

Using Empirical Marginal Cost to Measure Market Power in the US Economy*

Robert E. Hall

Hoover Institution and Department of Economics
Stanford University

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

November 6, 2018

Abstract

Market power arises in the case where a seller is aware that raising output will depress price. In the profit-maximizing equilibrium with market power, price exceeds marginal cost. The Lerner index—the ratio of price less marginal cost to the price—is a widely accepted measure of market power. Measuring marginal cost is a challenge. This paper develops and applies a direct empirical approach—marginal cost is measured as the ratio of the observed change in cost to the observed change in output. Because marginal cost is a partial derivative, both changes need to be adjusted for other sources of change. Thus marginal cost is the ratio of (1) the change in cost not associated with changes in input prices to (2) the change in output not associated with productivity change. I develop data for the 60 KLEMS industries for this measure. I find a typical Lerner index of 0.15. Lerner indexes grew moderately between 1988 and 2015.

JEL L1 D24

*This research was supported by the Hoover Institution. I am grateful to Susanto Basu, Dennis Carlton, Emmanuel Farhi, Oleg Itskhoki, Chad Syverson, James Traina, and Hal Varian for comments.

1 Introduction

The measurement of market power informs many branches of economics. This paper uses the Lerner index—the ratio of price minus marginal cost to price—as the measure of market power. 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 price-taking firm’s Lerner index is zero. A firm facing a constant-elastic residual demand, with elasticity ϵ , maximizes profit at the point where the Lerner index is $1/\epsilon$. In general, the Lerner index, designated \mathcal{L} 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 markup ratio, that is, the ratio of price to marginal cost, μ :

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

which maps the Lerner index from $\mathcal{L} \in [0, 1]$ to $\mu \in [1, \infty]$.

The literature on measurement of marginal cost has two main branches. The demand-side approach infers the residual elasticity ϵ , typically from a differentiated-products oligopoly model. The production-side approach uses data on price and cost from firms. Thirty years ago, Hall (1988) proposed a refinement of the production-side approach that measures marginal cost rather than average cost. De Loecker and Warzynski (2012) developed a quite different method that measures marginal cost from the production function. De Loecker and Eeckhout (2017) recently captured a great deal of attention with the finding that market power has risen substantially in the US in recent decades. See Traina (2018) and Gutiérrez and Philippon (2017) for critiques of that paper.

De Loecker and Eeckhout (2017) use data from individual firms, where the data on inputs are limited to the accounting measure of cost of goods sold. I use data on 60 fairly detailed industries, with measures of actual output (not value added) and five categories of inputs. Though the data have much higher quality than anything available for individual firms, the data are aggregated across heterogeneous firms. The results may suffer from aggregation bias. I examine this issue toward the end of the paper.

Since the publication of my 1988 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 \mathcal{L} . In addition, a number of recent papers on market power and profit, 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 market power. This paper responds to those developments. It finds support for the conclusion that market power has risen in recent decades, though by less than in De Loecker and Eeckhout's paper.

An older manifestation of the production-side approach studies profit margins in various ways. Some attempt to distinguish marginal from average cost, generally by measuring the margin by subtracting only costs that are deemed variable. Their accuracy depends entirely on the investigator's success in guessing how costs vary with output. The methods of this paper and of De Loecker and Eeckhout turn that job over to the data in systematic ways.

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

The text describes the many calculations underlying this paper in general terms. The calculations are fully documented in the computer files available from my website.

2 The Lerner Index of Market Power

2.1 Measuring marginal cost and the Lerner index

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

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

in obvious notation. The change in cost is

$$dc = \sum_i x_i dw_i + \sum_i w_i dx_i. \tag{3}$$

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:

$$\sum_i w_i dx_i. \tag{4}$$

The technology is

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

so output growth is

$$dy = A df(x) + f(x) dA = A df(x) + y \frac{dA}{A} \quad (6)$$

The desired component purged of effects from changing productivity is

$$A df(x) = dy - y \frac{dA}{A} \quad (7)$$

Marginal cost is the ratio of adjusted cost change to adjusted output change,

$$m = \frac{\sum_i w_i dx_i}{dy - y dA/A}. \quad (8)$$

The Lerner index is

$$\mathcal{L} = \frac{p - m}{p} = 1 - \frac{\sum_i w_i dx_i}{p(dy - y dA/A)}. \quad (9)$$

