CEWP 21-06

The Behavior of the Aggregate U.S. Wage Markdown

Hashmat Khan & Konstantinos Metaxoglou

Carleton University

July 6, 2021

CARLETON ECONOMICS WORKING PAPERS

Carleton University

Department of Economics

1125 Colonel By Drive
Ottawa, Ontario, Canada
K1S 5B6

The Behavior of the Aggregate U.S. Wage Markdown*

Hashmat Khan†
Carleton University

Konstantinos Metaxoglou‡
Carleton University

July 6, 2021

Abstract

We estimate the aggregate U.S. wage markdown for 1987–2018 using the KLEMS data and the approach in Hershbein, Macaluso and Yeh (2020). Building on the existing literature, our markdown estimates depend on output elasticities of inputs and their shares of gross output. We identify four salient features of the markdown. First, the markdown is above 1, implying an average wage below the competitive level. Second, the markdown has increased over time, mainly during the 2000–2015 period. Third, the variation of the markdown is mostly driven by the input shares and not the output elasticities. Fourth, the markdown is strongly procyclical, implying a larger deviation of wages from their competitive levels during expansions.

Key words: Employer Market Power, Monopsony, Markdowns, Business Cycles
JEL classification: E24, E32

*First draft: July 6, 2021.

†Department of Economics, Carleton University, Ottawa, Canada, hashmat.khan@carleton.ca.
‡Department of Economics, Carleton University, Ottawa, Canada, konstantinos.metaxoglou@carleton.ca.

1 Introduction

Recently, there has been renewed interest in quantifying employers' buyer power, namely, their ability to set wages below the marginal revenue product of workers. Several trends observed in the U.S. labor markets and the economy as a whole, such as stagnant or declining real wages, rising concentration in labor markets and income inequality, declining labor share of income, and deunionization of the workforce, provide the backdrop to this issue.

In a Washington Post Op-Ed published last June, Summers and Stansbury (2020) urged policy makers concerned with inequality, low pay, and poor work conditions, to pay attention to employers' increasing buyer power. A couple of months later, Manning (2020) expressed similar concerns for employers' buyer power arguing that it has been neglected for far too long and something needs to be done about it given its seriousness.[1]

Although the topic of market power in *input* markets is not new in economics and dates back to Joan Robinson's early work on the Economics of Imperfect Competition (Robinson (1933)), it has attracted far less attention than that of market power in *output* markets. Arguably, concentration, markups, demand elasticities, and monopoly power in output markets have attracted much more attention than concentration, markdowns, supply elasticities, and monopsony power in input markets. This has also been the case for Antitrust laws, where the focus has historically been on competition-related issues in output markets and not in input markets.

Recent "micro" developments in input markets driving the "macro" developments discussed above, such as increasing concentration due to the rise of mega employers in the big-tech (Naidu et al. (2018)) and healthcare markets (Prager and Schmitt (2021)), and buyer power in agricultural markets (Hafiz and Miller (2021)), coupled with the emergence of high-profile Antitrust cases (e.g., the "Sillicon Valley" cartel, Streitfeld (2014)) have renewed the interest of both academics and policy makers in assessing the firms' ability to exercise market power in the input markets. Lately, the prospect for a lift of the NCAA's ban on student-athletes earning money has attracted a lot of attention too (e.g., Blinder (2014)).

As a brief review of the fast-growing literature on topics related to market power in input markets (buyer power), Manning (2021) surveys the large body of work on estimating

[1]See Alan Manning's Marshall Lecture delivered on August 25, 2020 at the European Economic Association Congress available here.

markdowns *indirectly* using labor supply elasticities—an empirical exercise that resembles that of inferring markups using demand elasticities—while Benmelech et al. (2019), Berger et al. (2019), Goolsbee and Syverson (2019) and Azar et al. (2020), among others, are examples of recent work on increased concentration in labor markets. Naidu et al. (2018) discuss the role of Antitrust laws in mitigating buyer power. Another stream of the literature has been focusing on obtaining *direct* measures of buyer power using empirical tools developed to assess market power in the output markets; see, for example, Morlacco (2019) and Hershbein et al. (2020). Hershbein et al. (2020), whose work is particularly relevant for this paper, estimate plant-level wage markdowns for U.S. manufacturing plants from 1976–2014. Their estimated markdowns exhibit a downward trend until the early 2000s, and then increase sharply.[2]

We build on Hershbein et al. (2020) and advance the literature along two dimensions with a macro perspective. First, we estimate an aggregate wage markdown for the whole U.S. economy, as opposed to just manufacturing, and document its trend over time. Our analysis is based on U.S. KLEMS data for 1987–2018. Second, we are the first to study the markdown's behavior over the business cycle.

