

The Rise of Market Power and the Macroeconomic Implications: Reply to [Benkard, Miller, and Yurukoglu \(2025\)](#)*

JAN DE LOECKER
KU Leuven[†]

JAN ECKHOUT
UPF Barcelona[‡]

GABRIEL UNGER
Stanford University[§]

October 13, 2025

Abstract

[Benkard, Miller, and Yurukoglu \(2025\)](#) claim the results in [De Loecker, Eeckhout, and Unger \(2020\)](#) are not robust. In this note, we show that their results are exclusively driven by outliers in one four-digit industry (NAICS 3254). They re-estimate the production function including small (biotech) firms that have costs but negligible sales. This generates implausible output elasticities in sector 32, which affects economy-wide aggregate markups. We show that the results in [De Loecker, Eeckhout, and Unger \(2020\)](#) are robust to different sample restrictions, and to alternative methods to obtain output elasticities for these outliers.

1 Introduction

In their comment, [Benkard, Miller, and Yurukoglu \(2025\)](#) (BM_Y), claim that the results in [De Loecker, Eeckhout, and Unger \(2020\)](#) (DLEU) are not robust. We welcome theirs and others' additional analysis and robustness checks. We also value their contribution to helping uncover the facts about market power in the macroeconomy. The main objective of this joint endeavor is to understand the fine nuances of the data and nature of market power and how it has evolved since 1980. As will become evident, BM_Y's analysis together with our continuation to further lay bare the driving forces is akin to a forensic investigation that enriches our understanding of the evolution of economy-wide market power among publicly traded firms in the US.

The BM_Y critique. BM_Y argue that the results in DLEU are not robust for two reasons: 1. the inclusion of FIRE, which consists of two sectors, Finance and Insurance (NAICS 52) and Real Estate (NAICS 53); 2. sample selection due to the exclusion of firms not reporting SGA (Sales, General and

*Eeckhout acknowledges support from the ERC, Advanced grant 882499, from AEI (Barcelona School of Economics CEX2019-000915-S), and from PGC2018-096370-B-I00. Jan De Loecker presented a first version of these findings at the CRESSE conference in Crete on July 4, 2025 in a panel on Market Power.

[†]jan.deloecker@kuleuven.be

[‡]jan.eeckhout@upf.edu – ICREA-BSE-CREI

[§]gunger@stanford.edu

Administrative expenses). Regarding point 1., in DLEU, we already report all the results excluding FIRE and find no significant change in the pattern of markups.¹ We confirm this again here. Regarding point 2., as is common in the literature, DLEU construct a sample that excludes firms that have no observations on the variables used in the analysis, most notable sales, variable costs (COGS), capital, and overhead costs (SGA), as well as trimming of outliers (while assessing the impact of various trims), in order to connect markups to overhead costs and ultimately profit rates.² In an earlier version, [De Loecker and Eeckhout \(2017\)](#) (the 2017 NBER working paper) did not analyze overhead yet and therefore did not drop the firms with no observations on overhead. The results were very similar to those in DLEU. And below, we document that these results are not driven by the sample restrictions. Whether the full sample is used or not, we find a similar pattern of markups and of output elasticities.

Many tiny startups. Here we show that the findings in BMY are entirely driven by outliers in one four-digit NAICS industry 3254 (Pharmaceutical and Medicine Manufacturing). BMY estimate output elasticities that are affected by the inclusion of extremely small firms that have negligible revenue because they produce no output.³ This severely impact the production function estimation and biases the estimates for the entire sample of firms in the broader industry (here NAICS 32). One key requirement of production function estimation is that there be production.⁴ Instead, we find that if we either treat these outliers directly (by eliminating them from the production function estimation step), or consider alternative approaches to obtain output elasticities – for example using cost shares – the pattern of markups that DLEU document is extremely robust, regardless of whether FIRE or not is excluded. Our investigation to make sense of the BMY findings points to a remarkable feature of the data that is concentrated in one industry, the four-digit NAICS industry 3254.⁵

The key features of these small firms are: 1. They have very low sales, some costs, no SGA, and negative profits; 2. There is a substantial number of them; 3. There is a massive change over time, from virtually none in 1990 to more than half of the firms in this 4 digit industry; 4. These small firms jointly have a small market share in industry 3254, less than 3% of revenue, they are negligible in sector 32 (the sector that is uniquely responsible for the main findings in BMY), and even more so in the overall economy.

¹See Online Appendix, Appendix 10 on page 15. This immediately suggests that the change in the markup distribution takes place across the entire economy.

²In Appendix B on page 633, DLEU describe in detail how to construct the sample, and the table lists the number of observation on each variable.

³We refer to these small firms in short as biotech startups. More specifically 3254 is Pharmaceutical and Medicine Manufacturing and it consists of four six-digit industries: Medicinal and Botanical Manufacturing (325411), Pharmaceutical Preparation Manufacturing (325412), In-vitro Diagnostic Substance Manufacturing (325413), and Biological Product (except Diagnostic) Manufacturing (325414). The small biotech startups that are the outliers do not report SGA and they have negligible revenue.

