The Decline of Labor Share and New Technology Diffusion: Implications for Markups and Monopsony Power*

Shoki Kusaka[†]
Tetsuji Okazaki[‡]
Ken Onishi[§]
Naoki Wakamori[¶]

March 15, 2023

Abstract

We investigate the cause of the decline in labor share using highly-detailed plant-level data from the cement industry. Using the information on production technology at each plant, we show that the major source of the decline is due to the diffusion of new technology; The labor share decreases for the plants adopting more advanced capital-intensive technology, whereas it remains steady for the plants keeping the same technology. The use of plant-level technology data is key to rejecting other potential hypotheses, such as increasing market power and monopsony power, and our conclusion would be qualitatively different without it.

JEL Classification: D2; L1; E2.
Keywords: Technology adoption; Productivity; Labor share; Markup; Monopsony.

*We thank Daisuke Adachi, Jan Eeckhout, Jacob Gramlich, Mitsuru Igami, Ebehi Iyoha, Jordi Jaumandreu, Daiji Kawaguchi, Pete Klenow, Toshihiko Mukoyama, Amil Petrin, Devesh Raval, Michael Rubens, Gloria Sheu, Chad Syverson, and participants at MEA Conferences, Federal Reserve Board, Federal Reserve Bank of St. Louis, China, Australasian, Asian meeetings of Econometric Society, EEA, IIOC, EARIE, APIOC, Washington Area Labor Economics Symposium, Kyoto University, University of Tokyo, Hitotsubashi University and Waseda University. Okazaki and Wakamori gratefully acknowledge financial support from the JSPS KAKENHI Grant Numbers 19J00587 and 22H00837. Okazaki also acknowledges financial support from Center for Advanced Research in Finance (CARF), the University of Tokyo; Research Institute of Economy, Trade and Industry (RIET); and Canon Institute for Global Studies (CIGS).

[†]Yale University. Email: shoki.kusaka@yale.edu.

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

[§]Hitotsubashi University. Email: ken.t.onishi@gmail.com.

[¶](Corresponding Author) Hitotsubashi University. Email: naoki.wakamori@r.hit-u.ac.jp.

1 Introduction

The decline of the labor share over time has been globally observed, and many economists and policymakers have paid attention to this phenomenon. An enormous number of studies have investigated this issue, proposing hypotheses to explain the phenomenon, such as factor-biased technical changes, 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.[1] Among these hypotheses, technology plays a key role; Not only some hypotheses related to technical changes, but also other hypotheses—product market power and monopsony power—hinge on technology, because estimation of markups and marginal products of labor often require estimates of production function.[2] Despite its importance, technology is typically unobserved, in particular, in the census data that are commonly used to quantify economy-wide effects in the literature.[3] The existing studies that do not control for technology, therefore, may obscure the difference between technological changes and other factors.

This paper precisely addresses this issue arising from unobservability of technology by taking a distinct and complementary approach; We collect technology information for each plant, including the timing of new technology adoption, in a particular industry. We first provide a series of evidence that "technology diffusion" is the primary driver of labor share decline by appropriately controlling for technology in our analysis. We then examine other market-power-driven hypotheses with and without technology information and show theoretically and empirically that plant-level technology information is key to rejecting these hypotheses. Without technology information, the decline of labor

[1] See Grossman and Oberfield (2022) for a more detailed summary of the literature.

[2] De Loecker and Warzynski (2012) establishes the production function approach to estimate market power, and improvement to the approach has been proposed in recent studies such as Demirer (2022), Foster, Haltiwanger and Tuttle (2022), and Raval (2022). Also, 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, Murry and Yurukoglu (2021), Miller, Osborne, Sheu and Sileo (2022), and Azar, Berry and Marinescu (2022).

[3] Most existing studies need to assume that there is an industry-wide production function common to all plants or need to infer the state of technological progress from auxiliary data indirectly. For example, Acemoglu and Restrepo (2020) construct an industry-level exposure to robots, and Aghion, Antonin, Bunel and Jaravel (2020) use the balance sheet values of industrial equipment and plant-level records of the usage of electromotive force to proxy the degree of automation.

share coincides with the increase in markups and monopsony power. However, such a relationship vanishes once controlling for technology.

Our analysis focuses on the cement industry and the diffusion of different generations of kilns, specifically from the suspension preheater (SP) kiln to the 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 macro-level 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 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 gap 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 have adopted 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. We examine how labor share and employment respond to technology adoption, and we find that they both start to fall after 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 following the installation of NSP kilns, the capital-labor ratio discretely increases.

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 using the control function approach consolidated by Ackerberg, Caves and Frazer (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. We estimate plant-level markups using the methodology from De Loecker and Warzynski (2012) and 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 MPL and wages becomes increasingly disconnected. The former finding has recently attracted many researchers' attention and has been documented in existing studies, such as De Loecker, Eeckhout and Unger (2020) and Autor, Dorn, Katz, Patterson and Van Reenen (2020). We observe a similar pattern 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 occur 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 as 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 body of literature on labor share decline.[4]

[4] Grossman and Oberfield (2022) classify hypotheses for declining labor share into the following five categories: (i) technological changes, (e.g., Karabarbounis and Neiman, 2014; Acemoglu and Restrepo, 2020; Autor et al., 2020), (ii) increased exercise of product market power by large firms (e.g., Barkai, 2020; De Loecker et al., 2020), (iii) declining worker power in labor relations (e.g., Stansbury and Summers, 2020; Drautzburg et al., 2021), (iv) globalization and the rise of China (e.g., Abdih and Danninger, 2017; Sun,

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 technological changes, in particular, introduction of computers and industrial robots, as in Acemoglu and Restrepo (2020), Autor et al. (2020), and 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.

Furthermore, 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, including Doraszelski and Jaumandreu (2018), Raval (2022), Zhang (2019), and Demirer (2022), have recently considered departures from this standard assumption. These papers highlight the importance of labor-augmenting productivity and develop ways to estimate production functions with factor-augmenting productivity change. By contrast, our paper instead assumes that producers have an explicitly different production function according to their use of old or new types of kilns, aside from any productivity differences. There are few works with this approach. Examples are van Biesebroeck (2003), who models the choice between lean or mass production in the car industry, Rubens (2022), who features the introduction of mechanical coal cutters in the 19th-century coal mining industry, and Collard-Wexler and De Loecker (2015), who collect data on production technology directly as we do in this paper.

Finally, our paper contributes to the recent literature on the time-series changes in markups. There is a growing literature on market power from macroeconomics where the literature relies on the "production approach" (see Syverson (2019)). Drawing on De Loecker and Warzynski (2012) and production function estimation from the IO literature (e.g., Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Ackerberg et al., 2015), De Loecker, Eeckhout and Unger (2020) estimate markups during the period 1955-2016 for the U.S. economy and find that they have risen steadily. Similarly, Yeh et al. (2022)

2020), and (v) changes in the composition of the workforce (e.g., Glover and Short, 2020; Acemoglu and Restrepo, 2020).

develop 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. Our findings are consistent with the findings of Jaumandreu (2022). Given the rise of the macroeconomics approach, several studies, such as Grieco, Murry and Yurukoglu (2021) and Miller, Osborne, Sheu and Sileo (2022), re-examine findings using an IO approach by focusing on specific industries as we do in this paper.