Our main findings regarding the salient features of the time series of the aggregate markdown for the U.S. economy can be summarized as follows. First, the average markdown is 1.9, implying that the aggregate wage is 47% below the competitive level. Second, it exhibits substantial variation with an upward trend. Third, its variation is mainly due to the input shares and not due to the output elasticities. Fourth, the markdown is highly procyclical with a correlation of 0.5 with aggregate output, which suggests that aggregate wages deviate more from their competitive levels during expansions than during recessions.

2 Aggregate Wage Markdowns

Overview. To set the stage for our markdown estimates, consider the profit-maximization problem of a firm exercising market power in both the input (labor) and output markets. Optimal behavior implies a price markup over the firm's marginal cost in the output market, and a wage markdown relative to the marginal revenue product of labor in the

[2]Morlacco (2019) presents micro-level evidence on buyer power in input trade and evaluates its effects on the aggregate economy using trade and balance sheet data for French manufacturing importers over the period 1996–2007. Rubens (2021) studies how ownership consolidation affects productivity and market power in both input and output markets in Chinese cigarette manufacturing.

input market that satisfies the following:

$$\text{Markdown} \times \text{Markup} = \frac{\text{Output Elasticity of Labor}}{\text{Labor Share}}. \tag{1}$$

Although equation (1) may serve as a starting point for estimating the wage markdown, it showcases the main challenge behind this empirical exercise in the absence of additional assumptions. Simply put, knowing the right hand side of (1) is not enough to disentangle the markdown from the markup. Proposition 1 in Hershbein et al. (2020) however, allows us to overcome this empirical challenge. In particular, when a firm's production requires an additional flexible input that is not subject to buyer power, estimating the wage markdown is possible. Using such a "reference" input and an expression analgous to (1) in which the markdown is set equal to 1, we can now write the following equation:

$$\text{Markdown} = \frac{\text{Output Elasticity of Labor}}{\text{Output Elasticity of Reference Input}} \times \frac{\text{Reference Input Share}}{\text{Labor Share}}. \tag{2}$$

Our estimated wage markdown, a measure of the economy-wide labor market power in the U.S., is based on (2). We assume that materials (an intermediate input) is the reference input not subject to buyer power. We then proceed to first estimate output elasticities and subsequently combine these elasticty estimates with input shares to complete our empirical exercise using the KLEMS data for 1987–2018.

Before discussing our main findings, and calling the reader's attention to (1), we should note that the vast macro literature on price markups in output markets—(Bils (1987), Galí and Gertler (1999), Rotemberg and Woodford (1999), Galí et al. (2007), Bils et al. (2018), and, more recently, Burstein et al. (2020), and Nekarda and Ramey (2021)), among others —has been built on the assumption that the wage markdown equals 1. In other words, this strand of the macro literature assumes away employer market power.[3] Although the chapter in the Handbook of Macroeconomics by Rotemberg and Woodford (1999) provides a conceptual framework for the role of buyer power in making markups (profits) countercyclical (procyclical), Rotemberg and Woodford implicitly assume that the markdown is known and constant. Thus, markup dynamics in this literature also capture markdown dynamics complicating their interpretation.

[3]Following Erceg et al. (2000), households are assumed to have labor market power that generates wage markups over the marginal rate of substitution in most DSGE models (e.g., Christiano et al. (2005) and Smets and Wouters (2007), among others). Alpanda (2019) is the first to present a model incorporating oligopsonistic labor markets (hence, wage markdowns) in a DSGE setting.

Estimation. We estimate an aggregate wage markdown $\overline{\nu_t}$ of the form:

$$\overline{\nu_t} = \sum_{i=1}^{n} w_{it} \widehat{\nu}_{it}, \tag{3}$$

where w_{it} is the weight for industry i at time t, and $\widehat{\nu}_{it}$ is the industry-i markdown. Given the data in hand, possible candidates for reference inputs not subject to buyer power are materials and energy. We consider materials (m) as the reference input in our analysis as in Hershbein et al. (2020). We prefer materials over energy as the reference input for two reasons. First, materials account for a larger share of revenues compared to energy in most industries. Second, there is evidence for the presence of monopsony power in energy markets (Davis et al. (2013)).