⁴Biotech startups are publicly traded because it is a way to finance the highly risky and uncertain venture of FDA approval and whether or not the clinical trials will be successful. Typically, the expenses are high, often on outsourced lab work and Research & Development expenditures. And most importantly, there are negligible sales, because there is zero or very little production. These biotech firms are often acquired by the large pharmaceutical manufacturers when the trials are successful.

⁵This industry includes the well-known traditional pharmaceutical manufacturers such as Johnson & Johnson, Pfizer and Novartis, all in the DLEU sample, but it also includes a large and increasing number of small biotech startup firms.

The production approach to markup measurement leans on production function estimation to obtain the output elasticity of the variable input. But production function estimation heavily biases the output elasticities in the case of the small firms, reporting extremely low output to input ratios. First, these firms produce negligible quantities, and have revenue that is much smaller than their variable costs, which makes the production function estimation mute.⁶ By including firms that have revenue that is orders of magnitude smaller than their costs, the regression coefficients of the production function are unreliable. Using cost shares instead yields the usual output elasticities because calculating cost shares does not require sales. Second, these firms are negligible in sector 32, yet they have equal weight as large firms in the estimation because production function estimation typically does not weigh the observations. However, in the estimation, a small biotech startup for example has equal weight as Pfizer. This is in itself not a problem, production function estimation often pools small and large firms within a given industry (in fact this is the case across many industries in the DLEU sample). This is not the case here: these firms do not produce and therefore inject outliers into the analysis leading to implausible output elasticities, as their output to input ratios are minuscule.

2 Robustness: the role of outliers

We now investigate the claim that our results are not robust to the specific sample that is constructed (points 1. and 2. mentioned above). First, we replicate the BMY result and confirm that indeed when constructing the larger sample, including the firms reporting zero SGA,⁷ and then re-estimating the sector-year specific output elasticities, we obtain the aggregate markups series BMY finds. BMY show that when further excluding FIRE, this discrepancy is even stronger.

We find that their selection of the sample, coupled with re-estimating the production functions for this sample, is the one specification where indeed the markup trends looks different. We do not find it for any other sample, or when moving across samples holding fixed the set of output elasticities obtained without the inclusion of the outliers in the single 4 digit industry. We then show that the findings reported in BMY are solely due to the inclusion of a set of small firms, that start to appear in the Compustat sample in the early nineties. Those firms have almost no sales, while reporting COGS information (leading to extremely low values for the sales-to-COGS ratio), and those firms are uniquely located in a single 4 digit industry, NAICS 3254, during the period 1990-2016. These firms, about 200 of them, constitute less than 3 percent of total revenue inside this 4 digit industry, and are completely negligible in their 2 digit industry (NAICS 32), let alone in the aggregate sample.

So then how come for this slightly different sample, running the same code, we obtain such different results? The answer is that the production function estimation relies on statistical analysis that is designed for estimating industry-year specific production functions not including outliers. When those outliers are included, the statistical analysis can yield extreme results.

Let us remind the reader not familiar with the exact implementation of DLEU, how the main

⁶BMY argue that FIRE (Finance, Insurance and Real Estate) should not be included because financial firms do not have a conventional production function. For a similar reason, firms that do not produce output do not have a conventional production function. We show that there are viable alternatives to include these firms in the analysis.

⁷Whether it is missing, or actually zero, we cannot distinguish.

results are obtained. From the firm’s cost minimization decision, the first-order condition yields a simple expression for the markup, which is the product of the output elasticity of the variable input and the inverse expenditure share of that variable input, revenue divided by the expenditure on the variable input. The inverse expenditure share we can read off for each firm from the data. To obtain the output elasticity, we estimate a common production function for subsamples of firms (by year and sector), which gives us a common output elasticity for all those firms in the subsample.⁸ To strike a balance between number of observations, and in line with the tradition of analyzing productivity, DLEU’s main specification considers production functions at the level of a two digit industry.⁹

For what follows, the industry NAICS 32 (Manufacturing) is of interest, because it contains those small biotech firms that we mention above. While these biotech firms are negligible in sales, there are a few hundreds of them that emerge in the data around 1990. In essence, these firms do not produce, and therefore report extremely low sales-to-COGS ratios, hereby making any estimation of a production function impossible or at least not well-defined.¹⁰ In practice the presence of these small firms, present in a single 4 digit industry (NAICS 3254), greatly affects the estimated output elasticity in the entire two digit industry (NAICS 32). In particular, it leads to implausibly low output elasticities dropping from about 0.85 in 1990 to about 0.5 by the end of the sample. By imposing these two digit elasticities, including these outliers affect the aggregate markup in that industry, and since this industry is about 20 percent of the overall sample (in terms of total sales), also the overall aggregate implications.

