

Kawaguchi, Daiji; Okazaki, Tetsuji; Zhu, Xuanli

Working Paper

Factory Automation, Labor Demand, and Local Labor Market

IZA Discussion Papers, No. 16885

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Kawaguchi, Daiji; Okazaki, Tetsuji; Zhu, Xuanli (2024) : Factory Automation, Labor Demand, and Local Labor Market, IZA Discussion Papers, No. 16885, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/295908>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 16885

**Factory Automation, Labor Demand, and
Local Labor Market**

Daiji Kawaguchi
Tetsuji Okazaki
Xuanli Zhu

MARCH 2024

DISCUSSION PAPER SERIES

IZA DP No. 16885

Factory Automation, Labor Demand, and Local Labor Market

Daiji Kawaguchi

University of Tokyo, RIETI and IZA

Tetsuji Okazaki

CEPR

Xuanli Zhu

Keio University

MARCH 2024

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

Factory Automation, Labor Demand, and Local Labor Market

This study examines the impact of a technological change on employment and wages, focusing on the adoption of power looms in the silk-weaving industry. Exploiting plant-level panel data from 20th century Japan, we demonstrate that at the plant level, the power loom adaption increased the employment and wages of adult male workers, who likely conducted engineering tasks, and moderately increased wages of female adults, who were simultaneously displaced and reinstated to more non-routine tasks. The wage hike of adult workers induced the exit of less efficient plants and decreased female adult employment by 28 percent at the area level.

JEL Classification: J23, J24, J31, N35

Keywords: electrification, automation, employment, wage, second industrial revolution

Corresponding author:

Daiji Kawaguchi
University of Tokyo
Hongo 7-3-1
Bunkyo-ku
Tokyo
Japan

E-mail: kawaguchi@e.u-tokyo.ac.jp

1 Introduction

The rapid advancement of automation technologies, such as artificial intelligence (AI) and robotics, has surged the debate on how new technologies affect labor demand. Despite the fact that automation has penetrated the entire process of modern technological advance since the Industrial Revolution ([Acemoglu and Restrepo, 2018a, 2019](#); [Johnson and Acemoglu, 2023](#)), most detailed evidences so far has come from the studies focusing on recent decades where abundant micro-level data is available. On the other hand, recent empirical studies to estimate the impact of new technology on employment and wages often face difficulty in finding drastic technological change because technological progress becomes more costly and continuous ([Gordon, 2017](#); [Bloom et al., 2020](#)). Focusing on early stage of economic development overcomes the difficulty due to the abundance of episodes of discontinuous technological change.

Factory automation led by electrification was an epoch among the past technological changes. [Goldin and Katz \(1998\)](#) argue that electrification was the origin of technology-skill complementarity because it was along with an increased demand for skilled workers (e.g. engineers who maintain the machines) relative to unskilled workers. Despite the importance of factory electrification as a historical event, empirical evidence of its impact on labor demand is coarse due to the lack of granular data. [Goldin and Katz \(1998\)](#) found a positive correlation between the capital-labor ratio in 1909 (and 1919) and the fraction of high-school graduate workers in 1940 using the industry-level data. The data limitation did not allow them to correlate the electrification and the contemporaneous change in skill composition. Recently [Atask et al. \(2023\)](#) overcame the limitation by exploiting data from the State of Massachusetts to construct industry-level panel data of the wage distribution and the fraction of electrification in 1895 and 1920. They showed that electrification compressed the lower half of the wage distribution and interpreted that electrification substituted machines for unskilled labor. However, this evidence is not definitive of the effect of electrification on skill demand because the data lack the number of employment and wages by workers' characteristics. Moreover, the recent theoretical development of the task-based framework has shown that instead of a simplified dichotomy of technological change being skill-biased or unskill-biased, it is more about what types of tasks and workers are directly replaced by machines and about the net effect of displacement, productivity improvement, and new task creation induced by automation ([Acemoglu and Restrepo, 2018b, 2019](#)). Furthermore, more recent empirical literature on automation suggests that industry-level analysis often masks the direct effect of technological adoption at firm level since there are spillover effects through market competitions and thus reallocation of labor from non-adopting firms to adopting firms within an industry or a local labor market ([Aghion et al., 2022](#)). Detailed micro-data are necessary to discern the existence and importance

of these entangled effects.

To overcome the limitation of extant studies, this study examines the impact of factory electrification on labor demand using plant-level panel data of the silk-weaving industry from the early 20th century in Fukui Prefecture of Japan. Fukui Prefecture is one of Japan's major centers of silk fabric industry that experienced a transition from hand looms to power looms in the early 1900s. Existing studies demonstrate that the adoption of the power looms replaced old tasks with new tasks and remarkably increased labor productivity by 2.6-2.7 fold (Sanbe, 1961; Okazaki, 2021); With handlooms, a worker operates only one loom with her hands and legs, whereas with power looms, a worker can operate multiple looms (Uchida, 1960; Sanbe, 1961; Tsunoyama, 1983; Hunter, 2003). We study how this drastic transition of production technology affected labor demand.

Our data set is plant-level panel data of about 100 plants from 1904 to 1914 that record the number of employees and average daily wages by gender \times adult/childhood, along with the plant's power source. We exploit the staggered adoption of power loom to estimate the dynamic effects based on the event study design.

The plant-level analysis indicates that, male adult workers, allegedly involved in new engineering tasks, doubled after the adoption of power loom compared with the non-adopting plants in the same area (village and town). In contrast, female adult workers, the main workforce in weaving plants whose previous tasks of operating hand-looms were completely replaced under new technology, had insignificant changes along with power adoption, and the same goes to the child workers who have insufficient capacity adapting to new production mode compared to adults. Furthermore, power loom adoption increases adult males' average wage by 8% and adult females' by 10%, without changing the children's. Our results demonstrate that factory electrification changed the structure of labor demand as claimed by Goldin and Katz (1998) and Atack et al. (2023) with much granular data and credible identification strategy. More importantly, our results demonstrate that the diverged demand changes across different worker types are consistent with the integrated impact of displacement effects, productivity effects, and reinstatement effects as predicted by the recent task-based theories.

The impact of automation on plant-level employment and wage obtained from event study analysis does not necessarily aggregate into market-level results because adoption could potentially generate spillover effects under market competitions in the local labor market (Acemoglu et al., 2022; Aghion et al., 2022). The power-loom adoption by a firm can induce other competitors in the area to raise wage and reduce employment through local labor market competition. To identify the aggregate impact, we examine market-level response by switching the unit of analysis from plant to area. The analysis results show that the area-level automation reduced adult female employment by 28% without affecting the adult male and child employment.

It increased the wages of adult males and females by about 25% and 11% respectively. We also find that area-level electrification reduced the number of plants in an area and induced the concentration of employment to fewer plants.

We set up a theoretical model that combines task-based framework and imperfectly competitive labor market with perfectly competitive product and machine market. The model consistently explains our plant-level and area-level empirical results. Among male adults, the absence of an employment effect at the area level as opposed to the plant level, as well as more than doubled wage increases at the area level compared to the plant level, suggests the presence of significant spillover effects. When a firm adopts the power looms, the demand increase for male adult prompts competitors to reduce employment while bidding up wages. This spillover to the competitors amplifies the plant-level employment effect and attenuates the plant-level wage effect. On the other hand, for adult females, the power adoption moderately increases wages at both plant- and area- level reflecting productivity enhancement. This wage increase induces the exit of the less efficient plants under the pressure of more efficient competitors using the new technology. The attrition of less efficient plants attenuates the firm-level estimate of the impact on employment because the employment trends of adopters and surviving non-adopters are similar after the sample attrition. The exit of factories, at the same time, explains the substantial decrease in employment at the area level.

In sum, the power loom adoption caused a remarkable labor market transition and substantially increased wages among survivors through increased productivity and newly created tasks. However, the effects on employment at firm-level were largely dampened by the attrition of establishments, same as what have been found in the recent literature of modern automation ([Aghion et al., 2022, 2023](#)).

2 The Silk Weaving Industry in Fukui, Japan

The pivot of the Japanese Industrial Revolution, propelled by the Meiji Restoration in 1868 and the subsequent adoption of Western technologies and institutions, was the textile industry that accounted for over 30% of total manufacturing production until the early 1930s. Within the textile industry, the weaving industry contributed to around 40% of the production. The Japanese weaving industry marked a period of profound organizational and technological transformations from the beginning of the 20th century. As a result of these changes, while only 12.3% of weaving workers were employed in factories - defined as plants employing 10 or more workers - in 1905, this figure had more than doubled to 26.7% by 1914 ([Okazaki, 2021](#)). Simultaneously, the industry underwent a remarked technological shift as power looms became increasingly prevalent.