To illustrate our estimation strategy and using l to denote labor, we write:

$$\widehat{\nu}_{it} = \frac{\widehat{\theta}_{l,it}}{\widehat{\theta}_{m,it}} \times \frac{\alpha_{m,it}}{\alpha_{l,it}}, \tag{4}$$

where $\widehat{\theta}_{j,it}$ is the estimated output elasticity with respect to input j and $\alpha_{j,it}$ is the corresponding share calculated from the data in hand for $j \in \{l, m\}$. We estimate output elasticities assuming a translog production function using capital (k), labor (l), materials (m), and energy (e), while controlling for industry and time fixed effects to address the usual endogeneity concerns (Griliches and Mairesse (1999)):

$$\begin{aligned} y_{it} = {} & \gamma_i + \gamma_t + \beta_k k_{it} + \beta_l l_{it} + \beta_m m_{it} + \beta_e e_{it} + \beta_{kk} k_{it}^2 + \beta_{LL} l_{it}^2 + \beta_{mm} m_{it}^2 + \beta_{ee} e_{it}^2 \\ & + \beta_{kl} k_{it} l_{it} + \beta_{km} k_{it} m_{it} + \beta_{ke} k_{it} e_{it} + \beta_{ml} m_{it} l_{it} + \beta_{me} m_{it} e_{it} + \beta_{le} l_{it} e_{it} + \varepsilon_{it}. \end{aligned} \tag{5}$$

In terms of notation, we use γ_i and γ_t to denote cross-sectional and time fixed effects, and ε_{it} to denote the econometric error. The dependent variable is the real gross output. All the inputs are in real dollars except for labor which is in hours. Based on the assumed translog functional form, the implied output elasticity with respect to labor that exhibits both cross-sectional and temporal variation—analogous expressions can be derived for the remaining three inputs—is given by:

$$\widehat{\theta}_{l,it} = \frac{\partial y_{it}}{\partial l_{it}} = \widehat{\beta}_l + 2\widehat{\beta}_{ll} l_{it} + \widehat{\beta}_{kl} k_{it} + \widehat{\beta}_{ml} m_{it} + \widehat{\beta}_{le} e_{it}. \tag{6}$$

3 Evidence on the Trend and Cyclicality of Markdowns

We use the U.S. KLEMS dataset covering 63 industries for 1987–2018. Given the KLEMS data and the assumed functional form for the production function in (5), we estimate about 2,000 elasticities that exhibit both cross-sectional and temporal variation. These output elasticities can be aggregated using (3) and a variety of weighting schemes. Moreover, following common practice in the literature, the aggregation may be based on a summary statistic (e.g., a mean) of the estimated output elasticities to discern the salient features of the markdowns more easily.

3.1 Aggregate wage markdown trends

Figure 1 shows the low-frequency trend of the aggregate wage markdown for 1987–2018. The line shown in the figure corresponds to markdown estimates obtained with materials as the reference input not subject to buyer power. The markdowns shown are obtained in 3 steps. In the first step, we estimate the output elasticities. In the second step, we calculate the mean of the output elasticities, such that all the variation in the markdowns is due to the input shares of gross output; simply put, "θ ratio" in (4) is constant. In the third step, we calculate the average markdown in each year. Omitting the second step in this 3-step process—allowing the θ ratio to vary—introduces noise in our markdown estimates making it harder to discern its behavior over time. Calculating a weighted average markdown in the third step using real gross output as a weight has essentially no implications for the markdown's behavior over time. The figure also shows 95% confidence intervals for our markdown estimates obtained using the delta method. The lines shown in panel A of Figure 2 are based on alternative summary statistics of the output elasticities in the second step of the 3-step process described. These lines are also shown as indices in panel B of the same figure.

There are several observations regarding the salient features of the markdown series in Figure 1. First, the estimated aggregate markdown is always above 1. This pattern implies that the average aggregate wage has been below the competitive level over the sample period. Second, the estimated output elasticities have no material effect on both the level and behavior of the markdown over time. As Figure 2 clearly shows, the variation in the markdown is driven primarily by the ratio of the input shares in (4). Third, there is an upward trend in the markdown during the 2000s and the first half of the 2010s. A notable short-lived exception is the most recent financial crisis. We also see a drop between 2014 and 2016, followed by a nascent uptick.

The upward trend between 2002 and 2012 in our aggregate markdown for the U.S. economy is similar to the upward trend in the markdown for U.S. manufacturing in Figure 4 of Hershbein et al. (2020). Relatedly, Akcigit et al. (2021) use the same methodology as Hershbein et al. coupled with Orbis data for advanced (excluding the U.S.) and emerging market countries during 2000–2015 to estimate markdowns. In Akcigit et al., the markdown increased in manufacturing, but appears to have decreased in the finance and insurance industries.

3.2 The aggregate wage markdown over the business cycle

We next explore the behavior of the aggregate markdown over the business cycle. We start with Figure 3, which shows the cyclical component of the log aggregate wage markdown, along with its counterpart for the log real GDP, in the top two panels. We also show a scatterplot of the two cyclical components along a fitted regression line. For both series, we extract the cyclical components using the Hodrick-Prescott filter.[4]

Both graphs are consistent with a procyclical behavior of the aggregate markdown. The correlation between the cyclical markdown and output is about 0.5.[5] Based on a time-series regression of 32 observations, the cyclical elasticity (slope coefficient of the regression line in the scatterplot) is about 3 and the R-squared of the regression line is about 0.2. A similar analysis based on the aggregation of the output elasticities in Figure 2 leads to essentially the same conclusions as the ones drawn here.

