




2020

Labor Market Monopsony and Wage Inequality: Evidence from Online Labor Market Vacancies

Samuel I. Thorpe
Williams College, sit2@williams.edu

Follow this and additional works at: <https://digitalcommons.iwu.edu/uer>

 Part of the [Antitrust and Trade Regulation Commons](#), [Economic Policy Commons](#), [Income Distribution Commons](#), and the [Labor Economics Commons](#)

Recommended Citation

Thorpe, Samuel I. (2020) "Labor Market Monopsony and Wage Inequality: Evidence from Online Labor Market Vacancies," *Undergraduate Economic Review*: Vol. 17 : Iss. 1 , Article 11.

Available at: <https://digitalcommons.iwu.edu/uer/vol17/iss1/11>

This Article is protected by copyright and/or related rights. It has been brought to you by Digital Commons @ IWU with permission from the rights-holder(s). You are free to use this material in any way that is permitted by the copyright and related rights legislation that applies to your use. For other uses you need to obtain permission from the rights-holder(s) directly, unless additional rights are indicated by a Creative Commons license in the record and/ or on the work itself. This material has been accepted for inclusion by faculty at Illinois Wesleyan University. For more information, please contact digitalcommons@iwu.edu.

©Copyright is owned by the author of this document.

Labor Market Monopsony and Wage Inequality: Evidence from Online Labor Market Vacancies

Abstract

This paper estimates the effects of employer labor market power on wage inequality in the United States. I find that inequality as measured by interdecile range is 23.7% higher in perfectly monopsonistic labor markets than in perfectly competitive markets, even when controlling for commuting zone and occupation fixed effects. I also decompose these results into 50/10 and 90/50 ratios, finding much larger impacts on inequality among low earners. These results suggest that monopsony power has significant and policy-relevant impacts on wage inequality, and particularly harms the lowest earning subsets of the labor force.

Keywords

Labor Economics, Market Power, Monopsony, Inequality, Public Policy

Cover Page Footnote

I would like to express my deepest gratitude to Professor Owen Thompson for his guidance and support throughout the research process. This paper would not exist in its current form without his help. I would also like to thank Professors Jon Bakija and David Zimmerman for invaluable lessons on STATA and econometrics, and Alex Horne and Anna Kankkunen for their helpful comments. Special thanks as well to Professor Marshall Steinbaum, who provided essential data and crosswalks at an early stage of the project.

1. Introduction

Important frictions are present in the labor market. As Manning (2003) points out, “people go to the pub to celebrate when they get a job rather than greeting the news with the shrug of the shoulders that we might expect if labour markets were frictionless... [and] go to the pub to drown their sorrows when they lose their job rather than picking up another one straight away.” And since at least Manning and Boal and Ransom (1997), there has been a strong theoretical grounding for the idea that labor market monopsony is an important source of those frictions. The impacts of monopsony are theoretically clear: greater monopsony leads to more unemployment and lower wages. More precisely, employers under monopsony face an upward-sloping labor supply curve, and if they are profit-maximizing, will choose to set wages and employment at levels lower than those present in perfect competition. This relationship is fundamentally exploitative: workers under monopsony conditions are underpaid relative to the value they create for the firm.

Many recent empirical papers have attempted to quantify these claims, estimating the level of monopsony in specific labor markets and its impact on wages (e.g. Benmelech et al., 2020, Hershbein et al., 2019, Dube et al., 2020). Nearly all have found large and statistically significant negative impacts of monopsony power on wages, as suggested by theory, and provided evidence that an overwhelming share of U.S. labor markets exhibit at least some significant degree of monopsony. Yet despite this flood of recent research, very little empirical work has focused on how monopsony affects the wages of different subsets of the labor market in different ways or how monopsony creates and interacts with wage inequality. Understanding this issue is crucial for deriving appropriate policy responses, particularly in occupations and geographic areas which exhibit high levels of monopsony.

This paper aims to fill part of this gap in the literature. I introduce a model for why inequality and monopsony might be causally linked, and test the predictions of this model using a combination of high-quality data on monopsony power for all commuting zones and occupations in the United States for the year 2016 and individual microdata from the American Community Survey. I find that increased monopsony power in a particular labor market is associated with a significant and policy-relevant increase in the level of wage inequality in that market. I further break down these results into estimates for the lower and higher ends of the income distribution, finding that

increased inequality is primarily a result of falling wages among the lowest earners in the labor force. These findings provide evidence that the dual trends in the modern American economy of increasing wage inequality and increasing market power may be causally linked.

