

New Evidence on the Markup of Prices over Marginal Costs and the Role of Mega-Firms in the US Economy

Robert E. Hall
Hoover Institution and Department of Economics
Stanford University

rehall@stanford.edu; stanford.edu/~rehall

April 27, 2018

Abstract

The markup of price over marginal cost reveals market power. The distinction between marginal and average cost is key. Average cost is easy to measure, but the price/average cost ratio understates the price/marginal cost ratio when fixed costs are present. In particular, in free-entry equilibrium, where revenue equals cost, the price/average cost ratio is always one, while the price/marginal cost ratio may be above one. The idea here is to calculate marginal cost as the ratio of the adjusted expenditure on inputs to the adjusted change in output. The first adjustment is to remove the change in expenditure that arises from the changes in input costs. The second adjustment is to remove the change in output attributed to productivity growth. Application to KLEMS productivity data finds a typical markup ratio of 1.3. Markup ratios grew between 1988 and 2015. For mega-firms, the paper uses employment at firms with 10,000+ workers. Substantial heterogeneity occurs across sectors and in growth rates. There is no evidence that mega-firm-intensive sectors have higher price/marginal cost markups, but some evidence that markups grew in sectors with rising mega-firm intensity.

JEL L1 D24

1 Introduction

The ratio of price to marginal cost is an interesting quantity in many branches of economics. A profit-maximizing price-taking firm equates its marginal cost to the prevailing price of output. That price is invariant to the firm's output choice. The firm's price/marginal cost ratio is one. A firm facing a constant-elastic residual demand, with elasticity ϵ , maximizes profit at the point where the ratio of price to marginal cost is $\epsilon/(\epsilon - 1)$. In general, the ratio of price to marginal cost, designated μ in this paper, is a useful way to think about market power or monopoly power. It has a simple functional relationship to an equivalent measure, the profit margin on sales, or Lerner index,

$$\mathcal{L} = 1 - \frac{1}{\mu}, \tag{1}$$

which maps the price/marginal cost ratio from $\mu \in [1, \infty]$ to $\mathcal{L} \in [0, 1]$.

The literature on measurement of μ has two main branches. The demand-side approach infers the residual elasticity ϵ , typically from a differentiated-products oligopoly model. The supply side-approach uses data on price and cost from firms. Thirty years ago, Hall (1988) proposed a refinement of the supply-side approach that measures marginal cost rather than average cost. De Loecker and Warzynski (2012) proposed a related method also focusing on empirical marginal cost, and De Loecker and Eeckhout (2017) recently captured a great deal of attention with the finding that μ has risen substantially in the US in recent decades. See Traina (2018) and Gutiérrez and Philippon (2017) for critiques of that paper. Since the publication of my paper, much improved data have become available, thanks to the efforts of US statistical agencies in developing productivity data. Their compiled data feed directly into calculations of μ . In addition, De Loecker and Eeckhout's recent paper together with a literature on the rising importance of large firms and the decline in the labor share has generated great interest in the growth of the price/marginal cost ratio. This paper responds to those developments. It finds support for the conclusion that the ratio has risen in recent decades, though by less than in De Loecker and Eeckhout's paper.

One reason for the rise in market power revealed in the price/marginal cost literature may be increasing concentration, particularly the rising role of mega-firms in some US industries. Most work on this topic has used data from publicly traded firms. In data from the US Economic Census covering all firms (not establishments), which reports employment in firms by number of employees, including those with 10,000 or more, which I designate as

mega-firms, increases in the mega-firm employment fraction are modest where they occur. Further, important sectors including manufacturing experienced declines in the mega-firm employment fraction. I find no systematic relation between the level mega-firm fraction and the price/marginal cost ratio across 19 major sectors of the US economy. But there is moderately strong evidence that industries with growing mega-firm fraction have gained market power in the years since 1998.

This paper is self-contained and presumes no acquaintance with my earlier work or other work on this subject. Nonetheless, everything here is new, including a novel derivation of the basic idea of extracting marginal cost from time-series data.