so

$$1 - \mathcal{L} = \frac{\sum_i w_i dx_i}{p(dy - y dA/A)}. \quad (10)$$

Now let

$$\alpha_i = \frac{w_i x_i}{p y}, \quad (11)$$

the share of factor i in revenue, $p y$. The equation can then be written

$$(1 - \mathcal{L}) \left(dy + y \frac{dA}{A} \right) = y \sum_i \alpha_i \frac{dx_i}{x_i}. \quad (12)$$

Dividing by y and rearranging yields a useful result,

$$\frac{dy}{y} - \sum_i \alpha_i \frac{dx_i}{x_i} = \mathcal{L} \frac{dy}{y} + (1 - \mathcal{L}) \frac{dA}{A}. \quad (13)$$

With discrete time, the same equation is

$$\Delta \log y - \sum_i \alpha_i \Delta \log x_i = \mathcal{L} \Delta \log y + (1 - \mathcal{L}) \Delta \log A. \quad (14)$$

This formulation is useful because the left-hand side is the Solow residual, calculated meticulously in productivity accounts. Note that if $\mathcal{L} > 0$, the Solow residual does not measure actual technical progress, because it does not adjust for market power.

This derivation of the measurement of \mathcal{L} 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 \mathcal{L} 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, $a = (1 - \mathcal{L})\Delta \log A$, is a statistical residual in equation (14). It can only be measured with knowledge of the Lerner index, \mathcal{L} . The most basic approach is to treat \mathcal{L} 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 changes. Shifts in supply and demand in the output market or factor markets could be eligible. The residual based on the estimated value of \mathcal{L} is the estimated rate of true productivity growth, adjusted for market power.

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

$$\Delta \log y_t - \sum_i \alpha_i \Delta \log x_i = (\phi + \psi t)\Delta \log y + a_t. \quad (15)$$

Here ϕ controls the level of the Lerner index and ψ is the per-period growth of the index. With panel data, the function multiplying $\Delta \log y$ may capture differences in market power across industries as well.

2.2 Interpretation

Here I consider whether the procedure described earlier measures the Lerner index accurately or measures it with a bias of known sign in the presence of decreasing and increasing returns to scale, market power of factor suppliers, and monopsony power in factor markets. This discussion introduces the assumption of optimization by firms, an assumption deliberately omitted from the earlier derivation of the empirical marginal cost measure.

Differentiation of the production function,

$$y = Af(x), \quad (16)$$

yields

$$\frac{dy}{y} = \frac{dA}{A} + \sum_i \frac{x_i}{f(x)} \frac{\partial f}{\partial x_i} \frac{dx_i}{x_i}. \quad (17)$$

Now assume that the firm is a price-taker in all of its input markets, and the firm equates the marginal revenue product of a factor to its price:

$$(1 - \mathcal{L})pA \frac{\partial f}{\partial x_i} = w_i. \quad (18)$$

Use this equation to substitute out the $\frac{\partial f}{\partial x_i}$ in the previous equation and rearrange to get

$$\frac{dy}{y} - \sum_i \alpha_i \frac{dx_i}{x_i} = \mathcal{L} \frac{dy}{y} + (1 - \mathcal{L}) \frac{dA}{A}, \quad (19)$$

as before. Under the new assumptions, \mathcal{L} is the Lerner index of market power. Notice that the assumptions do not include constant returns to scale. But the second-order condition for profit maximization requires that the Lerner index exceed $1 - 1/\gamma$, where γ is the returns-to-scale index of the production function, the elasticity of $f(\theta x)$ with respect to θ , at $\theta = 1$. A firm with strong increasing returns and weaker market power will not satisfy the second-order condition.

To make some further progress on these issues, consider the simple case with only one factor, labor, n , paid wage w . The production function is

$$y = n^\gamma. \quad (20)$$

The elasticity γ is positive but may lie in either side of 1. In changes,

$$\frac{dy}{y} = \gamma \frac{dn}{n}. \quad (21)$$

The Solow residual uses the revenue share,

$$\alpha = \frac{wn}{py} = (1 - \mathcal{L})\gamma, \quad (22)$$

capturing the well-known depressing effect of market power on the measured share of labor.

