

3-4-2022

## Monopsony in Manufacturing

Chen Yeh

*Federal Reserve Bank of Richmond*, [chen.yeh@rich.frb.org](mailto:chen.yeh@rich.frb.org)

Claudia Macaluso

*Federal Reserve Bank of Richmond*, [claudia.macaluso@rich.frb.org](mailto:claudia.macaluso@rich.frb.org)

Brad J. Hershbein

*W.E. Upjohn Institute for Employment Research*, [hershbein@upjohn.org](mailto:hershbein@upjohn.org)

Upjohn Author(s) ORCID Identifier:

 <https://orcid.org/0000-0002-2534-8164>

Follow this and additional works at: [https://research.upjohn.org/up\\_policybriefs](https://research.upjohn.org/up_policybriefs)



Part of the [Labor Economics Commons](#)

---

### Citation

Yeh, Chen, Claudia Macaluso, and Brad J. Hershbein. 2022. "Monopsony in Manufacturing." Policy and Research Brief. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research. <https://doi.org/10.17848/pb2022-44>

This title is brought to you by the Upjohn Institute. For more information, please contact [repository@upjohn.org](mailto:repository@upjohn.org).

# RESEARCH BRIEF

## Monopsony in Manufacturing

Chen Yeh, Claudia Macaluso, and Brad Hershbein

### BRIEF HIGHLIGHTS

- *Recent research has suggested that labor markets may not be fully competitive, meaning workers are paid less than the value they add.*
- *We develop a new approach to measuring and estimating “markdowns”—the gaps between wages and the value a worker adds—in the U.S. manufacturing sector.*
- *Drawing on confidential survey data on U.S. manufacturers, we find that the typical worker at the average manufacturing plant earns 65 cents on each marginal dollar generated.*
- *The degree of monopsony varies greatly across and within manufacturing industries, with markdowns generally higher for larger and more productive plants.*
- *For U.S. manufacturing as a whole, markdowns shrank between 1980 and 2000 but have risen sharply more recently, suggesting manufacturing workers are increasingly underpaid relative to their value.*

For additional details, see the full working paper at [https://research.upjohn.org/up\\_workingpapers/364](https://research.upjohn.org/up_workingpapers/364).

Most textbook economic models assume that labor markets are perfectly competitive. In such a case, workers’ pay equals the marginal contributions to their employers’ revenues, or what economists call marginal revenue product of labor (MRPL). For example, a worker who generates \$20 more in employer revenues per hour should be compensated \$20 per hour. If an employer paid this worker less than that, the theory says the worker could simply switch to another employer who is willing to pay slightly more and still make a profit.

However, many researchers and policymakers have recently grown concerned that the textbook model is not accurate and that employers’ market power over workers has increased. This power could arise from ways in which employers may restrict their workers from looking for other job opportunities, such as noncompete agreements (Starr, Prescott, and Bishara 2021), a growing reluctance for workers to change locations (Molloy, Smith, and Wozniak 2014), or business mergers that result in fewer employers competing for workers (Prager and Schmitt 2021). However, it has been remarkably challenging for researchers to find direct and general evidence on labor market power, and this has complicated the policy debate. The reason for this lack of evidence is simple: While wages are observable in some data sets, firms’ MRPLs are hard to measure. Without the latter, it is almost impossible to determine whether MRPLs are equal to wages, as is predicted by perfect competition.

In a [recent paper](#), we develop a new technique to show that employer market power is substantial and widespread in the U.S. manufacturing sector. Using this approach, we find that over the past several decades, a worker in the average manufacturing plant receives only 65 cents on each dollar generated on the margin. Furthermore, we construct a novel aggregate measure for labor market power, across all manufacturing plants and workers in the United States, to understand how this phenomenon has evolved over time. We find that employers’ market power actually decreased between the late 1970s and the early 2000s, but that it sharply increased in the decade or so afterward. Our study thus provides direct evidence that many workers are paid less than their “fair share,” complementing earlier studies (as reviewed in Sokolova and Sorensen 2020) which often relied on indirect methods. Despite recent gains in worker wages since the pandemic, employer market power likely remains considerable.

