

# The Decline of Labor Share and New Technology Diffusion: Implications for Markups and Monopsony Power<sup>\*</sup>

Shoki Kusaka<sup>†</sup>  
Ken Onishi<sup>§</sup>

Tetsuji Okazaki<sup>‡</sup>  
Naoki Wakamori<sup>¶</sup>

November 25, 2022

## Preliminary and Incomplete

### Abstract

We study the mechanism behind the decline in labor share using highly detailed plant-level data on the cement industry in Japan. Using information on the production technology in place at each plant, we find that most of the labor share decline can be explained by new technology diffusion (introduction of “New Suspension Preheater kiln”). The labor share stays constant, or even slightly increases, over time within plants of the same technology, whereas the aggregate labor share declines as production shifts to plants with a new and more capital-intensive technology. We also find that the information on plant-level technology is key to rejecting other potential hypotheses and that we would reach a qualitatively different conclusion without this information. To show this, we examine, with and without technology information, two alternative hypotheses; (i) the decline in labor share is associated with an increase in markups, and (ii) firms exercise monopsony power in the labor market. We reject these two hypotheses with technology information but may not without it.

JEL Classification: D2; L1; E2.

Keywords: Technology adoption; Productivity; Labor Share; Markup; Monopsony.

---

<sup>\*</sup>We thank to Daisuke Adachi, Jan Eeckhout, Jacob Gramlich, Mitsuru Igami, Ebehi Iyoha, Jordi Jaumandreu, Daiji Kawaguchi, Pete Klenow, Toshihiko Mukoyama, Amil Petrin, Devesh Raval, Gloria Sheu, Chad Syverson, and seminar participants at MEA Conferences, Federal Reserve Board, Federal Reserve Bank of St. Louis, IIOC, and APIOC. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by other members of the staff, by the Board of Governors of the Federal Reserve System, or by the Federal Reserve Banks. Any remaining errors are our own.

<sup>†</sup>Yale University. Email: shoki.kusaka@yale.edu.

<sup>‡</sup>The University of Tokyo and the Research Institute of Economy, Trade, and Industry. Email: okazaki@e.u-toyko.ac.jp.

<sup>§</sup>Federal Reserve Board. Email: ken.t.onishi@gmail.com.

<sup>¶</sup>Hitotsubashi University. Email: naoki.wakamori@r.hit-u.ac.jp.

# 1 Introduction

The decline of the labor share has been globally observed, and many economists and policymakers have paid attention to this phenomenon. An enormous number of studies have investigated this issue and proposed possible explanations and hypotheses for why the labor share has declined over time, such as factor-biased technical changes (e.g., [Karabarbounis and Neiman, 2014](#); [Acemoglu and Restrepo, 2020](#); [Autor et al., 2020](#)), increased exercise of product market power by large firms (e.g., [Barkai, 2020](#); [De Loecker et al., 2020](#)), declining worker power in labor relations (e.g., [Stansbury and Summers, 2020](#); [Drautzburg et al., 2021](#)), globalization and the rise of China (e.g., [Abdih and Danninger, 2017](#); [Sun, 2020](#)), and changes in the composition of the workforce (e.g., [Glover and Short, 2020](#); [Acemoglu and Restrepo, 2020](#)).<sup>1</sup>

However, a large fraction of the existing studies take a macroeconomic approach to quantify the economy-wide effect, and the most detailed data available to researchers would be establishment-level census data. A limitation of empirical research using such datasets is that we do not observe the technological differences across plants. As a result, we need to assume that there is an industry-wide production function common to all the plants, or researchers need to infer the state of technological progress from auxiliary data indirectly.<sup>2</sup> Furthermore, these studies have had to rely on deflated monetary variables of sales as an output measure because of the absence of quantity data. This may obscure the difference between market power variation and production efficiency changes as drivers of labor share decline.

Among these hypotheses, technology plays an important role; Needless to say, the first set of hypotheses (consider better phrase) capture technical changes as factor-biased productivity changes of a production function. Moreover, other hypotheses, such as product market power and monopsony power, also hinge on technology, as markups

---

<sup>1</sup>See [Grossman and Oberfield \(2021\)](#) for a more detailed summary of the literature.

<sup>2</sup>For example, [Acemoglu and Restrepo \(2020\)](#) constructs an industry-level exposure to robots, and [Aghion, Antonin, Bunel and Jaravel \(2020\)](#) uses the balance sheet values of industrial equipment and plant-level records of the usage of electro-motive force to proxy the degree of automation.

and marginal products of labor often require the estimates of production function.<sup>3</sup> Despite the importance of technology, most existing studies that take a macroeconomic approach—using establishment-level census data to quantify economy-wide effects—cannot directly observe technology which may differ across plants due to technology adoption. As a consequence, they must assume that there is an industry-wide production function common to all the plants or infer the state of technological progress from auxiliary data indirectly. This may obscure the difference between technological changes and market power variation as drivers of labor share decline.

This paper precisely addresses this issue by taking a distinct and complementary approach to the existing studies. Our empirical strategy is to focus on a specific industry and collect exact plant-level technology information, including the timing of new technology adoption. We show that “technology diffusion” is the primary driver of labor share decline by directly controlling for technology in our analysis. We also show theoretically and empirically that the plant-level technology information is key to rejecting (or avoiding overemphasizing?? ) other market-power-driven hypotheses. Without the technology information, we may conclude that the markups and monopsony power have increased.

Our analysis focuses on the cement industry and the diffusion of different generations of kilns, specifically from Suspension Preheater (SP) Kiln to New Suspension Preheater (NSP) Kiln. This industry provides us with an ideal environment because we can observe the types of kilns that each plant uses for production. In addition, the quantity and price of physical units of (homogeneous) output are observable so that we can differentiate the market power variation from production efficiency changes as drivers of labor share decline. Admittedly, no such detailed dataset is available that covers all industries, but as documented in [Kehrig and Vincent \(2021\)](#), the decline of the labor share is driven by within-industry effects. Therefore, we believe that unraveling the mechanism of the phenomena in a specific industry would help us draw macroeconomic implications. To

---

<sup>3</sup>There are some recent papers that use an IO approach—demand estimation of differentiated products, instead of production approach, to estimate markups or monopsony power, e.g., Grieco et al. Miller et al. () Azar Berry...

ensure the generalizability of our analysis, we first confirm that we can replicate the patterns observed in existing studies: the decline of the labor share, an increase in industry-wide markups paired with the decline of the labor share, and an increased discrepancy between labor productivity growth and wage growth.

We find that the industry-level labor share declines over time, but the labor share slightly increases within the plants with the same old technology. Therefore, the industry-level labor share decline is largely explained by the fact that more and more plants adopt a new and more capital-intensive production technology.

To confirm that technology diffusion explains this trend, we employ an event study design using observed variation in the timing of technology adoption. (DO WE NEED THIS SENTENCE?? Specifically, we use Difference-in-Differences with leads and lags of the treatment variable.) We examine how labor share and employment respond to technology adoption, and we find that they both start to fall at the time of adoption, which confirms that the diffusion of new technology drives the phenomena. We also examine the evolution of the plant-level capital-labor ratio, and we find that right after the installation of NSP kilns, the capital-labor ratio discretely jumps up. (DO WE NEED THIS SENTENCE?? We do not find any statistically significant pre-trend in the variables, suggesting that our research design is valid for making causal inferences.)

These findings suggest that the introduction of NSP kilns embodies the explicit technological change with a different shape of production functions across plants rather than a simple increase in total factor productivity (TFP). We confirm this by estimating the production function for each technology by using the control function approach consolidated by [Akerberg et al. \(2015\)](#). We find that the new technology is indeed more capital-intensive.

We then use our production function estimates to evaluate other hypotheses for labor share decline and illustrate that our conclusion would be qualitatively different if we lacked data on production technology. (OR "if we did not take the heterogeneity in production technology into account" ?) We estimate plant-level markups using the methodology from [De Loecker and Warzynski \(2012\)](#) and marginal products of labor (MPL) for

each plant, and we find that without considering the differences in production technology, (i) the labor share decline is paired with an increase in the markup and (ii) the growth rate of the marginal products of labor (MPL) and wage becomes increasingly disconnected. The former finding has recently attracted many researchers' attention and has been documented in existing studies, such as [De Loecker et al. \(2020\)](#) and [Autor et al. \(2020\)](#). We observe a similar pattern exists when the technology information is missing. However, we theoretically demonstrate that the decline of the labor share and the increase in the industry-level markup happen simultaneously when production shifts from plants with relatively labor-intensive technology to plants with relatively capital-intensive technology. To confirm this prediction, we control for the plant-level technology in our analysis and show that a large part of the negative correlation between labor share and markups disappears. The latter pattern is well documented in the literature (for example, in [Stansbury and Summers \(2018\)](#)), and researchers and policymakers have debated whether it is a technology-driven phenomenon or caused by some other factors, such as decreased worker power. We find that this seemingly disconnected relationship results from production technology heterogeneity, and the discrepancy vanishes once we control for plant-level technology.

This paper aims to contribute to a large literature on labor share decline and technological change. The decline of the labor share has been observed in many countries ([Karabarbounis and Neiman, 2014](#)) and in many industries ([Kehrig and Vincent, 2021](#)), and many researchers ascribe it to some sort of technological changes, especially computers and industrial robots ([Acemoglu and Restrepo \(2020\)](#), [Autor et al. \(2020\)](#), [Humlum \(2021\)](#)). Our focus on the advancement of kilns in the cement industry during the 1970s can address the gap between the rise in automation and ICT in the 1990s and the labor share decline observed since the 1980s.

