DISCUSSION PAPER SERIES

DP16156

Monopolistic price setting behavior of IT firms

Coen N Teulings and Ellen Van 't Klooster

LABOUR ECONOMICS

ISSN 0265-8003

Monopolistic price setting behavior of IT firms

Coen N Teulings and Ellen Van 't Klooster

Discussion Paper DP16156
Published 14 May 2021
Submitted 07 May 2021

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

Monopolistic price setting behavior of IT firms

Abstract

De Loecker et al. (2020) have shown that markups of publicly traded firms have risen since 1980 in the US. They find that this rise cannot be attributed to a particular sector. Using the same data, this paper shows that the increase in markup is concentrated among IT firms. Firms can be classified as IT or non-IT based on industry codes, but this method ignores a number of IT firms outside specific IT industries, e.g. Amazon and Uber. We develop an alternative, firm-level classification method, by applying natural language processing (NLP) to the description of firms' activities in Compustat. After classifying firms as IT and non-IT, we show that markups in the period since 1980 fall apart in two episodes. In the first, from 1980 until 1996, non-IT firms recovered from the fall of markups in the seventies. In the second episode, since 1996, markups of IT firms exploded from 46% in 1996 to 94% in 2017, while the markup of non-IT firms was largely stagnant.

JEL Classification: N/A

Keywords: Information Technology, market power, markup

Coen N Teulings - c.n.teulings@outlook.com
Utrecht University and CEPR

Ellen Van 't Klooster - e.l.vantklooster@uu.nl
Utrecht University

MONOPOLISTIC PRICE-SETTING BEHAVIOR OF IT FIRMS

Comparing Information Technology (IT) producing firms to non-IT firms in the US

Ellen van 't Klooster and Coen Teulings[1]

May, 2021

ABSTRACT

De Loecker et al. (2020) have shown that markups of publicly traded firms have risen since 1980 in the US. They find that this rise cannot be attributed to a particular sector. Using the same data, this paper shows that the increase in markup is concentrated among IT firms. Firms can be classified as IT or non-IT based on industry codes, but this method ignores a number of IT firms outside specific IT industries. We develop an alternative, firm-level classification method, by applying natural language processing (NLP) to the description of firms' activities in Compustat. After classifying firms as IT and non-IT, we show that markups in the period since 1980 fall apart in two episodes. In the first, from 1980 until 1996, non-IT firms recovered from the fall of markups in the seventies. In the second episode, since 1996, markups of IT firms exploded from 46% in 1996 to 94% in 2017, while the markup of non-IT firms was largely stagnant.
JEL Codes: D2, D4, E2, L1.

[1]Ellen van 't Klooster: Utrecht University, e-mail: `e.l.vantklooster@uu.nl`; Coen Teulings: Utrecht University.

We thank Christiaan Behrens and SEO Amsterdam Economics for the research guidance.

1 Introduction

Today, the five firms with the highest market capitalization in the world are Microsoft, Apple, Amazon, Alphabet and Facebook: all IT firms. In 2019, the total profits of these firms were \$217 billion[2]. The price to earnings ratios of these firms tend to be around thirty, suggesting that the capital market expects these profits to grow further. Their current and expected profits suggest that these firms have substantial market power. There is a vast body of theoretical research on digital firms and common characteristics of the digital economy, such as network externalities, switching costs and two-sided markets (Rochet and Tirole, 2003; Katz and Shapiro, 1986; Laffont et al., 1998a,b). Varian (2001) identifies competitive strategies that are of particular importance to digital firms, including personalization of products and prices, versioning, bundling, switching costs, lock-in, economies of scale and network effects. These strategies are likely to lead to a concentrated market structure. However, it remains difficult to *measure* the effects of IT (Bessen and Righi, 2019). This paper researches whether IT firms have higher markups than non-IT firms in the US and whether these mark-ups have risen since the rise of internet.

De Loecker et al. (2020) have shown that markups of publicly traded firms in the US have risen between 1980 and 2016. Because there is no strong compositional pattern across industries and the increase occurs mainly within industries, the authors claim that the rise cannot be attributed to one particular industry. According to the authors, the rise in markups might be driven by information technology, but then it must have affected markups across all industries, both traditional and more modern IT-driven industries.

This paper shows that distinguishing IT from non-IT firms leads to significant differences in markups. The era of rising markups between 1980 and 2017 can be decomposed into two episodes. In the first episode, 1980 until 1996, the increase occurs mainly in traditional firms. These firms recover during that period from the decline in markups in the preceding decade. For the total group of firms, the average markup above marginal cost increased from 14% in 1980 to 39% in 1996, then remained more or less constant until 2014 and finally rises to 58% in 2017. However, in the second episode, from 1996 until 2017, the increase in markups is heavily concentrated among IT firms. IT markups were around 45% between 1980 and 1996 and then exploded to 94% in 2017. Similar to De Loecker et al. (2020), we find that for both episodes and for both IT and non-IT firms, the median markup does not change. The rise in

[2]Profits here are measured as EBITDA, or earnings before interest, taxes, depreciation, and amortization. Calculated from the Compustat North America database.

markups is driven by the increase in markups of firms in the top of the distribution. This is consistent with Autor et al. (2020) who document the rise of "superstar firms". Industries are increasingly dominated by the most productive firms.

The difference in conclusion compared to De Loecker et al. (2020) is due to the classification of IT firms that we apply. De Loecker et al. (2020) only distinguish industries at the 2-digit North American Industry Classification System (NAICS) code level. This paper distinguishes IT from non-IT firms more precisely.

Other research compares digital-intensive to non-digital intensive (Calligaris et al., 2018) or high tech from non high tech (Decker et al., 2020) at the 4-digit NAICS code level. Bessen (2017) identifies nine 4-digit NAICS industries involved in creating information technology products, such as 'Software publishers' and 'Computer and peripheral equipment manufacturing'. However, some IT companies are not included in these industries. Amazon and Netflix, for example, belong to the industries 'Electronic Shopping and Mail-Order Houses' and 'Video Tape and Disc Rental' respectively. Calligaris et al. (2018) show that markups are higher in digital intensive industries. However, they argue that their industry level taxonomy fails to capture within-industry heterogeneity in digital adoption. These examples illustrate that a sample based on industries is incomplete and imprecise, because IT firms are active in a broader variety of industries.

Brynjolfsson and Hitt (1995) look at firm-level investments in IT and find that these are associated with increases in firm-level productivity. However, their data contain at most 500 firms from exclusively the manufacturing and services sectors. Other studies that look at IT investments at the firm-level also use a limited number of firms and typically focus on one industry (Doms et al., 2004; Cline and Guynes, 2001).