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 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. We discuss the generalizability and robustness of our results in Section 5. 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 the Japanese Cement Association.[5] They are defined by the Japanese Industrial Standards and thus can be treated as homogeneous products. 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 popularity and took a dominant position. Most of the newly built kilns in the 1960s were SP kilns, and in the 1970s, continuing improvements were made by Japanese engineering firms and cement firms, and new suspension preheater (hereinafter NSP) kilns were developed. [6] The main innovation of NSP kilns is attaching pre-calciner to the SP kilns, which breaks down CaCO3 in limestone into CaO and CO2 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 also simplifies our analysis.

[5] See https://www.jcassoc.or.jp/cement/1jpn/jc.html (Last accessed: November 25, 2022).

[6] For interested readers, Shimoda (2016) has a detailed discussion and explanation of the history of technology evolution in the cement industry.

2.2 Data Sources

For this study, we combine two complementary plant-level data sources: (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 dataset 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 in the census data, the Yearbook dataset provides such kiln-level information. By contrast, 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 are unavailable. We combine these two data sources via some common variables in both data sources. We impute the plant-level wage and material inputs before 1980 using the census data and variables that we observe throughout the entire sample period. See Appendix A for details.

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 the years 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 was 53, whereas it was 30 in 2010, implying that the number of plants decreased by about 40% over 40 years. By contrast, the

Table 1: Summary Statistics

	Num. of Obs.	Mean	Std. Dev.	Min.	Max.
Panel (a): Plant-Level Statistics (All)					
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	1,748	.549	.429	0	1
Panel (b1): Plant-Level Statistics in 1970					
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
Panel (b2): Plant-Level Statistics in 1990					
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
Panel (b3): Plant-Level Statistics in 2010					
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

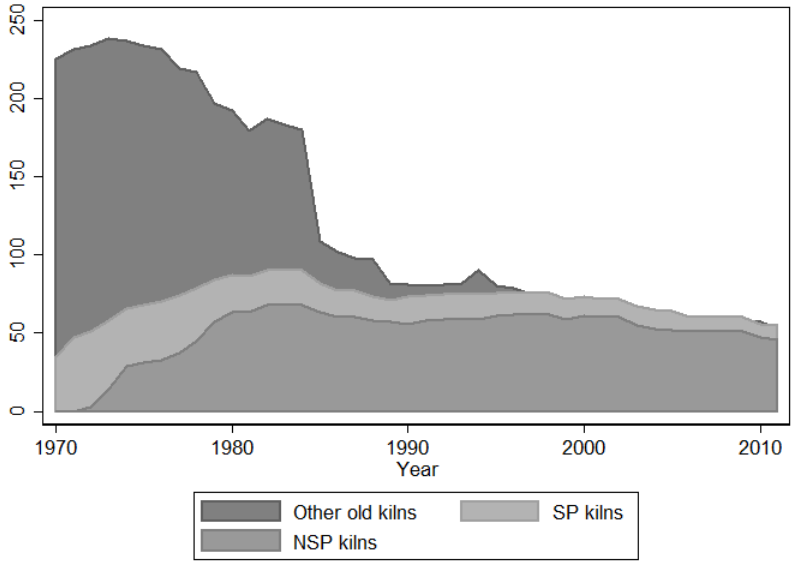
Notes: This table reports summary statistics for the following variables: Monthly (Production) Capacity, measured in ton; Annual Clinker Production, measured in ton; Average Cement Price, measured in JPY per ton; Number of Workers, measured in number of people; Average Wage per Worker, measured in JPY; and Share of NSP Kilns, defined as the fraction of NSP kilns out of total number of kilns in a given plant. Panel (a) shows summary statistics for all observations between 1970 and 2010, whereas Panels (b1) to (b3) show summary statistics for some selected years to see the variation within a year and changes in The number of observations for the number of workers is 1,672 in Panel (a), because the Cement Yearbook in 1976 does not provide the number of workers.

number of firms in 1970 and 2010 was 22 and 18, respectively, implying that most firms concentrated their production on a single plant or a small number of plants 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, by technology, i.e., types of kilns, over time. In 1970, the initial year of our sample period, there were about 220 kilns, the majority of which were kilns of old types. SP kilns accounted for less than 20%, and there were no NSP kilns. During the 1970s, 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.

Figure 1: Diffusion of Technology

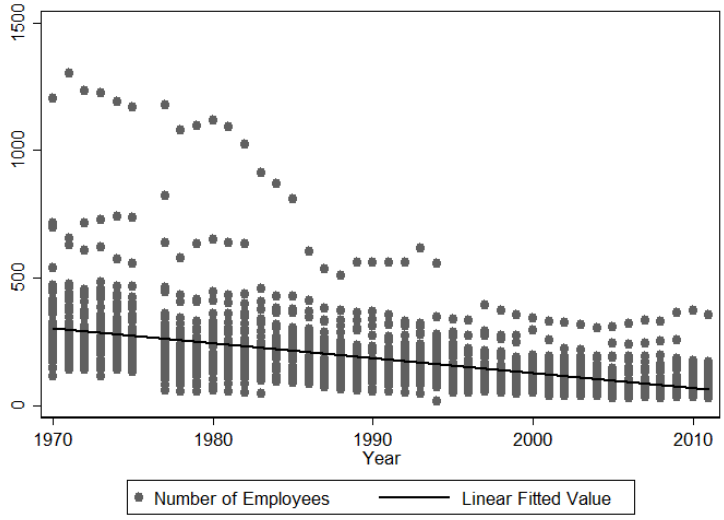


Notes: This figure graphically demonstrate that the time-series absolute number of kilns and share on the vertical axis by types of kilns.

Third, the number of workers has decreased sharply; the average number of workers per plant 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, the number of workers decreased for all plants 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. By contrast, though the average wage also in-

creased over time, the change in the average wage is not as large as the change in labor productivity. These facts raise a couple of questions: whether this reduction in the number of workers was driven by the adoption of new technology and whether the gap between growth in labor productivity and wages was due to increased monopsony power of firms in the labor market.

Figure 2: Number of Workers per Plant Over Time

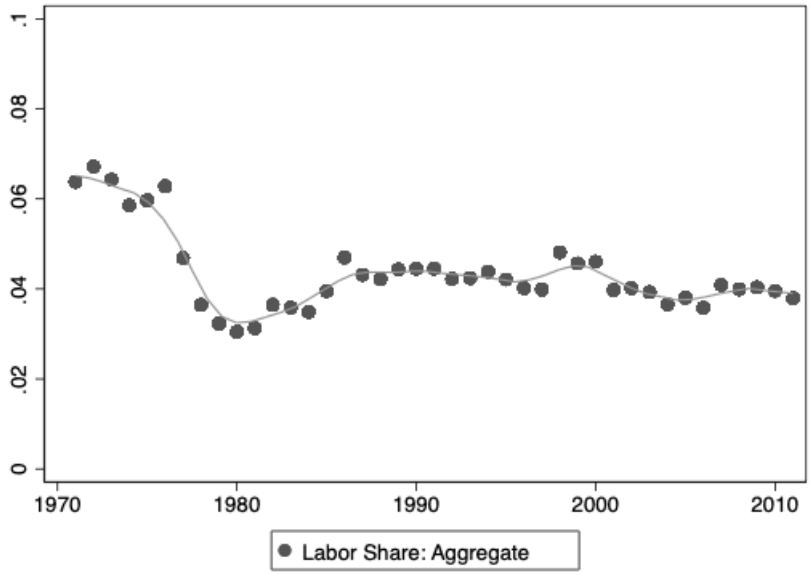


Note: This figure plots the plant-level number of employees over time together with a linear fitted value.