The rest of the Reply establishes this in further detail. First we replicate the BMY results. Then we show that the source of the problem is uniquely due to the inclusion of those small firms that can be thought of as outliers in the sales-to-COGS ratio, that perversely affect the output elasticities. We show no other industry has this feature when comparing the slightly bigger sample excluding the SGA requirement. We show that whichever sensible approach one adopts to correct for these outliers (dropping them all together, eliminating extreme values in the sales-to-COGS ratios, or using cost shares for the output elasticities), restores the main findings of DLEU. This is a useful observation and therefore we see the comment of BMY to underscore yet again that the evolution of markups are all coming from the change in the weighted sales-to-COGS ratio across the entire dataset, and are not due to a change in the estimated output elasticities.

A powerful result is that when moving across the various samples (with or without the SGA information, or excluding FIRE) but while keeping the elasticity fixed (to either a common calibrated level of say 0.85, or the ones reported in the original DLEU article), the patterns are virtually identical. This further confirms that all the action is in the sales-to-COGS ratio, and not the output elasticity.

Finally, regarding the inclusion or exclusion FIRE, we do not think there is a right answer, and surely the markup measurement can be refined as BMY suggests. What we do find is that even when

⁸When we use cost shares, we do obtain one output elasticity per firm, though this requires additional assumptions, most notably on the returns to scale of the technology.

⁹DLEU also report results using more disaggregated production functions, all the way to firm-specific cost-shares which allow for full heterogeneity in output elasticities across firms and time under the prevailing assumptions required to implement the cost-share approach, see [De Loecker and Syverson \(2021\)](#). We will also report these again in this Reply.

¹⁰We do think that their reported cost-shares could present some information about their underlying technology, and we will rely on these estimates to assess the robustness of our results. Note that the production function estimation is severely affected, regardless the technique to combat the well-known econometric challenges due to input endogeneity.

excluding FIRE as well as using the bigger sample, our results are robust, as long as we do not include the set of small firms that report no sales, and that perversely affect the estimated output elasticities.

This discussion points to the need to include more heterogeneity in the production function estimation. In recent work, [De Loecker and Eeckhout \(2025\)](#) rely on firm-specific cost shares, which further confirms the findings of DLEU.¹¹ Now each of these 200 firms accounting for less than 0.003 percent of total sales of the sample gets their own elasticity. This does not affect the entire sample.

We now proceed by providing more details about the above findings and statements, with the intent to corroborate the robust pattern of markups. We start by introducing four distinct samples, allowing us to locate the presence of those firms reporting extremely low sales-to-COGS ratios. We report the implications for the estimated output elasticities, and the next step is to show how these affect the overall aggregate markup patterns. Finally, we show that when adequately dealing with the outliers in a host of different ways, the main findings of DLEU are robust.

3 Robust Patterns of Market Power

We consider four different samples, presented in Table 1, where A and C are the samples used in BMY, while [De Loecker and Eeckhout \(2017\)](#) and DLEU have considered various samples (essentially A through D). We summarize the robustness of our findings across these samples, while varying the approach to estimate the output elasticities, and highlight that the sample selected by BMY, together with the practice of re-estimating the production functions, is the one case that jumps out. We then proceed by identifying the source behind these diverging result; and detect the outliers in a single 4 digit industry. Finally, we report the aggregate markups across the samples, and highlight that when relying on the original reported elasticities (either production function estimation or cost-share based ones), or using the novel firm-specific ones, the results are robust, as previously reported, indicating it is not the sample that generated the BMY finding. In addition, when treating the outliers in that specific 4 digit industry (NAICS 3254) using a range of what we think are sensible approaches, we restore the findings of DLEU.

3.1 Sample definitions

Table 1: Sample definition

Reported Variables	Entire Economy	Excluding FIRE
SALES, COGS	A	C
SALES, COGS, SGA	B	D

Notes: FIRE refers to sectors NAICS 52 and 53. All samples require information on the capital stock (ppeggt in the standard case), but this does not add a meaningful distinction across the samples introduced above.

Table 2 summarizes the combinations of samples and approaches to estimating the production function, and **X** indicates the combination that yields diverging results. Interestingly, most of the

¹¹Being aware that firm-specific cost-shares have all the well-known caveats.

combinations, listed in red, had already been presented by DLEU or earlier versions (De Loecker and Eeckhout (2017)). The main message from Table 2 is that there is one case, the one where the output elasticities are re-estimated, at the sector-year level, on the broader sample A (or C in the case of dropping FIRE) that yield diverging results, starting around 1990, and that is the case BMY reports and relies on to reject the results in DLEU and do not warrant the conclusion that aggregate markups have, on average, increased over time.