2. Contribution to the Literature

There is a substantial literature on increasing market concentration and markups alongside the decline of the labor share of national income in the United States (e.g. Karabounis and Neiman 2014, Autor et al. 2020, De Loecker et al. 2020). A related literature has documented the presence of employer labor market power in various markets and analyzed its wage effects under a variety of data sources, monopsony measures, and market definitions. Naidu et al. (2018) directly tie monopsony power to the falling labor share, estimating that high levels of monopsony reduce the labor share by as much as 22%, while Grullon et al. (2019) find that U.S. industries are becoming more concentrated over time and that this concentration results in higher profit margins but not increased operational efficiency. Azar et al. use online vacancies data to estimate concentration and wage effects and find that high concentration is associated with a significant decrease in hourly wages, particularly concerning given that more than 60% of U.S. labor markets are highly concentrated by their definition. Hershbein et al. (2019) use data from vacancies and total employment along with a variety of market definitions to calculate the wage effects of monopsony, finding that a 1% increase in local labor market concentration is associated with a 0.14% decrease in average hourly wages; Benmelech et al. (2020) find similar results and note that higher rates of unionization are associated with lower monopsony wage penalties. Perhaps the most surprising results come from Dube et al. (2020), who find that a substantial degree of monopsony power is present even in online labor markets with seemingly low switching and search costs. The authors suggest that based on these results, monopsony power is likely present in the vast majority of labor markets, even those which should theoretically be highly competitive. Taken together, this evidence suggests that monopsony power is a major factor in the modern American economy, and employer profit margins tend to be higher at the expense of workers in high-monopsony areas.

The implications of monopsony theory are broader than just negative wage effects: monopsony theory has been used as a tool to explain issues from worker exploitation in professional sports (Kahn, 2000) to the small effects of minimum wages on employment relative to those expected in neoclassical theory (Dustmann et al., 2019). Manning (2020) provides a review of the modern

literature and its implications for a wide variety of topics. But despite all the recent interest in monopsony, there remains a paucity of studies on the effects of monopsony power on wage inequality. To the best of my knowledge, only two papers have attempted to directly quantify the relationships between monopsony and inequality. Webber (2015) uses data from the Longitudinal Employer Household Dynamics to compute firm-level measures of monopsony power and analyzes their effects on the earnings distribution. He finds positive and statistically significant effects of employer labor market power on inequality, and that the negative impact of firm market power is strongest in the lower half of the earnings distribution. Rinz (2018) uses data from the Longitudinal Business Database and Form W-2 to calculate monopsony measures, finding somewhat idiosyncratic results on local market concentration. He finds modest but positive effects of labor market concentration on inequality and confirms the results in Webber on who is most affected, noting that earnings fall most at the bottom when concentration increases. Neither of these papers use data on vacancies, my preferred specification of the effective monopsony conditions facing workers.

3. Model

There is no immediately obvious reason to assume that monopsony would be associated with increased inequality. In its simplest form, monopsony theory simply posits that most employers face an upward-sloping labor supply curve. Under such conditions and assuming a Cournot model of oligopsony, a profit-maximizing firm faces the problem

$$\text{Max } R(L) - w(L + L^*)L \quad (1)$$

where L denotes employment for the particular firm, L^* denotes employment for all other firms in the market, $R(L)$ is revenue as a function of employment, and $w(L + L^*)$ is the inverse labor supply function where wages are a function of employment. The monopsonistic firm pays workers less than their marginal revenue product (MRP), and this can be written as a firm-specific Pigouvian “rate of exploitation” E such that

$$E = \frac{MRP - w}{w} \quad (2)$$

A more accurate way to model the firm's problem under monopsony is allowing for differentiated workers of particular types, each contributing differently to revenue on the margin. With differentiated workers, the firm's problem is similar but slightly distinct:

$$\text{Max} \sum_{t=1}^T R(L_t) - w_t(L_t + L_t^*)L_t \quad (3)$$

where L_t is employment for workers of type t , $w_t(L_t + L_t^*)$ is the inverse labor supply function for workers of type t , and T is the number of distinct types of workers. For each type of worker, our profit-maximizing firm faces a different choice: they must set employment and wages equal to the profit-maximizing level for that particular type, which will differ depending on the characteristics of workers of that type, including their propensity to quit. The inverse labor supply function $w_t(L_t + L_t^*)$ is type-specific in this model, which means that the rate of exploitation is also type-specific:

$$E_t = \frac{MRP_t - w_t}{w_t} \quad (4)$$

The fact that the rate of exploitation can vary based on worker type is a crucial insight which allows us to model the inequality effects of monopsony. *Ceteris paribus*, if the rate of exploitation increases faster on average for low-wage workers than high-wage workers as monopsony increases, then monopsony power will tend to result in greater wage inequality. There are a number of reasons to believe that this might be the case, particularly in the United States. If we incorporate the ability of workers to quit their jobs to the model of exploitation in equation (4), and define a set of high-wage workers and low-wage workers (for which the cutoffs are largely arbitrary, except that the sets must not overlap and all high-wage workers must earn more than the highest-earning low-wage worker), then we can model the average rate of exploitation of high-wage workers as

$$\bar{E}_{hw} = \frac{\overline{MRP}_{hw} - \bar{w}_{hw}(\bar{L}_{hw}, \bar{L}_{hw}^*, \bar{q}_{hw})}{\bar{w}_{hw}(\bar{L}_{hw}, \bar{L}_{hw}^*, \bar{q}_{hw})} \quad (5)$$

where the average quit rate \bar{q}_{hw} at a given wage is a function of various characteristics associated with the group of high-wage workers. The average rate of exploitation for low-wage workers will look essentially the same:

$$\bar{E}_{lw} = \frac{\overline{MRP}_{lw} - \bar{w}_{lw}(\bar{L}_{lw}, \bar{L}_{lw}^*, \bar{q}_{lw})}{\bar{w}_{lw}(\bar{L}_{lw}, \bar{L}_{lw}^*, \bar{q}_{lw})} \quad (6)$$

It is intuitively clear that the average quit rates \bar{q}_{lw} and \bar{q}_{hw} are inversely associated with wages. If a worker is unable or unwilling to quit as their wages and conditions worsen, a monopsonistic profit-maximizing employer will cut wages more deeply, all else equal. If workers are more willing to quit as wages stagnate or fall, a monopsonistic employer will be forced to keep wages higher to maintain their profit-maximizing rate of employment. Moreover, as the wage decreases below MRP, the rate of exploitation increases mechanically. Thus, if all else equal, low-wage workers have a lower propensity to quit than high-wage workers under conditions of monopsony ($\bar{q}_{lw} < \bar{q}_{hw}$), low-wage workers will be more highly exploited. The relevant question then becomes what structural frictions impede low-wage workers from quitting their jobs under conditions of monopsony.

In the U.S., it is likely that such frictions stem the lack of an adequate social safety net along with the lower occupational and geographic mobility of low-wage workers. The lack of safety nets and adequate unemployment insurance, coupled with the well-publicized fact that the majority of low-income American families lack adequate savings to pay for even an \$400 emergency expense (Federal Reserve, 2018), mean that low-wage workers rarely have the resources to quit their job without another job offer close at hand. In high-monopsony labor markets, these new jobs will be few and far between. Moreover, this effect is amplified by the uniquely American system of employer-based health insurance, which leads most individuals who quit their jobs to be at least temporarily at risk of catastrophic financial burdens if they get sick. The relative lack of affordable high-quality postsecondary education in the United States means that it is hard for workers without the financial means to change occupations, particularly later in their careers, so they are effectively “hit harder” by existing levels of monopsony than those with more financial resources. And the geographic immobility of low-wage workers (e.g. Purcell, 2020, Rodgers and Rodgers, 2000) further magnifies this effect, making it more challenging for low-wage workers to escape high-

monopsony labor markets and thus decreasing their quit rates and increasing their level of exploitation. If any of these observations are relevant to the decision-making of low-wage workers, the models of exploitation in equations (5) and (6) suggest that increased monopsony will also be associated with increased inequality. This model makes the additional prediction that much of this increase in inequality is caused by reduced wages at the bottom of the earnings distribution. In Section 5, we test these predictions and find them to be supported empirically.

