bins.

zones in a given time-period and census division and variation in the ex-ante local labor market concentration. We also include $\mathbf{X}_{c,1990}$, a vector of commuting zone baseline characteristics in 1990, to allow for differential trends due to observable differences in demographics (age, education, gender and ethnic composition), industry shares (manufacturing, light-manufacturing) or in the exposure to Chinese import competition and offshoring (share of routine employment).

We also explore the heterogeneous effect of local labor market exposure to robots on wages across labor markets with different initial employer concentration. We compute for each commuting zone the average hourly, weekly and yearly wages of workers within 250 demographic cells defined by gender, education, age and race.⁸ By looking at wages within defined demographics, we can control for changes in wages in the commuting zone that are driven by changes in the characteristics of the work force such as age. We estimate the following model at the demographic group by commuting zone level in two stacked differences over the periods from 1990 to 2000 and from 2000 to 2007:

$$y_{cd}(t_0, t_1) = \beta_1 \text{Robots}_{c(t_0, t_1)} + \beta_2 \text{Robots}_{c(t_0, t_1)} \overline{\text{HHI}}_{c,1990} + \beta_3 \overline{\text{HHI}}_{c,1990} + \mathbf{X}_{c,1990} + d + t + j + \epsilon_{cd}(t_0, t_1) \quad (3.6)$$

where y stands for the log average wage of workers in a demographic cell d in commuting zone c and year t . In addition to the dummies for census divisions j , the period indicators t , the commuting zone characteristics in 1990 $\mathbf{X}_{c,1990}$, we also include a dummy for each demographic cell that corrects for differential long-run trends in wages across the different demographics. Again, we are interested in the coefficient of the interaction term β_2 which reflects the heterogeneous effect of robots on average wages across commuting zones with differential initial levels of labor market concentration.

4 Results

In our empirical results we, first, document the negative effect of robot exposure on commuting zone employment consistent with Acemoglu and Restrepo (2020) and reveal that the negative employment effect is substantially more pronounced in initially more concentrated labor markets. Next, we find heterogeneous effects of robots on measures labor force participation. Last, we document that labor market concentration also moderates the negative effect of robots on average wages.

⁸See Appendix B.3 for more details on the computation of average wages.

4.1 Employment

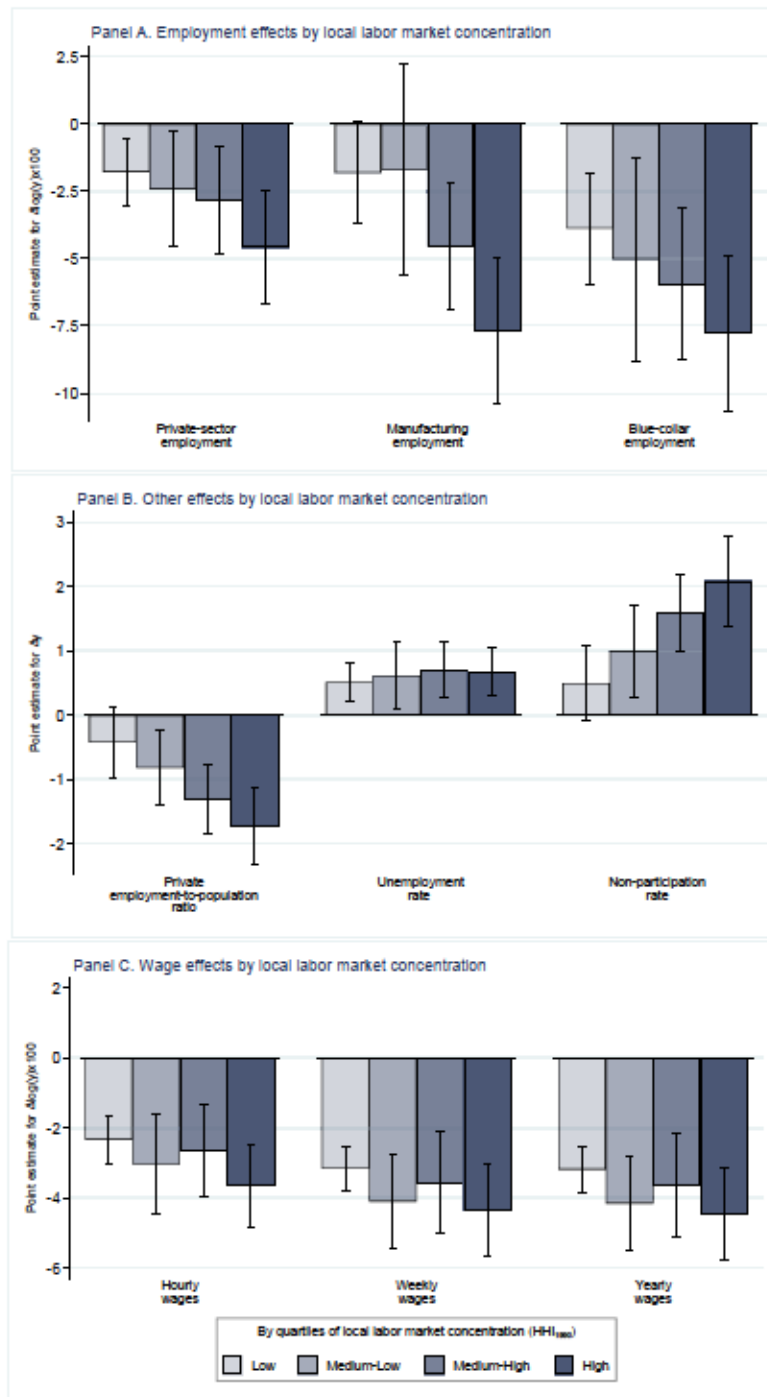
Consistent with previous evidence by Acemoglu and Restrepo (2020) and Faber et al. (2022), column (1) in Panel A of Table 1 shows that an increase of commuting zone exposure of 1 robot per thousand workers decreases total employment by about 2 percent. However, column (2) reveals significant heterogeneity in the effect by initial levels of labor market concentration in 1990. It shows that the coefficient of the interaction of robot exposure and the continuous variable of local labor market concentration in 1990 is significantly negative and sizable in magnitude. The estimates in column (2) imply that the effect of robots on total employment is about 60% stronger in commuting zones at the 75th percentile of labor market concentration distribution ($\overline{\text{HHI}}_{1990}^{75th} = 0.44$) compared to commuting zones at the 25th percentile ($\overline{\text{HHI}}_{1990}^{25th} = 0.22$). We can also observe the strong moderating effect of labor market concentration when looking at manufacturing and blue-collar occupation reported in columns (4) and (6) which are arguably most exposed to industrial robots.