Our firm-level classification method of IT provides a solution to these limitations and is able to capture within-industry heterogeneity of firms. To the best of our knowledge, such a method has not been developed before. We implement an algorithm called BERT. It is developed and published by Google AI Research and obtains state-of-the-art results on a variety of natural language processing tasks, including text classification (Devlin et al., 2018; Sun et al., 2019). First, we create a group of all firms in IT-producing industries and a group of all firms in non-IT industries. Subsequently, the algorithm compares business descriptions of firms in the IT sectors to those of firms in the non-IT sectors. This comparison enables the algorithm to identify whether business descriptions resemble those of firms in the IT or in the non-IT sectors. With this knowledge, the algorithm classifies all firms in the data as IT or non-IT using their business description.

A feature of this approach is that a firm-level *definition* of IT is not required. The algorithm simply selects firms that most resemble those in the IT producing industries. Such a classification is less disputable than a firm-level definition would be. Our firm-level classification method can also be useful for other empirical research into IT. It can be applied to any text that describes a firm, such as the firm's website or the description on Wikipedia.

The paper proceeds as follows. Section 2 discusses the firm-level classification algorithm. In Section 3, the empirical framework for measuring the markups is presented. Section 4 introduces the data used and presents estimates of the markups of all firms in the data from 1980 until 2019 and compares IT to non-IT firms. Section 5 concludes.

2 IT classification

To compare IT to non-IT firms, the firms in the database need to be classified as IT or non-IT. Because IT firms are active in both IT and different non-IT sectors, an industry-level classification is imprecise. A schematic illustration of the classification problem is depicted in figure 1. This section proposes a firm-level classification method, using the business descriptions of the firms. These business descriptions are one-sentence descriptions of the firm. The information is compiled by Standard Poor's (S&P) analysts and is included in the Compustat database.

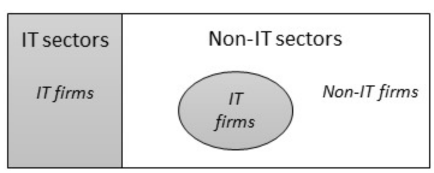

Figure 1: All firms in IT sectors are classified as IT firms. Additionally, IT firms are active in non-IT sectors.

First, we create a group of IT firms, based on GICS and NAICS codes for IT sectors. The GICS is an industry taxonomy developed by Morgan Stanley Capital International (MSCI) and S&P and is widely used by the global financial community. It contains the sector "Information Technology". NAICS codes are the standard sector codes used by Federal statistical agencies in the US for collecting, analyzing and publishing statistical data Bessen (2017). identifies nine NAICS sectors as information technology producing sectors[3]. The

Compustat database provides NAICS and GICS codes for each firm. We group all firms in either one of these nine NAICS sectors or in the GICS IT sector. We have thus obtained a group of firms in IT sectors and a group of firms in non-IT sectors.

An algorithm classifies firms at the firm-level as IT or non-IT firms. To do so, the algorithm first compares business descriptions of firms in the IT sectors to those of firms in the non-IT sectors. This comparison enables the algorithm to identify whether business descriptions resemble those of firms in the IT or in the non-IT sectors. We implement three algorithms, based on popular text mining and NLP techniques. The first is called tf-idf (term frequency-inverse document frequency), the second word2vec and the third BERT (Bidirectional Encoder Representations from Transformers). In this paper we show the results for BERT, because it performs our classification task best.

BERT uses two steps: pre-training and fine-tuning. The English BERT encoder block is pretrained by Google on BookCorpus and English Wikipedia.[4] During this pre-training it learns meaning of words and context of words. Once the pre-training is complete, the same model can be fine-tuned for a variety of downstream tasks (Devlin et al., 2018). We call this training. We train the model for classification of IT and non-IT firms.

The model is trained and its performance is tested using a training and test set, which both consist of 50% IT and 50% non-IT firms. The group of IT firms is obtained by taking all firms in the GICS and NAICS IT sectors, consisting of 5,512 firms (16% of all 33,751 firms in the dataset). 80% (4,409) of IT firms is used for the training set and 20% (1,103) for the test set.[5] We then add the same number of non-IT firms to the training and test set. For the training set, we randomly select 4,409 firms that are not in IT sectors and label these as non-IT. This selection may include IT firms, since IT firms are also active in non-IT sectors. However, since the selection predominantly consists of non-IT firms, we assume the composition is of sufficient quality to train the algorithm. For the test set we randomly

[4]BookCorpus, a dataset consisting of 11,038 unpublished books from 16 different genres and 2,500 million words from text passages of English Wikipedia.

[4]These sectors are NAICS 5112, Software publishers, 5181, Internet service providers and web search portals, 5182, Data Processing, Hosting, and Related Services, 5191 Other information services, 5415 Computer Systems Design and Related Services, 3341 Computer and peripheral equipment manufacturing, 3342 Communications Equipment Manufacturing, 3344 Semiconductor and Other Electronic Component Manufacturing, and 3345 Navigational, measuring, electromedical, and control instruments manufacturing.

[5]This ratio is often used in machine learning, see for instance the Machine Learning Crash Course published by Google (Google).

select around 1,600 firms in non-IT sectors. We check by hand whether these firms are really non-IT, by reading the corresponding firm descriptions on the Bloomberg website. IT firms are removed and from the remaining non-IT firms, 1,103 are randomly selected for the test set.

With the labelled training set, we train BERT to classify IT and non-IT firms. Subsequently, we evaluate the performance of the algorithm with the test set. We let the trained algorithm classify all firms in the test set as IT and non-IT and compare this classification to our labels for IT and non-IT. We calculate two commonly used measures of the effectiveness of a classification algorithm: recall and precision (Sokolova and Lapalme, 2009). Precision is the number of true positives over the sum of the total number of firms that are classified as IT. Recall is the number of true positives over the total number of IT firms in the test set. Figure 2 graphically illustrates precision and recall. Recall is a measure of type I errors, precision of type II errors. The test set contains 1,103 IT and 1,103 non-IT firms. BERT classifies 1,114 as IT and 1,092 as non-IT. 1,012 of the firms labelled as IT by BERT are true positives, 102 are false positives, 1,001 are true negatives and 91 are false negatives. Precision is 91% and recall is 98%. For comparison, we also look at the performance of the other two classification algorithms on the same test set. The results are shown in table 1. BERT outperforms the other classification algorithms.

