The Decline of Labor Share and New Technology Diffusion*

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

February 11, 2022

Abstract

We study the mechanism behind the decline in the 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) firms exercise monopsony power in the labor market and (ii) the decline in labor share is associated with an increase in mark-ups. We reject these two hypotheses with technology information but may not without it.

JEL Classification: D2; L1; E2.

Keywords: Technology adoption; Productivity; Labor Share; Mark-up; Monopsony.

*We thank to Daisuke Adachi, Jacob Gramlich, Daiji Kawaguchi, Pete Klenow, Toshihiko Mukoyama, Gloria Sheu, and seminar participants at Federal Reserve Board, IIOC, and etc. 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.

†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.

§Federal Reserve Board. Email: ken.t.onishi@frb.gov.

¶The University of Tokyo. Email: nwakamo@e.u-tokyo.ac.jp.

1 Introduction

The decline of the labor share is a phenomenon observed globally, and it has attracted attention from both policy makers and researchers. An enormous number of studies has investigated this issue and proposed explanations 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), globalization and the rise of China (e.g., Abdih and Danninger, 2017; Sun, 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), and changes in the composition of the workforce (e.g., Glover and Short, 2020; Acemoglu and Restrepo, 2020).[1]

We tackle this issue and propose "technology diffusion" as a main driver of this phenomenon, by taking a distinct and complementary approach to the existing studies: collecting plant-level data on production technologies and directly controlling for technology in our analysis. A large fraction of the existing studies take a macroeconomic approach which quantifies economy-wide effects. To implement an economy-wide analysis, the most detailed data available to researchers would be plant-level census data. However, with these data, researchers may still face the well-known difficulty of measuring technical change. As precise information on production technology is typically unobserved in the data, researchers need to infer the state of technological progress indirectly from auxiliary data.[2] We overcome this difficulty by using exact plant-level technology information, including the timing of new technology adoption.

We focus on a specific industry in order to take our approach because there are no such data that cover all industries. 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

[1] See Grossman and Oberfield (2021) for more detailed summary of the literature.

[2] For example, Acemoglu and Restrepo (2020) constructs an industry-level exposure to robots and Aghion et al. (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.

macroeconomic implications. For this purpose, the cement industry provides us an ideal environment because plant-level technology, the type of kilns, is easily observed, as well as the diffusion of different generations of kilns, specifically from Suspension Preheater (SP) Kiln to New Suspension Preheater (NSP) Kiln. 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 increased discrepancy between labor productivity growth and wage growth, and an increase in industry-wide markups paired with the decline of the labor share.

In the analysis, we find that technology diffusion explains most of the important patterns in the data; the industry-level labor share declines over time. However, the industry-level decline of the labor share is largely explained by the diffusion of a new and more capital-intensive technology. Within the plants with the same old technology, the labor share is slightly increasing. Consistent with the findings in Kehrig and Vincent (2021), we also find that the low labor share plants benefit from high labor productivity, not from low wages. These observations suggest that the new technology is more capital intensive and suggests a different shape of the production function rather than a simple increase in total factor productivity (TFP). We confirm that when we estimate the production function for each technology, we find that the new technology is indeed more capital intensive.

Also, we find that our conclusion would be qualitatively different if we lacked data on production technology, which highlights the importance of our approach. We find that without considering the differences in production technology (i) the growth rate of the marginal productivity of labor (MPL) and wage becomes increasingly disconnected and (ii) the labor share decline is paired with an increase in the markup. The former pattern is well documented in the literature (for example in Stansbury and Summers (2018)) and researchers and policy makers have been debating 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 is a result of production technology heterogeneity, and the discrepancy vanishes once we control for plant-level technology. The second finding has attracted attention recently and has been documented in existing studies, such as De Loecker et al. (2020) and Autor et al. (2020). We find that a similar pattern exists when the technology information is missing. 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 input-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 mark-ups disappears.

In order to confirm that technology diffusion explains our findings, we employ an event study design using observed variation in the timing of technology adoption. Specifically, we use Difference-in-Differences with leads and lags of the treatment variable. We first examine how the 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 do not find any statistically significant pre-trend in the variables, suggesting that our research design is valid to make a causal inference.

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 as an explanation for the decline in the labor share in Section 3 together with alternative hypotheses. We further confirm our hypothesis using a event study design in Section 4. Section 5 concludes.

2 Industry Backgrounds and Data

In this paper, we focus on a specific industry, namely the Japanese cement Industry. 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 complement these studies, we therefore take a distinct approach—focusing on one specific industry, the cement industry, and collecting production technology information that we typically cannot observe in the standard census data. Although one might worry about generalizability, we believe that accumulating micro-level hard evidence helps us understand macroeconomic phenomena. Moreover, in addition to the observability of production technology, there are two more advantages for studying the cement industry: (i) homogeneity of the product, which enables us to estimate markups accurately, and (ii) a simple production process, which enables us to estimate productivity easily through production function estimation. In the following subsections, we first describe the industry backgrounds and 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 ingredients for construction work, as concrete and mortar are made from cement. There are several types of cement. For example, Portland cement, the focus of this paper, is the most common type of cement, accounted for about 75% of cement product, according to Japan Cement Association. Though there are several different types of cement, each of them is standardized by the Japanese Industrial Standards and thus can be treated as 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. The final procedure of mixing ground clinker with gypsum produces cement. Cement kilns are the heart of this simple production process, and it is important for us to understand some technological aspects of cement kilns in Japan. Even though there are several types of kilns, we can roughly categorize them into two types: dry process kilns and wet process kilns.[3] Dry process kilns were developed in the late 19th century, and wet process kilns became dominant in subsequent periods. In the 1960s, the suspension preheater (SP) process, part of the dry 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, continuing improvements were made by Japanese companies, and new suspension preheater (hereinafter

[3] More precisely, there are also semi-dry and semi-wet process kilns. See Shimoda (2016).