2 Framework

2.1 Measuring the ratio of price to marginal cost, μ

In time-series data, a natural measure of marginal cost is the change in cost divided by the change in output. More precisely, the numerator is the change in cost not associated with changes in factor prices and the denominator is the change in output not associated with the change in Hicks-neutral productivity. Cost is

$$c = \sum_i p_i x_i, \tag{2}$$

in obvious notation. The change in cost is

$$dc = \sum_i x_i dp_i + \sum_i p_i dx_i. \tag{3}$$

The first summation is the component associated with changes in factor prices, while the second is the desired component purged of effects from changing factor prices:

$$\tilde{d}c = \sum_i p_i dx_i. \tag{4}$$

The technology is

$$y = A f(x), \tag{5}$$

so the part of output growth to leave out of $\tilde{d}y$ is

$$f(x) dA = y \frac{dA}{A}. \tag{6}$$

Thus

$$\tilde{d}y = dy - y \frac{dA}{A}. \quad (7)$$

Marginal cost is $\tilde{d}c/\tilde{d}y$. The price/marginal cost ratio is

$$\mu = \frac{p}{\tilde{d}c/\tilde{d}y} \quad (8)$$

or

$$p dy = \mu \tilde{d}c + \frac{dA}{A}. \quad (9)$$

Now let

$$\alpha_i = \frac{p_i x_i}{p y}. \quad (10)$$

the share of factor i in revenue, $p y$. The expression for output growth can then be written

$$\frac{dy}{y} = \mu \sum_i \alpha_i \frac{dx_i}{x_i} + \frac{dA}{A}. \quad (11)$$

With discrete time, the same equation is

$$\Delta \log y = \mu \sum_i \alpha_i \Delta \log x_i + \Delta \log A. \quad (12)$$

This expression for the price/marginal cost ratio is useful because the quantities $\Delta \log y$ and $\sum_i \alpha_i \Delta \log x_i$ are calculated meticulously in productivity accounts based on Solow (1957). These accounts publish the value of the Solow residual,

$$\Delta \log \hat{A} = \Delta \log y - \sum_i \alpha_i \Delta \log x_i, \quad (13)$$

so the measure of the change in total factor input, $\sum_i \alpha_i \Delta \log x_i$, can be calculated by subtracting the Solow residual from the rate of change of output. Note that if $\mu > 1$, the Solow residual does not measure actual technical progress, because it does not adjust for market power.

Notice that this derivation of the measurement of μ does not assume anything about optimal choice by the firm, apart from remaining on its production function. The firm is not necessarily satisfying its first-order conditions in the output market or any input market. The coefficient μ does not necessarily describe the residual demand function facing the firm, effects of market power by sellers of inputs including labor unions, or monopsony power of the firm in those input markets.

The growth rate of productivity, $a = \Delta \log A$, is a statistical residual in equation (12). It can only be measured with knowledge of the price/marginal cost ratio μ . The most basic approach is to treat μ as a parameter to be estimated in time-series or panel data, with suitable instrumental variables. Eligible instruments are variables that are uncorrelated with productivity growth but are correlated with output and inputs. The residual based on the estimated value of μ is the estimated rate of true productivity growth, adjusted for market power.

With a single time series, the specification for μ may capture changes over time, with a small number of parameters. For example, an equation that considers a linear trend is

$$\Delta \log y_t = (\phi + \psi t) \sum_i \alpha_{i,t} \Delta x_{i,t} + a_t. \quad (14)$$

Here ϕ controls the level of μ and ψ is the per-period growth of μ . With panel data, the specification of μ may capture differences in market power across industries as well.

2.2 Potential sources of bias in the measurement of the markup ratio

Here I consider a number of cases where the procedure described earlier either measures marginal cost accurately or measures it with a bias of known sign. The estimate of the ratio of price to marginal cost, μ , is the slope coefficient of the relation of $\Delta \log y$ to $\sum_i \alpha_i \Delta \log x_i$. The examples considered here all involve potential biases in one of the α_i s. In all of them, the input quantities are taken to be measured correctly. If an α_i is biased downward, the slope coefficient μ is also biased downward.

A firm makes excess profit. The productivity data used in this paper measure the price of capital services as a residual—the payments to factors $\sum_i p_i x_i$ add up to revenue $p y$. Consequently, the price of capital services is overstated for a firm that enjoys excess profit, and the share α_i is correspondingly overstated. Thus μ is overstated. This effect is likely to be small because changes in capital from year to year are usually quite small. To cure this problem, one could replace the residual price of capital with a price inferred from the formula for the rental price of capital.

An industry is in zero-excess-profit equilibrium with monopolistic competition. With zero excess profit, the measured level of profit equals the market rental cost of capital, so the

problem discussed above does not arise. The measured ratio of price to marginal cost reflects the actual market power.

The firm purchases an input in a market where sellers of the input have market power. The leading example is unionized labor. If a seller of an input exercises its market power by setting a higher price that reflects that power, the calculation described in this paper takes account of the true marginal cost associated with that input, and the calculation uncovers the true price/marginal cost ratio of the firm. Notice that such an arrangement is bilaterally inefficient. If the firm and the input seller use efficient two-part pricing, the average price paid exceeds the underlying marginal price. In that case, the calculation overstates μ . Because labor's input share is generally large and the change in labor input is also generally large, the bias is probably more important than the one arising from monopoly profits.