Our analysis focuses on the silk weaving industry in Fukui Prefecture, primarily

chosen due to its relatively simple product market structure and the availability of detailed plant-level data. We focus on *habutae*, a plain silk fabric, predominantly produced in Fukui Prefecture, a region slightly more industrialized than Japan overall in terms of manufacturing share (Kandachi, 1974; Hashino, 2007). Fukui Prefecture was distinguished by its unique manufacturing concentration, making it an ideally clean context for the examination of technological advancement impacts. Over 80% of Fukui’s manufacturing production in the early twentieth century was textiles, with the majority being silk fabric, particularly *habutae*, constituting 71%–74% of the manufacturing output.

The local labor markets of the industry was not perfectly integrated. Kandachi (1974) investigates in detail the source of workers in Harue Village of Sakai County, which was one of the centers of the silk fabric industry in Fukui Prefecture, to find that most workers were from within the village and neighboring villages in the same county (pp.260-261).

The adoption of power looms in Fukui’s *habutae* producers was swift and widespread, especially among those factories. From near zero in 1905, the ratio of power to total looms rose to over 60% in all producers by 1914 (Figure A1). Factors such as the availability of inexpensive domestic power looms and access to electricity have been argued to facilitate this transition.¹ The introduction of power looms led to a significant surge in labor productivity. As detailed by Okazaki (2021), an analysis of plant-level data for the *habutae* industry in Fukui Prefecture indicates a 2.62-fold increase in labor productivity in power loom factories compared to their nonpowered counterparts in 1913–1914, after controlling for plant scale and working hours.²

The introduction of power looms was accompanied by the automation of existing tasks conducted by workers and the creation of new tasks. Weaving, as a process, involves three basic tasks: shuttle manipulation, warp thread regulation, and the beating of weft threads (Bythell, 1969). The transition to power looms mechanized these tasks, freeing workers from the physical constraints of handloom operation and allowing them to manage multiple machines. The remaining tasks for a worker involved halting the loom for thread resupply or thread repair (Uchida, 1960; Sanbe, 1961; Tsunoyama, 1983; Hunter, 2003). Another set of newly created tasks are the installation and maintenance of powered machines. As such, adopting power looms mechanizes existing tasks and generates new tasks, entailing a reallocation of workers with distinct skills into different tasks.

¹Kandachi (1974) documents that many small- and medium-sized land owners founded *habutae* plants and the land rent revenue enabled them to finance investment in power looms.

²This increase aligns with the narratives for this period. Sanbe (1961) notes that while a foot-operated handloom could yield 1.5 tan of silk fabric daily, a power loom could produce 2 tan. Furthermore, a worker who previously could only operate one foot-operated handloom could now manage two to three power looms. These two changes resulted in a total increase of labor productivity by more than 2.67-fold—a figure closely matching Okazaki’s estimation.

3 Data and Summary Statistics

The annual Statistical Yearbook published by Fukui Prefectural Government provides data on factories with ten or more workers for the period spanning 1904 to 1917. It includes specifics such as plant name, location (city, town, or village), owner name, foundation year, major product, power source, total horsepower used, daily working hours, and the number and daily wage of workers along with their gender and age category. We restrict our analysis to the data before 1915 as a revision in the age category that year complicates comparability with earlier data. By doing so, we also avoid the potential distortion caused by the economic boom and inflation during World War I.

We use all the data on plants for which their products are recorded as silk fabric or *habutae*. While the dataset does not explicitly outline details on silk weaving technologies, we can identify those plants with power looms by exploiting the documented information on power source. That is, we regard those plants using inanimate power, water, steam, gas, or electricity, as power loom plants, and the other plants as handloom plants.

We constructed our plant-level panel data by linking individual plants across different years based on plant name, owner’s name, plant address, and foundation year.³ Our data compilation process yielded a dataset of 1,362 distinct plants, constituting 4,470 plant-year observations and spreading across 10 counties (including Fukui city) and 135 distinct areas (towns or villages). To tailor our dataset for the event study analysis, we made a series of exclusions, leading to a focused dataset of 697 plants and 3,231 plant-year observations distributed across 7 counties and 82 areas.⁴ Given the dynamic nature of the industry at that time, there were numerous entries and exits over time, and our resulting panel data is thus unbalanced, with the average observation years of a plant being around 4.6 years.

Table 1 presents the basic statistics derived from our data. A straightforward comparison of the means between powered and non-powered plants reveals that powered plants employed around 50 percent more workers and paid 0.22 log points higher wages. We also scrutinize three distinct worker groups within each plant as outlined in the original statistics: adult males, adult females, and children. Notably,

³Given the age of the data and the potential for documenting inaccuracies, we adopt a fuzzy-matching strategy. Specifically, we regard plants in different years as identical if they share the same plant addresses and at least two of the other three pieces of information—plant name, plant owner’s name, and foundation year—match.

⁴Specific exclusions involved 215 observations from plants that exhibited records of power discontinuation after initial adoption and 285 observations from plants that initially appeared in the dataset already using power looms, thereby lacking pre-treatment data. Additionally, 72 observations from three counties (Nanjo, Onyu, Oi), where the *habutae* industry was relatively underdeveloped, were removed. Lastly, we omitted plants with only one year of observation and other plants that did not align with our econometric specifications for the event study analysis in Section 4. For the area-level analysis discussed in Section 5, we included the plants excluded in this last step before aggregating the plant-level data to area-level data. Our exclusion criteria here also, more or less, aids in addressing potential measurement errors from historical sources.

female adults comprise most of the workforce in silk weaving plants, especially in non-powered ones, while adult male and child workers play comparatively marginal roles. The economic history literature suggests that demographic information can serve as a proxy for worker skill levels (Atack et al., 2004). Consistently, Table 1 reveals that the average daily wages (in sen= 1/100 yen) for adult workers significantly outstrip those for child workers. Perhaps more intriguingly, we observe a wage reversal between adult males and females when comparing powered and non-powered plants; Male adults earn less than their female counterparts in non-powered plants, while the opposite is true in powered plants. This inversion hints at the potential for adult males to possess skills more complementary to the tasks associated with power loom operation, such as installation and maintenance of power looms.⁵

4 Firm-level Analysis

This section examines the impact of the power loom adoption on labor demand and wage structure at the plant level. Specifically, we scrutinize how adopting the power loom affects a plant’s employment and wages by worker categories. To this end, we employ both event study and difference-in-differences (DiD) estimators.

Let Y_{iat} represent our dependent variable for plant i in area a at time t , which can be various plant-level outcomes of interest such as employment or the natural logarithm of the average wages. The specification of our event study analysis is as follows

$$Y_{iat} = \sum_{k=-10}^{-2} \gamma_k 1\{t - G_i = k\} + \sum_{k=0}^5 \gamma_k 1\{t - G_i = k\} + \alpha_i + \delta_{at} + \epsilon_{it}. \quad (1)$$

where G_i is the first year of adopting the motor power by the plant i , α_i represents plant fixed effects, and δ_{at} signifies area-by-year fixed effects. Our coefficients of interest are the lead-lag coefficients γ_k , which capture the dynamic effects pre- and post-event. In the estimation, we normalize γ_{-1} to 0 to serve as the baseline. We validate the causality of this specification through a standard pre-trend falsification test—namely, the absence of pre-trends indicated by statistically insignificant lead coefficients. Given the well-established problem that the estimation of Equation (1) through a two-way fixed effects regression can inadvertently introduce comparisons between treatment groups and undesired control groups, we employ the estimation method proposed by Sun and Abraham (2021) that solves this issue by using never-

⁵Given that our wage data is presented in terms of daily wage, which is susceptible to the influence of variations in working hours, one might be concerned about the simultaneous changes in working hours that may coincide with the adoption of power looms, potentially affecting our results. However, this concern appears to be less warranted as Table 1 shows that the average discrepancy in working hours between non-powered and powered plants is relatively minor—approximately a quarter of an hour—and both categories exhibit small standard deviations.

treated and last-treated cohorts as comparison groups.⁶ In complement to the event study analysis, we also estimate a DiD model using a post-treatment indicator of plant i 's power adoption at time t . This model provides a more parsimonious specification with fewer coefficients estimated and thus improved statistical power at the cost of collapsing all post-event dynamic effects into a singular, permanent effect. For both the event study and DiD analyses, we cluster our standard errors at the plant level.

