Should AI Be Taxed When Firms Have Labor Market Power?.

Felipe Balmaceda *

July 15, 2026

*Email: fbalmace@gmail.com.

Abstract

In a labor market with market power, firms automate based on the rent they appropriate from a job, not its output, since wages transfer part of human-job output to workers while automation output is fully appropriated. This generates excess automation relative to productive and constrained-efficient benchmarks, and appears to call for a corrective tax on automation capital. We show this prescription can reverse once firms can also invest in generative AI to augment human-skill jobs: because automation and generative AI draw on the same capital, a tax discourages both margins simultaneously. Calibrated to U.S. automation-adoption and wage-markdown data, the generative-AI margin dominates, so the welfare-maximizing policy is a *subsidy*, not a tax, reversing the standard robot-tax prescription. Competition policy produces a non-monotone welfare response through the same channel. We provide descriptive evidence consistent with the model's prediction that labor market power reduces substitutive automation.

JEL: J3, D2, J30

Key Words: Labor market power, automation, generative AI, wages, displaced workers, productive and allocative efficiency, human capital, taxes.

1 Introduction

How does labor market power distort firms' AI adoption? Should governments tax the technologies that replace workers? The idea has an old pedigree: Marx (1894) argued that firms invest in labor-saving machines to reduce wage costs, thereby displacing workers and weakening labor's bargaining power. Keynes, echoing Marx, warned of "technological unemployment" and Leontief made a similar prediction: "Labor will become less and less important... . Machines will replace more and more workers. I do not see that new industries can employ everybody who wants a job".

Why Labor Market Power and AI? Because market power has risen globally, allowing firms to suppress wages below marginal productivity even in online markets that, a priori, one would expect to be highly competitive (Manning, 2021). Simultaneously, technological change has accelerated, with automation and artificial intelligence (AI) transforming the nature of work. Some technologies fully replace human labor, whereas others—particularly generative AI—augment it. Empirical evidence shows that automation disproportionately displaces routine and mid-skill workers and contributes to wage inequality (Acemoglu and Restrepo, 2021, Webb, 2020), while AI technologies tend to complement higher-skilled workers and raise productivity (Autor, Chin, Salomons, and Seegmiller, 2022, De Souza and Li, 2023, Brynjolfsson, Li, and Raymond, 2025).

Variants of a tax proposal have been endorsed by the European Parliament, by prominent entrepreneurs, and academic experts (`https://www.wemustactnow.ai`),—and a growing theoretical literature has begun to formalize when a tax on automation capital is welfare-improving.[1] Guerreiro, Rebelo, and Teles (2022a) show, in a dynamic general-equilibrium model, that it is optimal to tax robots while the current generations of routine workers, who cannot move into non-routine occupations, are still in the labor force. Once these workers retire, the optimal robot tax is zero. This is because taxing robots mitigates automation by reducing the income loss and borrowing constraints of displaced workers, while increasing the labor supply of non-routine workers who benefit from automation. Their result, and the broader case for robot taxation, rests on a single margin: taxing automation capital makes the automation technology relatively more expensive, so firms substitute toward labor.

[1] For instance, Bill Gates has advocated for a tax on robots. See `https://qz.com/911968/bill-gates-the-robot-that-takes-your-job-should-pay-taxes/`.

This paper shows that the case for a robot tax can invert once firms have a *second* margin of technological adjustment: the choice, within jobs that remain assigned to human workers, of how intensively to augment those workers with generative AI. We develop a model in which firms with labor-market power choose, job by job, whether to automate, to hire a worker with no technological augmentation, or to hire a worker and invest in generative AI capital that raises that worker's output. A tax on technological capital raises the cost of *both* margins simultaneously: it discourages automation, as the standard robot-tax logic requires, but it also discourages the generative-AI investment that makes the surviving human-skill jobs more productive and, through the incomplete wage pass-through, more attractive for firms to keep staffed with workers in the first place. Whether the tax raises or lowers welfare depends on which margin is more sensitive to the tax rate—an object we characterize analytically and then discipline quantitatively.

At empirically plausible parameter values, we find that the gen-AI margin dominates: a tax on technological capital reduces automation only modestly while substantially depressing wages and productivity in the jobs it was meant to protect, and the welfare-maximizing policy is not a tax but a *subsidy* to technological capital. This is not a knife-edge or fine-tuned result—in our counterfactual analysis, the optimal instrument moves continuously from "no intervention" at low competition intensity to a sizeable subsidy as labor markets become more competitive, exactly the range of concentration our calibration targets. The reversal relative to the existing robot-tax literature is driven by a channel that is entirely absent from models without a labor-augmenting technology and where workers are paid their marginal product of labor: taxing capital is not merely a tax on the technology that substitutes for labor, it is also a tax on the technology that complements it, and a planner who ignores the latter will misjudge the sign of the optimal policy.

The mechanism generating this result is a specific and, we think, underappreciated distortion in how firms with labor-market power compare automation to human-staffed production. Because firms must share the output of a human-skill job with the worker through the equilibrium wage, whereas they fully appropriate the output of an automated job, firms compare the *rent* from staffing a vacancy– not its *output*—against the full payoff from automation. This rent-appropriation wedge means that automation can be privately profitable even when human-skill jobs are substantially more productive: in our calibra-

tion, disciplined by U.S. data on automation adoption and wage markdowns, human-skill jobs are roughly six times more productive than automated ones, yet the calibrated equilibrium features a 20% automation rate, and both productive and constrained-efficient allocation would require none. This wedge is what makes a corrective tax on automation appear attractive in the first place; our point is that the same wedge, once a complementary technology is present, also constrains what a capital tax can achieve, because the tax cannot target automation without also taxing the margin that makes non-automation attractive.

This has a direct and, we think, policy-relevant implication for the design of automation taxation. Proposals to tax "robots" or "automating technology" are typically written at the level of capital equipment or investment, not at the level of the specific task the capital replaces. Our results imply that such instruments are poorly targeted whenever the taxed capital base overlaps with the capital used to augment, rather than replace, human labor—as is increasingly the case for computing and AI capital, which can be deployed on either margin depending on how a firm designs a given job. A tax written broadly enough to capture automating robots will, in our model, mechanically also tax the generative-AI investments that expand employment and productivity in the jobs the tax was meant to protect. The policy implication is not that automation is efficiently provided—our model says the opposite, that automation is systematically overprovided relative to both productive and constrained-efficient benchmarks—but that the tax instrument required to correct this overprovision needs to be considerably better targeted than a uniform capital tax before it can be expected to raise welfare.

We also examine competition policy as an alternative instrument and find that it interacts with the tax margin in a way that has, to our knowledge, not been noted in the robot-tax literature: more competitive labor markets erode the human-job rents that make firms prefer staffing a vacancy over automating it, so competition policy that reduces labor market concentration can *increase* automation and worker displacement even as it raises wages for workers who keep their jobs. The welfare effect of competition policy is therefore non-monotone in our calibration, and antitrust enforcement that moves a market from monopsony toward moderate competition can lower welfare before it raises it. This means the optimal capital-tax or subsidy rate is not a single number but depends on the level of labor-market concentration a regulator is simultaneously targeting; we char-

acterize this joint optimum across competition levels and show that it moves from no intervention to a large subsidy as concentration falls, with no region in which a positive tax is optimal at the calibrated parameters. This paper builds on several strands of literature spanning labor economics, industrial organization, and the economics of technological change.

An extensive empirical literature studies the impact of automation on employment and wage structures. Acemoglu and Restrepo (2021) show that between 50% and 70% of recent changes in the US wage distribution are due to the declining relative wages of workers performing routine tasks. Webb (2020) finds that previous automation technologies have led to declines in employment and wages, particularly for low- and mid-skill occupations, while AI appears to target high-skill tasks. Similarly, Autor et al. (2022) show that automation reduces employment in exposed occupations, while augmenting technologies expand employment and wages in complementary jobs. Bessen, Goos, Salomons, and van den Berge (2020) find that while automating firms grow faster, they experience short-term job losses at the point of automation.

Dixon, Hong, and Wu (2021), using Canadian firms, find that the robot-adopting firms experience a subsequent increase in employment, but decreases in the total number of managers, and that the employment increase is predominantly from low- and high-skill workers and falls for middle-skill workers. Koch, Manuylov, and Smolka (2021), using Spanish firms, report a similar finding. These studies also find that firms that adopt robots experience higher performance (measured by firm-level total factor productivity or revenue). Interestingly, the studies find that non-adopting firms in the same industry as the robot-adopting firms experience employment declines (Acemoglu and Restrepo, 2020). Eggleston, Lee, and Iizuka (2021) study the effects of robots on workers in Japanese nursing homes and find that robots complement human labor and reduce labor turnover.

Recent studies examine the differential effects of AI and software adoption. Aghion, Antonin, Bunel, and Jaravel (2022) find that automation boosts employment, sales, and profits but does not significantly affect wages or wage inequality within firms. De Souza and Li (2023) find that robots have significantly decreased employment and wages of low-skill workers in operational occupations. However, tools, machines, and software that use AI to augment labor have led to an equally large increase in the reinstatement of these workers, raising their employment and wages.

It is increasingly common to examine imperfect competition in labor markets rather than in product markets. Manning (2021) reviews extensive evidence showing that labor market power is pervasive, affecting wage setting and employment decisions. Acemoglu and Restrepo (2024) show that automation is concentrated in high-rent jobs, dissipating worker rents and exacerbating wage losses. However, in their setting, rents arise from worker-side frictions, whereas in our model, firms hold the bargaining advantage.

Thuemmel (2023) shows that it is optimal to distort robot adoption through taxes or subsidies that exploit general equilibrium effects to compress wages, reducing income tax distortions and improving welfare. In the calibrated model, a robot subsidy is optimal when robots are expensive, but taxation becomes optimal as robots become cheaper. However, most welfare gains can be achieved by adjusting income taxes, with marginal gains from differentially taxing robots being negligible. Costinot and Werning (2023) investigate optimal taxes on new technology firms using sufficient statistics that can be estimated empirically with minimal structural assumptions. They show that small, positive robot taxes are optimal, though welfare gains are small when income taxes are optimally set. Since robots and new technologies are challenging to tax due to classification problems, Loebbing (2022) analyzes how progressive tax reforms can affect technological change when robots cannot be taxed separately. More progressive taxes give low-skilled workers stronger work incentives, inducing firms to adopt technologies suitable for low-skilled workers, making the optimal tax system more progressive.

Kina, Kindsgrab, and Krueger (2023) study optimal taxation of capital and labor income with capital-skill complementarities, showing it is optimal to rely more on capital income taxes when these complementarities are considered. The optimal capital income tax rate is 60%, significantly higher than 48% without capital-skill complementarity. Schulz, Tsyvinski, and Werquin (2022) examine how to reform income taxes in response to technological changes, deriving a formula for compensating tax reform that accounts for skill complementarities in production. They quantify income losses from robots and show how compensatory reforms can effectively mitigate automation's negative effects.

Our paper contributes to the literature on innovation, automation, and inequality by modeling technology adoption under imperfect labor markets. It provides a unified explanation for excessive displacement, wage suppression, and skill misallocation, offering a novel theoretical foundation for ongoing debates over taxing automation and regulat-

ing AI in the labor market. By dealing with both allocative and productive inefficiencies, we identify novel margins through which labor market power distorts technological transitions that has not been identify in the literature.

The rest of the paper is structured as follows. In the following section, motivational evidence concerning labor market power and AI adoption is provided. Section 2 develops the model of automation, human-skill production, and generative-AI investment under labor-market power. In the next section, Section 3, we derive the subgame-perfect equilibrium. Then, in Section 4, we discuss the efficiency of the equilibrium. Section 5 compares the equilibrium allocation to productive and constrained-welfare benchmarks and isolates the rent-appropriation wedge. Section 6 disciplines the model with U.S. data, reports the calibrated tax and competition-policy counterfactuals, and provides descriptive evidence, external to the calibration, consistent with the model's qualitative predictions. In Section 7 concludes.

2 The Model

2.1 Set-Up

Let us consider the following labor-market game. In the first period, firms decide whether or not to automate the job. Afterward, wages are chosen simultaneously by the firms that choose not to automate the job. In the third period, individuals learn the firm-specific non-pecuniary preference shocks and the wage for each firm and choose to supply their labor to the firm that offers the higher utility, provided that this exceeds the outside option payoff. Afterward, firms learn applicants' skill level t and decide whether to assign the worker to a job that combines generative AI with human skills, where the intensity of generative AI is endogenously determined.

To keep the analysis simple, firms and individuals are risk-neutral and do not discount the future. Firms separate workers into different human capital or credential groups or classes, denoted by s, where s could be, for instance, college degree workers, high-school graduates, etc. Firms believe that the skill, denoted by t, cumulative distribution function of workers belonging to class s is $F(t|s)$. This has full and bounded support $\mathcal{T} \subset \Re_+$, and density $f(t|s)$. The class to which a worker belongs and $F(t|s)$ are common knowledge.

However, the skill level is the workers' private information. A worker from skill class $s' > s$ has distribution $F(t|s')$ that dominates $F(t|s)$ in the sense of first-order stochastic dominance. Thus, $F_s(t|s) < 0$. A s-class worker who cannot find a job or chooses the outside option receives a payoff of $b(s)$, which is non-decreasing in s. Thus, workers in higher classes have a higher opportunity cost of labor market participation.

There are $n(s)$ firms in the labor market s which are horizontally differentiated from individuals' point of view, indexed by $j \in \{1, \ldots, n\}$. Firms produce tradeable goods that are perfect substitutes, and thus trade in a perfectly competitive market at a price p, normalized to 1.[2]

We model such occupational differentiation by adopting a random-utility framework in the spirit of Perloff and Salop (1985).[3] Let $\epsilon_l = (\epsilon_l^j, \ldots, \epsilon_n^j)$ be the match-specific utility shock of individual l in each of the $j \in \{1, \ldots, n\}$ possible firm/jobs. Thus, the utility of individual l in job j is given by: $w^j + \epsilon_l^j$. We assume that ϵ_l is i.i.d. across individuals, reflecting idiosyncratic tastes for different jobs, and that, for a given worker, it is also i.i.d. across jobs. These non-wage job characteristics may include hours of work, the distance of the firm from the worker's home, and the social environment in the workplace, among others. In the forthcoming analysis, we will suppress the index l. ϵ^j is distributed $G(\cdot)$ with compact and full support $[\underline{\epsilon}, \bar{\epsilon}] \subset \Re$, zero mean and twice differentiable density $g(\cdot)$.

We assume the existence of a local labor market for each worker class s. Within that market, competing firms simultaneously set wages without discriminating among individuals of the same type or class. Because firms know workers' class, a worker from class s cannot apply to a job in class s'. In that sense, the worker class determines the job market in which they can participate. For instance, a lawyer with skill level t will not be hired as a smartphone technician, because the skill level t for a lawyer has a very different meaning from that for the technician. From now on, we will refer to workers belonging to the class s as the labor market s.

The next assumption is crucial for characterizing the labor-market equilibrium.

Assumption 1. *$g(\epsilon)$ is log-concave.*

This ensures the existence of an equilibrium, and the markdowns go to zero as the

[2] To save on notational burden, we will omit the argument (s) whenever there is no risk of confusion.

[3] Balmaceda (2025) also studies a Perloff and Salop (1985) in the context of occupational choice, and Azar, Berry, and Marinescu (2019) estimate labor market power using a logit model, which is a particular case of Perloff and Salop's (1985) model.

number of firms goes to infinity.[4] Firms have access to constant returns-to-scale technology, i.e., the total output of a firm equals the sum of the outputs of each job.[5]

Jobs are of three different types: automated jobs (a) that fully substitute for human skills jobs; human-skills-only jobs (h) whose only input is human skills; and generative AI-human skills jobs (g) that use both human skills and generative AI. There is a finite number of jobs that can be automated.[6]

The expected output in an automated job in firm j in the labor market s when automation capital is a is

$$y^j(a;s) = \int_0^{\bar{y}} y dA(y|a,s).$$

where $A(y|a,s)$ is continuous CDF with support $[0,\bar{y}]$ and satisfying $A(y|a',s) \leq A(y|a,s)$ for all $a' > a$.[7] Thus, automation capital improves the automation output in the sense of first-order stochastic dominance.

The expected output of a human skills job in the labor market s when the investment in generative AI capital is h in firm j is given by

$$y^j(t,h;s) = \int_0^{\bar{y}} y dH(y|t,h,s).$$

where $H(y|t,h,s)$ is continuous CDF with support $[0,\bar{y}]$ and satisfying $H(y|t',h,s) \leq H(y|t,h,s)$ for all $t' > t$ and all h, $H(y|t,h',s) \leq H(y|t,h,s)$ for all $h' > h$ and all t. Thus, generative AI enhances productivity across all skill levels.

The technology satisfies the following properties.

Assumption 2. *For all $j \in \mathcal{J}$,*

i) $y^j(a;s)$ is strictly concave in a, $\lim_{h\to 0} y^j_a(a;s) > r$, and $\lim_{h\to\infty} y^j_a(a,s) < r$.

[4]See, Gabaix, Laibson, Li, Li, Resnick, and de Vries (2016) for details about markup convergence in random utility models.

[5]This assumption is not as restrictive as it appears at first glance. If the technology is of constant returns to scale and inputs can be freely adjusted, the marginal contribution of a worker will be independent of the other inputs. The reason is that a profit-maximizing firm will maintain a constant input ratio.

[6]This assumption is meant to avoid the solution of firms creating as many automated jobs as they wish whenever the rent from automation is positive. To address this, we could have assumed decreasing returns to automation technology, but that would complicate the algebra without gaining economic intuition.

[7]A potential interpretation is that there is a continuum of tasks in $[0,\bar{y}]$ that must be performed to produce the output. The amount of output produced depends on the corresponding technological capital in the case of automation, and on human skills in the case of a human job. The restriction would be that the tasks cannot be split between human skills and automation technology.

ii) For all $t \in \mathcal{T}$, $y^j(t,h;s)$ *is strictly concave in* h, $y^j(t,0;s) > 0$, *and* $\lim_{h\to\infty} y^j_h(t,h;s) < r$.

iii) For any $(t',h') > (t,h)$, *either* $y(t',h',s) + y(t,h,s) \lesseqgtr y(t,h',s) + y(t',h,s)$.

These are standard conditions to guarantee the uniqueness of the capital investments. Part ii) also establishes that human skills are productive even when no investment in generative AI is made. Part iii) says that generative AI capital and human skills can be either complements or substitutes.

Rosen (1987) was the first to highlight the importance of non-pecuniary job characteristics in the compensating wage differentials literature. Lamadon, Mogstad, and Setzler (2022) show that worker preferences over non-pecuniary job characteristics lead to imperfect competition in the US labor market. Maestas, Mullen, Powell, von Wachter, and Wenger (2018) find that high-wage workers and college-educated workers have uniformly better job characteristics, and Mas and Pallais (2017) argue that there is evidence that workers in the US are willing to give up part of their income compensation to avoid undesirable working conditions. Sullivan and To (2014) show that there are substantial gains to workers from job search based on non-pecuniary factors, workers sort into jobs with better non-pecuniary job characteristics, and are willing to pay for them. Sorkin (2018) shows a high prevalence of US workers who move to lower-paying firms in a way that cannot be accounted for by layoffs or differences in recruiting intensity to benefit from non-pecuniary job characteristics. He estimates that compensating differentials account for over half of the firm component of the earnings variance. These results provide a foundation for labor-market power driven by the horizontal differentiation of jobs.

In addition, it is highly plausible that individuals with identical productivity may choose different jobs due to their differing tastes. Accounting for job preferences is particularly important for understanding differences in job choices across groups.

3 The Equilibrium

3.1 Automation and Generative AI Capital Investments

Let us consider the labor market s and suppose firm j opens a vacancy. Because when a worker applies to a job in firm j, both the firm and the worker already know the worker's

skill t, firm j chooses generative AI capital to solve the following problem

$$\max_{h \in \Re_+} \{y^j(t, h^j; s) - w^j - rh^j\}. \tag{1}$$

The first-order condition is

$$\int_0^{\bar{y}} y dH_h(y|t, h^j, s) - r \leq 0.$$

Because of Assumption 2 part ii), if $\lim_{h \to 0} y^j_{h^j}(t, h^j; s) > r$, a unique interior solution exists. Otherwise, the optimal investment in generative AI is zero. Let's denote the optimal solution by $h^j(t, r, s)$. Thus, the output when the job is not automated is given by $y(t, r; s) \equiv y^j(t, h^j(t, r, s); s) - rh^j(t, r, s)$.

When the optimal solution is strictly positive, it is easy to check that if generative AI capital and skills are complements, $h^j(t, r, s)$ rises with t, whereas if they are substitutes, $h^j(t, r, s)$ falls with t. Furthermore, it readily follows from the implicit function theorem and first-order stochastic dominance with respect to s that $y(t, r; s)$ rises with s and falls with r.

Because the wage is already set, it does not affect the investment decision in generative AI. When $h^j(t, r, s) > 0$, the worker is allocated to a job enhanced with generative AI. Otherwise, the job is produced solely by human labor. If skills and generative AI are substitutes, if for any t, $h^j(t, r, s) = 0$, then $h^j(t', r, s) = 0$ for all $t' > t$, whereas if they are complements, then $h^j(t', r, s) = 0$ for all $t' < t$. Thus, whether the investment in generative AI is large or smaller for high-skill workers depends on whther skills and generative AI are complements or substitutes.

When the job is automated, firm j chooses automation capital to solve the following problem

$$\max_{a \in \Re_+} \{y^j(a^j; s) - ra^j\}. \tag{2}$$

The first-order condition is

$$\int_0^{\bar{y}} y dA_a(y|a^j; s) - r = 0.$$

Because of Assumption 2 part ii), a unique positive solution, denoted by $a^j(r,s)$, exists. Thus, the output when the job is automated is given by $y(r,s) \equiv y^j(a^j(r,s);s) - ra^j(r,s)$. The investment level is higher in the labor market, $s > s'$, whenever the marginal productivity of automation rises with the labor market qualifications. In other words, when automation is a stronger substitute for human skills.

3.2 Equilibrium Wages

Let us consider the labor market s. Let $d^j \in \{0,1\}$ be firm j's automation decision, where $d^j = 1$ means the job is assigned to human skills and $d^j = 0$ means the job is automated.[8] Then, for any profile d, let $\mathcal{J}(d) \subseteq \mathcal{J}$ be the set of firms that opens a vacant and $\mathcal{J}(d^{-j}) \equiv \{k \in \mathcal{J} : k \neq j \text{ and } d^k = 1\}$ be the set of firm j's competitors that open a vacant.

Because workers in the labor market s observe $(\epsilon, w(s))$ before choosing a firm to supply their labor, they will choose the firm that provides the highest expected utility among all those that have a vacant available $j \in \mathcal{J}(d)$ provided that this yields a higher utility than the outside option b. Thus, a worker chooses firm $j \in \mathcal{J}(d)$ whenever $w^j + \epsilon^j \geq \max\{b, w^{j'} + \epsilon^{j'}\}$. Hence, the probability that a worker chooses firm $j \in \mathcal{J}(d)$ instead of any other firm is given by

$$P^j(w) = P[w^j + \epsilon^j \geq \max_{k \in \mathcal{J}(d^{-j})} \{w^k + \epsilon^k, 0\}] = \int_{\max\{\underline{\epsilon}, b - w^j\}}^{\bar{\epsilon}} \prod_{k \in \mathcal{J}(d^{-j})} G^k\Big(w^j + \epsilon^j - w^k\Big) dG^j(\epsilon^j),$$

where the equality follows from the independence assumption about the G's distributions.

Proposition 1. *$P^j(w)$ is strictly positive, strictly increasing in w^j, strictly decreasing in $w^{j'}$ for all $j' \neq j$, log-concave in w^j, and log-supermodular in w.*

The log-concavity follows from the fact that G is log-concave and that the product of log-concave functions is log-concave. The log-concavity of the firm-specific labor supply implies that the price elasticity of supply increases with the wage.

Because total-positivity of order 2 is preserved under marginalization, the supply is

[8] When there is no risk of confusion, we will omit the arguments to keep the notation simpler, and we will omit the dependence of distributions and wages on the skill class s.

log-supermodular in w.[9] This means that the price elasticity of supply decreases as competitors' wages increase. The latter will imply an increasing best-response correspondence when jobs are gross substitutes.

Let's define the marginal product of labor by $\mathbf{E}_t[y^j(t,r,s)] = \mathbf{E}_t[y^j(t,h^j(t,r,s),s)]$. For any given wage profile w, firm j's profits when a vacancy is open are then given by:

$$\Pi^j(w) \equiv \mathbf{E}_t\left[\left(y^j(t,r,s) - w^j\right)P^j(w)\right]. \tag{3}$$

Thus, firm j chooses w^j, taken $w^{-j} \equiv (\ldots, w^{j-1}, w^{j+1}, \ldots)$ as given, to solve the following problem

$$\max_{w^j \in \Re_+} \Pi^j(w^j, w^{-j}).$$

In what follows, we will focus on parametric restrictions such that the case in which $\underline{\epsilon} \leq b - w^j$ holds for all j and, therefore, the outside-option payoff is chosen with positive probability for some workers in each labor market.[10] From here onward, let the subindex denote the derivative for the corresponding wage. Because G^k's are identically distributed, the first-order condition is given by

$$(\mathbf{E}_t[y^j(t,r,s)] - w^j)P_j^j(w) - P^j(w) \leq 0, \tag{4}$$

where,

$$\begin{aligned} P_j^j(w) = \int_{b-w^j}^{\bar{\epsilon}} \sum_{h \in \mathcal{J}(d^{-j})} \nu_g\left(w^j + \epsilon^j - w^h\right) \prod_{k \in \mathcal{J}(d^{-j})} G\left(w^j + \epsilon^j - w^k\right) dG(\epsilon^j) + \\ g(b - w^j) \prod_{k \in \mathcal{J}(d^{-j})} G(b - w^k), \end{aligned} \tag{5}$$

where $\nu_g(\cdot) \equiv g(\cdot)/G(\cdot)$ is the the reversed hazard rate of distribution G and the subindex j denotes the derivative with respect to wage w^j.