Secondly, our paper contributes to the macro market power literature that relies on the "production approach." (see [Syverson \(2019\)](#)) Drawing on [De Loecker and Warzynski \(2012\)](#) and production function estimation from the IO literature ([Olley and Pakes \(1996\)](#), [Levinsohn and Petrin \(2003\)](#), [Akerberg et al. \(2015\)](#)), [De Loecker et al. \(2020\)](#) es-

timate markups during the period 1955-2016 for the U.S. economy and find that it has risen steadily. Similarly, [Yeh et al. \(2022\)](#) developed a new way to characterize aggregate markdowns from production function estimation and quantify the long-term trends of monopsony powers in the US manufacturing sector. Our contribution to this literature is that we document the importance of estimation bias in market power due to the lack of production technology information.

Finally, our paper relates to the burgeoning literature on technological change and production function estimation. A common approach to production function estimation assumes productivity as a Hicks-neutral shifter. Several authors have recently considered departures from this standard assumption. ( [Doraszelski and Jaumandreu \(2018\)](#), [Raval \(2022\)](#), [Zhang \(2019\)](#), [Demirer \(2022\)](#) ) These recent papers highlighted the importance of labor-augmenting productivity and developed ways to estimate production functions with factor-augmenting productivity change. In contrast, our paper instead assumes that producers have an explicitly different production function according to their use of the old or new type of kilns, aside from any productivity differences. There are few works with this approach. Examples are [van Biesebroeck \(2003\)](#), which models the choice between lean or mass production in the car industry, and [Rubens \(2022\)](#), which features the introduction of mechanical coal cutters in the 19th-century coal mining industry.

This paper is organized as follows. Section 2 describes the industry and provides the historical background of the Japanese cement industry as well as the data used in our empirical analysis. We propose (PROPOSE?) technology diffusion to explain the decline in the labor share in Section 3. We further examine other hypotheses proposed in the literature in Section 4. Section 6 concludes.

## 2 Industry Backgrounds and Data

Though the majority of the studies in the literature take a “macroeconomics” approach which uses census data to quantify economy-wide effects, production *technology* is still

unobserved in such census data. To overcome the problems associated with the unobservability of technology, we take a distinct approach, focusing on one specific industry, namely the Japanese cement industry. There are three important advantages to studying the cement industry: (i) the observability of production technology, which is typically unobserved in the standard census data, (ii) the homogeneity of the product, which enables us to estimate markups accurately; and (iii) a simple production process, which enables us to estimate productivity easily through production function estimation. Of course, one might worry about the generalizability of our results, as we use only one particular industry from the manufacturing sector. This concern is discussed in Section 5.

In the following subsections, we first explain the industry backgrounds, elaborate on the aforementioned features and advantages of the industry, and describe the two data sources that we use in this paper. We then show some key statistics.

## **2.1 Industry Backgrounds: Cement and Its Production Technology**

Cement is one of the most important construction materials, as concrete and mortar are made from cement. There are several types of cement. For example, Portland cement is the most common type of cement, accounting for about 75% of cement products, according to Japan Cement Association. (WOULDNT WE NEED TO CITE A DOCUMENT ?) Though there are several different types of cement, each of them is standardized by the Japanese Industrial Standards and thus can be treated as a homogeneous product. To produce cement, crushed limestone, clay, and other minerals are mixed and put into a kiln to be heated. This process yields clinker, which is an intermediate cement product and the focus of this paper. The final procedure of mixing ground clinker with gypsum produces cement. As demonstrated, the production process of cement is simple.

Cement kilns are the heart of this simple production process, and it is important for us to understand some technical aspects of cement kilns in Japan. Prior to our sample period, in the 1960s, the suspension preheater (SP) process was imported from Germany, and due to its high energy efficiency, SP kilns gained in popularity and took a dominant position. Most of the newly built kilns in the 1960s were SP kilns, and in the 1970s, contin-

uing improvements were made by Japanese companies, and new suspension preheater (hereinafter NSP) kilns were developed. The main innovation of NSP kilns is attaching pre-calciner to the SP kilns, which breaks down  $\text{CaCO}_3$  in limestone into  $\text{CaO}$  and  $\text{CO}_2$  in an efficient way, and this feature enabled further mass production. In our data, after 1970, almost all newly built kilns were NSP kilns, and this homogeneity of investment simplifies our analysis, as well.

(To be added) TBA: Market definition - we divide the markets into eight areas. (BUT WE DIDNT EXPLOIT ANY MARKET VARIATIONS!)

(To be added) Import and export? (Brazil, US) Japan is geographically isolated from other countries, and we only import and export from/to Korea, Taiwan, and HK?

(To be added) There are some important no following Okazaki, Onishi, and Wakamori, 2022 Is this different from Nishiwaki? 201x?

Note that import and export in

Unlike AAA or BBB, which study Brazil and US cement industries, Japan is geographically isolated from other countries in recent years. up to 15%?

## 2.2 Data Sources

For this study, we combine two complementary plant-level data: (i) *Cement Yearbook (Cement Nenkan)*, published by the Cement Press Co. Ltd. (Cement Shinbunsha), and (ii) Census of Manufacture, collected by the Japanese Ministry of Economy, Trade, and Industry. The yearbook mainly provides plant-level information on monthly production capacity (in tons), production output (in tons), number of workers, and ownership and geographical location of the plants. In addition to these basic characteristics of the plants, the data also contains the types and the number of kilns that each plant owns, which makes this dataset special. Although the technology each plant employs is typically unobserved, this Yearbook data provides such kiln-level information. On the other hand, the Census of Manufacture provides a similar but slightly different set of information on the plants, i.e., the total shipment value (in JPY), material inputs (in JPY), number of employees, total wage (in JPY), investment (in JPY), and asset values (in JPY).



Note that the sample periods for these two data sources are slightly different. We obtain the former data from 1970 to 2010, whereas we obtain the latter data from 1980 to 2010 because the data from 1970 to 1979 for the latter data are unavailable. We combine these two data sources via location and plant name information, as most plants are located in a different municipality with very few exceptions.

## 2.3 Summary Statistics and Key Features

Summary statistics of our data are given in Table 1. Panel (a) presents plant-level summary statistics pooling all years, whereas Panels (b1), (b2), and (b3) present plant-level statistics for some selected years, namely 1970, 1990, and 2010, respectively. There are some important trends that can be found by comparing Panels (b1), (b2), and (b3).

First, the number of observations in 1970 is 53, whereas it is 30 in 2010, implying that the number of plants decreased by about 40% over 40 years. On the other hand, the number of firms in 1970 and 2010 was 22 and 18, respectively, implying that most firms concentrate their production on a single plant or a small number of plants, on average, in 2010. Although the number of plants decreased sharply, monthly capacity, defined as how much clinker a plant can produce when operating for 600 hours per month, and annual clinker production per plant have increased over 40 years so that industry-level capacity and production have decreased only slightly.

Second, the fraction of the number of NSP kilns has increased considerably. There were no NSP kilns in 1970, whereas the old kilns were mostly replaced by NSP kilns over 40 years. To further see the change in cement production technology, Figure 1 graphically shows the absolute number of kilns and share, depending on technology, i.e., types of kilns, over time. In 1970, the initial year of our sample period, there were about 220 kilns, though the majority of them were old types and SP kilns accounted for less than 20%. Note, again, that there were no NSP kilns in 1970. During the 70s, however, NSP kilns dramatically increased their popularity, maintaining their dominant position after the 1980s. In our main analysis, we explore the labor share with and without controlling for this technology information.


Table 1: Summary Statistics

	Num. of Obs.	Mean	Std. Dev.	Min.	Max.
<b>Panel (a): Plant-Level Statistics (All)</b>					
Monthly Capacity (tons)	1,748	188,716	124,935	18,180	696,250
Annual Clinker Production (tons)	1,748	1,742,626	1,293,791	0	8,082,269
Average Cement Price (JPY/ton)	1,748	10,375	2,580	5,800	17,075
# of Workers (person)	1,672	193	140	16	1303
Average Wage per Worker (JPY)	1,748	4.34	1.43	.928	13.77
Share of NSP Kilns	1,738	.549	.429	0	1
<b>Panel (b1): Plant-Level Statistics in 1970</b>					
Monthly Capacity (tons)	53	128,396	80,840	25,000	350,000
Annual Clinker Production (tons)	53	1,025,507	621,417	48,000	2,684,197
Average Cement Price (JPY/ton)	53	5,965	202.0	5,800	6,900
# of Workers (person)	53	318	175	114	1,205
Average Wage per Worker (JPY)	53	2.32	.527	.928	3.62
Share of NSP Kilns	53	0	0	0	0
<b>Panel (b2): Plant-Level Statistics in 1990</b>					
Monthly Capacity (tons)	41	178,472	111,121	30,000	553,417
Annual Clinker Production (tons)	41	1,836,281	1,160,588	255,000	5,428,197
Average Cement Price (JPY/ton)	41	11,550	1,375	9,600	13,200
# of Workers (person)	41	169	94.4	57	560
Average Wage per Worker (JPY)	41	4.47	.571	2.83	5.41
Share of NSP Kilns	41	.750	.379	0	1
<b>Panel (b3): Plant-Level Statistics in 2010</b>					
Monthly Capacity (tons)	30	165,567	127,285	36,167	557,083
Annual Clinker Production (tons)	30	1,561,800	1,321,220	276,000	6,169,000
Average Cement Price (JPY/ton)	30	10,076	471.0	9,000	10,900
# of Workers (person)	30	104	65.5	34	371
Average Wage per Worker (JPY)	30	5.98	.842	4.32	7.91
Share of NSP Kilns	30	.861	.300	0	1

*Note:* The number of observations for the number of workers is 1,672 in Panel (a), because Cement Yearbook in 1976 does not provide the number of workers.