We start by exploring the impact of power loom adoption on employment. Figure 1 illustrates the event study outcomes for overall employment changes and the heterogeneous effects across different worker categories. We find a significant rise in the demand for male adult workers following the power loom introduction, starting with approximately 1.5 additional workers at event time ($k = 0$), and extending to around 3 workers in succeeding post-event periods. Given the average employment of less than 2 male adult workers in nonpowered plants, this trend signifies a marked upsurge in high-skilled labor demand.⁷ In contrast, for female adult workers, we observe no statistically significant changes immediately after adoption. Despite the coefficients showing an upward trajectory in subsequent periods, culminating in an average addition of 9 female adult workers by period $k = 4$, we advise caution in interpreting these findings. The dynamic effects observed could be attributable to the sample variation across different treatment periods, with $k = 4$ data points likely representing a few resilient producers who successfully navigated the power loom transition. Hence, we posit that the power loom implementation did not substantially increase adult female employment compared with non-powered plant in the same area. For child workers, the analysis yields no statistically significant power adoption impact across all post-treatment periods. As female labor constitutes a significant proportion of the workforce in silk-weaving plants, our overall labor demand findings closely echo the patterns observed for female adult workers—initially insignificant alterations succeeded by some positive increase during later post-treatment periods.

Our DiD analysis, as shown in Table 2, corroborates the event study findings. The coefficient for male adults shows a significant increase of 2.20 workers, which is statistically significant, more than doubling the mean of the control groups at 1.77 workers. The coefficient for female adults is only 1.06 whereas the control mean is

⁶In the Appendix A.2, we further confirm the robustness of our findings by using alternative methods proposed in the literature, including those from Callaway and Sant'Anna (2021); De Chaisemartin and d'Haultfoeuille (2020); Borusyak et al. (2021). Notably, in the specification that follows Callaway and Sant'Anna (2021), we incorporate the never-adopted plants as control groups, as in our benchmark specification, and the not-yet-adopted firms. These alternative estimators generally produce similar results to our baseline findings.

⁷Inoue (1913) reports that in the habutae industry in Fukui, Ishikawa and Toyama Prefectures, there were three types of workers, namely workers for preparation, weavers and mechanics, and that mechanics were "those who took charge of all the maintenance works of weaving machines including lubrication and repair, and this type of workers newly emerged after the adoption of power looms" (pp.84-85, authors' translation).

16.55 and statistically insignificant. The coefficient for children is -1.26 as to the control mean of 4.10, but statistically insignificant.

To summarize, our employment analysis suggest a clear surge in the demand for adult male workers due to the power loom adoption, who possibly worked as engineers tasked with machine installation and maintenance. Conversely, demand for adult female workers, whose previous routine tasks were displaced by power looms but whose productivity more than doubled under the new tasks and new technology, does not show a noticeable increase or decrease. Likewise, the demand for low-skilled child workers does not show an increase and might even suggest a minor decrease.

We next consider the impacts on wages. Figure 2 show the corresponding event study results. The log wages for adult male and female workers show a similar increase after the adoption of power looms. The increase is approximately 0.1 log points and remains stable throughout the post-treatment periods. Crucially for causality, these labor categories exhibit no discernible pre-trends. In contrast, log wages for child workers do not exhibit statistically significant changes, if anything, they show a slight decreasing trend at post one and two periods. Furthermore, the standard deviations for this group are markedly larger. The overall plant-level increases by about 0.1 log points and is statistically significant. Our event study for both employment and wage analysis displays no signs of pre-trends.⁸

Our DiD findings in Table 2 Panel B corroborate these event study results. For the log wages of male adults, and female adults, overall workers, the estimated coefficients are 0.08, 0.10, and 0.09, respectively. In contrast, the coefficient for child log wage is negative at -0.16 and statistically insignificant. Therefore, power loom adoption brought about a wage increase of around 10 percent for both high-skilled male adult workers and medium-skilled female adult workers, but no positive wage effects are discernible for low-skilled child workers.

A note of caution when interpreting our findings is that the effects we've estimated are relative changes compared to the never-treated and last-treated cohorts within an area. As a result, the increased labor demand for male adult workers and the wage increases for male and female adult workers may not correspond to similar

⁸In our robustness tests, we also applied the method proposed by Roth (2022) to examine the linear violation of the parallel trend assumption. We find that, for all the cases where significant treatment effects are found, there are more than one post-treatment period in which the originally estimated confidence intervals deviate from the hypothesized linear trend at 50% pretest power: all post-treatment periods for male employment, and first two post-treatment periods (t=0 and t=1) for wages of adult males, adult females, and overall.

Notably, the lack of a pre-trend does not necessarily imply that the technological adoption is exogenous and devoid of selection. Instead, it suggests that our model specification effectively accounts for the plant heterogeneity inherent in the data-generating process before the treatment. Moreover, while the absence of a pre-trend does not entirely eliminate potential concerns about the effect of power adoption being conflated with other unobserved determinants, it does, in conjunction with our area-by-year control, limit the set of potential confounders to only those shocks that both coincide with power adoption and are not shared by other plants in the same area.

aggregate-level outcomes if the power loom adoption generates spillover effects on non-adopters. Consequently, we conduct area-level analyses in the next section to investigate the aggregate effects of power loom adoption.

5 Area-level Analysis

To investigate the impact of power loom at the area level, we aggregate plant-level employment and take an employment-weighted average of plant wages (i.e. $W_{at} \equiv \sum_{i \in a} s_{it} W_{it}$, where s_{it} is the employment share of plant i in area a in year t). We similarly define an employment-weighted degree of power adoption in an area, E_{at} . We then estimate the following area-level model,

$$Y_{at} = \mu E_{at} + \alpha_a + \delta_t + \epsilon_{at}, \quad (2)$$

where Y_{at} represents the area-level aggregate employment or the natural logarithm of area-level average wages. We control for the area fixed effects, α_a , and the year fixed effects, δ_t . We cluster our standard errors at the area-level.

Table 3 Panel A displays the regression results of employment. The results show that a shift in area power intensity from 0 (no adoption) to 1 (full adoption) does not affect male adults and children, but reduces female adults and overall employment by 28% in a statistically significant way. Since the area power intensity increased from 0 to 60% between 1904 and 1914, the total employment was reduced by 17%.

Table 3 Panel B shows the impact of area electrification on average wages. A unit increase in power intensity lifts log wage by 0.25 log points for male adults, 0.11 log points for female adults, none for children, and 0.11 log points for overall workers.

The combination of employment and wage effects shows the reaction of the local labor market to factory automation. No increase in employment, regardless of the 0.25 log-point increase in wages of male adults, implies that their labor supply is inelastic. This is probably due to the limited supply of engineers in the area. In contrast, for female adults, the combination of the decrease in employment and the increase in wages cannot be explained by a simple shift of labor demand in the local labor market. Instead, automation substantially changed the task content of female adults and increased mean wages among survivors through productivity enhancement and creating new tasks.

At the same time, the increase in the wages in the local labor market may well induce the exit of low-productivity plants and reduce area-level employment. The impacts of area power penetration on the number of plants and the employment concentration reported in Table 3 Panel C bolsters this argument. The transition of no power to full power reduces the the number of plants in the area by about 1.95 plants, where the average number of plants in a non-electrified area is 4.45.

This reduction of the number of plants entailed the concentration of employment. A full electrification increases the Herfindahl-Hirschman Index (HHI) by about 20%. Thus, power-loom adoption reduced aggregate employment in an area.

We next characterize the selection of the plants induced by the penetration of the power loom. For the purpose, we estimate the following equation:

$$Exit_{it+1} = \beta_1 \Delta E_{at+1} + \beta_2 1(W < Med)_{at} + \beta_3 E_{at+1} \cdot 1(W < Med_{at}) + u_{it}, \quad (3)$$

where $Exit_{it+1}$ indicates the disappearance of plant between t and $t + 1$, ΔE_{at+1} is the change in the power penetration between t and $t + 1$, $1(Wage < Med)_{at}$ is the indicator for the plants whose wage is below the median in area a in t , which is a proxy for the less efficient plant. The interaction term is defined using the deviation from the sample mean denoted by \ddot{X} . The results reported in the first column of Panel D shows that the penetration of power loom induces the exit of plants and the lower wage plants are more likely to exit. The impact of power loom adoption on exit is 2.63 times larger for low-wage plants than for high-wage plants. To address the possible reverse causality that plant exit induces the area-level power loom penetration, we try an alternative specification that uses the lagged variable of ΔE_{at} as an explanatory variable. The results reported in the second column indicate that the results are robust in this alternative specification. Overall, our results show that the penetration of new technology accelerates the exit of less efficient plants.

