

Monopolistic price-setting behavior of information technology firms

E. L. van 't Klooster and C. N. Teulings

Utrecht School of Economics, Utrecht University, Utrecht, Netherlands

ABSTRACT

Markups, which are when prices are greater than the marginal cost of production, have increased in the USA over the last decades. Our paper explores the differences in markups between Information Technology (IT) and non-IT firms. By doing so, we contribute to the understanding of the role of the IT industry in the increase in markups. We extend to De Loecker, Eeckhout, and Unger (2020), who find that markups of publicly traded firms have risen since 1980. They argue that no industry has systematically higher markups. We develop a novel firm-level classification method using natural language processing (NLP) to distinguish IT from non-IT firms. Our approach differs from the commonly used North American Industry Classification System (NAICS). Using our classification, we find that the increase in markup in the period since 1980 occurred in two separate episodes. In the first, from 1980 until 1996, firms recovered from the fall of markups in the 1970s. In the second episode, since 1996, markups of IT firms diverge enormously. Markups of IT firms surge from 47% in 1996 to 80% in 2018, while the markup of non-IT firms remains largely unchanged.

ARTICLE HISTORY

Received 14 February 2022
Accepted 14 February 2023

KEYWORDS

Markup; IT firms; natural language processing (NLP); market power

JEL CLASSIFICATION


D2; D4; E2; L1

1. Introduction

The five firms with the highest market capitalization in the world are Microsoft, Apple, Amazon, Google (or Alphabet) and Facebook (or Meta): all are IT firms. In 2020, the total profits of these firms were \$190 billion (calculated from the Compustat North America database using earnings before interest and taxes). Their profits suggest that these firms have substantial market power. Correspondingly, a vast body of theoretical research points out characteristics of IT firms that are likely to lead to market power and a more concentrated market structure. Examples of such characteristics are network externalities, personalization of products and prices, two-sided markets, switching costs and economies of scale (Varian 2001; Rochet and Tirole 2003; Katz and Shapiro 1986; Laffont, Rey, and Tirole 1998a, 1998b; Crouzet and Eberly 2019; De Ridder et al. 2019). Concurrent with the surge of IT firms, the average markup, an important indicator of market power, of US firms has increased.

This paper answers the question: how do the markups of IT firms compare to the markups of non-IT firms in the USA? IT firms are firms that are primarily engaged in producing and developing software, data processing and services, or producing hardware and semiconductors. We extend to De

CONTACT E. L. van 't Klooster  e.l.vantklooster@uu.nl  Utrecht School of Economics, Utrecht University, Kriekenpitplein 21-22, Utrecht, Netherlands

 Supplemental data for this article can be accessed online at <https://doi.org/10.1080/10438599.2023.2194641>.

© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

Loecker, Eeckhout, and Unger (2020) (henceforth DEU), who find that markups of publicly traded firms in the USA have risen between 1980 and 2016. On the one hand, DEU find a similar rise in markups across all 2-digit NAICS sectors, arguing that it is not likely that one specific industry is driving the rise in markup. On the other hand, the theoretical literature points towards high market power of IT firms and the recent surge of IT firms suggests the IT industry plays a role in rising markups. The comparison of IT and non-IT markups in this paper provides empirical insight into the role of the IT industry.

First, this paper contributes by developing a firm-level classification method for IT and non-IT firms, enabling a more precise analysis of the role of IT firms. Thus far, the empirical literature compares IT to non-IT firms at the sector level. Bessen (2017) identifies nine 4-digit NAICS sectors involved in creating information technology products, such as *Software publishers* and *Computer and peripheral equipment manufacturing*. However, this classification is imprecise, because IT firms are active in a broad range of sectors, including non-IT sectors. Amazon and Uber, for instance, belong to the sectors *Electronic Shopping and Mail-Order Houses* and *Taxi services* respectively, while they are primarily engaged in developing software and data services. While a sector-level classification of IT excludes these firms, our firm-level classification does label Amazon and Uber as IT firms. We are thus able to capture within-sector heterogeneity of firms. To the best of our knowledge, such a method has not been developed before. This firm-level classification method can also be useful for other empirical research. It can be applied to any text that describes a firm, such as the firm's website or the description on Wikipedia.

Second, this paper contributes by estimating and comparing indicators of the market power of IT and non-IT firms. We want to know what role IT firms specifically play in the rising market power of firms in the USA, due to the vast body of *theoretical* research pointing towards high market power of IT firms and the unprecedented market capitalization of big tech. Country-level empirical evidence for this is scarce and it remains difficult to *measure* the effects of IT (Bessen and Righi 2019). It requires isolating IT from non-IT firms. Our firm-level classification algorithm enables us to do this, enabling a comparison between IT and non-IT firms on key indicators of market power, such as markups, profits, and cost structures. This analysis of IT vs. non-IT firms provides insight into the role of IT firms in the rising market power as documented by De Loecker, Eeckhout, and Unger (2020). Previous research compares IT to non-IT at the sector level (Calligaris, Criscuolo, and Marcolin 2018; Brynjolfsson and Hitt 1995). Such sector-level taxonomies do not include IT firms outside of IT sectors and thus fail to capture within-sector heterogeneity in digital adoption (Calligaris, Criscuolo, and Marcolin 2018). Previous studies that look at IT investments at the firm-level use a limited number of firms and typically focus on one sector (Doms, Jarmin, and Klimek 2004; Cline and Guynes 2001).

To distinguish IT from non-IT at the firm level, we implement an algorithm called BERT (Bidirectional Encoder Representations from Transformers). 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 sectors and a group of all firms in non-IT sectors. 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 non-IT sectors. With this knowledge, the algorithm classifies all firms in the data as IT or non-IT, based on their business description.

To compare markups, profits and costs of IT to non-IT firms, we implement a methodology that is similar to the methodology of DEU. We take into account the recent studies by Bond et al. (2021) and Edmond, Midrigan, and Xu (2018) that point out vulnerabilities of this methodology. We use all firms in the Compustat North America database between 1980 and 2020 and estimate the markups, profits and costs for the group of IT and non-IT firms.

The results show that markups of IT firms are significantly higher than markups from non-IT firms and that the difference has doubled over the last decade. The era of rising markups between 1980

and 2018 can be decomposed into two episodes. In the first episode, 1980 until 1996, the increase occurs in all 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 50% in 2020. In the second episode, from 1996 until 2018, the increase in markups is heavily concentrated among IT firms. IT markups were around 45% between 1980 and 1996 and then surge to 80% in 2018. DEU find a rise in markup within all 2-digit NAICS sectors and find no sectors with significantly higher markups. However, this paper shows that distinguishing between IT and non-IT at both the firm level and the 4-digit NAICS sector level leads to pronounced differences in markups of IT and non-IT firms and sectors. We also look at profits that take fixed costs into account. Despite IT firms having higher fixed costs, their profits are still significantly higher than the profits of non-IT firms. Our findings shed a different light on the role of the information technology industry in rising market power.