Lemma 1. *Firm j's best response $B^j(w^{-j}) \in (0, \mathbf{E}_t[y^j(t,r,s)])$ exists and is unique.*

Profits are log-supermodular in w because the markdown depends only on w^j and

[9] A function $f : X \times Y \to \Re$ with $X, Y \subseteq \Re$ is *totally positive of order 2* (TP_2) if $\forall\, x_1 < x_2 \in X$ and $y_1 < y_2 \in Y$, $f(x_1,y_1)f(x_2,y_2) - f(x_1,y_2)f(x_2,y_1) \geq 0$.

[10] This assumption does not change the results. If we allow for $\underline{\epsilon} > b - w^j$ in some labor market, in the symmetric equilibrium, the markdown will be a constant depending only on the number of firms.

$P^j(w)$ is log-supermodular in w. The following result readily follows from this and Theorem 6 in Milgrom and Roberts (1990). It also follows from Theorem 5 in Milgrom and Roberts (1990) that each firm has only one serially undominated strategy. Hence, the original game is dominance solvable and the equilibrium is globally stable under any adaptive learning rule satisfying assumption A6 in Milgrom and Roberts (1990).

Proposition 2. *For each labor market s, the equilibrium set has the component wise largest and smallest elements, given by $w_H(t,r,s)$ and $w_L(t,r,s)$ respectively, with*

$$w_l^j(t,r,s) = \mathbf{E}_t[y^j(t,r,s)]\frac{\xi^j(w_l^j(t,r,s))}{1+\xi^j(w_l^j(t,r,s))},$$

for all $j \in \mathcal{J}(d)$, where $\xi_l^j(w_l^j(t,r,s))$ is the elasticity of the labor supply for $l \in \{H,L\}$.

Hence, a type-s worker is paid a lower wage than his expected marginal product of labor. The markdown as a percentage of the wage is the inverse of the labor-supply elasticity. The higher the elasticity, i.e., the more intense the competition, the higher the wage.

From here onward, we will focus on the symmetric equilibrium for each type or local labor market, which requires assuming that $\mathbf{E}_t[y^j(t,r,s)] = \mathbf{E}_t[y(t,r,s)],\ \forall\ j \in \mathcal{J}(s)$. Let the cardinality of $\mathcal{J}(d)$ in the labor market s be n. Then, it readily follows from the first-order condition in equation (4) and integration-by-parts that the equilibrium wage $w(r,s)$ in the labor market s is determined by a fixed point of the following equation

$$\mathbf{E}_t[y(t,r,s)] - w = m(b-w) \equiv \frac{1}{n}\frac{1-G(b-w)^n}{\underbrace{G(b-w)^{n-1}g(b-w)}_{\text{exclusion effect}} + \underbrace{\int_{b-w}^{\bar{\epsilon}} g(\epsilon)dG(\epsilon)^{n-1}}_{\text{competition effect}}}. \quad (6)$$

The numerator in equation (6) is the equilibrium labor supply since the workers choose the outside option with probability $G(b-w)^n$ (i.e., when each firm j has a valuation less than $b-w$). The denominator is the slope of the labor supply. This has two terms: (i) the market exclusion effect (equivalent to the exclusion effect in the goods market). When the valuations for all other firms are below $b-w$, which occurs with probability $G(b-w)^{n-1}$, firm j acts as a monopsony. Lowering its wage w by ϵ will exclude $\epsilon g(b-w)$ individuals from paid employment; and (ii) the competition effect (up to the adjustment that the

marginal individual's valuation for paid employment is given by $b - w$) considering that a wage increase lowers the probability to be hired, which entails loosing not only the pecuniary benefit of being employed (w) but also the non-pecuniary benefit ϵ. This term represents the density of a firm's marginal workers—those who are indifferent between the corresponding firm and the best outside option for them—times the loss from a lower probability of being hired.

If both sides of equation (6) are divided by $w(r,s)$, the left-hand side is the Learner's index, denoted by $L(s) \equiv (e_t[y] - w(r,s))/w(r,s)$, and the right-hand side is the inverse of the labor-supply elasticity, denoted by $\xi(w)$. Hence, in equilibrium, the Lerner's index is the inverse of the supply elasticity. The Lerner's index ranges from 0 to ∞. A perfectly competitive firm pays $w(r,s) = \mathbf{E}_t[y(t,r,s)]$, and therefore $L(y) = 0$ –such a firm has no market power. An oligopsonist firm pays $w(r,s) < \mathbf{E}_t[y(t,r,s)]$, so its index is $L(s) > 0$, but the extent of its markdown depends on the elasticity of labor supply, which in turn depends on the strategic interaction with competing firms as well as the outside option.

Let $w^m(t,r,s)$ be the wage when there is a monopsony ($n = 1$). In this case, the elasticity is equal to the hazard rate evaluated at $b - w^m(t,r,s)$ and this increases with $b - w^m$ due to the log-concavity of f. The following is proven in the appendix, where all proofs are placed.

Proposition 3. *For each s-type, there exists a unique symmetric equilibrium wage given by* $w(r,s) \in [w^m(t,r,s), \mathbf{E}_t[y(t,r,s)]$.

Uniqueness follows from the fact that f is log-concave which makes $m(b - w(r,s))$ increasing in $w(r,s)$, while the LHS in equation (6) is decreasing in $w(r,s)$. Log-concavity implies that the CDF of the second-order highest statistic increases in the sense of first-order stochastic dominance with $b - w(r,s)$. This explains why the RHS in equation (6) increases with $w(r,s)$ and is bounded. The LHS in equation (6) falls with w. Then by the Intermediate Value Theorem, there is a unique $w(r,s) \in [w^m(t,r,s), \mathbf{E}_t[y(t,r,s)]$ that solves equation (6). In other words, as the wage rises, labor supply increases at a decreasing rate. Since a higher wage implies a lower markdown, competing firms raise the wage until the marginal gain in the number of workers offsets the markdown reduction.

Proposition 4. *The equilibrium wage* $w(r,s)$ *increases with* (n,b)*, falls with* r*, and* $\lim_{n\to\infty} w(r,s) \to y$ *and* $\lim_{n\to\infty} P(w(r,s)) \to 0$.

The equilibrium wage increases with competition intensity, as workers are more likely to find another paid job that they prefer to the one offered by firm j. This induces firm j to set a higher wage to attract workers. In the limit, as the number of firms approaches infinity, the worker is paid their marginal product of labor. This is because log-concave distributions have either a fat or a thin tail. Otherwise, the wage markdown will not converge to zero as the number of firms approaches infinity (see Gabaix et al. (2016) for details).

Due to increased competition, holding wages constant reduces the market exclusion effect by increasing the number of jobs available and thereby raising wages. Employment at each firm decreases with competition intensity, and Lerner's index falls.

An increase in the outside-option payoff raises the wage. Wages increase because workers choose the outside option more often when firms keep wages constant. Thus, firms increase wages less than the b increase. The pass-trough from b to wages is equal to $-m'/(1-m')$, which is lower than 1.[11] Thus, a larger outside-option payoff decreases market power because, holding wages constant, the labor-supply elasticity rises as more workers find the outside option more attractive.

3.3 Automation Decision

We allow mixed strategies regarding the automation decision. Let $\alpha^j \in [0,1]$ be firm j's probability to open a vacant, i.e., $\alpha^j = \mathbb{P}(d^j = 1)$.

Let firm j's expected profits from opening a vacant when competitors choose the mixed strategy be α^{-j} be $\mathbf{E}_{\alpha^{-j}}\Pi^j(d^{-j})$, where $\mathbf{E}_{\alpha^{-j}}$ is the expectation with respect to d^{-j} under the mixed-strategy profile α^{-j} and

$$\Pi^j(d^{-j}, s) \equiv \left(\mathbf{E}_t[y^j(t,r,s)] - w^j(b,r,s)\right) \times \int_{b-w^j(b,r,s)}^{\bar{\epsilon}} \prod_{k \in \mathcal{J}(d^{-j})} G\Big(w^j(b,r,s) + \epsilon^j - w^k(b,r,s)\Big) dG(\epsilon^j).$$

[11] If firms could choose non-pecuniary benefits together with wages, they will also use them to compete against self-employment opportunities up to the point where the marginal return of increasing non-pecuniary benefits is equal to that from raising the wage.

Firm j's problem is as follows

$$\max_{\alpha^j \in [0,1]} \{\alpha^j \mathbf{E}_{\alpha^{-j}} \Pi^j(d^{-j}, s) + (1 - \alpha^j) y(r, s)\}.$$

Thus, firm j's best response is given by

$$BR^j(\alpha^{-j}) = \begin{cases} 1 & \text{if } y(r,s) \leq \mathbf{E}_{\alpha^{-j}} \Pi^j(d^{-j}), \\ [0,1] & \text{if } y(r,s) = \mathbf{E}_{\alpha^{-j}} \Pi^j(d^{-j}), \\ 0 & \text{if } y(r,s) > \mathbf{E}_{\alpha^{-j}} \Pi^j(d^{-j}). \end{cases} \tag{7}$$

Firm j chooses to automate the job, provided that the competitor's mixed-strategy profile is α_{-j}, whenever this is more profitable than the expected profits of allocating the job to the worker. Thus, what matters to the firm when deciding on automation is the rent a non-automated job generates relative to the rent an automated job yields when competitors choose α^{-j}.

The next result readily follows from this and the existence theorem for Nash equilibria.

Proposition 5. *For each labor market s, there exists a sub-game perfect equilibrium $(\alpha(r,s), w(r,s))$.*

Let us assume symmetric firms and focus on a symmetric equilibrium. There are three types of symmetric equilibrium: i) one where all firms choose to offer a vacancy, and thereby, $\mathbf{d} = \mathbf{1}$; ii) one where all firms choose to automate their jobs, and thereby, $\mathbf{d} = \mathbf{0}$; and iii) one where all firms use a non-degenerate mixed strategy where they offer a vacancy with probability $\alpha \in (0,1)$, and thereby, $\mathbf{d} = \mathbf{1}$ with probability α and $\mathbf{d} = \mathbf{0}$ with probability $1 - \alpha$.

First, let us consider the case where $\mathbf{d}^{-\mathbf{j}} = \mathbf{1}$. Using the first-order conditions for wages, we deduce that firm j chooses $d^j = 1$ whenever

$$y(r,s) \leq y(r,s;\mathbf{1}) \equiv m(b - w(r,s)) \frac{1 - G(b - w(r,s))^n}{n}.$$

Second, let's consider the case in which $d^{-j} = \mathbf{0}$. Using the first-order conditions for

wages, it is easy to see that firm j chooses $d^j = 0$ whenever

$$y(r,s) \geq y(r,s;\mathbf{0}) \equiv m(b - w(r,s))(1 - G(b - w(r,s))).$$

This means that the profit from automation exceeds the profit when the firm is a monopoly in the job market.

Thirdly, let's consider the case in which for all $j \in \mathcal{J}$, $d^{-j} = 1$ with probability α and $d^{-j} = 0$ with probability $1 - \alpha$. Let the probability that the cardinality of the set $\mathcal{J}(d^{-j})$ is $v \leq n - 1$ be

$$P(t,n,\boldsymbol{\alpha}) = \binom{n-1}{v} \alpha^v (1-\alpha)^{n-1-v}$$

Observe that $P(t,n,\boldsymbol{\alpha})$ increases with α for all $v > (n-1)\alpha$ and decreases otherwise.

Using the first-order conditions for wages, firm j chooses $d^j = 1$ with probability 1 whenever each competitor i is choosing $d^{-j} = 1$ with probability α if and only if

$$y(r,s) = y(r,s;\boldsymbol{\alpha}) \equiv \mathbf{E}_v \left[m(b - w(r,s)) \frac{1 - G(b - w(r,s))^{1+v}}{1+v} \right]. \tag{8}$$

Because $w(r,s)$ rises with the number of firms v and the firm's labor supply falls with the number of firms v, profits decrease with the number of firms that choose to post a vacancy instead of automating the job. This implies that $y(r,s;\mathbf{0}) > y(r,s;\mathbf{1})$. This, together with $P(t,n,\alpha)$ being decreasing in α for v small and increasing for v large, implies that $y(r,s;\alpha)$ falls with α, since as α increases, more weight is placed in states where profits are small and less in those with large profits. Thus, we have the following result.

Proposition 6. *For each labor market s, let's consider a symmetric equilibrium.*

i) If $\mathbf{E}_t[y(t,r,s)] \leq y(r,s)$, then $d(b,r,s) = \mathbf{0}$.

ii) If $\mathbf{E}_t[y(t,r,s)] > y(r,s)$, then

a) There exists a threshold $y(r,s;\mathbf{1})$ such that for all $y(r,s) \leq y(r,s;\mathbf{1})$, the equilibrium is given by $d(b,r,s) = \mathbf{1}$. The threshold $y(r,s;\mathbf{1})$ rises with (s,r) and falls with n.

b) There exists a threshold $y(r,s;\mathbf{0})$ such that for all $y(r,s) > y(r,s;\mathbf{0})$, the equilibrium is given by $d(b,r,s) = \mathbf{0}$. The threshold $y(r,s;\mathbf{0})$ rises with (s,r).

c) For all $y(r,s;\mathbf{0}) > y(r,s) \geq y(r,s;\mathbf{1})$, the equilibrium is a mixed strategy equilibrium given by $d(b,r,s) = \mathbf{1}$ with probability $\alpha(r,s)$, with $\alpha(r,s)$ being the unique solution to $y(r,s) = y(r,s;\alpha)$. $\alpha(r,s)$ rises with (r,s) and falls with n.

A key empirical question concerning automation is the relationship between its adoption rate and the intensity of competition, as measured by the number of firms

Proposition 7. *Suppose firms are symmetric within each labor market s. The automation adoption rate in the labor market s is lower when market power increases (n falls) and, conditional on the equilibrium being in mixed strategies, when s rises.*

The comparative statics in each part are due to the pass-through from y and b to wages being positive and lower than 1; the equilibrium wage rises with the number of firms that post a vacancy, and the labor supply faced by each firm, holding the wage constant, falls with the number of firms posting a vacancy.

As the number of firms rises, which is our measure of competition intensity, equilibrium wages increase, and thereby, the rents of human-skill jobs fall. Thus, the more intense the competition, the higher the adoption rate of automation. Similarly, as s increases, equilibrium wages rise, but the markup increases because the pass-through is less than 1. This implies a larger rent for a human-skill job and a lower probability of automation under the mixed-strategy equilibrium.

4 Efficiency and Displaced Workers

In this section, we compare: i) the equilibrium job assignments with the productively efficient ones; and ii) the equilibrium job assignments with the welfare-efficient ones. To facilitate comparisons, we focus on the symmetric equilibrium.

4.1 Productive Efficiency

Because at the time the automation decision is made, firms do not know the workers' realized skills, and they fully anticipate the productivity of automation, it is productively efficient to automate the job whenever the expected output in an automated job exceeds

that for a human-skill jobs, i.e., $y(r,s) > \mathbf{E}_t[y(t,r,s)]$. Productive efficiency differs from allocative efficiency because the former does not account for non-pecuniary benefits.

We deduce the following result from this and Proposition 6.

Proposition 8 (Productive Efficiency)**.** *Suppose a symmetric equilibrium.*

i) Suppose that $\mathbf{E}_t[y(t,r,s)] \geq y(r,s)$.

a) If $y(r,s) > y(r,s;\mathbf{0})$*, the job is inefficiently automated, and workers are inefficiently displaced.*

b) If $y(r,s;\mathbf{0}) \geq y(r,s) > y(r,s,\mathbf{1})$*, the job is inefficiently automated and workers are inefficiently displaced with probability* $1-\alpha$.

c) If $y(r,s) \leq y(r,s;\mathbf{1})$*, the job is efficiently assigned to human skills, and workers are efficiently employed.*

ii) Suppose that $\mathbf{E}_t[y(t,r,s)] < y(r,s)$*, the job is efficiently automated, and workers are efficiently displaced.*

The driving force behind this productive inefficiency is that firms choose automation over human-skill jobs based on the rents they receive from each option, rather than on the actual productivity of each option. In fact, whenever $\mathbf{E}_t[y(t,r,s)] \geq y(r,s) > y(r,s,\mathbf{1})$, there is too much automation from the perspective of productive efficiency, since the firm shares the revenues from human-skill jobs with the worker, whereas it fully appropriates the rent from automation. Furthermore, workers in the labor market s are inefficiently displaced, since whenever this occurs, firms should have posted a vacancy and hired workers. Instead, they automate the task with probability at least $1-\alpha(s,r)$. In this case, workers in the labor market s either take their best outside opportunity or receive unemployment benefits. In contrast, whenever $\mathbf{E}_t[y(t,r,s)] < y(r,s)$, the job is efficiently automated since the rent from a human-skill job is lower than the job's productivity when automated, and this exceeds the productivity of a human-skill job. From this, we deduce the following corollary.

Corollary 1. *When the productivity of human-skill jobs exceeds that of automated jobs, i.e.,* $\mathbf{E}_t[y(t,r,s)] > y(r,s)$*, labor market power results in an inefficient adoption of automation and work displacement relative to productive efficiency.*

Reallocating some jobs to human skills would increase output by the difference between the productivity in human skills and that of automation. Displaced workers lose the pecuniary benefits (wages) and the non-pecuniary benefits they would have gotten if those jobs had not been automated. The automation of these jobs creates an inefficiency because the rents the firms would have earned from hiring a worker are lower than the rents earned from automation, even though productivity would have been higher. The lower the rent on human skills, i.e., the more competitive the market, the higher the risk of automation.

There are three reasons why productive inefficiency driven by market power entails a welfare loss for the workers: firstly, they are paid less than the marginal product of labor; secondly, some workers lose the non-pecuniary benefits since they are inefficiently displaced by automation, and thirdly, workers are displaced and end up working in jobs where they are even less productive such as home production, self-employment, or unemployed.

4.2 Constrained Welfare Efficiency

Let us consider a benevolent social planner who chooses automation to maximize total welfare, i.e., the sum of profits and workers' utility, but does not intervene in the market structure so as to limit labor market power. Thus, wages are not equal to the marginal product of labor.

The central planner solve the following problem: $\max_{d\in\{0,1\}^n}\{W(d)\}$, where

$$W(d) = \max\{d(n\Pi + U) + (1-d)(ny(r,s) + b)\},$$

$$U \equiv n\int_{b-w(r,s)}^{\bar{\epsilon}} (w(r,s)+\epsilon)G^{n-1}\big(\epsilon\big)dG(\epsilon), \tag{9}$$

and

$$\Pi \equiv (\mathbf{E}_t[y(t,r,s)] - w(r,s))\frac{1-G^n(b-w(r,s))}{n}.$$

Thus, it is efficient to automate the jobs in the labor market s whenever

$$ny(r,s) + b \geq n\int_{b-w}^{\bar{\epsilon}} (\mathbf{E}_t[y(t,r,s)] + \epsilon)G^{n-1}(\epsilon)g(\epsilon)d\epsilon.$$

After integration-by-parts,[12] it is easy to see that it is efficient to automate the job if and only if

$$y(r,s) \geq y^*(r,s),$$

where

$$y^*(r,s) =: \frac{1}{n}\left(\mathbf{E}_t[y] - b + \bar{\epsilon} - (\mathbf{E}_t[y] + b - w)G^n(b-w) - \int_{b-w}^{\bar{\epsilon}} G^n(\epsilon)d\epsilon\right),$$

where this fully accounts for the worker's outside option and non-pecuniary benefits.

It readily follows that firms' automation decisions overlook workers' non-pecuniary benefits and the value of the outside option, which are only partially captured by the equilibrium wage. Thus, we conclude the following.

Proposition 9 (Constrained Welfare Efficiency)**.** *Suppose a symmetric equilibrium.*

i) If $y(r,s) \geq y^*(r,s)$, the job is efficiently automated and workers are efficiently displaced.

ii) If $y^*(r,s) \geq y(r,s) > y(r,s,\mathbf{0})$, the job is inefficiently automated, and workers are inefficiently displaced.

iii) If $y(r,s,\mathbf{0}) \geq y(r,s) > y(r,s;\mathbf{1})$, the job is inefficiently automated with probability $1 - \alpha(r,s)$ and efficiently assigned to human skills with probability $\alpha(r,s)$. Workers are inefficiently displaced with probability $1 - \alpha(r,s)$.

iv) If $y(r,s) < y(r,s;\mathbf{1})$, the job is efficiently assigned to human skills, and workers are efficiently employed.

Let's define $y^*(r,s)|_{w=e_t[y]}$ as the threshold above which the productivity of automation must be for this to be welfare efficient when workers are paid their marginal product of labor. Notice that $y^*(r,s)|_{w=y} > y^*(r,s)$.

Corollary 2. *When the productivity of human-skill jobs exceeds that of automated jobs.*

[12] Integration-by-parts implies the following

$$n\int_{b-w}^{\bar{\epsilon}} \epsilon G(\epsilon)^{n-1} g(\epsilon)d\epsilon = \bar{\epsilon} - (b-w)G^n(b-w) - \int_{b-w}^{\bar{\epsilon}} G(\epsilon)^n d\epsilon.$$

i) If $y(r,s) \in (y(r,s;\mathbf{1}), y^*(r,s))$, there is more automation and work displacement than constrained allocative efficiency warrants.

ii) If $y(r,s) \in (y(r,s;\mathbf{1}), y^*(r,s)|_{w=e_t[y]})$, there is more automation and work displacement than allocative efficiency warrants.

5 Taxing Automation

Keynes (1929) predicted that the rapid spread of technologies would bring "technological unemployment". Leontief made a similar prediction: "Labor will become less and less important... . Machines will replace more and more workers. I do not see that new industries can employ everybody who wants a job". These ideas are echoed by businesspeople and politicians, who argue about the potential benefits of taxing automation, believing it will lead to significant job losses and lower wages. In this section, we study technological capital taxation under two different scenarios: first, one sector exogenous technological capital cost r, and second, we introduce several sectors and a competitive labor market.

5.1 Exogenous Capital Cost and One Sector Economy

Because automation is inefficiently high, levying a tax $\tau \in \Re$ on automated jobs could mitigate productive and allocative inefficiencies, thereby avoiding the inefficient displacement of workers.[13] This policy is known as a "robot tax," a proposed policy under which companies would pay a tax for using robots or automated systems that replace human workers. This proposal was made by the European Parliament and by entrepreneurs, including Bill Gates. South Korea has addressed the issue indirectly by reducing the tax credit, thereby lowering investment in automation and increasing employment. Namely, Kang, Lee, and Quach (2024) find, using Korean data, that a reduction in the tax credit reduces investments in automation and increase employment, lowers wage inequality due to slower wage growth in the upper half of the income distribution, and has a positive fiscal externality, implying that behavioral responses to reductions in the tax credit increased the government's revenue beyond the direct mechanical impact of the policy.

[13]The same result can be obtained by taxing the investment in robots a. This will result in the capitalization of automation decreasing with the tax rate τ. However, an investment tax will lead to inefficient investment in jobs that can be automated.

When $\mathbf{E}_t[y(t,r,s)] \geq y(r,s)$, then a tax satisfying $y(r,s) - \tau = y(r,s,\mathbf{1})$ re-establishes productive efficiency. In contrast, when $\mathbf{E}_t[y(t,r,s)] < y(r,s)$, no tax is required since the job is efficiently automated. When $y(r,s) \in [y(r,s,\mathbf{1}), y^*(r,s))$, re-establishing allocative efficiency requires again finding a tax so that $y(r,s) - \tau = y(r,s,\mathbf{1})$ so that all firms prefer to open a vacancy instead of automating the job. In contrast, when $y(r,s) > y^*(r,s)$, no tax is needed since $y^*(r,s) > y(r,s,\mathbf{1})$ and, thereby, firms automate the job as efficiency requires.

The above argument assume that robots can be directly taxed. However, this is challenging due to classification problems, which means that taxes on robots or new technologies, while potentially attractive in theory, may be difficult for governments to implement. Therefore, it is important to also consider the optimal tax implications of automation and generative AI under the assumption that such specific taxes are not feasible, e.g., see (Kina et al., 2023, Loebbing, 2022) for the same assumption. The most obvious alternative is to tax technoogical capital investments; however, this has the drawback of lowering the marginal product of labor, as adopting generative AI becomes more expensive and as automated jobs that can be efficiently automated become less productive. The tax also reduces the equilibrium wage; therefore, the outside option is chosen more often, resulting in a corresponding loss of expected non-pecuniary benefits.

To see this, let's suppose that an ad-valorem tax $\tau \in [1, \infty)$ is levied on technological capital. Then, the output in a human skill job is $\mathbf{E}_t[y(t,r,s,\tau)]$ and that in an automated job is $y(r,s,\tau)$, where these result from firms facing a marginal capital cost of technological capital equal to τr.

Because of the envelope theorem,

$$\frac{\partial \mathbf{E}_t[y(t,r,s,\tau)]}{\partial \tau} = -r\mathbf{E}_t[h(t,r,s,\tau)] \text{ and } y_\tau(r,s,\tau) = -ra(t,s,\tau).$$

Because $\mathbf{E}_t h(t,r,s,\tau)$ rises with s, the impact of technological capital tax on human skill jobs is larger in labor markets when more advanced skills are needed.

If $\mathbf{E}_t[y(t,r,s)] > y(r,s) > y(r,s,\tau;\mathbf{0})$, then to reestablish productive efficiency requires that $y(r,s,\tau)$ to fall with τ at a faster rate than $y(r,s,\tau;\mathbf{0})$, while if $y(r,s,\tau;\mathbf{0}) \geq y(r,s) > y(r,s,\tau;\mathbf{1})$, $y(r,s,\tau)$ must fall with τ at a faster rate than $y(r,s,\tau;\mathbf{1})$. Thus, for $v \in \{1,n\}$,

the following must hold

$$-rA(r,s,\tau_s) < \frac{1}{v}\frac{r\mathbf{E}_t h(t,r,s,\tau)}{1-m'(b-w)}\Big(tm'(b-w)(1-G(b-w)^v) - m(b-w)G(b-w)^{v-1}g(b-w)\Big). \tag{10}$$

If $y^*(r,s) > y(r,s) > y(r,s,\tau;\mathbf{0})$, reestablishing allocative efficiency requires the condition in equation 10 to hold.