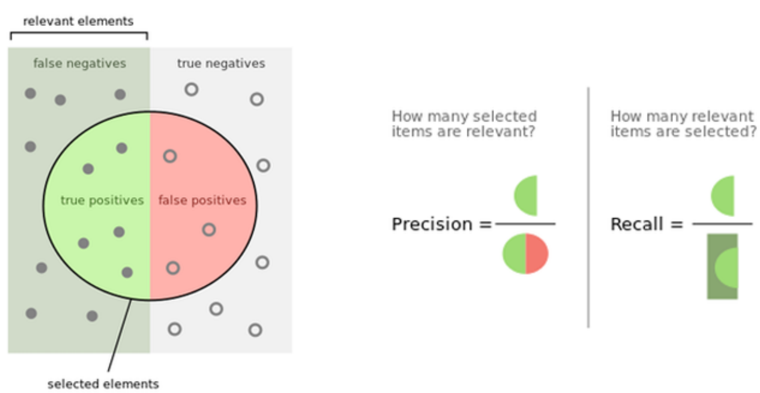

Figure 2: An illustration of recall and precision (Riggio, 2019).

A vulnerability of the above classification is that it is based purely on the business description of the firm. If the business description does not reflect how IT-savvy a firm is, the algorithm cannot properly classify the firm as IT or non-IT. To illustrate, Amazon is not classified as an IT-firm, which is consistent with its business description: "*Amazon.com, Inc. engages in the retail sale of consumer products and subscriptions in North America and internationally. The company operates through three segments: North America, Interna-*

	True positives	False positives	True negatives	False negatives	**Precision**	**Recall**
Test set	1103	–	1103	–	–	–
Tf-idf	587	128	975	516	**81%**	**57%**
Word2vec	557	138	965	546	**80%**	**54%**
BERT	1012	103	1000	91	**91%**	**98%**

Table 1: BERT's precision and recall are substantially higher than precision and recall of the other classification algorithms.

tional, and Amazon Web Services (AWS) segments." Uber, on the other hand, is classified as IT. The business description of Uber is: "*Uber Technologies, Inc. develops and supports proprietary technology applications that enable independent providers of ridesharing, and meal preparation and delivery services to transact with end-users worldwide. The company operates in two segments, Core Platform and Other Bets.*"

3 The empirical framework for measuring markups

Following De Loecker et al. (2020), we define the markup as price over marginal cost.[6] Due to lack of data, marginal cost is not directly observed. Hall (1988) proposed a method to indirectly measure the ratio price over marginal cost. We shall apply this method. Hall considers a variable factor used in the production process, for which we assume the adjustment costs are zero. Let $C_{it}(P_{vt}, Q_{it})$ be the cost of firm i at time t as a function of the price P_{vt} of its variable input and its output Q_{it}. The markup μ_{it} of the firm, defined as its price over its marginal cost, satisfies

$$\begin{aligned} \mu_{it} &\equiv \frac{P_{it}}{\partial C_{it}/\partial Q_{it}} = \theta_{it} \cdot \frac{P_{it}Q_{it}}{P_{vt}\partial C_{it}/\partial P_{vt}} = \theta_{it} \cdot \frac{P_{it}Q_{it}}{P_{vt}V_{it}}, \\ \theta_{it} &\equiv \frac{\partial C_{it}/\partial P_{vt}}{\partial C_{it}/\partial Q_{it}} \frac{P_{vt}}{Q_{it}}, \end{aligned} \tag{1}$$

where P_{it} is the output price of the firm and where V_{it} is the quantity of the variable input. In the third step, we use Roy's identity (cost minimization implies the quantity of an input to be equal to the derivative of the cost function with respect to the input price). θ_{it} is the output elasticity of the variable input of production; $P_{vt}V_{it}/P_{it}Q_{it}$ is the share of the variable input in

[6]In the literature, several versions of the markup are used. The most common is the Lerner index, which is equal to (price minus marginal cost) over price.

total revenue. Hence, the markup is equal to the output elasticity of the variable input divided by its cost as a share of total revenue. This cost share can be derived straightforwardly from the data, so that only the output elasticity remains to be estimated. The advantage of Hall's expression is that it relies only on cost minimization and hence does not require information on the demand for the firms product, as would be the case when the mark up were derived as the solution to a profit maximization problem.

Estimation of the production function

Estimating the production function requires non-trivial choices, such as its functional form, the instruments, and the estimation procedure. Here, we follow the choices made by De Loecker et al. (2020). We reiterate their argument here for the sake of transparency. We use a translog production function

$$q_{it} = \beta_t^v v_{it} + \beta_t^k k_{it} + \beta_t^{vv} v_{it}^2 + \beta_t^{kk} k_{it}^2 + \omega_{it} + \varepsilon_{it}. \tag{2}$$

where the β's are parameters and lower cases are the logs of the corresponding upper cases; k_{it} denotes the log capital stock.[7] The Cobb-Douglas production function is embedded in this specification as a special case for which $\beta_t^{vv} = \beta_t^{kk} = 0$. ω_{it} is an unobserved Hicks-neutral firm specific log productivity term, while ε_{it} captures random disturbances like measurement error in sales. The elasticity θ_{it} that we are interested in is the derivative of the log production function with respect to v_{it}. It satisfies $\theta_{it} = \beta_t^v + 2 \cdot \beta_t^{vv} v_{it}$.

The problem in estimating this equation is that ω_{it} is related to the firm's inputs since it responds to a shock in ω_{it} by re-optimizing its inputs, either downward (since the higher productivity allows the firm to produce the same output with lower inputs) or upward (since the higher productivity allows the firm to reduce its output price, thereby gaining market share, requiring more inputs); which of these two scenario's prevails depends the elasticity of the firm's demand curve. This problem is resolved by applying a two-stage estimation procedure, building on Olley and Pakes (1996).

[7]Following De Loecker et al. (2020), we omit the interaction term $\beta_t^{vk} \cdot v_{it} k_{it}$ to reduce the impact of measurement error of capital on the output elasticity.

The First Stage

The first stage aims to obtain a measure of total production ϕ_t that includes unobserved productivity, but not random disturbance ε_{it}. We estimate the following equation:

$$q_{it} = \phi_t(v_{it}, k_{it}) + \varepsilon_{it}, \tag{3}$$

where

$$\phi_t = \beta_t^v v_{it} + \beta_t^k k_{it} + \beta_t^{vv} v_{it}^2 + \beta_t^{kk} k_{it}^2 + \omega_{it}. \tag{4}$$

To approximate ϕ_t, output is modelled by a third degree polynomial in variable inputs v_{it} and k_{it}; the time dependent parameters of this polynomial are estimated by OLS. Let the estimate for this third order polynomial be denoted by $\hat{\phi}_t \equiv \phi_t\,(v_{it}, k_{it})$. This procedure ensures that all output directly or indirectly related to variable inputs and capital, or any combination of the two, including unobserved productivity, is isolated from the error term ε_{it}.