Similar to DEU, we find that for both episodes and both IT and non-IT firms, the median markup does not change. The rise in 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. Autor et al. (2020) find that sectors are increasingly dominated by *superstar firms* that have high markups and a low labor share of value-added. 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. Information technology can thus be a contributing factor to the emergence of ‘superstar’ firms. We find that, besides contributing to the emergence of ‘superstar’ firms, the role of IT is broader and IT firms lead to an increase in markups, most prominently through the low variable cost share that we find for IT firms.

This 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 discusses the results and their relation to recent literature. Section 6 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. We make use of the NAICS and GICS sector classifications that are included in our data. Both GICS and NAICS IT sectors contain only a selection of all IT firms. For instance, Amazon and Uber are not included. Because IT firms are active in both IT and non-IT sectors, the sector-level classifications are incomplete. Figure 1 presents a schematic illustration of the classification problem. This section proposes a firm-level classification method, using the business descriptions of the firms that are included in the Compustat database. These business descriptions are one-sentence descriptions of the firm, compiled by Standard & Poor’s (S&P) analysts. The algorithm simply selects firms with a business description that resembles the business descriptions of firms in the IT-producing sectors. At the end of this section, we use a test set to show the power of this approach.

First, we create a group of IT firms, based on NAICS codes for IT sectors. The Compustat database provides NAICS codes for each firm. NAICS codes are the standard sector codes used by Federal




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

statistical agencies in the US for collecting, analyzing and publishing statistical data. Bessen (2017) identifies nine NAICS sectors as information technology-producing sectors. 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*. We group all firms that are in one of these nine IT NAICS sectors and group all firms in the other, 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 non-IT sectors. We implement three algorithms, based on popular text mining and NLP techniques: tf-idf (term frequency-inverse document frequency), word2vec and BERT (Bidirectional Encoder Representations from Transformers). Because BERT outperforms the other algorithms, we use BERT for the final classification and explain it in more detail.

BERT uses two steps: pre-training and fine-tuning. The English BERT encoder block is pretrained by Google on 11,038 unpublished books from BookCorpus and 2,500 million words from text passages of English Wikipedia. During this pre-training, it learns the meaning of words and context of words. Once the pre-training is complete, the same model can be fine-tuned for a variety of tasks (Devlin et al. 2018). We call this training. We train the model for the classification of IT and non-IT firms.

We train the model and test its performance using a training and test set, which both consist of 50% IT and 50% non-IT firms. The group of IT firms for training is obtained by taking all firms in the NAICS IT sectors, consisting of 4610 firms (18% of all 25,527 firms in the dataset). We then add the same number of non-IT firms to the training set, by randomly selecting 4610 firms that are not in IT sectors and labeling these as non-IT. This selection may include IT firms, since IT firms are also active in non-IT sectors. However, the selection predominantly consists of non-IT firms. We assume the business descriptions in this sample sufficiently characterize 'non-IT' business descriptions to train the algorithm.

The objective of the algorithm is to classify IT firms that are not in IT sectors as IT. To evaluate the algorithm's performance, we use a test set with IT firms that are not in the NAICS IT sectors. We look at the GICS sectors, which are also provided in the Compustat database. The GICS is a sector 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*. We select all firms that are in the GICS Information Technology sector, but not in one of the NAICS IT producing sector and obtain a group of 815 IT firms that are not in the NAICS IT sectors. For the non-IT firms in this test set, we randomly select firms that are neither in the NAICS IT sectors, nor in the GICS Information Technology sector. We check by hand whether these firms are truly non-IT, by reading the corresponding firm descriptions on the Bloomberg website, and we remove the IT firms.

With the labeled 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 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. Recall is a measure of type I errors, precision of type II errors. The test set contains 815 IT and 815 non-IT firms. BERT classifies 714 as IT and 889 as non-IT. 683 of the firms labeled as IT by BERT are true positives, 58 are false positives, 757 are true negatives and 132 are false negatives. Precision is 96% and recall is 84%. For comparison, we also look at the performance of the other two classification algorithms on the same test set. Table 1 shows the results. BERT outperforms the other classification algorithms.

Table 1. Precision and recall measures of the different classification algorithms.

	True positives	False positives	True negatives	False negatives	Precision	Recall
Test set	815	–	815	–	–	–
Tf-idf	587	128	975	516	71%	50%
Word2vec	557	138	965	546	82%	53%
BERT	683	58	757	132	96%	84%

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. Two examples of IT firms that are not classified as IT by BERT are UpSnap and Link Motion. UpSnap is a digital advertising firm and Link Motion is a software and technology company focusing on the smart drive business. Their business descriptions are: *UpSnap, Inc. provides mobile advertising and direct mail solutions in the United States and internationally*, and *Link Motion Inc. operates as a smart car and smart ride company in the People's Republic of China and internationally*. These business descriptions are likely too short for proper classification.

Uber and Amazon are examples of IT firms that are not in IT sectors, but that BERT does classify 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*. Amazon's business description is: *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, International, and Amazon Web Services (AWS) segments*.

3. Measuring markups

Following DEU, we define the markup as price over marginal cost. (The literature uses several versions of the markup. The most common is the Lerner index, which is equal to the price minus marginal cost, over the price.) Due to a lack of data, marginal cost cannot be observed directly. Hall (1988) proposed a method to indirectly measure the ratio of price over marginal cost. We 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

$$\mu_{it} \equiv \frac{P_{it}}{\partial C_{it} / \partial Q_{it}} = \theta_{it} \cdot \frac{P_{it} Q_{it}}{P_{vt} \partial C_{it} / \partial P_{vt}} = \theta_{it} \cdot \frac{P_{it} Q_{it}}{P_{vt} V_{it}}, \quad (1)$$

$$\theta_{it} \equiv \frac{\partial C_{it} / \partial P_{vt} P_{vt}}{\partial C_{it} / \partial Q_{it} Q_{it}}, \quad (2)$$

where P_{it} is the output price of the firm and where Q_{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 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. We estimate the output elasticity using a two-stage GMM procedure, similar to DEU. Online appendix A explains the empirical framework in more detail.

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 firm's product, as would be the case when we derive the markup as the solution to a profit maximization problem. Because we do not have data on output quantity, we follow DEU and use revenue as a proxy of output.