6 Theoretical Explanations on the Firm-level and Area-level Results

To facilitate the interpretation of our estimation results, we build a theoretical model in [Appendix B](#). that integrates the task-based framework developed by [Acemoglu and Restrepo \(2018a,b\)](#) with the oligopsony framework proposed by [Berger et al. \(2022\)](#). Our model assumes local labor markets are not perfectly competitive, but product and machine markets are perfectly competitive. These assumptions are realistic given high commuting/immigration costs and integrated product and machine markets. The model comprehensively explains plant-level and area-level empirical results outlined as follows:

(i) The observed stronger labor demand for male adults, relative to female adults and children at the firm level, is consistent with the model's prediction of a negative displacement effect (see Equation (B12)). New automation technologies replace tasks typically performed by the low skilled female and children, while simultaneously generating a positive productivity effect that benefits male adults and the remaining female adults.

(ii) The muted employment response among adult females coupled with a rise in

their wages at the plant-level can be attributed to the model’s postulation of inelastic intra-market labor supply for this group (small η_L in Equation (B13)). In contrast, the significant employment growth and moderate wage escalation for adult males at the plant-level are likely to be a joint result of their more elastic intra-market labor supply and the spillover effects described next.

(iii) Assuming greater intra-market labor supply elasticity compared to inter-market elasticity, the model predicts that increased labor demand from an adopting plant will trigger both a decline in competitors’ employment and an increase in their wages ($\eta_H - \phi_H > 0$ in Equation (B14)). This result thus explains why we observe a stronger employment rise relative to wage increase for adult males at the plant-level and accounts for why the employment effect becomes more subdued and wage increase is larger at the area-level.

(iv) At the area-level, the model can also explain the negative employment impacts, particularly for female adults, as a consequence of the aforementioned spillover effects. The wage hike due to the spillover effect leads to the exit of the less efficient firms and yields a more concentrated market structure (see Equation (B15) and the description).

The consistency between the model’s predictions and the empirical findings demonstrates that the local labor market of our study can be effectively characterized by the task-based framework’s representation of technological evolution and the oligopsony model’s depiction of labor market dynamics.

7 Conclusion

This study examined the impact of automation of silk-weaving plants brought by electrification on employment and wages by demographic groups using plant-level panel data from early 20th-century of Japan. The plant- and area-level evidence picture a remarkable dynamism brought to the local labor market by automation. The technological change was skill-biased, as shown by the substantial increase in employment and wages for adult male workers who were capable in engineering jobs. The automation technology was also displacing and destructive at least in the short-run—the area-level employment for adult female workers decreased by 28 percent despite their productivity being doubled. This negative aggregate effect mainly stemmed from the increased local wages induced by the fast diffusion of new technologies and the exit of less efficient plants. Perhaps interestingly, our empirical findings about the relatively rudimentary technologies during the period of factory electrification largely resonate with evidence newly discovered by the emerging literature on more recent automation technologies (Aghion et al., 2022). This indicates that there may have been surprisingly few changes in the impact of technological change on the labor market over the past century and that historical

episodes could still be useful in shedding light on the future literature.

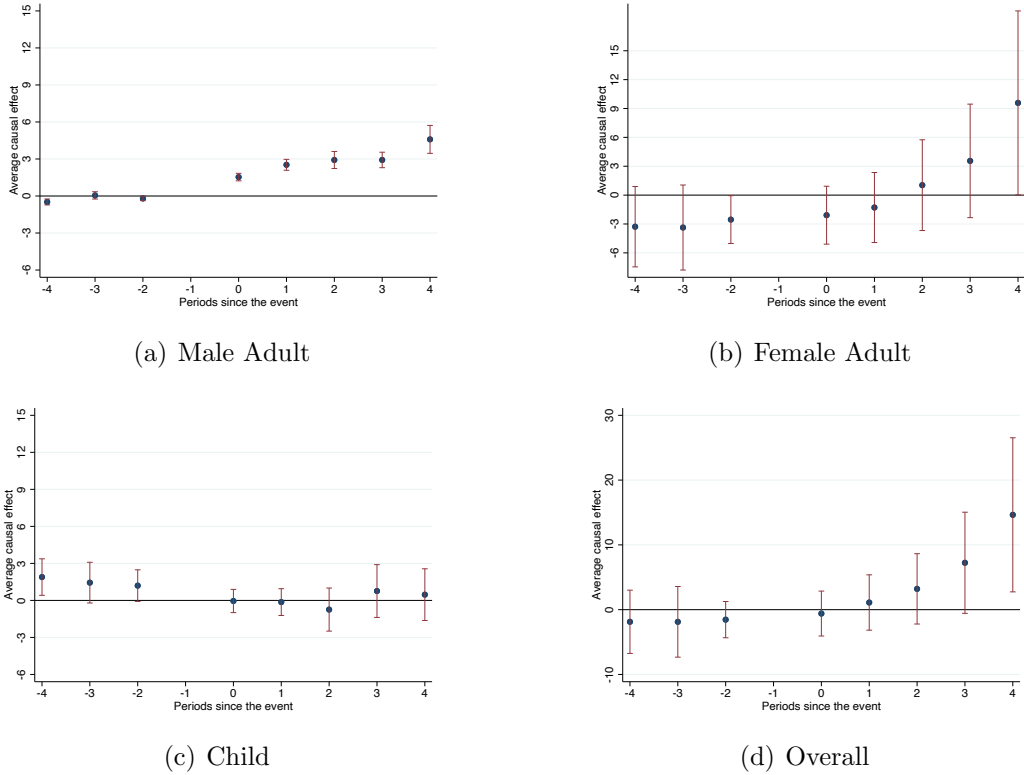
References

- Acemoglu, Daron and Pascual Restrepo**, “Artificial intelligence, automation, and work,” in “The Economics of Artificial Intelligence: An Agenda,” University of Chicago Press, 2018, pp. 197–236.
- **and** –, “The race between man and machine: Implications of technology for growth, factor shares, and employment,” *American Economic Review*, 2018, *108* (6), 1488–1542.
- **and** –, “Automation and new tasks: How technology displaces and reinstates labor,” *Journal of Economic Perspectives*, 2019, *33* (2), 3–30.
- , **Claire Lelarge, and Pascual Restrepo**, “Competing with robots: Firm-level evidence from France,” in “AEA papers and proceedings,” Vol. 110 American Economic Association 2014 Broadway, Suite 305, Nashville, TN 37203 2020, pp. 383–388.
- , **Gary W Anderson, David N Beede, Cathy Buffington, Eric E Childress, Emin Dinlersoz, Lucia S Foster, Nathan Goldschlag, John C Haltiwanger, Zachary Kroff et al.**, “Automation and the workforce: A firm-level view from the 2019 Annual Business Survey,” Technical Report, National Bureau of Economic Research 2022.
- Aghion, Philippe, Céline Antonin, Simon Bunel, and Xavier Jaravel**, “The effects of automation on labor demand: A survey of the recent literature,” in “Robots and AI,” Routledge, 2022.
- , – , – , **and** –, “Modern manufacturing capital, labor demand, and product market dynamics: Evidence from France,” Technical Report, CEPR Discussion Paper No. 1910 2023.
- Atack, Jeremy, Fred Bateman, and Robert A Margo**, “Skill Intensity and Rising Wage Dispersion in Nineteenth-Century American Manufacturing,” *Journal of Economic History*, 2004, *64* (1), 172–92.
- , **Robert A Margo, and Paul Rhode**, “Wage Inequality in American Manufacturing, 1820-1940: New Evidence,” Technical Report, National Bureau of Economic Research 2023.
- Berger, David, Kyle Herkenhoff, and Simon Mongey**, “Labor market power,” *American Economic Review*, 2022, *112* (4), 1147–1193.
- Bloom, Nicholas, Charles I Jones, John Van Reenen, and Michael Webb**, “Are ideas getting harder to find?,” *American Economic Review*, 2020, *110* (4), 1104–1144.

