Is Aggregate Market Power Increasing?
Production Trends Using Financial Statements

James Traina
Booth School of Business, University of Chicago

Is Aggregate Market Power Increasing?
Production Trends using Financial Statements

James Traina[*]

February 2018

Abstract

Recent work in macroeconomics argues that firm market power dramatically increased since the 1980s. Using financial statement data, I find that public firm markups increased only modestly over this time period, and are within historical variation. These estimates improve on earlier work by accounting for marketing and management expenses, which I document are a rising share of costs in firm production. Markups are increasing in firm size and vary by sector. Reasonable calibrations accounting for the representativeness of public firms show a flat or even decreasing aggregate markup.

Recent work in macroeconomics argues that firm market power dramatically increased since the 1980s. For example, Barkai (2016) argues that only an increase in markups can generate a simultaneous decline in the shares of both labor and capital. De Loecker and Eeckhout (2017) claim that "in 1980, average markups start to rise from 18% above marginal cost to 67% now."[1] These papers are compelling because they simultaneously address two key secular trends: (1) the steady rise in industrial concentration, and (2) the steady fall in the labor share. They also serve as preliminary explanations of the potential decline in the capital share, output growth, and business and labor market dynamism.

This paper challenges these explanations by providing new key facts on aggregate market power. My first contribution is to show that public-firm market power has not substantially increased in recent decades. Using financial statement data from public firm filings and the latest tools in production-based markup estimation, I calculate markups for the universe of non-utility, non-financial US public firms. Aggregating these markup estimates to annual

[1]This article was quickly picked up by a number of media outlets, including The Economist, The Wall Street Journal, the NY Times, Financial Times, Bloomberg, Reuters, Quartz, Harvard Business Review, Pro Market, Noahpinion, Growth Economics, Marginal Revolution, The Weeds, and Vox Podcast.

estimates, I find that public firm markups increased only modestly since the 1980s. Moreover, this increase is within historical variation – measured markups have increased from 1980 - 2010 as much as they have decreased from 1950 - 1980.

Delving deeper into the underlying causes of disagreement with prior work, I uncover and highlight new trends in firm production. My second contribution is to show that firms have increasingly devoted more of their inputs toward marketing and management costs. As a share of variable costs for firms, these components have increased from roughly 12% in 1950 to 22% today. This shift in production is consistent with the consequences of information and communication technology improvements in the broader economy.

These markup estimates rely on public firm data, so a remaining question is whether they are representative of aggregate markups in the real economy. My last contribution is to show that selection biases in extrapolating public firm markups to the aggregate economy tend to substantially overestimate aggregate market power. I begin this analysis by noting that measured markups are positively correlated with size. Davis et al. (2006) document that public firms make up only about 1/3 of US sales and employment. Since these public firms are often larger than their private firm counterparts, markup estimates using only public firm data would bias an aggregate estimate upwards. I also document considerable dispersion of markups by sector. This variation could further bias an extrapolated public markup measure insofar as certain sectors are overrepresented (Manufacturing) or underrepresented (Construction) in Compustat. Moreover, because of trends in listings and delistings, as well as mergers and acquisitions, the representativeness of public firms is markedly changing over time. Importantly, these sample changes are not just about the number of public firms, but in fact affect the distribution of public firm characteristics[2]. As a first step to address these issues, I use the economy-wide size and industry distributions from the Census's Business Dynamics Statistics to reweight firms toward an aggregate estimate of market power. Reasonable calibrations accounting for the representativeness of public firms show a flat or even decreasing aggregate markup.

The closest paper to this one is De Loecker and Eeckhout (2017), which substantially opens the research agenda of exploring how important market power is in the aggregate economy. While they use similar data sources and methods, a key difference is this paper uses a better accounting measure of variable cost, which includes important components of costs omitted by earlier work. Specifically, this measure includes indirect costs of production such as marketing and management, which are an increasingly vital share of variable costs for firms. Neglecting these costs meaningfully overstates both the level and growth in markups. A significant contribution of this paper is to inform the debate on aggregate market power by offering a starkly different conclusion from the main empirical result of De Loecker and Eeckhout (2017).

---

[2]See Fama and French (2004), Harford (2005), Davis et al. (2006), Brown and Kapadia (2007), or more recently Doidge et al. (2017) for a discussion of these trends.

This paper is also related to Dorn et al. (2017), who propose an alternative mechanism by which both industrial concentration rises and labor shares fall. If globalization or innovation differentially benefit large firms, economies of scale will increase for these firms. Since large firms have a higher market share and lower labor share, this effect leads to both macroeconomic secular phenomena. The subsequent implications for markups depend on the nature of competition: in Cournot competition, markups are positively correlated with market share, so aggregate markups would increase; in monopolistic competition, markups are equal within industries. My findings are consistent with the latter view.

These results have important economic and policy implications. For example, if market power is a substantial feature of the aggregate economy, we may need to re-evaluate implications of models that don't explicitly account for this economic phenomena through, say, allowing imperfect competition in the firm sector. Moreover, if rising concentration and falling labor shares are the result of increases in market power, calls for increased anti-trust enforcement could be first-order in economic policy reform, unleashing more jobs and higher growth. However, this paper's evidence suggests that these secular trends are a response to underlying technological change, supporting light-touch regulatory approaches and a focus of future research on how specific mechanisms might transform the organizational boundaries of firms to larger structures.

1   Data, Sample Selection, and Variable Construction

The primary dataset for this paper is the Fundamental Annual Compustat file from Wharton Research Data Services. These data span from 1950 to 2016 and cover private sector firms with public equity or debt. The Compustat data contain firm-level balance sheet information, which allows me to use the production approach to measure markups. In particular, it provides information on sales, variable input expenditures, capital, and industry classification.

I also download price deflators to convert nominal variables to real variables. I use the NIPA Table 1.1.9. GDP deflator (line 1) and nonresidential fixed investment good deflator (line 9)[3].