The firm has monopsony power in an input market. The leading example is a firm whose employment level is a substantial fraction of total employment in its labor market. The average price paid for the input understates the effective marginal price. The share α_i is understated and the estimate of μ is correspondingly understated.

3 Data

The data in the Solow productivity framework come from `klemscombinedbymeasure.xlsx`, available at bls.gov/mfp/mprdload.htm#Multifactor%20Productivity%20Tables. See bls.gov/mfp/#technotes for extensive technical descriptions of the data. The data are annual starting in 1987. I use the version of the data for 60 distinct non-overlapping industries. Some of the advantages of the data relative to data in earlier work on production-side measurement of the price/marginal cost literature are:

- Rigorous adherence to proper measurement of output—no reliance on value added
- Uniform use of the modern NAICS industry definitions
- Breakdown of inputs into 5 categories: capital, labor, energy, materials, and services
- Aggregation of capital and labor inputs from detailed underlying data using appropriate methods
- Use of Tørnqvist's refinement of the weights applied to log-changes in factor inputs




Figure 1: Instrumental Variables

For instrumental variables, I follow the identification strategy of Hall (1988), which treats an industry’s productivity growth as orthogonal to government purchases of military goods and services and to movements of the oil price. NIPA table 3.11.3 breaks down real military purchases into a variety of categories. I use FRED series ACOILWTICO, the market price of west Texas intermediate crude, as a measure of the oil price. The instruments are:

- Military purchases of equipment
- Military purchases of ships
- Military purchases of software
- Military expenditure on research and development
- The oil price

All of these enter as log differences.

Data on the fraction of employment by industry of very high employment firms (10,000 or more employees) come from the SUSB database compiled from business census data. See census.gov/data/tables/2015/econ/susb/2015-susb-annual.html, link to “U.S., NAICS sectors, larger employment sizes up to 10,000+” and similar for years back to 1998. These data

are for 19 NAICS sectors and begin in 1998. Data on the distribution of firm employment at more detailed industry levels is available, but not including the employment at mega-firms, thanks to disclosure restrictions. Apparently many more detailed industries have only one or a small number of mega-firms, so publication of the data would violate the anonymity of census data.

4 Results

From the perspective of research on total factor productivity in Solow’s framework, it is natural to write the estimating equation with output growth on the left and weighted total input growth on the right side. In principle, in an instrumental variables setting, the results of estimating μ in that normalization would be the same as flipping the variables, estimating μ^{-1} and then taking the reciprocal as the estimate of μ . With a single instrumental variable, the results are numerically identical. In the KLEMS data using the 5 instruments in this paper, the principle of reciprocity comes close to holding. But for the reason discussed in Hall (1988), it is better to estimate μ^{-1} . As noted earlier, μ ranges from 1 to infinity. Very high values will occur in cases where variable inputs are unimportant—software and proprietary pharmaceuticals are examples. The reported standard error for these cases will be high. The results in those cases make more sense where μ^{-1} is estimated, and the high and uncertain values of μ are mapped into a small region around 1.

4.1 First-stage results

With the normalization that output growth is on the right side of the equation, the first-stage regressions have output growth on the left and the five instruments, as log-changes, on the right. The KLEMS data form a panel with 60 industries and 28 years when stated as log differences. Because the instruments are all time series, cross-section regressions in single years, or in small groups of years, are not identified. I focus on the 60 first-stage regressions where each industry contributes a time-series OLS regression. The question at hand is whether the instruments have adequate power to support instrumental-variables estimation.

Table 1 describes the power of the instruments in terms of the ratios of the first-stage coefficients to their standard errors. The first column gives the percent of the 60 ratios that exceed two in absolute value. These numbers would be about 5 percent if the data

<i>Instrument</i>	<i>Percent of first-stage t-statistics > 2</i>
Military purchases of equipment	38
Military purchases of ships	7
Military purchases of software	60
Military expenditure on research and development	17
Oil price	35
Average	31

Table 1: Metrics for the Power of the Instruments

were purely random. All 5 of the instruments outperform that standard, by a considerable margin in all cases but one.

4.2 Estimates of the ratio of price to marginal cost, μ , by industry

The results for the 60 industries are too extensive to digest in a single table. Table 2 summarizes them in aggregates at the level of 19 NAICS sectors, sorted by the estimate of the estimated markup ratio μ . The table presents averages across the industries contained in the sectors. The standard errors of the coefficients are summarized as averages and should not be confused with the standard errors that would result from aggregating the underlying data and estimating a single coefficient. The coefficients for four sectors—information, utilities, finance–insurance, and agriculture–fishing–hunting—are sufficiently large to render the estimates questionable. In these cases, the instruments lack the power to identify the markup ratio with usable accuracy.