NSP) kilns were developed. In our data, after 1970, almost all newly built kilns were NSP kilns, and this homogeneity of investment simplifies our analysis.

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 ton), production output (in ton), number of workers, and ownership and geographical location of the plants. In addition to these basic characteristics of the plants, the data also contain the types and the number of kilns that each plant owns, which make this data special. Although technology each plant employs is typically unobserved, this Yearbook data provides such kiln-level information.

On the other hand, Census of Manufacture provides similar but slightly different 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 2011, whereas we obtain the latter data from 1980 to 2011, 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 of plants are located in a different city with very few exceptions.

2.3 Summary Statistics and Key Features

Summary statistics of our data are given in Table 1; Panel (a) presents a couple of year-level statistics, whereas Panel (b) presents plant-level statistics. In Panel (a), there are three variables: the number of plants observed in each year, the fraction of NSP kilns in each year, and the average cement prices. The number of plants observed in each year ranges from 30 to 54. There were initially 54 plants in 1970, whereas there were only 30 plants remain in 2011, implying that about a half of them exits from the market over the

Table 1: Summary Statistics

	Num. of Obs.	Mean	Std. Dev.	Min	Max
Panel (a): Year-Level Statistics					
Observation Year	42	–	–	1970	2011
# of Plants	42	42	7	30	54
Fraction of NSP Kilns (%)	42	62%	28%	0%	94%
Average Cement Price (JPY/t)	39	11,014	2,087	8,705	16,536
Panel (b): Plant-Level Statistics					
Monthly Capacity (tons)	1,748	188,716	124,935	18,180	69,6250
Annual Clinker Production (tons)	1,748	1,740,715	1,294,285	13,250	8,082,269
# of Workers (person)	1,673	193	140	16	1,303
Average Wage	1,673	3.77	1.50	1.27	23.84

Note: The number of observation for the average cement price is 39 because the price information is missing for 1970 to 1972. Monthly capacity, whereas clinker production is annualized.

sample period.

The fraction of NSP kilns in each year, listed in the second row, also varies, ranging from 0% to 94%. There were no NSP kilns in Japan in 1970, whereas the old kilns were mostly replaced to NSP kilns over 40 years, reaching to 94%. 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 less than 20%. Note, again, that there were no NSP kilns in 1970. During the 70s, however, NSP kilns dramatically increased its popularity, maintaining its dominant position after 1980s. In our main analysis, we explore the labor share with and without controlling for this technology information.

The average cement prices, listed in the third row in Panel (a) of Table 1 also varies across years, though the change is not monotonic unlike aforementioned two variables. To see the change in the average cement price, we plot the prices together with the total production quantity in Figure 2. Basically, there are three phases in the price movement. From 1970 to 1980, due to the two oil crises, the cement prices increased dramatically. From 1980 to 2000, the prices decreased shapely, even lower than the level of 1970. However, since 2000, the price mildly increased. The total quantity moved parallel with the prices from 1970 to 1990, whereas it has moved counter-cyclically since 1990.

Through the second to fourth rows of Panel (b), we show the summary statistics for

Figure 1: Diffusion of Technology

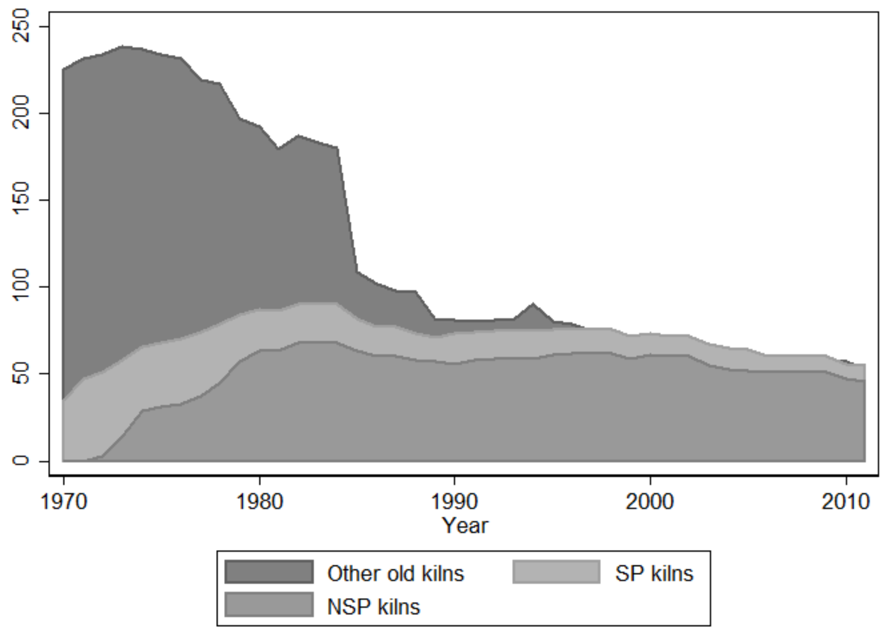

monthly capacity defined as how much clinker a plant can produce when operating for 600 hours per month, the number of workers, and average wage. One important pattern to note is that we see a dramatic decrease in the number of workers. To make this point clear, Figure 3 plots the plant-level number 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. For example, in 1970, the average number of workers was about 320, but in 1995 this had fallen to 150 which is a half of the number in 1970. This change indicates that there was substantial technological advancement in the form of automation and, consequently, labor productivity increased sharply.