When $y(r,s,\tau;\mathbf{0}) \geq y(r,s) > y(r,s,\tau;\mathbf{1})$, the central planner choose the tax to maximize total welfare

$$\alpha(r,s,\tau_s)n\int_{b-w}^{\bar{\epsilon}}(\mathbf{E}_t[y(t,r,s,\tau)]+\epsilon)G^{n-1}(\epsilon)g(\epsilon)d\epsilon + (1-\alpha(r,s,\tau_s))(ny(r,s,\tau)+b),$$

where $\alpha(r,s,\tau_s)$ is the mixed strategy when the tax rate is τ. It is useful to write this as follows

$$\alpha(r,s,\tau_s)ny^*(r,s,\tau) + (1-\alpha(r,s,\tau_s))ny(r,s,\tau) + b.$$

Because $\mathbf{E}_t[y(t,r,s,\tau)]$ falls with τ, and the pass-through from $\mathbf{E}_t[y(t,r,s,\tau)]$ to wages is lower than 1, the rent from human-skill jobs falls with τ. This implies that $y(r,s,\tau;\alpha(r,s))$ falls, holding $\alpha(r,s)$ constant, with τ. Thus,

$$\alpha_\tau(r,s,\tau) = \frac{y_\tau(r,s,\tau) - y_\tau(r,s,\tau;\alpha)}{y_\alpha(r,s,\tau;\alpha)}\Big|_{\alpha=\alpha(r,s,\tau_s)} \lesseqgtr 0,$$

where $y(r,s,\tau;\alpha)$ is defined in equation (8) and falls with α. Observe that

The first-order condition is as follows

$$\underbrace{n_s\,\alpha_\tau(r,s,\tau)\big(y^*(r,s,\tau)-y(r,s,\tau)\big)}_{\text{Allocation effect, } A(r,s,\tau)} + \underbrace{n_s\Big[\alpha(r,s,\tau_s)\,y^*_\tau(r,s,\tau) - (1-\alpha(r,s,\tau_s))\,r\,A(r,s,\tau_s)\Big]}_{\text{Resource-cost effect, } \Xi(r,s,\tau)>0} \lesseqgtr 0, \tag{11}$$

where

$$y^*_\tau(r,s,\tau) = -r\frac{\mathbf{E}_t[h(t,r,s,\tau)]}{n}\left(1 - G^n(b-w) + n(\mathbf{E}_t[y]+b-w)\frac{G^{n-1}(b-w)g(b-w)}{1-m'(b-w)}\right) < 0.$$

The allocation effect $A(r,s,\tau)$ has the sign of $\big(y^*(r,s,\tau) - y(r,s,\tau)\big)\big(y_\tau(r,s,\tau) - y_\tau(r,s,\tau;\alpha)\big)$, and the resource-cost effect $\Xi_s(r,s,\tau) < 0$, since $a(\cdot) > 0$ and $h(\cdot) > 0$. Consequently, it is optimal to tax AI if and only if the allocation effect dominates the resource-cost effect, $A(r,s,\tau) > -\Xi(r,s,\tau) > 0$, whereas it is optimal to subsidize it if and only if $A(r,s,\tau_s) < -\Xi(r,s,\tau)$. Thus, we have the following.

Proposition 10. *If $A(r,s,\tau) > -\Xi(r,s,\tau)$, a marginal tax increase from $\tau = 1$ strictly raises welfare; otherwise, it strictly lowers welfare.*

An increase in the technological capital tax lowers wages and productivity in both automated and human-skill jobs due to inefficient investment in artificial intelligence, and lowers the rent from a human-skill job due to the incomplete pass-through. Thus, holding the probability that the job is assigned to human skills constant, welfare falls with the tax rate. However, a marginal increase in the tax could eventually increase the probability that the vacancy is open, and this leads to an increase in welfare since a human job is more productive than an automated job. Thus, taxing capital results in the following trade-off. On the one hand, it improves the allocation of workers to jobs by reducing inefficient automation whenever $A(r,s,\tau) > -\Xi(r,s,\tau)$. Therefore, firms are more likely to open a vacancy when it is efficient to do so. On the other hand, those who keep their jobs are paid a lower expected wage than before, as generative AI capital falls with rising tax rates. This also induces more workers to take their outside options, harming efficiency. Thus, welfare falls. The optimal tax rate balances this trade-off.

In contrast, whenever $A(r,s,\tau) < -\Xi(r,s,\tau)$, taxing capital worsens the allocation of workers to jobs by increasing inefficient automation. This reinforces the effect on wages and productivity, making automation capital taxes a poor instrument for improving welfare. In this case, a subsidy on technological capital will be needed.

An alternative is to subsidize the hiring of workers, since this increases the rent for a human-skill job since the pass-through is incomplete while holding the rent for an automated job constant.

5.2 Capital Markets and Multiple Sectors

The analysis so far treats the capital cost r as a single exogenous parameter common to one labor market. Here, we embed the subgame-perfect equilibrium already derived for one industry into an economy with many strategically independent industries, indexed by $s \in \{1, \ldots, S\}$, that differ in every primitive of the model—the skill distribution $F(t|s)$, the number of competing firms n_s, the outside option b_s, and the automation and gen-AI production technologies—while leaving the *within-industry* game entirely unchanged.

Let's assume that industry s is populated by n_s symmetric firms as in Section 2, with skill distribution $F(t|s)$, non-pecuniary shocks distribution $G_s(\cdot)$ satisfying Assumption 1, outside option b_s, and automation and human-skill production technologies $y^j(a; s)$ and $y^j(t, h; s)$ satisfying Assumption 2.

As in the preceding subsection, we assume that automation capital and generative-AI capital are, physically, the same scarce input—compute, specialized hardware, and the engineering capacity to deploy it—drawn from a common, competitive market. Let $r > 0$ denote the *economy-wide, pre-tax* rental price of this capital, determined by market clearing below. The government can levy an ad-valorem wedge $\tau_s \geq 0$ on the use of technological capital in industry s, so that firms in industry s face an effective capital cost

$$r_s \equiv \tau_s r, \tag{12}$$

exactly as in Section 5, except that τ_s is now allowed to differ across industries and r is no longer exogenous.

Because r_s enters industry s's problem exactly as r entered the baseline problem, every result in Sections 3–4 applies verbatim to industry s upon substituting r_s for r and s for s. In particular: the optimal automation and gen-AI investments are $a_s(r_s, s)$ and $h_s(t, r_s, s)$, as characterized in Section 3.1; the equilibrium wage $w(r_s, s)$ is the unique solution to equation (6) evaluated at (r_s, s), by Proposition 3; the automation decision is the mixed-strategy equilibrium$\alpha(r, s, \tau_s)$ characterized in Proposition 6, with automation rate $1 - \alpha(r, s, \tau_s)$; welfare in industry s, $W(r_s, s) \equiv W(d(r_s, s))$, is given by the planner's problem in Section 4.2 evaluated at (r_s, s).

We collect these objects for later use: let

$$D(r_s, s) \equiv n_s \big[(1 - \alpha(r, s, \tau_s))\, a(r, s, \tau_s) + \alpha(r, s, \tau_s)\, \mathbb{E}_t[h(t, r_s, s)] \big]$$

denote industry s's total expected demand for technological capital at effective price r_s. By the comparative statics already established in Section 3.1—$a(\cdot)$ and $h(t, \cdot)$ are each strictly decreasing in their own argument, by Assumption 2(i)–(ii) and the implicit function theorem applied to the first-order conditions—$D(\cdot, s)$ is continuous and strictly decreasing in r_s.

Let $S(r)$ denote the aggregate, economy-wide supply of technological capital, non-decreasing and continuous in the pre-tax price r. The pre-tax capital price is determined by market clearing:

$$\sum_{s=1}^{S} D(\tau_s r, s) \;=\; S(r). \tag{13}$$

Definition 3 (Industry Equilibrium)**.** *Given a tax vector $\tau \equiv (\tau_1, \ldots, \tau_S)$, an* industry equilibrium *is a price $r^*(\tau) > 0$ together with, for each s, an industry-s subgame-perfect equilibrium $(\alpha(\tau_s r^*(\tau), s), w(\tau_s r^*(\tau), s))$ as characterized in Propositions 2–6, such that equation* (13) *holds at $r = r^*(\tau)$.*

Because each $D(\tau_s r, s)$ is continuous and strictly decreasing in r (Section 3.1), yhe aggregated demand $\sum_s D(\tau_s r, s)$ is continuous and strictly decreasing in r, from a finite value at $r = 0$ to 0 as $r \to \infty$ under the boundary conditions on $A(r, s, \tau_s)$ and $h(t, r_s, s)$ in Assumption 2. Because $S(r)$ is continuous, non-decreasing, and spans $[0, \infty)$, the Intermediate Value Theorem delivers a crossing point; strict monotonicity of S relative to the strictly decreasing aggregate demand delivers uniqueness. Thus, we deduce the following result

Proposition 11 (Existence and Uniqueness of the Industry Equilibrium)**.** *Suppose $S(\cdot)$ is non-decreasing and continuous, with $S(0) = 0$ and $\lim_{r\to\infty} S(r) = \infty$. Then, for any $\tau \gg 0$, an industry equilibrium exists. If, in addition, $S'(r) > 0$ wherever $S(r) < \sum_s D(r_s, s)$, the equilibrium price $r^*(\tau)$ is unique.*

Two special cases nest existing results and are worth noting explicitly. First, if $S(r) \to \infty$ for all r above some floor $\underline{r}$—the small-open-capital-market case—the market clears at $r^* = \underline{r}$ regardless of $\{\tau_s\}$, and each industry's equilibrium collapses exactly to the single-sector model of Sections 2–5 with $r = \underline{r}$ fixed, taxed at the industry-specific rate τ_s: this

is the baseline model, industry by industry, with no cross-industry linkage. Second, if $s = 1$, equation (13) simply relabels the baseline model's exogenous r as the market-clearing price against a single industry's demand, recovering Section 5 verbatim. The multi-industry link is therefore a strict generalization that becomes empirically relevant precisely when $S(\cdot)$ is not perfectly elastic—i.e., when technological capital (computing power, chips, specialized engineering talent, hardware production capacity) is scarce in the short-to-medium run, as the recent behavior of AI compute markets suggests.

Total welfare is $W(\tau) \equiv \sum_{s=1}^{S} W\big(\tau_s r^*(\tau), s\big)$, and the planner chooses τ to maximize $W(\tau)$ and, given the chosen taxes, the capital price is determined by the capital market clearing condition in (13). Because r^* depends on the entire vector τ, a change in τ_s affects welfare through two channels: the direct, within-industry channel already characterized in Section 5 (equation (11)), and an indirect, general-equilibrium channel operating through every other industry's effective capital price.

Differentiating equation (13) with respect to τ_s using the implicit function theorem to obtain the interest response to a tax change and differentiating $W(\tau)$ with respect to τ_s, at an interior optimum, the tax vector τ^* satisfies, for each s,

$$\underbrace{\frac{\partial W_s}{\partial r_s} r(\tau)\Big|_{r=r^*}}_{\text{Direct Effect ((11), industry } s)} + \underbrace{\frac{\partial r^*}{\partial \tau_s}}_{\text{Industry } s\text{'s leverage over } r} \underbrace{\sum_{s'} \frac{\partial W_{s'}}{\partial r_{s'}} \tau_{s'}\Big|_{\tau_{s'} r^*}}_{\text{Welfare Change due to Change in } r} = 0, \quad (14)$$

where

$$\frac{\partial r^*}{\partial \tau_s} = \frac{-\, r^*\, \partial D / \partial r_s\big|_{\tau_s r^*}}{\sum_{s'=1}^{S} \tau_{s'}\, \partial D' / \partial r_{s'}\big|_{\tau_{s'} r^*} \;-\; S'(r^*)} < 0. \quad (15)$$

Equation (14) is the natural multi-sector generalization of the single-sector optimal tax condition in Section 5. The direct term is exactly equation (10), industry by industry: it is positive or negative depending on whether the gen-AI margin or the automation margin is more sensitive to the industry's own tax, as already characterized. The second term is the reason taxation should generally differ across industries even when every industry's underlying technology and demography were identical: an industry with large leverage over the common price—large n_s, or a demand for capital that is highly sensitive to its own effective price—moves r^* more than a small industry does, and therefore imposes a larger pecuniary externality on every other industry through the shared capital market.

The planner internalizes this by taxing (or subsidizing) high-leverage industries differently from low-leverage ones, exactly as a Ramsey planner taxes goods with different demand elasticities at different rates to minimize the aggregate distortion from a limited set of instruments.

Multiplying each first-order condition for τ_s and summing over s, we get that

$$\sum_{s'} \frac{\partial W_{s'}}{\partial r_{s'}} \tau_{s'} \left(1 + \sum_{s} \frac{\tau_s}{r(\tau)} \frac{\partial r(\tau)}{\partial \tau_s} \right) = 0.$$

Substituting the capital market equilibrium condition, it is easy to see that if $S'(r(\tau) > 0$, then $\frac{\partial W_s}{\partial r_s} = 0$ for each s. We collect this in the next lemma.

Lemma 2. *Suppose that $S'(r) > 0$, then the optimal tax/subsidy in industry s satisfies $\frac{\partial W_s}{\partial r_s} = 0$. If industries differ in at least one of $\{n_s, b_s, F(t|s), y^j(\cdot; s)\}$. Then a uniform tax, $\tau_s = \tau$ for all s, satisfies the first-order condition* (14) *simultaneously for every s only for a measure-zero set of primitives.*

The corollary follows immediately from the fact that the direct term in equation (14) is industry-specific—it depends on $(\mu_{t,s}, \sigma_{t,s}, n_s, b_s)$, while the spillover term depends on each industry's demand elasticity relative to every other industry's welfare sensitivity to r. Generically, no single τ solves s distinct equations of this form simultaneously.

The lemma above establishes that a uniform tax is generically *suboptimal*. A stronger and more policy-relevant statement is that a uniform tax or subsidy can move welfare in *opposite directions* across industries, so that even the *sign* of the optimal uniform policy—whether the government should tax or subsidize automation capital economy-wide—is undetermined without further information on the cross-industry distribution of primitives.

Consider a uniform wedge, $\tau_s = \tau$ for all s, and let $r^*(\tau)$ denote the resulting equilibrium price from Proposition 11. Normalize the undistorted economy to $\tau = 1$, and let $r_0 \equiv r^*(1)$ denote the corresponding baseline capital cost. For each industry s, define

$$\Phi(r, s, \tau) \equiv \left. \frac{dW_s}{dr_s} \right|_{r_s = r}, \tag{16}$$

the marginal effect on industry s's welfare of the effective price it faces, evaluated us-

ing the planner's first-order condition in equation (11) applied to each market s. Consequently, as shown in Proposition 10, $\Phi(r, s, \tau) > 0$ if and only if the allocation effect dominates the resource-cost effect, $A(r, s, \tau_s) > \Xi(r, s, \tau_s)$, and $\Phi(r, s, \tau) < 0$ if and only if $A(r, s, \tau_s) < \Xi(r, s, \tau_s)$.

Proposition 12. *Suppose $S'(r_0) > 0$, and suppose there exist industries s and s' such that $A(r_0, s', \tau_{s'}) < \Xi(r_0, s', \tau_{s'})$ and $A(r_0, s', \tau_{s'}) > \Xi(r_0, s', \tau_{s'})$ both hold at r_0. Then:*

i) A marginal uniform tax increase from $\tau = 1$ strictly raises welfare in industry s and strictly lowers welfare in industry s':

$$\left.\frac{dW_s}{d\tau}\right|_{\tau=1} > 0 \qquad \text{and} \qquad \left.\frac{dW_{s'}}{d\tau}\right|_{\tau=1} < 0.$$

ii) The sign of the aggregate welfare effect,

$$\left.\frac{dW}{d\tau}\right|_{\tau=1} = \big(r_0 + r^{*\prime}(1)\big) \sum_{s=1}^{S} \Phi(r_0, s, \tau), \tag{17}$$

In particular, there exist economies—differing only in the competition intensity n_s or productivity of the industries satisfying $A(r_0, s', \tau_{s'}) < \Xi(r_0, s', \tau_{s'})$ versus those satisfying $A(r_0, s', \tau_{s'}) > \Xi(r_0, s')$—in which the welfare-maximizing uniform policy is a tax ($\tau^* > 1$), and other economies, with the same sets of $A(r_0, s', \tau_{s'}) < \Xi(r_0, s')$ and $A(r_0, s', \tau_{s'}) > \Xi(r_0, s', \tau_{s'})$ industries but different relative weights, in which the welfare-maximizing uniform policy is a subsidy ($\tau^* < 1$).

This says that the suboptimality of a uniform instrument is not even compatible with a determinate sign (e.g., some tax is better than no tax, even if the rate is wrong); what we show here is that even the sign can be indeterminate, i.e., even a central planner that observes $A(r_0, s)$ and $\Xi(r_0, s)$ for each industry s cannot infer the sign of the right economy-wide policy from this qualitative information alone; it must weigh the sectors by their relative size and the magnitude of each $\Phi(r, s, \tau)$. Thus, not only a uniform tax is the wrong *rate*, its *sign* is, absent further restrictions, undetermined, whereas the differentiated policy of Proposition 2 resolves the ambiguity industry by industry, taxing s and subsidizing s' simultaneously.

Thus, the *corrective instrument* should not be uniform across industries even when the government's only goal is to correct that wedge economy-wide. Two forces drive the industry-specific optimal rate. First, industries with a large gen-AI-to-automation productivity ratio have a larger direct term in equation (14) in the direction of *subsidizing* capital, since the gen-AI channel identified in Section 5 is stronger there. Second, industries that are large relative to the aggregate capital market, or whose capital demand is highly price-elastic, have larger $|\partial r^*/\partial \tau_s|$ and therefore a stronger incentive for the planner to use them, rather than smaller or less elastic industries, to move the common capital cost in the direction that benefits every other industry. In practice, this suggests that sectors such as manufacturing, where automation capital is a large share of a large, elastic capital stock, are natural candidates for a different, less generous subsidy rate than a smaller, human-capital-intensive sector such as professional services, where the gen-AI channel is likely stronger but the industry's leverage over the aggregate capital price is small.

6 Quantitative Analysis

In this section, we provide descriptive evidence consistent with the model's central mechanism (Proposition 7) and use it to discipline a calibrated structural exercise. We do not claim causal identification of the effect of labor market power on automation adoption; a credible design for that question would require plant-level entry/exit shocks or comparable quasi-experimental variation in labor-market concentration, which is not available at the resolution our model requires. Instead, we document two cross-sectional associations, in different data sources and geographies, that move in the direction the model predicts and are robust to the standard controls used in this literature. We then use the calibrated model, disciplined by these associations and by targeted moments, to quantify the mechanism directly. The calibration, not the reduced-form correlations, is the paper's main quantitative contribution; the regressions serve only to establish that the qualitative pattern the model requires is present in the data before we take the model to numbers.

6.1 Descriptive Evidence

We document two empirical regularities consistent with the model's central prediction (Proposition 7): higher labor market power is associated with less substitutive automation. We do not claim causal identification of the effect of labor market power on automation; rather, we show that across two data sources covering different geographies and measures, market power is robustly negatively associated with automation, even after controlling for the standard determinants emphasized in the literature. In both cases, we follow the empirical specifications of Acemoglu and Restrepo (2022), who study labor-force aging and automation at the commuting-zone and country levels, adapting their controls to our setting.

We use two complementary pieces of evidence. First, at the U.S. commuting-zone level, we show a negative association between labor market concentration—measured as the Herfindahl-Hirschman Index (HHI) of posted vacancies (Choi and Marinescu, 2024)—and the local presence of robotics activity, and the evidence suggests that this association is driven by concentration in the occupations most exposed to robots. Second, at the country level, we show a negative association between manufacturing markdowns and the (log) net imports of robots per one thousand industrial workers; the result extends to a broader set of automation technologies, while a placebo of complementary technology (computers) shows no such relationship.

6.2 Evidence at the Commuting-Zone Level

Appendix B.1 provides a complete description of the data and its sources. Following Acemoglu and Restrepo (2022), we measure automation adoption at the CZ level with a dummy indicator of the presence of robot integrators—a firm that installs, programs, and maintains industrial robots—in that CZ as of 2015, based on Leigh and Kraft (2018). Its CZ mean is $\bar{d} = 0.190$.

We define a local labor market as a commuting zone, because it is a closer representation of the options available to a worker searching for job in a particular labor market. Data on vacancy HHI come from Choi and Marinescu (2024), who construct upper- and lower-bound Herfindahl–Hirschman Index (HHI) estimates for online vacancy postings at the CZ$\times$6-digit SOC level using Lightcast (formerly Burning Glass Technologies) data.

This tracks online vacancy postings from over 45,000 online job boards, carefully removes duplicates, and cleans the data. The resulting dataset covers the near universe (70%) of all U.S. vacancy postings and comprises 250 million job vacancy postings for the years of 2007 and 2010-2019. We use the 2007 midpoint HHI as a pre-period baseline, free of contamination from the Great Recession.[14] We aggregate to an employment-weighted HHI at the commuting-zone level (weighting occupations by national OEWS employment; see Appendix B.1.1) and use its 2007 average. Our main specification uses the midpoint of the bounds; results by bound are in Appendix B.1.7.

We estimate,

$$integrators_c = \beta_0 + \beta_1 HHI_c + \Gamma X_{c,1990} + \nu_c,$$

where c indexes the commuting zone, $integrators_c$ is a dummy for the presence of robot integrators, HHI_c is the 2007 average HHI, and $X_{c,1990}$ is a rich set of controls (mostly measured in 1990) for demographics, industry composition and robot exposure, worker routinization, and Chinese import exposure, as in Acemoglu and Restrepo (2022).

The results are in Table 1. Our preferred specification is column (4), with the full set of controls. The standard deviation of the midpoint HHI in 2007 is 0.16, so the column (4) coefficient implies that a one-standard-deviation increase in concentration lowers the probability of integrator presence by about 4.4%. Column (5) shows the results considering only commuting zones with good Lightcast coverage across occupational groups[15]. Figure 1 plots the predicted probability of integrator presence against HHI from column (4). In Appendix B.1.4 we show that the 2007 HHI is highly autocorrelated over time, so the regressor captures structural labor market concentration.

The model's prediction concerns substitutive automation, so the relevant object is concentration in the occupations that robots can replace. Table 2 recomputes the HHI over occupation subsets defined by the robot-exposure index of Webb (2020): columns (1) and (2) use all occupations, columns (3) and (4) the top-35th-percentile (most exposed) and bottom-35th-percentile (least exposed) occupations. Column (1) runs the regression on the full sample while columns (2)–(4) exclude the bottom-quartile commuting zones with the worst occupational coverage (measured over all occupations). The negative as-

[14] See Choi and Marinescu (2024) for construction details.

[15] We consider a commuting zone to have good Lightcast coverage if it is in the upper three quartiles of employment-weighted occupational coverage—the share of national employment in the occupations observed in every quarter of 2007 in the respective commuting zone.

Table 1. Labor Market Concentration and the Local Presence of Robot Integrators: Baseline Specification (Midpoint HHI)

HHI 2007	-1.3131***	-0.2694***	-0.2660***	-0.2689***	-0.6734***
	(0.1183)	(0.0958)	(0.0838)	(0.0835)	(0.1451)
Census Region Dummies	✓	✓	✓	✓	✓
Demographic controls (1990)		✓	✓	✓	✓
Industry controls + robot exposure			✓	✓	✓
Other shocks controls				✓	✓
Good cov.					✓
Observations	721	721	721	721	541
R-squared	0.31	0.41	0.45	0.45	0.46

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Columns 1-4 progressively add controls. Column 1 includes only Census region dummies. Column 2 adds demographic controls. Column 3 adds industry employment shares and the commuting-zone-level exposure to robots over 1993-2007. Column 4 adds controls for exposure to Chinese imports and the routine occupation share. Column 5 repeats the fully-controlled specification of column 4 on the subsample of commuting zones above the 25th percentile of employment-weighted occupational coverage (the share of national employment in the occupations observed in every quarter of 2007 in the respective commuting zone). Demographic controls include the male share, the urban share, log per-capita income, log population, and shares of workers by five education levels and five racial groups, all measured in 1990. Industry controls include 1990 employment shares in agriculture, mining, construction, manufacturing, and finance and real estate. For a more detailed description of data sources refer to Appendix Section B.5.

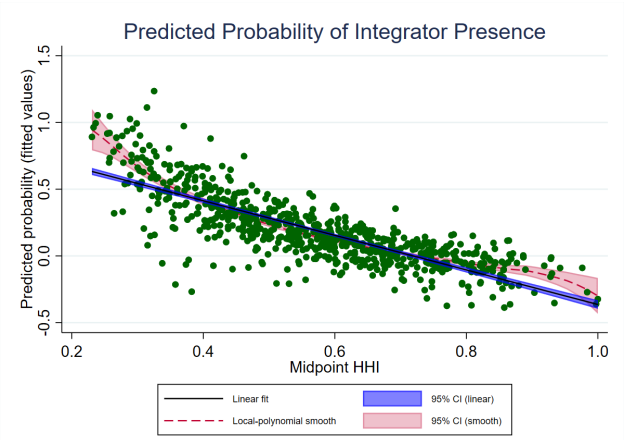


Fig. 1. Predicted Probability against HHI 2007.
Predicted probability of robot-integrator presence from the column (4) specification of Table 1. HHI is the 2007 employment-weighted vacancy HHI (Choi and Marinescu, 2024); the integrator outcome is from Leigh and Kraft (2018) as used by Acemoglu and Restrepo (2022). See Appendix B.1 for data sources and construction.

sociation is concentrated in the HHI of the high-exposure occupations and is weaker for low-exposure occupations, consistent with the pattern the model predicts: market power co-moves with less substitutive automation specifically in the occupations robots can replace. Because highly automatable occupations tend to have lower occupational coverage, we exclude commuting zones with poor occupational coverage. The results should therefore be interpreted with caution, as they are sensitive to this sample restriction.