To the best of our knowledge, there is very limited evidence on how employer market power varies with the business cycle. The available evidence is based on estimated labor supply elasticities using micro data.[6] Depew and Sørensen (2013) find that the labor

[4]We use a smoothing parameter of 6.25 following the recommendation of Ravn and Uhlig (2002) for annual data.

[5]The correlation between the cyclical components of the labor share and real output is −0.58. The correlation between the cyclical components of the materials' share and real output is 0.54.

[6]Using meta-analysis of labor supply elasticity estimates, Sokolova and Sorensen (2021) report a weighted median labor supply elasticity equal to 2 for the top journals in their table 1. A labor supply elasticity of this magnitude would generally require a markdown of 1.25. In Kroft et al. (2020), the point estimate of the firm-specific labor supply elasticity is 4.1. This indicates that, if an American construction firm aims to increase the number of employees by 10%, it needs to increase wages by around 2.4%. This implies wages are marked down 20% relative to the marginal revenue product of labor. Lamadon et al. (2019) estimate a labor supply elasticity of 4.6 using firm-level variation and Serrato and Zidar (2016) estimate a labor supply elasticity of 4.2 using state-level variation, while Card et al. (2018) pick 4.0 as the preferred value in their calibration exercise. A related literature using experimentally manipulated piece rates for small tasks

supply elasticity is procyclical using data for two U.S. manufacturing firms in the 1st half of the 20th century. Their estimates are 4 and 1.6 for expansions and recessions, respectively. Using the standard markdown expression for a monoposonist:

$$\text{Markdown} \equiv \frac{MRP_l}{\text{Wage}} = \frac{1}{\text{Labor Supply Elasticity}} + 1, \tag{7}$$

where MRP_l is the marginal revenue product of labor, their estimates imply a countercyclical markdown in the two manufacturing firms. Hirsch et al. (2018) use German administrative data for 1985–2010 and estimate a mean labor supply elasticity of 2.41, which suggests that workers obtain just $1/(1 + 1/2.4) = 0.7$ of their marginal revenue product. An increase in the unemployment rate by 1 percentage point—based on the results from model 3 in their Table 2—decreases the elasticity by 0.15 to 2.26, and workers receive only 69.3% of their marginal revenue product. These estimates also suggest a countercyclical markdown.

Using (7) and our estimated markdown in Figure 1, we obtain an *indirect* measure of the labor supply elasticity following the opposite (markdown to labor supply elasticity) route:

$$\varepsilon_t^s = \frac{1}{\overline{v}_t - 1}. \tag{8}$$

We examine the cyclical properties of the series ε_t^s in panels B and D of Figure 3 following the same approach as in the case of the wage markdown. Our indirect estimate of the labor supply elasticity is countercyclical, which is consistent with its inverse relationship with the wage markdown. Its cyclical correlation with output is −0.49. The elasticity of its cyclical component with respect to the cyclical component of output is −6.42.

Our finding for the procyclical behavior of the markdown may appear counterintuitve. One might think that jobs are plentiful during booms and the abundance of jobs may reduce employer buyer power as employers compete more on wage offers. Hence, real wages must rise during expansions. However, the data show exactly the opposite. The median real weekly earnings are countercyclical with a correlation of −0.21 with real output during 1987Q1-2018Q4.[7] Thus, a procyclical markdown can arise if the real wage is *less* procyclical than the real marginal revenue product of labor.

typically finds labor supply elasticities ranging from 3.0 to 5.0 (Caldwell and Oehlsen. (2018); Dube et al. (2021); Sokolova and Sorensen (2021)).

[7]See here and here for our wage and GDP series from FRED, respectively. We use filtered logged versions of both series with a smoothing parameter of 1,600 in the Hodrick-Prescott filter.

It is widely documented that the cyclicality of labor productivity has declined sharply since the mid-1980s; see Galí and van Rens (2020), Fernald and Wang (2016), and Galí and Gambetti (2009), among others. Galí and van Rens (2020) report a cyclical correlation between labor productivity (output per hour) and output equal to −0.09 in their Table 1. In the same table, the cyclical correlation between output per worker and output equals 0.32. To the extent that these labor productivity measures are proxies for the unobserved marginal revenue product of labor, real wages are less procyclical than productivity in the data.