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 in existing studies. Then we argue that the new technology adoption is the main driver of the decline of labor share and offer evidence by implementing our 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.

Figure 3: Industry-level Labor Share



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.[7] In Figure 3, each dot represents the labor share for each year and the gray line represents the smoothed nonparametric fit.

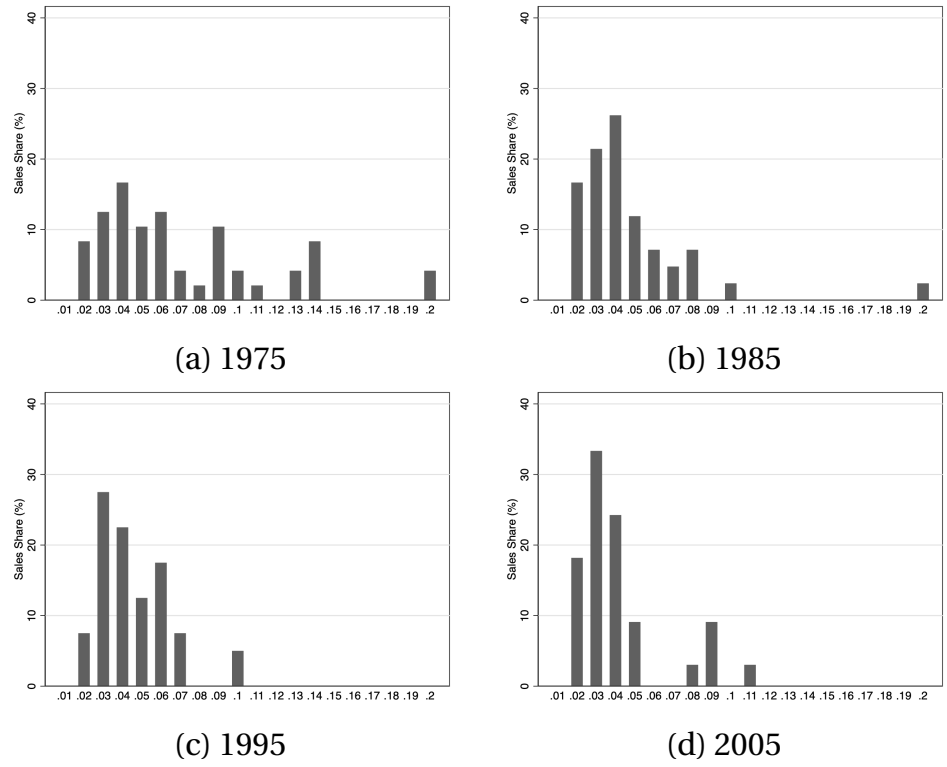
The industry-level labor share falls over our sample period with a sharper decline when the new technology diffuses between 1973 and the early 1980s, as we see in Figure 1. The presented labor share is very low. This is due to two factors. First, our definition of labor share is based on the total output value rather than the value added. We perform the same exercise with value-added as the denominator, and we confirm qualitatively the same results with a higher level of labor share, around 10%.[8] Second, the cement industry is a heavy equipment industry, and, by its nature, the labor share is lower compared with other industries.

At the firm level, output shifts from high labor-share plants to low labor-share plants.

[7] As the Census data are available only after 1980, we compute the labor share using the data in the 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.

[8] Here, value added is defined as the monetary value of total output minus the material expenditure. Since the material expenditure is only present in the census data, we use the imputed value for the 1970s.

Figure 4: Plant-level Production Share and Labor Share



(a) 1975 (b) 1985

(c) 1995 (d) 2005

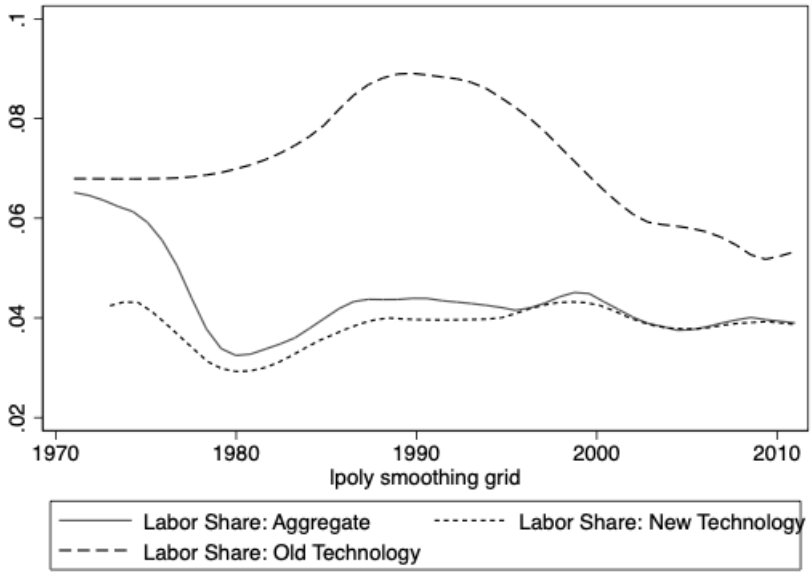
Notes: This figure shows the histogram of the output share (on the vertical axis) by plant-level labor share for the years 1975, 1985, 1995, and 2005.

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 the years 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 each plant. 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 plants with new technology (the dotted line), the average labor share within plants with old technology (the dashed line), and the industry-level labor share (the solid line). 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 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 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 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, βs are the parameters to be estimated, and ε_{it} is an independent error term. Here, we are interested in the estimated coefficient on t, i.e., β_1. We expect that β_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. By contrast, we expect that β_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 the results without control for technology, and the coefficient on year is estimated as negative and statistically significant. In the second column, once we control

Table 2: Time Trend of Labor Share

	(i)	(ii)	(iii)	(iv)
Dependent Var.	Labor Share	Labor Share	Labor Share	log(Labor Share)
Year (β_1)	-1.13$\times 10^{-3}$ ***	-2.42$\times 10^{-4}$	3.71$\times 10^{-4}$ ***	0.006***
	(0.157$\times 10^{-3}$)	(1.76$\times 10^{-4}$)	(0.880$\times 10^{-4}$)	(0.001)
NSP kiln dummy (β_2)		-0.0474***	-0.0244***	-0.389***
		(0.00469)	(0.00302)	(0.0301)
Constant	√	√	√	√
Plant fixed effects			√	√
N	1,673	1,673	1,673	1,673

Notes: Standard errors are reported in parentheses and significance levels are denoted by <0.1 (*), <0.05 (**), and <0.01 (***). The dependent variables for specifications (i) to (iii) are labor share, whereas the dependent variable for specification (iv) is logarithms of labor share.