Online job postings over-represent higher-skill occupations and industries (Choi and Marinescu, 2024, Hershbein and Kahn, 2018). In Appendix B.1.5 we document that less-covered occupations are more robot-exposed and more concentrated, so coverage error biases against finding our result. We also show the estimates are robust to excluding commuting zones with high concentration of low-Lightcast-coverage-industries and to recomputing the HHI over occupation subsets, also defined by Lightcast coverage (Tables 12 and 13).

Using the lower-bound HHI yields a similar but smaller pattern; the upper bound yields larger coefficients (Appendix B.1.7).

The key threat to identification is that there is a time-varying market-specific variable that is correlated with HHI and automation decisions. The key confound, according to the oligopsony theory discussed above, is labor productivity. What other confounds are most likely? According to search and matching theory, posted wages are determined by labor market tightness, productivity, and the worker's out-of-work benefit (Rogerson, Shimer, and Wright 2005). Since unemployment benefits are determined at the state level, we can control for workers' out-of-work benefits by including market fixed effects and, in some specifications, market-by-time fixed effects. Therefore, the main threat to identification remains time-varying, market-specific productivity changes.

A natural concern is that commuting-zone-level HHI may be endogenous to local shocks correlated with both labor-market structure and automation demand. We probe this concern through several approaches, though, as discussed below, our attempt at instrumental variation proves underpowered, and we do not treat it as resolving the concern. First, we progressively add controls for baseline industry composition, demographic characteristics, and local economic shocks, and verify that OLS estimates remain stable (Table 1). Second, we conduct falsification tests examining whether HHI predicts pre-sample trends in robot adoption and whether vacancy-concentration measures con-

Table 2. Mechanism: Labor Market Concentration in High- vs Low-Automation Occupations (Midpoint HHI)

HHI 2007	-0.2689***	-0.6734***	-0.6822***	-0.5959***
	(0.0835)	(0.1451)	(0.1674)	(0.1391)
Occupation set	All occ.	All occ.	Top 35pp	Bottom 35pp
CZ sample	Full	Good cov.	Good cov.	Good cov.
Base controls	✓	✓	✓	✓
Observations	721	541	541	541
R-squared	0.45	0.46	0.46	0.46

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Column 1 reports the aggregate HHI on the full sample. Columns 2–4 restrict to commuting zones above the 25th percentile of employment-weighted occupational coverage (share of national employment observed, over all occupations). Column 2 uses all occupations, column 3 the occupations in the top 35% of the robot-exposure index of Webb (2020), and column 4 those in the bottom 35%. A stronger negative coefficient for the high-exposure occupations is consistent with the model's mechanism: substitutive automation responds to labor market concentration in the occupations it can replace. All specifications include the full set of controls of Table 1.

structed in earlier periods exhibit similar relationships (results available upon request). Third, we instrument HHI using a Bartik-style index of national industry-level unionization declines interacted with 1990 local industry composition, motivated by the observation that secular union decline was driven by national policy shifts largely orthogonal to local automation demand. However, as reported in Appendix Table 20, this instrument exhibits weak first-stage power conditional on rich controls ($F < 2$ in most specifications), suggesting that residual variation in national union trends is limited after accounting for industry composition and robot-exposure measures. We therefore treat the IV estimates as a robustness check rather than a definitive causal identification and rely primarily on the OLS results, their stability across specifications, and the calibrated model (Section 6) to support our quantitative conclusions.[16]

We cannot measure employment-based labor market concentration or firm-level labor supply elasticities from public U.S. data: the sources commonly used for this[17] are released only in firm-censored form, with employer identifiers available solely through restricted-access research data centers (firm-level labor supply elasticities, for instance,

[16]The cross-country analysis in the online appendix confirms the CZ results for the US.

[17]For example, the Quarterly Census of Employment and Wages, the Longitudinal Employer–Household Dynamics (LEHD), the Longitudinal Business Database, and the establishment-level Economic Census and County Business Patterns.

require LEHD microdata, Webber, 2015).

The only public employment-based measure we found, the BLS "HHI Measures" file (Handwerker and Dey, 2022, 2024), defines markets as an occupation within an OEWS area (metropolitan areas plus balance-of-state aggregates) and groups occupations into roughly 392 coarse categories rather than detailed six-digit SOC. This is not comparable to our six-digit-SOC $\times$ commuting-zone definition—the coarser grouping mechanically lowers concentration and cannot be disaggregated, and the geography covers only metropolitan counties. We therefore judge it unreliable for our design and do not use it.

6.3 Calibration

We now turn to the quantitative analysis. First, we make specific functional form assumptions. Next, we calibrate the model to the US economy such that it reproduces the empirical average adoption rate of robots, the average labor supply, and the average wage markdowns. The calibrated model is then used to compute optimal taxes and to study the welfare impact of taxes and incentivizing competition.

6.3.1 Parametric Functional Forms and Structural Parameters

To simplify the calibration, we will assume that the expected output can is as follows:

$$\int_0^{\bar{y}} y dH(y|t,h,s) = \int_0^{\bar{y}} y dH(y|f(t,h,s))$$

where $f(t,h,s)$ is the production function and $h(y|\cdot)$ is log normal and $s = 1$. We will assume that $f(t,h,s)$ is Cobb-Douglas $A_h t^{\alpha} h^{1-\alpha}$ for the human skill job and $A_a a^{1-\alpha}$ for the automation job with $\alpha \in (0,1)$. The inner integral is the production function. This gives the expected output under automation. In fact, the optimal automation investment is $a^* = \left(A_a \frac{1-\alpha}{r}\right)^{\frac{1}{\alpha}}$ and the expected output minus the automation cost is given by

$$y(r,s) = A_a^{\frac{1}{\alpha}} \alpha \left(\frac{1-\alpha}{r}\right)^{\frac{1-\alpha}{\alpha}}$$

.

To get the expected output in a human job, we expand the term t^{α} around the mean μ_t

using a second-order Taylor series:

$$t^{\alpha} \approx \mu_t^{\alpha} + \alpha\mu_t^{\alpha-1}(t-\mu_t) + \frac{1}{2}\alpha(\alpha-1)\mu_t^{\alpha-2}(t-\mu_t)^2$$

Now, we integrate this approximation with respect to the skill distribution $G(t)$:

$$\int_0^{\bar{t}} t^{\alpha} dG(t) \approx \int_0^{\bar{t}} \left[\mu_t^{\alpha} + \alpha\mu_t^{\alpha-1}(t-\mu_t) + \frac{1}{2}\alpha(\alpha-1)\mu_t^{\alpha-2}(t-\mu_t)^2\right] dG(t)$$

By the definition of moments ($\int (t-\mu_t)dG(t) = 0$ and $\int (t-\mu_t)^2 dG(t) = \sigma_t^2$), the integral simplifies to:

$$E[t^{\alpha}] \approx \mu_t^{\alpha} + \frac{\alpha(\alpha-1)}{2}\mu_t^{\alpha-2}\sigma_t^2$$

Substitute this back into the production function aggregate:

$$\int_0^{\bar{t}} \left[A_h t^{\alpha} h^{1-\alpha}\right] dG(t) = A_h h^{1-\alpha}\left[\mu_t^{\alpha} + \frac{\alpha(\alpha-1)}{2}\mu_t^{\alpha-2}\sigma_t^2\right]$$

Let's define $A(\mu_t, \sigma)$ as the term multiplying $h^{1-\alpha}$. Then, it is easy to show that the optimal of gen-AI is $h^* = \left(A(\mu_t,\sigma)\frac{1-\alpha}{r}\right)^{\frac{1}{\alpha}}$ and the expected output minus gen-AI cost in a human job is

$$y(t,r,s) = \alpha A(\mu_t, \sigma_t^2)^{\frac{1}{\alpha}} \left(\frac{1-\alpha}{r}\right)^{\frac{1-\alpha}{\alpha}}$$

The parameter $A(\mu_t, \sigma_t^2)$ represents the productivity of the "representative" firm in the Commuting Zone. It is scaled by the mean skill μ_t, but adjusted downwards by the variance σ_t^2 (since $\alpha(\alpha-1) < 0$ for $\alpha \in (0,1)$). This reflects the fact that with diminishing marginal returns to skill ($\alpha < 1$), skill dispersion across workers reduces aggregate efficiency compared to a perfectly homogenous workforce.

In particular, we assume $\alpha = 0.59$, which is the estimated value from Autor, Dorn, Katz, Patterson, and Van Reenen (2020). This is the estimated value after a sharp decline starting in the early 2000s (dropping from 64% to 59%), the last decade has been characterized by relative "stagnant stability. This implies that gen-AI and human skills are complements.

6.3.2 Data and Empirical Moments

The calibration draws on the same data used for the regression analysis. Recall that the robot adoption measure is the mean of the robot-integrator measure given by $\bar{d} = 0.190$, which is our primary empirical moment for SMM.

We use the weighted average of the HHI from the Lightcast data as in the regression analysis (Choi and Marinescu (2024))given by

Table 3. Descriptive Statistics for HHI Indices (N=720)

Variable	Obs	Mean	Std. Dev.	Min	Max
HHI_L	720	0.4394	0.2074	0.0472	1.0000
HHI_{avg}	720	0.5720	0.1639	0.2307	1.0000
HHI_H	720	0.7045	0.1336	0.3816	1.0000

This gives a weighted average of 0.543. Because we consider a symmetric equilibrium in each market $n = 1/HHI \sim 2$.[18] An HHI of 0.543 exceeds the 9.25 threshold for high concentration under the Department of Justice–Federal Trade Commission (DOJ–FTC) horizontal merger guidelines. Concentration varies by occupation and city, with larger cities being less concentrated.

6.3.3 Simulated Method of Moments

The SMM estimator chooses $\theta = (\mu_t, \sigma_t, \sigma_\varepsilon)$—holding (b, r, k, n) fixed at values drawn from the literature and data—to minimize the weighted distance

$$\hat{\theta}_{\text{SMM}} = \text{argmin}_{\theta \in \Theta} \big[\hat{m}(\theta) - m_0\big]^\top \mathbf{W} \big[\hat{m}(\theta) - m_0\big], \tag{18}$$

where $\mathbf{W} = \text{diag}(1, 0.9, 1)$ weights each moment by the inverse of its squared sampling uncertainty.

Fixed parameters. Three parameters are calibrated outside the SMM to well-established empirical counterparts. The capital cost $r = 1.055$ corresponds to the average real cost of capital for equipment investment in the United States over the sample period. The outside

[18]This is higher than the average HHI=0.3157 calculated by (Azar, Marinescu, and Steinbaum (2017)) using the Burning Glass Data with the CZ×6-digit SOC level for similar years.

option $b = 3.00$ is set to match the replacement rate of unemployment insurance relative to average manufacturing wages, following the approach of Manning (2003); crucially, $b < \mathbf{E}_t[y(t,r,s)]$, so formal employment dominates the outside option in expectation, and workers participate in paid work for both pecuniary and non-pecuniary reasons. As discussed before, the number of firms $n = 2$ is set to $1/\overline{\text{HHI}}$ from Choi and Marinescu (2024), as in the regression analysis.

Skills are distributed log-normal with mean μ_t and variance σ_t^2. Both are estimated. The non-pecuniary utility shock is distributed truncated normal in $[-40, 40]$, with mean 0 and variance σ_ϵ^2, which is estimated.

Moment conditions. Three moment conditions identify the three free parameters $(\mu_t, \sigma_t, \sigma_\varepsilon)$.

Moment 1: Automation rate. In the mixed-strategy equilibrium, the probability that a job is automated os $1 - \alpha^*$. We target

$$\underbrace{1 - \alpha^*}_{\text{model}} = \underbrace{0.1903}_{\text{data}}, \tag{19}$$

where the data moment is the average, across commuting zones from Leigh and Kraft (2018) as used by Acemoglu and Restrepo (2022), probability that a robot-integrator is present.

Moment 2: Wage markdown. The equilibrium wage equation (6) implies a markdown equal to

$$\frac{\mathbf{E}_t[y(t,r,s)] - w^*}{\mathbf{E}_t[y(t,r,s)]} = \frac{m(b - w^*)}{\mathbf{E}_t[y(t,r,s)]} = 13\,\%, \tag{20}$$

the midpoint of the wage-markdown range for the United States documented by Manning (2003) (6–20 %). *Moment 3: Firm's supply* The firms' supply supply

$$\frac{1 - G(b - w^*)^n}{n} = 0.488$$

where this is the average acceptance job rate per worker derived by Faberman, Mueller, Şahin, and Topa (2022)).

Computational procedure. For each candidate θ, the objective is evaluated by: (i) computing $\mathbf{E}_t[y(t,r,s)]$; (ii) solving the equilibrium wage equation $\mathbf{E}_t[y] - w = m(b - w)$ for w^* via Brent's method, where $m(\cdot)$ is evaluated by adaptive numerical integration;

(iii) computing the profit thresholds $y(r, s; 1)$ and $y(r, s; 0)$; and (iv) if $y(r, s; 1) < y(r, k, s) < y(r, s; 0)$, recovering α^* by solving $y(r, k, \tau) = \mathbf{E}_t[\Pi(t; \alpha^*)]$ via Brent's method on the binomial mixture of firm profits. An initial grid search over $(\mu_t, \sigma_t, \sigma_\varepsilon)$ locates the basin of attraction; the final minimization uses L-BFGS-B with bounds and tolerance 10^{-8}.

The skill mean μ_t ad variance σ_t^2 are the primary drivers of $\mathbf{E}_t[y(t, r, s)]$: a higher μ_t raises the expected marginal product of labor, widening the human-job rent and pushing the equilibrium toward vacancy posting ($\alpha^* \to 1$). The skill dispersion σ_t shifts the variance of human-job output and therefore the sensitivity of the automation threshold $y(r, s; 1)$ to changes in competition intensity. The taste spread σ_ε controls the labor-supply elasticity: a larger σ_ε flattens labor supply, compressing the markdown and raising w^*. These three channels jointly identify $(\mu_t, \sigma_t, \sigma_\varepsilon)$, with the automation rate providing the binding moment and the markdown providing the cross-check.

6.3.4 Results.

Table 4 reports the estimated parameters, moment fit, and equilibrium outcomes.

Because $y(r, s; \mathbf{1}) < y(r, s, \tau) < y(r, s; \mathbf{0})$, the equilibrium is in mixed strategies (Proposition), with $\alpha^* = 0.797$ and automation rate $1 - \alpha^* = 0.203$.

A key feature of the calibrated equilibrium is the gap between human-job productivity and automation productivity: $\mathbf{E}_t[y(t, r, s)] = 24.04 \gg y(r, s, \tau) = 3.897$. Human-skill jobs are 6 times more productive than automated jobs at the calibrated capital cost, yet the equilibrium features 20 % automation. This is the productive inefficiency identified in Proposition 8(i)(b): since $\mathbf{E}_t[y] \geq y(r, s, \tau)$ and $y(r, s; \mathbf{0}) \geq y(r, s, \tau) > y(r, s; \mathbf{1})$, jobs are *inefficiently automated with probability* $1 - \alpha^*$ and workers are *inefficiently displaced* from a productive efficiency standpoint of view.

The comparison against the welfare threshold $y^*(r, s) > y(r, s, \tau) = 3.897$ further confirms that the equilibrium lies strictly below the constrained welfare benchmark (Proposition 9(iii)): 20.02 % of jobs are automated even though both productive and constrained allocative efficiency require none to be automated.

The mechanism is the rent-appropriation wedge of Corollary 1. Firms compare the *rent* from posting a vacancy, $(\mathbf{E}_t[y] - w^*)(1 - G^n(b - w^*))/n$, to the *full output* from automation, $y(r, s, \tau)$. Because wages transfer part of the human-job surplus to workers, the human-job rent falls far short of human-job output, so automation is privately profitable

Table 4. SMM Calibration: Parameters and Moment Fit

Parameter / Moment	Model	Data / Target	Source / Role
Free Parameters (estimated by SMM)			
Mean skill μ_t	6.029	—	SMM
Skill dispersion σ_t	1.000	—	SMM
Taste spread σ_ε	1.933	—	SMM
Automation Productivity A_a	4.488	—	SMM
Fixed Parameters (set outside SMM)			
Outside option b	3.000	—	Literature
Capital cost r	1.055	—	Average US Corporate Bonds Rates
Sustitution parameter α	0.59	—	Autor et al. (2020)
Number of firms n	2	$1/\overline{\text{HHI}}$	Choi and Marinescu (2024)
Moment Fit			
Frac. automated $1-\alpha^*$	0.202	0.1902	Leigh and Kraft (2018)
Wage markdown	17.7 %	13 %	Manning (2003)
Supply	0.5	0.488	Faberman et al. (2022)
Equilibrium Outcomes			
Exp. human output $\mathbf{E}_t[y(t,r,s)]$	19.309	—	Eq. result
Automation output $y(r,s,\tau)$	3.897	—	Eq. result
Equilibrium wage w^*	15.883	—	Eq. result
Total welfare W	18.457	—	Eq. result
Optimal Policy			
τ^* (Optimal Tax)	1.000	—	Eq. Result
$W(\tau^*)$ (Welfare at τ^*)	18.457	—	Eq. Result
ΔW (Welfare Gain)	0.000	—	Eq. Result
$1-\alpha^*(\tau^*)$ (Automation at τ^*)	0.203	—	Eq. Result
$w^*(\tau^*)$ (Wage at τ^*)	15.883	—	Eq. Result

Notes: SMM minimizes a three-moment criterion weighting the automation rate (weight 1), the wage markdown (weight 1), and the market tightness (weight=0.8). The loss function value at the optimum is 0.0022962994.

despite the large productivity gap. The outside option $b = 3.00 < \mathbf{E}_t[y(t,r,s)] = 19.309$ confirms that workers prefer formal employment in expectation; they accept the monopsonistic wage $w^* = 15.883$ because the equilibrium wage substantially exceeds b. Displacement is therefore a genuine welfare loss: displaced workers receive only $b = 3.00$, far below the wage they would have earned in a vacancy and the non-pecuniary benefit lost.

Figure 2 illustrates the equilibrium profit curves, the best-response correspondence, and the key efficiency thresholds.

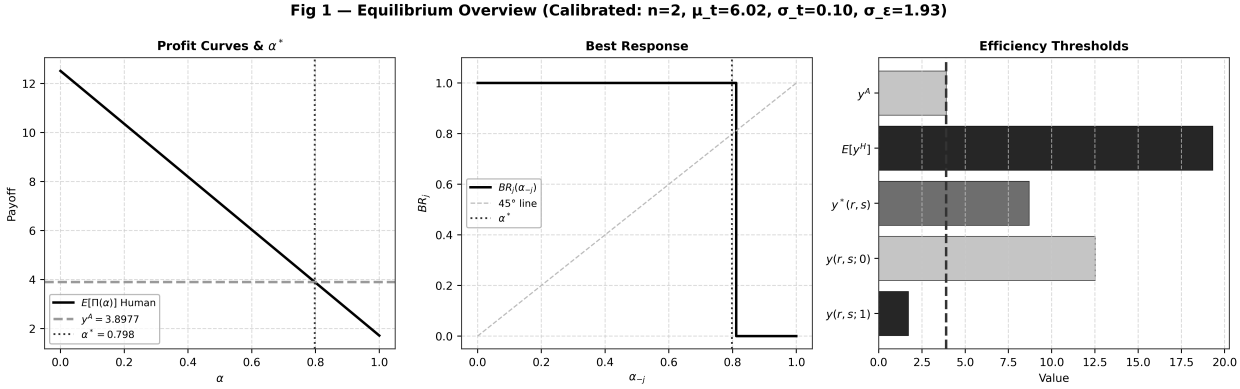


Fig. 2. Equilibrium overview at calibrated values ($n = 2$, $r = 1.055$, $b = 3.00$). *Left:* Expected profit $\mathbf{E}_t[\Pi(\alpha)]$ and automation payoff $y(r,s,\tau)$ as functions of α. The equilibrium $\alpha^* = 0.797$ (Tertical dotted line) equates the two. *Center:* Best-response correspondence $BR_j(\alpha_{-j})$ (step function) and the 45° line; their intersection gives α^*. *Right:* Key efficiency thresholds; the shaded region marks the range of $y(r,s,\tau)$ values for which the equilibrium is inefficiently automated.

6.3.5 Counterfactual Policies

We report two counterfactual experiments, each shocking one policy instrument while holding all other parameters at calibrated values. The purpose is to quantify the welfare and distributional effects of capital taxation, competition policy, and their interaction.

Competition Policy (n). Antitrust enforcement or market-entry subsidies that reduce vacancy-market concentration can be modeled as an increase in n, holding (r, b, τ, n) fixed.

Table 5 reports the equilibrium for $n \in \{1, \ldots, 11\}$, holding all other parameters at calibrated values.

Table 5. Competition Policy: Varying Number of Firms n

n	Regime	α^*	$1-\alpha^*$	w^*	Markdown	W	ΔW	$y^*(r,s)$
1	All Human	1.000	0.000	5.858	69.7%	18.223	+0.000	15.223
2	Mixed	0.798	0.202	15.883	17.7%	18.457	+0.233	8.700
3	Mixed	0.524	0.476	17.025	11.8%	17.970	-0.254	5.982
4	Mixed	0.385	0.615	17.431	9.7%	19.633	+1.409	4.575
5	Mixed	0.303	0.697	17.647	8.6%	22.206	+3.983	3.711
6	Mixed	0.250	0.750	17.784	7.9%	25.230	+7.007	3.126
7	Mixed	0.212	0.788	17.880	7.4%	28.508	+10.285	2.703
8	Mixed	0.185	0.815	17.951	7.0%	31.943	+13.720	2.383
9	Mixed	0.163	0.837	18.007	6.7%	35.482	+17.259	2.131
10	Mixed	0.146	0.854	18.053	6.5%	39.093	+20.870	1.928
11	Mixed	0.133	0.867	18.091	6.3%	42.757	+24.534	1.761

Notes: Parameters at calibrated values ($\mu_t = 6.025, \sigma_t = 0.100, \sigma_\varepsilon = 1.933, b = 3.00, r = 1.05, k = 4.49$). ΔW represents the change in total welfare relative to the monopsony case ($n = 1$). $y^*(r,s)$ denotes the constrained welfare threshold.

Proposition 4 establishes that, as n rises, wages increase toward the competitive benchmark $\mathbf{E}_t[y(t,r,s)]$ and the markdown falls. Because higher competition erodes the rent from human-skill jobs relative to the fixed automation payoff $y(r,s,\tau)$, the automation rate $1-\alpha^*$ rises (α^* falls) from zero to 86%. This is the central tension of the model: more competitive labor markets benefit employed workers through higher wages and better non-pecuniary benefits opportunities, but simultaneously increase firms' incentive to automate and thus raise the risk of displacement, as formalized in Corollary 1.

The results reveal a *non-monotone* welfare response to competition. This is the other central quantitative finding of the paper. Welfare rises from $n = 1$ to $n = 2$ ($\Delta W = +0.233$) as the transition from monopsony to duopsony substantially raises wages. It then *falls* for $n = 3$, to then sharply increase to 42.757 when there are 11 firms. The mechanism is the displacement channel formalized in Corollary 1.

As competition intensifies beyond duopsony, wages rise and the markdown shrinks, which compresses human-job rents. Firms respond by automating more jobs: the automation rate rises from 0 % at $n = 1$ to 69 % at $n = 5$ and 57 % at $n = 11$. Displaced workers lose the equilibrium wage $w^* \approx 18$ and receive only the outside option $b = 3.00$, a welfare loss of roughly 15 per displaced worker. For moderate competition levels, this displacement cost outweighs the wage gain enjoyed by the fraction of workers who retain their jobs.

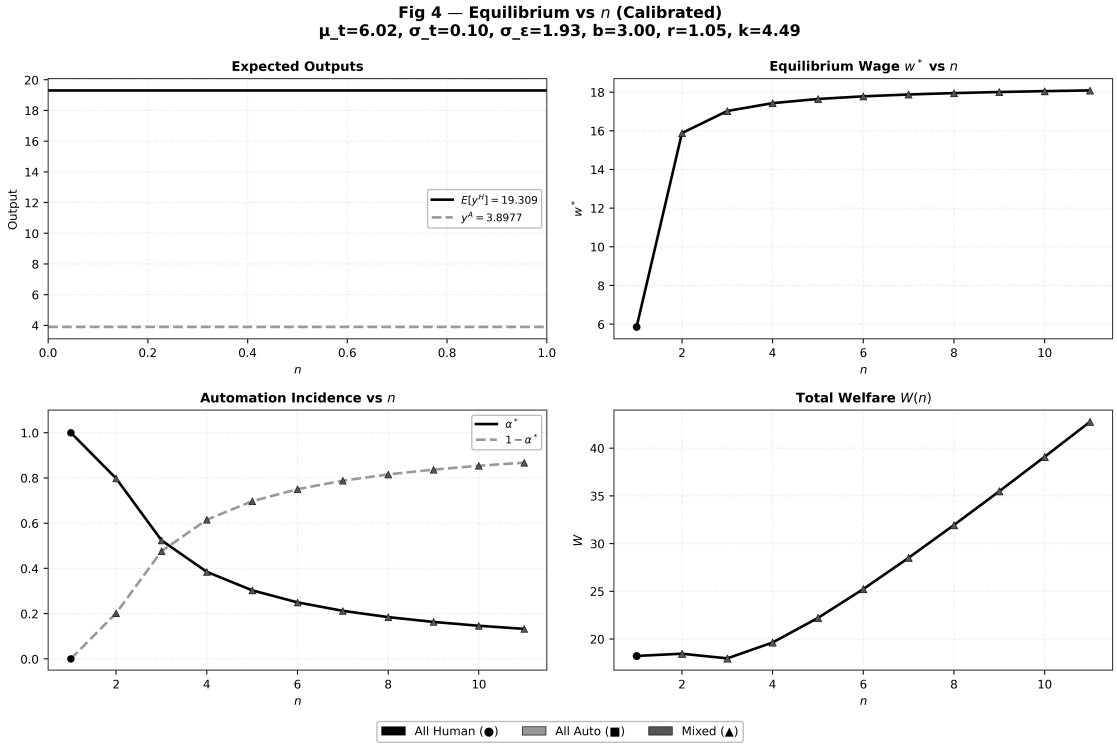


Fig. 3. The Impact of Competition Intensity

This non-monotonicity has a direct policy implication: antitrust enforcement that moves a concentrated labor market toward moderate competition may *reduce* welfare if the market already features some competition (say, $n = 2$), because the additional entry induces automation that displaces workers at a rate outpacing wage gains. Competition policy is welfare-improving only when it pushes the market away from the extreme of monopsony.