Table 2: Assessing Robustness of DLEU (2020)

Output elasticity specification	Sample A	Sample B	Sample C	Sample D
Calibrated ($\theta = 0.85$)	✓	✓	✓	✓
QJE (2020): PFsector	✓	✓	✓	✓
QJE (2020): PF sector-time	✓	✓	✓	✓
Re-estimated: PF sector	✓	✓	✓	✓
Re-estimated: PF sector-time	✗	✓	✗	✓
Re-estimated: Cost-share sector-time	✓	✓	✓	✓

Notes: ✓ indicates that headline markup trend is robust, and if in color red it indicates it is part of the original DLEU analysis, while green indicates it was part of De Loecker and Eeckhout (2017) where essentially sample A was used (no requirement to observe SGA) and an industry-specific translog production function specification was used to effectively generate time-varying production functions, and similar patterns were found (see Appendix A and B in De Loecker and Eeckhout (2017) for the sample size and the aggregate markup pattern for sample A). We only engage with the aggregate trend, for more analysis on the distribution, see the original DLEU article.

3.2 The Role of Small Pharma Companies

The one diverging finding, reported by BMY, can be traced back to a set of firms located in a single 4 digit industry (i.e., 3253), which we label as outliers in the sales-to-COGS ratio. Importantly, these firms show up in the sample exactly when the findings of BMY start to diverge from the battery of specifications (in Table 2) reported by DLEU or those we supply in this note, around 1990.

In Table 3 we identify the 2 digit sector where sales-to-COGS ratio contain extreme values, in particular we look for extremely low values. We do so by listing the unweighted and weighted (by sales) 5th percentile of this ratio in the entire sample A; we also list the sales weight of each industry in the last column. There is one sector, 32, where the unweighted 5th percentile is extremely low, 0.23. In other words, sales are 23% of the variable cost. However, the weighted analogue is 1.09 and does not stand out relative to the other sectors. This indicates that in terms of firm counts, there must be a set of firms in this sector only with extremely low values. We repeat the same exercise for sector 32, across its 3 digit industries (Table 4), and find that again only unit (325) is responsible for this fact. Finally, narrowing down to the 4 digit industry, in the lower panel of Table 4, we find that industry 3254 contains the extreme values: overall the 5th percentile is only 0.08 percent! The analogue, in sample B is 1.07, confirming that these outliers are not representative, even for this 4 digit industry alone (note that this pattern is absent in all other 4 digit industries in the entire database!). We provide further support to our conclusion that the difference between sample A and B is stark, and can only be found in that industry: In the Appendix we plot the number of firms and the share of sales and

COGS across samples A and B (in Figure A.2), as well as the distribution of sales-to-cogs ratio's across the various 4 digit industries (Figure A.1). In particular, the outliers in industry 3254 in sample A (in blue) jump out. Table 5 further indicates that this started to emerge by the end of the eighties, and accelerated during the nineties. In particular since 1990, the unweighted fifth percentile goes from 0.13 in the decade 1990-2000, to 0.05 in the subsequent decade.

Table 3: Sales-to-COGS ratio: 5th percentile across sectors (Sample A)

Sector	Unweighted	Sales Weighted	Share Ind
11	0.67	1.11	0.00
21	0.40	1.12	0.03
22	1.11	1.07	0.06
23	0.96	1.02	0.01
31	1.08	1.07	0.05
32	0.23	1.09	0.18
33	1.07	1.10	0.19
42	1.03	1.02	0.04
44	1.06	1.06	0.05
45	1.08	1.10	0.04
48	1.04	1.05	0.03
49	1.00	0.95	0.01
51	0.88	1.30	0.09
52	1.00	1.05	0.16
53	0.91	1.10	0.01
54	0.91	1.10	0.01
56	1.03	1.08	0.01
61	1.05	1.23	0.00
62	0.96	1.09	0.01
71	0.84	1.16	0.00
72	1.06	1.08	0.01
81	1.08	1.01	0.00
99	0.47	1.13	0.02

Notes: For each 2 digit NAICS industry, the unweighted 5th percentile sales-to-cogs ratio is listed, as well the sales-weighted 5th percentile. The last column lists the share of the 2 digit industry, in terms of sales, in the overall sample. Color blue indicates the outlier.

