The balance of power: monopsony, unions and wages in the United Kingdom

Will Abel*, Silvana Tenreyro† and Gregory Thwaites‡

August 2020

Abstract

We document the evolution of monopsony power in the UK private sector labour market from 1998-2018, how labour unions have counterbalanced this power, and the net effect on wages. Using linked employee-firm micro-data, we find that: (1) Measures of labour market concentration have not exhibited a time trend over the time period examined. (2) There is substantial cross-sectional variation in monopsony at the industry level. (3) Higher levels of labour market concentration are associated with lower pay amongst workers not covered by a collective bargaining agreement. (4) For workers covered by a collective bargaining agreement, the association between labour market concentration and pay disappears. (5) The effects of concentration and union coverage are generally larger for lower-paid workers, and workers in tradable industries. (6) Collective bargaining agreements weaken the impact of workers' outside options in other labour markets, which nonetheless remain strong.

*Bank of England

†Bank of England, CEPR and LSE

‡CfM and University of Nottingham

[1]We thank Maria Ludovica Ambrosino for excellent research assistance.

1 Introduction

Robinson (1933) founded the analysis of markets, including the labour market, in which buyers or sellers have market power. Manning (2003) outlines how we should expect monopsonistic labour markets to emerge as a result of frictions, the most important of which are lack of employee information about work opportunities, individual heterogeneity in job preferences, and mobility costs. A substantial literature (see, for example, Card et al. (2004)) analyses the impact that labour unions have on employment and wages.

In this paper we use linked worker-firm micro-data from the United Kingdom to study the evolution and effects of monopsony power in the UK private sector from 1998-2017. Using a 1% representative panel sample of all employee, along with unique firm identifiers and the business census, we construct industry-region measures of concentration in employment over time. Crucially, we are able to measure whether workers are covered by a union-negotiated collective bargaining agreement (CBA). The prevalence of these agreements fell sharply over our sample period, making the UK an ideal laboratory to study their effects.

We document a complex relationship between wages, monopsony power and the balancing effects of labour unions. We find that while there are large variations in labour market concentration across industries and regions, aggregate measures of labour market concentration have been relatively stable over the period we examine. Increasing monopsony power, in general, is associated with lower pay. For an employee covered by a CBA, increases in concentration have little relationship with pay levels. For the same worker without CBA coverage, increases in employer concentration are associated with lower pay. For this worker, moving from the 25th percentile of employer concentration to the 75th percentile would be related to a decline in pay of 1.1%. We find that these effects are stronger among lower-skilled workers, and within industries that produce tradable products. They are robust to the inclusion of a large array of fixed effects, and to different ways of measuring labour market concentration.

We follow Schubert et al. (2020) and construct a measure of workers' outside options in different labour markets and corroborate their finding that the worker's wage depends positively on the outside option. Our key results on the impact of monopsony and the attenuating effect of CBAs are robust to their inclusion. Furthermore, we show that the impact of the outside option is somewhat reduced when a worker is covered by a CBA. Our intepretation is that, from a worker's perspective, a CBA and a credible outside option are both means of capturing part of the surplus from an employment relationship. When a CBA is in place, the outside option is less necessary.

Our findings contribute to a growing literature on the extent, evolution and impact of employer market power. Our main contribution is to show that CBAs help workers to offset employer power across the entire private sector of a large industrialised economy. Azar et al. (2020) also look at the effects of monopsony power on wage growth in the US

- also finding a negative relationship between the two. Our finding are similar to and concurrent with those of Benmelech et al. (2018) who examine the effect of labour market concentration in the US economy using Census data of manufacturing firms. Similarly to us they find a negative relation between employer concentration and wages, that union membership weakens this relationship and that the link between productivity and wages is lower in more concentrated labour markets. The magnitude of their effects are comparable to those we find for the UK economy. However where they find that labour market concentration has been increasing over time in the United States, we find these measures relatively unchanged for the United Kingdom. Our findings cover the whole economy compared to Benmelech et al. (2018) who focus exclusively on manufacturing firms which represent around 9% of current US employment. Furthermore, we are able to control for individual employer and employee characteristics.

The remainder of our paper is structured as follows. We review the relevant literature in section 2. Section 3 explains how our data are measured and constructed. Section 4 presents some simple descriptive statistics. Section 5 sets out our econometric analysis and results. Section 6 concludes.

2 Literature

This paper builds on several strands of literature, stretching back many decades, covering the effect of labour unions on the wages of different groups, how unionisation has changed over time, how employers gain power in the labour markets they operate in, and how unions can counterbalance this power.