Third, the number of workers has decreased sharply; the average number of workers in 1970 was 318, whereas it was 104 in 2010. Figure 2 plots the plant-level number of employees over time together with linear fitted values. Though we observe substantial heterogeneity in plant size, all plants decrease the number of workers over time. This decrease in the number of workers means that the labor productivity—measured in output per worker—also increased over 40 years, as we see that the plant-level clinker production has increased. On the other hand, though the average wage also increased over time, the change in the average wage is not as large as the change in labor productivity. These facts raise some questions; whether this reduction in the number of workers is driven by the adoption of new technology and whether the gap between growth in labor produc-

Figure 1: Diffusion of Technology



tivity and in wage is due to increased monopsony power.

Figure 2: Number of Workers per Plant over Time



*Note:* This figure plots the plant-level number of employees over time together with a linear fitted value.

### 3 Decline of Labor Share and Technology Adoption

In this section, we first present the labor share patterns in the data, which exhibit similar features as the existing studies have documented. Then we argue that the new technology adoption is the main driver for the decline of labor share and offer evidence by implementing the event-study design analysis. Finally, we show that new technology adoption comes with a change in the shape of production function by estimating the production function with information on plant technology. More specifically, we show that cement production becomes more capital-intensive when plants adopt new technology.

#### 3.1 The Decline of Labor Share and New Technology Adoption

We first plot the industry-level labor share, defining the labor share as the total wage payment divided by the monetary value of total output.<sup>4</sup> In Figure 3, each dot represents the labor share for each year and the gray line represents the smoothed-nonparametric fit. The industry-level labor share falls over our sample period with a sharper decline when the new technology diffuses between 1973 and the early 1980s, as we see in Figure 1.

At the firm level, the output shifts from high-labor share plants to low-labor share plants. Figure 4 plots the histogram of the share of output on the vertical axis and the plant-level labor share on the horizontal axis for some selected years, i.e., in 1975, 1985, 1995, and 2005. From 1975 to 2005, the distribution shifted from the right to the left, which implies that the production shifted from plants with a high labor share to plants with a low labor share.

The virtue of our approach is that we observe the exact technology used at the plants. To quantify how much the diffusion of new technology contributes, we replicate the analysis in Figure 3 *conditional on* the plant-level technology. In Figure 5, we plot the average labor share within the plants with new technology (the dotted line), within the plants with old technology (the dashed line), and the industry-level labor share (the solid line).

---

<sup>4</sup>As the Census data is available only after 1980, we compute the labor share using the data in Cement Yearbook. More specifically, the total wage payment is computed as the number of employees multiplied by the average wage and the monetary value of total output is computed as the output multiplied by the average cement price in that region.

Figure 3: Industry-level Labor Share


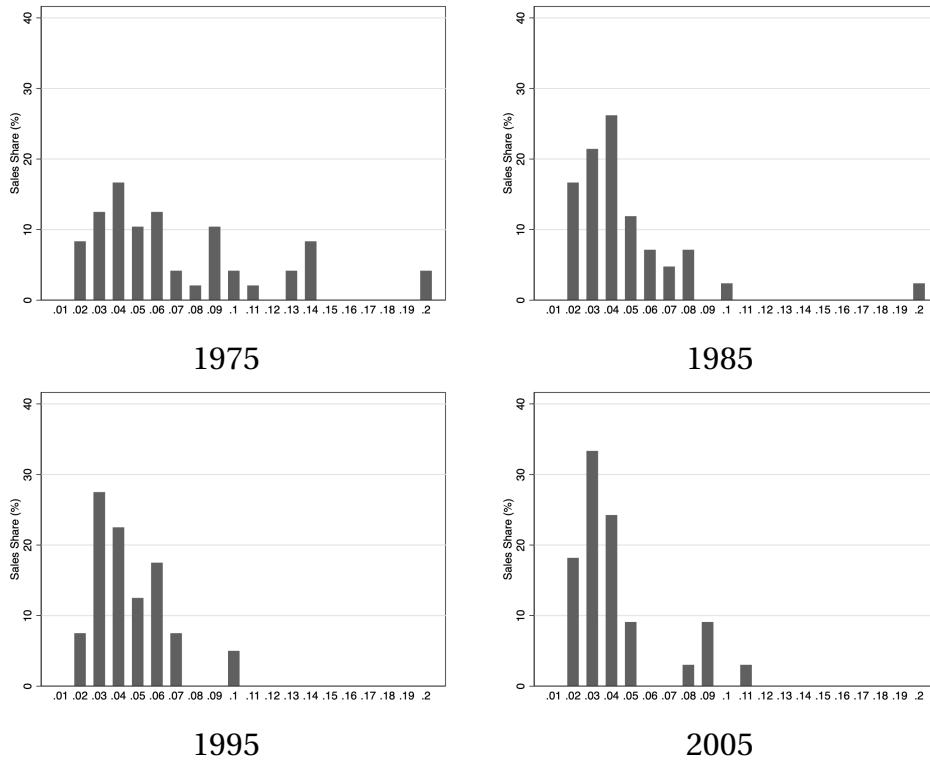



Figure 4: Plant-level Production Share and Labor Share



The last line, the industry-level labor share, corresponds to the solid line in Figure 3. Interestingly, the labor share does not fall within the same technology plants as the dashed

line and dotted lines stay relatively flat. However, the industry-level labor share, the solid line, falls rapidly as new technology diffuses because the new technology plants have a lower level of labor share. Figure 5 clearly shows that the decline of labor share is associated with the new technology diffusion.

Figure 5: Labor Share Conditional on the Plant-level Technology



To assess the argument more quantitatively in a descriptive manner, we estimate the following equations using the plant-level labor share by Ordinary Least Squares (OLS);

$$\text{LaborShare}_{it} = \beta_0 + \beta_1 t + \beta_2 \mathbf{1}_{\{\text{NSP Kilns}_{it}\}} + F_i + \varepsilon_{it},$$

where  $i$  is a plant index,  $t$  denotes year,  $\mathbf{1}_{\{\text{NSP Kilns}_{it}\}}$  is a dummy variable taking one if a plant owns at least one NSP kiln in year  $t$  and zero otherwise,  $F_i$  is a plant fixed effects,  $\beta$ s are the parameters to be estimated, and  $\varepsilon_{it}$  is an independent error term. Here, we are interested in the estimated coefficient on  $t$ , i.e.,  $\beta_1$ . We expect that  $\beta_1$  would be estimated as negative when we do not control for the plant-level technology because the industry-level labor share declines over time. In contrast, we expect that  $\beta_1$  would be estimated near zero or positive when we control for the plant-level technology. Table 2 summarizes the estimation results and confirms our expectations. The first column presents

Table 2: Time Trend of Labor Share

Dependent Var.	(i) Labor Share	(ii) Labor Share	(iii) Labor Share	(iv) log(Labor Share)
Year ( $\beta_1$ )	$-1.13 \times 10^{-3}$ *** (0.0157 $\times 10^{-3}$ )	$-2.42 \times 10^{-4}$ (0.0176 $\times 10^{-3}$ )	$0.371 \times 10^{-3}$ *** (0.0880 $\times 10^{-3}$ )	$0.006$ *** (0.001)
NSP Kiln dummy ( $\beta_2$ )		$-0.0474$ *** (0.00469)	$-0.0244$ *** (0.00302)	$-0.389$ *** (0.0301)
Constant	✓	✓	✓	✓
Plant fixed effects			✓	✓
$N$	1,673	1,673	1,673	1,673

Standard errors in parentheses

Other Controls includes a constant term.

\* ( $p < 0.10$ ), \*\* ( $p < 0.05$ ), \*\*\* ( $p < 0.01$ )

the results without controls for technology, and the coefficient on year is estimated as negative and statistically significant. In the second column, once we control for the technology, the significance of  $\beta_1$  disappears. However, we now find that the coefficient on an NSP kiln dummy,  $\beta_2$ , is estimated as negative and statistically significant, implying that a plant introducing NSP kilns has a lower labor share. When we further control for the plant fixed effects, the estimates become positive and statistically significant. These results are consistent with Figure 5. To quantify the economic significance of the results in the third column, we replace the left-hand-side variable with the logarithm of labor share, which allows us to quantify the percentage change easily. The result is presented in the fourth column, suggesting that the labor share increases at the plant level by 0.6% every year. The magnitude is not very large but not negligibly small.

Note that the labor share here is computed using the data from Cement Yearbook. Even when we use the value-added variables in the Census data, we obtain the qualitatively same results. See Appendix XXX for these results.

### 3.2 Evidence from Event Study Design

We further zoom into the plant-level changes in variables to confirm that our findings in the previous sections are driven by technology diffusion. To this end, we take advantage

of the richness of our data, i.e., we can observe the timing of new technology adoption. Using the variation in the timing of technology adoption, we employ an event study design, i.e., difference-in-differences with leads and lags of treatment variable. Formally, we adopt a method proposed by [Callaway and Sant’Anna \(2021\)](#). Here, the adoption of NSP kilns is the “treatment,” we estimate the average treatment effect on the treated (ATT) for each treatment cohort. ATT from  $\tau$  years from the treatment for the plants who adopted NSP kilns in year  $t$  is identified as:

$$\text{ATT}(t, \tau) = E \left[ \left( \frac{G_{it}}{E[G_{it}]} - \frac{\frac{p_t(X_{i,t-1})C_{it}}{1-p_t(X_{i,t-1})}}{E \left[ \frac{p_t(X_{i,t-1})C_{it}}{1-p_t(X_{i,t-1})} \right]} \right) (y_{i,t+\tau} - y_{i,t-1}) \right], \quad (1)$$

where  $G_{it}$  is one if plant  $i$  adopts NSP kilns in year  $t$  and zero otherwise,  $C_{it}$  is one if firm  $i$  never adopts or has not yet adopted NSP kilns and zero otherwise,  $p_t(X_{i,t-1})$  is the probability that plant  $i$  adopt NSP kilns in year  $t$  conditional on  $G_{it} = 1$  or  $C_{it} = 1$ , and  $y_{iu}$  is the outcome variable of plant  $i$  in year  $u$ .<sup>5</sup> We define ATT  $\tau$  years from the treatment as the weighted average of  $\text{ATT}(t, \tau)$  as:

$$\text{ATT}(\tau) = \sum_t w_t \text{ATT}(t, \tau),$$

where  $w_t$  denotes the weight, which is the number of plants treated in year  $t$  divided by the total number of treated firms.