4. Data and Empirical Strategy

This paper relies on a combination of data on monopsony power at the commuting zone-occupation level (from Azar, Marinescu, Steinbaum, and Taska, 2020) along with individual microdata on wages and demographic characteristics from the American Community Survey (ACS). Azar et al. use the near-universe of online job postings for the year 2016 collected by Burning Glass Technologies to construct measures of monopsony power for every occupation (occupations defined by BLS SOC code) and commuting zone pair throughout the United States. They calculate monopsony power using the Herfindahl-Hirschman Index (HHI) of market concentration, defined in market m and time t as

$$HHI_{m,t} = \sum_{j=1}^J s_{j,m,t}^2 \quad (7)$$

where $s_{j,m}$ represents the market share of firm j in market m (Azar et al., 2020). Crucially, the market share estimate $s_{j,m}$ is based on the number of *vacancies*, rather than total share of existing employment, of a particular firm in a particular market. The share of vacancies is arguably a better way to gauge the labor market faced by an individual worker than the share of total employment; a worker without available vacancies in their occupation and area effectively faces conditions of perfect monopsony, even if other firms employ a large share of workers but are not hiring new ones. The ACS data comes from the 2016 version of the American Community Survey; this data includes information on individual income from wages along with a variety of demographic characteristics. It also includes information on the occupation and county location of most individuals in the sample. I use the same crosswalk used by Azar et al. (2020) to convert from county to commuting zone, ensuring that my commuting zone definitions are identical to those in

the HHI data. I merge between the two datasets on the commuting zone-occupation level, using generalized (5-digit) SOC codes. See the Appendix for a discussion of why 5-digit SOC codes were preferred. I weight all analyses using the person-weights provided by ACS.

Table 1 shows summary statistics for the workers remaining in the sample after initial processing. HHI normally ranges from 0 to 10,000 but is rescaled linearly to fall between 0 and 1 for ease of interpretation (e.g. an HHI of 2500 becomes a rescaled HHI of 0.25.) Note that the means of the HHI and wage distribution statistics are implicitly weighted by the number of people working in each CZ-SOC combination, hence why the average HHI here is much lower than the average HHI of 4,378 in Azar et al., who consider each CZ-SOC combination equally. While the majority of *commuting zones* exhibit very high market concentration, the majority of *people* live in commuting zones with low to moderate market concentration, since cities and suburbs tend to both be more densely populated and have lower levels of monopsony compared to rural areas.

Table 1: Descriptive Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
HHI, rescaled, SOC5	1,105,094	.223	.277	0	1
HHI, rescaled, SOC6	631,446	.098	.179	0	1
Wage income	1,011,863	40,175	57,404	0	714,000
Unemployed (1=yes)	985,447	.048	.214	0	1
Self-employed (1=yes)	1,137,527	.088	.283	0	1
Top decile wage	691,871	97,273	76,168	18000	665,000
Median wage	691,871	45,299	29,719	1000	458,000
Bottom decile wage	691,871	16,305	13,050	1000	215,000
90/10	691,871	8.256	8.757	1	426
50/10	691,871	3.619	2.798	1	147.75
90/50	691,871	2.229	1.033	1	41.8

Note: Top, median, and bottom decile wages, and ratios constructed from them, are calculated at the commuting zone-occupation level. I exclude all individuals who are unemployed or below age 18. When calculating ratios, any cz-occupation pair with less than 5 respondents in it is dropped.

It is of some interest that there are both very low and extremely high interdecile ranges in the sample. A 90/10 ratio of 1 implies that wages are exactly the same between the top and bottom decile in that occupation-commuting zone combination, while a 90/10 ratio of 426 is clearly too high to accurately reflect the conditions faced by workers. Fortunately, the influence of these outliers is likely to be low. Only two interdecile ranges (out of 75,440 SOC5-CZ combinations) are equal to 1, while only 62 out of 75,440 are greater than 100. A wide variety of cutoffs (ranging

from 40 to 400) were used in preliminary analysis and did not influence the significance or direction of results. Because of this, I choose to keep the full sample rather than introducing an arbitrary cutoff for values that are “too high” or “too low.”