The reader may wonder about the feature(s) of the economy that might generate a procyclical wage markdown. Although developing a theoretical model that delivers a procyclical wage markdown is beyond the scope of our paper, the recent work of Fernández-Villaverde et al. (2021) appears to be a good starting point. The authors (FVMYZ, henceforth) emphasize the interaction between search complementarities and monopsony power for understanding rising market concentration over the past three decades. We think that these forces may also potentially help in accounting for the procyclicality of the aggregate markdown that we have documented. In particular, FVMYZ find that the gap between wages and productivity, *increases* with firm size. Thus, their model predicts that markdowns are increasing in firm size, which is supported by empiricial evidence. For example, Hershbein et al. (2020) show that markdowns are increasing in plant size in their Figure 1. Similarly, Akcigit et al. (2021) find that labor market power and firm size are positively related in their Figure 15(2). A procyclical (average) firm size, which is likely given the observed evolution of market concentration, may allow the the FVMYZ model to deliver a procyclical markdown.

4 Conclusion

We estimate the aggregate U.S. wage markdown during 1987–2018. Using the KLEMS dataset and an approach to disentangle the wage markdown from the price markup, we identify four salient features of the aggregate wage markdown. First, the average aggregate markdown is 1.9, implying that the average wage is 47% below its competitive level. Second, the aggregate markdown has increased over time, mainly during the 2000–2015 period. Third, the variation of the markdown is mostly driven by the input shares and not the output elasticities of the individual industries used to obtain our estimates. Fourth, the aggregate markdown moves closely with output, rising during expansions and contracting during recessions. This cyclical behavior suggests that aggregate wages may deviate more from their competitive benchmark during expansions. Understanding the behavior of aggregate wage markdown through the lens of macroeconomic models is a fruitful area for further research, and the framework developed in Fernández-Villaverde et al. (2021) appears to be a promising starting point.

References

Akcigit, Ufuk, Wendjie Chen, Federicao J. Díez, Romain Duval, Philipp Engler, Jiayue Fan, Chiara Maggi, Marina M. Tavares, Daniel Schwarz, Ippei Shibata, and Carolina Villegas-Sánchez. 2021. "Rising Corporate Market Power: Emerging Policy Issues." IMF Staff Discussion Note SDN/21/01, IMF.

Alpanda, Sami. 2019. "Business Cycles with Oligopsonistic Competition in Labor Markets." working paper, University of Central Florida.

Azar, José, Ioana Marinescu, and Marshall Steinbaum. 2020. "Labor Market Concentration." *Journal of Human Resources*.

Benmelech, Efraim, Nittai K. Bergman, and Hyunseob Kim. 2019. "Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?" working paper, Northwestern University.

Berger, David, Kyle Herkenhoff, and Simon Mongey. 2019. "Labor Market Power." manuscript, Duke University.

Bils, Mark. 1987. "The Cyclical Behaviour of Marginal Cost and Price." *American Economic Review*, 77(5): 838–855.

Bils, Mark, Peter J. Klenow, and Benjamin A. Malin. 2018. "Resurrecting the Role of the Product Market Wedge in Recessions." *American Economic Review*, 108(4-5): 1118–1146.

Blinder, Alan. 2014. "Tech Giants Settle Antitrust Hiring Suit." URL: https://www.nytimes.com/2021/07/01/podcasts/the-daily/ncaa-college-sports-student-athletes.html?rref=vanity.

Burstein, Ariel, Vasco M. Carvalho, and Basile Grassi. 2020. "Bottom-up Markup Fluctuations." manuscript, UCLA.

Caldwell, Sydnee, and Emily Oehlsen. 2018. "Monopsony and the Gender Wage Gap: Experimental Evidence from the Gig Economy." working paper, MIT.

Card, David, Ana Rute Cardoso, Joerg Heining, and Patrick Kline. 2018. "Firms and labor market inequality: Evidence and some theory." *Journal of Labor Economics*, 36(S1): S13–S70.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.

Davis, S. J., C. Grim, J. Haltiwanger, and M. Steitwieser. 2013. "Electricity Unit Value Prices and Purchase Quantities: U.S. Manufacturing Plants, 1963-2000." *Review of Economics and Statistics*, 95(4): 1150–1165.

Depew, Briggs, and Todd A Sørensen. 2013. "The elasticity of labor supply to the firm over the business cycle." *Labour Economics*, 24 196–204.

Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri. 2021. "Monopsony in Online Labor Markets." *American Economic Review: Insights*.

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin. 2000. "Optimal monetary policy with staggered wage and price contracts." *Journal of Monetary Economics*, 46 281–313.

Fernald, John G., and J. Christina Wang. 2016. "Why Has the Cyclicality of Productivity Changed? What Does It Mean?" *Annual Review of Economics*, 8(1): 465–496.