3 Decline of Labor Share and Existing Hypotheses

In this section, we first present the patterns in the data that the existing studies have documented. Then, we summarize the existing explanation/hypotheses on why we observe the decline of labor share and examine whether we can reject some of the hypotheses with and without the information on plant-level technology.

Figure 2: Industry Evolution over Time

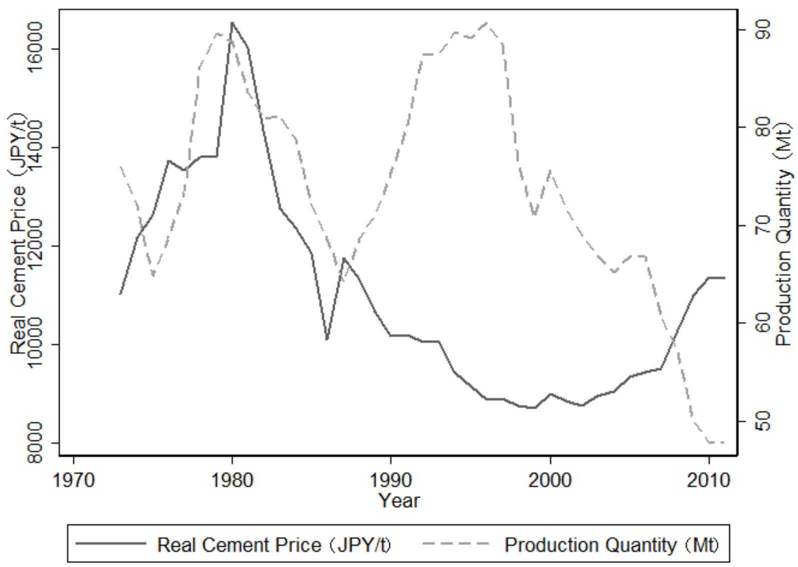

Note: This figure plots the annual average price of cement and the annual total production quantity of Clinker. The cement price is converted to the real price level in 2000 using GDP deflator.

3.1 The Decline of Labor Share

We first plot the industry-level labor share. Since the Census data is available only after 1980, we are not able to compute the value added in the earlier years. Therefore, we define the labor share by the total wage payment divided by the monetary value of total output. Figure 4 plots the actual data and local polynomial-smoothed version of industry-level labor share. The industry-level labor share falls sharply during the period when the new technology diffuses.

At the firm-level, the output shifts from high-labor share plants to low-labor share plants. Figure 5 plots the histogram of the share of output on the vertical axis and the plant-level labor share on the horizontal axis in different years. From 1975 to 2005, the distribution shits from the right to the left, which implies that the production shifts to plants with high labor share to plants with low labor share.

Figure 3: Number of Workers per Plant over Time

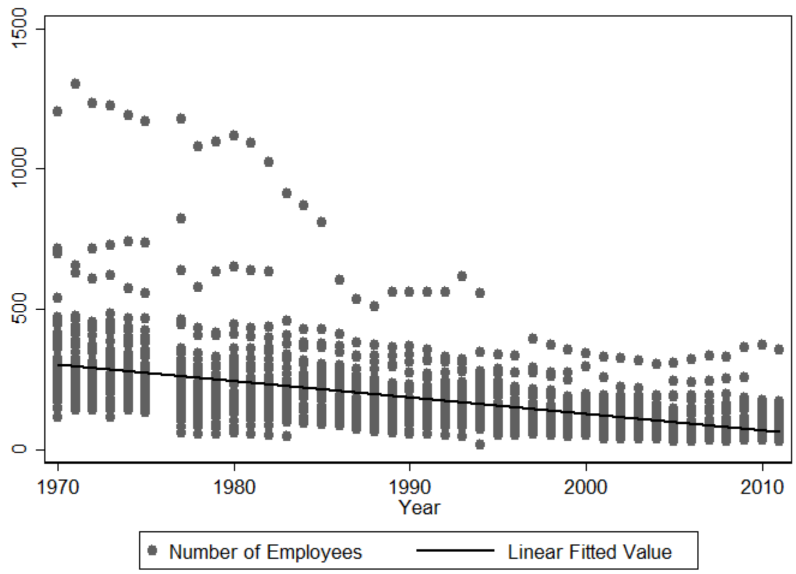

Note: This figure plots the plant-level number of employees over time together with a linear fitted value.

3.2 The Existing Hypotheses

In this section, we present the existing hypotheses and demonstrate how our new approach help us distinguish/reject them.

Technology Change

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 4 *conditional on* the plant-level technology. In Figure 6, we plots the average labor share within the plants with new technology, within the plants with old technology, and the industry-level labor share. 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 lower level of labor share. Figure 6 clearly shows that the decline of labor share is caused by the new technology diffusion.