The Second Stage

The second stage aims to find the optimal value for the parameters β_t^v and β_t^{vv}, by evaluating two moment conditions for different values of these parameters. The moment conditions are obtained as follows: we assume that the unobserved productivity term follows an AR(1) process; $\omega_{it} = g \cdot \omega_{it-1} + \xi_{it}$. Given (4), this AR(1) process can be written as follows:

$$\hat{\phi}_t - \beta_t^v v_{it} - \beta_t^k k_{it} - \beta_t^{vv} v_{it}^2 - \beta_t^{kk} k_{it}^2 = g\left(\hat{\phi}_{t-1} - \beta_t^v v_{it-1} - \beta_t^k k_{it-1} - \beta_t^{vv} v_{it-1}^2 - \beta_t^{kk} k_{it-1}^2\right) + \xi_{it} \tag{5}$$

However, since the input v_{it} is variable and fully flexible, it responds to current productivity shocks ξ_{it}. Therefore, β_t^v and β_t^{vv} are estimated through two moment conditions:

$$\mathbb{E}\left[\left(\xi_{it}(\beta_t^v, \beta_t^{vv}) \begin{pmatrix} v_{i,t-1} \\ v_{i,t-1}^2 \end{pmatrix}\right)\right] = 0 \tag{6}$$

These moment conditions assume that the correlation of the productivity process's error term and the lagged variable input use should be zero (see (5)). In other words; lagged variable input should not respond to current productivity shocks.

The GMM procedure evaluates the moment conditions (6) for different values of β_t^v and β_t^{vv}. As initial values for these parameters, we use the parameters from an estimate of the translog production function without the unobserved productivity term. The procedure returns the optimal values for the parameters β_t^v and β_t^{vv}, namely the values for which the moment conditions are minimized.

4 Comparing IT to non-IT firms

4.1 Data

This section applies this methodology to data for IT and non-IT firms. We use data from the Compustat Database North America. It contains financial statements of all US' publicly traded firms, with detailed breakdowns of revenues and costs. We use data between 1978 and 2019. We drop all firms with missing data on the crucial variables: sales, cost of goods sold, capital stock, SG&A (Selling, General, and Administrative expenses), and a NAICS code. Cost of goods sold represents all expenses directly allocated by the company to production, such as material, labor and overhead. All nominal values are deflated using the FRED GDP deflator[8]. In line with De Loecker et al. (2020), we eliminate firms with a ratio between cost of goods sold to sales or of SG&A to sales in the bottom or top 1% of the corresponding year. This results in a sample of 21,974 firms.

We then use the trained BERT algorithm to classify all firms as IT or non-IT. The share of IT firms in the total number of firms and sales over time is shown in figure 3.[9] Of all 21,974 firms in the data, 32% are labelled as IT, representing 16% of total sales. Of the 6,979 IT firms, 70% are active in IT industries and 30% in non-IT industries. 13% of firms in non-IT industries are labelled as IT and 88% of firms in IT industries are labelled as IT. A large portion of firms in our dataset is classified as IT. This can in part be explained by the fact that we drop mainly non-IT firms during preprocessing. During preprocessing, 44% of all observations are removed due to missing values on variables required to calculate the markup. Of the removed observations, 8% are classified as IT and 92% are classified as non-IT.

Summary statistics are provided in table 2. The mean sales, cost of goods sold, capital stock, wage bill and employment of the IT firms are lower than the means of non-IT firms. The mean wage per person is higher for IT firms at \$69,400, compared to \$66,210 for non-IT firms. These numbers may be biased to to a large number of missing values on the wage bill. The mean number of annual observations per firm is 10.3 for non-IT firms, with a median of 7.00 years. For the IT sample, the mean is 10.6, with again a median of 8.00 years.

[8] U.S. Bureau of Economic Analysis, Gross Domestic Product: Implicit Price Deflator [GDPDEF], retrieved from FRED, Federal Reserve Bank of St. Louis; https://fred.stlouisfed.org/series/GDPDEF, June 11, 2019.

[9] This dataset only includes the observations with complete data on the crucial variables.

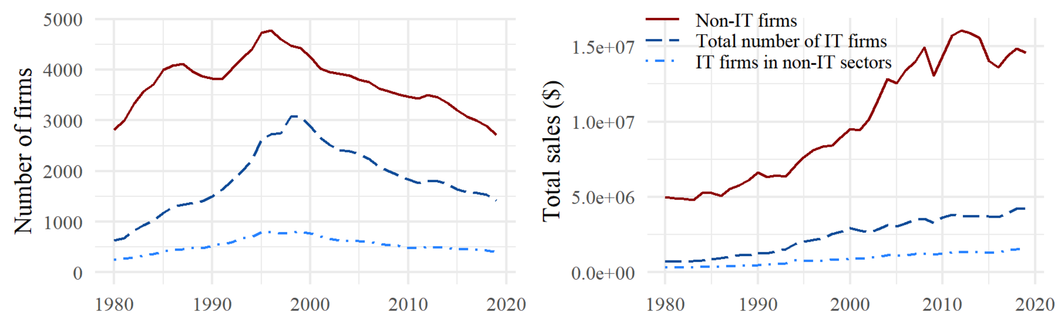

(a) Number of firms over time *(b) Total sales over time*

Figure 3: 32% of all publicly traded firms in the US are labelled as IT by the selection algorithm, representing 16% of total sales. 30% of all IT firms are active in non-IT sectors.

	Mean	Std. Dev.	25th Pct.	50th Pct.	75th Pct.	90th Pct.	Nr. Obs.
Non-IT firms							
Sales (mln $)	2,669	13,858	33	204	1,046	4,270	150,089
Cost of goods sold (mln $)	1,824	10,621	19	126	676	2,807	150,089
Capital stock (mln $)	2,166	12,995	15	104	630	3,048	150,089
SG&A (mln $)	445	2,099	7	33	166	709	150,089
Wage bill (mln $)	1,094	3,323	7	89	678	2,893	19,675
Employment (thousands)	9.6	41.3	0.2	1.1	5.1	19.7	133,571
IT firms							
Sales (mln $)	1,317	7,007	18	73	342	1,576	74,206
Cost of goods sold (mln $)	775	4,278	8	36	177	842	74,206
Capital stock (mln $)	686	5,190	5	19	101	582	74,206
SG&A (mln $)	329	1,787	8	28	106	412	74,206
Wage bill (mln $)	930	3,511	2	22	247	2,014	4,779
Employment (thousands)	5.5	25.1	0.1	0.4	1.7	7.9	66,944

Table 2: Summary statistics for the non-IT and IT sample, 1980-2019. Nominal figures are in millions of USD, deflated using the GDP Deflator with base year 2010. Employment is in thousands.