From Table 2, it is clear that the estimates specific to the industries have a good deal of sampling error. In particular, 31 percent of the industries have values of $\hat{\mu}_i$ below the minimal possible value of one. To disentangle the distribution of the true values of the price/marginal cost ratio across industries from the distribution of sampling error, I consider a simple statistical model that exploits the fact that sampling error must have a role sufficient to explain the 31 percent of values of the ratio that are estimated to be below one. The statistical model is

$$\hat{\mu} = \eta + \nu + 1, \tag{15}$$

<i>Weighted averages across industries</i>					
<i>Inverse of markup ratio</i>	<i>Standard error</i>	<i>Markup ratio</i>	<i>Percent of value added in sector</i>	<i>Number of industries in sector</i>	<i>Sector name</i>
1.13	(0.11)	0.89	5.1	3	Health Care and Social Assistance
1.05	(0.10)	0.95	0.2	1	Educational Services
1.02	(0.15)	0.98	6.5	1	Construction
0.97	(0.07)	1.04	3.9	2	Administrative and Support and Waste Management and Remediation Services
0.93	(0.12)	1.09	6.0	2	Real Estate and Rental and Leasing
0.90	(0.43)	1.12	2.7	1	Utilities
0.84	(0.26)	1.19	2.4	1	Management of Companies and Enterprises
0.91	(0.39)	1.26	1.4	3	Mining, Quarrying, and Oil and Gas Extraction
0.79	(0.10)	1.27	7.0	1	Wholesale Trade
0.77	(0.10)	1.31	9.0	3	Professional, Scientific, and Technical Services
0.81	(0.09)	1.31	4.1	8	Transportation and Warehousing
0.75	(0.17)	1.33	2.8	1	Other Services (except Public Administration)
0.92	(0.21)	1.39	5.2	4	Information
0.79	(0.06)	1.41	21.3	18	Manufacturing
0.69	(0.15)	1.45	8.0	1	Retail Trade
0.72	(0.28)	1.46	8.5	4	Finance and Insurance
0.71	(0.17)	1.52	1.0	2	Arts, Entertainment, and Recreation
0.65	(0.09)	1.55	3.1	2	Accommodation and Food Services
0.54	(0.64)	1.85	1.7	2	Agriculture, Forestry, Fishing and Hunting

Table 2: Estimates of the Ratio of Price to Marginal Cost by Industry, Stated as Sector Averages

where the sample error η is normal($0, \gamma$) and the log of the true markup's random component is normal (δ, σ). The two components are independent. Identification rests on the hypothesis that $\mu = \nu + 1 \geq 1$. To find the parameters, I match the first three uncentered moments of $\hat{\mu} - 1$. These are

$$\text{Mean}(\hat{\mu} - 1) = \exp\left(\delta + \frac{1}{2}\sigma^2\right), \quad (16)$$

$$\text{Mean}(\hat{\mu} - 1)^2 = \gamma^2 + \exp(2\delta + 2\sigma^2), \quad (17)$$

and

$$\text{Mean}(\hat{\mu} - 1)^3 = \exp\left(3\delta + \frac{9}{2}\sigma^2\right) + 3\gamma^2 \exp\left(\delta + \frac{1}{2}\sigma^2\right). \quad (18)$$

These equations form a well-conditioned system that is easy to solve with a Newton-type solver.

Table 3 shows the inputs to and results of these calculations based on the distribution of estimates $\hat{\mu}$. The upper panel shows the uncentered moments of the 60 estimates of the markup ratio. The middle panel translates those moments into the solved values of the underlying parameters. The standard deviation γ of the implied distribution of the sampling error, η , is 0.34. The main feature of the distribution of the estimated markup ratio that supports this finding is that 31 percent of the estimates are below one, which can only arise from the left tail of the distribution of the sampling error. The mean, δ , of the distribution of $\log \nu$, the measure of true $\mu - 1$, is -1.43, indicating that the distribution is bunched near zero—that is, the distribution of μ places high probability on values not far above the competitive value of one. But the standard deviation of $\log \nu$, labeled σ , is 0.70, so the distribution has an extended upper tail. The bottom panel gives the implied mean and standard deviation of the level of the true markup μ . The mean is one plus the mean of the distribution of the estimated $\mu - 1$, as shown in the top line of the table, 1.31, and the standard deviation is 0.24. Thus the distribution of the true markup is fairly tightly contained in the range between 1.0 and 1.5.