To assess the argument more quantitatively, we estimate the following equations us-

Figure 4: Industry-level Labor Share

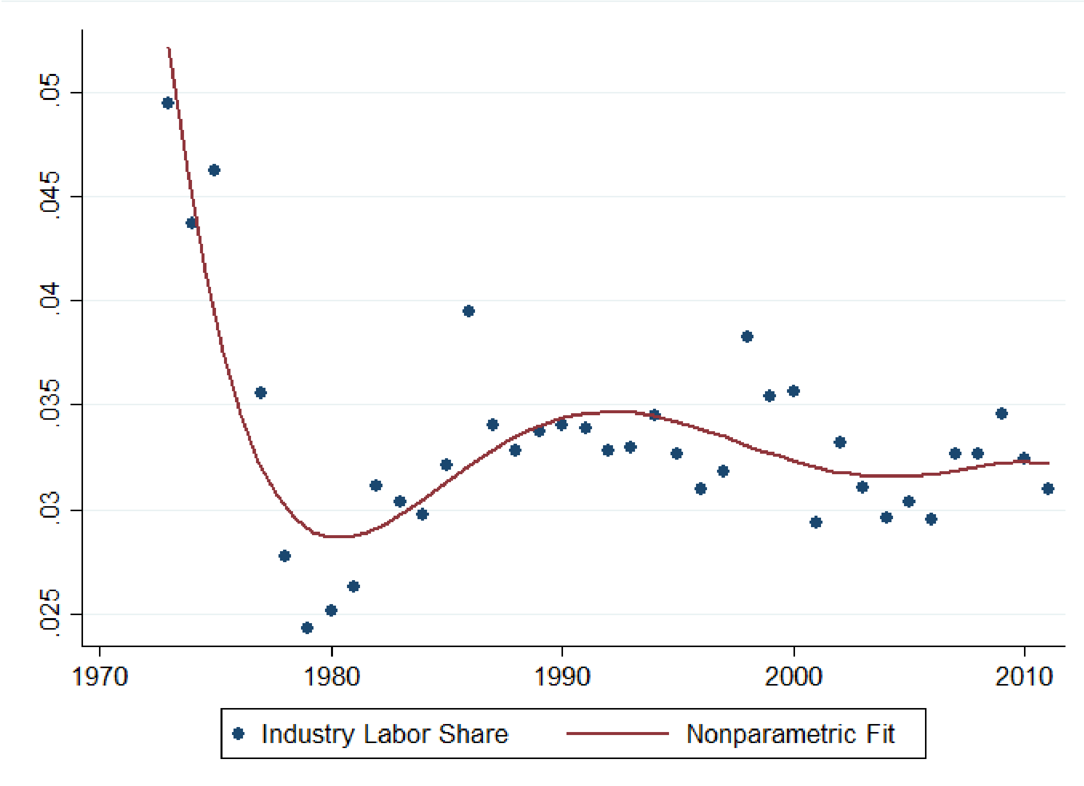

ing the plant level labor share by Ordinary Least Squares (OLS);

$$\text{LaborShare}_{it} = \beta_0 + \beta_1 t + X_{it}\gamma + \varepsilon_{it},$$

where i is a plant index, t denotes year, X_{it} is other plant-level control variables, βs and γ are the parameters to be estimated, and ε_{it} is an independent error term. Here, we are interested in the estimated coefficient on t. We expect that β_1 would be estimated negative when we do not control for the plant-level technology because the industry-level labor share declines over time. In contrast, we would 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 expectation. The first column presents the results without controls for technology, and the coefficient on year is estimated negative and statistically significant. In the second column, once we control for the technology, the significance disappears. When we further control for the plant fixed effects, the estimates

Figure 5: Plant-level Production Share and Labor Share

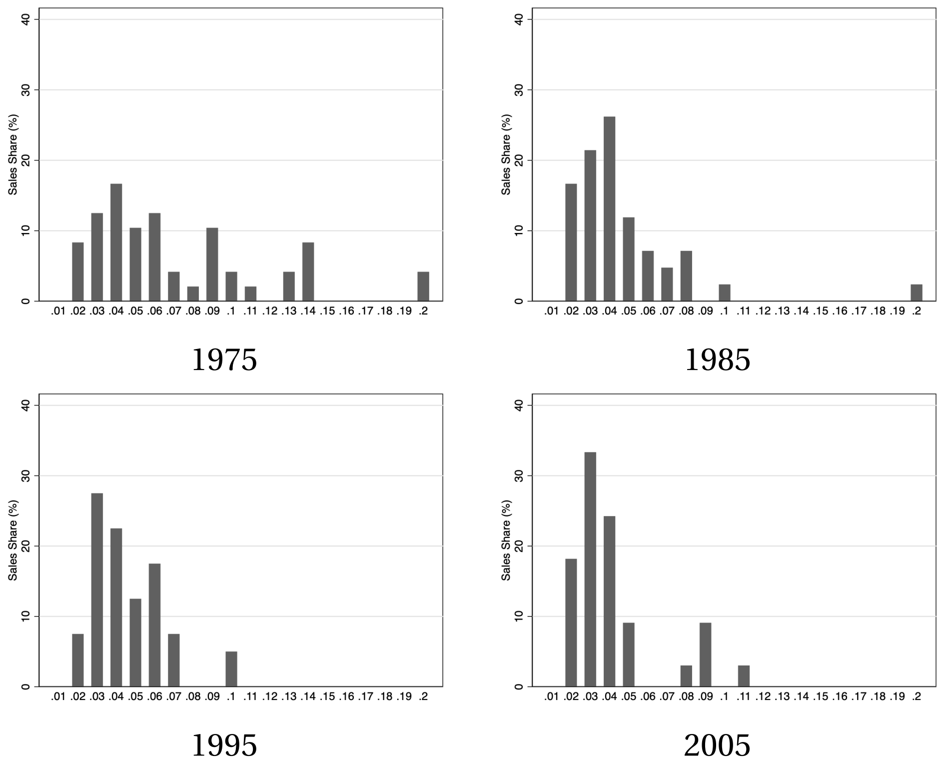

becomes positive and statistically significant. These results are consistent with Figure 6. To quantify the economic significance of the results in the third column, we replace the left hand side variable by the logarithm of labor share, which allows us to quantify the percentage change easily. The result is presented in the fourth column, suggesting that the labor share increases at the plant level by 0.7% every year. The magnitude is not very large but not small either.