for the technology, the significance of β_1 disappears. However, we now find that the coefficient on an NSP kiln dummy, β_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 the Cement Yearbook. Even when we use the value-added variables in the census data, we obtain the same qualitative 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 de-

sign, 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 after τ years from the treatment for the plants that adopted NSP kilns in year t is identified as:

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

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 adopts 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.[9] We define ATT τ years from the treatment as the weighted average of ATT(t,τ) as:

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

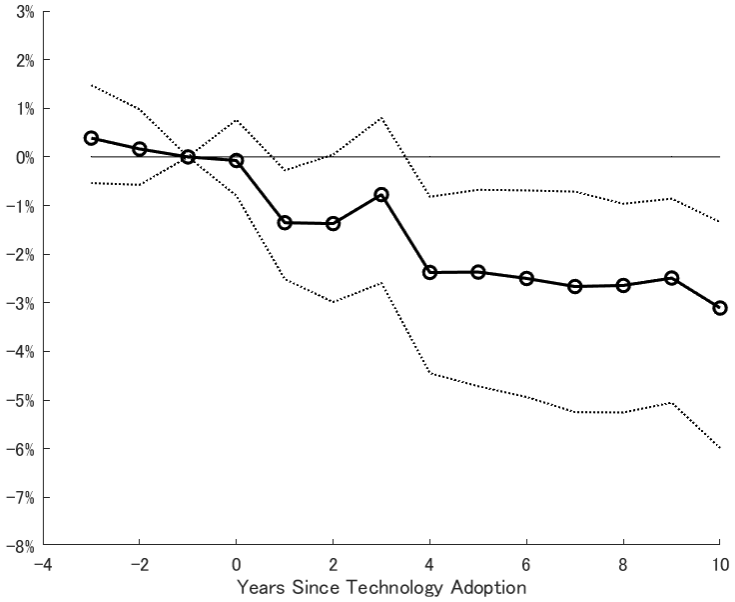
where w_t denotes the weight, which is the number of plants treated in year t divided by the total number of treated firms.

We estimate ATT(t,τ) 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 the 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 to the year of NSP adoption and the y-axis shows the estimated ATT. Here, the year

[9] Callaway and Sant'Anna (2021) 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. In addition, they provide computer codes for Stata and R to implement the estimation and 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



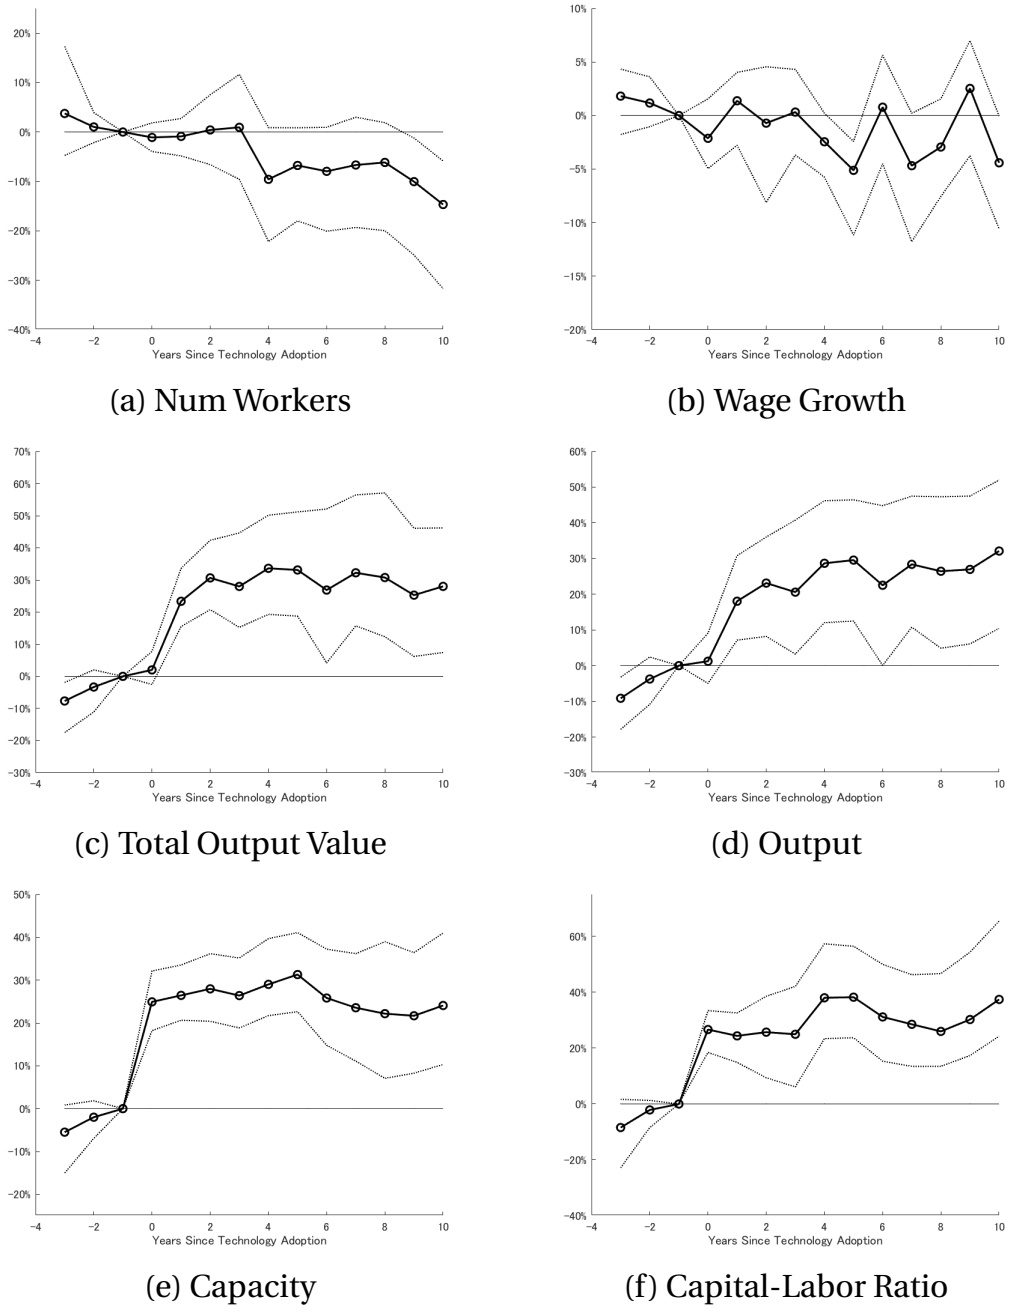
of NSP adoption is identified by the year the plant installed its first NSP kiln.[10] The solid line presents the estimated ATT and the gray dotted lines represent a 90% confidence interval. The confidence interval is constructed by the bootstrap method with 200 replications. When we look at the solid 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 preadoption 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 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 the decline in the labor share. By contrast, there is no clear difference between the treated group and the

[10] A plant typically has multiple kilns, and the adoption of NSP kilns is typically gradual, i.e., each plant replaces one or two of its kilns first and then replaces the remaining kilns over time.