- Borusyak, Kirill, Xavier Jaravel, and Jann Spiess**, “Revisiting event study designs: Robust and efficient estimation,” Technical Report, arXiv preprint arXiv:2108.12419 2021.
- Bythell, Duncan**, *The Handloom Weavers: A Study in the Cotton Industry during the Industrial Revolution*, Cambridge University Press, 1969.
- Callaway, Brantly and Pedro HC Sant’Anna**, “Difference-in-differences with multiple time periods,” *Journal of Econometrics*, 2021, 225 (2), 200–230.
- Chaisemartin, Clément De and Xavier d’Haultfoeuille**, “Two-way fixed effects estimators with heterogeneous treatment effects,” *American Economic Review*, 2020, 110 (9), 2964–2996.
- Goldin, Claudia and Lawrence F Katz**, “The origins of technology-skill complementarity,” *Quarterly Journal of Economics*, 1998, 113 (3), 693–732.
- Gordon, Robert**, *The rise and fall of American growth: The US standard of living since the Civil War*, Princeton University Press, 2017.
- Hashino, Tomoko**, “The Rise of the Power-Loom Factory in the Traditional Silk-Weaving District: Change in Demand and Labor Market,” in Tetsuji Okazaki, ed., *Production Organizations in Japanese Economic Development*, Abingdon: Routledge, 2007, pp. 14–38.
- Hunter, Janet**, *Women and the Labour Market in Japan’s Industrializing Economy: The Textile Industry before the Pacific War*, London: RoutledgeCurzon, 2003.
- Inoue, Tokunosuke**, *Yushutsu Habutae [Export Habutae]*, Tokyo: Dobunkan, 1913. in Japanese.
- Johnson, Simon and Daron Acemoglu**, *Power and Progress: Our Thousand-Year Struggle Over Technology and Prosperity*, Hachette UK, 2023.
- Kandachi, Haruki**, *Meiji Ki Noson Orimonogyo no Tenkai [Development of Rural Weaving Industry in Meiji Era]*, Tokyo: The University of Tokyo Press, 1974. in Japanese.
- Koch, Michael, Ilya Manuylov, and Marcel Smolka**, “Robots and firms,” *The Economic Journal*, 2021, 131 (638), 2553–2584.
- Okazaki, Tetsuji**, “Disentangling the effects of technological and organizational changes during the rise of the factory: the case of the Japanese weaving industry, 1905- 14,” *The Economic History Review*, 2021, 74 (4), 976–1005.

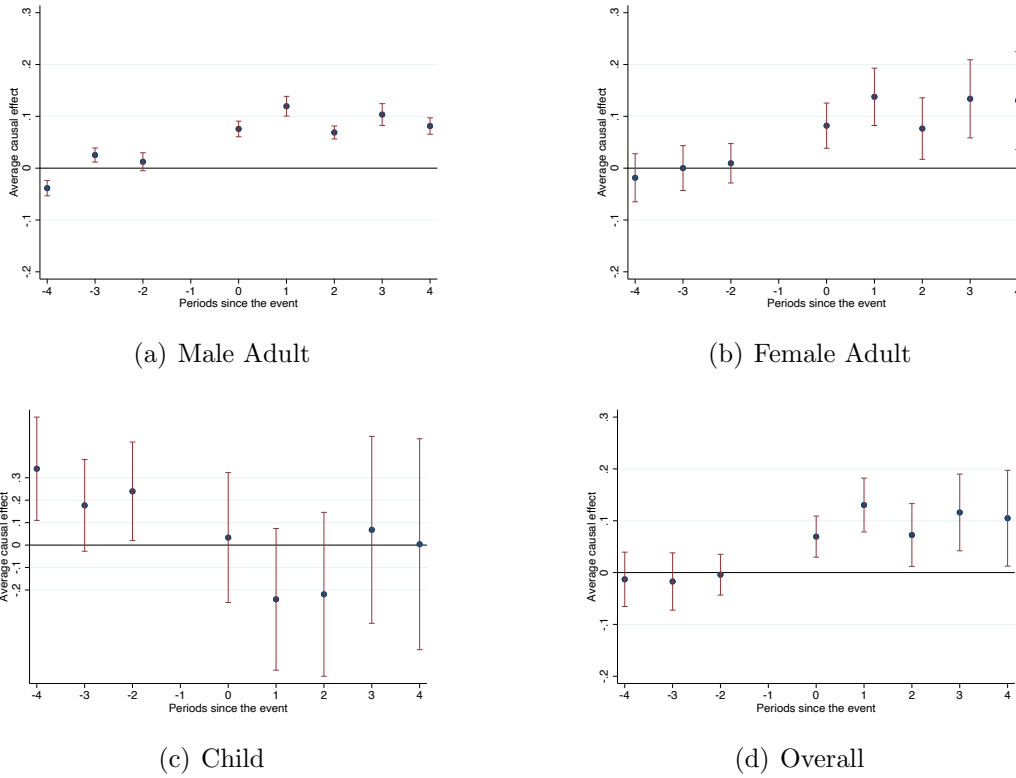
- Roth, Jonathan**, “Pretest with caution: Event-study estimates after testing for parallel trends,” *American Economic Review: Insights*, 2022, 4 (3), 305–322.
- Sanbe, Takako**, *Nihon Kigyo Shi (History of the Japanese Weaving Industry)*, Tokyo: Yuzankaku, 1961. in Japanese.
- Sun, Liyang and Sarah Abraham**, “Estimating dynamic treatment effects in event studies with heterogeneous treatment effects,” *Journal of Econometrics*, 2021, 225 (2), 175–199.
- Tsunoyama, Yukihiro**, “Nihon no Shokki” [Looms in Japan],” in Keiji Nagahara and Keiji Yamaguchi, eds., *Boshoku [Spinning and Weaving]*, Nihon Hyoron-sha, 1983, pp. 17–104. in Japanese.
- Uchida, Hoshimi**, *Nihon Boshoku Gijutsu no Rekishi [History of Spinning and Weaving Technologies in Japan]*, Tokyo: Chijin Shikan, 1960. in Japanese.

Figure 1: The Plant-Level Impacts of Power Adoption on Employment



Note: This figure reports the results of plant-level event studies on employment across different worker categories and overall. In particular, the dot plots are the estimated γ_k in Equation (1), i.e. the coefficients for the lead and lag event-time dummies, and the error bar indicates the 95% confidence intervals based on the standard errors clustered at the plant level. Both plant fixed effects and area-by-year fixed effects are controlled. The estimation follows the method proposed by [Sun and Abraham \(2021\)](#) and uses never-treated and last-treated cohorts as control groups.

Figure 2: The Plant-Level Impacts of Power Adoption on $\ln(\text{Wage})$



Note: This figure reports the results of plant-level event studies on employment across different worker categories and overall. In particular, the dot plots are the estimated γ_k in Equation (1), i.e. the coefficients for the lead and lag event-time dummies, and the error bar indicates the 95% confidence intervals based on the standard errors clustered at the plant level. Both plant fixed effects and area-by-year fixed effects are controlled. The estimation follows the method proposed by [Sun and Abraham \(2021\)](#) and uses never-treated and last-treated cohorts as control groups.

Table 1: Summary Statistics

	NonPowered Plants (Mean)	Powered Plants (Mean)	Powered - NonPowered
Total Worker Per Plant	22.42 [17.69]	34.14 [31.92]	11.71 (1.01)
- Male Adult Worker	1.77 [2.55]	4.69 [5.77]	2.93 (0.16)
- Female Adult Worker	16.55 [13.56]	26.91 [25.38]	10.36 (0.78)
- Child Worker	4.10 [6.41]	2.53 [5.35]	-1.57 (0.31)
Work Hour Per Day	11.46 [1.35]	11.68 [1.17]	0.22 (0.07)
Average Daily Wage Per Plant (Sen)	16.87 [4.14]	23.60 [4.20]	6.73 (0.20)
- Male Adult Worker	14.31 [11.47]	26.51 [11.17]	12.20 (0.56)
- Female Adult Worker	17.60 [4.03]	23.71 [4.22]	6.11 (0.20)
- Child Worker	1.85 [2.14]	0.91 [1.34]	-0.94 (0.10)
Observations	2,743	488	3,231

Note: Means are reported. Standard deviations are reported in square brackets; Standard errors are reported in parentheses

Table 2: The Plant-Level Effect of Power Introduction on Employment and Wages

	(1) Male Adult	(2) Female Adult	(3) Child	(4) Overall
Panel A: Effect on employment				
Power	2.202 (0.697)	1.059 (1.954)	-1.264 (0.757)	1.997 (2.365)
Control Means	1.77	16.55	4.10	22.42
N	3,231	3,231	3,231	3,231
Panel B: Effect on Ln(Wages)				
Power	0.084 (0.021)	0.096 (0.017)	-0.163 (0.123)	0.092 (0.018)
Control Means	3.15	2.94	0.85	2.89
N	2,001	3,220	1,882	3,231

Notes: This table reports the regression coefficients of plant-level employment (Panel A) or the natural logarithm of average wages (Panel B) on the indicator variable for adopting power. The unit of observation is plant. Clustering robust standard errors against the plant-level correlations are reported in parentheses. All specifications include firm and area \times year fixed effects.