However, Bond et al. (2021) show that, if the markup is measured with the revenue elasticity instead of the quantity elasticity, the markup is uninformative about the true markup. Firms with market power set markups as a function of the elasticity of demand, so that the *revenue* elasticity of input equals its *output* elasticity divided by the true markup. The markups estimated with the revenue elasticity equal the markup as measured with the *output* elasticity, divided by the true markup, and will therefore equal one. De Ridder et al. (2021) show that this holds on average so that the average markup of the ratio estimator with the revenue elasticity is one. However, we measure the revenue elasticity as an average for a group of firms, with firm-level differences. De Ridder et al. (2021) demonstrate that the estimator *will* be informative about the *dispersion of markups* within the group. They also find that on the firm-level, the correlation with profitability, market share and labor shares are similar for both quantity- and revenue-based markups.

A solution to the findings of Bond et al. (2021) is to measure the production function for each sector separately, and then compare the dispersion in markups of IT to non-IT firms within each sector. Measuring the production function per sector takes differences in production technologies between sectors into account. However, some sectors contain too few IT firms for a comparison of IT and non-IT. Therefore, we distinguish only between goods and services and assume that firms in goods and services sectors have comparable production technologies. We follow the U.S. Bureau of Labor Statistics in categorizing the NAICS sectors into goods and services (U.S. Bureau of Labor Statistics 2021). The results for the output elasticities for goods and services, with their corresponding standard errors, are shown in Table 3 in Appendix A. To obtain the correct standard errors, we block-bootstrap the estimation procedure while resampling within firms over time, with 100 iterations.

After estimating the production function for goods and services, we distinguish IT from non-IT firms. While the average markup that we find will not be informative according to Bond et al. (2021), the dispersion between the markup of IT and non-IT firms does reflect the true dispersion, as De Ridder et al. (2021) demonstrate. Because the object of this paper is to compare IT to non-IT firms, we aggregate the results of the goods and services sectors for both IT and non-IT firms, using the Cobb–Douglas specification. Online appendix E shows the results for the goods and services sectors separately. Online appendix F shows the results for the translog specification.

4. Comparing IT to non-IT firms

4.1. Data

This section applies the methodology from the previous two sections to data from the Compustat Database North America. This database contains financial statements of all US publicly traded firms, with detailed breakdowns of revenues and costs. We use data between 1978 and 2020. 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 the NAICS code. In line with DEU, 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 22,182 firms.

We then use the trained BERT algorithm to classify all firms as IT or non-IT. Figure 2 shows the number and share of IT firms as a share of the total for all firms. This dataset only includes the observations with complete data on the crucial variables. Of all 22,182 firms in the data, 30% are labeled as IT, representing 21% of total sales. Of the 6718 IT firms, 62% are active in IT sectors and 38% in non-IT sectors. 15% of firms in non-IT sectors are labeled as IT and 88% of firms in IT sectors are labeled as IT.

Table 2 provides the summary statistics. The cost of goods sold is a proxy for the variable costs and includes all expenses directly allocated by the company to production, such as material, direct labor and maintenance. Selling, General and Administrative Expense (SG&A) is a proxy for overhead costs and includes ‘all commercial expenses of operation incurred in the regular course of business pertaining to the securing of operating income’ (Standard & Poor’s 2003, 269). It includes

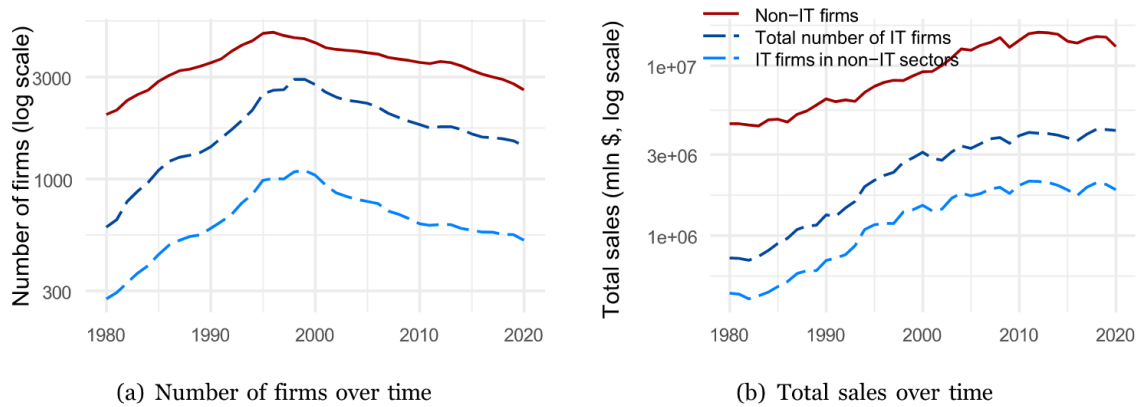


Figure 2. 30% of all publicly traded firms in the USA are labeled as IT by the selection algorithm, representing 21% of total sales. 38% of all IT firms are active in non-IT sectors.

expenses not directly related to product production. Examples of included items are expenses on accounting, marketing and R&D. All nominal values are deflated using the FRED GDP deflator (U.S. Bureau of Economic Analysis 2020). The mean sales, cost of goods sold, capital stock, wage bill and employment of the IT firms are lower than the mean of non-IT firms. 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. The data does not provide reliable information on the mean wage per person, because we can only calculate the mean wage per person for 6% of IT firms and 12% of non-IT firms, due to missing observations on wage bill and employment.

4.2. Markups

We use the methodology explained in Section 3 to estimate the markups. Figure 3 presents the main results of this paper. It compares the markups of the non-IT and IT samples, using the CD production function specification. The markups shown in this section are sales-weighted averages. Appendix D shows the cost-weighted markup. The period since 1980 falls apart in two separate episodes: first, the episode from 1980 to 1996, and then, the episode from 1996 to 2018. Markups of non-IT firms increased from 1.10 in 1980 to 1.37 in 1996 and remained relatively stable afterwards. IT markups are significantly higher for the full period. They were relatively stable between 1980 and 1996 at about 1.45. After 1996, the markups increase to 1.80 in 2018. The rise in markups of IT firms can explain part of the rise in markups of the full sample since 1996. If we look at the goods and services sectors separately, we find that the rise of IT relative to non-IT occurs in the services

Table 2. Summary statistics for the non-IT and IT sample, 1980-2020. Nominal figures are in millions of USD, deflated using the GDP Deflator with base year 2010. Employment is in the thousands.

