

A Flexible Cost Share Approach to Markup Estimation*

Devesh Raval

Federal Trade Commission

devesh.raval@gmail.com

March 3, 2023

Abstract

Under the production approach to markup estimation, the markup is the output elasticity of any variable input divided by the input's share of revenue. However, output elasticities vary across producers due to non-neutral technological differences. I develop a flexible cost share estimator to account for such heterogeneity when estimating the output elasticity. This estimator generates markups that are similar when estimated with different inputs.

*I would like to thank Chris Adams, John Asker, Emek Basker, Ben Bridgman, Thomas Chaney, Allan Collard-Wexler, Emin Dinlersoz, Liran Einav, Amit Gandhi, John Haltiwanger, Berthold Herrendorf, Dan Hosken, Liran Einav, Ethan Kaplan, Rob Kulick, Fernando Luco, Ildiko Magyari, Ryne Marksteiner, Aviv Nevo, Ezra Oberfield, Ariel Pakes, Ted Rosenbaum, Pierre-Daniel Sarte, Dave Schmidt, Marshall Steinbaum, Philippe Sulger, Andrew Sweeting, Nico Trachter, James Traina, Brett Wendling, Kirk White, Nate Wilson, and Mo Xiao. I also thank Ana Fernandes, Joep Konings, and Benni Moll for providing deflators for various datasets in this paper. Any opinions and conclusions expressed herein are those of the author and do not necessarily represent the views of the Federal Trade Commission, or its Commissioners.

The *production approach* to markup estimation (Hall, 1988; De Loecker and Warzynski, 2012) has allowed economists to evaluate how mergers and changes in trade barriers have affected market power, and to assess whether market power has been rising (De Loecker et al., 2020). It identifies the markup as *any* variable input's output elasticity divided by its share of revenue.

However, productivity differences across firms tend to augment labor (Doraszelski and Jaumandreu, 2018; Oberfield and Raval, 2021; Raval, 2019; Zhang, 2019). When productivity is not Hicks neutral, more productive firms will have different output elasticities than less productive firms. Econometric estimates that ignore such heterogeneity will get markups wrong. Raval (forthcoming) tests the production approach by using different inputs to measure the markup; markup estimates using different inputs have systematically different levels of dispersion and are negatively correlated in the cross section and time series.

In this paper, I develop an estimator for output elasticities that accounts for differences in labor augmenting productivity. I group plants into bins with similar labor augmenting productivities based upon the labor to materials cost ratio and estimate output elasticities as input cost shares within each group. This estimator does not require data on output quantities and so avoids biases from estimating output elasticities based upon revenue production functions. Across five datasets, as well as in Monte Carlo simulations, markups estimated with different inputs using this approach are positively correlated, have similar time trends, and similar dispersion.

1 Production Function

Production is CES with elasticity of substitution σ , neutral productivity A_{it} , labor augmenting productivity B_{it} , and distribution parameters α_l and α_m :

$$F_{it} = A_{it}((1 - \alpha_l - \alpha_m)K_{it}^{\frac{\sigma-1}{\sigma}} + \alpha_l(B_{it}L_{it})^{\frac{\sigma-1}{\sigma}} + \alpha_m M_{it}^{\frac{\sigma-1}{\sigma}})^{\frac{\sigma}{\sigma-1}}. \quad (1)$$

Input shares of revenue are equal to the output elasticity divided by the markup μ_{it} :

$$\frac{w_{it}L_{it}}{P_{it}F_{it}} = \frac{1}{\mu_{it}} \left(\frac{w_{it}}{\lambda_{it}A_{it}} \right)^{1-\sigma} (\alpha_l)^\sigma (B_{it})^{\sigma-1} \quad (2)$$

$$\frac{p_{it}^m M_{it}}{P_{it}F_{it}} = \frac{1}{\mu_{it}} \left(\frac{p_{it}^m}{\lambda_{it}A_{it}} \right)^{1-\sigma} (\alpha_m)^\sigma \quad (3)$$

where λ_{it} is the marginal cost, w_{it} the wage, and p_{it}^m the price of materials.

An increase in neutral productivity A_{it} does not affect input shares of revenue, as the marginal cost λ_{it} falls to exactly compensate. Labor augmenting productivity B_{it} , in contrast, does affect input shares of revenue. When σ is less than one (Doraszelski and Jaumandreu, 2018; Raval, 2019), a plant with higher labor augmenting productivity has a lower labor share, higher materials share, and lower labor to materials cost ratio.

Given competitive input markets, the markup is:

$$\mu_{it} = \frac{\beta_{it}^X}{s_{it}^X}, \quad (4)$$

where β_{it}^X is the output elasticity for variable input X and s_{it}^X is input X's share of revenue (De Loecker and Warzynski, 2012). Thus, the main object of interest is the output elasticity for the relevant input.