Table 3: The Area-Level Effect of Power Introduction on Employment and Wages


	(1)	(2)	(3)	(4)
	Male Adult	Female Adult	Child	Overall
Panel A: Effect on Employment				
Area Power Intensity	-1.173	-27.921	-1.650	-30.744
	(1.975)	(7.103)	(3.856)	(8.640)
Control Means	7.70	75.91	19.19	102.81
N	849	849	849	849
Panel B: Effect on $\ln(\text{Wage})$				
Area Power Intensity	0.251	0.109	0.037	0.113
	(0.070)	(0.031)	(0.162)	(0.035)
N	679	849	605	849
Panel C: Effect on Market Structure				
	# of Plants	HHI		
Area Power Intensity	-1.951	0.122		
	(0.526)	(0.044)		
Control Means	4.45	0.49		
N	849	849		
Panel D: Effect on Plant Exit between t and $t + 1$				
Δ Area Power Intensity	0.140	0.140		
	(0.051)	(0.061)		
Wage < Median	0.026	0.031		
	(0.013)	(0.015)		
Δ Area Power Intensity	0.229	0.181		
\times Wage < Median	(0.110)	(0.105)		
Δ Area Power Intensity	$t, t + 1$	$t - 1, t$		
N	2,826	2,582		

Notes: This table reports the regression coefficients of area aggregate employment (Panel A), employment weighted average of the natural logarithm of average wages (Panel B), or area-level market structure (Panel C) on the employment weight averaged indicator variable for adopting power. The unit of observations is area \times year. Clustering robust standard errors robust against the area-level correlations are reported in parentheses. All specifications include area and year fixed effects. Panel D reports the results of regressing an plant-level exit dummy at period $t+1$ (where exit is defined by the first year a plant disappeared in our data set) on the first difference of area power intensity, the dummy variable indicating that the plant wage is below the area-level median wage, the interaction between the mean deviations of two variables. The second column reports the results of the specification with the lagged first difference of area power intensity. All the observations in the first and the last year of panel data are dropped.

Appendix A. Additional Figures and Tables


A.1 Historical Background and Dataset

Figure A1: Trend of Power Adoption




A.2 Robustness on Event Study Estimation

Figure A2: Comparison of estimators for plant-level event-study estimation (employment)



Note: This figure reports the robustness check of the plant-level event studies on employment in Figure 1 under different estimators. In particular, we test with five estimators that has been used the literature: the OLS estimator, the [Sun and Abraham \(2021\)](#) estimator (the baseline one used in the main text), the [Callaway and Sant'Anna \(2021\)](#) estimator, the [De Chaisemartin and d'Haultfoeulle \(2020\)](#) estimator, and the estimator in [Borusyak et al. \(2021\)](#). In the case using [Callaway and Sant'Anna \(2021\)](#) estimator, we include also the not-yet-adopted firms in the control group, in addition to the never-treated or last-treated plants used in our baseline estimation. Since the estimators of [Callaway and Sant'Anna \(2021\)](#), [De Chaisemartin and d'Haultfoeulle \(2020\)](#), and [Borusyak et al. \(2021\)](#) require more data for statistical power, we replace the area-by-year fixed effects used in our main text with the county-by-year fixed effects, and cluster the standard errors at county level.

Figure A3: Comparison of estimators for plant-level event-study estimation (log wage)



Note: This figure reports the robustness check of the plant-level event studies on log wage in Figure 2 under different estimators. In particular, we test with five estimators that has been used the literature: the OLS estimator, the Sun and Abraham (2021) estimator (the baseline one used in the main text), the Callaway and Sant’Anna (2021) estimator, the De Chaisemartin and d’Haultfoeulle (2020) estimator, and the estimator in Borusyak et al. (2021). In the case using Callaway and Sant’Anna (2021) estimator, we include also the not-yet-adopted firms in the control group, in addition to the never-treated or last-treated plants used in our baseline estimation. Since the estimators of Callaway and Sant’Anna (2021), De Chaisemartin and d’Haultfoeulle (2020), and Borusyak et al. (2021) require more data for statistical power, we replace the area-by-year fixed effects used in our main text with the county-by-year fixed effects, and cluster the standard errors at county level. The estimated results under the Borusyak et al. (2021) estimator is not obtained and reported for the child group given its small sample size.

Appendix B. Theoretical Framework

In this section, we construct a theoretical framework to better interpret our empirical findings in the main text, by integrating the task-based framework of [Acemoglu and Restrepo \(2018a,b\)](#) with the oligopsony framework proposed by [Berger et al. \(2022\)](#). We assume local labor markets are imperfectly competitive but the product and machine markets are perfectly competitive. Our focus on an oligopsony setting in the local labor market instead of a monopoly or oligopoly one (see, e.g., [Acemoglu et al. \(2020\)](#); [Koch et al. \(2021\)](#)) is motivated by the fact that the silk-weaving plants under our study predominantly produced one homogenous raw good, *habutae*, and that wage dispersed significantly for different plants in our data. With a perfectly competitive labor market, which is likely to be infeasible especially in our historical context, plants would react to any technological or demand shocks exclusively by adjusting employment levels, leaving wages unaffected. Nevertheless, our model retains the business stealing effect highlighted in [Acemoglu et al. \(2020\)](#) and [Aghion et al. \(2022\)](#), manifested through the competition in the local labor market.

Production. Consider a firm i located in a local labor market j populated with number of firms n_j . All firms in this market produce one homogeneous good with its price normalized to one. As suggested by the task-based literature, the production is achieved by completing a set of different tasks ranging from $N - 1$ to N :

$$\ln Y_i = \ln z_i \int_{N-1}^N \ln y_i(x) dx, \quad (\text{B1})$$

where Y_i represents total production of the good, $y(i)$ denotes task-level production, and z_i stands for firm specific productivity. Each task is produced according to the following technological regime:

$$y_i(x) = \begin{cases} \gamma_l(x)l_i(x) + \gamma_m(x)m_i(x) & \text{if } x \in [N - 1, I] \\ \gamma_l(x)l_i(x) & \text{if } x \in (I, I') \\ \gamma_h(x)h_i(x) & \text{if } x \in [I', N], \end{cases} \quad (\text{B2})$$

where I and I' represent technological and skill thresholds, and γ are continuous functions indicating the productivity of three inputs: machine m , low-skilled labor l , and high-skilled labor h , across different tasks. Tasks lying between $N - 1$ and I can be carried out either by low-skilled labor l or machinery m in a perfectly substitutive manner. Beyond I , a task can be only produced by human labor, and thus I implies a technological constraint on current automation technologies. An additional constraint in Equation (B2) is that low-skilled worker and high-skilled worker conduct separate tasks with the boundary defined by I' . In our case, this constraint could arise from either distinct comparative advantages between men and

women or social and cultural norms that dictated the roles of men and women during that historical era. As a result, new automation technologies modeled as an increase in I would directly replace low-skilled labor l but not high-skilled labor h .⁹

To simplify, we assume that machinery m is competitively supplied by external producers at a fixed rate R . In contrast, both types of labor, l and h , are supplied elastically in the local labor market.