We further corroborate the relationship estimating the mean effects of robots in commuting zones in different quartiles of the distribution of labor market concentration in 1990.⁹ Panel A in Figure 3 shows the point estimates by concentration quartile and confirms the previous finding showing consistently stronger effects in commuting zones in the upper quartiles of the concentration distribution for total employment as well manufacturing and blue-collar employment.

The combined evidence suggests that improvements in automation technology seem to affect employment more negatively in more concentrated local labor markets. While there is no data on robot adoption across commuting zones allowing us to test mechanisms directly, the obtained results are in line with case 3 presented in Section 2.3. The observed pattern suggests that firms in more concentrated labor markets are bound by the automation threshold and are more likely to automate as automation technology improves, amplifying the effect on employment. On the other hand, firms in less concentrated labor markets are less likely to be bound and do therefore react less as automation technology improves.

⁹We order commuting zones by the level of labor market concentration in 1990 and group them into 4 bins that each contain one quarter of the total US population in 1990.

Figure 3: Local Labor Market Effects of Exposure to Robots by Local Labor Market Concentration



Note: The figure displays coefficient estimates of the impact of robots on labor market outcomes for each quartile of commuting zones based on the employment-weighted distribution of labor market concentration in 1990. For each outcome variable, we estimate a single regression model. The displayed coefficients are obtained from the interaction of quartile dummies with the main explanatory variable: exposure to robots. The capped lines indicate 95% confidence intervals. All estimates in Panel A and B are from specifications that include the full set of controls from Table 1. All estimates in Panel C are from specifications that include the full set of controls from Table 2.

4.2 Labor force participation

Next, we explore how lower employment growth due to robots might lead to higher unemployment rates and non-participation in the labor force in more concentrated regions. Column (1) in Panel B of Table 1 shows that the decrease in employment in the first part of the analysis translates into a lower employment to population ratio. We find that an increase of exposure by 1 robot per thousand workers reduced the share of employed individuals in the population of working age adults by 0.6 percent. This negative effect is again significantly more pronounced in more concentrated labor markets.

Yet, column (4) shows that the positive effect of robots on unemployment rates is not significantly different in more versus less concentrated labor markets. Column (6) shows that the reduction in the employment to population ratio due to robots leads to higher non-participation rates in more concentrated labor markets, thus providing explanation to absence of the differential effect on unemployment. This interesting pattern is consistent with recent evidence by Dodini et al. (2023) showing that more concentrated labor markets provide workers with fewer outside options which leads to higher non-participation rates and larger earnings declines for workers after involuntary job separation. Panel B in Figure 3 illustrates that the effect of robots on employment and non-participation rates is systematically more pronounced in commuting zones in the upper quartiles of the concentration distribution.