To select domestic firms, I use standard industry format observations in USD with Foreign Incorporation Codes (FIC) in the USA. I exclude utilities (SIC codes between 4900 - 4999) because they are heavily regulated on prices. I also exclude financials (SIC codes between 6000 - 6999) because their balance sheets are dramatically different from other firms. For data quality, I exclude observations with negative or missing assets, sales, cost of goods sold, operating expenses, or gross plants, property, and equipment (PPE). To avoid picking up merger and acquisition distortions, I also exclude observations in which acquisitions are

---

[3]See https://www.bea.gov/national/nipaweb/DownSS2.asp

larger than 5% of the value of total assets.[4]

Just under 9% of sample observations are missing observations of sales, cost of goods sold, operating expenses, gross PPE, or net PPE within a given firm. I replace these missing observations with a linear interpolation of their neighboring values.[5] To get a real measure of sales and variable inputs, I deflate sales, cost of goods sold, and operating expenses by the GDP deflator.

As is standard in the production estimation literature, I construct the measure of capital using the perpetual inventory method. Specifically, I initialize the capital stock using the first available entry of gross PPE. I then iterate forward on capital using the accumulation equation $k_{it} = k_{it-1} + i_{it} - \delta k_{it-1}$, where I compute net investment using changes to net PPE. Since I want a measure of the real capital stock, I deflate net investment by the investment goods deflator.

2   Markup Estimation

I rely on the proposed framework of De Loecker and Eeckhout (2017), originated by De Loecker and Warzynski (2012) based on work dating to Hall (1988), to estimate firm-level markups using firm financial statements. This method relies on cost minimization of a variable input of production. For all subsequent analysis, I use operating expenses as a direct measure of variable inputs.[6] Operating expenses include materials, labor, marketing, and management. I defer discussion of this measure and a leading alternative measure to the next section.

De Loecker and Eeckhout (2017) derive a simple expression for the markup:

$$\mu_{it} = \theta_{it}^V \frac{P_{it}^Q Q_{it}}{P_{it}^V V_{it}}$$

where θ_{it}^V is the output elasticity of a variable input, $P_{it}^Q Q_{it}$ is output (sales) and $P_{it}^V V_{it}$ is a variable input (operating expenses). Sales and the variable input are directly measurable, but the output elasticity requires estimation. Note that this markup equation holds for any variable input. However, it still requires an unbiased estimate of that variable input's output elasticity.

I estimate industry-specific Cobb-Douglas production functions, with variable inputs and capital:

$$q_{it} = \beta_v v_{it} + \beta_k k_{it-1} + \omega_{it} + \varepsilon_{it}$$

---

[4]My results are robust to including or excluding any combination of these data selection filters.

[5]My results do not depend on including these interpolated values.

[6]Unfortunately, public firm financial statements neither commonly nor consistently differentiate between labor and material inputs.

with q_{it} measuring log real sales, v_{it} log variable input, k_{it-1} log capital stock, and ω_{it} log productivity. Note that since firms record capital at the end of the period, capital used in production is k_{it-1}. I assume productivity follows an AR(1) process:

$$\omega_{it} = \rho\omega_{it-1} + \xi_{it}$$

The estimation proceeds in two stages. In the first stage, I remove idiosyncratic measurement error from the production process. Specifically, I calculate predicted sales from a sales-weighted regression of sales on the variable input and capital, with firm and year indicator variables:

$$q_{it} = \beta_v v_{it} + \beta_k k_{it-1} + \mu_i + \mu_t + \epsilon_{it}$$

In the second stage, I use these sales estimates to derive implied productivity as a function of elasticity parameters β. I project this function onto its lag to recover the function of innovation to productivity. I use this function to recover industry-specific output elasticities by assuming the following moment conditions:

$$E[\xi_{it}(\beta) \begin{matrix} v_{it-1} \\ k_{it-1} \end{matrix}] = 0$$

These conditions are valid under the assumption that the variable input and capital in production respond to productivity shocks, but their lagged values do not.

One potential concern is whether the production-based approach produces plausible estimates of markups based on the crucial assumption that firms minimize variable input costs. As a test of this concern, De Loecker and Scott (2016) use demand and production data in the US beer industry to estimate markups using both the more established demand-side and newer production-side approaches. They find statistically indistinguishable estimates when using retail and wholesale price instruments from a demand approach, although as in demand-side estimates, these results are somewhat sensitive to the choice of instrument set.

3   Aggregate Market Power for Publicly-Owned Firms

Applying this estimation routine provides annual firm-level markup estimates for each entity in Compustat. With these in hand, I calculate annual aggregate market power estimates by taking a sales-weighted average each year. The blue line in Figure 1 plots this aggregate market power time series from 1950 to the present. The shaded areas represent NBER end-of-year recession dates. In 1950, markups were about 15% over marginal cost. Over the next 30 years, they decreased approximately linearly, reaching just under 10% over marginal cost at the beginning of the 1980s. From then until today, they have increased approximately linearly, returning to the 1950 level. Consistent with Nekarda and Ramey

(2013) and contrary to the predictions of New Keynesian models, markups tend to decrease in recessions and are hence somewhat cyclical.

Figure 1: Aggregate Market Power in Compustat
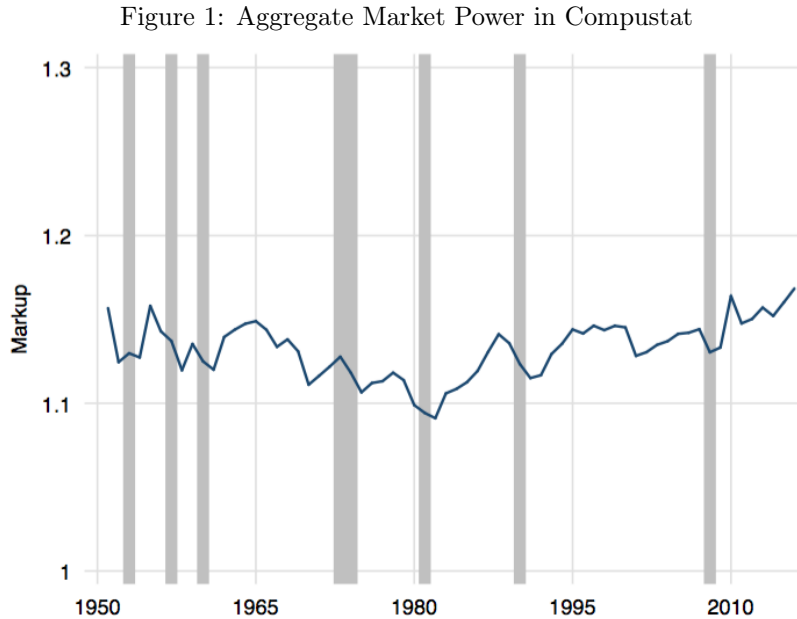
The key conclusion from this figure is that public firm markups increased only modestly from 1980 to the present. Moreover, this increase is well within historical variation. Together, these facts are strong evidence against the argument that firm markups have dramatically increase since 1980, and hence that market power can explain secular trends in industrial concentration or factor shares. Nevertheless, it should be noted that markups are still persistently above 1. This fact is direct evidence of market power in the US economy, contradicting typical macroeconomic assumptions of perfect competition among firms.