Household and local labor market. We assume that the representative household in the local labor market j faces the following problem,

$$\begin{aligned} & \max_{C_i, L_i, H_i} U_j \left(\mathbf{C} - \frac{\mathbf{L}^{\frac{\phi_L+1}{\phi_L}}}{\frac{\phi_L+1}{\phi_L}} - \frac{\mathbf{H}^{\frac{\phi_H+1}{\phi_H}}}{\frac{\phi_H+1}{\phi_H}} \right) \\ \text{s.t. } & \mathbf{C} = \sum_{i \in j} W_{iL} L_i + \sum_{i \in j} W_{iH} H_i + \Pi_j, \mathbf{L} = \left(\sum_{i \in j} L_i^{\frac{\eta_L+1}{\eta_L}} \right)^{\frac{\eta_L}{\eta_L+1}}, \text{ and } \mathbf{H} = \left(\sum_{i \in j} H_i^{\frac{\eta_H+1}{\eta_H}} \right)^{\frac{\eta_H}{\eta_H+1}}, \end{aligned} \quad (\text{B3})$$

where $L_i = \int_{N-1}^N l_i(x) dx$ and $H_i = \int_{N-1}^N h_i(x) dx$ are the firm-level labor inputs, and Π_j is the aggregated firm profits in location j . The use of the aggregate indexes \mathbf{L} and \mathbf{H} , which do not correspond to any real aggregates, is a convenient way to model the oligopolistic competition between m_j firms in the local labor market j . [Berger et al. \(2022\)](#) shows that this supply system can be derived from a microfoundation of heterogeneous agents making discrete job choices on heterogeneous firms. The elasticity parameters $\eta_L > 0$ and $\eta_H > 0$ captures the extent of competition in the local labor market, similar to the elasticity of substitution in a monopoly or oligopoly setting. In other words, η_L and η_H capture the degree of differentiation among employers within an area. The larger the η_L and η_H are, the less the employers are differentiated, and the more competitive are the local labor markets. In the extreme case where $\eta_L \rightarrow \infty$ or $\eta_H \rightarrow \infty$, the local labor market tends to perfect competition and marginal products are equalized at a single wage in an area. In addition, parameters $\phi_L > 0$ and $\phi_H > 0$ capture the labor supply elasticities at the market level, which could potentially reflect either the household's trade-off between work and leisure or housework, or the labor market competition across different locations, or both. As we will not explicitly model the between-market competition, hereafter we omit the subscript j and focus on the analysis of a particular labor market. Solving the household problem gives us the labor supply curve that firms

⁹While we follow the convention and use the terms “low-skilled” and “high-skilled,” one can equally use alternative terms to distinguish these two types of labor as “routine” and “non-routine,” or more accurately, “displaceable” and “non-displaceable.” The essential distinction lies in the nature that the tasks conducted by one type of labor to be supplanted by automation in the impending technological advance while the tasks of the other type remain unaffected or complemented. In fact, under certain conditions, we can have “low-skilled” labor being paid even higher wage than “high-skilled” labor before the introduction of the new automation technology.

face:

$$W_{iS} = \mathbf{S}^{\frac{1}{\phi_S} - \frac{1}{\eta_S}} S_i^{\frac{1}{\eta_S}} \text{ for } S \in \{H, L\}, \quad (\text{B4})$$

where W_{iS} is the wage of labor S setting by firm i .

Characterization. We now solve the firm's optimization problem in two steps and then characterize the market equilibrium.

Firstly, given machine price R and the firm-level optimal choices of wages and input uses, firms optimally allocate resources into the production of different tasks. Since low-skilled labor and machine are perfect substitutes for tasks between $N - 1$ and I , a firm's optimal input choices on these tasks are determined by comparing the marginal rate of substitution, $\gamma_l(x)/\gamma_m(x)$, with the ratio of marginal costs, mc_{iL}/mc_m . Given our setting, the marginal cost for machine, mc_m , is just the machine price R , while the marginal cost for low-skilled labor, mc_l , is larger than the firm wage, W_{iL} , due to the fact that firms face an upward-sloping labor supply curve under the existence of monopsony. In particular, we have $mc_{iL} = W_{iL}/\mu_{iL}$, where $\mu_{iL} \in [0, 1]$ represents the firm-specific markdown on labor l , with the expression we will derive below. If $\frac{mc_{iL}}{R} > \frac{\gamma_l(I)}{\gamma_m(I)}$, firm i 's input choice is technologically bounded, i.e. although using machines for tasks beyond I can be potentially more productive or cost-saving than using human labor, such technology is currently unavailable. Otherwise, firm i is not technologically constrained and will choose an interior threshold $I_i^* < I$. In the oligopolistic case where more productive firms will employ more workers, pay higher wages, and have lower markdowns, as we will show below, our model suggests that large firms, faced with higher marginal costs on labor, is more likely to be technologically constrained. Consequently, our model can predict that larger firms are more likely to adopt new automation technologies following a technological breakthrough, which is consistent with our data. However, to ease the analysis, we abstract from any ex-ante difference in the technological thresholds among firms by assuming $\frac{mc_{iL}}{R} > \frac{\gamma_l(I)}{\gamma_m(I)} \forall i$ (A1).¹⁰ In other words, we assume that, prior to the coming of new power loom technology, all firms in our case had their input choices technological bounded, that is, $I_i^* = I \forall i$, and thus the adoption of the new automation technology (an increase in I) will induce increased use of machinery and enhanced production efficiency.

Given the same input prices across different tasks at firm level, a firm will equalize the marginal product for tasks that utilize identical inputs (i.e. m , l , or h). Under our Cobb-Douglas form of task aggregation, this results the same amount of inputs

¹⁰While the study of endogenous technological adoption in our framework is itself interesting, it complicates our analysis by allowing additional adjustment on the endogenous technological threshold, and adds little additional insights into our primary objective—assessing the technological impact on labor demand. One interesting feature when we combine endogenous technology adoption and oligopsony framework is that the adoption by one firm could potentially increase the wages of other firms through oligopsonistic competition and thus force them to follow up in adoption, generating technological diffusion.

being used across tasks utilizing the same input type. Thus we can rewrite the firm production function in Equation (B1) as a function of firm level input uses (M_i , L_i , and H_i):

$$Y_i = B_i \left(\frac{M_i}{I - N + 1} \right)^{I - N + 1} \left(\frac{L_i}{I' - I} \right)^{I' - I} \left(\frac{H_i}{N - I'} \right)^{N - I'}, \quad (\text{B5})$$

$$\text{where } B_i = z_i \exp \left(\int_{N-1}^I \ln \gamma_m(x) + \int_I^{I'} \ln \gamma_l(x) + \int_{I'}^N \ln \gamma_h(x) dx \right).$$

As a typical result of the task-based framework, the technological threshold I directly enters the share term of the input that can be substituted (here L), leading to a direct displacement effect under advancement in automation technology. In comparison, factor-augmenting technological changes, represented by increasing in γ terms, only affect the factor-neutral productivity, B_i , and thus always result in a positive effect on labor demand.

Equation (B5) provides us the marginal product for each type of input use at firm level, through which we can link to the labor supply side and characterize the firm optimal choices. In particular, our second step is to solve the firm problem of profit maximization:

$$\Pi_i = \max_{H_i, L_i, M_i} Y_i - W_{iH} H_i - W_{iL} L_i - R M_i \quad (\text{B6})$$

$$\text{s.t. } W_{iS} (S_i, S_{-i}^*) = \mathbf{S}^{\frac{1}{\phi_S} - \frac{1}{\eta_S}} S_i^{\frac{1}{\eta_S}} \text{ and } \mathbf{S} (S_i, S_{-i}^*) = \left[S_i^{\frac{\eta_S + 1}{\eta_S}} + \sum_{k \neq i} S_k^{*\frac{\eta_S + 1}{\eta_S}} \right]^{\frac{\eta_S}{\eta_S + 1}} \text{ for } S \in \{H, L\}.$$

Here the firm takes the actions of its competitors as given and a Nash equilibrium is achieved when all firms in the local labor market are making their optimal choices.

The first order conditions are

$$\frac{\partial Y_i}{\partial M_i} = R \quad (\text{B7})$$

$$\underbrace{\frac{\partial Y_i}{\partial S_i}}_{\text{Marginal product: } mp_{Si}} = \underbrace{W_{iS} + \frac{\partial W_{iS}}{\partial S_i} \Big|_{S_{-i}^*}}_{\text{Marginal cost: } mc_{Si}} S_i \text{ for } S \in \{H, L\}. \quad (\text{B8})$$

Following the derivation in Berger et al. (2022), we can rewrite Equation (B8) as

$$mp_{iS} = W_{iS} / \mu_{iS}, \text{ where } \mu_{iS} = \frac{\varepsilon_{iS}}{\varepsilon_{iS} + 1},$$

$$\varepsilon_{iS} := \left[\frac{\partial \ln W_{iS}}{\partial \ln S_i} \Big|_{S_{-i}} \right]^{-1} = \left[(1 - e_{iS}) \frac{1}{\eta_S} + e_{iS} \frac{1}{\phi_S} \right]^{-1} \text{ and } e_{iS} = \frac{W_{iS} S_i}{\sum_i W_{iS} S_i}. \quad (\text{B9})$$

Here, μ_{iS} denotes the markdown of firm i on input $S \in \{H, L\}$, ε_{iS} is the inverse of the firm's wage elasticity of labor supply on S , and e_{iS} denotes the firm's share of input S 's wage bill in the local labor market. In the case where $\eta_S > \phi_S$, i.e. the substitution between firms within the market is more elastic than the market-level

labor supply elasticity, firms with higher marginal products will offer higher wages, employ more workers, attain a larger share of the labor market, and end up facing a less elastic labor supply curve and experiencing a lower markdown (i.e. greater market power).