	Mean	Std. Dev.	25th Pct.	50th Pct.	75th Pct.	90th Pct.	Nr. Obs.
<i>Non-IT firms</i>	2781	14,275	38	227	1117	4458	146,267
Sales (mln \$)							
Cost of goods sold (mln \$)	1894	10,948	22	139	721	2903	146,267
Capital stock (mln \$)	2263	13,425	17	116	686	3249	146,267
SG&A (mln \$)	468	2171	8	36	178	756	146,267
Wage bill (mln \$)	1064	3306	8	90	640	2821	19,871
Employment (thousands)	10	43	0	1	5	20	130,698
<i>IT firms</i>	1452	7452	19	78	380	1843	73,285
Sales (mln \$)							
Cost of goods sold (mln \$)	879	4624	9	38	195	993	73,285
Capital stock (mln \$)	799	6049	5	20	113	680	73,285
SG&A (mln \$)	344	1844	9	30	118	471	73,285
Wage bill (mln \$)	1121	3731	3	30	438	2776	5133
Employment (thousands)	6	27	0	0	2	9	66,121




Figure 3. The average (sales-weighted) IT markup is systematically higher than the average markup of non-IT firms.

sectors. More details can be found in the Online Appendix. Figure 3 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.

We calculate whether the growth differences between IT firms and non-IT firms are statistically significant. We regress markups on a linear time trend for both types of firms, IT and non-IT, and compare the coefficients. We find that the difference in growth rate is significant with a p -value smaller than 0.00001 and that the growth rate of the IT markup is strictly greater than the non-IT markup growth rate. For details, see Appendix B.

Between 1980 and 1996, there is a recovery of non-IT firms from the 1970s when oil increases, union radicalism and severe recessions eroded profitability. The first episode was the days of Ronald Reagan in the USA and Margaret Thatcher in the UK. They ran a program of deregulation, welfare state reduction and confrontation with trade unions. The runaway of IT firms during the second episode, from 1996 onward, coincides with the rise of the internet.

DEU have shown that the upper percentiles of the distribution of the markup are the main driver of the rise of markups in the USA. They claim that this result holds essentially within each sector on its own, suggesting that the rise in the average markup is due to intra-sector and not inter-sector heterogeneity. Figure 4 shows that the same applies for the distributions of the markup of IT and non-IT firms: the increase in markup is concentrated in the upper end of the distribution. Similarly to DEU,




Figure 4. The distribution of the markups (sales weighted) shows that the rise in markups is mainly due to the rise of the 10% highest markups.

the lower percentiles do not explain the increase in markups. The difference between the 90th and 75th percentile is greater for IT than for non-IT firms. For non-IT firms, the 75th percentile remains relatively stable between 1990 and 2010, and decreases slightly after 2010. For IT firms, the 75th percentile shows an increase between 2002 and 2018. The rise in markups of the last decade is more widespread among IT than among non-IT firms.

Figure 5 shows the markup decomposition at the industry level, with IT and non-IT as the two industries. The markup decomposition at the industry level is calculated as follows:

$$\Delta\mu_t = \underbrace{\sum_s m_{s,t-1} \Delta\mu_{st}}_{\Delta\text{within}} + \underbrace{\sum_s \tilde{\mu}_{s,t-1} \Delta m_{s,t}}_{\Delta\text{between}} + \underbrace{\sum_s \Delta\tilde{\mu}_{s,t} \Delta m_{s,t}}_{\Delta\text{cross term}} \quad (3)$$

where $m_{s,t}$ is the market share of the industry s at the time t and μ_{st} the markup of industry s at the time t . The first part is the part of the change in markup that is due to an average increase in markup within the IT and non-IT industries. The reallocation shows how much of the change in markup is due to the reallocation between the non-IT and the IT industries. The cross term is the change in markup due to the joint change in markup and firm composition.

Figure 5 shows that the main driver of the change in markup comes from the change in markup within the non-IT industry, accounting for 56% of the total change in the markup from 1980 until 2018. The change in markup within the IT industry accounts for 17% of the total markup change. Taking into account that IT firms only represent 21% of total sales, this implies that the IT industry contributes more than proportional to the growth in markups. Reallocation represents the reallocation from the non-IT to the IT industry and accounts for 27% of the total change in markup. Thus the reallocation of market share from non-IT to IT firms is a significant driver of the increase in markup that took place in the last decades.

A rise in markup could be due to either rising output elasticity or falling variable cost share, see Equation (1). Because we only distinguish IT from non-IT firms after estimating the output elasticity, the output elasticity of IT and non-IT firms is the same. Figure 6(a) shows the output elasticity. Therefore, the difference in the markup between IT and non-IT firms is due to the differences in variable cost shares.

Figure 6(b) shows that the variable cost share is lower for IT than for non-IT firms. This corresponds with the theoretical literature: information and technology-related firms have a cost structure with a large fixed cost and a small marginal cost (Varian 2001; De Ridder et al. 2019). De Ridder et al. (2019) argue that the increased use of information technology and software can drive an increase in




Figure 5. The industry-level markup decomposition. Reallocation represents the reallocation from the non-IT to the IT industry and accounts for 27% of the total change in markup.




Figure 6. The two components of the markup: output elasticity and the variable cost share. IT firms have a lower sales-weighted expenditure share.

intangible inputs used in production. A key characteristic of intangibles is scalability: they can produce additional output at zero marginal cost. Therefore, a shift towards intangible inputs among IT firms can cause a shift from variable to fixed costs and help explain the low variable costs of IT versus non-IT firms. This cost structure has implications for market concentration: significant economies of scale lead to a cost advantage for the biggest firms. When Amazon was trying to build market share, it believed that scale economies were very important in online retailing and charged low prices to build market share (Varian 2001).

Since the year 2001, we see a decline in the variable cost share of IT firms. This could be due to increased efficiency in production, but also due to globalization. 2001 is the year China accessed the WTO, which was followed by a Chinese export boom to the USA (Handley and Liñao 2017; Acemoglu et al. 2016). Increasing imports from China is likely to decrease the marginal cost due to a fall in the cost of IT components, which are increasingly imported from China. Between 2005 and 2010, the import of computer equipment and parts from China increased with 69% and was the largest category of US imports from China (Morrison 2011). The variable cost share of non-IT firms does not decrease. Possibly, mainly IT firms benefit from increased competition caused by the surge in China–US trades.

4.3. Profitability

A higher markup does not necessarily imply that this group of firms has more market power. It is possible that a firm charges a higher markup to cover increased overhead (or fixed) costs or R&D expenditures. To explore whether the higher markup reflects higher market power, we estimate firm profitability, defined as the sales of a firm minus its total costs. This reflects the conceptual difference between rents and quasi-rents.

In line with DEU, we calculate the profit rate as

$$\text{profit rate} = 1 - \frac{\text{output elasticity}}{\text{markup}} - \frac{r * \text{capital stock}}{\text{sales}} - \frac{\text{overhead}}{\text{sales}}, \quad (4)$$

where the output elasticity over the markup substitutes the expenditure on variable inputs as a share of sales. Overhead is included in Compustat as Selling, General, and Administrative Expense (SG&A), representing ‘all commercial expenses of operation (such as, expenses not directly related to product production) incurred in the regular course of business about the securing of operating income’ (Standard & Poor’s 2003, 269). Examples of included items are expenses on accounting, marketing and R&D.