The decline of labor share is very difficult to rationalize if the new technology is simply an increase in TFP. Rather, it is natural to assume that the shape of production function changes as plants adopt new technology and the new technology is more capital intensive. A natural starting point to address this issue is to estimate production functions for both new and old technology separately. Here, we assume that the production function is Cobb-Douglas form;

$$PY_{it} = A_{it} K_{it}^{\alpha} L_{it}^{\beta},$$

where PY_{it} is the monetary value of the output, A_{it} is the TFP, K_{it} is the physical capacity,

Figure 6: Labor Share Conditional on the Plant-level Technology

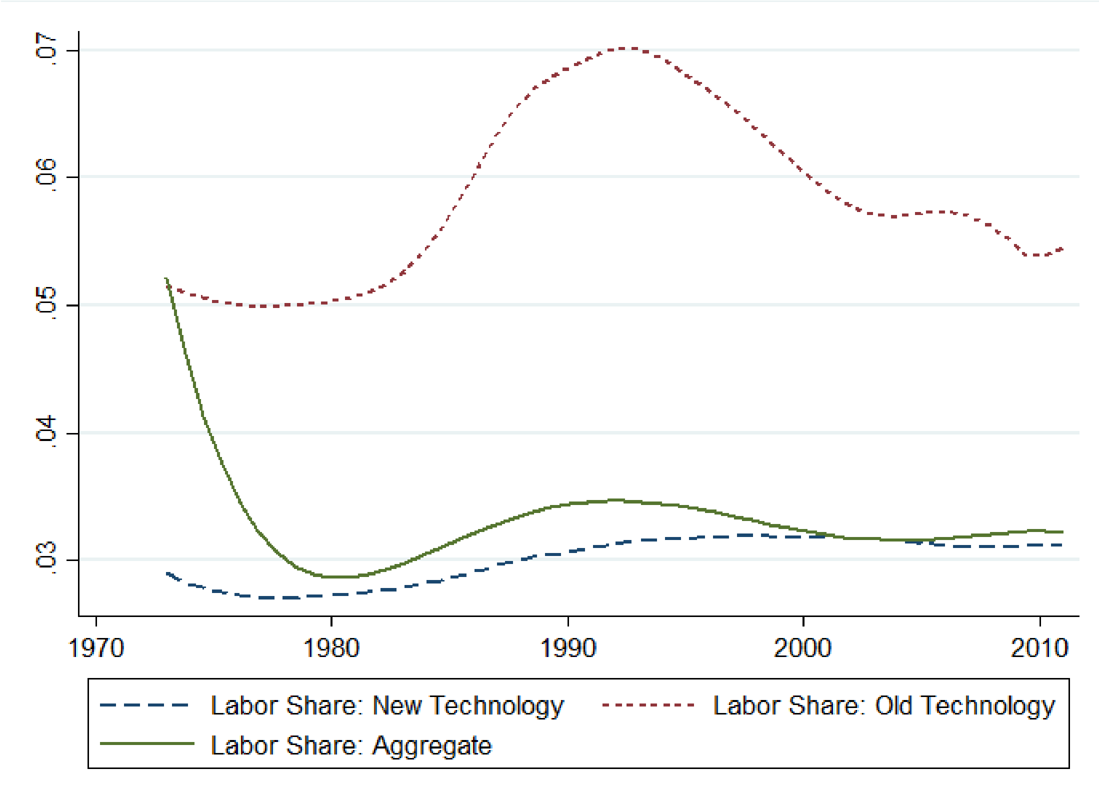

L_{it} is the total wage payment, and α and β are the parameters to be estimated. Table 3 summarizes the estimation results. As we expect, the new technology is more capital intensive. Therefore, profit maximizing plants would need less labor, which results in the lower level of labor share. When more plants adopt new technology, as a result, the industry-level labor share falls.

Growing Dispersion between Labor Productivity and Wage

As documented in Stansbury and Summers (2018), several studies find that the growth rate of wage and the growth rate of labor productivity and/or marginal productivity of labor has been increasingly disconnected. The literature proposes a technology-driven explanation and an explanation related to the worker power such as monopsony power of employers and decreased bargaining power of workers.