Our final step of characterization is to integrate the labor demand and labor supply sides by using wage rate, W_{iS} . In particular, from Equations (B4), (B5) and (B9), we have

$$\begin{aligned} W_{iL} &= \mu_{iL} mp_{Li} = \mu_{iL}(I' - I)Y_i/L_i \text{ (labor demand of } L) \\ W_{iL}(L_i, L_{-i}^*) &= \mathbf{L}^{\frac{1}{\phi_L} - \frac{1}{\eta_L}} L_i^{\frac{1}{\eta_L}} \text{ (labor supply of } L), \end{aligned} \quad (\text{B10})$$

and

$$\begin{aligned} W_{iH} &= \mu_{iH} mp_{Hi} = \mu_{iH}(N - I')Y_i/H_i \text{ (labor demand of } H) \\ W_{iH}(H_i, H_{-i}^*) &= \mathbf{H}^{\frac{1}{\phi_H} - \frac{1}{\eta_H}} H_i^{\frac{1}{\eta_H}} \text{ (labor supply of } H). \end{aligned} \quad (\text{B11})$$

The market equilibrium for this local economy is thus defined as a set of input uses $\{M_i, L_i, H_i\}_{i \in j}$ and firm-specific wages $\{W_{iL}, W_{iH}\}_{i \in j}$ such that, given the machine price R , the Equations (B7), (B10) and (B11) are satisfied for each firm i in the local labor market j .

Technological impact. With above framework in hand, we are now poised to examine the influence of new technology adoption on both a firm's own labor demand and that of its competitors, providing context for interpreting our empirical results. Specifically, we characterize a technology adoption event as an increase in the I for a particular firm i while keeping other competing firms' I unchanged. We first study the own effect without considering any interactions from labor market competition. Using the labor demand equations in Equations (B10) and (B11), we have

$$\begin{aligned} \frac{d \ln W_{iL}}{dI} &= \underbrace{\frac{d \ln \mu_{iL}}{dI}}_{\text{Markdown effect } \leq 0} + \underbrace{\frac{d \ln(I' - I)}{dI}}_{\text{Displacement effect } < 0} + \underbrace{\frac{d \ln(Y_i/L_i)}{dI}}_{\text{Productivity effect } > 0} \\ \frac{d \ln W_{iH}}{dI} &= \underbrace{\frac{d \ln \mu_{iH}}{dI}}_{\text{Markdown effect } \leq 0} + \underbrace{\frac{d \ln(Y_i/H_i)}{dI}}_{\text{Productivity effect } > 0}. \end{aligned} \quad (\text{B12})$$

If both types of labor are perfectly inelastic, demand shocks would entirely manifest as changes in wages, and through Equation (B12) it is clear that an increase in I would generate a more substantial positive impact on high-skilled labor H compared to low-skilled labor L since the latter faces an additional, negative direct displacement effect. This could explain why we observe a stronger demand impact on male adults compared to female adults post power-looms adoption in our estimation results and indicate the skill-biased nature of the automation technologies during the

historical time.¹¹ Given that elastic labor is inherent in our oligopsony setting (i.e. even if the labor supply is perfectly inelastic at market level firms can still adjust workers as long as $\eta > 0$), the extent of technological impact on labor employment versus wages depends on the wage elasticity of labor supply from within-market competition. To see this clear, combining the equations in Equations (B10) and (B11) we have

$$\begin{aligned} W_{iL} &= (\mu_{iL}(I' - I)Y_i)^{\frac{1}{1+\eta_L}} \mathbf{L}^{\frac{\eta_L - \phi_L}{\phi_L(1+\eta_L)}} \\ L_i &= (\mu_{iL}(I' - I)Y_i)^{\frac{\eta_L}{1+\eta_L}} \mathbf{L}^{\frac{\phi_L - \eta_L}{\phi_L(1+\eta_L)}}, \end{aligned} \quad (\text{B13})$$

and

$$\begin{aligned} W_{iH} &= (\mu_{iH}(N - I')Y_i)^{\frac{1}{1+\eta_H}} \mathbf{H}^{\frac{\eta_H - \phi_H}{\phi_H(1+\eta_H)}} \\ H_i &= (\mu_{iH}(N - I')Y_i)^{\frac{\eta_H}{1+\eta_H}} \mathbf{H}^{\frac{\phi_H - \eta_H}{\phi_H(1+\eta_H)}}. \end{aligned} \quad (\text{B14})$$

By comparing the equations in Equations (B13) and (B14), it is clear that the impact of technological change would be larger on employment than wage when η is large, i.e. when the labor market is more competitive. Therefore, the empirical findings in our event study analysis that indicate a relative increase in wages for female adult workers at adopting firms compared to their non-adopting counterparts contrasted with a muted effect on employment, could stem from a substantially inelastic labor supply of female adult workers. However, for male adult workers, the observed disproportionate increase in employment relative to wages does not necessarily imply a highly elastic labor supply and substantially more competitive labor market. This is because the event study results could also be influenced by strategic responses under oligopsonistic competition, as explored below.

To analyze the strategic responses of competitor firms upon the technology adoption by firm i , we can again use Equations (B13) and (B14) but now substituting i with $k \neq i$. Now, there is no direct effect from shifting the technological thresholds for both types of labor and it turns out that the changes of aggregate index \mathbf{L} and \mathbf{H} would dominate other indirect effects for the competitors. If the adopted firm increased employment to meet higher labor demand, and if $\eta_S > \phi_S$, this would induce its competitors to reduce employment and increase wage levels, as firms are strategic substitutes in employment and strategic complements in wage setting under Cournot competition. Consequently, in our context, the increase in male adult worker employment post-power loom adoption observed in our event study analysis is likely to be overestimated as their competitors would respond simultaneously by reducing employment. Conversely, the observed wage increase for male adult workers is likely underestimated, as their competitors were correspondingly bidding up workers. This hypothesis finds support in our area-level results, which show insignificant changes in male adult worker employment alongside power loom prolif-

¹¹While it is difficult to prove formally, the markdown effect due to the induced change in the market share is an indirect and second order effect following the direct demand changes stemmed from the combination of displacement effect and productivity effect.

eration, while area-level wages exhibit a more pronounced increase than plant-level results—all aligning with our theoretical predictions. For female adult workers in our context, the strategic responses of competitor firms are likely to be modest due to a limited labor demand increase of the adopted firm or due to that η_L is close to ϕ_L . The strategic movement in the case where $\eta_S > \phi_S$ would only understate the wage changes and overstate the employment change, unaligned with our event study results. Consistently, in the area-level results we find a wage increase extent comparable to plant-level findings.

Finally, we show that our framework also accommodates business stealing effects, akin to what is highlighted in monopoly or oligopoly settings. The intensive margin of this effect follows our previous discussions on competitor’s strategic response on labor employment. Facing a left-shifted labor supply curve, those competitors would have to reduce labor inputs and production. In addition, the extensive margin of business stealing under a similar mechanism also comes into play. The profit of a firm is

$$\Pi_i = Y_i [1 - (1 + \mu_{iL})(I' - I) - (1 + \mu_{iH})(N - I')]. \quad (\text{B15})$$

This profit is greater for more productive firm with a higher production, Y_i , and lower markdowns, μ_{iL} and μ_{iH} . For “luddite” firms that refuse power adoption, their market shares dwindle and markdowns shrink (i.e. increase in μ ’s) as competitors automate and raise labor demand, compelling them to confront costly labor and resulting in decreased production. As a result, non-powered firms would have their profits keep decreasing along with intensified mechanization in the local market. This attrition could thus lead to the exit of the least productive firms when their profits turn negative or plummet below operational fixed costs. This mechanism corroborates our area-level observation of a decreasing firm count concurrent with power diffusion, and also potentially explains the significant reduction in female adult employment at the area level.