Compustat includes a measure for the gross capital stock, namely Property, Plant and Equipment (PPEGT). r is a measure for the user cost of capital and is calculated as follows: $r_t = (I_t - \Pi_t) + \Delta$,




Figure 7. The average, sales-weighted, profit of IT firms is high compared to non-IT firms. Since 2010, the gap has increased.

where I_t is the nominal interest rate (we use PIRIC from FRED), Π_t the inflation rate (we use FEDFUNDS from FRED) and Δ the depreciation rate (set at 12%).

Figure 7 shows the profit rates of IT and non-IT firms. IT firms are more profitable than non-IT firms, except during the dot com crisis in 2002, and the difference is increasing over time. The high markup of IT firms is therefore not a quasi-rent, compensation for prior investment and high overhead, but genuine rent derived from barriers to entry.

To analyze the difference between markups and profits in more detail, we isolate overhead and capital costs, and plot them in Figures 8(a) and 8(b), respectively. IT firms have higher overhead than non-IT firms. This corresponds with the cost structure of IT firms as discussed before, namely relatively high fixed and low variable costs. The costs of capital share are higher for non-IT than IT firms. The costs of capital are calculated by multiplying the rate for the user cost of capital with the capital stock, see Equation (4). We assume the user rate of capital to be the same for IT and non-IT firms, so the difference in the user cost of capital is due to a higher capital stock of non-IT firms. Appendix C shows the profits of IT firms as a percentage of total profits and the reported EBIT of IT and non-IT firms.

5. Discussion

The results show a substantially higher markup for IT than for non-IT firms. In this section, we examine various explanations for the findings and how they relate to recent research.

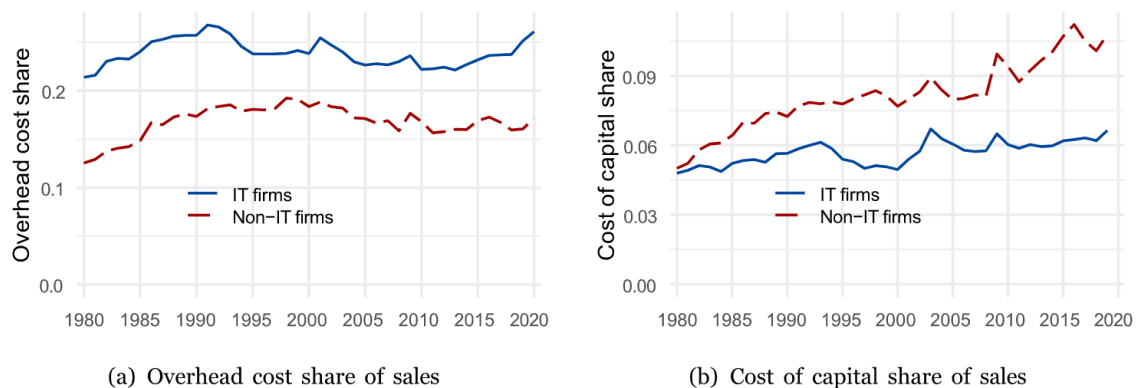


Figure 8. The average overhead cost share of IT firms is higher for IT firms than for non-IT firms, while the capital cost share of IT firms is lower.

Autor et al. (2020) explore the increasing markups and declining labor share in the USA and use a model that can explain these phenomena by the emergence of highly productive ‘superstar’ firms. If globalization or technological changes lead to increased sales for the most productive firms within each industry, industries become dominated by a few highly successful firms, known as superstar firms, which have high markups and a low proportion of labor costs in their value added. They find that information technology can contribute to the emergence of superstar firms. Our results are in line with their hypothesis but show that the role of IT is broader.

Similarly to Autor et al. (2020), our research finds that the rise in markups in the USA is primarily driven by the top percentiles of high markup firms. However, since 2000 the markups in the non-IT group have only risen for the firms with the 10% highest markups, while for IT firms the increase is more widespread. In line with Autor et al. (2020), Figure 9(a) shows that the average firm size is increasing, concurrent with the rising markups. This is the case for both non-IT and IT firms. However, the average IT firm is significantly smaller than the average non-IT firm.

Autor et al. (2020) also find that the ‘superstar firms’ have a low labor share of value added. In addition to the U.S. Economic Census Data, they use the Compustat database to estimate the labor share. Only a small number of firms (13%) report payroll data, which is not a required item for reporting. We therefore cannot draw conclusions, but we can explore the labor share for the firms that do report the payroll data. Autor et al. (2020) define the labor share as the ratio of wage bill to the value added. The best proxy for value added is the sum of the wage bill and EBITDA (earnings before interest, tax, depreciation and amortization). Among the firms that report payroll data, Autor et al. (2020) find a significant drop in the labor share from almost 60% in the early 1980s to 47% in 2015. Using the same data and estimates, Figure 9(b) shows that IT firms have a relatively high labor share. Over the last decade, both the labor share and the markup have increased for IT firms.

The relatively small size of the IT firms, the higher labor share and the more widespread increase in markup among IT firms suggest that, besides superstar firms, IT plays its distinct role in the rising markups. Different phenomena could explain the role of IT firms. Figure 6(b) showed that the higher markup for IT firms is driven by lower variable costs. Economies of scale in IT, with large fixed costs and low variable costs, are stressed both in early literature (Shapiro et al. 1999; Varian 2001), as well as in recent literature (De Ridder et al. 2019; Eeckhout 2022). Information goods can be expensive to create, but once they have been produced, it is relatively inexpensive to make additional copies. This is because the cost of creating the first copy includes significant fixed costs, such as research and development, but the cost of reproducing additional copies and distribution is minimal. The increasing average firm size, shown in Figure 9(a), indicates that IT firms are increasingly benefiting from economies of scale, resulting in declining variable costs.