Fernández-Villaverde, Jesús, Federico Mandelman, Yang Yu, and Francesco Zanetti. 2021. "The "Mathew Effect" and Market Concetration: Search Complementarities and Monopsony Power." manuscript, University of Pennsylvania.

Galí, Jordi, and Luca Gambetti. 2009. "On the Sources of the Great Moderation." *American Economic Journal: Macroeconomics*, 1(1): 26–57.

Galí, Jordi, and Mark Gertler. 1999. "Inflation dynamics: A structural econometric analysis." *Journal of Monetary Economics*, 44(2): 195–222.

Galí, Jordi, Mark Gertler, and J. David López-Salido. 2007. "Markups, Gaps, And The Welfare Costs of Business Fluctuations." *Review of Economics and Statistics*, 89(1): 44–59.

Galí, Jordi, and Thijs van Rens. 2020. "The Vanishing Procyclicality of Labour Productivity." *The Economic Journal*, 131(633): 302–326.

Goolsbee, Austan, and Chad Syverson. 2019. "Monopsony Power in Higher Education: A Tale of Two Tracks." working paper, NBER.

Griliches, Zvi, and Jacques Mairesse. 1999. *Production Functions: The Search for Identification*. 169–203, Econometric Society Monographs: Cambridge University Press.

Hafiz, Hiba, and Nathan Miller. 2021. "Competitive Edge: Big Ag's monopsony problem: How market dominance harms U.S. workers and consumers."Technical report, Washington Center for Equitable Growth.

Hershbein, Brad, Claudia Macaluso, and Chen Yeh. 2020. "Monopsony in the U.S. Labor Market." manuscript, W. E. Upjohn Institute.

Hirsch, Boris, Elke J Jahn, and Claus Schnabel. 2018. "Do employers have more monopsony power in slack labor markets?" *ILR Review*, 71(3): 676–704.

Kroft, Kory, Yao Luo, Magne Mogstad, and Bradley Setzler. 2020. "Imperfect Competition and Rents in Labour and Products Markets: The Case of the Construction Industry." Working Paper 27325, NBER.

Lamadon, Thibault, Magne Mogstad, and Bradley Setzle. 2019. "Competition, Compensating Differentials, and Rent Sharing in the U.S. Labor Market." Working Paper 2019-84, Becker Friedman Institute.

Manning, Alan. 2020. "Why we need to do something about the monopsony power of employers." *The LSE US Center Blog*.

Manning, Alan. 2021. "Monopsony in Labor Markets: A Review." *ILR Review* 3–26.

Morlacco, Monica. 2019. "Market power in input markets: Theory and evidence from french manufacturing." *From Micro to Macro: Market Power, Firms' Heterogeneity and Investment, Luxembourg. EIB*.

Naidu, Suresh, Eric A Posner, and Glen Weyl. 2018. "Antitrust Remedies for Labor Market Power." *Harvard Law Review*, 132 536–601.

Nekarda, Christopher, and Valerie Ramey. 2021. "The Cyclical Behaviour of the Price-Cost Markup." *Journal of Money, Credit and Banking (50th Anniversary Edition)*.

Prager, Elena, and Matt Schmitt. 2021. "Employer consolidation and wages: Evidence from hospitals." *American Economic Review*, 111(2): 397–427.

Ravn, M., and H. Uhlig. 2002. "On Adjusting the HP Filter for the Frequency of Observations." *Review of Economics and Statistics*, 84(3): 71–75.

Robinson, Joan. 1933. *The Economics of Imperfect Competition*.: Macmillan and Co..

Rotemberg, Julio J., and Michael Woodford. 1999. "The Cyclical Behaviour of Prices and Costs." In *Handbook of Macroeconomics*. eds. by John B. Taylor, and Michael Woodford, Vol. 1B: Amsterdam: North-Holland, 1051–1135.

Rubens, Michael. 2021. "Market Structure, Oligopsony Power, and Productivity." *Working paper*.

Serrato, Juan Carlos Suárez, and O. Zidar. 2016. "Who Benfits from State Corporate Tax Cuts? A Local Labor Markets Approach with Heterogeneous Firms." *American Economic Review*, 106(9): 2582–2624.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Sokolova, Anna, and Todd Sorensen. 2021. "Monopsony in Labor Markets: A Meta-Analysis." *ILR Review* 27–55.

Streitfeld, David. 2014. "Tech Giants Settle Antitrust Hiring Suit." URL: https://nyti.ms/1mEK3tE.

Summers, Lawrence, and Anna Stansbury. 2020. "Opinion: U.S. workers need more power." URL: https://www.washingtonpost.com/opinions/2020/06/28/us-workers-need-more-power/.