Figure 2 plots the inferred distribution of the true value of the markup ratio μ and the inferred distribution of the measured value, which is the convolution of the distributions of the true value and the sampling error. All of the distribution of the measured value below one is the result of the sampling error and a fair amount of the distribution above 1.5. Figure 3 compares the calculated cumulative distribution to the cumulative distribution of the 60

	1	0.31
Moments of estimated $\mu - 1$	2	0.27
	3	0.23
Standard deviation of sampling error	γ	0.34
Mean of log of (true markup ratio - 1)	δ	-1.43
Standard deviation of log of (true markup ratio - 1)	σ	0.70
Mean of true markup ratio		1.31
Standard deviation of true markup ratio		0.24

Table 3: Moments of the Distribution of the Estimated Ratio of Price to Marginal Cost, and Inferred Properties of the Distribution of the True Ratio

estimates. The fit is pretty good. The model smooths the distribution without seriously altering its general features.

4.3 The change in the markup coefficient over time

To study the widely discussed hypothesis of growth in market power, I extend the specification to include an industry-specific linear time trend over the sample period from 1988 through 2015:

$$\sum_i \alpha_{i,t} \Delta x_{i,t} = (\phi_i - \psi_i t) \Delta \log y_t - a_t. \quad (19)$$

Here t advances by one each year and crosses zero in the middle of the sample period, 2001. The implied functional form for the ratio of price to marginal cost is

$$\mu_{i,t} = \frac{1}{\phi_i - \psi_i t}. \quad (20)$$

I extend the set of instruments to include the product of the log-changes and the time-trend variable, so there are 10 instruments.

Table 4 shows the growth coefficients ψ_i in the same sector groupings as in Table 2 earlier. They are sorted from highest to lowest. Though the ranking is plausible—for example, the information sector has relatively rapid growth in market power—there is substantial sampling error. There is no neat way to separate sampling variation from heterogeneity in

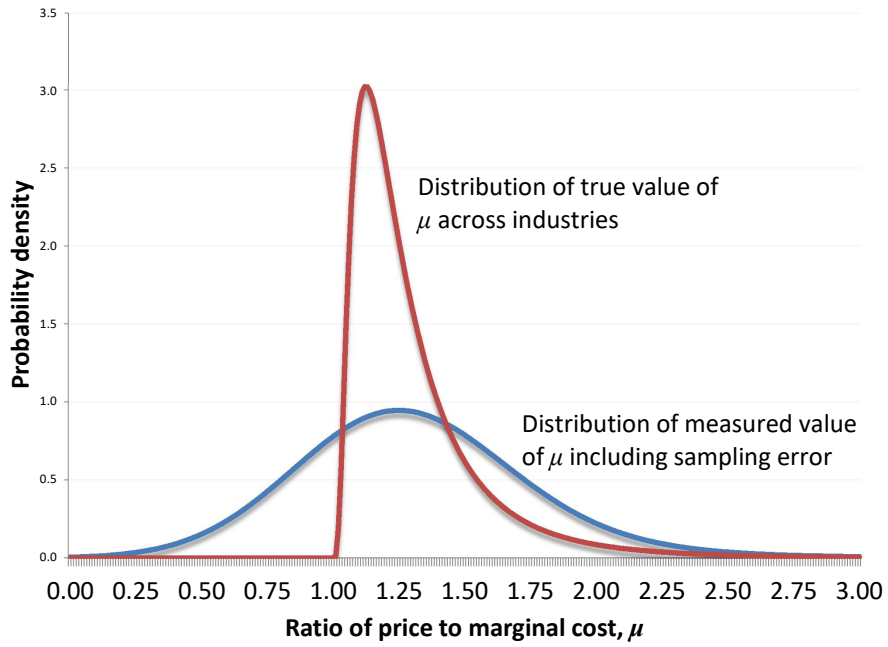


Figure 2: Inferred Distributions of True Ratios of Price to Marginal Cost across Industries, and Distribution Including Sampling Errors