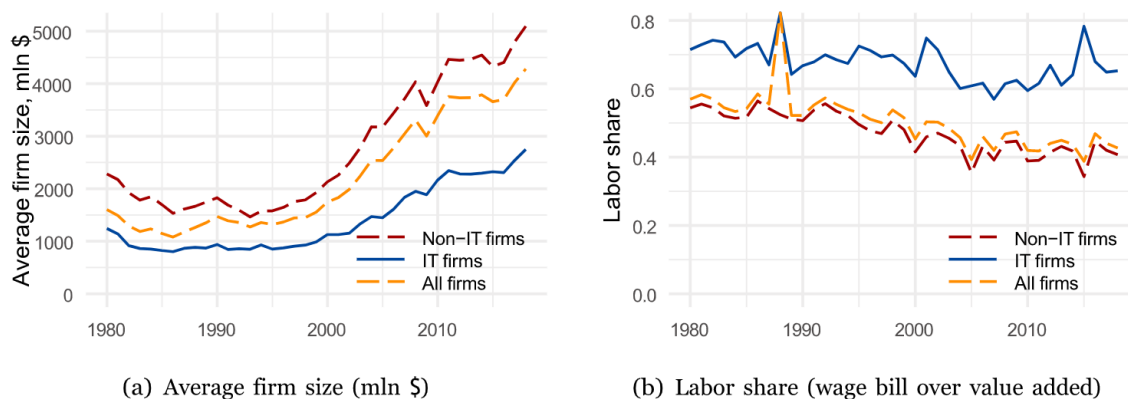


Figure 9. The average firm size of IT firms is small compared to non-IT firms. The labor share for IT firms has slightly increased over the last decade.

Figure 7 demonstrated that IT firms are more profitable than non-IT firms, even when taking fixed costs into account. This suggests that IT firms can generate rent on top of their investment and overhead costs, which may be due to a combination of factors such as lessened competition and higher-quality products. According to Berry, Gaynor, and Scott Morton (2019), the higher fixed costs for firms that rely heavily on information technology can lead to barriers to entry and higher markups. Accordingly, we find lower variable costs, higher fixed costs and higher markups for the group of IT firms, suggesting that IT firms can achieve higher profits due to a combination of factors including lessened competition, higher quality products and lower marginal costs.

The drop in variable costs of IT firms could also be an effect of globalization. Companies with an international supply chain have the opportunity to source cheaper materials, resulting in higher profit margins. The increased globalization may also contribute to an unequal distribution of this increase in profit margins (Berry, Gaynor, and Scott Morton 2019). IT in particular may have benefited from globalization, when in 2001 China accessed the WTO. Imports of information technology components particularly from China surged, which could explain the fall in costs for IT firms that we find.

Finally, data can play a role in increasing markups. Antitrust authorities have raised concerns that the data collected by digital platforms gives incumbents a competitive advantage in digital markets (Calvano and Polo 2021). Data could provide an advantage to incumbent firms and a barrier to entry for smaller competitors. This advantage can also be based on economies of scale that allow firms with more customers to amass more data, improve their service and attract even more customers. Eeckhout and Veldkamp (2022) examine how the use of data by firms could potentially lead to market power. They develop a model in which economies of scale in data lead to a data-rich firm investing in lower marginal cost and larger-scale production. In their setting, firms are utilizing data to shift production towards high-markup goods, so that markups at the firm- and industry-level increase. Other researchers argue that big data is not necessarily rare or inimitable and that there are many alternative data sources available to new entrant (Lambrecht and Tucker 2017). Further research is needed to better understand to what extent data access forms a barrier to entry.

While our study does not establish a causal relationship, we find a significant divergence in outcomes for IT and non-IT firms. Characteristics that contribute to market power, as identified in previous research, are particularly relevant to the information technology industry. Our results and previous research find that IT firms play a significant and distinct role in the rising market power.

6. Conclusion

This paper shows that the average markup and profits are significantly higher for IT than for non-IT firms. The period of the markup since 1980 falls apart in two separate episodes: first, the episode from 1980 to 1996, and then, the episode from 1996 to 2018. In the first episode, there is a recovery of all firms from the seventies when oil increases, union radicalism, and severe recessions eroded profitability. The first episode was 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. During the second episode, from 1996 onward, there is a runaway of IT firms. This runaway coincides with the rise of the internet. We find the runaway of the markup can be explained by a fall in the variable cost share of IT firms since 2001. This could be due to increased efficiency in production, but also due to globalization and economies of scale. In 2001, China accessed the WTO and imports of IT components from China surged, which could explain the fall in costs for IT firms. Strong economies of scale for information technology goods can mean that especially IT firms benefit from the increasing average firm size.

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 2018. For IT firms, both the 75th and the 90th percentile markups increase since 2002,

while for non-IT, only the 90th percentile increased over the last two decades. This implies that a rise in markups is more widespread among IT firms than among non-IT firms.

To analyze if the higher markup implies that IT firms have more market power, we also look at profitability. If, despite the higher markup, IT firms are not more profitable than non-IT firms, it is possible that a firm charges a higher markup to cover increased overhead (or fixed) costs or R&D expenditures. However, we find that IT firms are more profitable than non-IT firms, except during the dot com crisis in 2002, and that the difference is increasing over time. The high markup of IT firms is therefore not a quasi-rent, compensation for prior investment and high overhead, but genuine rent derived from barriers to entry.

DEU state that rising markups mainly occur *within* sectors and that the difference in the development of markups *between* sectors 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 DEU.

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*. Sectors 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 sectors in the USA. Information technology can be a contributor to the rise of 'superstar' firms. Additionally, our research finds that IT firms play a broader role in driving up markups, especially through their low variable cost share. It would be interesting to apply the firm-level classification of IT and non-IT to the research by Autor et al. (2020), to further distinguish the role of IT from superstar firms, and the contribution of IT to the rise of superstar firms.

Acknowledgement

We thank Maarten de Ridder for the valuable remarks.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Data availability statement

The data that support the findings of this study are available from WRDS (Wharton Research Data Services) Compustat North America database. Restrictions apply to the availability of these data, which were used under license for this study. Data are available from the authors with the permission of WRDS.

References

- Acemoglu, D., D. Autor, D. Dorn, G. H. Hanson, and B. Price. 2016. "Import Competition and the Great US Employment Sag of the 2000s." *Journal of Labor Economics* 34 (S1): S141–S198. doi:[10.1086/682384](https://doi.org/10.1086/682384)
- Akerberg, D. A., K. Caves, and G. Frazer. 2015. "Identification Properties of Recent Production Function Estimators." *Econometrica* 83 (6): 2411–2451. doi:[10.3982/ECTA13408](https://doi.org/10.3982/ECTA13408)
- Arellano, M., and S. Bond. 1991. "Some Tests of Specification for Panel Data: Monte Carlo Evidence and an Application to Employment Equations." *The Review of Economic Studies* 58 (2): 277–297. doi:[10.2307/2297968](https://doi.org/10.2307/2297968)
- 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. doi:[10.1093/qje/qjaa004](https://doi.org/10.1093/qje/qjaa004)
- Berry, S., M. Gaynor, and F. Scott Morton. 2019. "Do Increasing Markups Matter? Lessons from Empirical Industrial Organization." *Journal of Economic Perspectives* 33 (3): 44–68. doi:[10.1257/jep.33.3.44](https://doi.org/10.1257/jep.33.3.44)
- Bessen, J. E. 2017. "Information Technology and Industry Concentration." *Law & Economics Paper Series*.
- Bessen, J., and C. Righi. 2019. "Shocking Technology: What Happens When Firms Make Large IT Investments?" *Boston Univ. School of Law, Law and Economics Research Paper*.
- Bond, S., A. Hashemi, G. Kaplan, and P. Zoch. 2021. "Some Unpleasant Markup Arithmetic: Production Function Elasticities and Their Estimation from Production Data." *Journal of Monetary Economics* 121: 1–14. doi:[10.1016/j.jmoneco.2021.05.004](https://doi.org/10.1016/j.jmoneco.2021.05.004)