4.3 Wages

Finally, we find that local labor market concentration also moderates the effect of robots on average wages within demographic cells across commuting zones. Table 2 reports the estimates following the model specification of equation 3.6 with demographic cell fixed effects. Consistent with the results in Acemoglu and Restrepo (2020), we find in column (1) that an increase in exposure to robots by 1 robot per thousand workers decreases average hourly wages by more than 2 percent. Again, we find that this average effect masks significant heterogeneity along the dimension of labor market concentration. Estimates in column (2) imply that effect of robots on average hourly wages is 35 percent larger in commuting zones at the 75th percentile relative to commuting zones in the 25th percentile of local labor market concentration. This pattern is also true for average weekly and yearly wages. Panel C in Figure 3 corroborates again a systematic pattern showing that the mean effect on wages tends to be stronger among commuting zones in the upper quartiles of the concentration distribution.

Table 1: Local labour market concentration and the effect of exposure to robots on changes in employment, unemployment and non-participation: stacked differences from 1990 to 2000, 2000 to 2007 and 2007 to 2015

Panel A. Change in log employment 100						
	Total employment		Manufacturing employment		Blue-collar employment	
	(1)	(2)	(3)	(4)	(5)	(6)
Robots	-2.253 (0.494)	-1.240 (0.778)	-2.832 (0.659)	-0.556 (1.338)	-4.598 (0.772)	-2.979 (1.251)
Robots $\overline{\text{HHI}}_{1990}$		-12.855 (4.829)		-29.001 (7.956)		-20.061 (6.678)
$\overline{\text{HHI}}_{1990}$		0.561 (4.779)		-0.229 (7.407)		6.506 (6.215)
Observations	2166	2166	2166	2166	2166	2166
R-squared	0.44	0.44	0.37	0.38	0.42	0.43
Panel B. Change in rate						
	Employment to population ratio		Unemployment rate		Non-participation rate	
	(1)	(2)	(3)	(4)	(5)	(6)
Robots	-0.684 (0.117)	-0.125 (0.329)	0.554 (0.130)	0.505 (0.148)	0.815 (0.143)	0.144 (0.339)
Robots $\overline{\text{HHI}}_{1990}$		-6.729 (1.977)		0.559 (0.585)		8.130 (2.113)
$\overline{\text{HHI}}_{1990}$		4.630 (1.281)		-0.776 (0.518)		-4.975 (1.294)
Observations	2166	2166	2166	2166	2166	2166
R-squared	0.30	0.34	0.22	0.22	0.51	0.55
Demographics	✓	✓	✓	✓	✓	✓
Industry Shares	✓	✓	✓	✓	✓	✓
Routineness & Trade	✓	✓	✓	✓	✓	✓
Census Divisions	✓	✓	✓	✓	✓	✓
Period FE	✓	✓	✓	✓	✓	✓

Note: All specifications control for the following commuting zone characteristics in 1990: demographic characteristics of commuting zones in 1990 (log population; the share of females; the share of the population over 65 years old; the shares of the population with no college, some college, college or professional degree, and masters or doctoral degree; and the shares of whites, blacks, Hispanics, and Asians), the shares of employment in manufacturing and light manufacturing and the female share of manufacturing employment in 1990, as well as the exposure to Chinese imports and the share of employment in routine jobs in 1990. Regressions also control for census division and period dummies. Standard errors are robust against heteroskedasticity and correlation within states are given in parentheses. Regressions are weighted by commuting zone population in 1990. Standard errors are clustered at the state level. Coefficients with ***, **, and * are significant at the 1%, 5% and 10% confidence level, respectively.

Table 2: Local labour market concentration and the effect of exposure to robots on changes in average wages: stacked differences from 1990 to 2000 and 2000 to 2007

	Change in log wages 100					
	Hourly Wages		Weekly Wages		Yearly Wages	
	(1)	(2)	(3)	(4)	(5)	(6)
Robots	-2.553 (0.341)	-2.132 (0.354)	-3.379 (0.334)	-3.018 (0.311)	-3.416 (0.331)	-3.031 (0.308)
Robots $\overline{\text{HHI}}_{1990}$		-5.224 (2.372)		-4.513 (2.509)		-4.796 (2.546)
$\overline{\text{HHI}}_{1990}$		-2.634 (2.847)		-2.746 (2.958)		-2.653 (2.935)
Observations	158254	158254	156402	156402	156402	156402
R-squared	0.21	0.21	0.26	0.26	0.26	0.26
Demographics	✓	✓	✓	✓	✓	✓
Industry Shares	✓	✓	✓	✓	✓	✓
Routineness & Trade	✓	✓	✓	✓	✓	✓
Census Divisions	✓	✓	✓	✓	✓	✓
Period FE	✓	✓	✓	✓	✓	✓

Note: We estimate regressions at the demographic cell – commuting zone level where we define demographic cells by age, gender, education and race. The outcome variables are log changes in the average wage by demographic cell multiplied by 100. All specifications include a dummy for each demographic cell and control for the following commuting zone characteristics: demographic characteristics in 1990 (log population; the share of females; the share of the population over 65 years old; the shares of the population with no college, some college, college or professional degree, and masters or doctoral degree; and the shares of the Asian, Black, Hispanic and White population), the shares of employment in manufacturing and light manufacturing and the female share of manufacturing employment in 1990, as well as the exposure to Chinese imports and the share of employment in routine jobs in 1990. Regressions also control for census division and period dummies. Standard errors are robust against heteroskedasticity and correlation within states are given in parentheses. Regressions are weighted by commuting zone population in 1990. Standard errors are clustered at the state level. Coefficients with ***, **, and * are significant at the 1%, 5% and 10% confidence level, respectively.