4.2 Markups

We use the methodology explained in Section 3 to estimate the markups. First, we measure the markups for the full sample from 1970 onwards. We estimate output elasticities for each industry separately, in line with De Loecker et al. (2020). Figure 4a shows the estimates of the markups obtained by De Loecker et al. (2020), figure 4b the results in this paper, both using a Cobb-Douglas (CD) specification. All markups shown in this paper are sales weighted averages. The estimates are very similar. The difference is likely to be due to the fact that De Loecker et al. (2020) use a different deflator for the cost of goods sold.

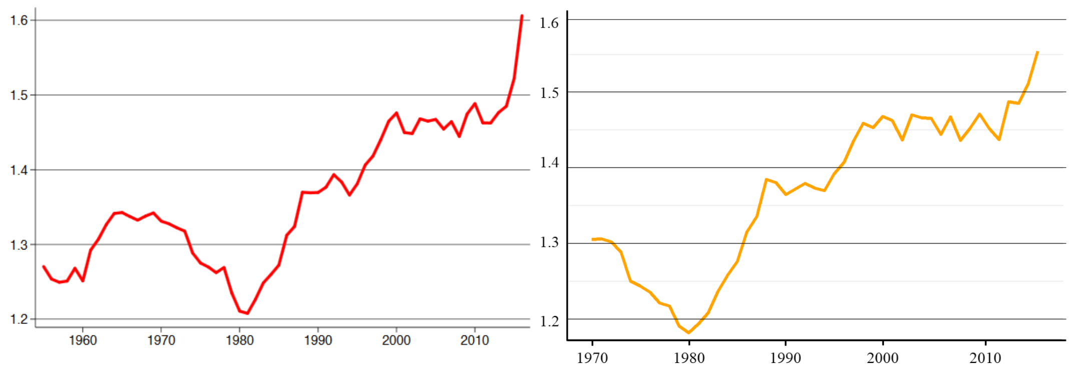

(a) Markups from De Loecker et al. (2020) *(b) Markup estimates of this paper*

Figure 4: The markup estimates in this paper are similar to the markups estimated by De Loecker et al. (2020).

The markups can be estimated using the CD or the translog production function. Figure 5a depicts the outcome for the non-IT sample and figure 5b for the IT sample. The parameters are estimated for each year separately but not for each industry and from 1980 onward, due to limited observations of IT firms in some industries. For the non-IT sample, the markups derived from the translog function are lower after 1985. Calligaris et al. (2018) also find lower markups for the translog production function. For the IT sample, the markups derived from the translog function are higher after 1993.

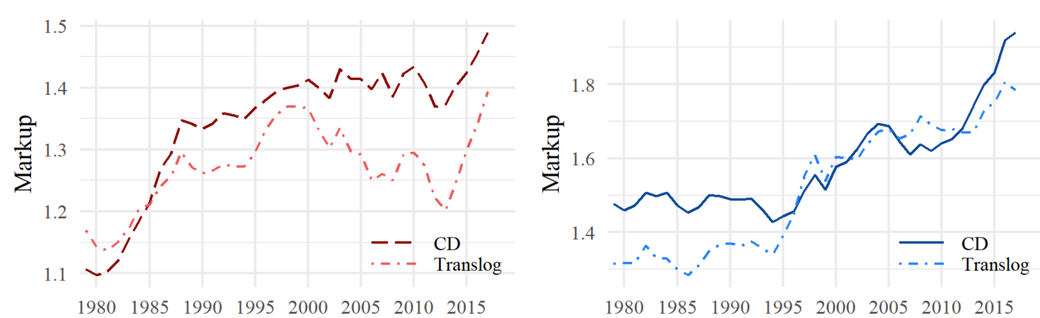

(a) Markups of the non-IT sample *(b) Markups of the IT sample*

Figure 5: A comparison between CD and translog markups over 1980-2017. For the non-IT sample, the translog markups are lower after 1985. In the IT sample, the translog markups are lower until 1996.

Figure 6 presents the main results of this paper. It compares the markups of the non-IT and the IT sample, using the CD production function specification. The evolution of the markups differs between the two samples. IT markups are significantly higher. They were relatively stable between 1980 and 1996 at about 1.45. After 1996, the markups increase to 1.94 in 2017. In contrast to the IT markups, markups of non-IT firms increased from 1.10 in 1980 to 1.38 in 1996 and remained relatively stable afterwards. The rise in markups of the full sample since 1996 can in part be explained by the rise in markups of IT firms. Figure 6 shows the sales-weighted average, in which the non-IT firms weigh more heavily, because they represent a larger part of total sales. Therefore, the difference in markups of the full sample compared to non-IT firms is limited.

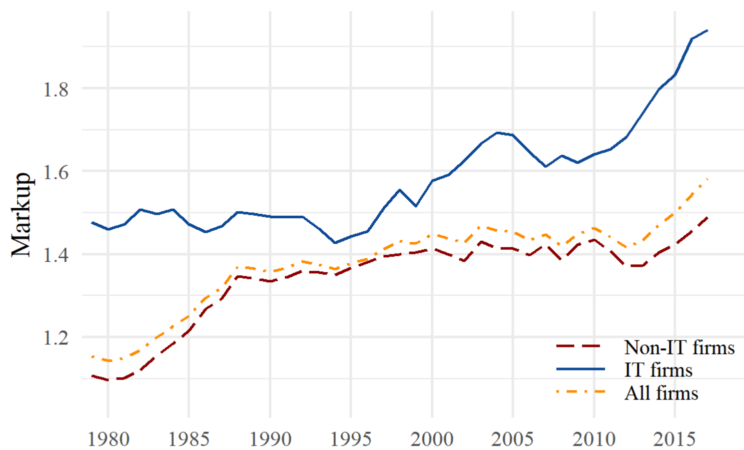

Figure 6: Average (sales weighted) markups between 1980 and 2017. The average IT markup is systematically higher than the average markup of the full sample.

The yellow line for the full sample is slightly different from the yellow line shown in figure 4b since here the β-parameters are not separately estimated for each industry for the sake of comparison with the IT and non-IT group, where we lack the data to estimate the β-parameters separately for each industry. The small difference between both lines suggests that estimating the β-parameters separately for each industry does not matter much.