The Solow residual is

$$\frac{dy}{y} - \alpha \frac{dn}{n} = \frac{dy}{y} - (1 - \mathcal{L})\gamma \frac{dn}{n}. \quad (23)$$

From above, two of the terms on the right net to zero, so

$$\frac{dy}{y} - \alpha \frac{dn}{n} = \mathcal{L}\gamma \frac{dn}{n}, \quad (24)$$

and, as before,

$$\frac{dy}{y} - \alpha \frac{dn}{n} = \mathcal{L} \frac{dy}{y} \quad (25)$$

The assumption that the firm is a price taker in its input markets does not mean that those markets are competitive. That property is sufficient but not necessary. The price-taking assumption would apply if a labor union or dominant seller of another input chose to exercise its market power by sticking to a fixed non-negotiable price quote.

On the other hand, if a firm has monopsony power in an input market and perceives that increasing its purchase volume will drive up the price, a downward bias in the estimate of the firm's Lerner index will result. Suppose the elasticity of the wage with respect to the firm's level of employment is λ . Then the observed labor share is further depressed by the fact that the average wage understates the marginal wage:

$$\alpha = \frac{wn}{py} = (1 - \mathcal{L}) \frac{\gamma}{1 + \lambda}, \quad (26)$$

which propagates through the rest of the math to the conclusion,

$$\frac{dy}{y} - \alpha \frac{dn}{n} = \frac{\mathcal{L} - \lambda}{1 + \lambda} \frac{dy}{y}. \quad (27)$$

Thus the coefficient on the right side of the equation is $\frac{\mathcal{L} - \lambda}{1 + \lambda}$, which is less than \mathcal{L} for any positive value of the monopsony parameter λ .

An important case deviating from the assumptions stated earlier is an omitted variable in the productivity calculation. A leading example is the firm's stock of intangible capital of a type not included in its measured capital stock. Let x_o designate the omitted factor quantity and α_o be the elasticity of the production function with respect to the omitted factor. Also let $\tilde{\alpha}_k$ designate the true elasticity of capital, on the assumption that the measured elasticity α_k is overstated because its revenue share includes the firm's earnings from the omitted factor along with the earnings of capital. The growth of output becomes

$$\frac{dy}{y} = \sum_i \alpha_i \frac{dx_i}{x_i} + \tilde{\alpha}_k \frac{dx_k}{x_k} + \alpha_o \frac{dx_o}{x_o} + \mathcal{L} \frac{dy}{y} + (1 - \mathcal{L}) \frac{dA}{A}. \quad (28)$$

The estimating equation becomes

$$\frac{dy}{y} - \sum_{i \text{ not } o} \alpha_i \frac{dx_i}{x_i} = (\tilde{\alpha}_k - \alpha_k) \frac{dx_k}{x_k} + \alpha_o \frac{dx_o}{x_o} + \mathcal{L} \frac{dy}{y} + (1 - \mathcal{L}) \frac{dA}{A}. \quad (29)$$

The omitted variable is $(\tilde{\alpha}_k - \alpha_k) \frac{dx_k}{x_k} + \alpha_o \frac{dx_o}{x_o}$. If it is correlated with the instruments, the IV estimates will be biased. Under the hypothesis that an instrument is positively correlated with $\frac{dx_k}{x_k}$ and with $\frac{dx_o}{x_o}$, the two components will contribute offsetting effects to the bias, because, by assumption, $\tilde{\alpha}_k - \alpha_k < 0$.

2.3 Conclusions about the applicability of the approach

Increasing returns to scale. The approach is robust to increasing returns. A fixed cost of continuing operations is a commonly considered source of increasing returns. Increasing

returns must be coupled with market power. A leading example of a firm with a combination of increasing returns and market power is one satisfying the assumptions of the monopolistic competition model. In that model, entry of firms with differentiated products (giving them market power) occurs to the point of zero profit. In the simplest model, increasing returns takes the particular form of a fixed cost of operation, independent of scale but avoidable by shutting down. In equilibrium, the firm makes just enough excess profit from market power to offset the fixed cost.