4 De Loecker and Eeckhout (2017): Cost of Goods Sold or Operating Expenses?

So far, this empirical methodology has closely followed De Loecker and Eeckhout (2017). However, an important missing feature of their markup calculation is the part of operating costs largely attributable to getting the product to the consumer. This part of the cost is important because the output measure uses the final sale, which includes these costs in production.

Accounting data classifies a firm's total expenses as either Capital Expenses or Operating Expenses (OPEX). The former directly builds the capital stock of the firm, whereas the latter

is an expenditure a business incurs through its normal business operations. OPEX is divided into Cost of Goods Sold (COGS) and Selling, General and Administrative Expenses (SGA)[7]. COGS measures direct inputs to production, such as materials and most of labor. SGA measures indirect inputs to production and mostly consists of marketing and management.

De Loecker and Eeckhout (2017) focus on the COGS part of OPEX as a measure of variable cost and conclude that markups have risen significantly since 1980. However, Figure 2 shows this increase largely disappears when we include SGA costs. The red line displays the markups estimate using COGS as variable costs. COGS markups show a small hump shape from 1950 to 1980, and then a dramatic rise of 15% above marginal cost in 1980 to 40% above marginal cost today. The blue line displays the markups estimate using OPEX as variable costs as in Figure 1.
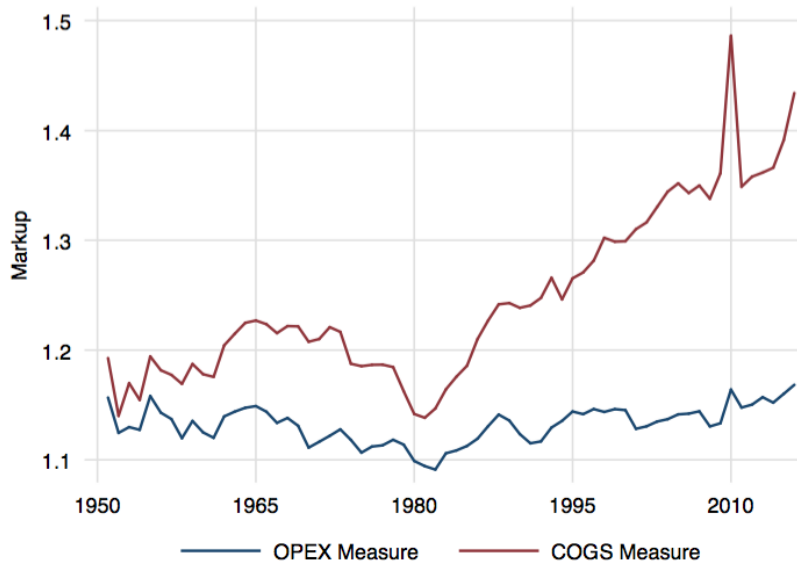
Figure 2: COGS vs. OPEX Markups

Figure 2 shows that correctly measuring variable costs is vitally important to getting the facts about markups right. With all components of variable cost, the red line tells a dramatically different story than the earlier claims. First, public firm market power is not substantially high, with a typical markup of 10% over marginal cost. Second, and more importantly, market power has not meaningfully increased since 1980.

The qualitative difference in the trends of these time series is not driven by choices in

---

[7]OPEX may include a residual term, Other Operating Expenses. For ease of exposition, I categorize these expenses in SGA. All results are robust to using OPEX only or the sum of COGS and SGA. The median observation in my Compustat sample has 0 Other Operating Expenses, with an average share of OPEX attributed to this category of .00003%.

estimating production functions, and therefore output elasticities of variable inputs, beyond the argument that omitting SGA increasingly biases the output elasticity on a COGS-only variable input production model. I report the raw data of variable input margins $\frac{P_{it}^Q Q_{it}}{P_{it}^V V_{it}}$ for both COGS and OPEX in Figure 11 in the Appendix, demonstrating this point. Although the level is mechanically higher since the margins are no longer downweighted by output elasticities θ_{it}^V below 1, the trend in adjusted markups is flat.

To delve into the details of how these time series can disagree so meaningfully, I expand on evidence that the earlier work was largely driven by omitting SGA expenses. Figure 3 shows the binscatter of the COGS markup measure plotted against the SGA share of OPEX. To control for idiosyncratic differences within firms and across time, both variables are residualized on firm and year fixed effects.[8] The blue dots represent 5% quantiles of the data, and the red line represents the corresponding line of best fit.