Figure 7 shows the markups derived from the translog production function. The differences between the markups of the IT, the non-IT, and the full sample are slightly larger after 1996 than for the CD specification.

Figure 8 splits IT firms in the subset of IT firms in the IT sectors and IT firms in non-IT sectors. It shows that the average markup of just the firms in the IT sectors is higher. The lower average markup of IT firms in non-IT sectors can be explained by the fact that almost

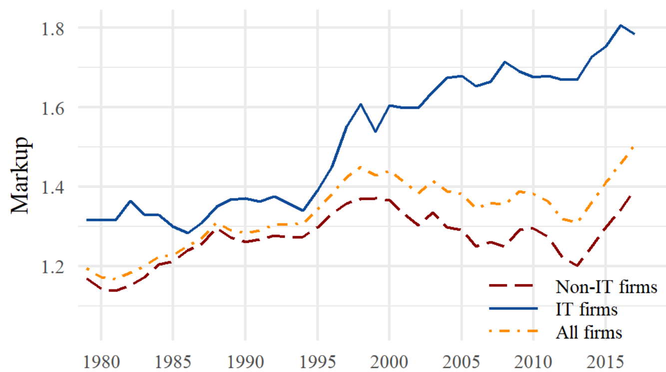

Figure 7: Markups according to the translog specification. Also for the translog specification, the average IT markup is systematically higher than the average non-IT markup.

half of these are active in manufacturing industries, of which many in the semiconductor industry. Markups in manufacturing are generally lower. Markups of IT firms in non-IT sectors are still higher than markups of non-IT firms.

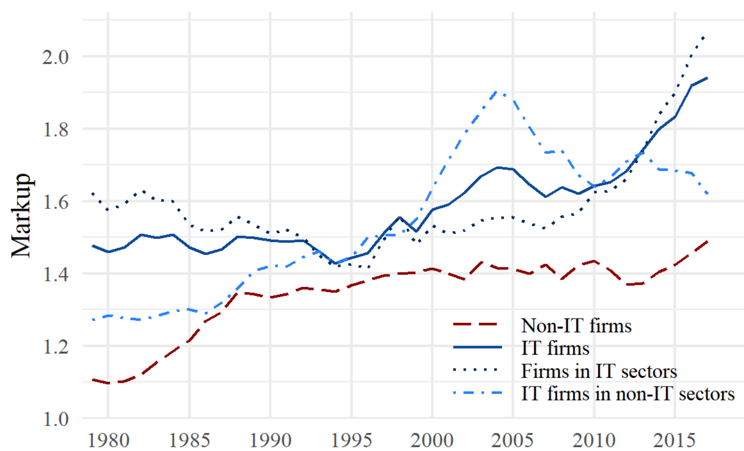

Figure 8: The average markup of the subset of firms in the IT sectors lies above the markup of IT firms in non-IT sectors.

A rise in markup could be due to either rising output elasticity or falling expenditure shares of variable inputs, see equation (1). Figure 9a shows the Cobb-Douglas output elasticities of the non-IT and the IT sample, while figure 9b shows the expenditure shares. As expected, both the output elasticity and the expenditure share of variable inputs are lower for IT than for non-IT firms: IT firms' inputs are less sensitive to their sales and the share of variable inputs in total sales is lower. Clearly, the effect of the share of variable input in total sales dominates, leading to a higher markup for IT firms.

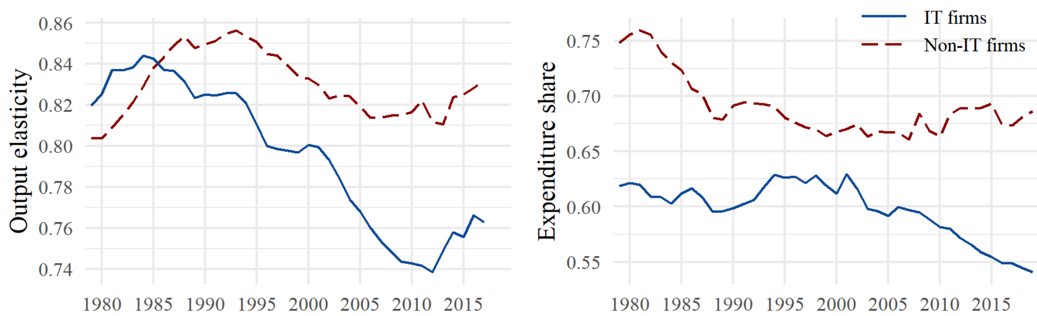

(a) Cobb-Douglas output elasticity *(b) Expenditure shares*

Figure 9: The sales weighted average output elasticity and expenditure share are systematically higher in the full sample.

4.3 The distribution of the markups

De Loecker et al. (2020) have shown that the rise of markups in the US is mainly driven by the upper percentiles of the distribution of the markup. They claim that this result holds essentially within each industry on its own, suggesting that the rise in the average markup is due to intra-industry and not inter-industry heterogeneity. Figure 10 shows that the same applies for the distributions of the markup for IT and for non-IT firms: the increase in markup is concentrated in the upper end of the distribution, though the increase in the upper tail of the distribution is much more pronounced in for IT than for non-IT firms. Similarly to De Loecker et al. (2020), the lower percentiles do not explain the increase in markups. The difference between the 90$^{\text{th}}$ and the 75$^{\text{th}}$ percentile is greater for non-IT than for IT firms. For non-IT firms, the 75$^{\text{th}}$ percentile remains relatively stable between 1990 and 2010, and decreases slightly after 2010. For IT firms, the 75$^{\text{th}}$ percentile shows a steady increase between 1996 and 2016, especially between 2006 and 2016. The rise in markups is more widespread among IT than among non-IT firms.