Decreasing returns to scale. Decreasing returns occur when factors, notably capital, involve delays, adjustment costs, or permanent restrictions on inputs. The approach is also robust to decreasing returns, which will be accompanied by profit in excess of factor costs.

Omitted input. As noted earlier, if the effect of the instrument on the growth of the omitted factor has the same sign as its effect on capital growth, there are two biases in opposite directions. In the best case, the net effect would be zero.

Market power held by a seller of an input. 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 Lerner index 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 \mathcal{L} .

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 employment share is understated and the estimate of \mathcal{L} is correspondingly understated.

2.4 Data for measuring the Lerner index

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

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. Figure 1 displays the instrumental variables.

2.5 Results for the Lerner index

Some other investigations, including my earlier work, measure market power with μ , the ratio of price to marginal cost. I noted earlier that $\mu = 1/(1 - \mathcal{L})$. There is no problem going back and forth between the two measures. In fact, for values in the center of the distribution of estimates in this paper, the linear approximation $\mu \cong 1 + \mathcal{L}$ is pretty good. 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.




Figure 1: Instrumental Variables

The reported standard error for these cases will be high. The results in those cases make more sense where \mathcal{L} is estimated, and the high and uncertain values of μ are mapped into a small region around 1 for μ^{-1} or zero for \mathcal{L} . In those cases, the linear approximation implicit in a standard error is completely unreliable.

2.5.1 First-stage results

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 p values corresponding to the F -statistics of the first-stage regressions. In 59 percent of the KLEMS industries, the hypothesis of no relation between the instruments and output growth can be rejected at the 5-percent level or better. In the industries where the instruments are weak, sampling

<i>p-value range</i>	<i>Percent of KLEMS industries</i>
Greater than 0.50	8
0.20 to 0.50	10
0.20 to 0.05	23
0.05 to 0.01	32
0.01 to 0.001	20
Less than 0.001	7

Table 1: Metrics for the Power of the Instruments Across Industries

error of the estimates of the Lerner index are high; in a few industries, the estimates are uninformative.

2.5.2 Estimates of the Lerner index, \mathcal{L} , 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 estimated value of the Lerner index. The first column of the table presents averages across the industries contained in the sectors, weighted by the 1987 distribution of industry values of output. The second column reports bootstrap standard errors. The standard errors for the three sectors with inadmissible negative values of the Lerner index are sufficiently large to render the estimates meaningless. In these cases, the instruments lack the power to identify the markup ratio with usable accuracy. The table also reports the percent of the value of output arising from the sector and the number of KLEMS industries in the sector.

2.6 Inference about the sources of dispersion of the measured Lerner index across industries

The estimates of the Lerner index specific to the industries have a good deal of noise. In particular, 30 percent of the industries have negative values of L_i despite the fact that the true value of \mathcal{L} cannot be negative. To disentangle the distribution of the true values of the Lerner index across industries from the distribution of the sampling error, I consider a simple statistical model that exploits the fact that sampling error must have a role sufficient to explain the 30 percent of values of the ratio that are estimated to be negative. The

<i>Weighted averages across industries</i>				
<i>Lerner index</i>	<i>Bootstrap standard error</i>	<i>Percent of value of output in sector</i>	<i>Number of industries in sector</i>	<i>Sector name</i>
-0.11	(0.32)	5.3	3	Health Care and Social Assistance
-0.05	(0.79)	0.5	1	Educational Services
-0.02	(0.45)	6.7	1	Construction
0.03	(0.21)	3.3	2	Administrative and Support and Waste Management and Remediation Services
0.05	(0.05)	1.3	3	Mining, Quarrying, and Oil and Gas Extraction
0.07	(0.10)	5.4	4	Information
0.08	(0.10)	6.1	2	Real Estate and Rental and Leasing
0.10	(0.07)	3.0	1	Utilities
0.16	(0.10)	1.9	1	Management of Companies and Enterprises
0.18	(0.19)	27.5	18	Manufacturing
0.21	(0.10)	5.4	1	Wholesale Trade
0.21	(0.11)	4.0	8	Transportation and Warehousing
0.23	(0.22)	7.2	3	Professional, Scientific, and Technical Services
0.24	(0.06)	8.0	4	Finance and Insurance
0.25	(0.08)	2.4	1	Other Services (except Public Administration)
0.29	(0.07)	0.8	2	Arts, Entertainment, and Recreation
0.31	(0.13)	6.3	1	Retail Trade
0.34	(0.08)	3.0	2	Accommodation and Food Services
0.46	(0.14)	1.7	2	Agriculture, Forestry, Fishing and Hunting