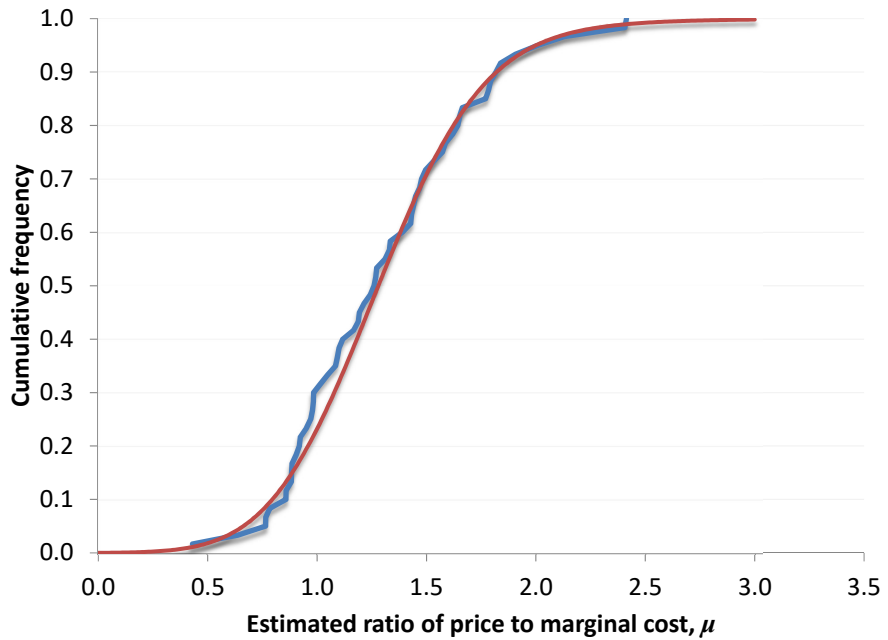


Figure 3: Actual Cumulative Frequencies of Estimates and Calculated Cumulative Distribution Functions from the Statistical Model




Figure 4: Implied Values of μ by Year

the true coefficients. There are indications that sampling variation dominates the observed heterogeneity of the estimates. For example, if there were substantial variation in the true values, the t -statistics would have more dispersion across industries than the standard t distribution has. In fact, the average of the squared t -statistics is 0.97, whereas it would be 1.13 for a t distribution with 15 degrees of freedom.

Despite the preponderance of sampling variation, the estimates give reasonable support to the hypothesis that the overall price/marginal-cost ratio rose over the period from 1998 through 2015. Table 5 shows the weighted average of the 60 estimates of ψ , which is 0.0061 decrease in μ^{-1} per year or a 0.0078 increase in μ per year, as of 2001. The weights are the shares of the industries in total value-added. The t -statistic for the hypothesis that ψ is actually zero, and that sampling error accounts for the increase, is 1.20. The p -value for the one-tailed test is 0.11, which is reasonably strong evidence against the null hypothesis.

Figure 4 plots the growth of the markup of price to marginal cost at the weighted averages of the parameters ϕ_i and ψ_i . The markup ratio grew from 1.12 in 1988 to 1.38 in 2015. This finding indicates substantial growth in market power, though rather less than the economy-wide increase reported by De Loecker and Eeckhout (2017).

<i>Weighted averages across industries</i>		
<i>Growth coefficient, ψ</i>	<i>Standard error</i>	<i>Sector name</i>
0.064	(0.035)	Finance and Insurance
0.036	(0.109)	Utilities
0.019	(0.007)	Health Care and Social Assistance
0.018	(0.010)	Real Estate and Rental and Leasing
0.017	(0.016)	Information
0.017	(0.016)	Construction
0.015	(0.024)	Management of Companies and Enterprises
0.013	(0.014)	Arts, Entertainment, and Recreation
0.006	(0.007)	Administrative and Support and Waste Management and Remediation Services
0.006	(0.015)	Other Services (except Public Administration)
0.004	(0.028)	Agriculture, Forestry, Fishing and Hunting
0.001	(0.008)	Accommodation and Food Services
0.001	(0.007)	Manufacturing
0.001	(0.009)	Transportation and Warehousing
-0.001	(0.010)	Educational Services
-0.010	(0.011)	Professional, Scientific, and Technical Services
-0.021	(0.011)	Wholesale Trade
-0.021	(0.011)	Retail Trade
-0.111	(0.061)	Mining, Quarrying, and Oil and Gas Extraction

Table 4: Estimates of the Growth in the Ratio of Price to Marginal Cost by Industry, Stated as Sector Averages

Weighted average of estimate of trend ψ	0.0061
Standard error	0.0051
t -statistic for hypothesis $\psi = 0$	1.20
p -value, one-tailed	0.11

Table 5: Evidence about the Statistical Reliability of the Finding of an Upward Trend in the Markup Ratio

5 Mega-Firms and the Price/Marginal Cost Ratio

Autor, Dorn, Katz, Patterson and Van Reenen (2017) describe superstar firms that grow to account for large fractions of sales in their output and input markets and presumably high ratios of price to marginal cost. The rise of these firms is thought to account for the decline in the labor share in the US and many other countries in recent decades. Here I investigate the relation between the employment share of high-employment mega-firms and the price-marginal cost ratios found using the method of this paper. I study the level and growth of the shares of mega-firms using the data described earlier. Table 6 shows the employment shares of firms with 10,000 or more workers in the 19 NAICS sectors in the first year the data are available, 1998, and the most recent year, 2015.