My main regression specification is

$$y_{ic} = \beta_0 + \beta_1 HHI_{ic} + X\beta_2 + v_i + v_c + \varepsilon_{ic}$$

where y_{ic} is the variable of interest (either average wages or average interdecile range of wages P_{90}/P_{10}) in occupation i and commuting zone c , HHI_{ic} is the HHI in occupation i and commuting zone c , v_i represents individual SOC fixed effects, v_c represents individual commuting zone fixed effects, and X is a vector of controls for race, sex, and educational attainment. These controls measure the average values of each category at the CZ-SOC5 level. With the included occupation and commuting zone FE, this model effectively compares workers with the same occupation across commuting zones that have varying levels of employer concentration for that occupation. Or stated conversely, it compares workers in the same commuting zone who are in occupations with varying levels of employer concentration. The identification strategy employed would be stronger with multiple years of data, which would allow for three dimensions of variation in HHI, but data limitations prevent this; access to the full Burning Glass Technologies data is prohibitively expensive. However, the data available remains comprehensive and high-quality, and several previously published papers by Azar, Marinescu, and Steinbaum have used the same data from just 2016 to estimate the overall wage impacts of monopsony.

5. Results

While not the primary focus of the current paper, I begin by estimating effects on overall wages, which is useful for comparisons to the existing literature. Table 2 displays these results. All models for wage estimation include commuting zone and occupation fixed effects; columns 2 and 4 include controls for race, education, and sex, and columns 3 and 4 only include full-time workers (those who work at least 40 hours a week and worked at least 40 weeks of the year.) These estimates of the effect of monopsony power on wages are similar to the existing literature. The coefficients on HHI are slightly larger in magnitude than Hershbein et al. (2019), and larger than Benmelech, Bergman, and Kim (2020) along with most of the estimates presented in the review of literature in Manning (2020). However, since these authors define monopsony power in terms

of absolute share of employment rather than vacancies, it is unsurprising that the results would be noticeably distinct. The results are more similar to those in Azar, Marinescu, and Steinbaum (2017), who use the same HHI data but wage data from online job postings.

Table 2: The Effect of Monopsony Power on Overall Wages

	Model 1	Model 2	Model 3	Model 4
HHI, rescaled	-14327.95*** (348.83)	-12972.94*** (328.57)	-15805.35*** (385.21)	-13260.91*** (361.12)
Constant	46411.64*** (1724.36)	44932.05*** (1633.93)	46710.66*** (1974.76)	45376.56*** (1850.88)
Commuting zone FE	X	X	X	X
Occupation FE	X	X	X	X
Race FE		X		X
Education FE		X		X
Sex FE		X		X
Full-time workers only			X	X
N	719708	719708	568982	568982
R-sq	0.137	0.247	0.144	0.262

Note: The dependent variable is income from wages. Commuting zone and occupation FE are CZ and SOC5. Standard errors are clustered at the commuting zone level.
* p<0.05 **p<0.01 ***p<0.001

The coefficients on HHI in all columns are very large in real terms: moving from a perfectly competitive labor market to a perfectly monopsonistic market is associated with a decrease in wages of over \$12,900 in all specifications. If this sounds too extreme, note that the 90th percentile for rescaled HHI in this sample is just 0.39; markets with some degree of monopsony are widespread, but perfect monopsony is very rare. Moving from the 25th to 75th percentile of monopsony is associated with a 0.164-point increase in rescaled monopsony, and thus with a drop in wages of between \$2,127 and \$2,592 depending on the model used.

The main results of the paper on monopsony and inequality are no less compelling. Results in all specifications are clear: increased monopsony power is associated with increased inequality, and these effects are large and highly statistically significant. My first set of results in Table 3 shows the effect of monopsony power on overall inequality as measured by the interdecile range of wages (P_{90}/P_{10}). Since the rescaled HHI ranges from 0 to 1, the coefficients on HHI can be interpreted as the change in interdecile range associated with moving from a perfectly competitive to a

perfectly monopsonistic market. The results in column 1 include only commuting zone and occupation (5-digit SOC code) fixed effects, while column 2 adds controls for race, education, and sex. Column 3, my preferred specification, includes all sets of controls and fixed effects and excludes all part-time workers (defined as those who work less than 40 hours a week on average or who worked less than 40 weeks of the year.) The results for all three models are significant at the 0.1% level.