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



(a) Num Workers

(b) Wage Growth

(c) Total Output Value

(d) Output

(e) Capacity

(f) Capital-Labor Ratio

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 output value relative to the timing of new technology adoption. Unlike the number of workers, output responds to the adoption relatively quickly and increases substantially.

It also grows gradually over time after the initial jump. The growth of output, together with the decline in the number of workers implies that labor productivity has increased. As the NSP kiln 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 Panels (c) and (d) are 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 increases right after the adoption and stays at a higher level compared with the preadoption 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 Panels (b) and (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 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 increases TFP. Then, what are the implications of these findings for production function? The recent literature on factor-augmenting technical changes, such as Doraszelski and Jaumandreu (2018), Zhang (2019), Raval (2022), and Demirer (2022), suggests that our findings may be explained by incorporating labor-augmenting technical changes. Although we could employ such an approach, we adopt a more straightforward one. Taking advantage of the data where we can directly observe technology for each plant, we estimate production functions separately by technology because it is more natural to assume that the shape of production functions differs across technologies.

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 τ, $\tau \in \{\text{old, new}\}$. We estimate the value-added production function that is considered by De Loecker and Scott (2016), Ackerberg 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, ω_{it} is an unobserved productivity shock, and ε_{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 Ackerberg et al. (2015):

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

The estimation procedure consists of two stages. First, we nonparametrically 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 θ including β_k^{τ} and β_l^{τ} 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 to obtain output elasticities by kiln types:

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

Standard errors are calculated by the bootstrap method with 200 replications.

When we estimate the model by pooling all plants, β_k is close to 1, and β_l is about 0.18, implying that technology exhibits economies of scale. By contrast, 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, respectively, implying that both technologies no longer exhibit economies of scale. One of the reasons for technology exhibiting economies of scale when estimating the model by pooling

Table 3: Production Function Estimates With and Without Technology Information

	(i) Pooling Both Technologies	(ii) Separately Old Tech	(ii) Separately New Tech	(iii) Pooling Both Technologies
β_k	0.971	0.778	0.907	0.872
	(0.110)	(0.110)	(0.085)	(0.071)
β_l	0.184	0.259	0.099	0.237
	(0.140)	(0.103)	(0.096)	(0.094)
β_0 (TFP Gain)	-	-	0.106	0.060
	-	-	(0.710)	(0.103)
N	1,408	1,408		1,408

both 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 $H_0 : \beta^k_{old} = \beta^k_{new}$ 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, the industry-level labor share falls consequently.

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 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 an upward bias in the scale parameter as the plant size is seemingly correlated with efficiency. With tech-

nology 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 at least partly mitigated. 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 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 (2022) classify 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 show that the increase in markups is paired with the decline of labor share. The literature follows the methods proposed by De Loecker and Warzynski (2012) and De Loecker et al. (2020) and estimates markups using a production function approach, assuming the optimality of variable inputs. Recently, a few studies (e.g., Raval, 2022; Doraszelski and Jaumandreu, 2019) have questioned whether the markup implied from cost minimization captures the actual product-level markups accurately. In this paper, we find another potential factor that may bias the estimated markups: the lack of information on plant- or firm-level technology. We find that, in the absence

of technology information, 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 a production technology characterized by $Y_i = A_i K_i^{\beta_k} L_i^{\beta_l}$. In addition, 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 wage level w, and each firm maximizes its profit by choosing its optimal level of labor input. 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 abovementioned profit maximization is

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

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

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

Note that this markup is constant and solely depends on the demand elasticity ε. In this environment, researchers can easily estimate the markup when β_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 β_l, $\hat{\beta}_l$, by

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

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 ω_i is an appropriate weight (such as the share of sales).

Now, furthermore, 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 technologies, 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 β_k and β_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

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

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, $\widetilde{\text{Markup}}_i^O < \widetilde{\text{Markup}}_i^N$, even though the markups in this environment must be identical and only depend on the demand elasticity, ε. 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 technologies, as firm-level markups are constant across firms, the industry-level markup would also be constant. When production shifts from plants with labor-intensive technology to 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 $\widetilde{\text{Markup}}_i^O < \widetilde{\text{Markup}}_i^N$. If researchers had the firm- 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 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 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

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

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

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

and the implied markup is

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

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 Ackerberg et al. (2015) and Gandhi et al. (2020). Under this specification, the markup that takes into account the material input can be expressed as

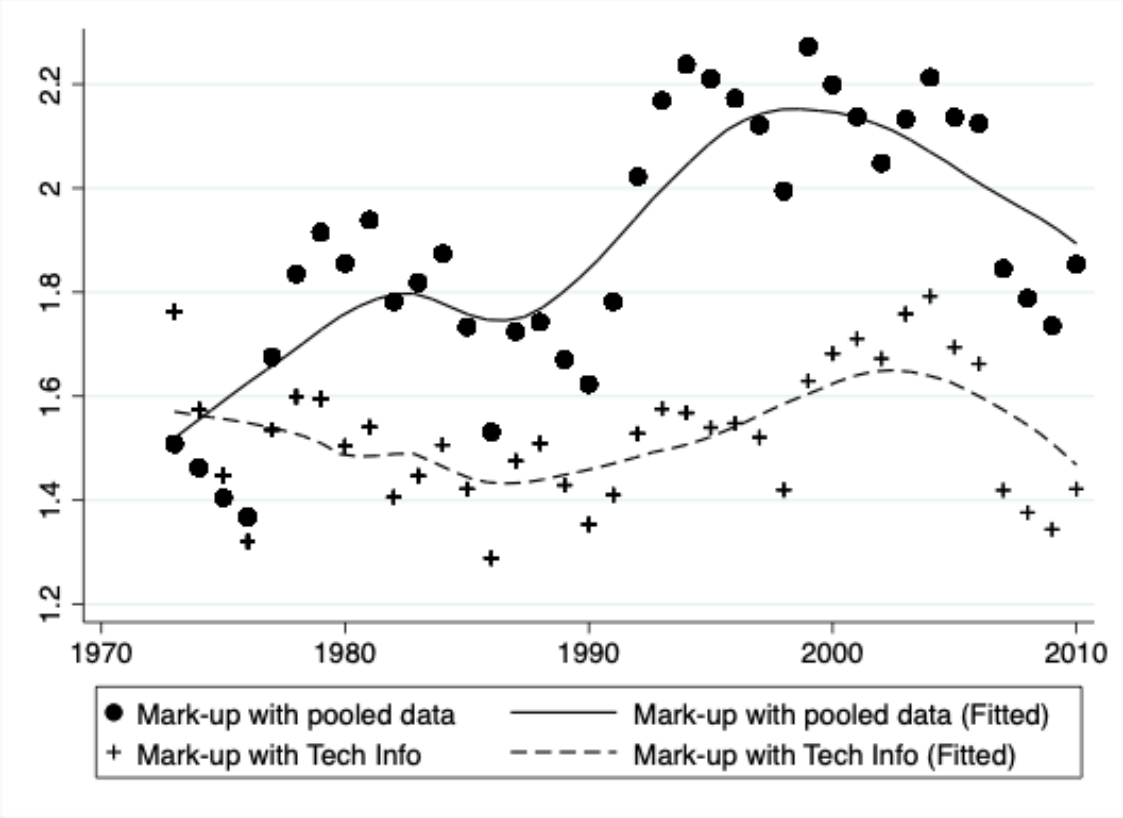
$$\text{Markup}_{it}^{M} = \frac{1}{\text{Markup}_{it}^{-1} + \frac{P^M M_{it}}{P_t Y_{it}}}, \tag{5}$$

where Markup_{it} is the markup estimates from Equation (4) and $P^M M_{it}$ is 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 production shifts to plants with new technology. By 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 implications qualitatively. Note that allowing the production function to depend on time does not help control the technology difference. [11]