All this is clearly seen on the following plot.

This non-monotonicity is the key policy implication of the calibration: reducing labor-market concentration is *not* unambiguously welfare-improving when the labor market is already moderately competitive. Competition policy that moves the market from duopsony toward an intermediate level of competition may reduce welfare by inducing excessive automation, even as it raises wages for the workers who retain their jobs.

The welfare threshold $y^*(r, s)$ declines with n, from 15.23 at $n = 1$ to 1.76 at $n = 11$. Because the planner's threshold for justifying automation shrinks as competition intensifies and human-job rents erode, the region of efficient automation expands. For all $n \geq 5$, $y(r, s, \tau) = 3.897 < y^*(r, s)$, so the calibrated equilibrium lies below the constrained welfare benchmark at every competition level above 5.

Technological capital tax policy Taxing technological capital has been much debated. This counterfactual reveals the benefit of doing so.

A tax $\tau \geq 1$ on technology capital raises the effective marginal cost of both automation and gen-AI investment from r to τr. Because both the automation payoff $y(r, s, \tau)$ and the human-job rent $\mathbf{E}_t[y(t, r, s)]$ decrease in τ, the effect on the automation rate $1 - \alpha^*$ is a priori ambiguous: it depends on whether the tax reduces automation profits faster than it reduces human-job rents. Equation (10) provides the necessary and sufficient condition for the tax to reduce automation.

Table 6 shows the equilibrium for tax rates ranging from $\tau = 1$ (no tax) to $\tau = 2$ (100 % surcharge). The tax delivers a clear *allocation cost* rather than a benefit. A tax on automation capital simultaneously reduces the automation payoff $y(r, s, \tau)$ *and* the human-job rent $\mathbf{E}_t[y(t, r, s)]$ by depressing gen-AI investment. At the calibrated parameters, the human-job rent falls more slowly than the automation payoff under the tax ($y_\tau(r, s, \tau; \alpha) < y_\tau(r, k, s)$ for all τ), so a higher τ makes automation *relatively less* attractive. Consequently, the automation rate falls monotonically from 20 % at $\tau = 1$ to 19 at $\tau = 2$, which is small drop. Welfare, however, falls drastically from 18.5 to 15 as the tax reduces wages and both human and automation productivity.

Table 6. Technological Capital Tax Polciy (Optimal $\tau^* = 1.000$)

Tax $\tau-1$	Regime	α^*	$1-\alpha^*$	w^*	$\mathbf{E}_t[y(t,r,s)]$	y(r,s,τ)	W	ΔW
0%	Mixed $\star$	0.798	0.202	15.883	19.3089	3.8977	18.4567	+0.0000
2%	Mixed $\star$	0.798	0.202	15.555	18.9804	3.8314	18.1719	-0.2848
5%	Mixed $\star$	0.799	0.201	15.240	18.6652	3.7678	17.8988	-0.5579
10%	Mixed $\star$	0.800	0.200	14.646	18.0715	3.6479	17.3849	-1.0718
12%	Mixed	0.800	0.200	14.366	17.7914	3.5914	17.1428	-1.3139
25%	Mixed	0.803	0.197	13.110	16.5354	3.3379	16.0591	-2.3976
27%	Mixed	0.804	0.196	12.884	16.3094	3.2922	15.8646	-2.5921
30%	Mixed	0.804	0.196	12.665	16.0908	3.2481	15.6766	-2.7801
32%	Mixed	0.805	0.195	12.454	15.8792	3.2054	15.4949	-2.9618
35%	Mixed	0.806	0.194	12.249	15.6743	3.1640	15.3190	-3.1377
38%	Mixed	0.806	0.194	12.050	15.4757	3.1239	15.1487	-3.3080
40%	Mixed	0.807	0.193	11.858	15.2831	3.0851	14.9838	-3.4730

Notes: $\tau = 1$: baseline (no tax). Optimal tax $\tau^* = 1.000$ (**starred row**) maximizes total welfare W. ΔW: change in welfare relative to $\tau = 1$ baseline.

The welfare-maximizing policy is therefore no technology capital tax, $\tau^* = 1$.

This finding stands in contrast to Guerreiro, Rebelo, and Teles (2022b), who find a

positive optimal robot tax in a dynamic general-equilibrium model. The key difference is the gen-AI channel: in our model, taxing capital simultaneously discourages automation *and* gen-AI investment in human-skill jobs. At the calibrated parameters, the second effect dominates, so the net welfare impact of a tax is negative.

Figure 4 illustrates the equilibrium, welfare, and distributional responses to the tax.

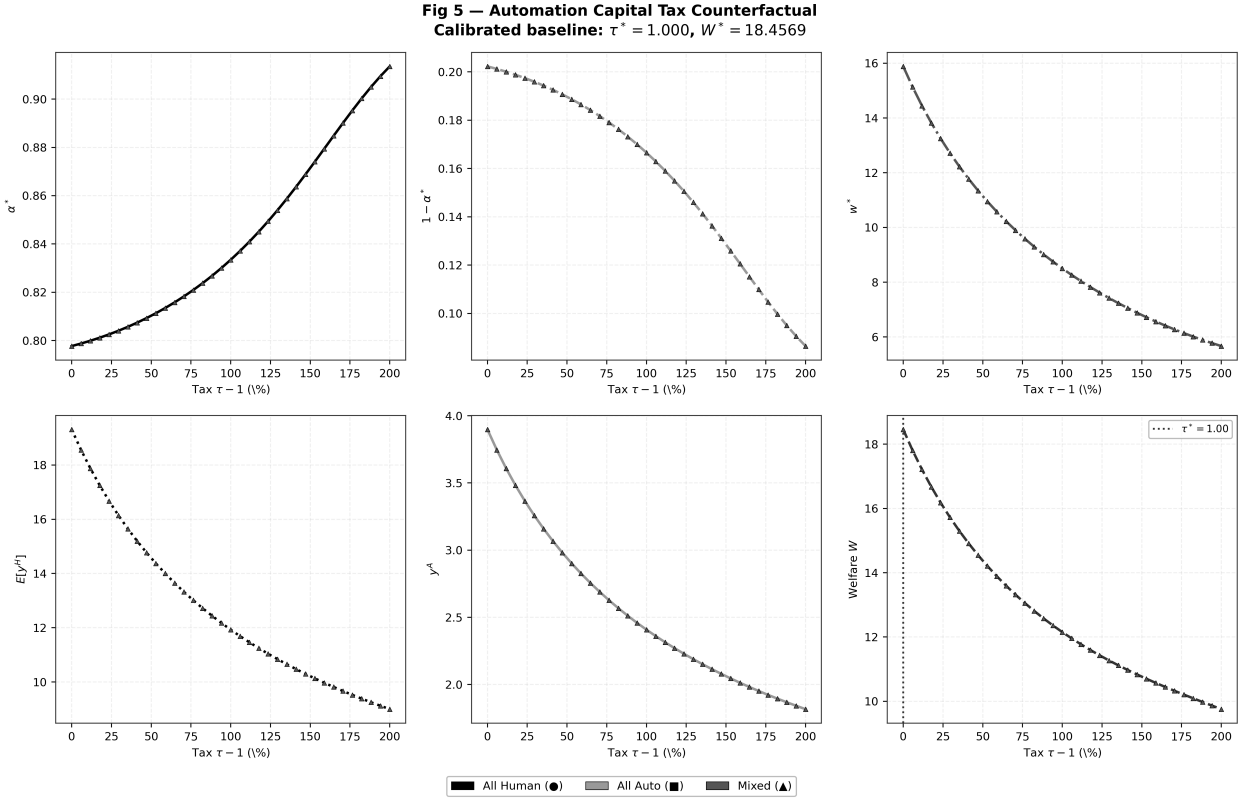


Fig. 4. Automation capital tax counterfactual ($n = 3$, $r = 1.055$, $b = 3.00$, $k = 1.50$). As τ rises, $1 - \alpha^*$ rises monotonically (upper left), wages and human-job productivity fall (lower left and upper right), and welfare falls monotonically (lower right). The welfare peak at $\tau^* = 1$.

The uniform automation-tax or robot-tax proposals, including our own single-industry optimal rate in Table 8, are correct as a first approximation only insofar as the economy can be treated as a single industry drawing on a perfectly elastic capital supply. Once technological capital is scarce and shared across industries with heterogeneous automation and gen-AI intensities, the planner's problem is intrinsically one of differentiated, not uniform, taxation.

Optimal Tax by Competition Level. Table 7 reports the welfare-maximizing tax $\tau^*(n) = \text{argmax}_\tau W(n, \tau)$ for each $n \in \{1, \ldots, 11\}$. Here, we allow for both taxes and subsidies.

Table 7. Optimal Automation Tax τ^* and competition Policy n

n	$\tau^*(n)$	Tax rate	Regime	$W^*(n)$	ΔW	$1-\alpha^*$ at τ^*
1	1.000	-0.0%	All Human	18.2235	+0.0002	0.000
2	1.000	-0.0%	Mixed	18.4569	+0.0002	0.202
3	0.100	-90.0%	All Auto	60.9204	+42.9508	1.000
4	0.100	-90.0%	All Auto	80.2271	+60.5945	1.000
5	0.100	-90.0%	All Auto	99.5339	+77.3278	1.000
6	0.100	-90.0%	All Auto	118.8407	+93.6106	1.000
7	0.100	-90.0%	All Auto	138.1475	+109.6394	1.000
8	0.100	-90.0%	All Auto	157.4543	+125.5112	1.000
9	0.100	-90.0%	All Auto	176.7610	+141.2790	1.000
10	0.100	-90.0%	All Auto	196.0678	+156.9746	1.000
11	0.100	-90.0%	All Auto	215.3746	+172.6178	1.000

Notes: Optimal capital tax $\tau^*(n)$ that maximizes total welfare W for each number of firms n. ΔW = welfare gain relative to the no-tax baseline at the same n.

For $n \leq 2$, the optimal policy is not tax at all. According to the calibrated model, neither taxing nor subsidizing technological capital is necessarily. However, when competition intensity rises ($n \geq 3$), the optimal is to uniformly *subsidy* technology capital at a high rate of 90%. The subsidy raises $\mathbf{E}_t[y(t, r, s)]$ more slowly than it raises automation productivity. Because the equilibrium switched to all automation for $n \geq 3$, there is no cost of subsidizing technological capital because it does not create a displacement effect. Welfare rises since we have not considered the shadow cost of funds.

Figure 5 illustrates the joint welfare and automation-rate surface across (n, τ).

The result that the optimal policy is a *subsidy* rather than a tax across all verified competition levels beyond $n = 2$ reinforces the main finding from the tax counterfactual: in this calibration, the gen-AI channel dominates. Taxing capital depresses gen-AI investment in human-skill jobs more than discouraging automation, worsening the allocation. The optimal response is instead to subsidize technological capital investment, exploiting the complementarity between gen-AI and human skills to expand the set of jobs that firms find it worthwhile to fill with human workers

Three broad conclusions emerge from this analysis: first, the optimal automation policy is a capital subsidy, not a tax; second, competition policy has non-monotone welfare; and the model supports the illustrative evidence provided.

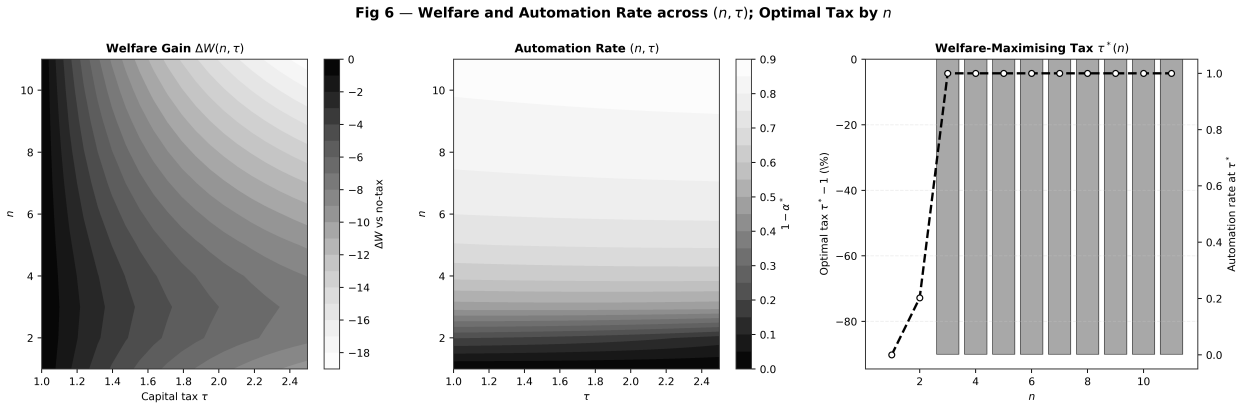


Fig. 5. Welfare and automation rate across (n, τ) (*left* and *centre* panels) and optimal tax/subsidy $\tau^*(n)$ (*right panel*, bars, left axis) with the corresponding automation rate at τ^* (dashed line, right axis).

The calibration reproduces the stylized fact: the HHI is negatively correlated with robot-integrator presence across CZs. The competition policy results show that this pattern arises endogenously from the model's equilibrium, where more concentrated markets (lower n, higher HHI) sustain higher wages through market power, which in turn makes human-skill jobs more profitable relative to automation. Despite the fact that increasing competition intensity increases work displacement, welfare rises because the wage plus the non-pecuniary benefits gain those who remain employed, together with the automation productivity, more than compensate for the loss of productivity from human-skill jobs.

The higher the outside option, the stronger the welfare gain from displacement. However, a higher b raises wages, but less than the increase in b because the pass-through is lower than 1. This implies that the rent from a human-skill job decreases, which makes automation more attractive. However, as b rises, displaced workers receive a higher payoff from losing their jobs. Thus, the welfare effect is a priori ambiguous.

7 Conclusions

This paper studies the taxation of automation capital in a labor market where firms exercise monopsony power and can adjust technology on two margins: full automation, which substitutes for human labor, and generative AI, which augments it. The central finding is that these two margins interact in a way that overturns the standard case for a

positive robot tax. A tax on technological capital does not act selectively on automation; it raises the cost of the same capital when deployed to augment human-skill jobs, and at empirically disciplined parameter values this second effect dominates. The welfare-maximizing policy in our calibration is a subsidy to technological capital, not a tax, despite an equilibrium in which automation is simultaneously excessive relative to both productive and constrained-efficient benchmarks.

This is not a claim that automation policy is unnecessary. The model identifies a specific and quantitatively large distortion—firms automate based on the rent they can appropriate from a job, not on that job's output, so automation is privately attractive even when far less productive than the human-staffed alternative it displaces. In our calibration this wedge is large: human-skill jobs are roughly six times more productive than automated ones, yet the equilibrium automates 20% of jobs, all of it inefficient by both criteria we consider. The implication for policy is not that this distortion should be left alone, but that the natural-seeming instrument—a capital tax aimed at automating technology—is poorly suited to correcting it whenever the tax base cannot distinguish automating capital from augmenting capital. Since generative AI and other modern computing technologies are, almost by construction, deployable on either margin depending on how a firm designs a job, this is not a knife-edge concern specific to our functional-form choices; it is a structural feature of any policy that taxes capital rather than the task the capital performs. A tax instrument that could be conditioned on the specific margin—for instance, a tax on job designs that eliminate the human input entirely, as opposed to a tax on computing or robotics capital more broadly—would not face this trade-off, but such an instrument is far harder to implement and enforce than the capital taxes currently under policy discussion.

Competition policy compounds rather than substitutes for this targeting problem. Because more competitive labor markets erode the rents that make firms prefer staffing a vacancy to automating it, antitrust enforcement that reduces labor-market concentration raises the equilibrium automation rate and, in our calibration, produces a non-monotone welfare response: moving away from extreme monopsony raises welfare, but intermediate increases in competition can lower it by inducing automation and worker displacement at a rate that outpaces the wage gains accruing to workers who keep their jobs. Competition policy and automation-capital policy are therefore not independent instruments;

the welfare-maximizing capital tax or subsidy depends on the level of labor-market concentration a regulator is simultaneously pursuing, and we characterize this joint optimum across competition levels. At every competition level we consider beyond the most concentrated, the optimal policy is a subsidy, not a tax, reflecting the same gen-AI channel that drives our main result.

Two limitations bound how far these conclusions should be pushed. First, our welfare and policy results are derived and calibrated under a symmetric equilibrium; we do not characterize whether asymmetric equilibria—in which otherwise-identical firms make different automation decisions off the mixed-strategy path—exist or would alter the sign of the tax result, and we leave this to future work. Second, our calibration relies on a specific Cobb-Douglas production structure and a second-order approximation to aggregate the skill distribution; while the qualitative mechanism does not depend on these choices, the quantitative magnitude of the optimal subsidy does, and a fuller robustness exercise across functional forms is warranted before treating the specific subsidy rate as a policy target rather than a qualitative result.

Subject to these caveats, the paper's central policy message is that the case for taxing automation is considerably more fragile than a single-margin model suggests, once the same capital stock can be redeployed to complement, rather than replace, human labor. Any automation tax proposal should be evaluated not only on how much it discourages the substitution margin it targets, but on how much collateral discouragement it imposes on the complementary margin it does not.

Robot taxes are a specific form of capital taxation that has been discussed to slow down the excessively labor-saving aspects of AI and automation technologies. Some optimal tax models propose taxes on specific capital inputs, robots, to achieve this goal. However, it remains unclear how a robot tax should be implemented in real life, as it involves drawing dividing lines between similar technologies. Moreover, to the extent that AI and automation increase the overall productivity of the economy, there is a general growth argument against implementing such taxes. There is an important interplay here between taxation, competition policy, and how such tax and policy are coordinated across countries. The conclusion is therefore that robot taxes do not seem to be a plausible part of a well-designed tax system in the near future.

References

Acemoglu, D. and Restrepo, P. Robots and Jobs: Evidence from US Labor Markets. *Journal of Political Economy*, 128(6):2188–2244, 2020. doi: 10.1086/705716.

Acemoglu, D. and Restrepo, P. Tasks, Automation, and the Rise in US Wage Inequality. NBER Working Papers 28920, National Bureau of Economic Research, Inc, June 2021.

Acemoglu, D. and Restrepo, P. Demographics and automation. *The Review of Economic Studies*, 89(1):1–44, 2022.

Acemoglu, D. and Restrepo, P. Automation and rent dissipation: Implications for wages, inequality, and productivity. Technical report, National Bureau of Economic Research, 2024.

Aghion, P., Antonin, C., Bunel, S., and Jaravel, X. The effects of automation on labor demand. *Robots and AI: A new economic era*, pages 15–39, 2022.

Armangué-Jubert, T., Guner, N., and Ruggieri, A. Labor market power and development. *American Economic Review*, 2024.

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

Autor, D., Chin, C., Salomons, A. M., and Seegmiller, B. New Frontiers: The Origins and Content of New Work, 1940–2018. NBER Working Papers 30389, National Bureau of Economic Research, Inc, August 2022.

Autor, D. H. and Dorn, D. The Growth of Low-Skill Service Jobs and the Polarization of the US Labor Market. *American Economic Review*, 103(5):1553–97, August 2013.

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

Azar, J., Marinescu, I., and Steinbaum, M. I. Labor market concentration. Working Paper 24147, National Bureau of Economic Research, December 2017.

Azar, J., Berry, S., and Marinescu, I. E. Estimating labor market power. *Available at SSRN 3456277*, 2019.

Balmaceda, F. Market Power and Occupational Mismatch. *Journal of Development Economics*, R&R, 2025.

Barro, R. J. and Lee, J.-W. A new data set of educational attainment in the world, 1950–2010. *Journal of Development Economics*, 104:184–198, 2013.

Bessen, J., Goos, M., Salomons, A., and Bergevan den , W. Firm-level automation: Evidence from the netherlands. *AEA Papers and Proceedings*, 110:389–93, May 2020. doi: 10.1257/pandp.20201004.

Brynjolfsson, E., Li, D., and Raymond, L. Generative ai at work. *The Quarterly Journal of Economics*, 140(2):889–942, 05 2025. ISSN 0033-5533. doi: 10.1093/qje/qjae044.

Choi, H. and Marinescu, I. Data for labor market concentration using lightcast (formerly burning glass technologies). *Data in Brief*, 55:110647, August 2024. ISSN 23523409. doi: 10.1016/j.dib.2024.110647.

Costinot, A. and Werning, I. Robots, trade, and luddism: A sufficient statistic approach to optimal technology regulation. *[Journal Name]*, 2023.

De Loecker, J. and Warzynski, F. Markups and firm-level export status. *American Economic Review*, 102(6):2437–71, May 2012. doi: 10.1257/aer.102.6.2437.

De Loecker, J., Eeckhout, J., and Mongey, S. Quantifying Market Power and Business Dynamism in the Macroeconomy. NBER Working Papers 28761, National Bureau of Economic Research, Inc, May 2021.

De Souza, G. and Li, H. Robots, Tools, and Jobs: Evidence from Brazilian Labor Markets. CESifo Working Paper Series 10813, CESifo, 2023.

Dixon, J., Hong, B., and Wu, L. The robot revolution: Managerial and employment consequences for firms. *Management Science*, 67(9):5586–5605, 2021. doi: 10.1287/mnsc.2020.3812.

Eckert, F., Fort, T. C., Schott, P. K., and Yang, N. J. Imputing missing values in the US Census Bureau's county business patterns. NBER Working Paper 26632, National Bureau of Economic Research, 2021.

Eggleston, K., Lee, Y. S., and Iizuka, T. Robots and labor in the service sector: Evidence from nursing homes. NBER Working Papers 28322, National Bureau of Economic Research, Inc, 2021.

Eslava, M., García-Marín, A., and Messina, J. Inequality and Market Power in Latin America and the Caribbean. IDB Publications (Working Papers) 13192, Inter-American Development Bank, October 2023.

Faberman, R. J., Mueller, A. I., Şahin, A., and Topa, G. Job search behavior among the employed and unemployed. *Econometrica*, 90(4):1487–1522, 2022.

Feenstra, R. C., Inklaar, R., and Timmer, M. P. The next generation of the Penn World Table. *American Economic Review*, 105(10):3150–3182, 2015.

Gabaix, X., Laibson, D., Li, D., Li, H., Resnick, S., and Vriesde , C. The impact of competition on prices with numerous firms. *Journal of Economic Theory*, 165(C):1–24, 2016.

Guerreiro, J., Rebelo, S., and Teles, P. Should robots be taxed? *The Review of Economic Studies*, 89(1):279–311, None 2022a. doi: None.

Guerreiro, J., Rebelo, S., and Teles, P. Should robots be taxed? *The Review of Economic Studies*, 89(1):279–311, 2022b.

Hall, R. E., Blanchard, O. J., and Hubbard, R. G. Market structure and macroeconomic fluctuations. *Brookings papers on economic activity*, 1986(2):285–338, 1986.

Handwerker, E. W. and Dey, M. Some facts about concentrated labor markets in the United States. BLS Working Paper 550, U.S. Bureau of Labor Statistics, 2022.

Handwerker, E. W. and Dey, M. Some facts about concentrated labor markets in the United States. *Industrial Relations: A Journal of Economy and Society*, 63(2):132–151, April 2024.

Hershbein, B. and Kahn, L. B. Do recessions accelerate routine-biased technological change? evidence from vacancy postings. *American Economic Review*, 108(7):1737–1772, 2018.

Kang, D., Lee, J. H., and Quach, S. The welfare effects of a robot tax: Evidence from a tax credit for automation technologies in korea. October 2024. Available at SSRN: `https://ssrn.com/abstract=5005128` or `http://dx.doi.org/10.2139/ssrn.5005128`.

Kina, Ø., Kindsgrab, B., and Krueger, D. Optimal capital and labor taxation with capital-skill complementarity. *Journal of Monetary Economics*, 139:84–103, 2023.

Koch, M., Manuylov, I., and Smolka, M. Robots and Firms. *Economic Journal*, 131(638): 2553–2584, 2021.

Lamadon, T., Mogstad, M., and Setzler, B. Imperfect competition, compensating differentials, and rent sharing in the us labor market. *American Economic Review*, 112(1):169–212, January 2022. doi: 10.1257/aer.20190790.

Leigh, N. G. and Kraft, B. R. Emerging robotic regions in the United States: Insights for regional economic evolution. *Regional Studies*, 52(6):804–815, 2018.

Loebbing, J. Tax policy and human capital: The role of endogenous skill-biased technical change. *Unpublished Manuscript*, 2022.

Maestas, N., Mullen, K. J., Powell, D., Wachtervon , T., and Wenger, J. B. The value of working conditions in the united states and implications for the structure of wages. Working Paper 25204, National Bureau of Economic Research, October 2018.

Manning, A. *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press, 2003.

Manning, A. Monopsony in labor markets: A review. *ILR Review*, 74(1):3–26, 2021. doi: 10.1177/0019793920922499.

Manson, S., Schroeder, J., Van Riper, D., and Ruggles, S. IPUMS national historical geographic information system: Version 12.0 [database]. Minneapolis: University of Minnesota, 2017.

Marx, K. *Das Capital Bd. 3*. Verlag von Otto Meissner, 1894.

Mas, A. and Pallais, A. Valuing alternative work arrangements. *American Economic Review*, 107(12):3722–59, December 2017. doi: 10.1257/aer.20161500.

Milgrom, P. and Roberts, J. Rationalizability, learning, and equilibrium in games with strategic complementarities. *Econometrica*, 58(6):1255–77, November 1990.

Perloff, J. M. and Salop, S. C. Equilibrium with Product Differentiation. *Review of Economic Studies*, 52(1):107–120, 1985.

Rosen, S. The theory of equalizing differences. In Ashenfelter, O. and Layard, R., editors, *Handbook of Labor Economics*, volume 1 of *Handbook of Labor Economics*, chapter 12, pages 641–692. Elsevier, 1987.

Schulz, K., Tsyvinski, A., and Werquin, N. Reforming us income taxes in the era of frictionless finance and technological change. *Brookings Papers on Economic Activity*, pages 303–368, 2022.