Figure 3: COGS Markups and SGA Costs

Bin scatter of variables of interest residualized on firm and year fixed effects.
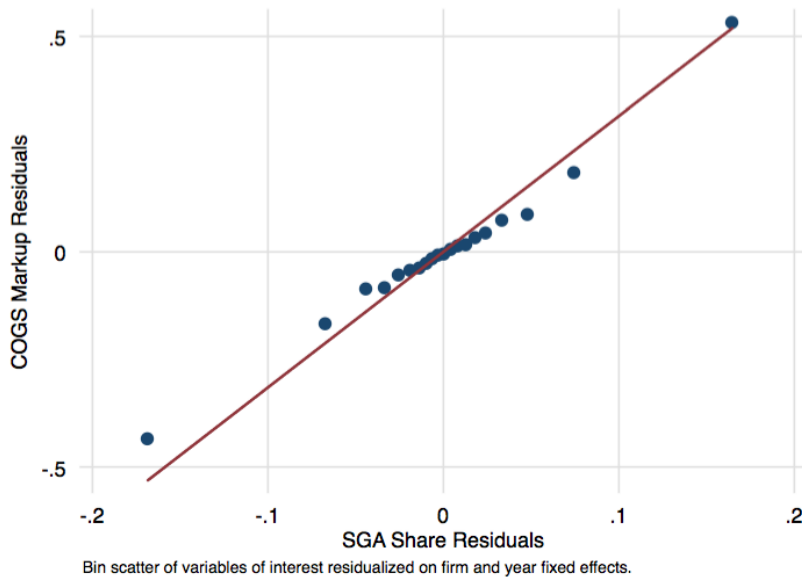
As the figure shows, COGS markup is highly correlated with the SGA share of OPEX at the firm-level. This illustration confirms that a significant part of the original markup estimation is misattribution of markups to SGA variable costs. However, if this omission were the entire story, then only the level of markups would be overestimated, not the trend.

---

[8]This figure remains largely unchanged from the raw binscatter of SGA share against OPEX, and hence this relationship holds both within and between firms and time.

To see this, recall the markups equation

$$\mu_{it} = \theta_{it}^V \frac{P_{it}^Q Q_{it}}{P_{it}^V V_{it}}$$

If $P_{it}^V V_{it}$ is undermeasured but consistently through time, then both θ_{it}^V and $P_{it}^V V_{it}$ would be stable across t. By omitting an input in the production function estimation, the output elasticity of interest θ_{it}^V would be biased upward. However, although firms would seem more productive with higher markups, this bias effect alone would not change since 1980. For such a bias to affect the trend in markups, we would need the share of omitted costs to increase over time.

Figure 4 shows exactly this trend, plotting the share of OPEX attributable to COGS over time. The share is largely decreasing from about 88% of OPEX in 1950 to about 78% of OPEX today. There is a notable inverse hump from 1950 to 1980 as with the COGS markup measure in Figure 2.

Figure 4: COGS Share of OPEX
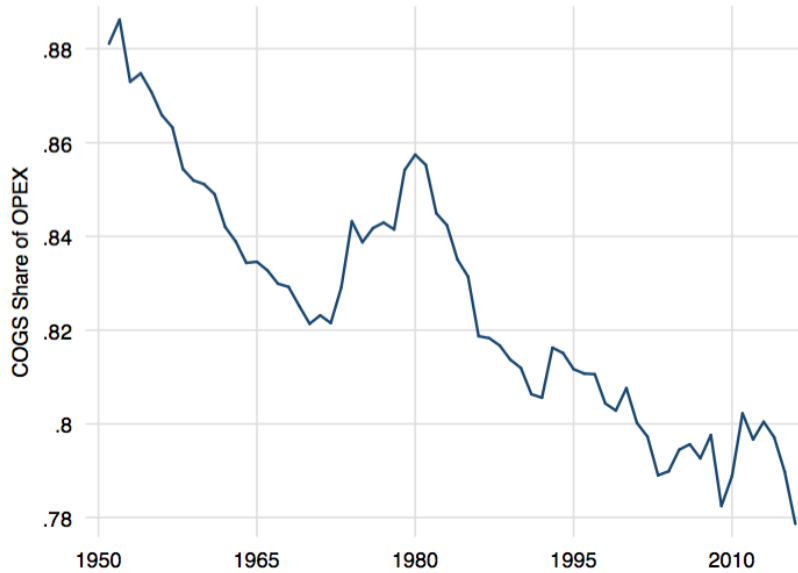
As is clear, this part of costs is a decreasing share of operating expenses. Equivalently, SGA is an increasingly vital share of variable costs for firms in the US economy. This trend highlights that if we neglect SGA operating expenses in markup estimation, we would see an increasingly undermeasured variable cost and so an increasingly overestimated markup. I further address implications of this trend for firm production in Section 6 below.

Note that in principle, we could alternatively estimate 3-input production functions of

capital, COGS, and SGA in rolling windows. Based on Figure 4, we can infer that the COGS elasticity will decline over time as more production shifts to SGA. This alternative requires a more complex production function on two dimensions: (1) three inputs, instead of two; and (2) time-varying output elasticities, which are more difficult to estimate precisely. Another alternative is to capitalize SGA. This addresses the 3 input problem if you lump it together with capital, but you'd still have to estimate time-varying output elasticities for COGS, which based on the cost share figure are decreasing over time.

For a final sanity check, I recalculate the annual aggregate markup series by only considering observations for which COGS exceeds SGA. These observations will have the smallest omitted variable bias in a COGS-only production function, and are perhaps more reliable in our assumptions on variable inputs. Figure 12 in the Appendix shows that markups calculated on this subset of firms are flat regardless of accounting for the omitted variable bias.

5   Longer Time Horizons

One concern with this alternative measure of costs is that some part of SGA represents fixed, rather than variable, costs. Accrual accounting follows the matching principle whereby firms record the timing of expenses to match either the revenue they generate or the period in which they're consumed. Qualitatively, a good feature of OPEX is that its components are largely marketing and management. Both features might reasonably be variable in that they affect output largely through the current year's production process, as confirmed by the matching principle. However, when both connections are infeasible, such as with research and development, costs are immediately expensed. This rule disconnects the timing of some expenses of SGA, implying that they are fixed costs instead of variable costs.

However, all fixed costs become variable in the long run, so this problem reduces to whether annual production function estimates are sufficient for any fixed cost components of OPEX to become variable costs. As a sanity check to determine the quantitative size of the problem, I repeat the above estimation with a 3-year production process. If fixed costs in OPEX are not quantitatively important, the time series aggregate markup of the original 1-year production process and the 3-year production process should look the same.

Figure 5 plots the original series (blue line) and the 3-year production process series (red line) together. The new line looks approximately the same as the old line, with a gradual decrease from about 1.15 in 1950 to 1.10 in 1980, and a gradual increase back to 1.15 today.