Table 2: Estimates of the Lerner Index by Industry, Stated as Sector Averages




Figure 2: The Family of Beta Distributions with Second Shape Parameter = 8

statistical model is

$$L = \mathcal{L} + \eta, \tag{30}$$

where \mathcal{L} is distributed as $\text{beta}(\nu, \beta)$, with density proportional to $\mathcal{L}^{\nu-1}(1 - \mathcal{L})^{\beta-1}$. The measurement error η_i accounts for the residual distribution of the measured index.

Four assumptions identify the model:

1. The true value of the Lerner index obeys the beta distribution, so it is between zero and one: $\mathcal{L} \in [0, 1]$
2. The second shape parameter of the beta distribution of the true Lerner index is $\beta = 8$
3. The two components are statistically independent
4. The mean of the measurement error η is zero

Independence and zero mean of measurement error are standard assumptions in models of this type. The restriction $\beta = 8$ defines a reasonable family indexed by the first shape parameter, ν . Figure 2 shows several members of the family.

The following result establishes the principle that the desired untangling is possible:

Identification Theorem: The mean of the measured Lerner index identifies the first shape parameter of the beta distribution of the true Lerner index; the distribution of the measurement error η is identified by solving a convolution problem.

Proof: By assumption 4, the mean of the estimated Lerner index is the mean of the true index. That mean is the mean of the beta distribution,

$$M = \frac{\nu}{\nu + \beta}. \quad (31)$$

Thus

$$\nu = \frac{\beta M}{1 - M}. \quad (32)$$

The distribution of the estimates of the Lerner index is the convolution of the distributions of \mathcal{L} and η . By the convolution theorem—see Cramér and Wold (1936)—the characteristic function of the convolution of two random variables is the product of the two characteristic functions. Thus the characteristic function of η is the ratio of the characteristic function of the observed random variable L to the characteristic function of the random variable \mathcal{L} . Distributions are one-to-one with characteristic functions, so the distribution of η is identified. \square

Although manipulating characteristic functions might seem to be the natural way to calculate the distribution of η , it appears to be unworkable in this application, so I proceeded by a direct solution—representing the distribution in a flexible parametric form and solving the convolution by minimizing the distance between the actual distribution of the estimates L and the distribution calculated as the convolution of the parametric distribution and the known distribution of the true Lerner index \mathcal{L} . I use a discrete 32-point support η_j so the probabilities π_j serve as the parameters, subject to the natural restrictions $\sum_j \pi_j = 1$ and $\pi_j \geq 0$ and small penalties encouraging smoothness of the π_j s. The convolution defining the calculated cdf is

$$\sum_j \pi_j \int_0^{\text{plus}(L_i - \eta_j)} f(\mathcal{L}) d\mathcal{L}. \quad (33)$$

Here $\text{plus}(\cdot)$ is the positive part function. With the estimated values sorted in increasing order, the observed value of the cdf at L_i is $i/60$. Let $D_i(\pi)$ be the discrepancy between the calculated and actual cdfs.

Moments of estimated Lerner indexes across industries	Mean	0.15
	Stan. dev.	0.31
	Skewness	-1.84
Shape parameter of true Lerner index	α	1.36
Moments of true Lerner indexes across industries	Mean	0.15
	Stan. dev.	0.11
	Skewness	1.14
Moments of measurement errors	Mean	0.00
	Stan. dev.	0.29
	Skewness	-2.30

Table 3: Moments of the Distribution of the Estimated Lerner Index, and Inferred Properties of the Distributions of the True Index and the Error in Measurement

The problem becomes

$$\min_{\pi_j} \left[D_i(\pi)^2 + \left(1 - \sum_j \pi_j \right)^2 + \sum_j \text{minus}(\pi_j)^2 + \omega \sum_2^j (\pi_j - \pi_{j-1})^2 \right]. \quad (34)$$

Here $\text{minus}(\cdot)$ is the negative part function. The weight ω on the smoothness term is taken as 0.01, so the much higher weight on the matching conditions results in values of π_j that come close to satisfying the matching conditions.