[11] 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 weigh 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 of 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 as those of 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 heterogeneity explicitly into account 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

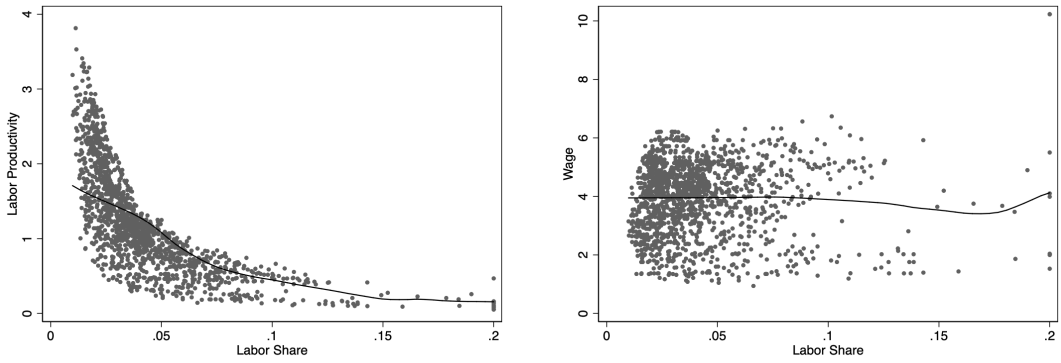
As documented in Stansbury and Summers (2018), several studies find that wedges between the growth rate of wages and the growth rate of the MPL have been increasing. The literature proposes a few explanations, e.g., a technology-driven explanation and an explanation related to worker power, such as increased monopsony power of firms 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 descriptive evidence following Kehrig and Vincent (2021), which claims that monopsony power is not likely to drive the decline of labor share, by examining the relationship between labor share, labor productivity, and wages. 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

productivity 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. Putting 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



Labor Productivity Wage

Notes: The left panel shows the relationship between the plant-level labor share on the horizontal axis and plant-level labor productivity, defined as total output value divided by the total wage payment, on the vertical axis, pooling all years between 1970 and 2010. The right panel shows the relationship between the plant-level labor share on the horizontal axis and the plant-level average wage on the vertical axis, pooling all years between 1970 and 2010.

The analysis presented above is largely based on a simple measure of labor productivity, which is affected by various factors. Because the wage equals MPL in a competitive environment, a more appropriate and economic theory-oriented approach is to estimate MPL from an economic model and compare it with 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 with 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 wages 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 β_{kt} and β_{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 first-order condition 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 β_{lt} with an estimate, $\hat{\beta}_{lt}$. The industry-level MPL is then estimated by the weighted average of the firm-level MPL as

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

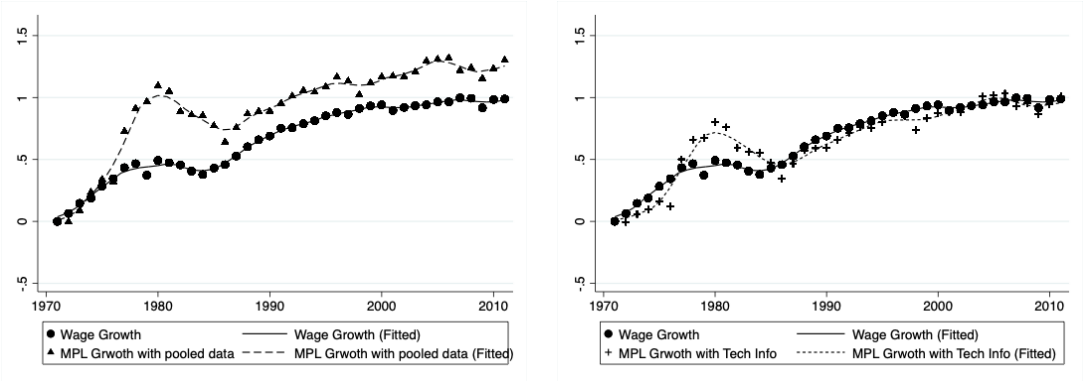
In this environment, as long as the labor market is competitive, the MPL, both at the firm level and the industry level, and the wage should grow at the same rate. However,

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 independently of 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 diverge during the period when the new technology diffuses in the industry. In a typical dataset where we do not observe technology clearly, we would reach the same observation as in the literature and find wage and MPL diverging.

On the other hand, Panel (b) plots the same variables, but the production function is estimated using plants with the same technology. The plots differ from those of Panel (a). After controlling for plant-level technology, wage growth and MPL growth are more closely aligned. When production shifts from labor-intensive plants to 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. By contrast, in Panel (b), there is still some difference 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



(a) With Pooled Data (b) With Technology Controlled

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

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

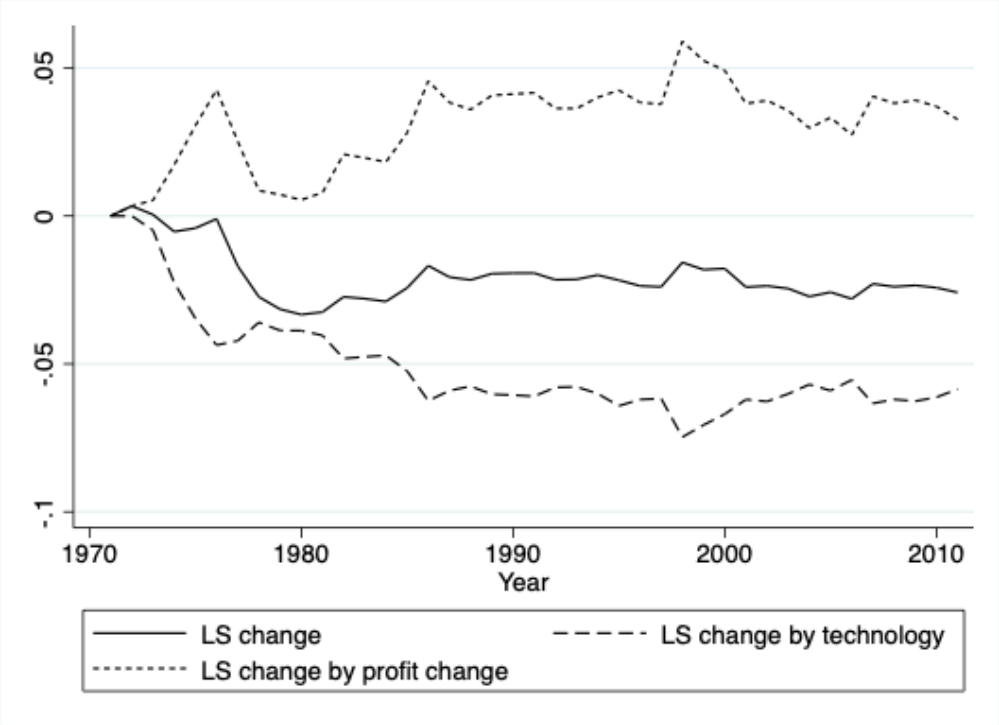
The first term corresponds to the change in labor share due to the change in technology (the change in labor coefficient in production function), whereas the second term corresponds to the change in labor share due to the change in market power. Moreover, β_l corresponds to the labor share for old technology in our production function estimation, whereas β_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 2-4 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 Discussion

In the previous sections, we highlight the importance of technology information. We show 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 (2022), 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 to this paper, we have not yet discussed the change in worker composition and globalization.