We estimate  $\text{ATT}(t, \tau)$  by replacing the expectation with the empirical average, and  $p_t(X_{i,t-1})$ , the propensity score, by estimating a probit model. For  $X_{i,t-1}$ , we use the logarithm of plant  $i$ ’s total capacity, production quantity, and monetary value of total output, and we estimate a separate probit model for each year.


First, [Figure 6](#) plots the evolution of plant-level labor share relative to the timing of new technology adoption. The estimated coefficient for the year before the adoption ( $t = -1$ ) is normalized to be zero by construction. The x-axis shows the years relative

---

<sup>5</sup>In [Callaway and Sant’Anna \(2021\)](#), the authors propose to use never treated individuals as the control group. However, it is not feasible in our context because the number of plants that never adopted NSP kilns is too small to derive any meaningful inference. Also, the authors provide computer codes for Stata and R to implement the estimation, and the authors provide the option to use never-treated and not-yet-treated individuals as the control group.



Figure 6: The Effects of Adoption of New Technology on Plant-Level Labor Share



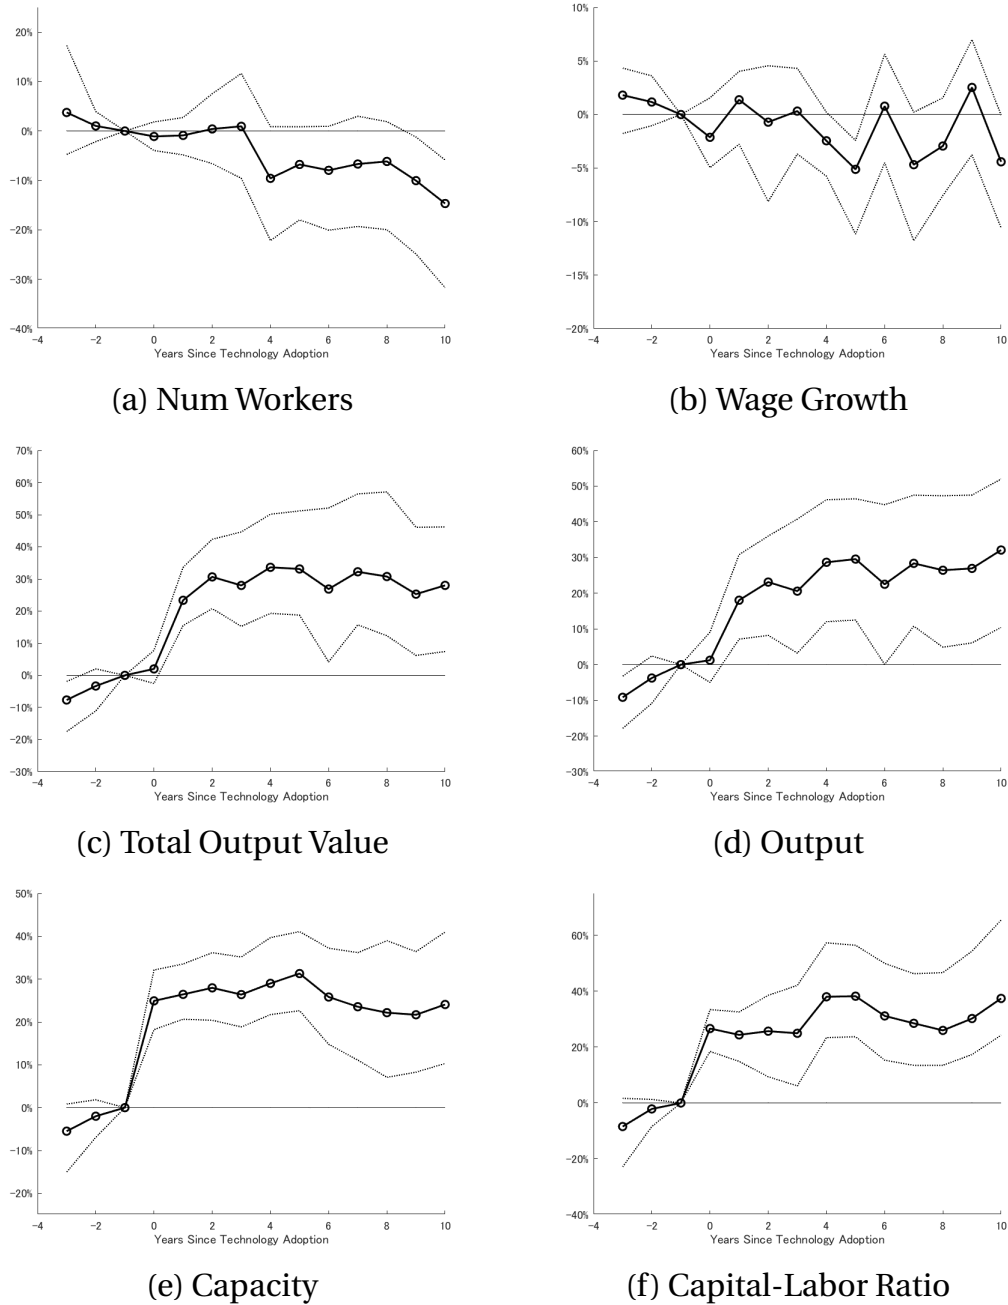
to the year of NSP adoption and the y-axis shows the estimated ATT. Here, the year of NSP adoption is identified by the year the plant installed its first NSP kiln.<sup>6</sup> The solid line presents the estimated ATT and the gray dotted lines represent a 90% confidence interval. The confidence interval is constructed by the bootstrap method with 200 replications. When we look at the solid line, the labor share starts to decline after the technology adoption. However, the decline is not immediate. Rather, it takes several years. After four years of adoption, the labor share remains below the pre-adoption level with statistical significance.

To decompose the effects into the changes in the numerator and denominator of labor share calculation, we first look at employment and wages. Panel (a) of Figure 7 plots the evolution of plant-level (log) employment relative to the timing of new technology adoption. In contrast to the labor share, the number of employees decreases even more gradually than the labor share and the estimated effect becomes statistically significant after 9 years of the new technology adoption. In the long run the number of workers decreases by 15% implying that the change in the number of workers is one of the drivers of

<sup>6</sup>A plant typically has multiple kilns and the adoption of NSP kilns are typically gradual, i.e., each plant replace one or two of its kiln first and then replace the remaining kilns over time.

the decline in the labor share. On the other hand, there are no clear differences between the treated group and the control group in terms of wage growth as in Panel (b) of Figure 7.

Figure 7: The Effects of Adoption of New Technology on Plant-Level Outcomes



Another driver of the labor share decline is the total output, the denominator of labor share calculation. Panel (c) of Figure 7 plots the evolution of plant-level (log) total out-

put value relative to the timing of new technology adoption. Unlike the number of workers, the output responds to the adoption relatively quickly and increases substantially. It also grows gradually over time. The growth of the output together with the decline in the number of workers implies that labor productivity has increased. As the NSP adoption has no effect on wage growth, such an increase in labor productivity translates into the labor share decline. One may worry that the change is driven by the change in the cement price. If a new technology produces higher-quality output, this may be a valid concern. However, in the cement industry, the output quality is homogeneous and it is hard to believe the cement price varies based on the technology. To confirm that the change is attributed to output quantity rather than price, we estimate the same equation with output quantity in Panel (d). The estimates in Panel (c) and Panel (d) is almost identical, suggesting that the denominator of labor share calculation increases due to higher real labor productivity.

Why does labor productivity, defined as output per worker, increase without any change in wage growth? To answer this question, we plot the evolution of plant-level (log) capacity relative to the timing of new technology adoption in Panel (e) of Figure 7. Capacity jumps up right after the adoption and stays at a higher level compared to the pre-adoption period. Similarly, Panel (f) of Figure 7 plots the evolution of plant-level capital-labor ratio relative to the timing of new technology adoption. As implied by the results in Panel (b) and Panel (e), the capital-labor ratio increases, which suggests that the production technology and optimal capital-labor ratio are different between non-NSP kilns and NSP kilns.

Overall, all the results are consistent with the hypothesis that the adoption of NSP kilns causes a decline in labor share. They also suggest the mechanism behind the decline. The production technology of NSP kilns is different from the old-type kilns and it is more capital-intensive. More capital-intensive production technology increases labor productivity without affecting wage growth, and, as a result, a plant has more output with fewer wage bills, which translates into a lower labor share.

### 3.3 Technology Adoption and the Shape of the Production Function

Our findings in the previous subsection—the decline of labor share in tandem with the increase in capital-labor ratio after new technology adoption—are very difficult to rationalize if the new technology simply brings an increase in TFP. Then, what are the implications of these findings for production function? The recent literature on factor-augmenting technical changes, such as [Doraszelki and Jaumandreu \(2018\)](#), [Zhang \(2019\)](#), [Raval \(2022\)](#), and [Demirer \(2022\)](#), suggest that our findings may be explained by incorporating labor-augmenting technical changes. Although we could employ such an approach, taking advantage of the data where we can directly observe technology for each plant, we adopt a more straightforward approach—estimating production functions *separately* for each technology, as it is more natural to assume that the shapes of the production function for each technology are different.