Our finding is consistent with the technology-driven explanation and does not sup-

Table 2: Time Trend of Labor Share

	(1) LS	(2) LS	(3) LS	(4) LogLS
Year	-0.365×10^{-3} ***	-4.36×10^{-5}	0.243×10^{-3} ***	0.007***
	(0.0854×10^{-3})	(0.0872)	(0.0730×10^{-3})	(0.001)
A dummy variable for new technology		-0.0291***	-0.0278***	-0.390***
		(0.00243)	(0.00302)	(0.0352)
Plant fixed effects	No	No	Yes	Yes
Constant	Yes	Yes	Yes	Yes
N	1513	1513	1513	1513

Standard errors in parentheses
Other Controls includes a constant term.
* $(p < 0.10)$, ** $(p < 0.05)$, *** $(p < 0.01)$

Table 3: Production Function Estimates via Olley and Pakes (1996)

	Pooling	Separately	
	Both Tech	Old Tech	New Tech
α (K)	0.438	0.278	0.639
	(0.034)	(0.075)	(0.034)
β (L)	0.162	0.200	0.140
	(0.019)	(0.056)	(0.021)
N	1,408	229	1,179

port monopsony power or decreased bargaining power. Figure 7 provides direct evidence to support our claim. The left Panel of Figure 7 plots plant-level labor productivity (defined as total output value divided by the total wage payment) on vertical axis and plant-level labor share on 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 7 plots plant-level average wage on vertical axis and plant-level labor share on horizontal axis together with a nonparametric fitted line. In contrast to the left Panel, the fitted line are mostly flat and there are no clear relationship between these two variables. The "no-relationship" indicates that the low-labor share plants do not suppress wages of their employees. Both Panels together, the data do not support the view that the decline of labor share is caused by suppressed wage due to monopsony power or decreased bargaining power of employees.

Figure 7: Plant-level Labor Share, Labor Productivity and Average Wage

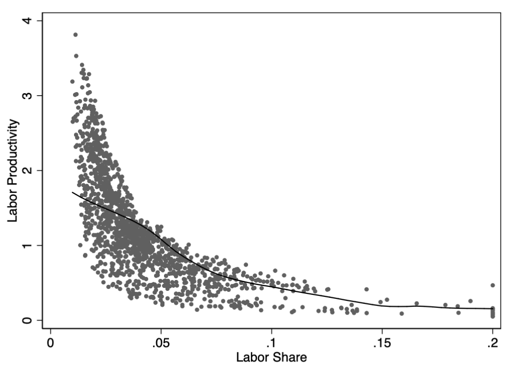

Labor Productivity

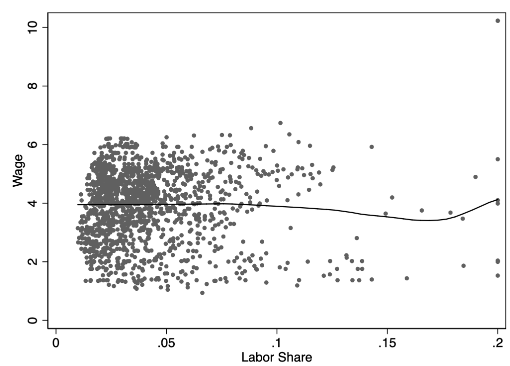

Wage

The analysis above is largely based on the simple measure of labor productivity. Theoretically, wage equals to marginal productivity of labor(MPL) in a competitive environment. To further assess whether wage and MPL has become increasingly disconnected, we estimate the production function and quantify the evolution of MPL over time. Formally, we consider the following production function;

$$Y_{it} = A_{it} K_{it}^{\alpha} L_{it}^{\beta}, \tag{1}$$

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 α and β are the parameters to be estimated. The profit maximizing plant solve 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

$$W_t = \beta \frac{P_t Y_{it}}{L_{it}}.$$

When we estimate Equation (1), we use the control function approach as in Olley and Pakes (1996).

Figure 8 plots the average real wage and MPL. On 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 average wage and MPL becomes apart during the period when new technology diffuses in the industry. In a typical dataset where we do not observe the exact technology, we would reach to the same observation as in the literature and find the dispersion between wage and MPL.

On the other hand, Panel (b) and Panel (c) 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 aligns 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 the seemingly disconnected relationship. In contrast, in Panel (b) and Panel (c), there are still some dispersion between the two variables, but these two variables grow together with similar rate overall. These results highlights the importance of controlling for the technology to draw implication from data and the usefulness of our complementary approach.

Figure 8: Growth of Real Wage and MPL

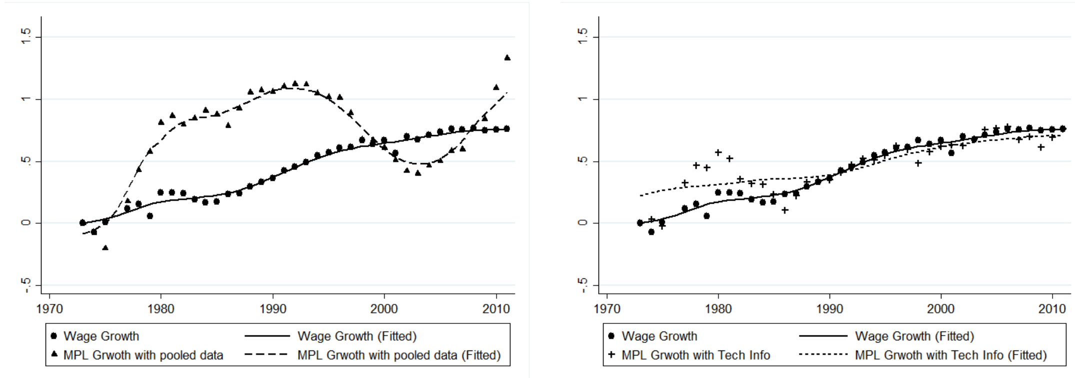

(a) with Pooled Data (b) with Technology Controlled

The increase in Mark-up

There is growing interest in how concentration affects the macroeconomic conditions. In our context, there are a number of studies that document the increase in mark-up is paired with the decline of labor share. We follow the method proposed De Loecker and Warzynski (2012) and De Loecker et al. (2020). Regarding the methodology, a few studies (e.g., Raval, 2020; Doraszelski and Jaumandreu, 2019) question whether the mark-up implied from cost minimization well captures the actual product-level markups. Independent of these studies, we find another potential factor that may bias the estimated markup; when production function is heterogeneous within the industry and when production is concentrated at plants with more capital intensive production function, industry-level mark-up would be overestimated.