Table 3: The Effect of Monopsony Power on Inequality

	Model 1	Model 2	Model 3
HHI, rescaled	2.2805*** (0.5652)	2.2861*** (0.5641)	1.9557*** (0.4936)
Constant	5.0948*** (0.4967)	4.8472*** (0.5036)	4.8614*** (0.4858)
Commuting zone FE	X	X	X
Occupation FE	X	X	X
Race FE		X	X
Education FE		X	X
Sex FE		X	X
Full-time workers only			X
N	671959	671959	530624
R-sq	0.190	0.191	0.192

Note: The dependent variable in this regression is interdecile range, calculated as the ratio between the 90th percentile wage in the CZ-SOC5 pair and the 10th percentile wage in the same pair. All specifications exclude CZ-SOC5 pairs with fewer than 5 eligible workers. Commuting zone and occupation FE are CZ and SOC5. Standard errors are clustered at the commuting zone level.

* p<0.05 ** p<0.01 *** p<0.001

The results from column 3 indicate that moving from a perfectly competitive to perfectly monopsonistic labor market is associated with a 1.956-point increase in interdecile range. The average interdecile range in this sample is 8.256, so this represents a 23.7% increase in inequality in perfectly monopsonistic labor markets, *ceteris paribus*. The specifications in models 1 and 2 yield even higher estimates of the inequality effects of moving from competition to monopsony, 27.6% and 27.7% respectively. Moving from the 25th to 75th percentile of HHI would be associated with an increase in interdecile range of between 3.9% and 4.5% depending on the preferred specification. These results are not just statistically significant; they are also meaningful in real terms. A 1.956-point increase in interdecile range looks roughly like moving

from Canada (with an average interdecile range of 4.1, according to the World Bank) to the United States (with a range of 6.2). Consider a market where under perfect competition, the mean bottom decile wage in some labor market is \$20,000 and the mean top decile wage is \$100,000. In this case, moving from perfect competition to perfect monopsony without broader wage effects would mean wages for the top decile increasing to \$139,120 without any rise in earnings at the bottom, or wages for the bottom decile falling to \$14,376 without any decrease in earnings at the top. In either case, this is a dramatic spike in the gap between rich and poor.

Of course, the real inequality effects of monopsony will be more complex and fall unevenly on different income groups. Fortunately, this can also be measured empirically. In Table 4, I specifically analyze upper and lower tail inequality by breaking the main measure of inequality into two partial ranges: the ratio between 50th and 10th percentile wages in a particular labor market, and the ratio between 90th and 50th percentile wages. Both models include commuting zone and occupation fixed effects along with controls for race, education, and sex.

	50/10 Ratio	90/50 Ratio
HHI, rescaled	0.6200*** (0.1708)	0.1678* (0.0742)
Constant	2.0070*** (0.0906)	2.4054*** (0.0540)
Commuting zone FE	X	X
Occupation FE	X	X
Race FE	X	X
Education FE	X	X
Sex FE	X	X
N	671959	671959
R-sq	0.137	0.301

Note: 50/10 and 90/50 ratios are calculated as the ratio between the given percentiles of the wage distribution for each CZ-SOC5 pair. All specifications use 5-digit generalized SOC codes and exclude pairs with fewer than 5 eligible workers. Standard errors are clustered at the commuting zone level.

* p<0.05 ** p<0.01 *** p<0.001

These results indicate that the increase in inequality due to monopsony is driven largely by a significant increase at the low end of the income distribution. Inequality between median earners

and top earners increases slightly, but by a far smaller magnitude. In percentage terms, moving from perfect competition to perfect monopsony is associated with an increase of inequality of 17% between median earners and the bottom decile, and an increase of 7.5% between top and median earners, relative to the mean of each ratio. These findings are significant at the 0.1% level for the 50/10 ratio, but only at the 5% level for the 90/50.

This is identical to the result predicted by the model in section 3, which postulated that low-wage workers would be subject to a greater degree of exploitation as monopsony increases relative to their high-wage counterparts. The bulk of the empirical evidence shows that all employees in a monopsonistic labor market, including high- and medium-wage workers, suffer some degree of monopsony wage penalty, but these results provide evidence that low-wage workers are hit the hardest. Even if one does not see inequality as inherently undesirable, a 17% increase in the gap between the median worker and the bottom decile driven by wage decreases will clearly tend to lead to greater poverty and more reliance on social safety nets among low earners. This is a policy-relevant effect, particularly given the magnitude of the change in inequality levels at the low end of the earnings distribution.