Here, we assume that the production function takes a Cobb-Douglas form:

$$Y_{it} = A_{it} K_{it}^{\beta_k^\tau} L_{it}^{\beta_l^\tau},$$

where  $Y_{it}$  is the quantity of the output,  $A_{it}$  is the TFP,  $K_{it}$  is the physical capacity,  $L_{it}$  is the total wage payment, and  $(\beta_k^\tau, \beta_l^\tau)$  is a set of parameters to be estimated for technology  $\tau$ ,  $\tau \in \{\text{old, new}\}$ . We estimate the value-added production function which is considered in [De Loecker and Scott \(2016\)](#), [Akerberg et al. \(2015\)](#), and [Gandhi, Navarro and Rivers \(2020\)](#). This usage of the value-added production function can avoid potential identification problems regarding intermediate inputs.

The specification is written as

$$y_{it} = \beta_k^\tau k_{it} + \beta_l^\tau l_{it} + \omega_{it} + \varepsilon_{it}$$

where each lowercase variable is in the form of a logarithm and  $\omega_{it}$  is an unobserved productivity shock and  $\varepsilon_{it}$  is the unanticipated shock to output. We control the unobserved productivity shock with a control function with the value of investment  $i_{it}$  as in [Olley and](#)

Pakes (1996) and Akerberg et al. (2015):

$$\omega_{it} = h^\tau(k_{it}, i_{it})$$

The estimation procedure consists of two stages. First, we non-parametrically estimate

$$y_{it} = \phi^\tau(k_{it}, l_{it}, i_{it}) + \epsilon_{it},$$

where  $\phi^\tau(k_{it}, l_{it}, i_{it}) = \beta_k^\tau k_{it} + \beta_l^\tau l_{it} + h(k_{it}, i_{it})$ . Given the productivity process  $\omega_{it} = g(\omega_{it-1}) + \xi_{it}$  and  $\omega_{it} = \phi^\tau(k_{it}, l_{it}, i_{it}) - \beta_k^\tau k_{it} - \beta_l^\tau l_{it}$  from the first stage, we estimate a set of parameters  $\theta$  including  $\beta_k^\tau$  and  $\beta_l^\tau$  by using the following moment condition:

$$E[\xi_{it}(\theta)(i_{it-1}, k_{it}, l_{it-1})'] = 0.$$

Table 3 summarizes the estimation results. Column (i) demonstrates the results when we estimate the labor and capital coefficients by pooling all plants regardless of their technology:

$$y_{it} = \beta_k k_{it} + \beta_l l_{it} + \omega_{it} + \varepsilon_{it},$$

whereas Column (ii) demonstrates the results when we estimate them separately for each technology via introducing kiln-type dummies and their interaction terms with other variables in order to get output elasticities by kiln types:

$$y_{it} = \beta_k^{old} k_{it} + \beta_l^{old} l_{it} + \mathbf{1}_{\{\text{NSP Kilns}_{it}\}}(\beta_0 + \beta_k^{new} k_{it} + \beta_l^{new} l_{it}) + \omega_{it} + \varepsilon_{it}. \quad (2)$$

Standard errors are calculated ...

When we estimate the model by pooling all plants,  $\beta_k$  is close to 1 and  $\beta_l$  is about 0.18, implying that technology exhibits economies of scale. On the other hand, when we estimate the model separately for each technology as in Column (ii), capital and labor coefficients are 0.778 and 0.259 for old technology and 0.907 and 0.099 for new technology, implying that both technology no longer exhibits economies of scale. One of the reasons why technology exhibits economies of scale when estimating the model by pooling both

Table 3: Production Function Estimates with and without Technology Information

	(i)	(ii)		(iii)
	Pooling Both Technology	Old Tech	New Tech	Pooling Both Technology
$\beta_k$	0.971 (0.110)	0.778 (0.110)	0.907 (0.085)	0.872 (0.071)
$\beta_l$	0.184 (0.140)	0.259 (0.103)	0.099 (0.096)	0.237 (0.094)
$\beta_0$ (TFP Gain)	- -	- -	0.106 (0.710)	0.060 (0.103)
$N$	1,408	1,408		1,408

technologies is omitted variable bias. As mentioned in Section 2, NSP kilns tend to be larger in size and have higher TFP (more efficient) than the older types of kilns. Thus, if we do not control for the TFP gain of new technology,  $\mathbf{1}_{\{\text{NSP Kilns}_{it}\}}\beta_0$  in Equation (2), we would have an upward bias for capital and labor coefficients, as there are positive correlations between the TFP gain and labor input and between the TFP gain and capital input.

Coming back to the results in Column (ii), as we expect, the new technology is more capital intensive, whereas the old technology is more labor intensive. We indeed test a hypothesis that  $H_0 : \beta_{old}^k = \beta_{new}^k$  and reject the null hypothesis at the 10% significance level. Therefore, profit-maximizing plants would need less labor, which results in a lower level of labor share. When more plants adopt new technology, as a result, the industry-level labor share falls.

One natural concern is that we may reach the same conclusion by just including the technology dummy in the production function. To address this concern, we check how the estimated production function would change by including technology-fixed effects and the the estimated results are presented in Column (iii). First, note that the scale parameter, i.e.,  $\beta_k + \beta_l$ , is about 1.15 in Column (i), whereas the scale parameter is close to 1 in Column (ii) of Table 3 . Because new technology plants are more efficient (higher TFP) and have a larger capacity, ignoring the technology information creates upward bias in the scale parameter as the plant size is seemingly correlated with efficiency. With

technology-fixed effects but with a single input elasticity for capital and labor (Column (iii)), the scale parameter is in between these, about 1.1, suggesting that the bias is partly mitigated but not completely. Interestingly, both the estimated capital coefficient and labor coefficient also fall between the estimated coefficients for old and new technology reported in Column (ii), suggesting that the Argue that it is still different from our specification. These results indicate that technology-fixed effects alone are not enough to capture the difference in technology.

## 4 Alternative Hypotheses and the Role of Technology Information

As the decline of the labor share has attracted huge attention from both researchers and policymakers, many alternative explanations have been proposed in the literature. [Grossman and Oberfield \(2021\)](#) classifies the existing hypotheses into five types: factor-biased technical changes, the increased exercise of product market power by large firms, declining worker power in labor relations, globalization and the rise of China, and changes in the composition of the workforce. In this section, we examine these alternative hypotheses *in the presence of* technology information.

### 4.1 The Increase in Markups

There is a growing interest in how concentration affects macroeconomic conditions, and there are a number of studies that document the increase in markups is paired with the decline of labor share. The literature follows the method proposed [De Loecker and Warzynski \(2012\)](#) and [De Loecker et al. \(2020\)](#) and estimates the markups by a production function approach, using the optimality of variable inputs. Recently, a few studies (e.g., [Raval, 2022](#); [Doraszelki and Jaumandreu, 2019](#)) have questioned whether the markup implied from cost minimization well captures the actual product-level markups. In this paper, we find another potential factor that may bias the estimated markup: the lack of information on the plant- or firm-level technology. We find that, in the absence of technology infor-

mation, the adoption of more capital-intensive technology at some plants would lead to an overestimation of their plant-level markups implied by cost minimization. Thus, as more and more plants switch to the new production technology, the industry-level markup would be overestimated, as if the labor share decline is caused by the increasing markups.

#### 4.1.1 The Role of Technology Information

Let us first provide an example to highlight the mechanism by which the lack of information on the technology would lead to a bias in the industry-level markup. Consider an environment where firm  $i$  has production technology characterized by  $Y_i = A_i K_i^{\beta_k} L_i^{\beta_l}$ . Also suppose firm  $i$  faces a demand curve characterized by  $P_i(Q_i) = \xi_i Q_i^{-\varepsilon}$  where  $\varepsilon < 1$ . The labor market is competitive with the wage level of  $w$  and each firm maximizes its profit by choosing its optimal level of labor input. Namely, firm  $i$  solves the following maximization problem:

$$\max_{L_i} P_i(Q_i)Q_i - wL_i \quad \text{subject to} \quad Q_i \leq A_i K_i^{\beta_k} L_i^{\beta_l}.$$

In this environment, we can analytically solve for the markup firm  $i$  charges. The corresponding cost minimization problem to the above profit maximization is

$$\min_{L_i} wL_i \quad \text{subject to} \quad Y_i \geq \bar{Q}.$$

The first-order condition of this problem gives us an analytical expression of the markup, which is given by

$$\text{Markup}_i = \beta_l \frac{P_i Y_i}{w L_i} = \frac{1}{1 - \varepsilon}.$$

Note that this markup is constant and solely depends on the demand elasticity  $\varepsilon$ . In this environment, researchers can easily estimate the markup when  $\beta_l$  is estimable. As  $P_i$ ,  $Y_i$ ,  $L_i$ , and  $w$  are in the data, firm  $i$ 's markup can be estimated with an estimate of  $\beta_l$ ,  $\hat{\beta}_l$ , by

$$\widehat{\text{Markup}}_i = \hat{\beta}_l \frac{P_i Y_i}{w L_i}.$$



Also, the industry-level markup can be estimated by the weighted average of the firm-level markups by

$$\widehat{\text{Markup}} = \sum \omega_i \widehat{\text{Markup}}_i,$$

where  $\omega_i$  is an appropriate weight (such as the share of sales).