Figure 5: Fixed Cost Robustness
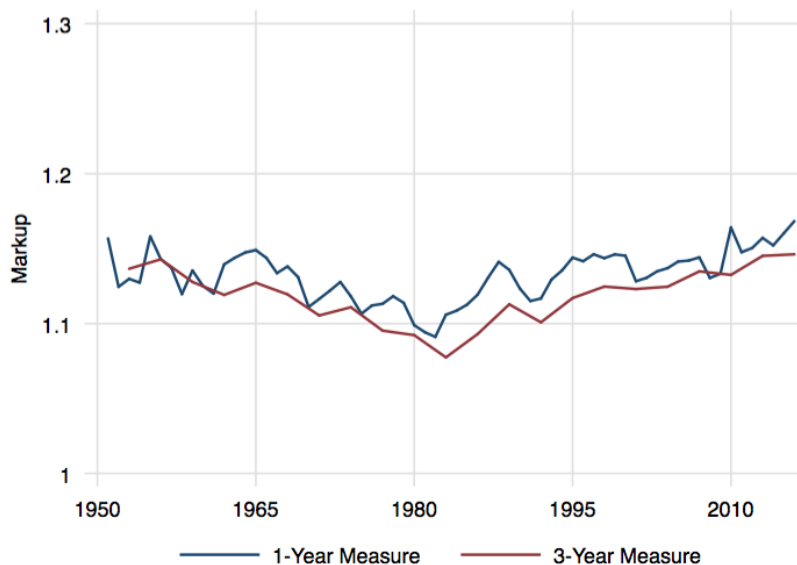
This figure confirms that OPEX is altogether a variable cost in production. This fact is important – if SGA is largely a fixed cost, and fixed costs go up, markups need to go up to pay off fixed costs in equilibrium. However, although it might first appear as a qualitative judgment call, the quantitative results strongly support the view that SGA is not a fixed cost in annual production processes.

As an alternative robustness test, I also show in Figure 13 in the Appendix that SGA expenses are smooth within firm. I subset on firms which form a balanced panel from 1981 to 1990, and rank their changes in log SGA. I then plot the average of these changes within each rank. If SGA were a fixed cost, these changes would look lumpy – the distribution would be skewed toward large changes, i.e. the 1st rank would be substantially large. On the contrary, the distribution across ranks is fairly linear, supporting the hypothesis that SGA costs are routine expenses that vary with output.

6   Interpreting the Rise in SGA

The changing nature of production in the US economy is consistent with theories of the consequences of information and communication technology improvements, and in particular the rise of intangible capital. As suggested by Dorn et al. (2017) (among others), these improvements may increase the efficiency of marketing and management. As a result, firms will shift production from COGS to SGA. These technology improvements also facilitate

economies of scale, leading to an efficient shift toward large firms, and to concentrated industries. This also matches the rise in the college skill premium, insofar as these technologies are complements to high skilled labor and substitutes for low skilled labor. Although this paper does not yet offer causal evidence in this direction, my findings are consistent with this view of the world relative to rent-seeking returns to scale.

Perhaps the most notable interpretation of the rise in SGA is the with the rise in intangible capital. Peters and Taylor (2017) emphasize this point, capitalizing a fraction of SGA expenses to create a measure of organizational capital. Including intangible capital in their investment and capital measures, they find a stronger investment-q and investment-cash flow relationship in firm-level data. Notably for this study, these relationships are robust to varying the fraction of capitalized SGA from 0% to 100%. Indeed, the raw correlation of log SGA to log intangible capital as measured by Peters and Taylor (2017) is 0.92. Hence, there is good evidence that this rise in the SGA share of expenses is in fact not just an accounting phenomena, but an economic one.

7 Dispersion in Markups

I next explore patterns of heterogeneity in the markup estimates.[9] Figure 6 plots the binscatter of markups against log real sales. The blue dots represent 5% quantiles of the data, and the red line represents the corresponding line of best fit.

Figure 6: Markups and Firm Size

---
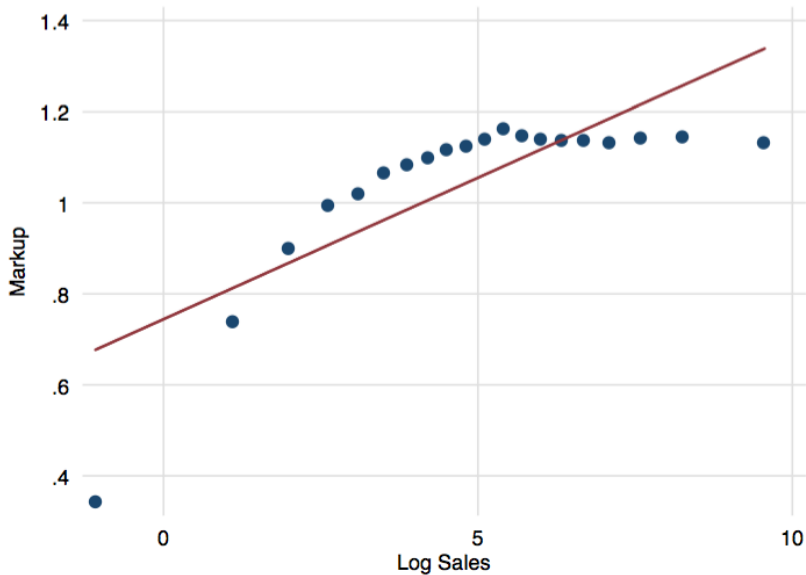
[9]For the raw distribution of markups in the data, see Figure 14 in the Appendix

Figure 6 shows a strong positive correlation between markups and firm size. Note that this correlation is not obvious – although Cournot competition grants such a correlation, many standard models of monopolistic competition do not. Nevertheless, this positive correlation is especially important for drawing inferences about aggregate market power using this public firm sample. Since public firms are often larger than private firms, markup calculations using only public firm data would overestimate aggregate market power.

Figure 7 shows the substantial heterogeneity of markups by sector. The horizontal axis sorts firms into bins based on the sector division of their SIC code. The vertical axis plots the sales-weighted aggregate markup of firms within these sectors across the entire sample. The corresponding estimates range from 2% over marginal cost for Mining and Wholesale Trade to 20% over marginal cost for Transportation, Communication, and Services.