Sorkin, I. Ranking firms using revealed preference. *The Quarterly Journal of Economics*, 133 (3):1331–1393, 2018.

Sullivan, P. and To, T. Search and nonwage job characteristics. *Journal of Human Resources*, 49(2), 2014.

Thuemmel, U. Optimal taxation of robots. *[Journal Name]*, 2023.

United Nations Industrial Development Organization. INDSTAT industrial statistics database. UNIDO, Vienna.

United Nations Statistics Division. UN Comtrade international trade statistics database. United Nations.

U.S. Bureau of Labor Statistics. Occupational employment and wage statistics (OEWS). U.S. Bureau of Labor Statistics.

Webb, M. The impact of artificial intelligence on the labor market. *Available at SSRN: https://ssrn.com/abstract=3482150*, pages 15–39, 2020.

Webber, D. Firm market power and the earnings distribution. *Labour Economics*, 35(C): 123–134, 2015.

World Bank. Enterprise surveys. World Bank, Washington, DC.

Yeh, C., Macaluso, C., and Hershbein, B. Monopsony in the us labor market. *American Economic Review*, 112(7):2099–2138, July 2022. doi: 10.1257/aer.20200025.

Zhou, J. Competitive bundling. *Econometrica*, 85(1):145–172, 2017. doi: 10.3982/ECTA14251.

A Main Proofs

Proof of Proposition 1. Because G^k are identically distributed, for all $j \in \mathcal{J}(d)$

$$P^j(w) = \int_{b-w^j}^{\bar{\epsilon}} \prod_{k \in \mathcal{J}(d^{-j})} G^k\Big(w^j + \epsilon^j - w^k\Big) dG^j(\epsilon^j),$$

and for all for all $j \in \mathcal{J}(d)$,

$$P^j_j(w) = \int_{b-w^j}^{\bar{\epsilon}} \sum_{h \in \mathcal{J}(d^{-j})} \nu_g\Big(w^j + \epsilon^j - w^h\Big) \prod_{k \in \mathcal{J}(d^{-j})} G\Big(w^j + \epsilon^j - w^k\Big) dG(\epsilon^j) + g(b - w^j) \prod_{k \in \mathcal{J}(d^{-j})} G(b - w^k),$$

where $\nu_g(\cdot) \equiv g(\cdot)/G(\cdot)$,

$$P^j_h(w) = -\int_{b-w^j}^{\bar{\epsilon}} \nu_g\Big(w^j + \epsilon^j - w^h\Big) \prod_{k \in \mathcal{J}(d^{-j})} G\Big(w^j + \epsilon^j - w^k\Big) dG(\epsilon^j) < 0.$$

$P^j(w)$ is strictly increasing in w^j and is strictly decreasing in $w^{j'}$.

$P^j(w)$ is log-concave in w^j if and only if the following holds

$$\frac{1}{P^j(w)} P^j_{j,j}(w) - \frac{1}{(P^j(w))^2} (P^j_j(w))^2 \leq 0$$

This holds because the multiplication of log-concave functions is log-concave and $G()$ is log-concave in w^j.

Also observe that for all $j \in \{1, \ldots, n\}$, we have that for all $j \in \mathcal{J}(d)$

$$P^j_{jj}(w) = \int_{b-w_j}^{\bar{\epsilon}} \left(\sum_{h \in \mathcal{J}(d^{-j})} \nu'_g\Big(w_j + \epsilon_j - w_h\Big) + \left[\sum_{h \in \mathcal{J}(d^{-j})} u_g\Big(w_j + \epsilon_j - w_h\Big) \right]^2 \right) \times \prod_{k \in \mathcal{J}(d^{-j})} G\Big(w_j + \epsilon_j - w_k\Big) dG(\epsilon_j) - h'(b - w_j) \prod_{k \in \mathcal{J}(d^{-j})} G(b - w_k).$$

Observe that for all $j, h \in \mathcal{J}(d)$, log-supermodularity implies

$$\frac{1}{P^j(w)} P^j_{j,j'}(w) - \frac{1}{(P^j(w))^2} P^j_j(w) P^j_{j'}(w) \geq 0.$$

Observe that

$$P^j_{jh}(w) = \int_{b-w_j}^{\bar{\epsilon}} \Big(-\nu'_g\Big(w_j+\epsilon_j-w_h\Big) - \nu_g\Big(w_j+\epsilon_j-w_h\Big) \sum_{h\in\mathcal{J}(d^{-j})} \nu_g\Big(w_j+\epsilon_j-w_h\Big)\Big)\times \prod_{k\in\mathcal{J}(d^{-j})} G\Big(w_j+\epsilon_j-w_k\Big)dG(\epsilon_j) + \nu_g(b-w^h) \prod_{k\in\mathcal{J}(d^{-j})} G\big(-w^k\big)g(b-w^j)$$

since $\nu'_h < 0$ because $g(\cdot)$ is log-concave.

□

To save on notation burden, in this appendix, we will denote $\mathbf{E}_t[y^j(t,r,s)]$ by y^j.

Proof of Lemma 1. Existence follows from Weiertress' Theorem and uniqueness from the log-concavity of the profit functions: the best response is the solution to

$$P^j_j(w) - \frac{1}{y^j - w^j} = 0. \tag{A1}$$

It follows from this that profits are log-concave whenever

$$\frac{P^j(w)P^j_{jj}(w) - (P^j_j(w))^2}{(P^j(w))^2} - \frac{1}{(y^j-w^j)^2} \leq 0, \tag{A2}$$

where the inequality follows from the fact that $P^j(w)$ is log-concave in w □

Proof of Proposition 3. The proof of this result follows closely Zhou (2017).

Recall that the first-order condition is given by

$$y - w = \frac{1}{n}\frac{1-G(b-w)^n}{G(b-w)^{n-1}g(b-w) + \int_{b-w}^{\bar{\epsilon}} g(\epsilon)dG(\epsilon)^{n-1}}. \tag{A3}$$

Let's define the CDF of the second-highest order statistics by

$$G_{n-1}(b-w) = G(b-w)^n + nG(b-w)^{n-1}(1-G(b-w))$$

Observe that at $w = y$, the right-hand side equation (A3) is greater than the left-hand side since it is strictly positive. Let $\lambda_g(\cdot) \equiv g(\cdot)/(1-G(\cdot))$ be the hazard rate. Let's define $w^m(y)$ as the wage when there is only one firm. It is easy to check that this is the unique

(due to log-concavity of $g(\cdot)$) solution to the following equation

$$y - w^m = \frac{1}{\lambda_g(y - w^m)}.$$

Observe that at $w = w^m$, the LHS is larger than the RHS. To see this notice that

$$\begin{aligned} m(b-w) =& \frac{1 - G(b-w)^n}{nG(b-w)^{n-1}g(b-w) + \int_{b-w}^{\bar{\epsilon}} \lambda y_g(\epsilon) dG_{n-1}(\epsilon)} \\ <& \frac{1 - G(b-w)^n}{nG(b-w)^{n-1}g(b-w) + \lambda_g(b-w)(1 - G_{n-1}(b-w))} \\ =& \frac{1}{\lambda_g(b-w)}, \end{aligned}$$

where the inequality follows from the fact λ_g is increasing. One deduces from this and the definition of $w(r,s)$, that the left-hand side is greater than the right-hand side at $w(r,s) = w^m(y,b)$. Thus, the first-order condition in equation (A3) has a solution $w \in [w^m(y,b), y]$.

Because the left-hand side of the first-order condition in equation (A3) falls with w at a rate equal to -1, the solution is unique if $m'(y-w) < 0$. To show that this is the case, observe that

$$m'(b-w) = -n\frac{G(b-w)^{n-1}g(b-w)}{1 - G(b-w)^n} m(b-w)\left(1 + m(b-w)\frac{g'(b-w)}{g(b-w)}\right).$$

Thus, $m'(b-w) < 0$ if and only if

$$ng(b-w)\Big(G(b-w)^{n-1}g(b-w) + \int_{b-w}^{\bar{\epsilon}} g(\epsilon) dG(\epsilon)^{n-1}\Big) + (1 - G(b-w)^n)h'(b-w) > 0.$$

Using the fact that log-concavity implies that $(1-G)h' + g^2 > 0$, one deduces that the inequality above holds if

$$n\int_{-w}^{\bar{\epsilon}} g(\epsilon) dG(\epsilon)^{n-1} > (1 - G(b-w)^n)\lambda_g(b-w) - nG(b-w)^{n-1}g(b-w).$$

Using the definition of G_{n-1}, we re-write this as follows

$$\int_{-w}^{\bar{\epsilon}} \lambda_g(\epsilon) dG(\epsilon)_{n-1} > (1 - G_{n-1}(b-w))\lambda_g(b-w).$$

The inequality holds because the hazard rate is increasing.

□

Proof of Proposition 4. It follows from the first-order condition in equation (6) and uniqueness that the equilibrium wage increases with $x \in (y, b)$ if and only if

$$\frac{\partial w}{\partial x}(1 - m'(b-w)) = \frac{\partial y}{\partial x} - m'(b-w)\frac{\partial b}{\partial x} > 0. \tag{A4}$$

Because $m'(\cdot) < 0$, we deduce that $w(r,s)$ rises with (y, b).

Next, we show that $w(r,s)$ increases with n and converges to y as n goes to infinity. Observe that (6) rewrites as follows

$$\begin{aligned}\frac{1}{y-w} &= \frac{g(\bar{\epsilon}) - \nu_g(b-w)G(b-w)^n - \int_{b-w}^{\bar{\epsilon}} h'(\epsilon)G(\epsilon)^{n-1}d\epsilon}{(1-G(b-w)^n)/n} \\ &= n\frac{g(\bar{\epsilon}) - \nu_g(b-w)G(b-w)^n}{1-G(b-w)^n} - \int_{b-w}^{\bar{\epsilon}} \frac{h'(\epsilon)}{g(\epsilon)} d\frac{G(\epsilon)^n - G(b-w)^n}{1-G(b-w)^n},\end{aligned}$$

where the first step follows from integration by parts. One can show that the first term rises with n. Second, log-concavity of $g(\cdot)$ implies that $-\frac{h'}{g}$ is increasing. Third, $\frac{G(\epsilon)^n - G(b-w)^n}{1-G(b-w)^n}$ is the distribution of the highest order statistics conditional on this being greater than $b-w$ and, therefore, it increases in n in the sense of first-order stochastic dominance. The result follows from these three facts.

Observe that

$$\lim_{n\to\infty} \frac{g(\bar{\epsilon}) - \nu_g(b-w)G(b-w)^n - \int_{b-w}^{\bar{\epsilon}} h'(\epsilon)G(\epsilon)^{n-1}d\epsilon}{(1-G(b-w)^n)/n} \to \infty$$

and therefore $\lim_{n\to\infty} w(r,s) \to y$. This follows from the fact that the numerator goes to $g(\bar{\epsilon})$, while the denominator goes to 0 due to the fact that $\lim_{n\to\infty} G(b-w)^n \to 0$.

Observe that

$$\begin{aligned}\frac{\partial w}{\partial n} =& \frac{m(b-w)}{1-m'(b-w)}\Bigg(\frac{1}{n} + \frac{\ln G(b-w)}{1-G(b-w)^n} + \\ & n\frac{m(b-w)}{1-G(b-w)^n}\int_{b-w}^{\bar{\epsilon}} \nu_G(\epsilon)g(\epsilon)G(\epsilon)^{n-1}\big(1-(n-1)(\ln G(b-w) - \ln G(\epsilon))\big)d\epsilon\Bigg) > 0\end{aligned}$$

Observe that the sum of the first two terms can be re-written as follows: $\frac{1-G^n}{n} + G^n \log F$, When evaluated at $\bar{\epsilon}$, this is zero, while at $-\bar{\epsilon}$, it is equal to 1 since $G^n \log G$ goes to zero (by L'Hopital). The derivative of this with respect to ϵ is given by $n^2 G^{n-1} g \log G$, which is 0 at $-\bar{\epsilon}$ and strictly negative in $(-\bar{\epsilon}, \bar{\epsilon}]$ and therefore $\frac{1-G^n}{n} + G^n \log G \geq 0$. This, together with the fact that $\ln G(b-w) - \ln G(\epsilon) \leq 0$ for all $\epsilon \geq b - w$, proves the result. □

Proof of Propostion 12. (i) By the chain rule and Lemma **??**,

$$\left.\frac{dW_j}{d\tau}\right|_{\tau=1} = \Phi_j(r_0) \cdot \left.\frac{dr_j}{d\tau}\right|_{\tau=1} = \Phi_j(r_0) \cdot \big(r_0 + r^{*\prime}(1)\big) \qquad \text{for every } j.$$

Since $S'(r_0) > 0$ by hypothesis, Lemma **??** gives $r_0 + r^{*\prime}(1) > 0$. By $A(r_0, s', \tau_{s'}) < \Xi(r_0, s', \tau_{s'})$, $\Phi_s(r_0) > 0$, so $dW_s/d\tau|_{\tau=1} = \Phi_s(r_0) \cdot (r_0 + r^{*\prime}(1)) > 0$. By $A(r_0, s', \tau_{s'}) > \Xi(r_0, s', \tau_{s'})$, $\Phi_{s'}(r_0) < 0$, so $dW_{s'}/d\tau|_{\tau=1} = \Phi_{s'}(r_0) \cdot (r_0 + r^{*\prime}(1)) < 0$.

(ii) Summing the expression above over all s industries,

$$\left.\frac{dW}{d\tau}\right|_{\tau=1} = \sum_{j=1}^{s} \left.\frac{dW_j}{d\tau}\right|_{\tau=1} = \big(r_0 + r^{*\prime}(1)\big) \sum_{j=1}^{s} \Phi_j(r_0),$$

which is (17). Because $r_0 + r^{*\prime}(1) > 0$ is a strictly positive scalar common to every industry, the sign of the left-hand side equals the sign of $\sum_j \Phi_j(r_0)$. snowing only that $\Phi_s(r_0) > 0$ for one industry s and $\Phi_{s'}(r_0) < 0$ for another, s', places no restriction on the sign of the sum over all s industries: the sum can be made positive, negative, or exactly zero by choosing the remaining $\{\Phi_j(r_0)\}_{j \neq s, s'}$, or by rescaling the sizes $\{n_j\}$ that enter each Φ_j through (**??**), freely and independently of $A(r_0, s') < \Xi(r_0, s', \tau_{s'})$ and $A(r_0, s', \tau_{s'}) > \Xi(r_0, s', \tau_{s'})$.

(iii) Construct two economies that agree on the sets of $A(r_0, s', \tau_{s'}) < \Xi(r_0, s', \tau_{s'})$ and $A(r_0, s', \tau_{s'}) > \Xi(r_0, s', \tau_{s'}, \tau_{s'})$ industries (in particular, both contain s satisfying $A(r_0, s', \tau_{s'}) < \Xi(r_0, s', \tau_{s'})$ and s' satisfying $A(r_0, s', \tau_{s'}) > \Xi(r_0, s', \tau_{s'})$, with $\Phi_s(r_0)$ and $\Phi_{s'}(r_0)$ identical across the two economies), but differ in the industry sizes $n_s, n_{s'}$. In the first economy, choose $n_s \Phi_s(r_0) > -\sum_{j \neq s} n_j \Phi_j(r_0)$, so that $\sum_j \Phi_j(r_0) > 0$ and, by part (ii), $dW/d\tau|_{\tau=1} > 0$; since $W(\tau)$ is continuous, this implies the welfare-maximizing τ^* in a neighborhood of $\tau = 1$ satisfies $\tau^* > 1$, a tax. In the second economy, choose n_s smaller (or $n_{s'}$ larger) so that $\sum_j \Phi_j(r_0) < 0$; by the same argument, $\tau^* < 1$, a subsidy. Both economies satisfy $A(r_0, s', \tau_{s'}) < \Xi(r_0, s', \tau_{s'})$ and $A(r_0, s', \tau_{s'}) > \Xi(r_0, s', \tau_{s'})$ identically;

they differ only in a primitive, n_j, that $A(r_0, s', \tau_{s'}) < \Xi(r_0, s', \tau_{s'})$ and $A(r_0, s', \tau_{s'}) > \Xi(r_0, s', \tau_{s'})$ do not restrict. This establishes the claim. □

B Online Appendix: Empirical Methodology and Robustness

B.1 Robot Integrators and Labor Market Concentration at the Commuting-Zone Level

B.1.1 Herfindahl-Hirschman Index

We use data from Choi and Marinescu (2024). They use job vacancy data from Lightcast covering 2007Q1 to 2021Q2 and measure labor market concentration with the Herfindahl-Hirschman Index (HHI) at the six-digit SOC occupation, commuting-zone-2000 (cz-2000), and quarter levels. The HHI of a market m (occupation $\times$ commuting zone) in quarter t is the sum of squared firm vacancy shares,

$$HHI_{m,t} = \sum_{j \in \mathcal{J}} S_{j,m,t}^2,$$

where $S_{j,m,t}$ is the share of vacancies posted by firm j in market m and quarter t. The authors provide a lower bound (all missing employer names are distinct firms) and an upper bound (all missing names are a single firm); we use the midpoint as our main measure and report the bounds as robustness. We rescale the published series (scaled by 1,000) to lie in $[0, 1]$.

We aggregate the occupation-by-commuting-zone HHI to the commuting-zone level. Because vacancy shares by occupation are not public, we cannot weight occupations by their local vacancy counts; instead we weight each occupation by its national employment share from the 2012 OEWS (detailed six-digit SOC), which keeps thinly-posted occupations—whose HHI is mechanically near one—from dominating the average.[19] We then merge with Acemoglu and Restrepo (2022) data on robot integrators, robot exposure, and demographic and economic characteristics (the latter originally county-level NHGIS data, Manson, Schroeder, Van Riper, and Ruggles, 2017; see Appendix B.1.3), which are defined at the commuting-zone-1990 (cz-1990) level. As the cz-2000 to cz-1990

[19] A simple (unweighted) occupation average inflates the commuting-zone HHI to roughly 0.70, well above the employment-weighted values reported in the labor-concentration literature; our weighted average is about 0.55.

crosswalk is not one-to-one, we map cz-2000 markets to cz-1990 using an employment-weighted average, where the weights are total County Business Patterns employment (2007) in the counties shared by each cz-2000 and cz-1990 pair (imputed CBP panel of Eckert, Fort, Schott, and Yang, 2021). Robot integrators in 2015—firms that install, program, and maintain robots—are originally compiled by Leigh and Kraft (2018) and used by Acemoglu and Restrepo (2022).

B.1.2 Exposure to robots

The measure of exposure-to-robots between 1993 and 2007 from Acemoglu and Restrepo (2020) captures the extent to which a commuting zone houses industries adopting robots at higher rates, and is defined as

$$\text{exposure-to-robots}_c^{t_0,t_1} = \sum_{i \in \mathcal{I}} l_{ci}^{1970} \, \overline{APR}_i^{t_0,t_1},$$

a weighted average of the adjusted penetration of robots in industry i between t_0 and t_1 ($\overline{APR}_i^{t_0,t_1}$), weighted by the 1970 labor share of industry i in commuting zone c (l_{ci}^{1970}). In turn,

$$\overline{APR}_i^{t_0,t_1} = \frac{1}{5} \sum_{j \in \mathcal{J}} \left[\frac{M_{i,t_1}^j - M_{i,t_0}^j}{L_{i,1990}^j} - g_{i,(t_0,t_1)}^j \frac{M_{i,t_0}^j}{L_{i,1990}^j} \right],$$

where $M_{i,t}^j$ is the stock of industrial robots in industry i in country j in year t, $L_{i,1990}^j$ is employment in industry i in country j in 1990, and $g_{i,(t_0,t_1)}^j$ is industry i's output growth in country j. The set $\mathcal{J}$ comprises five European countries ahead of the U.S. in robot adoption (Denmark, Finland, France, Italy, Sweden), used to avoid idiosyncratic U.S. variation. Acemoglu and Restrepo (2020) use IFR data for the robot stock and EU KLEMS for output and employment, and the 1970/1990/2000 Censuses and the ACS for U.S. industry employment shares; they publish the adjusted-penetration and exposure-to-robots series by industry and commuting zone.

B.1.3 Demographic, industry, and other controls

The control vector $X_{c,1990}$ replicates the commuting-zone controls of Acemoglu and Restrepo (2022); we take the variables directly from their replication data, which in turn

build on the "Robots and Jobs" commuting-zone files of Acemoglu and Restrepo (2020). They group into four blocks:

- **Demographic controls** (all measured in 1990): the male share, the urban share, log per-capita income, log population, and the shares of the working-age population by five education levels and five racial groups. These are county-level characteristics from the IPUMS National Historical Geographic Information System (NHGIS) (Manson et al., 2017), aggregated to commuting zones.

- **Industry-composition controls**: 1990 commuting-zone employment shares in agriculture, mining, construction, manufacturing, and finance and real estate, also from the NHGIS (Manson et al., 2017).

- **Other labor-market shocks**: the exposure to Chinese imports between 1990 and 2007 (Autor, Dorn, and Hanson, 2013) and the 1990 share of employment in routine occupations (Autor and Dorn, 2013), both as compiled in the Acemoglu and Restrepo (2020) commuting-zone data.

- **Regional dummies**: indicators for the four Census regions (Midwest, Northeast, South, West).

All controls are pre-determined relative to the 2015 robot-integrator outcome. Robot exposure (Appendix B.1.2) enters together with the industry-composition controls.

B.1.4 HHI Autocorrelation

Our 2007 commuting-zone HHI is highly persistent over time, so it captures structural—rather than transitory—labor market concentration. Table 8 reports cross-quarter correlations and Table 9 the first-order autoregression of the commuting-zone HHI in 2007.

Table 8. Persistence of CZ-level HHI: bivariate correlations

	Lower	Midpoint	Upper
2007-2011	0.937	0.959	0.902
2007-2014	0.915	0.946	0.881
2007-2018	0.910	0.945	0.903

Notes: Each cell reports the Pearson correlation across commuting zones between the annual HHI in 2007 and the average HHI over the indicated period. Lower and upper bounds correspond to the lower-bound and upper-bound HHI measures from Choi and Marinescu (2024). Midpoint corresponds to the simple average between the lower-bound and upper-bound HHI measures. Sample: 708 commuting zones.

Table 9. AR(1) estimates of CZ-level HHI persistence

	Pooled OLS		
	Lower	Midpoint	Upper
HHI_{t-1}	0.950***	0.963***	0.957***
	(0.002)	(0.002)	(0.002)
Implied half-life (quarters)	13.63	18.37	15.93
Observations	33943	33943	33943
Commuting zones	708	708	708

Standard errors in parentheses

Notes: The dependent variable is the quarterly commuting-zone vacancy HHI from Choi and Marinescu (2024), built from Lightcast (formerly Burning Glass) data. Standard errors clustered at the CZ level in parentheses. All specifications include quarter fixed effects.

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

B.1.5 Lightcast coverage and the direction of the bias

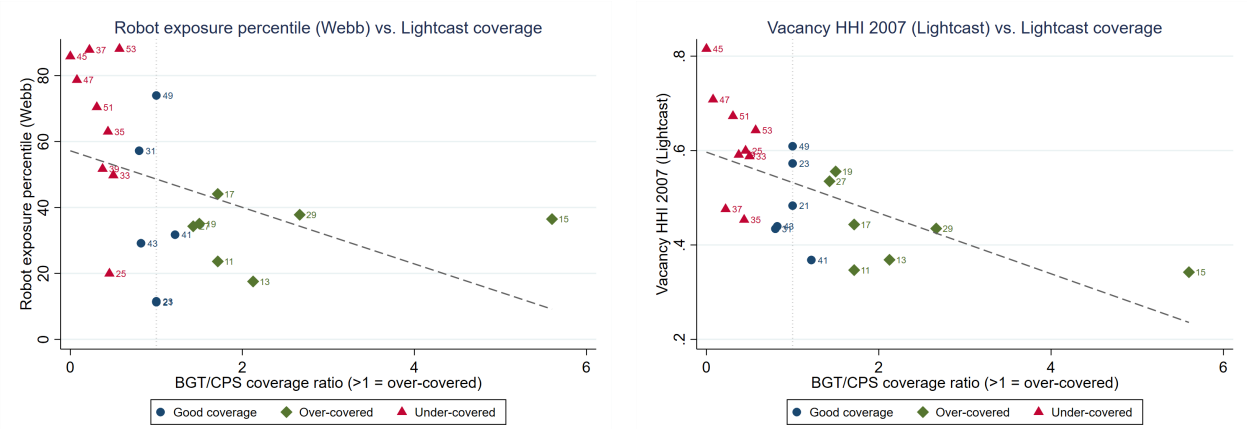
Online job postings do not cover all jobs uniformly: they over-represent higher-skill occupations and industries (Choi and Marinescu, 2024, Hershbein and Kahn, 2018). This matters only if coverage is correlated with our regressor or outcome. Hershbein and Kahn (2018) figures A1 and A2 document the Lightcast share of vacancy postings by occupation compared to the share of new hires by occupation from the Current Population Survey (CPS), and compared to the share of total employment by occupation from Occupational Employment Statistics (OES); and the share of posted vacancies by industry compared to the share of job openings from the Job Openings and Labor Turnover Survey (JOLTS).

In Figure 6 and in Table 10, we compare the BGT coverage ratio by occupation—the share of vacancies by occupation (Lightcast) to the share of new hires by occupation (CPS) (Hershbein and Kahn, 2018)—to the robot-exposure index of Webb (2020), and the weighted-average HHI in 2007. Less-covered occupations are more exposed to robots and face higher vacancy HHI. Indeed, the higher vacancy HHI in less-covered occupations could be a data artifact, as vacancy shares are biased upward if the number of firms is under-reported; nevertheless, if we were over-estimating the HHI for more robot-exposed occupations we would be biasing against finding a negative HHI–automation relationship, so our estimates are, if anything, conservative.