### Measuring Employers’ Labor Market Power

Under a perfectly competitive labor market, marginal gains in employer revenues generated by workers should go fully to workers. The intuition for this is straightforward: Employers that don’t do this would see many or most of their workers depart to competitors.

The presence of labor market power, on the other hand, implies that employers can withhold some of these marginal gains. Hence, an employer with labor market power compensates its workers with wage rates below their MRPLs. Typically, economists have expressed labor market power through the gap (or ratio) between a firm’s MRPL and the wages paid to workers, also known as the markdown.

The main problem with measuring markdowns is that their components are often not directly observable. MRPLs are never reported (and are, indeed, hard to measure), and

## Monopsony in Manufacturing

**An employer with labor market power compensates its workers with wage rates below the marginal revenue they add.**

even hourly wages are not always available, even in firm-level data. However, we show that by making some relatively weak assumptions about a firm's production function—how it combines labor, capital, and other inputs to make products—we can estimate plant-level markdowns. To do so, one requires a few additional pieces of information. The first is the revenue share of each input (that is, the effective spending on that input as a share of a plant's revenues). Many firm-level data sets have revenue share of inputs readily available, including the confidential Censuses of Manufacturing and Surveys of Manufacturing we use for our analysis. The second component needed is a measure of a plant's "output elasticities"—the percentage increase in physical output when a particular input increases by 1 percent. Because these quantities are not available, we estimate them using the rest of the data and a flexible form for a plant's production function, adapted from studies of industrial organization.

This "production function" approach also allows us to distinguish *labor* market power from *product* market power. In other words, a firm or plant could both pay workers below their MRPL (a wage markdown) and charge a price for its output product above the competitive rate (a product markup). Our approach allows us to isolate wage markdowns from product markups, an advantage over some other methods used to estimate labor market power. Moreover, it allows us to characterize labor market power for each plant (or firm) in the U.S. manufacturing sector.

### Labor Market Power in the U.S. Manufacturing Sector

The key takeaway from our main results, summarized in the table below, is that labor market power is substantial for U.S. manufacturing plants. At the average plant, workers

**Table 1 Estimated Plant-Level Markdowns in U.S. Manufacturing**

Industry	Median	Mean	Typical range
Petroleum refining	0.42	0.39	0.26 – 0.78
Computers and electronics	0.44	0.39	0.28 – 0.67
Plastics and rubber	0.55	0.52	0.40 – 0.76
Food and kindred products	0.57	0.52	0.37 – 0.92
Paper and allied products	0.59	0.56	0.41 – 0.85
Chemicals	0.62	0.55	0.37 – 1.06
Lumber	0.65	0.62	0.47 – 0.91
Primary metals	0.69	0.67	0.50 – 0.98
Motor vehicles	0.73	0.70	0.54 – 1.01
Printing and publishing	0.74	0.67	0.47 – 1.16
Electrical machinery	0.76	0.71	0.52 – 1.11
Fabricated metals	0.80	0.76	0.60 – 1.05
Nonelectrical machinery	0.80	0.76	0.56 – 1.16
Miscellaneous manufacturing	0.83	0.80	0.62 – 1.12
Textiles	0.83	0.79	0.58 – 1.23
Furniture	0.87	0.86	0.66 – 1.24
Nonmetallic minerals	0.88	0.82	0.58 – 1.44
Apparel and leather	0.97	0.87	0.59 – 1.65
<b>All manufacturing</b>	<b>0.73</b>	<b>0.65</b>	<b>0.45 – 1.22</b>

NOTE: The numbers shown represent the fraction of a dollar of revenue generated by workers that is paid in wages at the median plant, the average plant, and in a range of plants within one standard deviation of the mean, all for different manufacturing industries. If workers are paid their full value added, the value should be 1.00. Values below 1.00 thus represent wage markdowns. Values can be above 1.00 if workers are paid above their revenue value added. The sample size underlying the estimates is approximately 1.4 million plant-year observations. The industries shown approximately follow 3-digit NAICS categories. See the full paper for details of the estimation strategy.