Figure 7: Markups by Sector
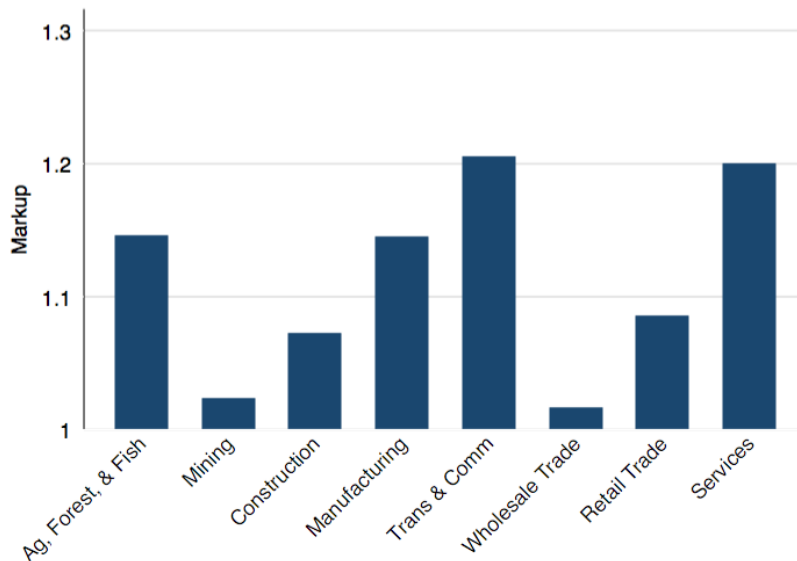
This figure suggests some sectors are in fact extremely close to perfect competition, while others are not. Moreover, although the industry mix is certainly important, it does not obviously interfere with inference in mapping public firm markups to aggregate market power in either direction of under or overestimation. Notably, the sectors that are overrepresented (Manufacturing) or underrepresented (Construction) in Compustat are roughly in the middle of the markups distribution. Of course, this point does not hold if either the industry mix or the markup by industry change substantially over time.

8   Representativeness of Compustat

The previous section casts important doubt in drawing inferences from public firm markups to aggregate market power estimates. Public firms represent about 1/3 of the US economy, with a substantially different industry mix and a strong bias towards larger firms. In addition, these differences are likely dynamically changing over time.

Figure 8 provides some flavor of the Compustat sample's stability by plotting firm size of public firms over time. The blue line tracks the average size of an entity in Compustat, as measured by its real log sales. The average firm size was high in 1950, and decreased substantially until the early 1980s, at which point it started increasing again.

Figure 8: Compustat Sample Firm Size
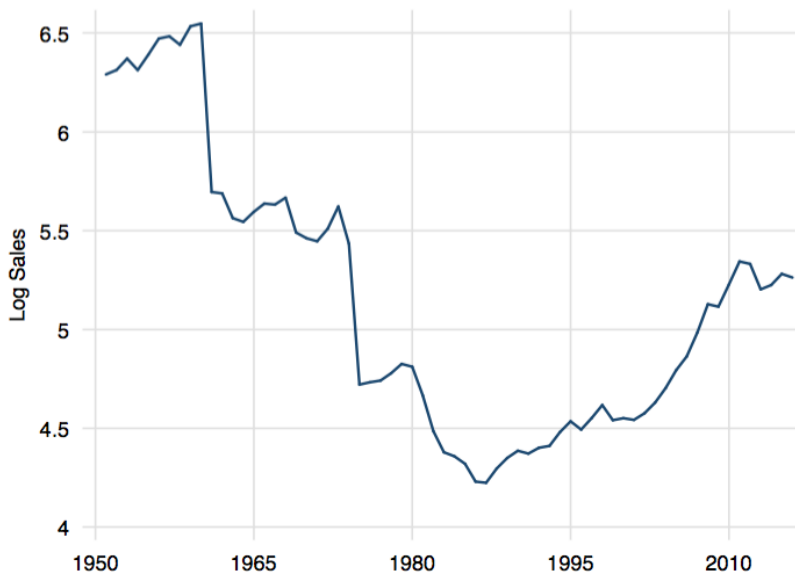
This size figure displays a remarkably similar pattern as the public firm markup time series. Together with the positive correlation between markups and size as documented in Figure 6, a possible hypothesis to the small trends we see in Figure 1 is that the changes in markups are the result of changes in the composition of public firms. As the sample of Compustat firms move towards smaller entities from 1950 to the early 1980s, the average markup declines. This trend reverses as the average firm size begins to increase from 1980.

To push this suggestive evidence further, we need a measure of what the size distribution should be in Compustat to effectively represent the aggregate economy. The Census's Business Dynamics Statistics provides a first pass at this measure, offering information on the annual share of employment in different employee size thresholds for the entire US pri-

vate sector starting in 1977. Employment information is also available in Compustat, and has a correlation coefficient with real sales of 0.72. By calibrating the Compustat firm size distribution to the Census's firm size distribution, we can improve our estimate of market power to better account for the size bias of public firms.

Figure 9 shows the aggregate market power estimate using Compustat markups reweighted to match the Census size distribution. The blue line illustrates a full reweighting, whereas the red line represents a reweighting that excludes firms below 100 employees. The former is decreasing since 1977, and the latter is completely flat.