Why is this relevant for the findings of BMY? They re-estimate the time-varying production function on sample A (including firms not reporting SGA). The problem is that these small firms with outliers in the sales-to-COGS ratio dramatically affect the output elasticities for the entire sector 32, and remember this sector in itself is almost 20 percent of the overall sample, and this will be even larger in the C sample when excluding the FIRE sectors. The facts reported in Tables 3 through 5 therefore indicate that we expect the re-estimated output elasticity series to be seriously affected

Table 4: Sales-to-COGS ratio: median and 5th percentile across industries in sector 32

ind3d	Sample A		Sample B		ind4d	Sample A		Sample B	
	p50	5th	p50	5th		p50	5th	p50	5th
321	1.26	1.06	1.27	1.08	3251	1.38	1.05	1.56	1.25
322	1.33	1.13	1.34	1.15	3252	1.37	1.09	1.39	1.08
323	1.43	1.14	1.45	1.17	3253	1.46	0.84	1.49	1.05
324	1.31	1.03	1.32	1.04	3254	1.61	0.08	2.25	1.07
325	1.56	0.12	1.79	1.12	3255	1.55	1.26	1.55	1.30
326	1.36	1.13	1.37	1.15	3256	2.04	1.10	2.09	1.24
327	1.39	1.14	1.39	1.16	3259	1.64	1.05	1.68	1.14

Notes: For each 3 digit NAICS industry, the median and 5th percentile sales-to-cogs ratio is listed, for samples A and B. Color blue indicates the outlier.

Table 5: Sales-to-COGS ratio: median and 5th percentile across periods (ind 3254)

Period	Sample A		Sample A (weighted)		Sample B		Sample B (weighted)	
	p50	5th	p50	5th	p50	5th	p50	5th
<1980	2.15	1.26	2.21	1.43	2.20	1.32	2.22	1.45
1980-1990	1.81	0.37	2.73	1.43	2.04	1.14	2.74	1.48
1990-2000	1.38	0.13	3.47	1.63	2.13	1.05	3.51	1.68
>2000	1.56	0.05	4.46	1.92	2.42	0.99	4.47	2.01

Notes: The 5th percentile and median sales-cogs ratio is listed across four periods, for sample A and B, unweighted and weighted by sales. Color blue indicates the outliers in sample A.

from the year 1990 onwards, and this only for sector 32. Key here is that these outliers are small (in terms of sales or cost share), yet in the estimation, they get the same weight as the large firms.


3.3 Impact on the re-estimated output elasticities and markups: sector 32

We plot various estimates of the output elasticity, across time, for each 2 digit sector in the Appendix (Figure A.3). In particular, we plot the original time-varying output elasticities from DLEU, the prevailing median cost-share (as used and reported in DLEU) and the BMV ones obtained by re-estimating the exact same production function estimation routine to sample A, including the small set of extreme sales-to-COGS ratio firms in industry 3254.¹² We plot the one sector with diverging trends – industry 32 as expected – in Figure 1.

This has the expected effect: the output elasticity series obtained by BMV for sector 32 looks

¹²We identify 1,916 firm-year observations in industry 3254 with sales-cogs ratios below 0.25, where 1,851 of these are not present in sample B (including non-missing/zero SGA information). All the other 4 digit industries in sector 32, have 15 or fewer firm-year outliers (in sample A, but not in B). Interestingly, the second and third highest number of outliers across all 4 digit industries in the data, are industries 5112 (Software publishers) and 5239 (Other financial investment activities), with 142 and 121 outliers, respectively.

Figure 1: Output elasticities and Markup in sector 32: role of outliers



Notes: All series are for sector NAICS 32, either for sample A or B, as indicated between parentheses – e.g., B32 refers to sector 32 sample B (as defined in Table 1). Red lines indicate series from DLEU, blue lines are associated with BMY, green lines refer to output elasticities measured by cost shares, and DLEU(.,1) refers to outlier treatment of firms in industry 3254 by eliminating these firms when estimating production functions for sector 32, DLEU(., 2) refers to outlier treatment of firms in industry 3254 located in sample A: eliminating firms in $3254 \in A \notin B$ when estimating the production functions for sector 32; but in both cases keep these observations for computing aggregate markups. The orange markup series relies on the output elasticities of DLEU(B32) (in red in the left panel), to compute aggregate markups for sample A – i.e., DLEU(A32).

completely different than that of all the other estimates, including the cost-shares for the sample A, or the estimates obtained on sample B. In particular, the elasticity almost drops in half between 1990 and the end of the sample (2016). Observe that the cost shares do not reflect this pattern, as no sales data is required. Small companies with almost no sales are shaping the entire output elasticity pattern.

We plot the aggregate markup for industry 32 next to the elasticity plot. The solid blue represents the BMY aggregate markup for sector 32, while the other series represent the original DLEU series for sector 32 (in red), applying the DLEU output elasticity estimates to the A sample of sector 32 (in orange), computing the aggregate markup (for the A sample of sector 32) using cost-share based elasticities (in green), and finally the series using the re-estimated DLEU estimates on sample A, like BMY, but correcting for the outliers in industry 3254 (in black). For this plot we simply remove firms from industry 3254 in sample A, but not in B, from the sample over which we estimate the production functions, we do however, use these observations when aggregating the industry-wide markup. It is quite clear, and as expected from the above, that the BMY estimate is an outlier in this plot, while all the other series are very similar.¹³

A final step in recovering the headline result of BMY is to then observe that sector 32 is about 20 percent (or more in sample D) of overall sales in the sample, and by relying on sector-wide output elasticities that almost drop in half, leads to sharply declining industry-wide markups for this sector, ultimately affecting the aggregate markup, leading BMY to conclude that the results in DLEU are not robust. We complete this argument in the next subsection.