Table 11 and Figure 7 panel (a), compare the BGT coverage ratio by industry—the share of vacancies by industry (Lightcast) to the share of job openings by industry (JOLTS) (Hershbein and Kahn, 2018)—to the labor-supply elasticity to the firm from Webber (2015) (LEHD 1985-2008). There is practically no correlation between labor supply elasticity and Lightcast coverage. Figure 7 panel (b) shows the adjusted penetration of robots from Acemoglu and Restrepo (2020)[20] by Lightcast coverage status. As IFR coverage is concentrated in manufacturing industries, and Hershbein and Kahn (2018) report good coverage in the manufacturing industry, most industries in Figure 7 are marked as having good coverage. One caveat is that, even though the manufacturing industry has good coverage in Lightcast, the production occupation (SOC major group 51) is reported to

[20]This variable corresponds to the stock variation, between 2004 and 2014, of industrial robots per one thousand workers by IFR industry definition adjusted by output growth. In this case we use the variable defined from the stock variation of robots in the US from Acemoglu and Restrepo (2020), so it is different from the one described in Appendix section B.1.2, as it is from the U.S. and not from the 5 European countries described there.

have bad coverage. This is an issue that we address in one of our robustness checks (see Appendix B.1.6).



(a) Coverage vs. robot exposure

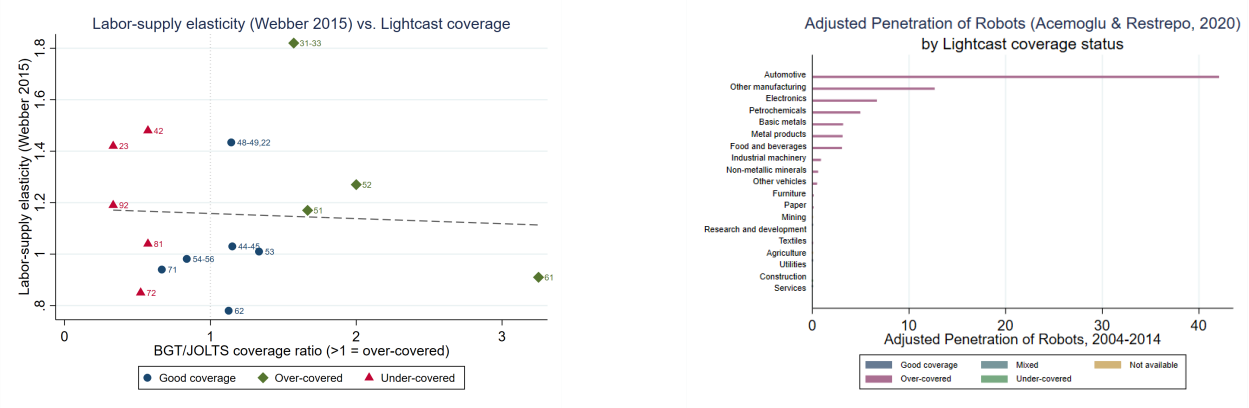
(b) Coverage vs. vacancy HHI

Fig. 6. Occupation coverage in Lightcast.
Each point is a SOC two-digit major group (marker = SOC code). Coverage in panel (a) is the ratio between share of Lightcast vacancy posting vs CPS share of new hires. Coverage in panel (b) is the ratio between share of Lightcast vacancy posting vs JOLTS share of job postings. Coverage data is from Hershbein and Kahn (2018). Robot exposure is from Webb (2020).

Table 10. Lightcast coverage, robotization and concentration across SOC major groups

	Coverage	Robot	VacancyHHI
Coverage	1.000	.	.
Robot	-0.412	1.000	.
VacancyHHI	-0.618	0.559	1.000

Pearson correlations across the 22 SOC two-digit major groups (vacancy HHI: 2007). Coverage is the BGT/CPS new-jobs posting ratio (Hershbein and Kahn, 2018). Robot exposure from Webb (2020); vacancy HHI is the author based occupation weighted average constructed with data from Choi and Marinescu (2024). Lower Lightcast coverage is associated with higher robot exposure and higher vacancy concentration.



(a) Industry coverage vs. robot intensity

(b) Robot adoption by coverage group

Fig. 7. Industry coverage in Lightcast.
In panel (a), each point is a NAICS two-digit sector (marker = NAICS code). Coverage in panel (a) is the ratio between share of Lightcast vacancy posting vs JOLTS share of job postings. Coverage data is from Hershbein and Kahn (2018). Labor supply elasticity to the firm is from Webber (2015).
In panel (b), each bar correspond to an IFR industry definition. Adjusted penetration of robots is from Acemoglu and Restrepo (2020).

Table 11. Lightcast coverage and labor-supply elasticities across NAICS sectors

	Correlation
Coverage vs. labor-supply elasticity	-0.054

Pearson correlation across NAICS two-digit sectors with non-missing coverage, between the BGT/JOLTS posting ratio (Hershbein and Kahn, 2018) and the labor-supply elasticity of Webber (2015) (estimated from LEHD/QCEW microdata for 1985–2008; a higher elasticity indicates lower monopsony power).

B.1.6 Robustness: coverage-restricted samples and occupation-subset HHI

We probe the coverage and outlier concerns in two complementary ways. First, Table 12 re-estimates the preferred specification on restricted commuting-zone samples: dropping rural zones, dropping zones dominated by under-covered industries, and excluding the automotive-intensive zones. Second, Table 13 re-estimates it after recomputing the HHI over occupation subsets: excluding under-covered occupations, and excluding over-covered occupations. As one of the under-covered occupation is production (SOC code 51), and this is one of the occupations heavily replaced by robots, we perform robustness excluding under-covered occupations with and without it. The negative coefficient is stable across both exercises. The author-built variables underlying these samples and subsets are described in Appendix B.5.

Table 12. Robustness: Excluding Commuting Zones with Limited Lightcast Coverage or Outlier Automation (Midpoint HHI)

HHI 2007	-0.2689***	-0.2468*	-0.3826***	-0.2205***	-0.3187*
	(0.0835)	(0.1359)	(0.1166)	(0.0806)	(0.1789)
Base Controls	✓	✓	✓	✓	✓
Excl. top-rural CZs		✓			✓
Excl. top u-cov. industry CZs			✓		✓
Excl. top u-cov. Automotive CZs				✓	✓
Observations	721	541	541	665	379
R-squared	0.45	0.49	0.47	0.41	0.47

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Each column re-estimates the baseline specification (HHI 2007 over all occupations) on a restricted commuting-zone sample. Column 2 drops the top quartile of 1990 rural population share. Column 3 drops the top quartile of employment share in under-covered industries (NAICS 11 Agriculture, 23 Construction, 42 Wholesale, 72 Accommodation and Food, 81 Other Services, 92 Public Administration; NAICS 92 is absent from County Business Patterns). Column 4 drops the top decile of automotive employment share. Column 5 imposes the rural, industry, and automotive restrictions jointly. All restrictions here are on commuting-zone composition. The under-covered industry classification follows Hershbein and Kahn (2018). All specifications include the full set of controls of Table 1.

Table 13. Robustness: HHI Recomputed over Occupations Grouped by Lightcast Coverage (Midpoint HHI)

HHI 2007	-0.2689***	-0.2252***	-0.2125**	-0.2193***
	(0.0835)	(0.0828)	(0.0794)	(0.0809)
Occupation set	All occupations	Excl. under (keep 51)	Excl. under (incl 51)	Excl. over
Base Controls	✓	✓	✓	✓
Observations	721	721	721	721
R-squared	0.45	0.45	0.45	0.45

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Each column re-estimates the baseline specification with HHI 2007 recomputed over a subset of occupations, on the full commuting-zone sample. Column 1 uses all occupations. Columns 2 and 3 exclude the SOC major groups under-represented in Lightcast (Hershbein and Kahn, 2018), keeping (column 2) or also excluding (column 3) production occupations (SOC 51); column 4 instead excludes the over-represented major groups. Stability of the coefficient whether under- or over-covered occupations are removed indicates the result is not driven by uneven Lightcast coverage. All specifications include the full set of controls of Table 1.

B.1.7 Results by HHI bound

Our main specification uses the midpoint of the Choi and Marinescu (2024) lower- and upper-bound HHI. Here we re-estimate the baseline and the mechanism specifications using each bound separately. The negative association is stable in sign and significance across bounds: the lower-bound HHI yields a similar but somewhat smaller coefficient, while the upper bound yields a larger one. Tables 14 and 15 report the baseline by bound, and Tables 16– 19 the robustness checks by bound.

Table 14. Labor Market Concentration and the Local Presence of Robot Integrators: Baseline Specification (Lower-Bound HHI)

HHI 2007	-0.9818***	-0.1264	-0.1452**	-0.1364**	-0.4142***
	(0.0923)	(0.0827)	(0.0694)	(0.0665)	(0.1075)
Census Region Dummies	✓	✓	✓	✓	✓
Demographic controls (1990)		✓	✓	✓	✓
Industry controls + robot exposure			✓	✓	✓
Other shocks controls				✓	✓
Good cov.					✓
Observations	721	721	721	721	541
R-squared	0.28	0.41	0.45	0.45	0.46

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Columns 1-4 progressively add controls. Column 1 includes only Census region dummies. Column 2 adds demographic controls. Column 3 adds industry employment shares and the commuting-zone-level exposure to robots over 1993-2007. Column 4 adds controls for exposure to Chinese imports and the routine occupation share. Column 5 repeats the fully-controlled specification of column 4 on the subsample of commuting zones above the 25th percentile of employment-weighted occupational coverage (the share of national employment in the occupations observed in every quarter of 2007 in the respective commuting zone). Demographic controls include the male share, the urban share, log per-capita income, log population, and shares of workers by five education levels and five racial groups, all measured in 1990. Industry controls include 1990 employment shares in agriculture, mining, construction, manufacturing, and finance and real estate. For a more detailed description of data sources refer to Appendix Section B.5.

B.2 Endogenity and Bartik Instruments

The instrument rests on two key premises. First, the dramatic decline in U.S. unionization between 1990 and 2006 was driven primarily by national policy, legal, and institutional developments—including the spread of right-to-work laws, pivotal NLRB rulings, and structural shifts in product-market competition—that operated largely orthogonal to commuting-zone-specific automation demand or local robot adoption decisions. Sec-

Table 15. Labor Market Concentration and the Local Presence of Robot Integrators: Baseline Specification (Upper-Bound HHI)

HHI 2007	-1.5944***	-0.3910**	-0.3401**	-0.3732**	-0.6514***
	(0.1289)	(0.1669)	(0.1481)	(0.1549)	(0.2208)
Census Region Dummies	✓	✓	✓	✓	✓
Demographic controls (1990)		✓	✓	✓	✓
Industry controls + robot exposure			✓	✓	✓
Other shocks controls				✓	✓
Good cov.					✓
Observations	721	721	721	721	541
R-squared	0.30	0.41	0.45	0.46	0.46

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Columns 1-4 progressively add controls. Column 1 includes only Census region dummies. Column 2 adds demographic controls. Column 3 adds industry employment shares and the commuting-zone-level exposure to robots over 1993-2007. Column 4 adds controls for exposure to Chinese imports and the routine occupation share. Column 5 repeats the fully-controlled specification of column 4 on the subsample of commuting zones above the 25th percentile of employment-weighted occupational coverage (the share of national employment in the occupations observed in every quarter of 2007 in the respective commuting zone). Demographic controls include the male share, the urban share, log per-capita income, log population, and shares of workers by five education levels and five racial groups, all measured in 1990. Industry controls include 1990 employment shares in agriculture, mining, construction, manufacturing, and finance and real estate. For a more detailed description of data sources refer to Appendix Section B.5.

Table 16. Robustness: Excluding Commuting Zones with Limited Lightcast Coverage or Outlier Automation (Lower-Bound HHI)

HHI 2007	-0.1364**	-0.0933	-0.2050**	-0.1246**	-0.1842
	(0.0665)	(0.1008)	(0.0930)	(0.0576)	(0.1271)
Base Controls	✓	✓	✓	✓	✓
Excl. top-rural CZs		✓			✓
Excl. top u-cov. industry CZs			✓		✓
Excl. top u-cov. Automotive CZs				✓	✓
Observations	721	541	541	665	379
R-squared	0.45	0.48	0.47	0.41	0.47

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Each column re-estimates the baseline specification (HHI 2007 over all occupations) on a restricted commuting-zone sample. Column 2 drops the top quartile of 1990 rural population share. Column 3 drops the top quartile of employment share in under-covered industries (NAICS 11 Agriculture, 23 Construction, 42 Wholesale, 72 Accommodation and Food, 81 Other Services, 92 Public Administration; NAICS 92 is absent from County Business Patterns). Column 4 drops the top decile of automotive employment share. Column 5 imposes the rural, industry, and automotive restrictions jointly. All restrictions here are on commuting-zone composition. The under-covered industry classification follows Hershbein and Kahn (2018). All specifications include the full set of controls of Table 14.

Table 17. Robustness: Excluding Commuting Zones with Limited Lightcast Coverage or Outlier Automation (Upper-Bound HHI)

HHI 2007	-0.3732**	-0.3740*	-0.4585**	-0.2840*	-0.3296
	(0.1549)	(0.2169)	(0.2043)	(0.1586)	(0.2716)
Base Controls	✓	✓	✓	✓	✓
Excl. top-rural CZs		✓			✓
Excl. top u-cov. industry CZs			✓		✓
Excl. top u-cov. Automotive CZs				✓	✓
Observations	721	541	541	665	379
R-squared	0.46	0.49	0.48	0.42	0.47

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Each column re-estimates the baseline specification (HHI 2007 over all occupations) on a restricted commuting-zone sample. Column 2 drops the top quartile of 1990 rural population share. Column 3 drops the top quartile of employment share in under-covered industries (NAICS 11 Agriculture, 23 Construction, 42 Wholesale, 72 Accommodation and Food, 81 Other Services, 92 Public Administration; NAICS 92 is absent from County Business Patterns). Column 4 drops the top decile of automotive employment share. Column 5 imposes the rural, industry, and automotive restrictions jointly. All restrictions here are on commuting-zone composition. The under-covered industry classification follows Hershbein and Kahn (2018). All specifications include the full set of controls of Table 15.

Table 18. Robustness: HHI Recomputed over Occupations Grouped by Lightcast Coverage (Lower-Bound HHI)

HHI 2007	-0.1364**	-0.1470**	-0.1339**	-0.0994
	(0.0665)	(0.0673)	(0.0653)	(0.0639)
Occupation set	All occupations	Excl. under (keep 51)	Excl. under (incl 51)	Excl. over
Base Controls	✓	✓	✓	✓
Observations	721	721	721	721
R-squared	0.45	0.45	0.45	0.45

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Each column re-estimates the baseline specification with HHI 2007 recomputed over a subset of occupations, on the full commuting-zone sample. Column 1 uses all occupations. Columns 2 and 3 exclude the SOC major groups under-represented in Lightcast (Hershbein and Kahn, 2018), keeping (column 2) or also excluding (column 3) production occupations (SOC 51); column 4 instead excludes the over-represented major groups. Stability of the coefficient whether under- or over-covered occupations are removed indicates the result is not driven by uneven Lightcast coverage. All specifications include the full set of controls of Table 14.

Table 19. Robustness: HHI Recomputed over Occupations Grouped by Lightcast Coverage (Upper-Bound HHI)

HHI 2007	-0.3732**	-0.2615*	-0.2617**	-0.3178*
	(0.1549)	(0.1365)	(0.1287)	(0.1608)
Occupation set	All occupations	Excl. under (keep 51)	Excl. under (incl 51)	Excl. over
Base Controls	✓	✓	✓	✓
Observations	721	721	721	721
R-squared	0.46	0.45	0.45	0.46

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Each column re-estimates the baseline specification with HHI 2007 recomputed over a subset of occupations, on the full commuting-zone sample. Column 1 uses all occupations. Columns 2 and 3 exclude the SOC major groups under-represented in Lightcast (Hershbein and Kahn, 2018), keeping (column 2) or also excluding (column 3) production occupations (SOC 51); column 4 instead excludes the over-represented major groups. Stability of the coefficient whether under- or over-covered occupations are removed indicates the result is not driven by uneven Lightcast coverage. All specifications include the full set of controls of Table 15.

ond, declining unionization weakened workers' bargaining power and increased effective labor-market concentration by eroding countervailing-power institutions, thereby allowing employers to capture greater monopsony rents and to adjust labor-market structure.

For each commuting zone c, we construct the instrument as

$$Z_c = \sum_j \omega_{jc,1990} \cdot \Delta\text{Union}_{j,1990\to2006},$$

where $\omega_{jc,1990}$ denotes the 1990 employment share of industry j (classified at the SIC 2-digit level) in commuting zone c, and $\Delta\text{Union}_{j,1990\to2006}$ is the national-level change in the industry's unionization rate. The exclusion restriction requires that national industry-level union trends, when interacted with 1990 local industry composition, affect 2007 local HHI *only* through shifts in labor-market structure—not directly through local automation adoption or through unobserved commuting-zone-level shocks correlated with robot demand.

The credibility of this exclusion restriction depends on controlling for confounders that might simultaneously drive both union decline and automation. We therefore progressively add demographic controls, baseline industry composition, and robot-exposure measures, and indicators of other local economic shocks. If the instrument remains weak after such controls, it suggests that the residual variation in national union trends is lim-

ited, or that union decline affects labor-market outcomes through channels not fully captured by vacancy-based HHI.

Table 20. Instrumenting HHI 2007 with a Bartik Index of Changes in Industry Unionization Rates (1990-2006) (Midpoint HHI)

HHI 2007	-2.1114***	1.9048	1.4259	1.1344
	(0.6763)	(3.4811)	(4.4796)	(4.5954)
Census Region Dummies	✓	✓	✓	✓
Demographic controls (1990)		✓	✓	✓
Industry controls + robot exposure			✓	✓
Other shocks controls				✓
Observations	720	720	720	720
First-stage F stat.	3.4	1.4	0.6	0.5
Anderson-Rubin Wald test p-value	0.14	0.46	0.71	0.78

Robust standard errors clustered by state in parentheses.

*** $p < 0.01$, ** $p < 0.05$, *$p < 0.1$.

Specification follows the baseline OLS in Table 1. The first-stage F-statistic is below conventional thresholds for instrument strength in all specifications, indicating a weak instrument. Anderson-Rubin Wald test p-values are reported as the inference-robust alternative to the t-test on the second-stage coefficient.

Table 20 presents the IV estimates across four nested specifications. In column (1), with only Census region dummies, the instrumented coefficient on HHI is large and highly significant: $\hat{\beta}_{\text{IV}} = -2.11$ (robust s.e. 0.68). This implies that a unit increase in instrumented HHI reduces robot-integrator presence by approximately 2.1 percentage points, a stark contrast to the OLS estimate and consistent with the theory that labor-market concentration deters substitutive automation. However, this estimate must be viewed with caution: the first-stage F-statistic is only 3.4, substantially below the Stock–Yogo critical value of approximately 10 for 10% maximal IV relative bias, indicating a *weak instrument*.

As we add richer controls—demographic characteristics in column (2), industry composition and baseline robot exposure in column (3), and other local economic shocks in column (4)—the first-stage F-statistic deteriorates sharply to 1.4, 0.6, and 0.5, respectively. Correspondingly, the second-stage coefficient shrinks toward zero, loses statistical significance, and becomes unstable: estimates range from 1.9 to 1.1, and the Anderson–Rubin Wald test p-values rise from 0.14 to 0.78, indicating we cannot reject a null effect at conventional levels. This pattern of deteriorating first-stage strength with the addition of controls is a hallmark of weak-instrument bias, where the estimated coefficient reflects

sample variation in the weaker residual instrument rather than a true causal relationship.

The collapse in instrument strength occurs for several plausible reasons. First, secular union decline may influence local labor-market structure and wages through mechanisms not fully captured by vacancy-based HHI—for example, through effects on within-firm wage-setting institutions, worker morale, and skill investment—rather than through changes in the number or concentration of employers posting vacancies. Second, the 1990–2006 national unionization trends may be largely absorbed by industry fixed effects or by our baseline robot-exposure controls, which themselves depend on industry composition. This leaves minimal residual variation in the instrument to identify the effect of HHI conditional on rich controls. Third, measurement error in HHI constructed from vacancy data—which may not fully represent the effective concentration facing workers—could amplify weak-instrument problems if the Bartik shock correlates only weakly with true labor-market concentration.

These results highlight the empirical difficulty of instrumenting HHI with macro-level shocks. While the univariate IV estimate (column 1) is suggestive and theoretically aligned with our model predictions, it lacks credibility as causal evidence because it cannot rule out omitted-variable bias from local shocks correlated with both initial industry composition and robot adoption. The sharp deterioration in first-stage strength upon adding standard controls—rather than showing stability—indicates that the instrument is not reliably identifying local labor-market concentration after conditioning on observable confounders. We therefore treat these IV results primarily as a *robustness check* rather than as definitive causal estimates.

To strengthen identification, alternative approaches would require either (i) quasi-experimental variation such as plant-level regulatory shocks to entry or labor-market access in specific commuting zones; (ii) historical instrumental variables exploiting long-lagged variation in HHI not subject to recent endogeneity; or (iii) direct estimation of labor-market power using structural methods (e.g., production-function or markup estimation). We leave these extensions to future work and note that the OLS results, combined with sensitivity analysis and the calibrated model, provide the primary empirical foundation for our quantitative conclusions.

B.3 Markdowns and Automation Imports at the Country Level

B.3.1 Markdowns by country

We estimate markdowns following Eslava, García-Marín, and Messina (2023), who apply the production approach of Hall, Blanchard, and Hubbard (1986) and De Loecker and Warzynski (2012)—extended to markdowns by Yeh, Macaluso, and Hershbein (2022), among others. The De Loecker and Warzynski (2012) markup for firm i is

$$\mu_i \equiv \frac{P_i}{MC_i} = \left(\frac{\partial F(\cdot)}{\partial X_i^{k'}} \frac{X_i^{k'}}{Q_i}\right) \left(\frac{V_i^{k'} X_i^{k'}}{P_i Q_i}\right)^{-1},$$

where $X_i^{k'}$ is a fully flexible, static input not subject to monopsony and $V_i^{k'}$ its unit price. If labor is fully flexible but the firm has monopsony power, the first-order condition gives the wage markdown

$$\nu_i \equiv \frac{MPL_i}{w_i} = \frac{1}{\mu_i} \left[\left(\frac{\partial F(\cdot)}{\partial l_i} \frac{l_i}{Q_i}\right) \left(\frac{\omega_i l_i}{P_i Q_i}\right)^{-1}\right].$$

Assuming all firms in an economic sector (two-digit ISIC Rev. 3) share a CRS technology, the input elasticity is the average cost share,

$$\hat{\eta}_{sector} \equiv \frac{1}{N} \sum_{i \in sector} \frac{\omega_i l_i}{P_i Q_i}.$$

The only source for the required firm-level sales and input/labor cost data is the World Bank Enterprise Surveys (WBES) (World Bank), designed to be representative of manufacturing (ISIC Rev. 3 Group D) and retail (ISIC 52). We follow the cleaning and imputation of Eslava et al. (2023): keep observations taken from book records; impute missing labor or input costs from weighted country-specific regressions of costs on sales with two-digit industry-year fixed effects; drop cost-share outliers (below the 5th and above the 95th percentile); truncate costs exceeding sales while holding the labor-to-input cost ratio fixed; and drop country-years with fewer than 250 firms. This yields markdowns for 108 countries surveyed between 2006 and 2024 (excluding 2020 for COVID-19), or 202 country-year observations; we use the simple average across waves by country.

B.3.2 Automation-technology imports

We classify imported capital goods following Acemoglu and Restrepo (2022), who group six-digit trade codes into (i) industrial automation, (ii) non-automated ("manual") industrial capital, and (iii) non-industrial technologies. From these we build the following per-worker import measures (UN Comtrade (United Nations Statistics Division), accumulated value net of re-exports, 1996–2015):

- **Net robot imports** (HS-1996 code 847950): the narrow, ISO-defined industrial robots used by Acemoglu and Restrepo (2022).
- **Broad automation imports** (Acemoglu and Restrepo (2022) set 1): robots plus numerically-controlled machines, dedicated industrial machinery, automatic welding equipment, and related categories.
- **Computer imports** (placebo): a complementary technology drawn from Acemoglu and Restrepo (2022)'s non-industrial set; the model predicts no relationship with market power.
- **Non-industrial imports, ex agriculture** (placebo): computers, laundry, and vending machinery.
- **Agricultural-machinery imports**: isolated because agriculture is itself being automated, so this is an automation outcome rather than a clean placebo.

We normalize each measure by industrial employment in 1995 (ILO, measured in thousands of workers and adjusted by hours per worker from PWT 9.0)—so each outcome is expressed per one thousand industrial workers—rather than by total intermediate imports as in Acemoglu and Restrepo (2022). We restrict to countries with positive net robot imports and exclude Germany (a major producer) and Luxembourg (a major entry port), following Acemoglu and Restrepo (2022). We further exclude India as being an outlier in Net Imports Robots per one thousand industrial workers.

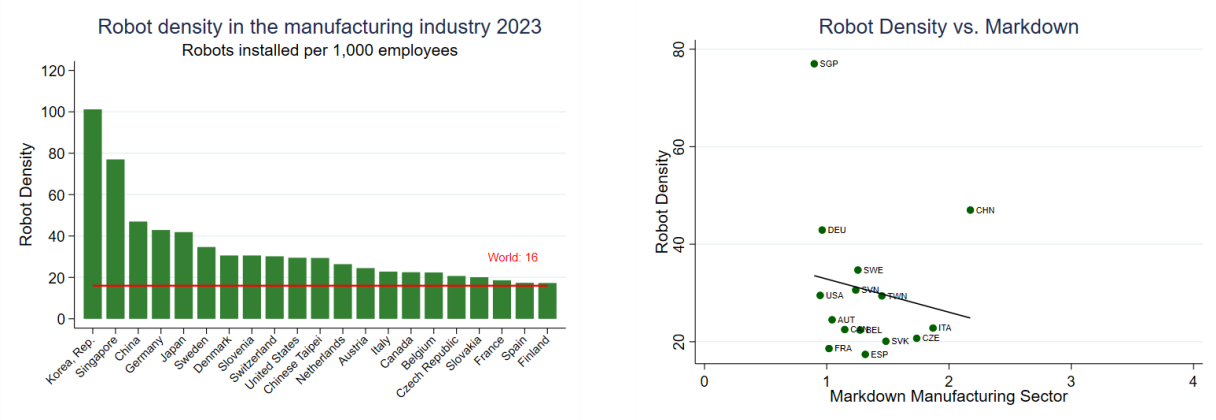
B.3.3 Controls and the bad-control taxonomy

Region dummies cover seven groups: six non-OECD regions (Sub-Saharan Africa, East Asia and Pacific, Eastern Europe and Central Asia, Latin America and the Caribbean,

Middle East and North Africa, South Asia) and OECD. Pre-determined country covariates (all 1995) are log population and the human-capital index from PWT 9.0 (Feenstra, Inklaar, and Timmer, 2015), the latter based on average years of schooling from Barro and Lee (2013). Three further variables—log GDP per capita (PPP, PWT 9.0), log manufacturing value added (constant 2015 USD, UNIDO (United Nations Industrial Development Organization)), and log total intermediate imports (HS 82, 84, 85, 87, 90; UN Comtrade)—are potential bad controls: Armangué-Jubert, Guner, and Ruggieri (2024) show that labor market power lowers aggregate output, so these development proxies are endogenous to the markdown. We therefore exclude them from the preferred specification (column 2) and add them one at a time in columns 3–5 only to show how much they attenuate the coefficient.