SOURCE: Authors' calculations using Annual Surveys of Manufacturing and Censuses of Manufacturing, 1976–2014.

**Labor market power is substantial for U.S. manufacturing plants: on average, workers collect only 65 cents of every dollar they generate on the margin for their employers.**

collect only 65 cents of every dollar they generate on the margin for their employers. Because the distribution is skewed, the markdown at the median plant is not quite as severe, but it still implies that half of plants pay their workers below 73 cents on the dollar.

Labor market power is extensive across manufacturing industries, with the mean (and median) plant paying wages less than MRPL in each industry. However, these markdowns vary considerably across industries; they are greatest for plants in the petroleum refining and the computer and electronics industries, and they are smallest in the apparel and leather industry. Furthermore, markdowns also vary greatly *within* each industry, as evidenced by the last column, which shows the middle range across plants. While it is certainly plausible that industry-level factors (such as unionization rates) can explain variation in markdowns, our results indicate that factors specific to individual plants can be important.<sup>1</sup>

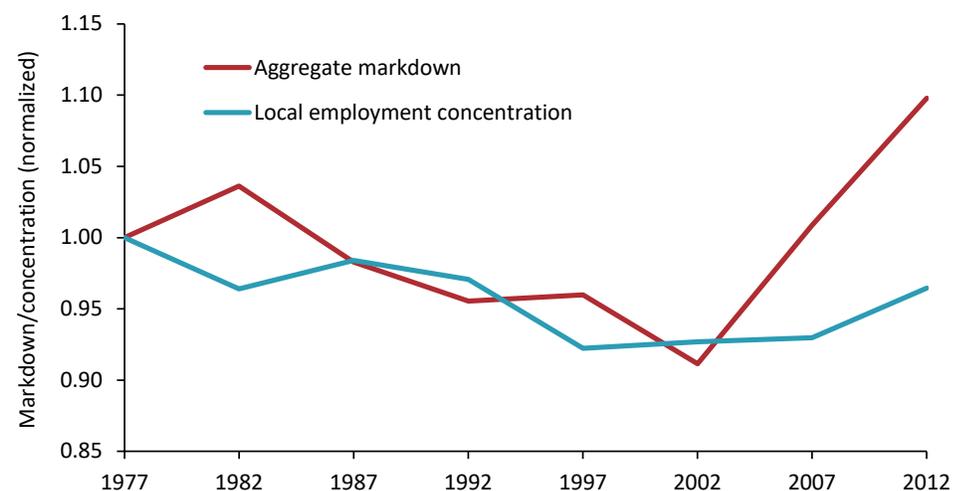
### Aggregate Trends in Labor Market Power

The markdowns described above apply at the plant level over our whole sample period of 1976 through 2014. As mentioned earlier, however, some policymakers and academics have suggested that labor market power has increased over time for the whole economy. Despite our markdown estimates by plant and year, it is not obvious how to combine these plant-level measures to obtain a statistic that reflects the aggregate economy—a simple (or even weighted) average doesn't quite work.

To calculate aggregate markdown, we note that plant-level markdowns are a function of output elasticities and revenue shares. When we apply a similar logic to the aggregate markdown, we can derive a specific equation that properly weights the plant-level markdowns, accounting for their differences in output elasticities, revenue shares, and total revenues. This equation for the aggregate markdown is flexible, in that it does not depend on *how* plants are assumed to combine inputs to make a product, nor on any particular mechanism for employer market power.