Growth in the shares of mega-firms has been anything but uniform across sectors. In four of the 19 sectors, very high-employment firms declined in importance over the 17-year span of the data. These sectors include all of manufacturing, which is the third-largest sector. The largest positive growth was in utilities, where mega-firms rose from 33.5 percent of employment in 1998 to 46.0 percent in 2015. Retail trade was another sector with a large increase in concentration by this metric. The weighted-average increase across all sectors was only 1.8 percentage points, from 25.3 percent to 27.1 percent. Thus it seems unlikely that rising concentration played much of a role in the general increase in market power that probably occurred over the 17 years. Figure 5 shows the movements of the high-firm-employment share by sector in the intervening years.

Figure 6 shows that there is essentially no systematic relation between the mega-firm employment ratio, on the horizontal axis, and the ratio of price to marginal cost found earlier in the paper, measured as the parameter ϕ . Over the wide range of variation in the employment ratio, sectors with low market power and with high market power are found, with essentially the same average values. There is no cross-sectional support for the hypothesis of higher markup ratios in sectors with more very large firms and thus more concentration in the product markets contained in those sectors. In Section 2.2, I observed that monopsony power in the labor market results in a downward bias in the estimated markup ratio. Thus the finding of no relation between labor-market concentration and the markup ratio could reflect the canceling of the offsetting upward effect on true markups and the downward bias from monopsony power.

<i>NAICS</i>	<i>Description</i>	<i>Employment, 2015, millions</i>	<i>Megafirm ratio in 1998</i>	<i>Megafirm ratio in 2015</i>	<i>Change</i>
11	Agriculture, Forestry, Fishing and Hunting	0.2	0.045	0.019	-0.026
21	Mining, Quarrying, and Oil and Gas Extraction	0.7	0.208	0.188	-0.020
22	Utilities	0.6	0.335	0.460	0.125
23	Construction	6.0	0.027	0.039	0.012
31-33	Manufacturing	11.6	0.271	0.248	-0.023
42	Wholesale Trade	6.1	0.156	0.179	0.023
44-45	Retail Trade	15.7	0.416	0.482	0.066
48-49	Transportation and Warehousing	4.6	0.369	0.408	0.039
51	Information	3.4	0.491	0.499	0.008
52	Finance and Insurance	6.1	0.418	0.435	0.016
53	Real Estate and Rental and Leasing	2.1	0.132	0.143	0.011
54	Professional, Scientific, and Technical Services	8.8	0.161	0.196	0.034
55	Management of Companies and Enterprises	3.3	0.542	0.509	-0.033
56	Administrative and Support and Waste Management and Remediation Services	11.1	0.296	0.325	0.030
61	Educational Services	3.6	0.141	0.161	0.020
62	Health Care and Social Assistance	19.2	0.190	0.200	0.010
71	Arts, Entertainment, and Recreation	2.2	0.118	0.124	0.006
72	Accommodation and Food Services	13.2	0.210	0.214	0.004
81	Other Services (except Public Administration)	5.4	0.052	0.051	0.000
Weighted average			0.253	0.271	0.018

Table 6: Ratio of Employment in Mega-Firms to Total Sectoral Employment, 1998 and 2015




Figure 5: Ratio of Employment in Mega-Firms to Total Sectoral Employment, 1998 through 2015

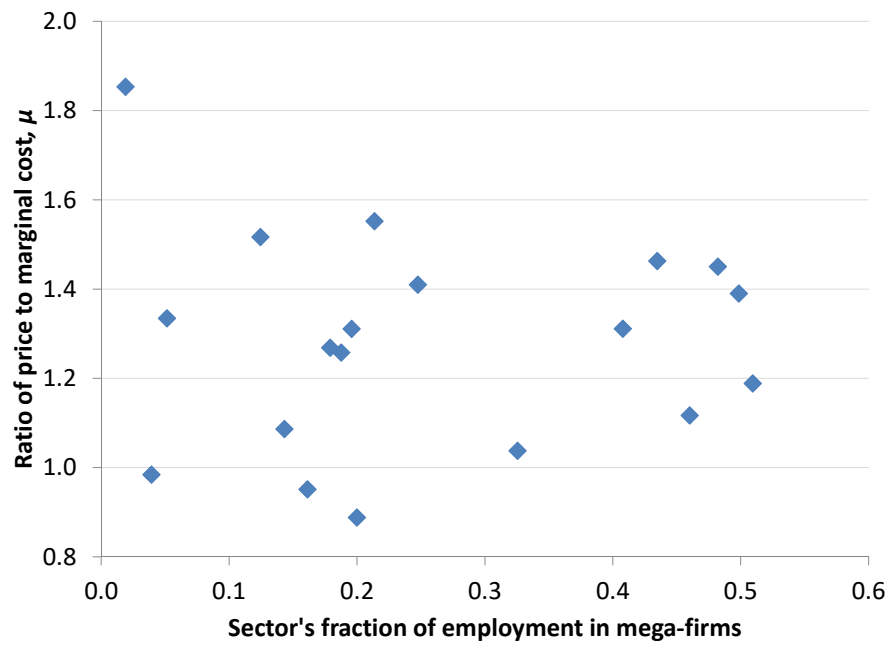


Figure 6: Relation between Employment in Mega-Firms and the Ratio of Price to Marginal Cost, μ