B.3.4 Baseline Analysis

Appendix B.3 provides a complete description of the data and its sources. As suggestive cross-country corroboration, Figure 8, panel (b), shows a negative correlation between the stock of robots installed per thousand workers (IFR World Robotics 2024) and manufacturing markdowns, for countries above the world average.



(a) Robot Density Manufacturing Sector 2023 **(b)** Correlation Robot Density and Markdowns

Fig. 8. Robot Density and Markdowns, Manufacturing Sector.
Robot density is robots installed per 1,000 manufacturing employees in 2023. Markdowns are estimated following Eslava et al. (2023) on the World Bank Enterprise Surveys (WBES); for countries with several waves we take the simple average. Panel (b) excludes the Republic of Korea (KOR) as an extreme outlier (101 robots/1,000, markdown 1.77). The data on robots installed is a sample of the 2024 World Robotics Report of the International Federation of Robotics (IFR) and was made publicly available on its webpage, (IFR World Robotics 2024, `https://ifr.org/wr-industrial-robots/`).

We then adapt the cross-country specification of Acemoglu and Restrepo (2022), with the manufacturing markdown as the regressor of interest:

$$log\left(\frac{\Delta Im_R_c^{1996–2015}}{L_{c,1995}}\right) = \beta_0 + \beta_1\, markdown_c + \Gamma X_{c,1995} + \mu_c,$$

where $Im_R_c^{1996–2015}$ is the accumulated value of robot imports (HS-1996 code 847950), net of re-exports, between 1996 and 2015 (UN Comtrade), divided by industrial employment in 1995 (measured in thousands of workers) adjusted by hours worked ($L_{c,1995}$), so the outcome is net robot imports per one thousand industrial workers; $markdown_c$ is the manufacturing markdown estimated as in Eslava et al. (2023)[21]; and $X_{c,1995}$ are 1995 controls. We normalize by employment (rather than by total intermediate imports as in Acemoglu and Restrepo (2022)) as this measure has a strong positive correlation with the stock of industrial robots, in logs[22].

Table 21. Log Net Robot Imports vs. Markdown (Manufacturing)

Markdown	-1.2837***	-1.1877***	-0.3877	-0.5527	-0.1563
	(0.3777)	(0.3948)	(0.4561)	(0.4364)	(0.3228)
Region dummies	✓	✓	✓	✓	✓
Pre-determined covs (1995)		✓	✓	✓	✓
+ GDP pc (bad control)			✓		
+ Manuf. VA (bad control)				✓	
+ Total imports (bad control)					✓
Observations	100	93	93	93	93
R-squared	0.63	0.68	0.76	0.77	0.85

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Dependent variable correspond to net imports per one thousand industrial workers, in logs.

Column 2 (region dummies plus pre-determined 1995 controls: log population and human capital) is the preferred specification; column 1 omits the pre-determined controls. Columns 3–5 each add ONE potentially endogenous control to column 2 (GDP per capita, log manufacturing value added, and log total intermediate imports, respectively). For a more detailed description of data sources refer to Appendix Section B.5.

Table 21 reports the results. Our preferred specification is column (2), which adds

[21] A main assumption in the Eslava et al. (2023) methodology is imposing a common technology for industries in the same sector across countries. As high-automation countries should have lower labor elasticities, this assumption will overestimate their labor elasticity and markdown. Nevertheless, overestimating markdowns for high-automating countries should bias the results against finding our result.

[22] In Figure A6 in Acemoglu and Restrepo (2022) they show a strong positive correlation between the log of robot stock variation per one thousand workers and log net imports of robots per one thousand workers.

the pre-determined 1995 controls (log population and human capital) to region fixed effects. The markdown coefficient is negative and significant (-1.19); given a markdown standard deviation of 0.52, this implies that a one-standard-deviation increase in market power is associated with roughly a 62% reduction in net robot imports per one thousand industrial workers. Columns (3)–(5) add GDP per capita, manufacturing value added, and total intermediate imports one at a time; we view these variables as potential bad controls because labor market power may affect productivity, industrial development, and firms' adoption decisions. If markdowns reduce productivity and economic development, as suggested by Armangué-Jubert et al. (2024), then GDP per capita, manufacturing value added, and intermediate imports may lie on the causal path between labor market power and robot adoption. Controlling for them would therefore absorb part of the effect of interest, mechanically attenuating the estimated markdown coefficient.

The result is not specific to ISO-defined robots. Table 22 shows that imports of a broader set of automation technologies—numerically-controlled machines, dedicated machinery, automatic welding, and the like, following Acemoglu and Restrepo (2022)—also fall with markdowns. As a placebo, Table 23 repeats the exercise for imports of computers, a complementary technology that the model predicts should be unrelated to market power: the markdown coefficient is small and insignificant, in contrast to the negative coefficients for robots and broad automation.

Figure 9 is an added-variable (partial-regression) plot for the preferred specification: the slope of the fitted line equals the markdown coefficient. Appendix B.4 shows the analogous plots by income group; the negative relationship is clearest among developing countries. Additionally, as seen in table 24, the coefficient is not attenuated when adding GDP per capita and manufacturing value added, as these countries are in a similar stage of development, so the causal path between markdown and development is less important.

B.3.5 Data limitations

The World Bank Enterprise Surveys cover only formal, registered private firms in manufacturing and services, excluding the informal sector, agriculture, and state-owned firms. They are also fielded in different years for different countries, with a changing country set across waves, so they do not form a country panel; we therefore use a cross-section,

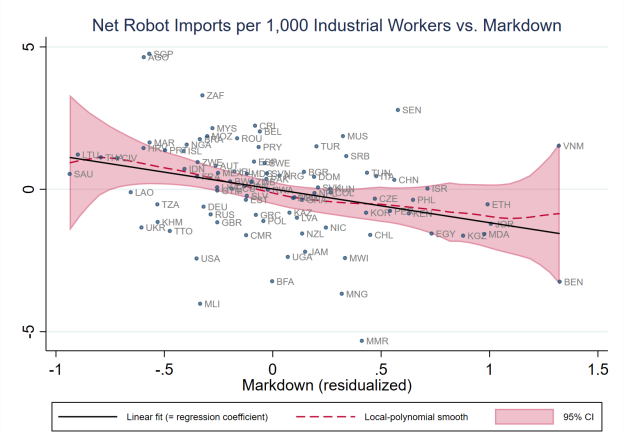


Fig. 9. Log Net Robot Imports per 1,000 Industrial Workers vs. Markdown—preferred specification.
Added-variable (partial-regression) plot: the outcome (log net robot imports per one thousand industrial workers) and the markdown are residualized on the preferred controls (region fixed effects, log population and human capital, 1995). The linear-fit slope equals the regression coefficient; the dashed line is a local-polynomial smooth with a 95% confidence band.

averaging each country's markdown across its available waves. We see this as a limited concern because market power is slow-moving and persistent (De Loecker, Eeckhout, and Mongey, 2021), so a wave-averaged markdown should capture a country's structural market power rather than transitory shocks.

B.3.6 Robustness checks

B.3.7 Full results

Tables 25 and 26 report the remaining outcomes (the non-industrial placebo and agricultural machinery). Tables 24–30 interact the markdown with income-group, showing the negative relationship is clearest among developing countries.

Table 22. Log Net Automation-Technology Imports vs. Markdown (Manufacturing)

Markdown	-0.5673**	-0.5314**	-0.1636	-0.3109	-0.0524
	(0.2362)	(0.2391)	(0.2898)	(0.2650)	(0.1566)
Region dummies	✓	✓	✓	✓	✓
Pre-determined covs (1995)		✓	✓	✓	✓
+ GDP pc (bad control)			✓		
+ Manuf. VA (bad control)				✓	
+ Total imports (bad control)					✓
Observations	102	94	94	94	94
R-squared	0.54	0.60	0.69	0.66	0.87

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Dependent variable correspond to net imports per one thousand industrial workers, in logs.

Column 2 (region dummies plus pre-determined 1995 controls: log population and human capital) is the preferred specification; column 1 omits the pre-determined controls. Columns 3–5 each add ONE potentially endogenous control to column 2 (GDP per capita, log manufacturing value added, and log total intermediate imports, respectively). For a more detailed description of data sources refer to Appendix Section B.5.

Table 23. Log Net Computer Imports (Placebo) vs. Markdown (Manufacturing)

Markdown	-0.4921	-0.4573	0.0958	-0.1082	0.2024
	(0.3482)	(0.3374)	(0.3954)	(0.3640)	(0.2437)
Region dummies	✓	✓	✓	✓	✓
Pre-determined covs (1995)		✓	✓	✓	✓
+ GDP pc (bad control)			✓		
+ Manuf. VA (bad control)				✓	
+ Total imports (bad control)					✓
Observations	101	94	94	94	94
R-squared	0.49	0.58	0.68	0.65	0.85

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Dependent variable correspond to net imports per one thousand industrial workers, in logs.

Column 2 (region dummies plus pre-determined 1995 controls: log population and human capital) is the preferred specification; column 1 omits the pre-determined controls. Columns 3–5 each add ONE potentially endogenous control to column 2 (GDP per capita, log manufacturing value added, and log total intermediate imports, respectively). For a more detailed description of data sources refer to Appendix Section B.5.

Table 24. Marginal Effect of Markdown on Net Robot Imports per 1,000 Industrial Workers, by Income Group

Markdown					
Developed Countries	-0.5123	-0.4897	0.3153	0.2045	0.8127*
	(0.4257)	(0.5485)	(0.6126)	(0.4898)	(0.4157)
Developing Countries	-1.1901***	-1.1925***	-0.8141*	-0.8743**	-0.2160
	(0.3950)	(0.4056)	(0.4561)	(0.3925)	(0.4355)
Other Countries	-0.3237	-0.2828	-0.2539	-0.3927	-0.4332
	(0.6260)	(0.6981)	(0.7003)	(0.7610)	(0.4926)
Region	✓	✓	✓	✓	✓
Pre-determined covs		✓	✓	✓	✓
+ GDP pc			✓		
+ Manuf. VA				✓	
+ Total imports					✓
Observations	100	93	93	93	93
R-squared	0.72	0.75	0.77	0.78	0.86

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The estimated regression is `outcome` $\sim$ `c.markdown##i.group` plus controls; each cell reports the implied average marginal effect of the markdown within each group. Columns follow the main tables: column 1 region; column 2 adds pre-determined 1995 controls (log population, human capital) and is preferred; columns 3–5 each add ONE potential bad control to column 2 (GDP per capita; manufacturing VA; total imports). Markdowns follow Eslava et al. (2023).

Income groups from 1995 GDP per capita (PPP), from the Penn World Table: Developed (log GDP pc $\geq$ 9.5), Developing (8.0–9.5), Other ($<$ 8.0). Region fixed effects (A&R regions) included.

Table 25. Log Net Non-Industrial-Tech Imports (Placebo, ex Agriculture) vs. Markdown (Manufacturing)

Markdown	-0.5165*	-0.4874*	-0.0030	-0.1876	0.0856
	(0.2919)	(0.2809)	(0.3231)	(0.3020)	(0.1779)
Region dummies	✓	✓	✓	✓	✓
Pre-determined covs (1995)		✓	✓	✓	✓
+ GDP pc (bad control)			✓		
+ Manuf. VA (bad control)				✓	
+ Total imports (bad control)					✓
Observations	101	94	94	94	94
R-squared	0.53	0.63	0.72	0.69	0.88

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Dependent variable correspond to net imports per one thousand industrial workers, in logs.

Column 2 (region dummies plus pre-determined 1995 controls: log population and human capital) is the preferred specification; column 1 omits the pre-determined controls. Columns 3–5 each add ONE potentially endogenous control to column 2 (GDP per capita, log manufacturing value added, and log total intermediate imports, respectively). For a more detailed description of data sources refer to Appendix Section B.5.

Table 26. Log Net Agricultural-Machinery Imports vs. Markdown (Manufacturing)

Markdown	-0.5949***	-0.5593***	-0.5104*	-0.5666**	-0.3669*
	(0.1989)	(0.2047)	(0.2600)	(0.2323)	(0.2166)
Region dummies	✓	✓	✓	✓	✓
Pre-determined covs (1995)		✓	✓	✓	✓
+ GDP pc (bad control)			✓		
+ Manuf. VA (bad control)				✓	
+ Total imports (bad control)					✓
Observations	102	94	94	94	94
R-squared	0.54	0.64	0.65	0.65	0.69

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Dependent variable correspond to net imports per one thousand industrial workers, in logs.

Column 2 (region dummies plus pre-determined 1995 controls: log population and human capital) is the preferred specification; column 1 omits the pre-determined controls. Columns 3–5 each add ONE potentially endogenous control to column 2 (GDP per capita, log manufacturing value added, and log total intermediate imports, respectively). For a more detailed description of data sources refer to Appendix Section B.5.

Table 27. Marginal Effect of Markdown on Net Automation-Technology Imports per 1,000 Industrial Workers, by Income Group

Markdown					
Developed Countries	-0.6908*	-0.6851*	-0.2465	-0.4972	0.2502
	(0.3822)	(0.3689)	(0.3485)	(0.3311)	(0.2098)
Developing Countries	-0.1164	-0.0857	0.0521	-0.0231	0.1247
	(0.5062)	(0.5118)	(0.5043)	(0.5151)	(0.2039)
Other Countries	-0.2778	-0.3097	-0.2972	-0.3391	-0.4165**
	(0.2358)	(0.2897)	(0.3318)	(0.3289)	(0.1780)
Region	✓	✓	✓	✓	✓
Pre-determined covs		✓	✓	✓	✓
+ GDP pc			✓		
+ Manuf. VA				✓	
+ Total imports					✓
Observations	101	94	94	94	94
R-squared	0.65	0.67	0.69	0.68	0.89

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The estimated regression is `outcome` $\sim$ `c.markdown##i.group` plus controls; each cell reports the implied average marginal effect of the markdown within each group. Columns follow the main tables: column 1 region; column 2 adds pre-determined 1995 controls (log population, human capital) and is preferred; columns 3–5 each add ONE potential bad control to column 2 (GDP per capita; manufacturing VA; total imports). Markdowns follow Eslava et al. (2023).

Income groups from 1995 GDP per capita (PPP), from the Penn World Table: Developed (log GDP pc $\geq$ 9.5), Developing (8.0–9.5), Other ($<$ 8.0). Region fixed effects (A&R regions) included.

Table 28. Marginal Effect of Markdown on Net Computer Imports per 1,000 Industrial Workers (Placebo), by Income Group

Markdown					
Developed Countries	-1.3148**	-1.3744**	-0.7811*	-1.0822**	-0.1762
	(0.5529)	(0.5323)	(0.4260)	(0.4669)	(0.3467)
Developing Countries	0.2894	0.3223	0.5087	0.4197	0.5919
	(0.8269)	(0.8142)	(0.8204)	(0.8143)	(0.4245)
Other Countries	0.0706	-0.0016	0.0153	-0.0473	-0.1384
	(0.2642)	(0.2835)	(0.3264)	(0.3265)	(0.1538)
Region	✓	✓	✓	✓	✓
Pre-determined covs		✓	✓	✓	✓
+ GDP pc			✓		
+ Manuf. VA				✓	
+ Total imports					✓
Observations	101	94	94	94	94
R-squared	0.62	0.67	0.70	0.69	0.87

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The estimated regression is `outcome` $\sim$ `c.markdown##i.group` plus controls; each cell reports the implied average marginal effect of the markdown within each group. Columns follow the main tables: column 1 region; column 2 adds pre-determined 1995 controls (log population, human capital) and is preferred; columns 3–5 each add ONE potential bad control to column 2 (GDP per capita; manufacturing VA; total imports). Markdowns follow Eslava et al. (2023).

Income groups from 1995 GDP per capita (PPP), from the Penn World Table: Developed (log GDP pc ≥ 9.5), Developing (8.0–9.5), Other (< 8.0). Region fixed effects (A&R regions) included.

Table 29. Marginal Effect of Markdown on Net Non-Industrial-Tech Imports per 1,000 Industrial Workers (Placebo, ex Agriculture), by Income Group

Markdown					
Developed Countries	-1.2187**	-1.2582***	-0.7689*	-1.0339**	-0.2382
	(0.5038)	(0.4595)	(0.3910)	(0.3994)	(0.3196)
Developing Countries	0.0960	0.1376	0.2914	0.2124	0.3671
	(0.6466)	(0.6283)	(0.6247)	(0.6284)	(0.2876)
Other Countries	0.0402	-0.0330	-0.0190	-0.0680	-0.1494
	(0.2589)	(0.2840)	(0.3254)	(0.3181)	(0.1597)
Region	✓	✓	✓	✓	✓
Pre-determined covs		✓	✓	✓	✓
+ GDP pc			✓		
+ Manuf. VA				✓	
+ Total imports					✓
Observations	101	94	94	94	94
R-squared	0.66	0.72	0.74	0.73	0.89

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The estimated regression is `outcome` $\sim$ `c.markdown##i.group` plus controls; each cell reports the implied average marginal effect of the markdown within each group. Columns follow the main tables: column 1 region; column 2 adds pre-determined 1995 controls (log population, human capital) and is preferred; columns 3–5 each add ONE potential bad control to column 2 (GDP per capita; manufacturing VA; total imports). Markdowns follow Eslava et al. (2023).

Income groups from 1995 GDP per capita (PPP), from the Penn World Table: Developed (log GDP pc ≥ 9.5), Developing (8.0–9.5), Other (< 8.0). Region fixed effects (A&R regions) included.

Table 30. Marginal Effect of Markdown on Net Agricultural-Machinery Imports per 1,000 Industrial Workers, by Income Group

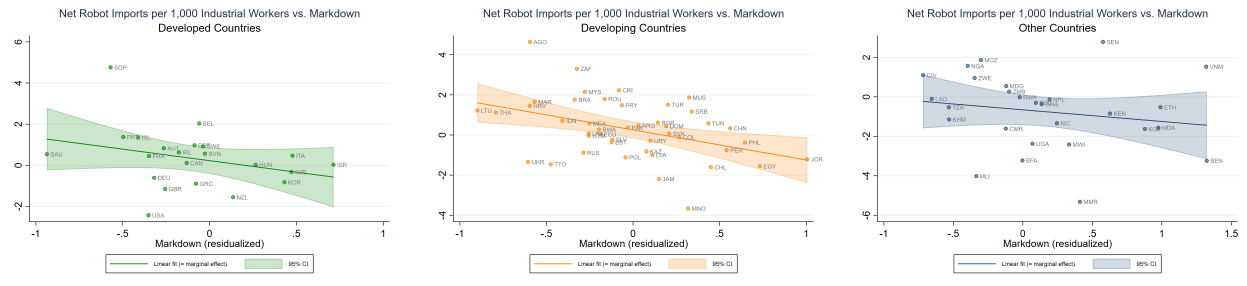
Markdown					
Developed Countries	-1.3927***	-1.3393***	-1.2117**	-1.3659***	-0.8293*
	(0.4988)	(0.4959)	(0.4992)	(0.4990)	(0.4684)
Developing Countries	-0.3287	-0.2463	-0.2062	-0.2551	-0.1315
	(0.4605)	(0.4222)	(0.4407)	(0.4385)	(0.2971)
Other Countries	-0.5372**	-0.6036**	-0.6000*	-0.5994**	-0.6618***
	(0.2381)	(0.2921)	(0.3064)	(0.2861)	(0.2411)
Region	✓	✓	✓	✓	✓
Pre-determined covs		✓	✓	✓	✓
+ GDP pc			✓		
+ Manuf. VA				✓	
+ Total imports					✓
Observations	101	94	94	94	94
R-squared	0.58	0.67	0.67	0.67	0.74

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. The estimated regression is `outcome` $\sim$ `c.markdown##i.group` plus controls; each cell reports the implied average marginal effect of the markdown within each group. Columns follow the main tables: column 1 region; column 2 adds pre-determined 1995 controls (log population, human capital) and is preferred; columns 3–5 each add ONE potential bad control to column 2 (GDP per capita; manufacturing VA; total imports). Markdowns follow Eslava et al. (2023).

Income groups from 1995 GDP per capita (PPP), from the Penn World Table: Developed (log GDP pc ≥ 9.5), Developing (8.0–9.5), Other (< 8.0). Region fixed effects (A&R regions) included.

B.4 Added-Variable Plots: Markdowns and Net Robot Imports

The figures below are added-variable (partial-regression) plots for the preferred specification: net robot imports per one thousand industrial workers and the markdown are each residualized on the preferred 1995 controls, and the slope of the linear fit equals the regression coefficient. The dashed line is a local-polynomial smooth with a 95% confidence band.



(a) Developed **(b)** Developing **(c)** Other

Fig. 10. Net robot imports per 1,000 industrial workers vs. markdown, by income group. Income groups from 1995 GDP per capita (PPP), from the Penn World Table: Developed (log GDP pc $\geq$ 9.5), Developing (8.0–9.5), Other ($<$ 8.0).

B.5 Construction of Author-Built Variables

This appendix lists the variables we construct, beyond those taken directly from the sources above. Occupation classifications use the SOC 2010 major groups; industry classifications use NAICS two-digit sectors.

Commuting-zone level.

- `hhi07_{mp,lw,hg}`: the 2007 commuting-zone vacancy HHI (midpoint, lower, and upper bounds). We average the Choi and Marinescu (2024) occupation×cz-2000 HHI over the four quarters of 2007, weight the six-digit SOC occupations by their national employment in the 2012 OEWS (detailed level, U.S. Bureau of Labor Statistics), and map cz-2000 to cz-1990 with the Autor and Dorn (2013) crosswalk using 2007 County Business Patterns employment as weights (imputed CBP panel of Eckert et al., 2021).

- `efsy_sh_ucov_{core,broad}1990` (also `2000`, `2007`): the share of a commuting zone's employment in industries under-covered in Lightcast, from the imputed CBP panel (Eckert et al., 2021). The "core" set is NAICS {11 Agriculture, 23 Construction}; the "broad" set adds {42 Wholesale, 72 Accommodation and Food, 81 Other Services, 92 Public Administration}. Because CBP excludes the public sector, NAICS 92 carries essentially zero employment, so the broad measure effectively captures {11, 23, 42, 72, 81}. The under-covered classification follows Hershbein and Kahn (2018). Only the broad measure is used in the tables.

- `covg_all`, `covg_t65`, `covg_b35`: employment-weighted occupational coverage of a commuting zone—the fraction of national OEWS-2012 employment in the occupations observed in every quarter of 2007 in the respective commuting zone—computed over all occupations (`all`), over the top-35% most robot-exposed occupations (`t65`), and over the bottom-35% least exposed (`b35`), with robot exposure from Webb (2020). `min_n_socs_{all,t65,b35}` are the corresponding counts of observed six-digit SOC occupations (a coarser, reference measure).

- `good_coverage_all`: indicator equal to one if the commuting zone lies above the 25th percentile of `covg_all` (i.e. not in the worst-covered quartile); used to define

the well-covered subsample in the baseline and mechanism tables.

- `hhi07_*_{exU,exU51,ex0}`: the 2007 HHI recomputed after dropping occupations by Lightcast coverage (Hershbein and Kahn, 2018). `exU` drops the under-covered SOC major groups {25, 33, 35, 37, 39, 45, 47, 53}; `exU51` additionally drops production (SOC 51); `ex0` instead drops the over-covered groups {11, 13, 15, 17, 19, 27, 29}.

- `hhi07_*_at_t65`, `hhi07_*_at_b35`: the 2007 HHI recomputed over the top-35% most robot-exposed and the bottom-35% least robot-exposed occupations, respectively (exposure from Webb, 2020); these are the mechanism regressors.

- `tag_urban`, `tag_industry`, `tag_auto`, `tag_all`: commuting-zone sample indicators for the robustness tables. `tag_urban` keeps zones above the bottom quartile of 1990 urban share (drops the most rural); `tag_industry` keeps zones below the top quartile of employment share in the broad under-covered-industry set; `tag_auto` keeps zones below the top decile of automotive employment share; `tag_all` imposes the three jointly.

Country level. All import measures are accumulated values net of re-exports over 1996–2015 (UN Comtrade, United Nations Statistics Division), in logs, per one thousand industrial workers (industrial employment in 1995 from the ILO, hours-adjusted with PWT 9.0). Technology categories defined and aggregated at the commuting zone level by Acemoglu and Restrepo (2022).

- `ln_net_imports_emp`: log net imports of ISO-defined industrial robots (HS-1996 code 847950) per one thousand industrial workers (primary outcome).

- `ln_auto_emp`: broad industrial automation (Acemoglu and Restrepo (2022) set 1)—robots, automatic welding, numerically-controlled machines, dedicated food and textile machinery, weaving-textile machinery, machine tools, transfer equipment, and control instruments.

- `ln_comp_emp`: computers (placebo, a complementary technology).

- `ln_nonind_emp`: non-industrial technologies excluding agriculture (placebo)—computers, laundry, and vending machinery.

- `ln_agri_emp`: agricultural machinery, isolated because agriculture is itself being automated.
- `markdown_man`: the simple average across WBES waves of the country-year manufacturing markdown estimated as in Eslava et al. (2023).