Table 3 shows the inputs to and results of these calculations based on the distribution of estimates of the estimated Lerner index L . The upper panel shows the moments of the 60 estimates of L . The next lower panel reports the value of the first shape parameter of the beta distribution assumed to describe the distribution of the true, $\nu = 1.36$. The third panel down gives the implied mean and standard deviation of the level of the true Lerner index, \mathcal{L} . The mean is 0.15 and the standard deviation is 0.11. The distribution of the true Lerner index is fairly tightly contained in the range between 0 and 0.4. As Figure 2 indicates, the distribution is skewed to the right, with a skewness coefficient of 1.14. The bottom panel gives moments of the solved distribution of the measurement errors. The mean is zero by assumption. The standard deviation of the implied distribution of the sampling error, η , is 0.29. The main feature of the distribution of the estimated Lerner index that supports this finding is that 31 percent of the estimates are negative, which can only arise from the left tail of the distribution of the sampling error.




Figure 3: Inferred Distributions of True Lerner Index across Industries

Figure 3 plots the inferred distribution of the true value of the Lerner index \mathcal{L} . The density of the beta distribution is

$$Bx^{\nu-1}(1-x)^{\beta-1}. \quad (35)$$

Figure 4 is a bar chart summarizing the inferred distribution of the measurement error η . The actual distribution over the 32 points of its support is quite jagged, reflecting the randomness in the 60 draws of the underlying data on the observed Lerner index L . Because the assumed distribution of \mathcal{L} is smooth, all of the randomness in the observed index maps into corresponding randomness in the inferred measurement error.

All of the distribution of the measured value L below zero is the result of the sampling error, and a fair amount of the distribution above 0.5. Figure 5 compares the calculated cumulative distribution—the convolution of the distributions of \mathcal{L} and η —to the cumulative distribution of the 60 estimates. The fit is pretty good. Its imperfections arise entirely from the assumption that the distribution of η has the assumed 16 discrete points in its support, rather than being a continuous distribution.




Figure 4: Distribution of the Measurement Error




Figure 5: Actual Cumulative Frequencies of Estimates and Calculated Cumulative Distribution Functions from the Statistical Model

2.7 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:

$$\Delta \log y_t - \sum_i \alpha_{i,t} \Delta \log x_{i,t} = (\phi_i + \psi_i t) \Delta \log y_t - a_t. \quad (36)$$

Here t advances by one each year and crosses zero in the middle of the sample period, 2001. The implied functional form for the Lerner index is

$$\mathcal{L}_{i,t} = \phi_i + \psi_i t. \quad (37)$$

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 lowest to highest. 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 the true coefficients.

Despite the 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 increase in \mathcal{L} per year. The weights are the shares of the industries in total value of output. 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 6 plots the growth of the Lerner index at the weighted averages of the parameters ϕ_i and ψ_i . The index grew from 0.11 in 1988 to 0.28 in 2015. This finding indicates substantial growth in market power, though less than the economy-wide increase reported by De Loecker and Eeckhout (2017). The figure also shows lines one standard error above and below the estimated Lerner index for the year.

<i>Weighted averages across industries</i>		
<i>Growth coefficient, ψ</i>	<i>Standard error</i>	<i>Sector name</i>
-0.122	(0.074)	Mining, Quarrying, and Oil and Gas Extraction
-0.021	(0.009)	Retail Trade
-0.021	(0.012)	Wholesale Trade
-0.009	(0.009)	Professional, Scientific, and Technical Services
-0.001	(0.009)	Educational Services
0.000	(0.008)	Accommodation and Food Services
0.001	(0.008)	Manufacturing
0.001	(0.009)	Transportation and Warehousing
0.006	(0.010)	Other Services (except Public Administration)
0.006	(0.007)	Administrative and Support and Waste Management and Remediation Services
0.007	(0.025)	Agriculture, Forestry, Fishing and Hunting
0.013	(0.012)	Arts, Entertainment, and Recreation
0.015	(0.024)	Management of Companies and Enterprises
0.016	(0.006)	Health Care and Social Assistance
0.017	(0.010)	Construction
0.018	(0.013)	Information
0.019	(0.008)	Real Estate and Rental and Leasing
0.036	(0.073)	Utilities
0.057	(0.024)	Finance and Insurance