5 Conclusion

The extent to which employers exercise monopsony power in labor markets has wide-ranging implications for workers, firms and labor markets. In this paper, we argue that labor market power can in fact drive excessive demand for automation as automation at firms with labor market power affects the total wage bill in two ways: it reduces the wage bill as fewer workers are being hired, *and* it reduces the wages of the remaining workers. Therefore, if monopsony power is high enough, the firm could have an incentive to automate with a technology that is less efficient than the workers it replaces and still obtain net cost efficiencies through the reduction in the wages of the remaining employees.

We formalize this idea by incorporating labor market power from the differentiated jobs model by Berger et al. (2022) into the task-based theory of automation of Acemoglu and Restrepo (2018c). When the labor market is competitive, some automatable tasks are not automated because it is still more cost-effective for them to be done by humans. However, it may be privately optimal for the firm to automate them under monopsony, precisely in order to exploit its monopsony power. For this reason, marginal increases in automatable tasks may not reduce labor demand when the labor market is competitive, while reducing it in the case of a monopsonistic labor market. On the other hand, when the automation threshold is binding in both high and low labor market power economies, it is ambiguous whether labor market power amplifies or mitigates the effect on automation.

We examine this question in an empirical setting of Acemoglu and Restrepo (2020) in studying industrial automation in the US. We replicate their results and explore the heterogeneity with respect to labor market concentration, measured by Herfindahl-Hirschman index of employer concentration for each industry. We show that automation is associated with considerably larger reductions in employment and wages in more concentrated labor markets. This provides first evidence that labor market power affects firms' automation decisions.

One policy implication of our model is that minimum wage policies could decrease the incentive to automate in monopsonistic labor markets where the minimum wage is binding. As minimum wages prevent the monopsonist to pay below the wage floor, the policy flattens the marginal cost of labor curve of the monopsonist. In this way, the minimum wage alleviates the incentive to automate beyond what the competitive firm would do. Therefore, perhaps surprisingly, minimum wage policies could reduce the level of automation in monopsonistic labor markets.

Our paper raises a number of questions that call for more research. First, there is a need for a

more thorough exploration of the relationship between labor market power and automation. This requires more detailed data at the local industry or firm level. Second, we have focused on industrial automation of production tasks mostly in manufacturing industries. However, with the advent of artificial intelligence, it would be useful to extend the technological scope as companies will likely automate more and more non-production tasks outside of manufacturing.

References

- Acemoglu, Daron and Pascual Restrepo (2018a). “Artificial intelligence, automation, and work”. In: *The economics of artificial intelligence: An agenda*. University of Chicago Press, pp. 197–236.
- (2018b). “Modeling automation”. In: *AEA papers and proceedings*. Vol. 108. American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203, pp. 48–53.
- (2018c). “The race between man and machine: Implications of technology for growth, factor shares, and employment”. In: *American economic review* 108.6, pp. 1488–1542.
- (2020). “Robots and jobs: Evidence from US labor markets”. In: *Journal of Political Economy* 128.6, pp. 2188–2244.
- (2021). “Tasks, Automation, and the Rise in US Wage Inequality”. In: *National Bureau of Economic Research Working Paper Series* No. 28920.
- (2022). “Demographics and automation”. In: *The Review of Economic Studies* 89.1, pp. 1–44.
- Adachi, Daisuke, Daiji Kawaguchi, and Yukiko Umeno Saito (2020). “Robots and employment: Evidence from Japan, 1978-2017”. In:
- Alekseeva, Liudmila, José Azar, Mireia Giné, Sampsa Samila, and Bledi Taska (2021). “The demand for AI skills in the labor market”. In: *Labour Economics* 71, p. 102002.
- Azar, José, Ioana Marinescu, and Marshall Steinbaum (2022). “Labor Market Concentration”. In: *Journal of Human Resources* 57.S, S167–S199.
- Benmelech, Efraim, Nittai K Bergman, and Hyunseob Kim (2022). “Strong Employers and Weak Employees How Does Employer Concentration Affect Wages?” In: *Journal of Human Resources* 57.S, S200–S250.
- Berger, David, Kyle Herkenhoff, and Simon Mongey (2022). “Labor market power”. In: *American Economic Review* 112.4, pp. 1147–93.
- Bessen, James, Erich Denk, and Chen Meng (2022). “The Remainder Effect: How Automation Complements Labor Quality”. In:
- Brynjolfsson, Erik, Cathy Bunton, Nathan Goldschlag, J Frank Li, Javier Miranda, and Robert Seamans (2023). *The characteristics and geographic distribution of robot hubs in US manufacturing establishments*. Tech. rep. National Bureau of Economic Research.
- Danzer, Alexander, Carsten Feuerbaum, and Fabian Gaessler (2020). “Labor supply and automation innovation”. In: *Max Planck Institute for Innovation & Competition Research Paper* 20-09.
- David, H and David Dorn (2013). “The growth of low-skill service jobs and the polarization of the US labor market”. In: *American Economic Review* 103.5, pp. 1553–97.
- Dechezleprêtre, Antoine, David Hémous, Morten Olsen, and Carlo Zanella (2021). “Induced automation: evidence from firm-level patent data”. In: *University of Zurich, Department of Economics, Working Paper* 384.