4.4 Fixed costs and profits

Markups do not take fixed costs into account. Therefore, a markup does not necessarily imply market power. Figure 11 shows the fixed costs as a percentage of total costs. In line with De Loecker et al. (2020), this is calculated as

$$\text{fixed cost share} = \frac{\text{xsga}}{\text{variable cost} + r \cdot \text{capital} + \text{xsga}},$$

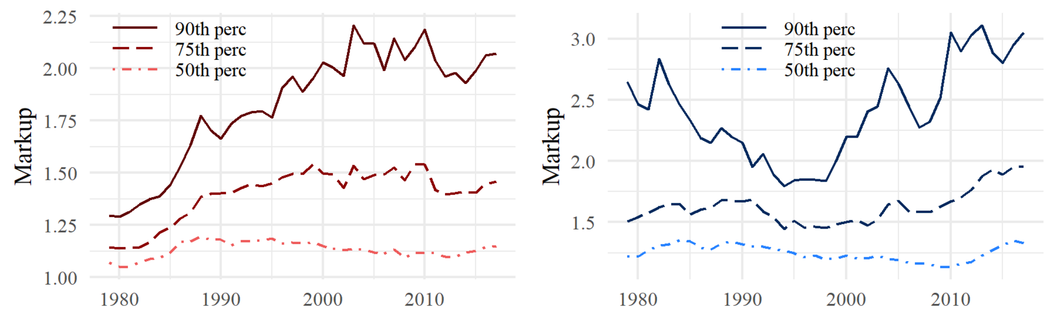

(a) Markup distribution, non-IT firms *(b) Markup distribution, IT firms*

Figure 10: The distribution of the markups (sales weighted) show that the rise in markups is mainly due to the rise of the 10% highest markups.

where xsga represents the fixed costs (sales, general cost, administration) and where r denotes the user cost of capital. Fixed costs are higher for IT than for non-IT firms.

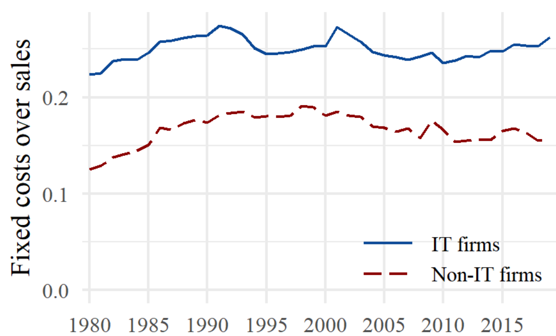

Figure 11: The average fixed costs of IT firms are high compared to non-IT firms.

Figure 12 shows the profit rate of IT and non-IT firms. In line with De Loecker et al. (2020), it is calculated as

$$\text{profit rate} = 1 - \frac{\text{output elasticity}}{\text{markup}} - \frac{r * \text{ppegt}}{\text{sales}} - \frac{\text{xsga}}{\text{sales}}, \tag{7}$$

where the output elasticity over the markup substitutes the expenditure on variable inputs as a share of sales. There is a conceptual difference between the markup and the profit rate: the latter takes the role of quasi rents as compensation for fixed cost and the cost of capital into account. Figure 12 shows the profit rates of IT and non-IT firms. The profit rate for non-IT firms looks similar to the rate shown by De Loecker et al. (2020). Contrary to De Loecker et al. (2020), it is slightly negative in 1981-1982 in our calculation. This difference is caused by the fact that markups reported in De Loecker et al. (2020) are slightly higher than markups reported here, which is due to two factors. First, De Loecker et al.

(2020) estimate the markup for each industry separately, which results in slightly higher markups. Secondly, they use a different deflator for the variable costs. The more important message from figure 12 is that, despite the higher fixed cost, the profit rate is substantially higher for IT than for non-IT firms, except for the dot com crisis in 2002, and that this difference is increasing rapidly over time. The high markup for IT firms is therefore not a quasi-rent, a compensation for prior investment and high fixed cost, but genuine rent derived from barriers to entry.

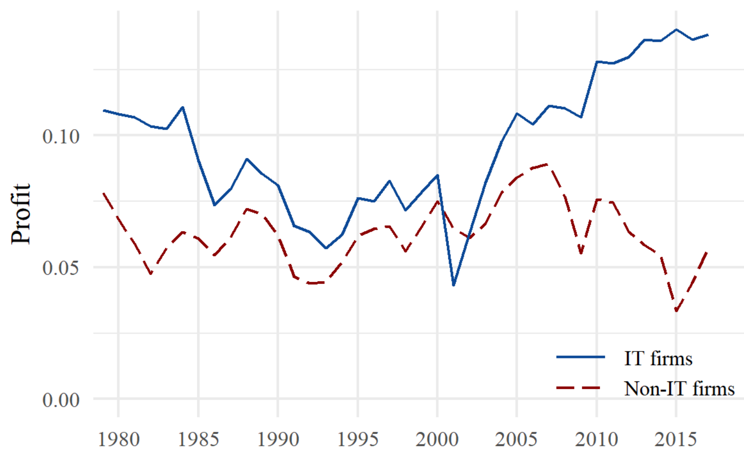

Figure 12: The average, sales weighted, profit of IT firms is high compared to non-IT firms. Since 2010, the gap has increased.

Figure 13 shows the profits of IT firms as a percentage of total profits. For each year, we summed the profits of all IT firms and divided these by the total profits of all firms. We did the same for the sales of IT firms. The figure shows that IT profits represent a large share of total profits, except during the dot com crisis. This is striking, considering the relatively small share in sales of IT firms. The figure starts in 1992. Before 1992, the estimated profits were unrealistic, with values below or very close to zero.

Figure 14 shows the EBIT as reported in the Compustat database. The sales-weighted reported EBIT of IT firms increases between 2001 and 2019. The EBIT of non-IT firms follows a more irregular path. Roughly, it remains stable between 1980 and 1999, then rises steeply until 2008 and then declines. It lies above the EBIT of IT firms until 2014.

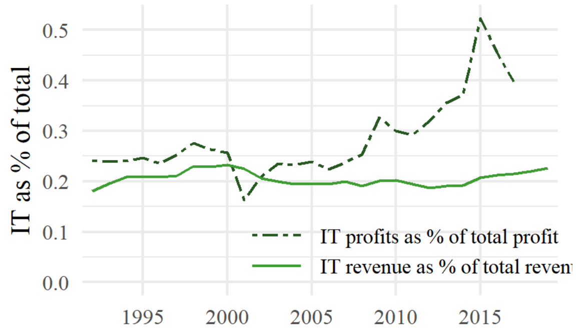

Figure 13: IT profits represent a large share of total profits, while IT sales are relatively small.

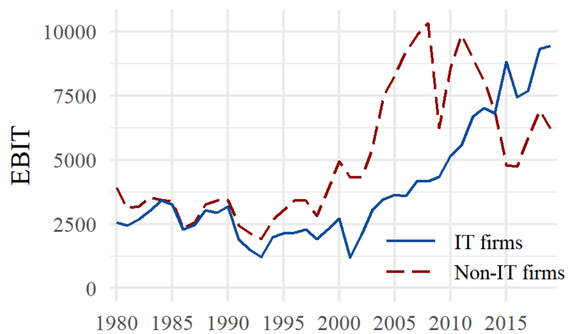

Figure 14: Reported EBIT (sales-weighted) of non-IT firms is higher until 2014, when it is overtaken by the EBIT of IT firms.