Figure 1. Trends in the aggregate wage markdown

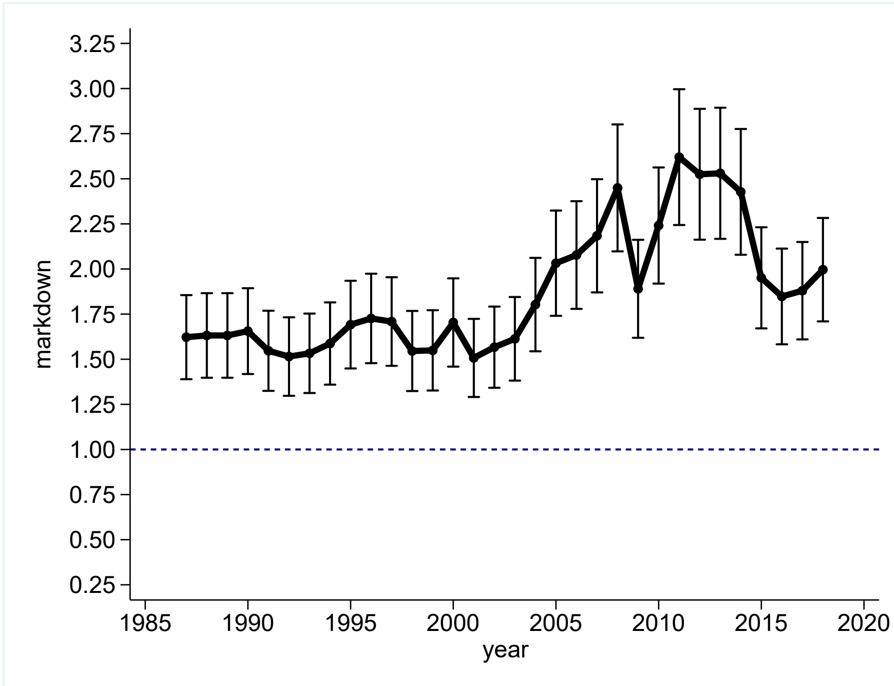

Note: We show our baseline aggregate wage markdown estimates using mean output elasticities. Hence, all the time variation in the markdowns is due to the input shares. The output elasticities are estimated using the translog production function in (5). We also show 95% confidence intervals calculated using the delta method.