For the former point, we test whether the change in worker composition occured in the period of our focus, taking advantage of the census data that contain worker composition for some years. More specifically, the Japanese census collected the number of blue-collar and white-collar workers and the total payment bill for these workers for the years 1981, 1984, 1987, and 1990. We tabulate the employment share and payment share of blue-collar workers at non-NSP plants and NSP plants over time in Panels (a) and (b) of Table 4, respectively. These numbers immediately suggest that worker composition did not change over time, at least for these years, because both employment and payment shares at non-NSP and NSP plants are not statistically different from each other. Therefore, the changes in worker share composition cannot be a persuasive explanation for the labor share decline in our context.

In terms of a globalization hypothesis, we believe it cannot explain the labor share decline in this specific industry, as the import and export of cement were not important in the period of our focus. Even though there are several papers that focus on the cement industry and emphasize the importance of international competition, including Miller and

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

Osborne (2014) who show that import competition affects prices and Salvo (2010) who shows that the potential "threat" of import competition restricts market power, Japan is geographically isolated from other countries. Less than 10%, at maximum, of total cement production was exported to other Asian countries, and there is almost no import from other countries in the period of our focus, according to Okazaki et al. (2022). Therefore, globalization cannot be a major concern in the Japanese cement industry.

5.2 Generalizability

One might worry about the generalizability of our analysis and results. 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 macroeconomic 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 unavailable, recent methodological development, such as by 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) note that the macroeconomic patterns in data largely come from within industry rather than across industries, 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 Grieco et al. (2021), Miller et al. (2022), and Ganapati (2022). 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 of the cement industry in Japan. Using the information on 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 production shifts to plants with new and more capital intensive technology. We also find that the information on plant-level technology is key to rejecting other potential hypotheses, and we would reach a qualitatively different conclusion without that 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," International Monetary Fund Working Papers 2017/167, 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," CEPR Press Discussion Paper No. 14443, 2020.

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*, 2020, *135* (2), 645–709.

Azar, José A, Steven T Berry, and Ioana Marinescu, "Estimating Labor Market Power," NBER Working Paper No. 30365, 2022.

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.

Collard-Wexler, Allan and Jan De Loecker, "Reallocation and Technology: Evidence from the US Steel Industry," *American Economic Review*, 2015, *105* (1), 131–171.

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 No. 22957, 2016.

_ **, Jan Eeckhout, and Gabriel Unger**, "The Rise of Market Power and the Macroeconomic Implications," *The Quarterly Journal of Economics*, 2020, *135* (2), 561–644.

Demirer, Mert, "Production Function Estimation with Factor-Augmenting Technology: An Application to Markups," Unpublished Manuscript, 2022.

Doraszelski, Ulrich and Jordi Jaumandreu, "Measuring the Bias of Technological Change," *Journal of Political Economy*, 2018, *126* (3), 1027–1084.

_ **and** _ , "Using Cost Minimization to Estimate Markups," CEPR Press Discussion Paper No. 14114, 2019.

Drautzburg, Thorsten, Jesús Fernández-Villaverde, and Pablo Guerrón-Quintana, "Bargaining shocks and aggregate fluctuations," *Journal of Economic Dynamics and Control*, 2021, *127*, 104121.

Foster, Lucia S., John C. Haltiwanger, and Cody Tuttle, "Rising Markups or Changing Technology?," NBER Working Paper No. 30491, 2022.

Ganapati, Sharat, "The Modern Wholesaler: Global Sourcing, Domestic Distribution, and Scale Economies," U.S. Census Bureau, Center for Economic Studies Working Paper 18-49, 2022.

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, *225* (2), 254–277.

Grieco, Paul L. E, Charles Murry, and Ali Yurukoglu, "The Evolution of Market Power in the US Auto Industry," NBER Working Paper No. 29013, 2021.

Grossman, Gene M. and Ezra Oberfield, "The Elusive Explanation for the Declining Labor Share," *Annual Review of Economics*, 2022, *14* (1), 93–124.

Humlum, Anders, "Robot Adoption and Labor Market Dynamics," Unpublished Manuscript, 2021.

Jaumandreu, Jordi, "The Remarkable Stability of the US Manufacturing Markups," CEPR Press Discussion Paper No. 17490, 2022.

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," ESRI Discussion Paper Series No. 368, 2022.

Kehrig, Matthias and Nicolas Vincent, "The Micro-Level Anatomy of the Labor Share Decline," *The Quarterly Journal of Economics*, 2021, *136* (2), 1031–1087.

Levinsohn, James and Amil Petrin, "Estimating Production Functions Using Inputs to Control for Unobserbables," *The Review of Economic Studies*, 2003, *70*, 317–341.

Miller, Nathan H. and Matthew Osborne, "Spatial Differentiation and Price Discrimination in the Cement Industry: Evidence from a Structural Model," *RAND Journal of Economics*, 2014, *45* (2), 221–247.

Miller, Nathan, Matthew Osborne, Gloria Sheu, and Gretchen Sileo, "The Evolution of Concentration and Markups in the United States Cement Industry," Unpublished Manuscript, 2022.

Okazaki, Tetsuji, Ken Onishi, and Naoki Wakamori, "Excess Capacity and Effectiveness of Policy Interventions: Evidence from the Cement Industry," *International Economic Review*, 2022, *63* (2), 883–915.

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," NBER Working Paper No. 30586, 2022.

Salvo, Alberto, "Inferring market power under the threat of entry: the case of the Brazilian cement industry," *RAND Journal of Economics*, 2010, *41* (2), 326–350.

Shimoda, Takashi, "History of Cement Manufacturing Technology," in Center of the History of Japanese Industrial Technology, ed., *Survey Reports on the Systematization of Technologies*, Vol. 23, Tokyo: National Museum of Nature and Science, 2016, pp. 1–114.

Stansbury, Anna and Lawrence H. Summers, "Productivity and Pay: Is the link broken?," NBER Working Paper No. 24165, 2018.

_ **and** _ , "The Declining Worker Power Hypothesis: An explanation for the recent evolution of the American economy," NBER Working Paper No. 27193, 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*, 2019, *140*, 145–168.