6. Conclusions and Policy Recommendations

This paper documents the relationship between monopsony and inequality, adding to the limited literature on the topic by showing that labor markets with high monopsony power as measured in terms of vacancies tend to also exhibit higher inequality. I demonstrate that this increase is driven primarily by a drop in wages among the lowest-earning subset of the labor force, indicating that firms with a higher degree of labor market power have a greater tendency to exploit low-wage workers. These results also contribute to a parallel literature: the dramatic increase in wage inequality in the United States (e.g. Piketty, 2015) has occurred at the same time as a similarly dramatic increase in market concentration (e.g. Benmelech et al., 2020, Autor et al., 2020), and the results provide some evidence that these trends are causally connected.

These results suggest that policymakers must make monopsony power a key consideration in any plan to address inequality in the United States. Historically, regulators have been far more concerned with concentration in the product market, rarely making judgements on the basis of labor market concentration or employee welfare. The inequality effects of monopsony, along with

its overall wage effects, make this an increasingly untenable standard. Regulators will need a set of analytic methods for judging the effects of mergers on labor markets; if used properly, they will create more efficient markets while also addressing a significant degree of exploitation and inequality in the United States. Azar et al. (2020) suggest that regulators might simply apply the same HHI standards to labor markets as they do in the product market, which would consider markets with an HHI above 2,500 to be highly concentrated and thus subject to additional scrutiny, while Naidu et al. (2018) propose more specific analytic methods. More research will certainly be produced to suggest better methods and standards, but it is important to take action quickly rather than waiting for the perfect answer, particularly in light of the historic collapse of small businesses during the Covid-19 pandemic. The results of this paper suggest that monopsonistic labor markets are characterized by an unnecessary degree of inefficiency and inequality, issues which are preventable if addressed at their source. Regulators and legislators should act to address the widespread and pervasive nature of monopsony in the United States; doing so is likely to reduce poverty and income inequality across the country.

Appendix: SOC Codes

It was necessary to use generalized 5-digit SOC codes rather than the more specific 6-digit version in conducting analyses on this data. 5-digit SOC codes can be thought of as identical to the categories (ending in 0) listed on the Bureau of Labor Statistics website which defines SOC codes. For example, both 27-1022 (Fashion Designers) and 27-1025 (Interior Designers) would be counted under the category 27-1020 (Designers). This method necessarily loses some degree of specificity in the data, but attempting to use 6-digit codes would have systematically biased the results. In an attempt to anonymize responses as fully as possible, the ACS data cuts the last digit of occupation off of occupations that are relatively uncommon in their geographic area and replaces it with an X (e.g. displaying occupation as 27102X or 13115X).¹ These represented almost 40% of the sample, and were disproportionately located in rural and suburban areas where fewer people worked in each occupation. Simply dropping these would create a systematic bias in the sample towards more densely populated (and thus generally lower-monopsony) areas, and we observe this difference in Table 1. The mean rescaled HHI when using 5-digit SOC codes is 0.223, while with 6-digit codes the mean HHI is just 0.098.

In order to address this source of bias, I used a weighted average of the HHI facing each subcategory in SOC6, weighted by the share of total employment, to calculate the HHI facing employees at the SOC5 level. This is the main measure of monopsony power used in analyses throughout the paper. It would be comforting to know that results are fundamentally similar when using the existing (albeit biased) SOC6 measures of inequality, and fortunately, that is the case. Table 3 displays the results of our main inequality specification using HHI data at the SOC6 level. Results are broadly similar; they have the same sign, similar magnitude, similar trend between models (Model 3 being the lowest under both the SOC5 and SOC6 measures). Results are statistically significant at the $P < 0.05$ level in two of the three specifications and at the $P < 0.10$ level in all three.

¹ A small number of occupation codes were shortened even further to 4 digits (e.g. 2710XX), but these represented less than 1% of the sample, were close to identical to the full sample on all dimensions of monopsony, and were dropped without much risk of bias.