- 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. doi:[10.1080/10438599500000002](https://doi.org/10.1080/10438599500000002)
- Calligaris, S., C. Criscuolo, and L. Marcolin. 2018. "Mark-Ups in the Digital Era." *OECD Science, Technology and Industry Working Papers* 2018/10.
- Calvano, E., and M. Polo. 2021. "Market Power, Competition and Innovation in Digital Markets: A Survey." *Information Economics and Policy* 54: 100853. doi:[10.1016/j.infoecopol.2020.100853](https://doi.org/10.1016/j.infoecopol.2020.100853)
- 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. doi:[10.1080/08874417.2001.11647002](https://doi.org/10.1080/08874417.2001.11647002)
- Crouzet, N., and J. C. Eberly. 2019. "Understanding Weak Capital Investment: The Role of Market Concentration and Intangibles." Technical Report, National Bureau of Economic Research.
- 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. doi:[10.1093/qje/qjz041](https://doi.org/10.1093/qje/qjz041)
- De Ridder, M., et al. 2019. "Market Power and Innovation in the Intangible Economy." *Cambridge Working Papers in Economics*.
- De Ridder, M., B. Grassi, G. Morzenti, et al. 2021. "The Hitchhiker's Guide to Markup Estimation." Technical Report.
- 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*.
- 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. doi:[10.1080/1043859042000201911](https://doi.org/10.1080/1043859042000201911)
- Edmond, C., V. Midrigan, and D. Y. Xu. 2018. "How Costly Are Markups?" Technical Report, National Bureau of Economic Research.
- Eeckhout, J. 2022. "The Profit Paradox." In *The Profit Paradox*, edited by Joe Jackson and Josh Drake, 42–67. Princeton: Princeton University Press.
- Eeckhout, J., and L. Veldkamp. 2022. "Data and Market Power." Technical Report, National Bureau of Economic Research.
- Hall, R. E. 1988. "The Relation Between Price and Marginal Cost in us Industry." *Journal of Political Economy* 96 (5): 921–947. doi:[10.1086/261570](https://doi.org/10.1086/261570)
- Handley, K., and N. Limão. 2017. "Policy Uncertainty, Trade, and Welfare: Theory and Evidence for China and the United States." *American Economic Review* 107 (9): 2731–2783. doi:[10.1257/aer.20141419](https://doi.org/10.1257/aer.20141419)
- Katz, M. L., and C. Shapiro. 1986. "Technology Adoption in the Presence of Network Externalities." *Journal of Political Economy* 94 (4): 822–841. doi:[10.1086/261409](https://doi.org/10.1086/261409)
- Laffont, J. J., P. Rey, and J. Tirole. 1998a. "Network Competition: I. Overview and Nondiscriminatory Pricing." *The RAND Journal of Economics*, 1–37. doi:[10.2307/2555814](https://doi.org/10.2307/2555814)
- Laffont, J. J., P. Rey, and J. Tirole. 1998b. "Network Competition: II. Price Discrimination." *The RAND Journal of Economics*, 38–56. doi:[10.2307/2555815](https://doi.org/10.2307/2555815)
- Lambrecht, A., and C. E. Tucker. 2017. "Can Big Data Protect a Firm from Competition?" *Competition Policy International CPI Antitrust Chronicle*.
- Morrison, W. M. 2011. "China-US Trade Issues." *CRS Report for Congress*.
- Olley, G. S., and A. Pakes. 1996. "The Dynamics of Productivity in the Telecommunications Equipment Industry." *Econometrica* 64 (6): 1263–1297. doi:[10.2307/2171831](https://doi.org/10.2307/2171831)
- Rochet, J. C., and J. Tirole. 2003. "Platform Competition in Two-Sided Markets." *Journal of the European Economic Association* 1 (4): 990–1029. doi:[10.1162/154247603322493212](https://doi.org/10.1162/154247603322493212)
- Shapiro, C., H. R. Varian, S. Carl, et al. 1999. *Information Rules: A Strategic Guide to the Network Economy*. Cambridge: Harvard Business Press.
- Sokolova, M., and G. Lapalme. 2009. "A Systematic Analysis of Performance Measures for Classification Tasks." *Information Processing & Management* 45 (4): 427–437. doi:[10.1016/j.ipm.2009.03.002](https://doi.org/10.1016/j.ipm.2009.03.002)
- Standard & Poor's. 2003. "Standard & Poor's Compustat User's Guide." Technical Report. Standard & Poor's/Compustat (2020). Compustat North America. Retrieved from Wharton Research Data Service. Accessed 5 September 2020; URL: <http://wrds-web.wharton.upenn.edu/wrds/>.
- 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.
- U.S. Bureau of Economic Analysis. 2020. "Gross Domestic Product: Implicit Price Deflator [GDPDEF]." Technical report. Accessed: 2020-07-01.
- U.S. Bureau of Labor Statistics. 2021. "Industries by Supersector and NAICS Code." Accessed: 2021-10-11.
- Varian, H. R. 2001. *Economics of Information Technology*. Berkeley: University of California.

Appendices

Appendix A. The empirical framework for measuring the output elasticity

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 DEU. We reiterate their argument here for the sake of transparency. We use a translog production function

$$q_{it} = \beta_t^v v_{it} + \beta_t^k k_{it} + \beta_t^{vv} v_{it}^2 + \beta_t^{kk} k_{it}^2 + \omega_{it} + \int_{it}, \quad (\text{A1})$$

where the β_t 's are parameters and lower cases are the logs of the corresponding upper cases; k_{it} denotes the log capital stock. (Following DEU, 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 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 scenarios prevails depends on the elasticity of the firm's demand curve. We resolve this problem by applying a two-stage estimation procedure, building on Olley and Pakes (1996) and Akerberg, Caves, and Frazer (2015).