<i>Left-hand variable</i>	<i>Right-hand variable</i>	<i>Slope, standard error, and 1-tail p value</i>
Estimated markup, μ	Level of mega-firm ratio	-0.05 (0.36) 0.55
Estimated markup trend coefficient, ψ	Change in mega-firm ratio	0.045 (0.049) 0.18

Table 7: Slope Coefficients for the Relation between Employment in Mega-Firms and the Trend Coefficient for the Markup Ratio

The upper row of Table 7 describes the relation between the level of the mega-firm employment ratio and the estimated markup ratio. It confirms the lack of a relationship. But the standard error of the coefficient is high enough that the results do not refute the hypothesis of a meaningful positive relationship.

Figure 7 shows the relation between the change in the mega-firm employment ratio, on the horizontal axis, and the trend coefficient ψ in the price/marginal cost ratio. The plot suggests some upward-sloping relation. The point at the lower left, for mining, quarrying, and oil and gas extraction, is both influential and suspicious. Extraction industries present a challenge to the measurement of productivity and its cousin, the markup ratio measured in this paper. The other influential observation supporting an upward slope, for utilities, is at the upper right. The lower row in Table 7 confirms that there is moderate evidence in favor of an upward slope. Thus the hypothesis that a movement toward higher labor-market concentration has a role in rising market power receives some support here.

6 Concluding Remarks

Direct measurement of market power using high-quality annual time-series productivity data for 60 industries yields good information of the heterogeneous incidence of positive market power in US industries. The mean of the markup ratio is 1.31 and the standard deviation of the ratio across industries is 0.24. The mean grew from 1.12 in 1988 to 1.38 in 2015, though sampling uncertainty is present in the growth rate. There is no evidence that concentration, measured by the fraction of workers in a sector employed at firms with 10,000 or more




Figure 7: Relation between the Change in Employment in Mega-Firms and the Trend Coefficient for the Markup Ratio, ψ

workers, is related to market power, but some evidence that growth of superstar mega-firms is associated with rising market power.

There is a good deal of noise in the calculations at the individual industry level. The noise is interpreted as the annual growth of Hicks-neutral technology. The paper tries to state the precision in its estimates using standard statistical tools. There is a good deal of cross-industry heterogeneity in the estimated parameters. All of the results are interpreted in a framework of heterogeneity.

The choice to use modern productivity data has advantages and disadvantages. The alternative is to use data from individual firms, such as Compustat for publicly traded firms or confidential survey or administrative data. The advantage of the productivity data is the care with which the BEA and BLS measure inputs and outputs. No data on individual firms comes close to the accuracy and detail of the productivity data. The advantage of the data on individual firms is much more variation in growth rates of inputs and outputs and thus lower sampling variation in the estimated coefficients.

References

- Autor, David, David Dorn, Lawrence F. Katz, Christina Patterson, and John Van Reenen, “The Fall of the Labor Share and the Rise of Superstar Firms,” Working Paper 23396, National Bureau of Economic Research May 2017.
- De Loecker, Jan and Frederic Warzynski, “Markups and Firm-Level Export Status,” *American Economic Review*, May 2012, *102* (6), 2437–71.
- and Jan Eeckhout, “The Rise of Market Power and the Macroeconomic Implications,” Working Paper 23687, National Bureau of Economic Research, August 2017.
- Gutiérrez, Germán and Thomas Philippon, “Comments on ‘The Rise of Market Power and the Macroeconomic Implications by De Loecker and Eeckhout’,” October 2017. New York University.
- Hall, Robert E., “The Relation between Price and Marginal Cost in U.S. Industry,” *Journal of Political Economy*, 1988, *96* (5), 921–947.
- Solow, Robert M., “Technical Change and the Aggregate Production Function,” *The Review of Economics and Statistics*, 1957, *39* (3), 312–320.
- Traina, James, “Is Aggregate Market Power Increasing? Production Trends Using Financial Statements,” New Working Paper Series No. 17, Stigler Center for the Study of the Economy and the State, University of Chicago Booth School of Business, February 2018.