- Deng, Liuchun, Verena Plümpe, and Jens Stegmaier (2021). *Robot adoption at German plants*. Tech. rep. IWH Discussion Papers.
- Dodini, Samuel, Michael F Lovenheim, Kjell G Salvanes, and Alexander Willén (2023). *Monopsony, Job Tasks, and Labor Market Concentration*. Tech. rep. National Bureau of Economic Research.
- Eckert, Fabian, Teresa C. Fort, Peter K. Schott, and Natalie J. Yang (2021). *Imputing Missing Values in the US Census Bureau's County Business Patterns*. Tech. rep. National Bureau of Economic Research.
- Faber, Marius, Andrés P Sarto, and Marco Tabellini (2022). *Local Shocks and Internal Migration: The Disparate Effects of Robots and Chinese Imports in the US*. Tech. rep. National Bureau of Economic Research.
- Fan, Haichao, Yichuan Hu, and Lixin Tang (2021). “Labor costs and the adoption of robots in China”. In: *Journal of Economic Behavior & Organization* 186, pp. 608–631.
- Firooz, Hamid, Zheng Liu, and Yajie Wang (2022). “Automation, Market Concentration, and the Labor Share”. In: *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 01–37.
- Koch, Michael, Ilya Manuylov, and Marcel Smolka (2021). “Robots and firms”. In: *The Economic Journal* 131.638, pp. 2553–2584.
- Leduc, Sylvain and Zheng Liu (2022). “Automation, Bargaining Power, and Labor Market Fluctuations”. In: *Federal Reserve Bank of San Francisco, Working Paper Series*, pp. 01–43.
- Manning, Alan (2021). “Monopsony in labor markets: A review”. In: *ILR Review* 74.1, pp. 3–26.
- Rinz, Kevin (2020). “Labor Market Concentration, Earnings, and Inequality”. In: *Journal of Human Resources* 57.S, S251–S283.
- Robinson, Joan (1933). “Imperfect Competition”. In: *An Essay on Marxian Economics*. Springer, pp. 73–81.
- Sokolova, Anna and Todd Sorensen (2021). “Monopsony in labor markets: A meta-analysis”. In: *ILR Review* 74.1, pp. 27–55.
- Stehrer, Robert, Alexandra Bykova, Kirsten Jäger, Oliver Reiter, and Monika Schwarzhappel (2019). “Industry level growth and productivity data with special focus on intangible assets”. In: *Vienna Institute for International Economic Studies Statistical Report* 8.
- Timmer, Marcel P, Mary O Mahony, Bart Van Ark, et al. (2007). “EU KLEMS growth and productivity accounts: an overview”. In:
- Tolbert, Charles M and Molly Sizer (1996). “US commuting zones and labor market areas: A 1990 update”. In:
- Yeh, Chen, Claudia Macaluso, and Brad Hershbein (2022). “Monopsony in the US Labor Market”. In: *Available at SSRN 4049993*.
- Zolas, Nikolas, Zachary Kroff, Erik Brynjolfsson, Kristina McElheran, David N Beede, Cathy Buffington, Nathan Goldschlag, Lucia Foster, and Emin Dinlersoz (2021). *Advanced technologies adop-*

tion and use by us firms: Evidence from the annual business survey. Tech. rep. National Bureau of Economic Research.

Appendix A Proofs

PROOF OF PROPOSITION 2.1

Proof. Conditional on a level of automation, equation (2.16) always has a solution, since the left-hand side is decreasing in L , and the right-hand side is increasing in L and its range is $[0, \infty)$. We can calculate the equilibrium level of employment for every possible level of automation, calculate the ratio of the marginal product of labor and the marginal product of capital, and find I as the level where the ratio of MPL to MPK is equal to the ratio of w to r . The ratio of MPL to MPK goes to infinity in the limit when zero tasks are automated, and to zero in the limit when all tasks are automated, which ensures an intersection with the ratio of w to r curve. The intersection of these two curves determines the unrestricted equilibrium level of automation I . If the resulting I is higher than the threshold of automatable tasks I , then the threshold is binding and the equilibrium I is equal to the threshold I .