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:

$$q_{it} = \phi_t(v_{it}, k_{it}) + \varepsilon_{it}, \quad (\text{A2})$$

where

$$\phi_t = \beta_t^v v_{it} + \beta_t^k k_{it} + \beta_t^{vv} v_{it}^2 + \beta_t^{kk} k_{it}^2 + \omega_{it}. \quad (\text{A3})$$

To approximate ϕ_t , we model the output 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 (A3), this AR(1) process can be written as follows:

$$\hat{\phi}_t - \beta_t^v v_{it} - \beta_t^k k_{it} - \beta_t^{vv} v_{it}^2 - \beta_t^{kk} k_{it}^2 = g(\hat{\phi}_{t-1} - \beta_{t-1}^v v_{it-1} - \beta_{t-1}^k k_{it-1} - \beta_{t-1}^{vv} v_{it-1}^2 - \beta_{t-1}^{kk} k_{it-1}^2) + \xi_{it}. \quad (\text{A4})$$

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:

$$E \left[\begin{pmatrix} \xi_{it}(\beta_t^v, \beta_t^{vv}) \begin{pmatrix} v_{it-1} \\ v_{it-1}^2 \end{pmatrix} \end{pmatrix} \right]$$

These moment conditions assume that the correlation of the productivity process's error term and the lagged variable input use should be zero (see (A4)). In other words; lagged variable input should not respond to current productivity shocks.

The GMM procedure evaluates the moment conditions (A5) 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 β_t^v and β_t^{vv} , namely, the values for which the moment conditions are minimized. Table A1 shows the resulting output elasticities for goods and services. To obtain the correct standard errors, we block-bootstrap the estimation procedure while resampling within firms over time, with 100 iterations.

Table A1. The output elasticities are shown separately for goods and services with their corresponding standard errors.

	Services output elasticity	Standard error	Goods output elasticity	Standard error
1980	0.824	0.002	0.799	0.002
1981	0.829	0.003	0.804	0.002
1982	0.836	0.002	0.805	0.002
1983	0.842	0.002	0.809	0.002
1984	0.849	0.002	0.814	0.002
1985	0.853	0.002	0.824	0.002
1986	0.853	0.002	0.828	0.002
1987	0.852	0.002	0.839	0.002
1988	0.849	0.002	0.845	0.002
1989	0.842	0.002	0.843	0.002
1990	0.842	0.002	0.846	0.002
1991	0.842	0.002	0.850	0.002
1992	0.844	0.002	0.852	0.002
1993	0.842	0.002	0.855	0.002
1994	0.837	0.002	0.851	0.002
1995	0.829	0.002	0.846	0.002
1996	0.822	0.002	0.840	0.002
1997	0.818	0.002	0.840	0.002
1998	0.818	0.002	0.835	0.002
1999	0.815	0.002	0.831	0.002
2000	0.812	0.002	0.834	0.003
2001	0.812	0.002	0.827	0.003
2002	0.805	0.002	0.819	0.003
2003	0.797	0.003	0.822	0.003
2004	0.787	0.003	0.821	0.003
2005	0.783	0.002	0.814	0.003
2006	0.776	0.002	0.808	0.003
2007	0.770	0.002	0.809	0.004
2008	0.765	0.003	0.813	0.003
2009	0.765	0.003	0.809	0.004
2010	0.767	0.003	0.812	0.004
2011	0.768	0.003	0.817	0.004
2012	0.760	0.003	0.808	0.005
2013	0.769	0.003	0.803	0.005
2014	0.778	0.003	0.815	0.005
2015	0.773	0.003	0.817	0.005
2016	0.777	0.003	0.824	0.004
2017	0.777	0.003	0.834	0.005
2018	0.765	0.004	0.819	0.005
2019	0.759	0.004	0.820	0.005

As a robustness check for the productivity estimation, we re-estimated the production function using a dynamic panel estimator, following Arellano and Bond (1991). The results for the output elasticity of the variable input from the Arellano–Bond estimator are comparable to the estimates in our paper. The average output elasticity for the variable input is 0.74, although the variance is greater. Detailed results are available upon request.

To compare if our results are consistent with DEU, we measure the output elasticities for the full sample from 1970 onwards and for each sector separately, in line with DEU. Figure 13 compares the estimates of the markups obtained by DEU to the estimates obtained 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 because DEU use a different deflator for the cost of goods sold.

Appendix B. The difference in the markup growth rates

We calculate whether the growth differences between IT firms and non-IT firms are statistically significant. We regress markups on a linear time trend for both types of firms, IT and non-IT, and compare the coefficients. To be precise, first we regress the markup on a time trend, once with the data of IT firms, and once with the data of non-IT firms. We thus obtain the coefficients of the growth rates of the markups (the time trend) for both IT and non-IT firms. Then we calculate the difference between the two as follows: growth rate IT minus growth rate non-IT. To obtain the correct standard errors, we block-bootstrap the entire procedure, and obtain 500 values for the difference between the two growth rates. Figure B1 shows the histogram of the difference in growth rates. The average growth rate for the IT markup is 0.006114 per year, while the average growth rate for the non-IT markup is 0.000146. The minimal value of the difference is 0.054, meaning that the growth rate of the IT markup is strictly greater for all 500 samples.




Figure B1. The histogram shows significant differences in markup growth rate (growth rate IT minus growth rate non-IT).

We calculate whether the difference is significant by calculating the z-score and corresponding p -value. The mean of the difference is 0.00597, with a standard error of 0.000008001 and sample size 500. This implies that the p -value is smaller than 0.00001, meaning the difference is highly significant.

Appendix C. Profits

Figure C1 shows the profits of IT firms as a percentage of the 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.




Figure C1. IT profits represent a large share of total profits, while IT sales are relatively small.

Figure C2 shows the EBIT (as reported in the Compustat database) as a share of firm sales and looks similar to the profits we calculate and shown in Figure 7. The sales-weighted EBIT shares of IT firms show a slump during the dot com bubble and then increase between 2001 and 2020. The EBIT share of non-IT firms remains roughly stable between 1980 and 2020. It lies above the EBIT share of IT firms between 1986 and 2009.




Figure C2. Reported EBIT (sales-weighted) of non-IT firms is higher between 1986 and 2014, when it is overtaken by the EBIT of IT firms.

Appendix D. Cost-weighted markups

Edmond, Midrigan, and Xu (2018) find that the rise in markups largely disappears for *cost-weighted* (using the cost of goods sold from Compustat) aggregates of the firm-level markups, as opposed to *sales-weighted* averages used by DEU. We check if our results are robust when we cost-weight the markups; see Figure D1. We find that the rise in markups for all firms indeed largely disappears, as found by Edmond, Midrigan, and Xu (2018). However, the dispersion in markups between IT and non-IT persists for cost-weighted averages.




Figure D1. Cost-weighting instead of sales-weighting the markups leads to a similar dispersion of markups between IT and non-IT firms.