Now, further, consider a case where there are two different types of firms. One type of firm has labor(material) intensive production technology characterized by  $Y_i = A_i K_i^{\beta_k^N} L_i^{\beta_l^N}$  and the other type of firm has capital intensive technology characterized by  $Y_i = A_i K_i^{\beta_k^O} L_i^{\beta_l^O}$ , where  $\beta_k^N > \beta_k^O$  and  $\beta_l^O > \beta_l^N$ . Even in this case of heterogeneous technology, the markup is constant,  $1/(1 - \varepsilon)$ , regardless of production technology at each plant. Suppose the researchers do not have direct information on the production technology each firm uses and estimate a single production function, a single value for  $\beta_k$  and  $\beta_l$ , by pooling all the observations. Let  $\tilde{\beta}_l$  be an estimate from such a misspecified model. When  $\tilde{\beta}_l$  is used to estimate the firm-level markup, the estimated firm-level markups would be biased because

$$\widehat{\text{Markup}}_i^t = \tilde{\beta}_l \frac{P_i Y_i}{w L_i} = \tilde{\beta}_l \frac{\beta_l^t P_i Y_i}{\beta_l^t w L_i} = \frac{\tilde{\beta}_l}{\beta_l^t} \frac{1}{1 - \varepsilon},$$

where  $t \in \{N, O\}$  denotes the type of firms. As  $\beta_l^O > \beta_l^N$ , the estimated markups for each technology under this misspecification would be different and have the relationship,  $\widehat{\text{Markup}}_i^O < \widehat{\text{Markup}}_i^N$ , even though the markups in this environment must be identical and only depends on the demand elasticity,  $\varepsilon$ . In addition, if  $\tilde{\beta}_l \in (\beta_l^N, \beta_l^O)$ , then the markup is downward biased for labor-intensive firms and upward biased for capital-intensive firms.

In an environment with heterogeneous technology, as the firm-level markups are constant across the firms, the industry-level markup would also be constant. When the production shifts from the plants with labor-intensive technology to the plants with capital-intensive technology, the misspecified model would lead to an increase in the estimated industry-level markup because the estimated industry-level markup is a weighted average of the estimated firm-level markups and  $\widehat{\text{Markup}}_i^O < \widehat{\text{Markup}}_i^N$ . If researchers had the firm-level (or plant-level) technology information, such an issue would not arise, i.e.,

if the model is correctly specified and a production function is separately estimated for each technology, the estimated markups for both firm-level and industry-level would be constant.

This example matches the data pattern observed in the Japanese cement industry well; there are labor-intensive (old types of kilns) and capital-intensive (NSP kilns) technologies, and the production has shifted to plants with NSP kilns because more and more plants adopt NSP kilns. Therefore, a natural concern arises that we would reach a qualitatively different conclusion as to whether a rise in markups is the main driver of labor share decline when we have or do not have plant-level technology information.

#### 4.1.2 The Estimation of Markup

Given the aforementioned potential concern, we examine how the estimated markups change over time with and without controlling for the plant-level technology. For this purpose, we first hypothetically assume that we do not observe plant-level technology and follow [De Loecker et al. \(2020\)](#) to estimate the industry-level markups. Then, we use the estimation results in Section 3 and estimate the markups taking into account the plant-level technology. The difference between these two tells us how the estimated markups are affected by the technology information.

For the case without technology information, again, we assume a Cobb-Douglas production function as

$$Y_{it} = A_{it} K_{it}^{\beta_{kt}} L_{it}^{\beta_{lt}}, \quad (3)$$

where we allow the shape of the production function to change over time as in [De Loecker et al. \(2020\)](#), i.e.,  $\beta_k$  and  $\beta_l$  now depend on time  $t$  as well. The corresponding cost minimization problem is written as

$$\min_{K,L} r_t K_{it} + w_t L_{it} \text{ subject to } Y_{it} \geq \bar{Q},$$

and the implied markup is

$$\text{Markup}_{it} = \beta_{lt} \frac{P_t Y_{it}}{W_t L_{it}}. \quad (4)$$

Further following [De Loecker et al. \(2020\)](#), we modify the production function in Equation (3) and take the material input into the production process as a fixed-proportion (Leontief) technology. Formally, we consider the following production technology:

$$Y_{it} = \min\{\beta_{mt}M_{it}, A_{it}K_{it}^{\beta_{kt}}L_{it}^{\beta_{lt}}\},$$

where  $\beta_{mt}M_{it}$  captures the material contribution to the final output. This specification is used commonly in the literature and is employed by not only [De Loecker et al. \(2020\)](#) but also other studies, including [Akerberg et al. \(2015\)](#) and [Gandhi et al. \(2020\)](#). Under this specification, the markup that takes into account the material input can be expressed as

$$\text{Markup}_{it}^M = \frac{1}{\text{Markup}_{it}^{-1} + \frac{P^M M_{it}}{P_t Y_{it}}}, \quad (5)$$


where  $\text{Markup}_{it}$  is the markup estimates from Equation (4) and  $P^M M_{it}$  is the total material spending. For the case with technology information, we follow the same steps and use the estimates in Table 3.

Figure 8 plots the industry-level markups with and without controlling for the plant-level technology. When we do not control for the technology as in the solid line, the estimated markup increases by 60 percentage points, from 1.6 in 1973 to 2.2 in 2000, during the period when the new technology diffuses and productions shift to plants with new technology. In contrast, the estimated markup after controlling for the plant-level technology stays around 1.6 for the corresponding period. These contrasting plots, again, highlight that availability of information on technology could change the result and its implication qualitatively. Note that allowing the production function to depend on time does not help control the technology difference.<sup>7</sup>

---

<sup>7</sup>We may expect that estimating a time-dependent production function captures the “average” technology in a given moment of time. When we compute the industry-level markups, we weight each observation based on output weights. For markup calculation with a single production function, the production function needs to capture the “weighted” average technology (weighted by the same weight as in the markup calculation). However, with a moment-based estimation method, each observation has equal weight when computing the moment condition. Therefore, the estimated production function would capture the average technology based on a simple average (based on the number of plants with old technology and the number of plants with new technology). This discrepancy results in the discrepancy in the estimated markups with and without technology information.

Figure 8: Markups with and without Technology Information



Our results are consistent with the recent findings by [Demirer \(2022\)](#), [Raval \(2022\)](#), and [Jaumandreu \(2022\)](#); Though they do not directly observe technology information, they account for technological differences across firms indirectly through labor-augmenting productivity and find relatively stable markups over time. In particular, we have similar quantitative results with [Demirer \(2022\)](#); Using the manufacturing industries in the US, he finds that the aggregate markup has risen from 1.3 in 1960 to 1.45 in 2012, though the aggregate markup has increased further without controlling for labor-augmenting productivity.

These studies and our study complement each other. Our results give support for their findings by highlighting the importance of technological changes in production, and their studies provide ways to reconcile these technological changes when we do not have access to technology information to derive implications on markups. An alternative way to take explicitly into account heterogeneity as a latent variable and incorporate it into a structural model is considered by [Kasahara, Schrimpf and Suzuki \(2022\)](#).

## **4.2 Growing Dispersion between Labor Productivity and Wage**

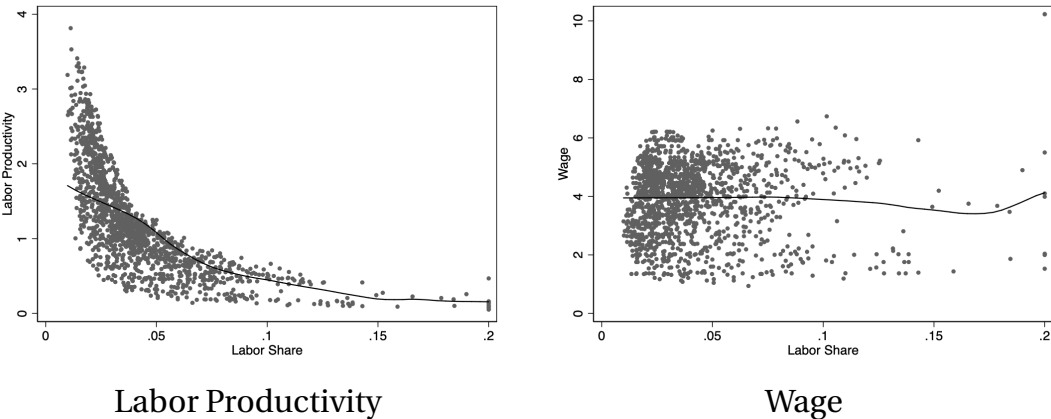
As documented in [Stansbury and Summers \(2018\)](#), several studies find that wedges between the growth rate of wage and the growth rate of marginal products of labor have been increasing. The literature proposes a few explanations, e.g., a technology-driven explanation and an explanation related to worker power, such as increased monopsony power and decreased bargaining power of workers. Although our findings so far are consistent with the technology-driven explanation, this subsection examines whether monopsony power exists in the industry.

We first present a series of descriptive evidence following [Kehrig and Vincent \(2021\)](#), which claims that monopsony power is not likely to exist by examining the relationship among labor share, labor productivity, and wage. On the one hand, if our main hypothesis—technological diffusion drives the changes in the labor share—is true, then the labor share should be explained by technology. As new technology is more capital intensive, labor productivity is higher for the new technology firms, and thus, labor pro-

ductivity should be negatively correlated with labor share, whereas the wage and labor share should have no correlation. On the other hand, if monopsony power exists, firms suppress wages, which results in a negative correlation between the wage and labor share and no correlation between labor productivity and labor share.

Figure 9 provides direct evidence to test these hypotheses. The left panel of Figure 9 plots plant-level labor productivity (defined as total output value divided by the total wage payment) on the vertical axis and plant-level labor share on the horizontal axis together with a nonparametric fitted line. There exists a clear and negative relationship between labor productivity and labor share, suggesting that the low-labor share plants benefit from higher labor productivity. The right Panel of Figure 9 plots plant-level average wage on the vertical axis and plant-level labor share on the horizontal axis together with a nonparametric fitted line. In contrast to the left Panel, the fitted line is mostly flat, and there is no clear relationship between these two variables. The “no-relationship” indicates that the low-labor share plants do not suppress the wages of their employees. Both panels together, the data do not support the view that the decline of labor share is caused by suppressed wages due to monopsony power or decreased bargaining power of employees.