Table 5: SOC6 Measures of Inequality

	Model 1	Model 2	Model 3
HHI, rescaled, SOC6	1.9836*	1.9753*	1.2565
	(0.9130)	(0.9122)	(0.8426)
Constant	5.3178***	5.1834***	5.1931***
	(0.4959)	(0.5240)	(0.5188)
Commuting zone FE	X	X	X
Occupation FE	X	X	X
Race FE		X	X
Education FE		X	X
Sex FE		X	X
Full-time workers only			X
N	436934	436934	349760
R-sq	0.166	0.166	0.176

Note: The dependent variable in this specification is interdecile range, calculated as the ratio between the 90th percentile wage in the CZ-SOC6 pair and the 10th percentile wage in the same pair. All specifications use 6-digit SOC codes, which may suffer from biases discussed previously, and exclude CZ-SOC6 pairs with fewer than 5 eligible workers. Commuting zone and occupation FE are CZ and SOC6.

* $p < 0.05$

** $p < 0.01$

*** $p < 0.001$

While it would be preferable to see results which look identical to those in Table 3, these are nonetheless reassuring. Given the downward bias in HHI for the remaining sample, we might expect some degree of change in the effects, and the fact that these estimates are so similar means we can feel more confident in drawing conclusions from the results in Section 5.

References

- Autor, D., Dorn, D., Katz, L. F., Patterson, C., & Van Reenen, J. (2020). The fall of the labor share and the rise of superstar firms. *The Quarterly Journal of Economics*, 135(2), 645-709.
- Azar, J., Marinescu, I., & Steinbaum, M. Labor market concentration, 2017. *nBER Working Paper*, (24147).
- Azar, J., Marinescu, I., Steinbaum, M., & Taska, B. (2020). Concentration in US labor markets: Evidence from online vacancy data. *Labour Economics*, 66, 101886.
- Benmelech, E., Bergman, N. K., & Kim, H. (2020). Strong employers and weak employees: How does employer concentration affect wages?. *Journal of Human Resources*, 0119-10007R1.
- Boal, W. M., & Ransom, M. R. (1997). Monopsony in the Labor Market. *Journal of economic literature*, 35(1), 86-112.
- De Loecker, J., Eeckhout, J., & Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2), 561-644.
- Dube, A., Jacobs, J., Naidu, S., & Suri, S. (2020). Monopsony in online labor markets. *American Economic Review: Insights*, 2(1), 33-46.
- Dustmann, C., Lindner, A., Schönberg, U., Umkehrer, M., & Vom Berge, P. (2019). *Reallocation effects of the minimum wage: Evidence from Germany*.
- Federal Reserve. (2019). Report on the Economic Well-Being of US Households in 2018. *Board of Governors of the Federal Reserve System, Washington, DC*.
- Grullon, G., Larkin, Y., & Michaely, R. (2019). Are US industries becoming more concentrated?. *Review of Finance*, 23(4), 697-743.
- Hershbein, B., Macaluso, C., & Yeh, C. (2019). Monopsony and Concentration in the Labor Market: Evidence from Vacancy and Employment Data.
- Kahn, L. M. (2000). The sports business as a labor market laboratory. *Journal of Economic Perspectives*, 14(3), 75-94.

Karabarbounis, L., & Neiman, B. (2014). The global decline of the labor share. *The Quarterly journal of economics*, 129(1), 61-103.

Manning, A. (2003). *Monopsony in motion: Imperfect competition in labor markets*. Princeton University Press.

Naidu, S., Posner, E. A., & Weyl, G. (2018). Antitrust remedies for labor market power. *Harv. L. Rev.*, 132, 536.

Piketty, T. (2015). About capital in the twenty-first century. *American Economic Review*, 105(5), 48-53.

Purcell, P. J. (2020). Geographic Mobility and Annual Earnings in the United States. *Social Security Bulletin*, 80(2), 1-24.

Rinz, K. (2018). Labor market concentration, earnings inequality, and earnings mobility. *Center for Administrative Records Research and Applications Working Paper*, 10.

Rodgers, J. R., & Rodgers, J. L. (2000). The effect of geographic mobility on male labor-force participants in the United States. *Journal of Labor Research*, 21(1), 117-132.

Webber, D. A. (2015). Firm market power and the earnings distribution. *Labour Economics*, 35, 123-134.