There is a large body of literature that relates to the impact of unionisation on wages. Card (1992) looks at the effect of unions on the wages in US. Using a longitudinal estimation to account for misclassification errors in union status, the author finds that variation in union wage gap represents a combination of a larger wage effect for lower-paid workers, and differential selection biases. It also look at how changes in unionization account for one-fifth of the increase of the variance of adult male wages between 1973 and 1987.

Previously, Hendricks (1975) used disaggregated union wage rates to measure the impact of labor and product market structure on the wage levels of union members. His findings suggest that the employer's ability to "whipsaw" the union may not be an important determinant of the union's power over wages or may be offset by other factors. Card et al. (2004) looks at how the impact of unions on the wage structure depends on the industrial relations system - the social, political, legal, institutional and economic environment in which unions operate. The author finds that unions tend to reduce wage inequality among men across countries (U.S., U.K. and Canada) and time periods (from the early 1970s to 2001), but that unions do not reduce wage inequality among women.

Freeman and Medoff (1979) show that the union differential in the 1970s was 20-30% using cross-section data. They reported that unions raise wages most for the young, the least tenured, whites, men, the least educated, blue collar workers and in the largely unorganized South and West. They also pointed out that age differentials tend to fall with size of workplace and that the magnitude of the differential varies among people, markets and time periods. Blanchflower and Bryson (2004) explore the various claims made by Freeman and Medoff. From there, Blanchflowe and Bryson show that private sector union wage premium is lower in thr 90s early 2000s than it was in the 1970s and that the union wage premium is counter-cyclical. They also finds evidence of a secular decline in the private sector union wage premium and that there is big variation in industry-level union wage premia. Union workers remain better able than non-union workers to resist employer efforts to reduce wages when market conditions are unfavorable.

More recently, in Forth and Millward (2002), the Employee Relations Survey shows that show that, in the private sector, 40% of employees work in establishments where some employees are covered by collective bargaining, but one fifth of these are not covered by the arrangements. Unions do achieve a wage premium of around 9% for about half of employees covered by collective bargaining arrangements but the effect of CBA spills over to other employees in the same workplaces. The authors also show that only where bargaining covers between 70% and 99% of employees is there a demonstrable pay premium over employees in similar non-union workplaces.

The study of labour market oligopsony also has a long history. In Boal and Ransom (1997), comparing models with few employers to models based on labor market frictions such as moving costs and search, they find that monopsony power in the former case is rare but occasionally large, while monopsony power based on frictions is probably widespread but small on average. Manning (2003) is based on two main assumptions: presence of frictions in the market and that employers set the wages. The former means the employers have market power over the workers, and the latter means that such power is exercised. These two assumptions allow us to recognise the theoretical and empirical importance of monopsony to understand distribution of wages, unemployment and human capital. Indeed, the author underlines how the presence of monopsony power of employers on the workers needs to be acknowledge to assess more realistically policies as minimum wage,equal pay legislation, and caps on working hours, and analyses these issues in a dynamic set-up.

In Azar et al. (2020), the authors also look at how concentration affects wages by quantifying concentration at occupation level for commuting zones in the US. In accordance with the literature, they find that higher concentration in the labor market is associated with lower posted wages. MARTINS (2018) contributes to the literature on monopsony as it looks at how the market is concentrated and what is the impact on wages, by exploiting rich matched employer-employee data, which include the full population of workers and occupations in Portugal. He finds that less than 9% of workers are exposed to concentration levels thought to raise market power concerns but he also underlines how this measure is sensitive to methodological choices (occupation and geographical area analysis).

Rinz et al. (2018) examines the trends in concentration in local labor markets in US (1976-2015) and its effect on earnings outcomes for various groups of workers, as well as for the whole workforce. He finds declining trends in local industrial employment concentration have differed substantially from national ones - which have been rising from 1990. This divergence is not sensitive to the industry, geographic area, or use of employment weights. The author uses the variation in local industrial concentration over time to estimate its effects on earnings, inequality, and mobility. Consistently with the overall literature, increased concentration reduces earnings. The estimates imply that moving from the median to the 75th percentile industrial concentration distribution would reduce earnings by about ten percent - a very large effect.

Hershbein et al. (2018) confirm the negative correlation between local labor market concentration and average wages. They estimate that a 1% increase in local labor market concentration is associated with a 0.14% decrease in average hourly wages. Further, they also show that monopsony power positively affect demand for skills. When labor is heterogeneous, monopsony potentially affects both the quantity and the quality of labor. Specifically, they find that in the last decade, at most 5% of new U.S. jobs are in moderately concentrated local markets; and that local labor market concentration decreased over time, dropping by at least 25% since 1976.