To see that an increase in labor market power increases the equilibrium level of automation when the automation threshold is not binding, we need to show that the MPL/MPK schedule shifts to the right. To see why this is the case, note that the derivative of $(\log) \text{MPL/MPK}$ with respect to $\log(1 + \tau)$ taking the level of automation as given is

$$\frac{d \log(F_L / F_K)}{d \log(1 + \tau)} = \frac{-\tau}{1 + \tau - (1 - s_L)(1 + \tau)} > 0 \quad (\text{A.1})$$

□

Proposition A.1 (Displacement Effect). *The derivative of the log labor share with respect to I when the automating threshold is binding is always negative.*

Proof. We start by obtaining an expression for the labor share. Combining the first-order condition for labor and the production function in equation 2.13, yields the following expression:

$$s_{l_j} = \frac{1}{1 + \frac{A_K k_j^{\frac{\sigma-1}{\sigma}}}{A_L l_j}} \frac{1}{1 + \frac{1}{\sigma}} \quad (\text{A.2})$$

Since all firms are symmetric, in equilibrium the labor share is the firm-level labor share evaluated

at the aggregate labor supply and the capital endowment:

$$s_L = \frac{1}{1 + \frac{A_K K}{A_L L}^{\frac{\sigma-1}{\sigma}}} \frac{1}{1 + \frac{1}{\sigma}} \quad (\text{A.3})$$

We rewrite the equation for equilibrium employment as

$$\log(L) = \log(\bar{L}) + \log(A_L) - \frac{1}{1} \log(s_L) - \frac{1}{1} \log \left(1 + \frac{1}{\sigma} \right) \quad (\text{A.4})$$

This equation does not have a closed-form solution. However, using the implicit function theorem we can take derivative with respect to I :

$$\frac{d \log(L)}{dI} = \frac{d \log(A_L)}{dI} + \frac{1}{1} \frac{d \log(s_L)}{dI} \quad (\text{A.5})$$

The expression for this derivative is

$$\begin{aligned} \frac{d \log(s_L)}{dI} &= s_K \frac{1}{A_K} \frac{d \log A_K}{dI} - \frac{d \log A_L}{dI} - \frac{d \log L}{dI} \\ &= \frac{s_K}{A_K} \frac{(I)^{-1}}{A_K^{-1}} + (1 + \frac{1}{\sigma}) \frac{(I)^{-1}}{A_L^{-1}} - \frac{s_K}{dI} \frac{d \log(s_L)}{dI} \\ &= \frac{\frac{s_K}{1 + \frac{1}{\sigma}}}{A_K} \frac{(I)^{-1}}{A_K^{-1}} + \frac{1 + \frac{1}{\sigma}}{A_L} \frac{(I)^{-1}}{A_L^{-1}} \end{aligned}$$

which is always negative. □

Proposition A.2 (Productivity Effect). *The derivative of the log labor productivity with respect to I when the automation threshold is binding is always positive.*

Proof. Labor productivity is

$$\begin{aligned} \frac{Y}{L} &= A_L \left(1 + \frac{A_K K}{A_L L} \right)^{\frac{\sigma-1}{\sigma}} \frac{\sigma}{\sigma-1} \\ &= A_L s_L \left(1 + \frac{1}{\sigma} \right)^{\frac{\sigma}{1-\sigma}} \end{aligned}$$

The expression for the derivative with respect to the automation threshold, when it is binding, is

$$\begin{aligned}
\frac{d \log(Y L)}{dI} &= \frac{d \log(A_L)}{dI} + \frac{1}{1} \frac{d \log(s_L)}{dI} \\
&= \frac{1}{1} \frac{(I)^{-1}}{A_L^{-1}} + \frac{d \log(s_L)}{dI} \\
&= \frac{1}{1} \frac{(I)^{-1}}{A_L^{-1}} \frac{s_K}{1 + s_K} \frac{(I)^{-1}}{A_K^{-1}} + (1 + \frac{s_K}{1 + s_K}) \frac{(I)^{-1}}{A_L^{-1}} \\
&= \frac{1}{1} \frac{1}{1 + \frac{s_K}{1 + s_K}} (1 - \frac{s_K}{1 + s_K} + \frac{s_K}{1 + s_K} \frac{(I)^{-1}}{A_K^{-1}}) \frac{(I)^{-1}}{A_L^{-1}} + \frac{s_K}{1 + s_K} \frac{(I)^{-1}}{A_K^{-1}} \\
&= \frac{1}{1} \frac{1}{1 + \frac{s_K}{1 + s_K}} (1 - \frac{s_K}{1 + s_K}) \frac{(I)^{-1}}{A_L^{-1}} + \frac{s_K}{1 + s_K} \frac{(I)^{-1}}{A_K^{-1}} + \frac{s_K}{1 + s_K} \frac{(I)^{-1}}{A_L^{-1}}
\end{aligned}$$

which is always positive. □

Proposition A.3. *The derivative of the equilibrium log wage with respect to I when the automation threshold is binding is the sum of the productivity effect and the displacement effects, with the sign of the overall effect being ambiguous.*