Figure 9: Reweight to the Aggregate Size Distribution
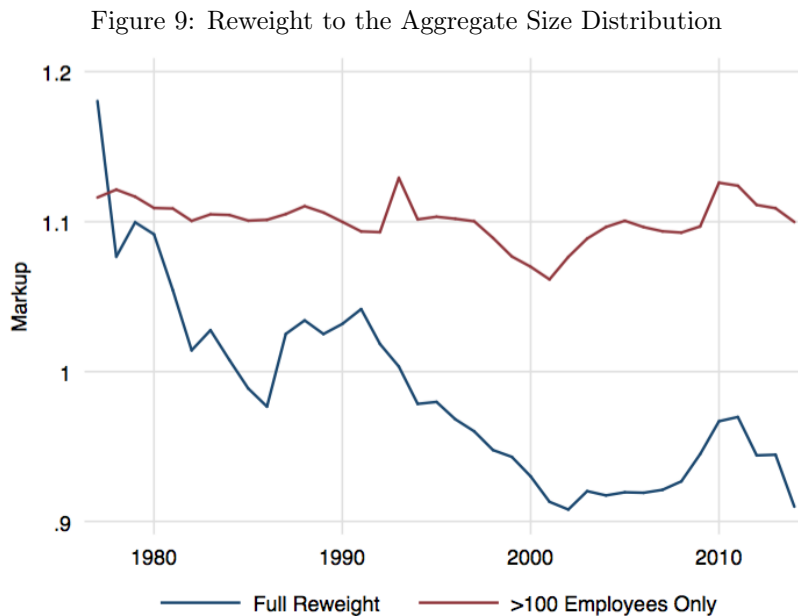
This evidence suggests that even selection into Compustat can completely account for the the small increase in measured public firm market power. The full calibration is decreasing nearly linearly, even falling below 1. However, this decrease may be misleading since it is primarily driven by small public firms, which are relatively rare and may poorly represent small private firms. The calibration that truncates these smallest firms is likely close to the actual aggregate market power estimate, showing a persistent markup of about 10% over marginal cost.

As noted in the previous section, these estimates may be further contaminated by an unrepresentative industry mix. To test this hypothesis, I collect data from the Census's Business Dynamics Statistics that similarly bins the share of US employment by size and SIC sector division. Figure 10 shows the same exercise as Figure 9 with this reweighting.

15

Figure 10: Reweight to the Aggregate Size-Sector Distribution
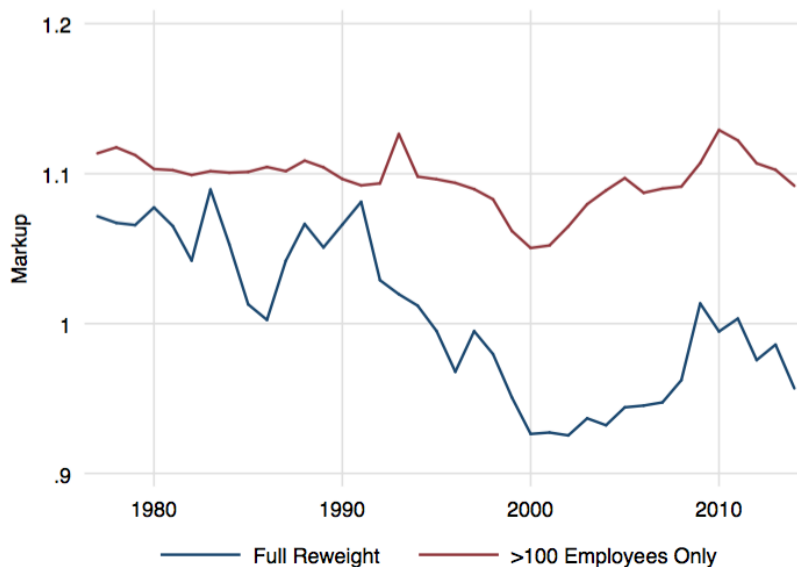

Figure 10 looks remarkably similar to Figure 9, showing that the representativeness of the public firm industry mix is of relatively little quantitative importance. Altogether, these estimates provide strong evidence that aggregate market power has not dramatically increased since 1980.

There are, of course, limitations to this analysis since we do not have information on markups (either levels or trends) for privately-owned firms. Publicly-owned firms may differ in other ways that affect the representativeness even within size-sector cells. For example, public firms tend to be older than private firms, as firms typically go public after operating as privately-owned. However, it's somewhat reassuring that within Compustat, markups are about constant throughout the firm life cycle (see Figure 15 in the Appendix). Nevertheless, public firms may have different markup patterns for other reasons of selection, so the results of this paper are indicative but not definitive for the aggregate economy.

9 Conclusion

This paper provides evidence on a fundamental question of the modern US economy: how has firm market power changed over time? This paper challenges the current albeit developing view that market power has risen dramatically since 1980. By proper measurement and consideration of public firm representativeness, I find that firm market power has either remained flat or declined. These results suggest that the rise in industrial concentration,

and the corresponding fall in the labor share, are likely the result of changing production technologies that are capital-biased and have increased economies of scale.

References

Barkai, S. (2016). Declining labor and capital shares. *Job Market Paper*.

Brown, G. and Kapadia, N. (2007). Firm-specific risk and equity market development. *Journal of Financial Economics*, 84(2):358–388.

Davis, S., Haltiwanger, J., Jarmin, R., Miranda, J., Foote, C., and Nagypal, E. (2006). Volatility and dispersion in business growth rates: Publicly traded versus privately held firms. *NBER Macroeconomics Annual*, 21:107–179.

De Loecker, J. and Eeckhout, J. (2017). The rise of market power and the macroeconomic implications. *National Bureau of Economic Research*.

De Loecker, J. and Scott, P. (2016). Estimating market power: Evidence from the US brewing industry. *National Bureau of Economic Research*.

De Loecker, J. and Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6):2437–2471.

Doidge, C., Karolyi, A., and Stulz, R. (2017). The US listing gap. *Journal of Financial Economics*, 123(3):464–487.

Dorn, D., Katz, L., Patterson, C., and Reenen, V. (2017). The fall of the labor share and the rise of superstar firms. *National Bureau of Economic Research*.

Fama, E. and French, K. (2004). New lists: Fundamentals and survival rates. *Journal of Financial Economics*, 73(2):229–269.

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

Harford, J. (2005). What drives merger waves? *Journal of Financial Economics*, 77(3):529–560.

Nekarda, C. and Ramey, V. (2013). The cyclical behavior of the price-cost markup. *National Bureau of Economic Research*.

Peters, R. and Taylor, L. (2017). Intangible capital and the investment-q relation. *Journal of Financial Economics*, 123(2):251–272.