Given this potential concern, we examine how the estimated mark-up change over time with and without controlling for the plant-level technology. Here, again, we assume a Cobb-Douglas production function as

$$Y_{it} = A_{it} K_{it}^{\alpha_t} L_{it}^{\beta_t}, \tag{2}$$

where we allow the shape of production function to change over time as in De Loecker et al. (2020). The corresponding cost minimization problem is

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

and the implied mark-up is

$$\text{Mark-up}_{it} = \beta_t \frac{P_t Y_{it}}{W_t L_{it}}.$$

Figure 9 plots the industry-level mark-up with and without controlling for the plant-level technology. When we do not control for the technology, the mark-up seemingly increases during the period when the new technology diffuses and productions shift to plants with new technology. In contrast, the estimated mark-up after controlling for the plant-level technology does not vary as much over time. These contrasting plots, again, highlight that availability of information on technology changes the result and its implication qualitatively.

Figure 9: Markups and Labor Share

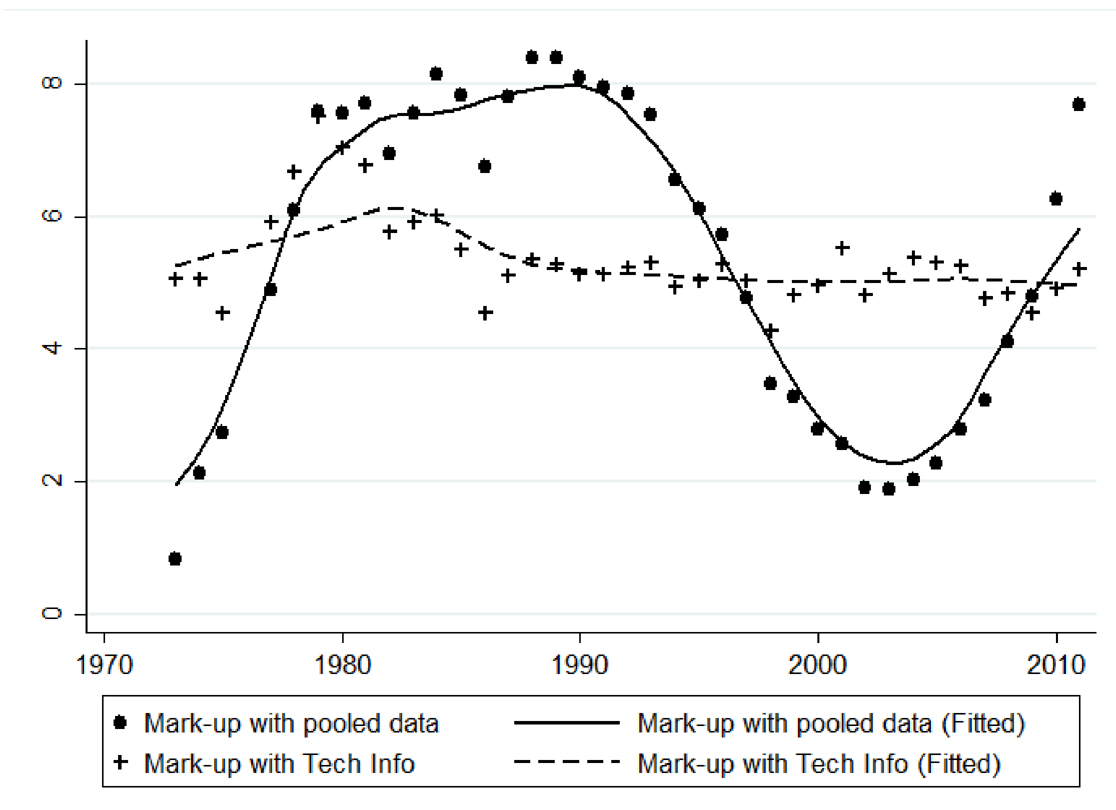

4 Event Study Design

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

$$y_{jt} = \sum_{\tau=\tau_{min}}^{\tau_{\max}} \mathbf{1}[t = t_j^* + \tau]\beta_\tau + \xi_j + \xi_t + \epsilon_{jt}, \tag{3}$$

where j is an index for plant, t is an index for year, t_j^* is the year plant j adopt 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 a event study design is often called as Two-Way Fixed Effect (TWFE) estimator. For the estimator to have meaningful interpretation, the treatment effect needs to be homogeneous across different cohort based on the treatment timing. See Goodman-Bacon (2021) for 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 pre-treatment period, we drop observation more than τ_{min} years before the treatment.