¹³In the Appendix, A.4, we replicate this finding using aggregate data from the census of manufacturing (recorded at the 6 digit NAICS industry code), providing further support to the series of output elasticities presented in the main text, for sector 32, and therefore for the sector's aggregate markup series (as implied by DLEU).

3.4 Impact on Aggregate Markups

Excluding FIRE: sample B. Before we turn to the full sample A, we consider samples C and D for the aggregate results. We do not think that dropping FIRE is the right decision, rather we want to highlight that the single reason for the findings of BMJ comes from the subsample of firms in industry 3254 that are in sample C and not D, and have extremely low sales-to-COGS ratio's. Our observation is the same when including the FIRE sectors. As highlighted in DLEU and in [De Loecker and Eeckhout \(2025\)](#), the key observation is the increase in the dispersion of the markup distribution over time.

Before we report weighted average markups, in Figure 2 we present the distribution of markups in 1980 and 2016 and this for sample A, comparing BMJ based markups to the DLEU ones (left panel), and to the DLEU re-estimated while correcting for outliers in industry 3254 (right panel). Both panels paint the same picture, confirming the original DLEU results. However, it is also noticeable how the BMJ markup distribution, in 2016, has much more mass at the lower end of the markup distribution, which is as expected from the discussion above. In Figure A.4 in the Appendix we plot the left panel again on sample A, but we omit sector 32 to confirm this point.


Figure 2: Changes in the markup distribution



Notes: Distribution of markups for years 1980 and 2016 are plotted in blue for sample A, and in red for sample B. The left panel compares the distribution of DLEU (on sample B) with that of BMJ (on sample A), including the outliers in 3254. The panel on the right repeats this but computes markups for sample A by correcting for the outliers. Treatment consists of eliminating firms in $3254 \in A \notin B$ when estimating the production functions for sector 32, but these observations are retained to compute markups.

In Figure 3 we present the aggregate markups for Sample C (excluding FIRE), with again the BMJ series in blue, and compare it to the same series as we did above, for sector 32. The BMJ series is the only combination of sample restrictions and specifications that leads to a different pattern. The aggregate series are in broad agreement when comparing the cost-share based series, the one relying on the re-estimated elasticities on sample A, but dealing with the outliers in industry 3254, to the original DLEU series (in red), when deployed on the sample excluding FIRE.

Figure 3: The impact of outliers on the Aggregate Markup Series
Excluding FIRE (Sample C and D)




Notes: All series are for sample C or D, as indicated between parentheses – e.g., C32 refers to sector 32 sample C (as defined in Table 1). Red lines indicate series from DLEU, blue lines are associated with BMY, green lines refer to output elasticities measured by cost shares, and DLEU(., 1) refers to treatment of outliers of firms in industry 3254 located in sample A but not B – i.e., no information on SGA. Treatment consists of eliminating firms in $3254 \in A \notin B$ when estimating the production functions for sector 32, but these observations are retained to compute aggregate markups.

All sectors: sample A. In Figure 4, we perform the same analysis, now comparing sample A versus B, thus including FIRE, highlighting again that the aggregate markup series of BMY is affected by the outliers in industry 3254, as discussed above. The aggregate markup series obtained by imposing the output elasticities from DLEU (in orange) is virtually identical to the series obtained by re-estimating the elasticities on sample A but excluding sector 3254 (the orange line). These two series are yet again very similar to the aggregate series first produced by DLEU, using sample B (in red).

To aid in seeing the full effect of the outliers when re-estimating the output elasticities, we reproduce the aggregate markup series but exclude sector 32 all together, in the right panel of Figure A.4 in Appendix A.

Finally, we briefly demonstrate how the outliers also impact subsequent analysis carried out in BMY, when revisiting the markup growth decomposition. In Appendix A.5, in Figure A.6 and A.7, we show how the outliers impact not only the aggregate markup series, as we already documented above (including and excluding the FiRE sectors). In addition, and as expected, the outliers impact the decomposition by suggesting a within-markup component that flatlines, and eventually starts to decline after the mid 2000s. Sector 32, with a relative big share is responsible for this, and ultimately explains why the aggregate markup series does not trend up as much. Interestingly, the reallocation component is almost unaffected, simply because the dramatic drop in the output elasticity of sector 32 in BMY implies negative markup growth, while the reallocation terms still capture the covariance of relative weights and markups across firms and sectors.

Figure 4: The impact of outliers on the Aggregate Markup Series:
Including FIRE (Sample A and B)



Notes: All series are for sample A or B, as indicated between parentheses (as defined in Table 1). Red lines indicate series from DLEU, blue lines are associated with BMY, and DLEU(., 1) refers to treatment of outliers of firms in industry 3254 located in sample A but not B – i.e., no information on SGA. Treatment consists of eliminating firms in $3254 \in A \notin B$ when estimating the production functions for sector 32, but these observations are retained to compute aggregate markups.