Appendix A Imputing Missing Variables

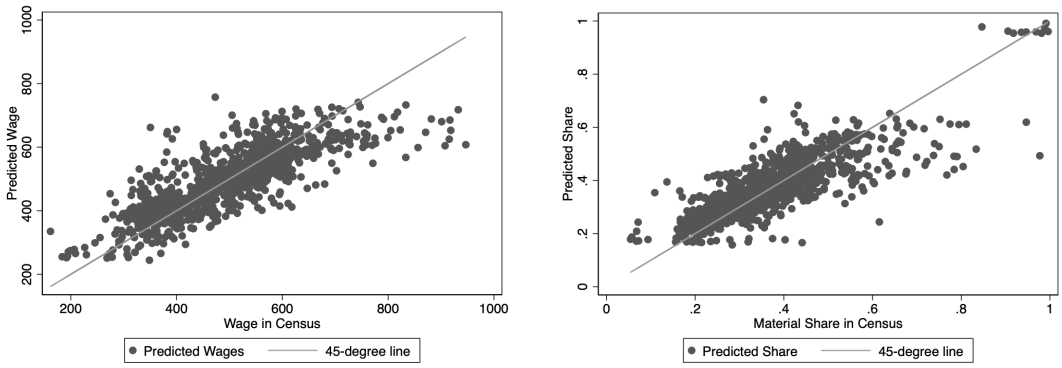
In our analyses, we combine two plant-level data sources: (i) Cement Yearbook(Cement Nenkan), published by the Cement Press Co. Ltd. (Cement Shinbunsha), and (ii) Census of Manufacture, collected by the Japanese Ministry of Economy, Trade, and Industry. The sample periods for these two data sources are 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 are unavailable. We impute the plant-level wage and intermediate inputs before 1980 using the census data and variables that we observe throughout the entire sample period.

Plant-level wages from 1970 to 1979 are imputed using prefecture-level wages in the industry, which are available for 1970 to 2010, and plant fixed effects. We regress census wages on prefecture-level wage and plant fixed effects using the period between 1980 and 2010 and predict census wages from 1970 to 1979. We confirm that the prediction matches actual values for 1980-2010 well. The results of our main analysis do not change when we use prefecture-level wages for the entire sample period.

We imputed intermediate input expenditure between 1970 and 1979 as follows. First, we calculate the sales share of the expenditure of intermediate materials, including energy expenses. Then, we take the logit function of this share, $\log(\frac{s}{1-s})$. We regress it on the set of explanatory variables, plant fixed effects, the indicator function whether a plant uses NSP kilns, the number of kilns in the plant, the share of NSP kilns in all the kilns a plant uses, and oil prices. We also control time trends flexibly. After the regression, we predict $\log(\frac{s}{1-s})$ and recover the predicted material share $\hat{s}$ for 1970-1979. This procedure guarantees that the predicted material expenditure does not exceed the value of cement produced.

Figure A1 shows the fit of the prediction for wages and material shares. The x-axis of Panel (a) is wage levels in the census and that of Panel (b) is material shares in the census. The y-axis indicates the predicted value. For both two variables, the dots concentrate on the 45-degree line, which implies that the performance of imputation is good enough.

Figure A1: The Prediction of Imputed Variables



(a) Wages (b) Material Share

Appendix B Event Study Design with Two-Way Fixed-Effects Estimators

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:

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

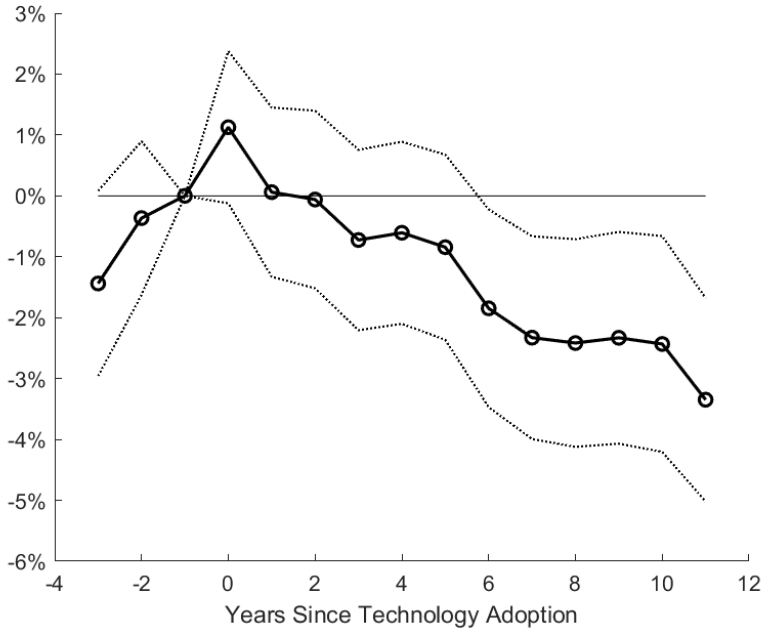
where j is an index for plant, t is an index for year, t_j^* is the year plant j adopts the new technology, ξ_j is a plant fixed effect, ξ_t is a year fixed effect, and ϵ_{jt} is an independent error term. Estimating an event study design in this way is often called a Two-Way Fixed Effect (TWFE) estimator. 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 remove plants that had already adopted the new technology at the beginning of our data period.

To balance the pretreatment period, we drop observation more than τ_{min} years before the treatment.

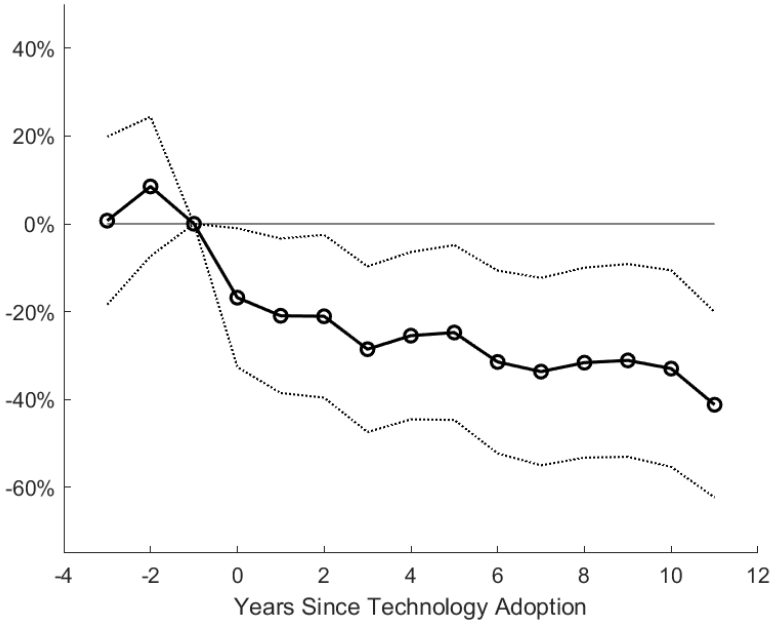
First, Figure B1 plots the evolution of plant-level labor share relative to the timing of new technology adoption. The estimated coefficient for the year before adoption is normalized to zero. The labor share starts to decline after the technology adoption. However, the decline is not immediate. Rather, it takes several years.

Figure B1: Evolution of Labor Share



Second, to decompose the effects into the changes in employment and wages, we now look at the change in employment. To this end, Figure B2 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 B2: Evolution of Number of Workers



Finally, Figure B3 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 B3: Evolution of Capital-Labor Ratio