Figure 2. Trends in the aggregate wage markdown

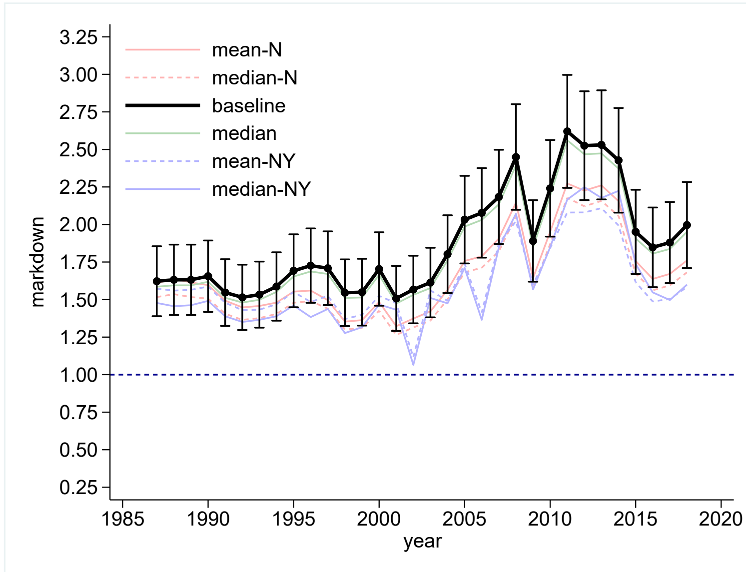

A. Robustness to summary statistics of output elasticities

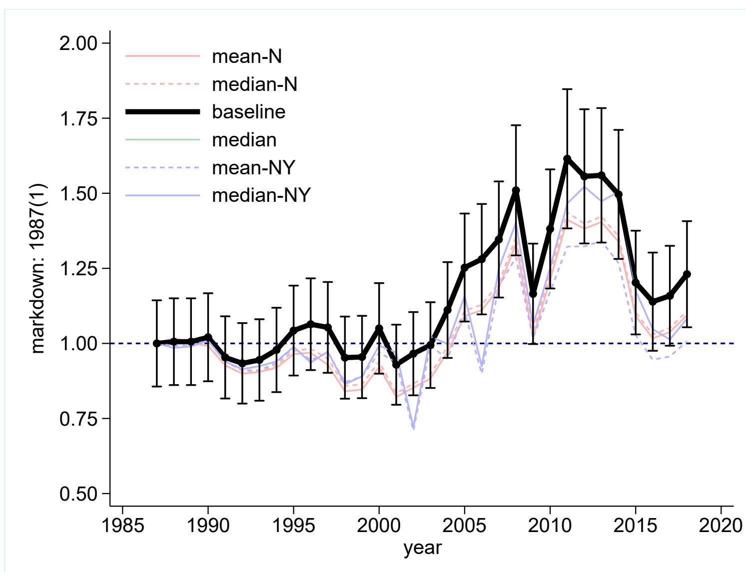

B. Aggregate markdowns as indices (1987=1)

Note: In panel A, we show alternative estimates of the aggregate markdown using the following elasticities: mean by NAICS sector (**mean-N**), median by NAICS sector (**median-N**), median (**median**), mean by NAICS sector and year (**mean-NY**), median by NAICS sector and year (**median-NY**). In the first two cases, we introduce *cross-sectional* variation in the output elasticities. In the last two, we introduce both *cross-sectional* and *time* variation in the output elasticities. In panel B, we replicate the series from panel A scaled such that their values are equal to 1 in 1987, which is the first year in our sample. In both panels, we also show 95% confidence intervals of our baseline estimates calculated using the delta method. We use **baseline** to refer to the line shown in Figure 2.

Figure 3. Cyclical aggregate wage markdown and labor supply elasticity

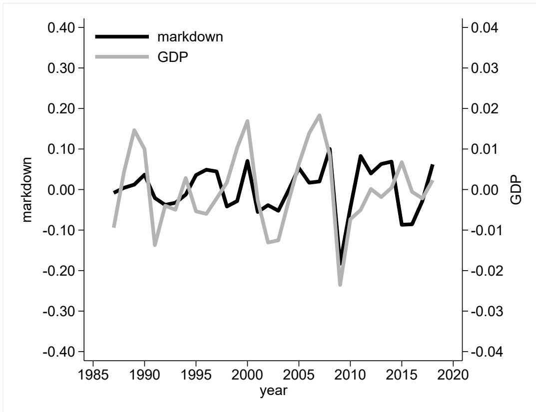

A. Aggregate wage markdown

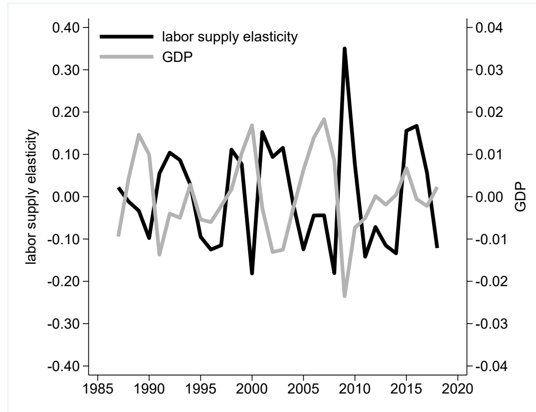

B. Labor supply elasticity

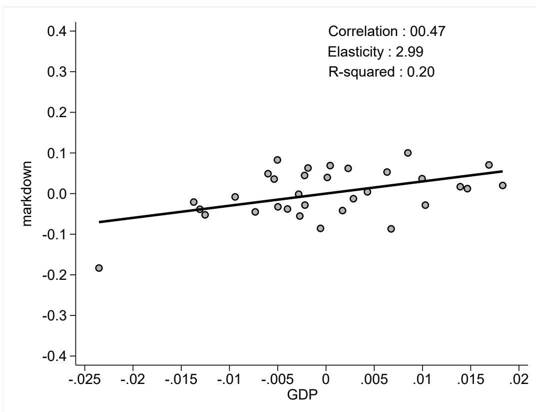

C. Aggregate wage markdown

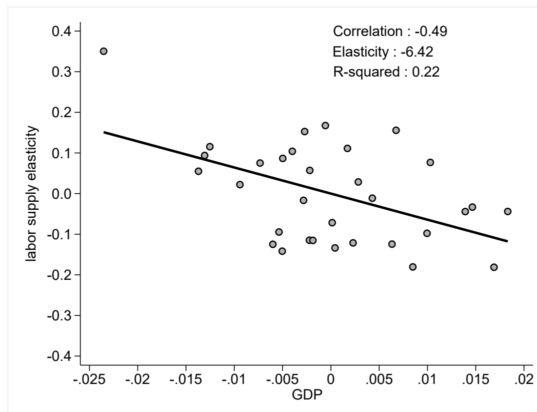

D. Labor supply elasticity

Note: In panels A and B, we show the time series of the cyclical component of real GDP along with its counterpart of the aggregate markdown and labor supply elasticity assuming materials serve as the reference input not subject to buyer power. In both cases, we extract the cyclical component of the series using the HP filter. In panels C and D, we show a scatterplot of the series' cyclical components and report, the correlation, slope coefficient and R-squated of the fitted regression line.