4 Conclusion

We conclude that the results in DLEU are robust to considering a larger sample with firms not reporting SGA (zero or missing). Moreover, we conclude that when including small firms that do not report SGA and that report COGS that is about 5 times larger than sales – in itself a very small share of the industry –, the results are affected as one would expect. It is not recommended to estimate production functions on firms that do not produce output or generate sales. We agree with BMY that it is useful to consider the broader sample A versus B, and that is important to scrutinize the results, and how they are obtained. However, we would not want to base our findings for the US economy based on firms constituting less than 0.002 percent of total sales in the overall sample A. We do find remarkable robustness across a variety of specifications and samples.

Many other observers had already independently replicated DLEU’s findings using a host of different approaches, essentially all confirming that the main action is in the weighted sales-to-COGS ratio, leaving very little action in the estimated output elasticity. We are grateful to BMY for bringing this to the attention yet again. We are the first to agree that obtaining credible estimates of production function coefficients using standard data sources, while allowing for firms to have market power, can be quite challenging (see [Akerberg and De Loecker \(2024\)](#)). However, this proves to be not of first-order concern when it comes to the main set of results in the DLEU analysis.

In this Reply we do not engage with the judgement call that sample A is superior to B; one can argue both ways. In DLEU we use a sample with firms reporting all variables required to measure

markups and profits, since our key insight is that both markups and profits are rising.¹⁴ To establish the latter, we need overhead costs, which we show are increasing, but not enough to offset the rise in markups, and hence profits are rising too. We are grateful to BMY for highlighting the role of excluding firms that have no observations on SGA, and we find that the patterns first presented by DLEU hold. In fact, in an earlier draft (De Loecker and Eeckhout (2017)) we used what is BMY's preferred sample. Regarding omitting FIRE or fine-tuning the markup estimates for this sector, again we do not feel strongly either way. We find that the DLEU results hold up in both samples, regardless of whether FIRE sectors are included or not.

Prior to BMY, DLEU had already investigated the robustness of the results with respect the exact sample composition and various measurement challenges. In light of the specific concern expressed by BMY, we revisit the properties of different samples, and while we learned some new facts on the role of small, loss-making companies located in a particular market in the economy, we confirm the robustness of the DLEU findings.

First, output elasticities – measured in different ways and on different subsamples – are remarkably *constant* over a long time horizon. The rise in aggregate markups is driven by the increase in the inverse expenditure share of the variable inputs, not by changing output elasticities. Second, the rise of aggregate markups cannot be seen without the massive increase in the *dispersion* of markups, and the dispersion of inverse expenditure shares in particular. Not all firms are equal: a few dominant firms have seen a massive increase in their revenue to variable cost ratio, while the vast majority of firms has seen it constant or decline. Third, *overhead* costs play an increasingly important role in the determination of market power. Fourth, also profits – not just markups – have increased and have become more dispersed. The rise in average profits rates is even more remarkable if one recognizes that the fraction of firms with *negative profits* has grown substantially. The comment by BMY sheds light on the fact that biotech startups are an important contributor to the rise of firms with negative profits: they are numerous and small, and they have negligible revenue.

¹⁴Appendix B of DLEU lists the number of observations per variable (Table B1), we copy it in the Appendix. The table indicates that sales, cogs and SG&A information are required to be in the sample for the main analysis, in order to link markups to profits and overhead costs. In the same paper DLEU verify robustness with respect to excluding the FIRE sectors.

References

- ACKERBERG, D. A. AND J. DE LOECKER (2024): “Production function identification under imperfect competition,” Tech. rep., Working paper, University of Texas, Austin.
- BENKARD, L., N. MILLER, AND A. YURUKOGLU (2025): “The Rise of Market Power and the Macroeconomic Implications: Comment,” Stanford mimeo.
- DE LOECKER, J. AND J. EECKHOUT (2017): “The Rise of Market Power and the Macroeconomic Implications,” NBER working paper 23687.
- (2025): “The Macroeconomics of Market Power,” Mimeo.
- DE LOECKER, J., J. EECKHOUT, AND G. UNGER (2020): “The Rise of Market Power and the Macroeconomic Implications,” *Quarterly Journal of Economics*, 135, 561–644.
- DE LOECKER, J. AND C. SYVERSON (2021): “An industrial organization perspective on productivity,” in *Handbook of industrial organization*, Elsevier, vol. 4, 141–223.

Appendix

A Additional figures

A.1 Industry 3254: sample characteristics

To demonstrate that the outliers in the sales-to-cogs ratio (as an indication how the production function estimation will be distorted) are uniquely located in industry 3254, we plot, in Figure A.1, the sales-to-cogs ratio for both samples, A (in blue) and B (in red), for each 4 digit industry in sector 32. Again, industry 3254 jumps out, with sample A having a mode at incredibly low values of sales to cogs ratios.