Applying this method, Figure 1 shows how the aggregate markdown for the U.S. manufacturing sector has evolved over time. (We normalize the aggregate markdown in the first year, 1977, at a value of 1.00, with larger values intuitively representing increases in employer market power.) After moderately declining in the late twentieth century, the

**Figure 1 Aggregate Markdowns and Local Employment Concentration**



NOTE: The red line plots our measure of the aggregate markdown, normalized to have a value of 1.00 in 1977. The turquoise line plots an aggregate of local labor market concentration, also normalized to have a value of 1.00 in 1977. Details of the construction of both measures are in the full paper (Yeh, Macaluso, and Hershbein 2022). SOURCE: Authors' calculations using Censuses of Manufacturing, 1977–2012.

## Monopsony in Manufacturing

**After moderately declining in the late 20th century, the aggregate markdown reversed course, and employer market power sharply increased in the new millennium.**

aggregate markdown reversed course, and employer market power sharply increased in the new millennium.

Because of the difficulties of directly measuring markdowns, previous studies have drawn conclusions on employers' market power based on the concentration of employment across firms. Concentrated labor markets—those with only a few firms—naturally limit alternative employment options for workers, which could lead to employers exercising their labor market power. Although this argument is attractive, employer concentration does not *necessarily* lead to a wage markdown. Indeed, we find that markdowns and concentration at the labor market level are only weakly correlated, and that their time paths at the national, aggregate level also diverge, especially recently (as illustrated by the turquoise line in Figure 1). Consequently, it is important to obtain independent estimates of markdowns to draw meaningful conclusions about labor market power.

### Conclusion

Our recent research shows that labor market power is pervasive in the U.S. manufacturing sector. Workers are not fully compensated for their marginal contributions to their employers' revenues. Instead, a worker at the average U.S. manufacturing plant receives only 65 cents for every dollar of revenue generated on the margin. While labor market power is widespread, there is tremendous variation at plants both across and within industries.

Moreover, our aggregate measure of markdowns for the U.S. manufacturing sector as a whole indicates that employers' market power has switched course over time, falling between the late 1970s and early 2000s but then sharply rising over the next decade or so. This suggests that rising employer market power is unlikely to be the driver behind the declining share of total income going to labor, which started its [downward trend](#) decades earlier.

*Note: The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.*

### Note

1. In the full paper (Yeh, Macaluso, and Hershbein 2022), we examine relationships between the magnitude of markdowns and plants' age, size, and productivity. While plant age and productivity do not appear to be strongly related to the size of the markdown, plant size does. In particular, the larger the plant in terms of its share of employment in the local labor market (defined by county and industry), the higher the markdown. Intuitively, this should make sense: workers' outside options are more limited whenever the local labor market is controlled by only a handful of plants.

### References

- Molloy, Raven, Christopher L. Smith, and Abigail K. Wozniak. 2014. "[Declining Migration within the U.S.: The Role of the Labor Market.](#)" NBER Working Paper No. 20065. Cambridge, MA: National Bureau of Economic Research.
- Prager, Elena, and Matt Schmitt. 2021. "[Employer Consolidation and Wages: Evidence from Hospitals.](#)" *American Economic Review* 111(2): 397–427.
- Sokolova, Anna, and Todd Sorensen. 2020. "[Monopsony in Labor Markets: A Meta-Analysis.](#)" *ILR Review* 74(1): 27–55.
- Starr, Evan P., J.J. Prescott, and Norman D. Bishara. 2021. "[Noncompete Agreements in the U.S. Labor Force.](#)" *Journal of Law and Economics* 64(1): 53–84.
- Yeh, Chen, Claudia Macaluso, and Brad Hershbein. 2022. "[Monopsony in the U.S. Labor Market.](#)" Upjohn Institute Working Paper No. 364. Kalamazoo, MI: W.E. Upjohn Institute for Employment Research.



W.E. Upjohn Institute for  
Employment Research



@UpjohnInstitute

WEBSITE

upjohn.org