Proof. The expression for the derivative is:

$$\begin{aligned}
\frac{d \log(W)}{dI} &= \frac{d \log(Y L)}{dI} + \frac{d \log(s_L)}{dI} \\
&= \frac{1}{1} \frac{\frac{s_K}{1 + s_K}}{1 + \frac{s_K}{1 + s_K}} \frac{1}{1} \frac{(I)^{-1}}{A_K^{-1}} + \frac{1}{1} \frac{1}{1} \frac{\frac{s_K}{1 + s_K}}{s_K} \frac{(I)^{-1}}{A_L^{-1}}
\end{aligned}$$

□

Proposition A.4. *The derivative of equilibrium log employment with respect to I when the automation threshold is binding is simply the elasticity of labor supply times the derivative of the log wage, and therefore it has the same sign as the latter.*

Proof. The expression for this derivative is

$$\begin{aligned}
\frac{d \log(L)}{dI} &= \frac{d \log(W)}{dI} \\
&= \frac{\frac{s_K}{1 + s_K}}{1 + \frac{s_K}{1 + s_K}} \frac{1}{1} \frac{(I)^{-1}}{A_K^{-1}} + \frac{1}{1} \frac{1}{1} \frac{\frac{s_K}{1 + s_K}}{s_K} \frac{(I)^{-1}}{A_L^{-1}}
\end{aligned}$$

□


Appendix B Data

B.1 Local labor market concentration

Table 3: Size brackets for number of employees by establishment as reported in County Business Patterns

Size bracket	Median
1-4	3
5-9	7
10-19	15
20-49	35
50-99	75
100-249	175
250-499	375
500-999	750
1000-1499	1250
1500-2499	2000
2500-4999	3750
5000-more	imputation

Figure 4: Local labor market concentration across US commuting zones in 1990



B.2 Local labor market exposure to robots

We follow Acemoglu and Restrepo (2020) and construct a measure of commuting zone exposure using the following data sources:

Industrial robots: We use data on the operational stock of industrial robots from the International Federation of Robotics (IFR) for the United States and six European countries (Denmark, Finland, France, Italy, Sweden, Germany) from 1993 to 2016.¹⁰ We classify the IFR data into 13 manufacturing industries, and 6 broad industries outside manufacturing.¹¹ To obtain the 19 IFR industries as in Acemoglu and Restrepo (2020), we perform the following adjustments to the original data: First, we keep the industry “all other manufacturing branches” and label it as “Miscellaneous manufacturing”. Second, “All other non-manufacturing branches” are considered as “Services”. Third, the residual category “Metal (unspecified)” is allocated proportionally to all industries in the “Metal industries” (Basic Metals, Metal Products, Electronics, Industrial Machinery) and 4.) the residual “Unspecified”, which is allocated proportionally over all 19 IFR industries. The IFR data comes with two drawbacks: first, it groups the US together with Canada as Northern America before 2011 and second, it doesn't provide a split-up by industries for the Northern America before 2004. Given that the US accounts for about 90 percent of the North American robot stock, we accept the first limitation. To deal with the second limitation, we apply an algorithm that attributes the total stock in each year before 2004 according to an industry's share in the total stock in 2004, the first year with disaggregated information on the industry level. We apply this solution also to Denmark, which similarly lacks data by industry before 1996.

Industry employment and output: Furthermore, we use data on employment and output from the 2007 and 2019 EU KLEMS releases (Stehrer et al., 2019; Timmer et al., 2007).¹² As in Acemoglu and Restrepo (2020), we translate the numbers of persons employed in each European country-industry in 1990 into “US equivalent workers” by dividing the total number of hours worked in a European industry by the hours per worker in the corresponding US industry. This is to account for the fact that European workers work on average less hours and to make employment numbers comparable. To adjust for the growth in robot stock due to output growth, we compute an output growth rate and use the output deflators provided by EU KLEMS to correct for inflation.

¹⁰These selected European countries exhibit levels and an evolution of the number of robots per 1000 workers that mirror the US over the sample period from 1993 to 2015 and will be used to construct an instrumental variable.

¹¹Manufacturing industries include Food and Beverages, Textiles, Wood and Furniture, Paper and Printing, Plastics and Chemicals, Minerals, Basic Metals, Metal Products, Electronics, Industrial Machinery, Automotive, Shipbuilding and Aerospace, Miscellaneous Manufacturing; Non-Manufacturing industries include Agriculture, Mining, Utilities, Construction, Education and Research, Services.

¹²We use both releases as the 2019 release in NACE 2 only covers the period 2000 to 2018, while the 2007 NACE1 release only provides data from 1970 to 2005. To obtain industry employment and output data for multiple countries from 1990 to 2016 we do therefore need to combine both the 2007 NACE 1 and the 2019 NACE 2 releases. The mapping of NACE 1/2 to IFR industries is available upon request.