Table 4: Estimates of the Growth in the Lerner Index 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




Figure 6: Implied Values of the Aggregate Lerner index by Year, with Bootstrapped Error Bands

3 Aggregation Issues: Are There Biases from Using Averages across Heterogeneous Firms?

In his discussion of this paper at the NBER conference, “The Rise of the Megafirm: Causes and Consequences for Labor and Product Markets”, October 19, 2018, Emmanuel Farhi suggested that applying the method of this paper to industry aggregates might encounter problems of aggregation. He cited results on aggregation from Baqaee and Farhi (2017). He observed that a failure of the assumption of orthogonality of the disturbance (characterized as productivity growth) and the instruments could occur in aggregates even though orthogonality held at the firm level. The failure would result from a component of measured productivity growth induced by the instruments arising from heterogeneity in markups within the units included in the aggregate. This problem would be less serious if the data came from individual firms, though still present with aggregation across the products of the individual firms. However, the data available from individual firms is vastly inferior to the KLEMS data in many other respects.

The potential aggregation bias would also arise in data aggregated from KLEMS industries, so one way to get at the magnitude of the bias is to compare results from aggregates of the industry-level results discussed earlier. Aggregation of the KLEMS data involves removal of transactions of firms within the aggregate. The agencies create 10 aggregates on this basis. Table 6 compares the results of (1) applying the method of this paper to those aggregates to (2) the weighted averages of the underlying KLEMS industries. It uses the values of the estimated coefficients ϕ_i for the industries and similar coefficients for the aggregates. These estimates are interpreted as the Lerner Index in the middle year of the sample. Note that there are two levels of aggregates for manufacturing, so the upper level, total manufacturing, is not independent of the two lower levels, non-durables and durables.

The results in Table 6 do not suggest any important overall bias from the further aggregation—5 of the aggregates have higher estimates and 5 have lower estimates. And there is almost no evidence of differences for individual aggregates from the aggregation of the underlying industries—among the 8 aggregates with reasonably small standard errors, only total manufacturing has a discrepancy of more than one standard error of the aggregate.

<i>Aggregation</i>	<i>Number of KLEMS industries</i>	<i>Estimate of φ based on aggregate data</i>	<i>Standard error of estimate</i>	<i>Weighted mean of estimated industry φs</i>
Agriculture, Forestry, and Fishery	2	0.97	(0.44)	1.05
Mining	3	0.22	(0.64)	1.63
Manufacturing	18	0.38	(0.10)	0.25
Non-Durable Manufacturing	8	0.14	(0.16)	0.12
Durable Manufacturing	10	0.42	(0.09)	0.34
Trade	2	0.29	(0.07)	0.31
Transportation and Warehousing	8	0.29	(0.09)	0.29
Information	4	0.24	(0.14)	0.27
Finance, Insurance, and Real Estate	6	-0.16	(0.16)	-0.19
Services	15	0.13	(0.07)	0.15

Table 6: Comparisons of Estimates of the Lerner Index for Aggregates to Averages across Detailed Industries

4 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. 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.

All of the results in this paper support the hypothesis that the sellers in many industries in the US economy have substantial market power. According to Figure 3, the distribution of the Lerner index across industries includes many with Lerner indexes above 0.3.

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

- Baqae, David Rezza and Emmanuel Farhi, “Productivity and Misallocation in General Equilibrium.,” Technical Report, National Bureau of Economic Research 2017.
- Cramér, Harald and Herman Wold, “Some Theorems on Distribution Functions,” *Journal of the London Mathematical Society*, 1936, 1 (4), 290–294.
- 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.
- 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.