Figure 9: Plant-level Labor Share, Labor Productivity and Average Wage



The analysis above is largely based on a simple measure of labor productivity, which is affected by various factors. Since the wage equals marginal products of labor (MPL) in a competitive environment, a more appropriate and economic theory-oriented approach

is to estimate MPL from an economic model and compare it to the wage. In the following analysis, we take this approach; we estimate the production function and quantify the evolution of MPL over time and compare it to the evolution of wage growth.

Here, the technology information is key. As discussed in Section 4.1, production function estimation without technology information may cause bias. With a similar mechanism as in the markup discussion, such bias may further result in a qualitatively different conclusion on MPL. To address such concerns, it is crucial to examine the relationship between MPL and wage with and without technology information.

Formally, we again consider the following production function;

$$Y_{it} = A_{it} K_{it}^{\beta_{kt}} L_{it}^{\beta_{lt}},$$

where  $Y_{it}$  is the physical unit of the output,  $A_{it}$  is the TFP,  $K_{it}$  is the physical capacity,  $L_{it}$  is the total number of employees, and  $\beta_{kt}$  and  $\beta_{lt}$  are the parameters to be estimated. The profit-maximizing plant solves the following problem;

$$\max_{L_{it}} P_t Y_{it} - W_t L_{it},$$

where we assume the labor input is the only variable input. The FOC of the problem induces

$$W_t = \beta_{lt} \frac{P_t Y_{it}}{L_{it}} = \text{MPL}_{it}.$$

Here, researchers can estimate MPL by substituting  $\beta_{lt}$  with an estimate,  $\hat{\beta}_{lt}$ . The industry-level MPL is then estimated by the weighted average of the firm-level MPL as

$$\text{MPL}_t = \sum_i \omega_{it} \text{MPL}_{it}.$$

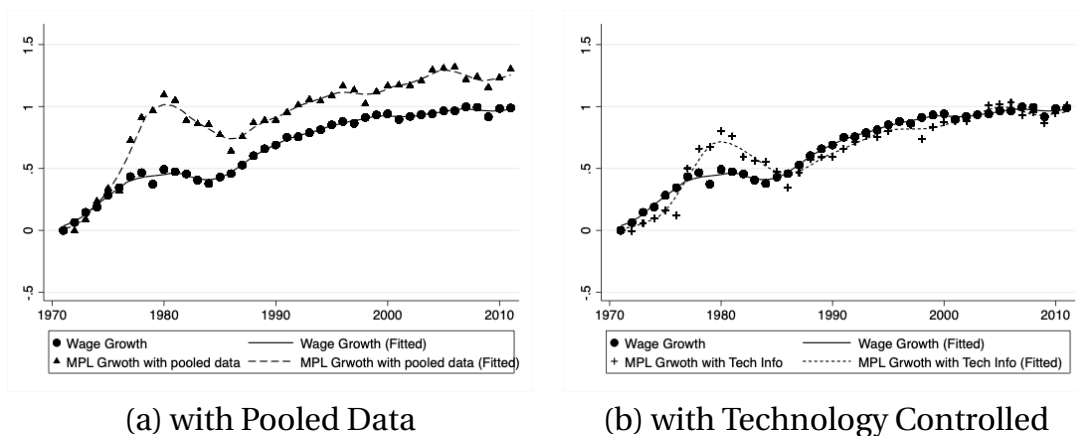
In this environment, as long as the labor market is competitive, the MPL, both at the firm level and the industry level, and the wage should grow at the same rate. However, as in the example in Section 4.1, the estimated MPL would be biased if different technologies co-exist and researchers do not have direct information on the firm-level technology. As production reallocation occurs, the industry-level MPL would fluctuate inde-

pendently from the wage growth.

Figure 10 plots the growth of industry-level real wage and MPL. On the one hand, In Panel (a), we plot them using all the data pooled and not controlling for the technology at each plant. As is clear from the plot, the growth rate of the real wage and MPL becomes apart during the period when the new technology diffuses in the industry. In a typical dataset where we do not observe the exact technology, we would reach the same observation as in the literature and find the dispersion between wage and MPL.

On the other hand, Panel (b) plots the same variables, but the production function is estimated using plants with the same technology. The plots well contrast that of Panel (a); After controlling for the plant-level technology, the wage growth and MPL growth align much closer. When production shifts from plants with labor-intensive plants to plants with capital-intensive plants, if we do not control for the technology of the plants, the growth of MPL is overestimated, which leads to a seemingly disconnected relationship. In contrast, in Panel (b), there is still some dispersion between the two variables, but these two variables grow together at a similar rate overall. These results highlight the importance of controlling for the technology to draw implications from data and the usefulness of our complementary approach.

Figure 10: Growth of Real Wage and MPL





### 4.3 Labor Share Decomposition

Thus far, we test two main alternative hypotheses on labor share decline proposed in the literature and show that, in the presence of technology information, these hypotheses can be rejected. To examine our hypothesis—the labor share decline is caused by technology diffusion—from a different angle, we now quantify the impact of technology adoption on the labor share by decomposing the change in the labor share into a technology-related component and a market-power-related component that includes market power in the product market and monopsony power (market power in labor markets).

The labor share can be expressed as

$$LS \equiv \frac{wL}{PQ} = \frac{wL}{wL + rK + \pi} = \frac{\frac{wL}{wL+rK}}{1 + \frac{\pi}{wL+rK}} = \frac{\beta_l}{1 + \frac{\pi}{wL+rK}},$$

where  $\pi$  is total profit, defined as  $\pi = PQ - (wL + rK)$ , and  $\beta_l$  is a labor coefficient of production function, Then, the change in labor share is given as

$$\begin{aligned} LS' - LS &= \frac{(wL)'}{(wL)' + (rk)' + \pi'} - \frac{wL}{wL + rK + \pi}, \\ &= \frac{\frac{(wL)'}{(wL)' + (rk)'}}{1 + \frac{\pi'}{(wL)' + (rk)'}} - \frac{\frac{wL}{wL+rK}}{1 + \frac{\pi}{wL+rK}}, \\ &= \frac{\beta_l'}{1 + \frac{\pi'}{(wL)' + (rk)'}} - \frac{\beta_l}{1 + \frac{\pi}{wL+rK}}, \\ &= \left( \frac{\beta_l'}{1 + \frac{\pi'}{(wL)' + (rk)'}} - \frac{\beta_l}{1 + \frac{\pi'}{(wL)' + (rk)'}} \right) + \left( \frac{\beta_l}{1 + \frac{\pi'}{(wL)' + (rk)'}} - \frac{\beta_l}{1 + \frac{\pi}{wL+rK}} \right). \end{aligned}$$


The first term corresponds to the change in labor share due to the change in technology (the change in labor coefficient in production function), whereas the second term corresponds to the change in labor share due to the change in market power. Moreover,  $\beta_l$  corresponds to the labor share for old technology in our production function estimation, whereas  $\beta_l'$  corresponds to the labor share for new technology.

Figure 11 demonstrates the result of labor share decomposition. The solid line plots the actual evolution of the labor share which coincides with Figure 3. The dashed line

plots the contribution of the technology adoption to the change in the labor share, whereas the dotted line plots the contribution of the change in profit. First, as discussed in Section 2 and demonstrated in Figure 11, the labor share was about 7% in the early 1970s and about 3% in the 2010s. This decomposition indicates that the labor share could have been even smaller if there were no other factors affecting the labor share. Second, we find that the other factors, including monopsony power or market power in the product market, contribute to an increase in the labor share.

These observations are also consistent with our descriptive analysis in Table 2 in two ways. First, the results in Columns (ii) and (iii) of Table 2 indicate that the labor share decreased by 2 to 4 percentage points due to technological adoption. This magnitude is identical to our findings in Figure 11. Second, when controlling for other factors through plant-fixed effects in Table 2, we find a statistically significant time trend of labor share in Columns (iii) and (iv). The magnitude is again identical to our findings in Figure 11.

Figure 11: Labor Share Decomposition



## 5 Discussions

In the previous sections, we highlight the importance of technology information. We show the technological change as the main driver of labor share decline, and only with the technology information do we reject two major alternative hypotheses proposed in the literature, such as increasing market power in the product or labor market. According to [Grossman and Oberfield \(2021\)](#), there are still two remaining hypotheses in the literature: (i) globalization and the rise of China, and (ii) changes in the composition of the workforce. In this section, we first discuss these hypotheses and then discuss the generalizability and robustness of our results.

### 5.1 Other Hypotheses: Worker Composition and Globalization

Among the hypotheses listed in the introduction of this paper, two of them are still not yet discussed; the change in worker composition and globalization. For the former,

Table 4: Employment and Payment Shares of Blue-Collar Workers

	Non-NSP Plants		NSP Plants	
	Mean	Std. Dev.	Mean	Std. Dev.
Panel (a): Employment Share of blue-collar workers				
1981	.714	.103	.681	.140
1984	.717	.100	.656	.128
1987	.697	.093	.683	.113
1990	.647	.107	.649	.121
Panel (b): Payment Share of blue-collar workers				
1981	.731	.118	.666	.125
1984	.695	.082	.646	.137
1987	.671	.101	.677	.122
1990	.592	.170	.673	.113

### 5.2 Generalizability

One might worry about how generalizable our analysis and results are. There are two aspects of generalizability: (1) the general applicability of our analysis and methodology, and (2) the generalizability of our insights to other industries and/or macroeco-

nomic analysis. For the first aspect, our analysis and methodology can be easily applied and extended to other industries as long as we can observe the technology employed by each plant or firm. Even when such technology information is not available, the recent methodological development, such as [Raval \(2022\)](#) and [Kasahara, Schrimpf and Suzuki \(2022\)](#), would still allow researchers to conduct a similar analysis.