Commuting zone employment: Finally, we compute industry employment shares in each commuting zone in 1970 and 1990 as well as changes in labor market outcomes using micro-data from the US Decennial Census for the years 1970, 1990 and 2000 as well as from the American Community Survey in 2006, 2007 and 2008 and 2014, 2015 and 2016 provided by the *Integrated Public Use Microdata Series* (IPUMS). We use the crosswalks by **autor2013growth** to map geographies provided in the IPUMS data to 722 continental commuting zones.

To compute the industry employment in each commuting zone in a given year, we sum over working individuals aged 16 or older by industry using person weights from IPUMS multiplied with probability weights from the geographical crosswalks. We calculate the total commuting zone employment simply as the sum of employment across all industries.¹³

B.3 Local labor market outcomes

Employment, unemployment and non-participation: Following Acemoglu and Restrepo (2020), we calculate averages for demographic groups within commuting zones using micro-data from the US Decennial Census for 1970, 1990, and 2000, the American Community Survey for 2006, 2007, and 2008, and for 2014, 2015, and 2016 provided by *Integrated Public Use Microdata Series* (IPUMS). We focus on individuals aged 16 to 65 employed in the private sector, specifically in manufacturing or blue-collar occupations. Unemployment rates are computed relative to the commuting zone's total labor force, and non-participation rates are relative to the total working-age population.

Average wages: To calculate average wage across demographic groups within commuting zones, we utilize micro-data from the US Decennial Census for 1970, 1990, and 2000, along with data from the American Community Survey for 2006, 2007, and 2008. Our analysis focuses on individuals aged 16 to 65 employed.

To handle top-coded wage incomes, we cap them at 1.5 times the respective annual top-coded wage for each year and deflate wages using the 1999 consumer price index. Average weekly income is computed by dividing total annual wage income by weeks worked, while hourly wages are derived by dividing the average weekly wage by the usual number of hours worked per week, as indicated in the micro-data. We winsorize hourly wages at \$2 USD, in line with Acemoglu and Restrepo (*ibid.*).

¹³The mapping of 1990 Census Bureau industry classes to corresponding IFR industries is also available upon request.

Individuals are categorized into one of 250 demographic cells within each commuting zone, defined by age groups (16-25, 26-35, 36-45, 46-55, 56-65), educational attainment (less than high school, high school degree, some college, college/professional degree, and masters/doctoral degree), sex (male/female), and race (Hispanic, Black, White, Asian, Other). We calculate average yearly, weekly, and hourly wages for each demographic group within each commuting zone cell.

Table 4: Descriptive Statistics of Commuting Zone Data

	Means by quartiles of exposure to robots			
	All (1)	Q1 (2)	Q4 (3)	Q4-Q1 (4)
Changes in outcomes, 1990-2015:				
Log private sector employment	21.72	28.23	14.23	-14.00***
Log manufacturing employment	-13.87	4.25	-28.48	-32.73***
Log blue-collar employment	1.80	16.42	-11.37	-27.79***
Employment to population ratio	2.00	4.23	0.22	-4.01***
Unemployment rate	-0.45	-0.85	-0.31	0.55***
Non-participation rate	1.84	0.40	3.14	2.74***
Changes in outcomes, 1990-2007:				
Log average yearly wage	10.69	15.66	6.31	-9.35***
Log average weekly wage	5.00	9.94	1.00	-8.93***
Log average hourly wage	2.80	8.41	-0.98	-9.39***
Share of population, 1990:				
Female	0.51	0.51	0.51	0.01***
Less than college	0.71	0.69	0.74	0.05***
Some college or more	0.25	0.28	0.23	-0.05***
White	0.87	0.87	0.89	0.03**
Black	0.08	0.03	0.09	0.05***
Asian	0.00	0.00	0.00	-0.00
Hispanic	0.06	0.10	0.02	-0.09***
Above 65 years old	0.13	0.14	0.13	-0.00
Share of employment, 1990:				
Manufacturing	0.17	0.08	0.24	0.16***
Light manufacturing (within manufacturing)	0.22	0.21	0.22	0.01
Female employment (within manufacturing)	0.33	0.32	0.32	0.00
Routine employment	0.36	0.33	0.38	0.05***
Labor market concentration, 1990:				
Employment \overline{HHI}	0.33	0.38	0.29	-0.10***
Observations	722	181	180	361

Note: Columns (1) to (3) display unweighted means of changes in outcomes *multiplied by 100* as well as unweighted means of commuting zone characteristics in 1990. For each commuting zone we compute the average exposure to robots the periods 1990 to 2000, 2000 to 2007 and 2007 to 2015. Columns (2) and (3) display unweighted means within commuting zones in the first and last quartile of the exposure distribution, respectively. Column (4) displays the difference in the mean commuting zone characteristics between means fourth and the first quartile of robot exposure and reports statistical significance of the underlying test. Coefficients with ***, **, and * are significant at the 1%, 5% and 10% confidence level, respectively.