Appendix

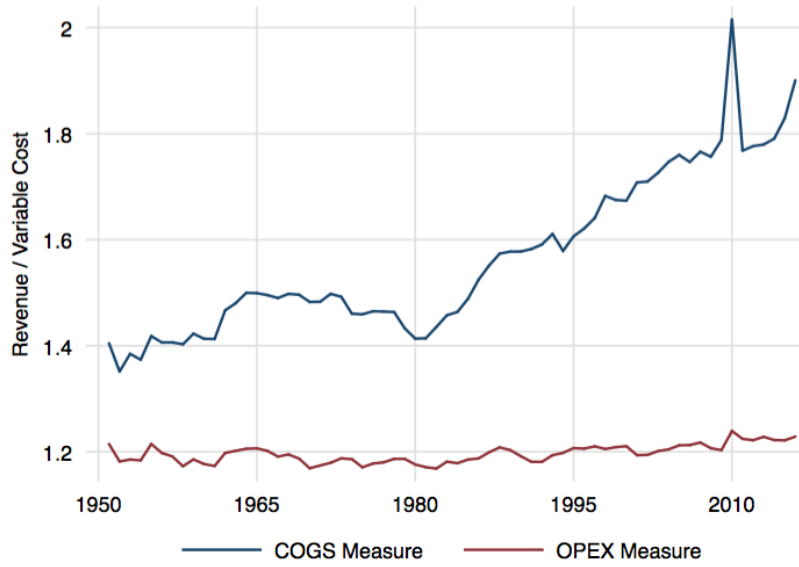Figure 11: COGS vs. OPEX Margins
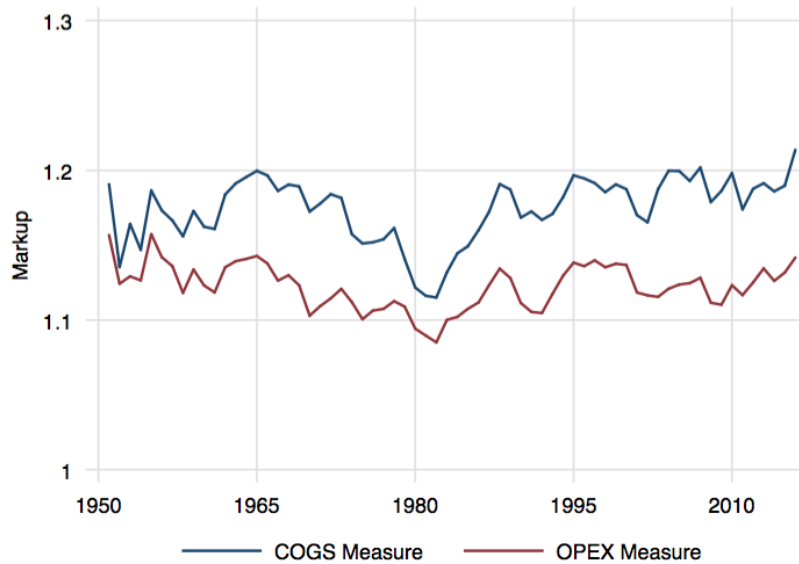

Figure 12: Market Power for COGS-Intensive Firms
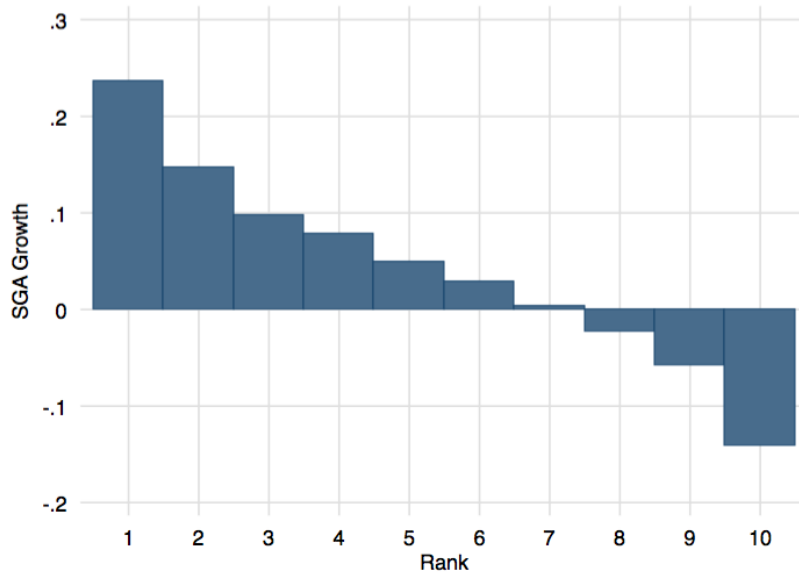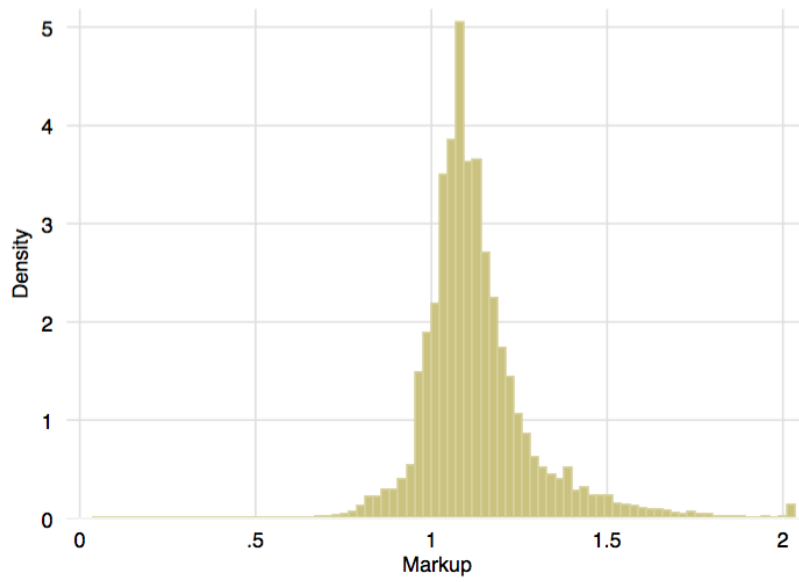
Figure 13: Smoothness of SGA Growth

Figure 14: Raw Markup Distribution

Figure 15: Markups by Age