5 Conclusion

This paper shows that the average markup is significantly higher for IT than for non-IT firms in the US. The period since 1980 falls apart in two separate episodes: first, the episode from 1980 to 1996, and then, the episode from 1996 to 2017. In the first episode, there is a recovery of non-IT firms from the seventies when oil increases, union radicalism, and severe recessions eroded profitability. The first episode were the days of Ronald Reagan in the US and Margaret Thatcher in the UK. They ran a program of deregulation, the welfare state reduction, and confrontation with trade unions. DiNardo et al. (1996) claim that these institutional changes, like freezing minimum wages and the decline of union power, have contributed substantially to increasing wage inequality. Several authors have shown that the fall in minimum wages can explain a large share of the increase in wage inequality at the bottom of the labour market (Teulings, 2003; Lee, 1999). During the second episode, from

1996 onward, there is a run away of IT firms. This run away coincides with the rise of the internet.

For both IT and non-IT firms, the rise of the average markup is mainly driven by the increase in markups of a few firms with the highest markups. The median markup does not change between 1980 and 2017. For IT firms the difference between the 75$^{\text{th}}$ and the 90$^{\text{th}}$ percentile is smaller than for non-IT firms in 2016. This implies that rising markups is more widespread among IT firms than among non-IT firms.

De Loecker et al. (2020) state that rising markups mainly occurs *within* industries and that the difference in the development of markups *between* industries is limited. The rise in markups cannot therefore be attributed to a single industry, in their view. To the contrary, our research shows that the difference between IT and non-IT firms is substantial. Rising IT markups can explain part of the rising markups found by De Loecker et al. (2020).

Researchers point out vulnerabilities of the method imposed by De Loecker et al. (2020). They question whether the absolute level of the markup is correctly measured (Syverson, 2019; Edmond et al., 2018; Traina, 2018). These vulnerabilities also affect the absolute levels of the markups of both IT and non-IT firms in this paper, but we find that the *divergence* of the markups of IT and non-IT firms is robust to these vulnerabilities. For instance, Edmond et al. (2018) find that the rise in markups largely disappears for *cost-weighted* aggregates of the firm level markups, as opposed to *sales-weighted* averages used by De Loecker et al. (2020). But also for *cost-weighted* aggregates, we find a surge in the aggregate markup of tech firms during the last decade and a divergence from the non-IT markups.

Autor et al. (2020) document the fall in the labor share and an interpretation of this fall based on the simultaneous rise of “superstar firms”. Industries are increasingly dominated by the most productive firms, that have high markups and a low labor share of value-added. Micro-economic evidence is provided for six major industries in the US. The authors identify the growth of platform competition, advances in information technology and strong network effects as potential causes for the “winner take most” mechanism. It would be interesting to apply the firm-level classification of IT and non-IT to the research by Autor et al. (2020), to analyze the role of IT firms in the falling labor share and the rise of superstar firms.

References

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

Bessen, J. E. (2017). Information technology and industry concentration. *Boston University School of Law Law Economics Paper Series*.

Bessen, J. E. and C. Righi (2019). Shocking technology: What happens when firms make large it investments? *Boston Univ. School of Law, Law and Economics Research Paper* (19-6).

Brynjolfsson, E. and L. Hitt (1995). Information technology as a factor of production: The role of differences among firms. *Economics of Innovation and New technology 3*(3-4), 183–200.

Calligaris, S., C. Criscuolo, and L. Marcolin (2018). Mark-ups in the digital era. *OECD Science, Technology and Industry Working Papers 2018/10*.

Cline, M. K. and C. S. Guynes (2001). A study of the impact of information technology investment on firm performance. *Journal of Computer Information Systems 41*(3), 15–19.

De Loecker, J., J. Eeckhout, and G. Unger (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics 135*(2), 561–644.

Decker, R. A., J. Haltiwanger, R. S. Jarmin, and J. Miranda (2020). Changing business dynamism and productivity: Shocks versus responsiveness. *American Economic Review 110*(12), 3952–90.

Devlin, J., M.-W. Chang, K. Lee, and K. Toutanova (2018). Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.

DiNardo, J., N. M. Fortin, and T. Lemieux (1996). Labor market institutions and the distribution of wages, 1973-1992: A semiparametric approach. Technical report.

Doms, M. E., R. S. Jarmin, and S. D. Klimek (2004). Information technology investment and firm performance in us retail trade. *Economics of Innovation and new Technology 13*(7), 595–613.

Edmond, C., V. Midrigan, and D. Y. Xu (2018). How costly are markups? Technical report, National Bureau of Economic Research.

Google. Training and test sets: Splitting data | machine learning crash course.

Hall, R. E. (1988). The relation between price and marginal cost in us industry. *Journal of Political Economy 96*(5), 921–947.

Katz, M. L. and C. Shapiro (1986). Technology adoption in the presence of network externalities. *Journal of Political Economy 94*(4), 822–841.

Laffont, J. J., P. Rey, and J. Tirole (1998a). Network competition: I. Overview and nondiscriminatory pricing. *The RAND Journal of Economics*, 1–37.

Laffont, J. J., P. Rey, and J. Tirole (1998b). Network competition: II. Price discrimination. *The RAND Journal of Economics*, 38–56.

Lee, D. S. (1999). Wage inequality in the united states during the 1980s: Rising dispersion or falling minimum wage? *The Quarterly Journal of Economics 114*(3), 977–1023.

Olley, G. S. and A. Pakes (1996). The dynamics of productivity in the telecommunications equipment industry. *Econometrica 64*(6), 1263–1297.

Riggio, C. (2019, Nov). What's the deal with accuracy, precision, recall and F1?

Rochet, J. C. and J. Tirole (2003). Platform competition in two-sided markets. *Journal of the European Economic Association 1*(4), 990–1029.

Sokolova, M. and G. Lapalme (2009). A systematic analysis of performance measures for classification tasks. *Information Processing & Management 45*(4), 427–437.

Sun, C., X. Qiu, Y. Xu, and X. Huang (2019). How to fine-tune bert for text classification? In *China National Conference on Chinese Computational Linguistics*, pp. 194–206. Springer.

Syverson, C. (2019). Macroeconomics and market power: Context, implications, and open questions. *Journal of Economic Perspectives 33*(3), 23–43.

Teulings, C. N. (2003). The contribution of minimum wages to increasing wage inequality. *The Economic Journal 113*(490), 801–833.

Traina, J. (2018). Markup estimation using financial statements: Cost of goods sold vs. operating expense. Technical report, University of Chicago Working Paper.

Varian, H. R. (2001). Economics of information technology. *University of California, Berkeley*.