2 Flexible Cost Share Approach

The traditional cost share method used in productivity analysis (Foster et al., 2001, 2008) and markup estimation (De Loecker et al., 2020) assumes a Cobb-Douglas production function, so $\sigma = 1$ and output elasticities are the same for all firms.

This estimator requires two assumptions:

Assumption 1 (Cost Minimization) *On average, first order cost minimization conditions hold for all inputs:*

$$E[p_{it}^X X_{it}] = \beta^X E[\lambda_{it} F_{it}]$$

Assumption 2 (Returns to Scale) *Returns to scale are constant.*

A sufficient condition for **Assumption 1** is that all inputs are flexibly determined. However, **Assumption 1** also holds if capital faces time to build adjustment conditions, under which the capital first order condition will hold on average. **Assumption 2** implies that the marginal cost is equal to the average cost, so one can use data on input costs to compute $\lambda_{it} F_{it}$. For capital, a measure of the rental rate of capital r_{it} would be required. An estimate

for the output elasticity of labor is then:

$$\beta^L = \frac{E(w_{it}L_{it})}{E(r_{it}K_{it} + w_{it}L_{it} + p_{it}^m M_{it})}. \quad (5)$$

While [Assumption 1](#) and [Assumption 2](#) are strong, the cost share estimator relaxes other common assumptions. First, it does not require data on firm quantities, which are typically unobserved in production datasets. Thus, it is robust to criticism that estimating revenue production functions can lead to biased output elasticities when markups vary across plants ([Doraszelski and Jaumandreu, 2019](#); [Bond et al., 2020](#)).

Second, this estimator only requires data on labor costs and not on labor input. When workers vary in quality, measures of labor input such as the number of workers may not reflect labor efficiency units. Finally, since only expectations of input costs enter, it is robust to measurement errors in inputs.

I adapt the cost share estimator to the model in [Section 1](#) in which labor augmenting productivity varies across plants. As [\(2\)](#) and [\(3\)](#) demonstrate, the output elasticities of labor and materials depend upon B_{it} – directly for labor, and indirectly for both through the marginal cost λ_{it} .

I estimate cost shares within groups with similar labor augmenting productivity. After dividing [\(2\)](#) and [\(3\)](#), labor augmenting productivity B_{it} is proportional to the labor to

materials cost ratio $\frac{w_{it}L_{it}}{p_{it}^m M_{it}}$:

$$B_{it} = \left(\frac{w_{it}}{p_{it}^m}\right)\left(\frac{\alpha_l}{\alpha_m}\right)^{\frac{-\sigma}{\sigma-1}}\left(\frac{w_{it}L_{it}}{p_{it}^m M_{it}}\right)^{\frac{1}{\sigma-1}}. \quad (6)$$

Plants with a similar labor to materials cost ratio have similar values of B_{it} and so similar output elasticities. I divide plants into groups based upon their labor to materials cost ratio to create groups with a similar value of B_{it} . I then estimate output elasticities as the input share of total cost within each group. The output elasticities for labor and materials for a plant in group g would be:

$$\beta_g^L = \frac{E(w_{it}L_{it}|G = g)}{E(r_{it}K_{it} + w_{it}L_{it} + p_{it}^m M_{it}|G = g)} \quad (7)$$

$$\beta_g^M = \frac{E(p_{it}^m M_{it}|G = g)}{E(r_{it}K_{it} + w_{it}L_{it} + p_{it}^m M_{it}|G = g)}. \quad (8)$$

By using quintiles, five groups approximate the differences in B_{it} across plants; output elasticities are the input share of total cost within the industry quintile.

The standard cost share approach allows only one group, while having every observation be its own group would set the markup to the revenue to cost ratio. By averaging across groups, I allow for measurement errors in capital, as well as for less strict assumptions on the flexibility of capital such as time to build adjustment frictions, while accounting for differences in labor augmenting technology.

The two polar cases above illustrate a major advantage of the grouping approach. The

econometrician can easily vary the size of the group to examine sensitivity to this tuning parameter. In addition, one can estimate production functions at the subindustry or product level at which the number of plants is small.

A potential concern is that wages and materials prices may also vary across plants, as the equation for B includes the ratio of factor prices as well. Such differences in factor prices could misclassify plants into the wrong groups. In that case, one can construct groups based on factor prices, or variables correlated with factor price differences such as year or location, as well as the labor to materials cost ratio.

Production functions could also vary across plants due to different production distribution parameters. Differences in the labor and materials distribution parameters would also affect the labor to materials cost ratio. Thus, the groups in the flexible cost share estimator would approximate differences in the distribution parameters.

In [Appendix A.1](#), I examine the performance of this estimator in Monte Carlo simulations in which manufacturing plants have different productivities A and B and different markups. I find that markups estimated using labor and capital are highly correlated with each other and with the true markup; these findings still hold after introducing plant specific input prices in [Appendix A.2](#).