First, Figure 10 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 the employment and wages, we now look at the change in employment. To this end, Figure 11 plots the evolution of plant-

Figure 10: Evolution of Labor Share

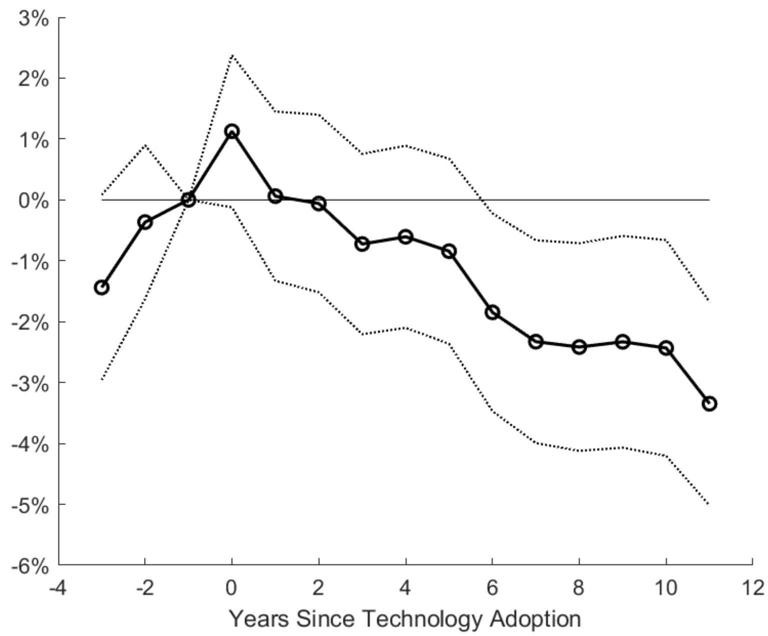

level employment relative to the timing of new technology adoption. In contrast to the labor share, the employment decreases immediately in the year of adoption, implying that the decline in the labor share is mainly driven by the change in the number of workers.

Finally, Figure 12 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 afterwards.

5 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 ex-

Figure 11: Evolution of Number of Workers

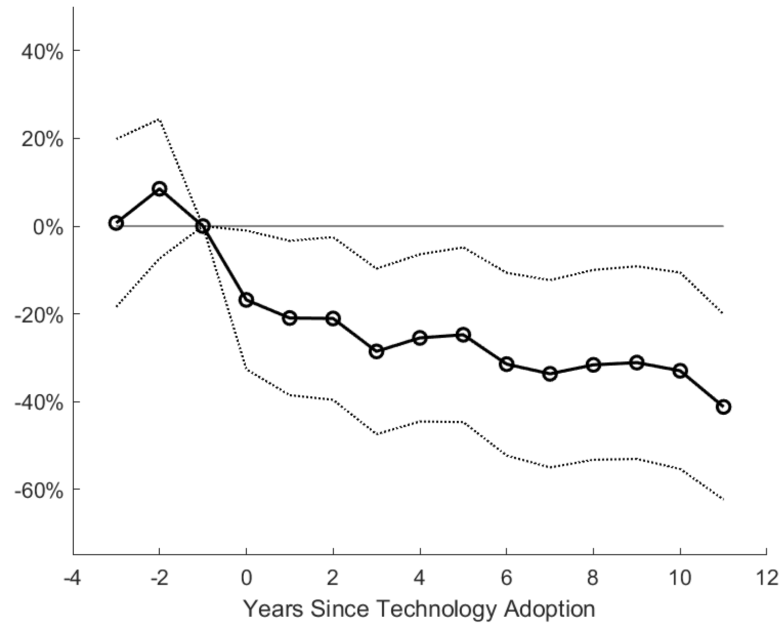

Figure 12: Evolution of Capital-Labor Ratio

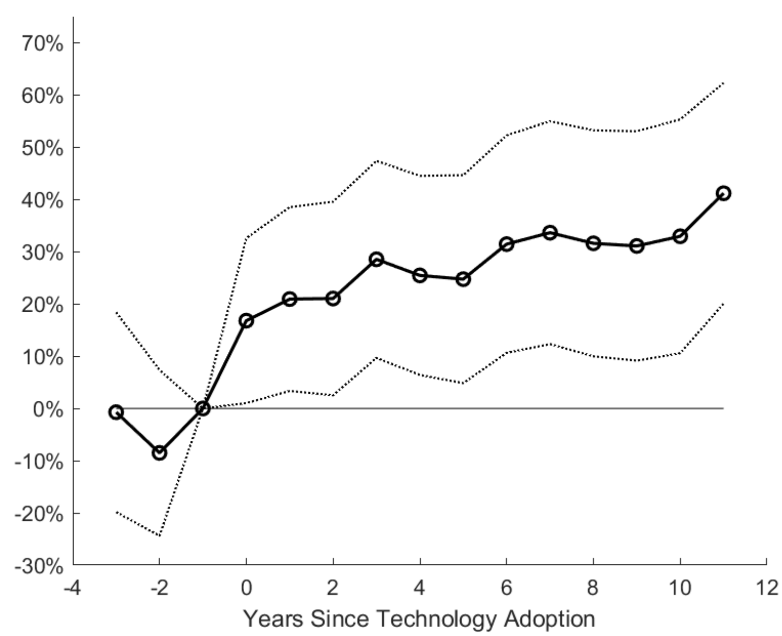

plained by the new technology diffusion: the labor share stays constant or even slightly increases overtime within the same technology plants, whereas the aggregate labor share declines because the production shifts to plants with new and more capital intensive technology. We also find that the information on the plant-level technology is a key to reject 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.

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.

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

Doraszelski, Ulrich and Jordi Jaumandreu, "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.

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.

Karabarbounis, Loukas and Brent Neiman, "The Global Decline of the Labor Share," *The Quarterly Journal of Economics*, 2014, *129* (1), 61–103.

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.

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," *Available at SSRN: https://ssrn.com/abstract=3324849*, 2020.

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 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.