For the second aspect, we believe that our results have generalizable insights outside the Japanese cement industry. [Kehrig and Vincent \(2021\)](#) notes that the macroeconomic patterns in data are largely coming from within industry rather than across industry, which suggests that accumulating industry-level insights helps us derive macroeconomic insights. In fact, such industry-level research has attracted attention from researchers recently such as (Automobil, Trade paper, and Cement paper (Nate and Gloria)). Though the insights from this paper might be limited, our results add one piece of solid evidence to this strand of literature, which further helps us understand macroeconomic phenomena.

## 6 Conclusion

We study the mechanism that causes the decline of labor share by investigating unusually detailed plant-level data in the cement industry in Japan. Using the exact information on the plant-level technology, we find that most of the labor share decline can be explained by the new technology diffusion: the labor share stays constant or even slightly increases over time within the same technology plants, whereas the aggregate labor share declines because the production shifts to plants with new and more capital intensive technology. We also find that the information on the plant-level technology is key to rejecting other potential hypotheses, and we would reach a qualitatively different conclusion without the information.

## References

**Abdih, Yasser and Stephan Danninger**, “What Explains the Decline of the U.S. Labor

Share of Income? An Analysis of State and Industry Level Data,” IMF Working Papers 2017/167, International Monetary Fund July 2017.

**Acemoglu, Daron and Pascual Restrepo**, “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 2020, 128 (6), 2188–2244.

**Ackerberg, Daniel A., Kevin Caves, and Garth Frazer**, “Identification Properties of Recent Production Function Estimators,” *Econometrica*, 2015, 83 (6), 2411–2451.

**Aghion, Philippe, Celine Antonin, Simon Bunel, and Xavier Jaravel**, “What are the Labor and Product Market Effects of Automation? New Evidence from France,” 2020. CEPR Discussion Paper No. DP14443.

**Autor, David, David Dorn, Lawrence F Katz, Christina Patterson, and John Van Reenen**, “The Fall of the Labor Share and the Rise of Superstar Firms,” *The Quarterly Journal of Economics*, 02 2020, 135 (2), 645–709.

**Barkai, Simcha**, “Declining Labor and Capital Shares,” *The Journal of Finance*, 2020, 75 (5), 2421–2463.

**Callaway, Brantly and Pedro H. C. Sant’Anna**, “Difference-in-Differences with Multiple Time Periods,” *Journal of Econometrics*, 2021, 225(2), 200–230.

**De Loecker, Jan and Frederic Warzynski**, “Markups and Firm-Level Export Status,” *American Economic Review*, 2012, 102 (6), 2437–2471.

– **and Paul T. Scott**, “Estimating market power Evidence from the US Brewing Industry,” *NBER Working Paper 22957*, 2016.

**Demirer, Mert**, “Production Function Estimation with Factor-Augmenting Technology: An Application to Markups,” 2022, p. 106.

**Doraszelski, Ulrich and Jordi Jaumandreu**, “Measuring the Bias of Technological Change,” *Journal of Political Economy*, 2018, 126 (3), 1027–1084. Publisher: The University of Chicago Press.

- **and** –, “Using Cost Minimization to Estimate Markups,” *CEPR Discussion Paper No. DP14114*, Available at SSRN: <https://ssrn.com/abstract=3496603>, 2019.
- Drautzburg, Thorsten, Jesús Fernández-Villaverde, and Pablo Guerrón-Quintana**, “Bargaining shocks and aggregate fluctuations,” *Journal of Economic Dynamics and Control*, 2021, 127, 104121.
- Gandhi, Amit, Salvador Navarro, and David A. Rivers**, “On the Identification of Gross Output Production Functions,” *Journal of Political Economy*, 2020, 128 (8), 2973–3016.
- Glover, Andrew and Jacob Short**, “Can capital deepening explain the global decline in labor’s share?,” *Review of Economic Dynamics*, 2020, 35, 35–53.
- Goodman-Bacon, Andrew**, “Difference-in-differences with variation in treatment timing,” *Journal of Econometrics*, 2021.
- Grossman, Gene M. and Ezra Oberfield**, “The Elusive Explanation for the Declining Labor Share,” 2021. NBER Working Papers Series.
- Humlum, Anders**, “Robot Adoption and Labor Market Dynamics,” *Working Paper*, 2021.
- Jaumandreu, Jordi**, “The Remarkable Stability of the US Manufacturing Markups,” 2022. mimeo.
- Karabarbounis, Loukas and Brent Neiman**, “The Global Decline of the Labor Share,” *The Quarterly Journal of Economics*, 2014, 129 (1), 61–103.
- Kasahara, Hiroyuki, Paul Schrimpf, and Michio Suzuki**, “Identification and Estimation of Production Functions with Unobserved Heterogeneity,” 2022. ESRI Discussion Paper Series No. 368.
- Kehrig, Matthias and Nicolas Vincent**, “The Micro-Level Anatomy of the Labor Share Decline,” *The Quarterly Journal of Economics*, 03 2021, 136 (2), 1031–1087.
- Levinsohn, James and Amil Petrin**, “Estimating Production Functions Using Inputs to Control for Unobservable,” *Review of Economic Studies*, 2003, 70, 317–341.

- Loecker, Jan De, Jan Eeckhout, and Gabriel Unger**, “The Rise of Market Power and the Macroeconomic Implications,” *The Quarterly Journal of Economics*, 01 2020, 135 (2), 561–644.
- Olley, Steven G. and Ariel Pakes**, “The Dynamics of Productivity in the Telecommunications Equipment Industry,” *Econometrica*, 1996, 64, 1263–1297.
- Raval, Devesh**, “Testing the Production Approach to Markup Estimation,” *Review of Economic Studies* (forthcoming), 2022.
- Rubens, Michael**, “Oligopsony Power and Factor-Biased Technology Adoption,” 2022.
- Stansbury, Anna and Lawrence H. Summers**, “Productivity and Pay: Is the link broken?,” *NBER Working Paper Series*, 2018.
- **and** — , “The Declining Worker Power Hypothesis: An explanation for the recent evolution of the American economy,” *NBER Working Paper Series*, 2020.
- Sun, Chang**, “Multinational production with non-neutral technologies,” *Journal of International Economics*, 2020, 123, 103294.
- Syverson, Chad**, “Macroeconomics and Market Power: Context, Implications, and Open Questions,” *Journal of Economic Perspectives*, 2019, 33 (3), 23–43.
- van Biesebroeck, Johannes**, “Productivity Dynamics with Technology Choice: An Application to Automobile Assembly,” *The Review of Economic Studies*, 2003, 70(1), 167–198.
- Yeh, Chen, Claudia Macaluso, and Brad Hershbein**, “Monopsony in the US Labor Market,” *American Economic Review*, 2022, 112 (7), 2099–2138.
- Zhang, Hongsong**, “Non-neutral Technology, Firm Heterogeneity, and Labor Demand,” *Journal of Development Economics*, June 2019, 140, 145–168.

## Appendix: Event Study Design with Two Way Fixed Effects Estimator

We further zoom into the plant-level changes in variables to confirm that our findings in the previous sections are driven by technology diffusion. To this end, we take advantage of the richness of our data, i.e., we can observe the timing of new technology adoption. Using the variation in the timing of technology adoption, we employ an event study design, i.e., difference-in-differences with leads and lags of treatment variable. Formally, we estimate the following regression estimation:

$$y_{jt} = \sum_{\tau=\tau_{min}}^{\tau_{max}} \mathbf{1}[t = t_j^* + \tau] \beta_{\tau} + \xi_j + \xi_t + \epsilon_{jt}, \quad (6)$$

where  $j$  is an index for plant,  $t$  is an index for year,  $t_j^*$  is the year plant  $j$  adopt the new technology,  $\xi_j$  is a plant fixed effect,  $\xi_t$  is a year fixed effect, and  $\epsilon_{jt}$  is an independent error term. Estimating an event study design is often called a Two-Way Fixed Effect (TWFE) estimator. For the estimator to have meaningful interpretation, the treatment effect must be homogeneous across different cohorts based on the treatment timing. See [Goodman-Bacon \(2021\)](#) for a more detailed discussion.


Here, our data structure is a typical situation of “staggered treatment timing.” One difficulty we have in our data structure is that we do not observe the timing of new technology adoption for plants that already have the new technology at the beginning of our sample period. To avoid potential bias caused by this missing data issue, we drop plants that already adopt the new technology at the beginning of our data period. Also, to balance the pre-treatment period, we drop observation more than  $\tau_{min}$  years before the treatment.

First, [Figure 12](#) plots the evolution of plant-level labor share relative to the timing of new technology adoption. The estimated coefficient for the year before the adoption is normalized to be zero. The labor share starts to decline after the technology adoption. However, the decline is not immediate. Rather, it takes several years.

To decompose the effects into the changes in employment and wages, we now look




Figure 12: Evolution of Labor Share



at the change in employment. To this end, Figure 13 plots the evolution of plant-level employment relative to the timing of new technology adoption. In contrast to the labor share, employment decreases immediately in the year of adoption, implying that the decline in the labor share is mainly driven by the change in the number of workers.

Figure 13: Evolution of Number of Workers



Finally, Figure 14 plots the evolution of plant-level capital-labor ratio relative to the timing of new technology adoption. As we see in the production function estimation results, the new technology is more capital-intensive. Therefore, we expect the capital-labor ratio to increase as plants adopt new technology. As we expect, right after the installation of NSP kilns, the capital-labor ratio jumps up by about 10% and increases slowly afterward.

Figure 14: Evolution of Capital-Labor Ratio