3 Application

I then estimate markups using the flexible cost share estimator for five production datasets – Chilean manufacturing plants from 1979 to 1996, Colombian manufacturing plants from 1978 to 1991, Indian manufacturing plants from 1998 to 2014, Indonesian manufacturing firms from 1991 to 2000, and thousands of retail stores from a US nationwide retailer for three years. [Appendix C](#) provides more details on the datasets. Output elasticities are estimated as the cost share for each industry quintile across all years. I use labor, materials, or a composite of both as inputs to estimate markups.

Features of the markup distribution are quite similar regardless of the input used to estimate the markup. I examine the same tests as [Raval \(forthcoming\)](#) by comparing dispersion, time trends, and cross sectional correlations among markups estimated with different inputs.

First, I measure markup dispersion through the ratio of the 90th percentile markup to the 50th percentile markup. For the manufacturing datasets, markup dispersion is quite similar across different countries and inputs, with the 90th percentile markup between 36% and 74% higher than the median. The retailer has much less dispersion across stores, with the 90th percentile between 5% and 7% above the median.

Second, I estimate correlations between markup measures by estimating the following regression:

$$\log(\mu_{it}^L) = \alpha + \beta \log(\mu_{it}^M) + \gamma_t + \delta_n + \epsilon_{it} \quad (9)$$

where μ_{it}^L and μ_{it}^M are the markups using labor and materials for establishment i in year t , and γ_t and δ_n are year and industry fixed effects. The labor and materials markups are highly correlated with each other. An establishment with a 10% higher materials markup has, on average, a 7.5% higher labor markup for Chile, 3.4% higher for Colombia, 6.8% higher for India, 7.2% higher for Indonesia, and 8.9% higher for the retailer.

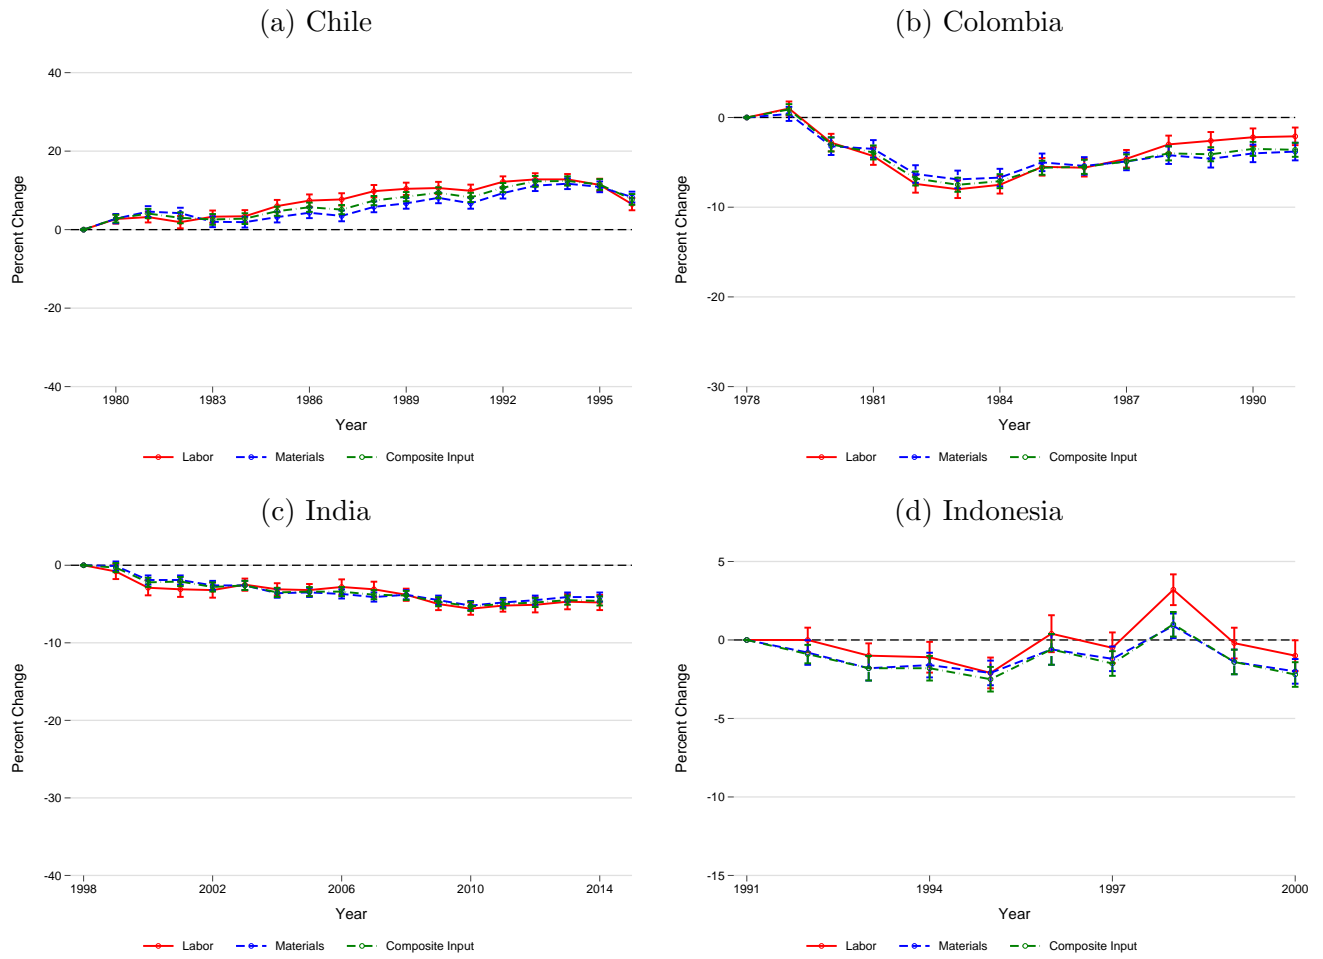
Finally, I examine time trends by estimating the following specification:

$$\log(\mu_{it}^X) = \alpha + \gamma_t + \delta_n + \epsilon_{it} \tag{10}$$

, where μ_{it}^X is the markup using input X for establishment i in year t , and γ_t and δ_n are year and industry fixed effects. I depict the time trends for markups in [Figure 1](#). Time trends in markups are very similar across inputs; the largest difference in markup trends for any year across all the datasets is 4.2 percentage points.

I then examine stylized facts for markups in [Appendix B](#). Markups are positively correlated with size, exporting, and profit shares, as would be expected from theory. For the retailer, I find little relationship between the degree of competition faced by a retail store and markups.

Figure 1 Markup Time Trends



Note: Estimates based on (10), and include 95% Confidence Intervals based on clustering at the establishment level. All estimates relative to the first year.

4 Conclusion

In this paper, I have developed a flexible cost share estimator to account for differences in labor augmenting technology across firms. Using this estimator, markups estimated with different flexible inputs have similar time trends and cross-sectional correlations, and exhibit stylized facts consistent with theory. This paper complements work by [Doraszelski and Jaumandreu \(2019\)](#) and [Demirer \(2020\)](#) developing a dynamic panel and non-parametric control function estimator, respectively, to estimate markups when productivity is non-neutral.

References

- Bond, Steve, Arshia Hashemi, Greg Kaplan, and Piotr Zoch**, “Some Unpleasant Markup Arithmetic: Production Function Elasticities and Their Estimation from Production Data,” Technical Report, NBER Working Paper 2020.
- Demirer, Mert**, “Production Function Estimation with Factor-Augmenting Technology: An Application to Markups,” Technical Report, Mimeo 2020.
- Doraszelski, Ulrich and Jordi Jaumandreu**, “Measuring the Bias of Technological Change,” *Journal of Political Economy*, 2018, 126 (3), 1027–1084.
- and –, “Using Cost Minimization to Estimate Markups,” 2019.
- Foster, Lucia, John C Haltiwanger, and Cornell John Krizan**, “Aggregate Productivity Growth: Lessons from Microeconomic Evidence,” in “New Developments in Productivity Analysis,” University of Chicago Press, 2001, pp. 303–372.
- , **John Haltiwanger, and Chad Syverson**, “Reallocation, Firm Turnover, and Efficiency: Selection on Productivity or Profitability?,” *American Economic Review*, 2008, 98 (1), 394–425.
- Hall, Robert E**, “The Relation Between Price and Marginal Cost in US Industry,” *Journal of Political Economy*, 1988, 96 (5), 921–947.
- Loecker, Jan De and Frederic Warzynski**, “Markups and Firm-Level Export Status,” *American Economic Review*, 2012, 102 (6), 2437–71.
- , **Jan Eeckhout, and Gabriel Unger**, “The rise of market power and the macroeconomic implications,” *The Quarterly Journal of Economics*, 2020, 135 (2), 561–644.
- Oberfield, Ezra and Devesh Raval**, “Micro Data and Macro Technology,” *Econometrica*, 2021, 89 (2), 703–732.
- Raval, Devesh**, “The Micro Elasticity of Substitution and Non-Neutral Technology,” *RAND Journal of Economics*, 2019, 50 (1), 147–167.
- , “Testing the Production Approach to Markup Estimation,” *Review of Economic Studies*, forthcoming.
- Zhang, Hongsong**, “Non-Neutral Technology, Firm Heterogeneity, and Labor Demand,” *Journal of Development Economics*, 2019.