Figure A.1: The arrival of small firms

We plot the number of firms in industry 3254 by sample, A (in blue) versus B (in red), in Figure A.2. The number of firms in sample A jumps up starting in 1990. However, as indicated in the second panel, these firms constitute less than 3 percent of total revenue, on average, in the industry 3254 (and therefore are negligible in their sector, let alone in the overall economy).




Figure A.2: The arrival of small firms in industry 3254 (Pharmaceutical and Medicine Manufacturing)

A.2 Output elasticities across sectors: sample A


A.3 Excluding sector 32 from BMY

A.4 Census of Manufacturing: aggregate analysis

We consult the publicly provided census of manufacturing database, listing output and input, for each unique 6 digit NAICS code. We compute the cost share (using the same methodology as in DLEU) for each 6 digit industry and year, as well as a calibrated measure of the markup: using 0.85 and the ratio of the value of shipments over the total variable cost bundle (in the database $\text{mat} + \text{cost} + \text{pay}$). In Figure A.8, we report the analogue to Figure 1 using the industry-wide census data and plot the sales-weighted aggregate markup for sector 32, alongside the prevailing average cost shares at that level of aggregation in the left panel. In the right panel we plot the relevant cost shares at both the 2 digit and the 4 digit for industry 3254. Both plots confirm the DLEU series (and its variations) discussed in the main text, both the rising aggregate markup series, and equally important a rather flat cost share series in the period 1990-end of the sample. This is again in stark contrast to the BMY output elasticities, and therefore BMY aggregate markup series for this sector (32). Note, Figure VI in De Loecker et al. (2020) reports the aggregate markup series using the confidential census of manufacturing database, where aggregation is done across plants, and similarly 4 digit cost share are computed to measure output elasticities; and similar results are obtained.

A.5 Decomposing Markup Changes: role of outliers

We revisit the main decomposition result of DLEU, Figure IV (Decomposition of Markup Growth at the Firm Level) of De Loecker et al. (2020), where the aggregate markup series is decomposed into 1) within-markup changes, 2) reallocation and 3) net-entry. BMY claim that this analysis is completely altered when considering the sample of firms not reporting SGA (i.e., sample A). We present their decomposition, including the outliers in industry 3254, and reproduce the same analysis for sample A with the outlier correction performed throughout the main text. I.e., the treatment consists of



Graphs by ind2d

Figure A.3: Output elasticities across sectors

eliminating firms in $3254 \in A \notin B$ when estimating the production functions for sector 32, but keeps these observations for computing aggregate markups, and the decomposition.

As expected the only reason for the diverging results in BMY are due to these outliers. Below we present the decomposition, on sample A, once following the strategy of BMY and once correcting for the outliers as discussed above, and throughout the main text. The figure on the right panel is almost identical to the first decomposition in Figure IV of DLEU. BMY, however, does only generate this reallocation plot of DLEU excluding FIRE, and we repeat the analysis on sample C, presented in Figure A.7, and the same message is obtained: the outliers severely impact the reallocation analysis, and even if once wishes to exclude FIRE, the same conclusions are drawn as in DLEU.

A.6 Main sample of DLEU




Figure A.4: Markup (distribution): BMY excluding sector 32




Figure A.5: Aggregate Markup in Manufacturing: Census data (DLEU)




Figure A.6: Decomposition of Markup Growth: role of outliers in BMY




Figure A.7: Decomposition of Markup Growth: role of outliers in BMY:
Excluding FIRE




Figure C.4: Decomposition of Estimated Markup Growth

Notes: The figure recreates Figure IV of DEU using the full sample without the finance, insurance, and real estate (FIRE) sectors. The black dot-dash line shows the sales-weighted average markup from our main results. The red dashed line represents the contribution of within-firm markup changes, the green solid line represents the contribution of entry and exit, and the blue dotted line represents the contribution of reallocation.




FIGURE IV

Decomposition of Markup Growth at the Firm Level

Figure A.8: Decomposition of Markup Growth: original BMY vs DLEU

TABLE B.1
SUMMARY STATISTICS (1955–2016)

		Sample A		
	Acronym, var.	Mean	Median	No. obs
Sales	<i>SALE, PQ</i>	1,922,074	147,806	247,644
Cost of goods sold	<i>COGS, V</i>	1,016,550	55,384	247,644
Capital stock	<i>PPEGT, K</i>	1,454,210	57,532	247,644
SG&A	<i>XSG&A, X</i>	342,805	29,682	247,644
Wage bill	<i>XLR, WL</i>	1,093,406	130,486	28,116
Employment	<i>EMP, L</i>	8,363	863	221,121

Figure A.9: DLEU Table B1: sample definition