Schubert et al. (2020) measure outside options across occupations and regions, and estimate their effect on wages. Outside options is composed by job options within workers' occupation - local labor market concentration - and job options outside workers' occupation - identified using data on occupational mobility. They find that moving from the 75th to the 95th percentile in employer concentration across workers, or the 25th to 75th percentile in employer concentration across occupation-city labor markets, results in a 5% lower wage. Differences in labor market concentration could explain 21% of the interquartile wage variation within a given occupation across cities. They also show that a further 13% of the interquartile wage variation within an occupation across cities can be attributed to differences in the quality of outside-occupation options.

The paper that addresses questions most similar to our is Benmelech et al. (2018), which focuses on labor market concentration in manufacturing and its effects on wages, within localised geographic areas in US. The data shows that local employer concentration has increased considerably over time, with the employment weighted mean HHI increasing by 5.8%, from 0.698 during 1977–1981 to 0.756 during 2002–2009. Employers operating in areas with more concentrated labor markets thus appear able to exploit monopsony power in order to reduce employee wages. They also find a negative relation between the HHI and wages, and that this is significantly weaker among plants in industries with high unionization rates. Finally, the paper shows that high levels of concentration allow employers to use their monopsony power to impede the translation of productivity growth into wage increases.

3 Data sources construction

In this section of the paper we describe our key data sources and how we transform them to produce the variables we use in later sections of the paper.

Our key data source is the National Earnings Survey - Annual Survey of Hours and Earnings (NES-ASHE) panel dataset. This is a 1% weighted sample of all employees sampled from National Insurance numbers from 1975-2017, and includes employee-level information on weekly pay, industry, occupation, union coverage, and the location and and size of each individual's place of work. We can use this survey to measure wages at the individual level, and also aggregate it to measure labour-market concentration at the level of the local labour market.

The NES-ASHE data records whether an employee is covered by a collective bargaining agreement (CBA). A CBA is an agreement between one or more employers and one or more labour unions or workers' committess concerning aspects of employment such as pay and conditions. In the UK, employers can voluntarily recognise labour unions for the purposes of negotiating such an agreement, or can be required to do so if a sufficiently large fraction of eligible workers vote in favour of this. Agreements can be negotiated at the national, sub-national, employer or workplace level. Such agreements often cover non-union employees within the same group. As a result, union coverage (the fraction of workers covered by a union-negoatiated agreement) exceeds union density (the fraction who are members of a union). In our dataset 50.4% of private sector employees are covered by a CBA in 1998, declining at a relatively constant rate to 21.4% in 2017.

We merge the NES-ASHE worker-level data with firm data from the Business Survey Database (BSD), which has annual turnover and employment for the universe of UK firms from 1997-2017. Prior to 1998 the firm identifiers which we require to construct our desired concentration indexes are no longer available and as such we are limited to the period of 1998-2017. Furthermore the firm identifiers are distorted in 2000 so this year of data must be dropped. In total we have 3 million observations across 19 years.

From the BSD we construct firm-level measures of turnover per head, to proxy for firm productivity. We cross check this using detailed data from the Annual Business Survey (ABS) for the same time period. From this we construct value added per employee (a truer measure of worker productivity than turnover per head) for a sample of firms. The correlation between firm turnover and value added per head is 0.5, giving us sufficient confidence in our use of turnover per head measures as a proxy for productivity. Unfortunately the firm identifiers used by the BSD and ABS differ from those used in the NES-ASHE prior to 2002, meaning that for certain pieces of analysis we will be forced to use the smaller sub-sample of 2002-2017.

We construct a Herfindahl-Hirschman index (HHI) to measure the concentration of employment across firms. An HHI is a typical measure of concentration which ranges from 0 to 1 where a score of 1 indicates a completely concentrated market with only one employer. Lower scores indicate higher levels of competition in the market. We construct our HHI at an industry-year-region level, where industry is measured at the 2-digit SIC level and location with NUTS2 regions (approx. 1m jobs). This gives us a total of 85 industries, 39 regions and 19 years leading to 62,985 concentration measures in our dataset.

Specifically we calculate:

$$HHI_{ind,t,region} = \sum_{j=1}^{J} s^2_{j,ind,t,region}$$

where $s_{j,ind,t,region}$ is the employment share of firm j in a given industry-year-region cell.

In our baseline measure of concentration, we construct the HHI using total employment (measured with the BSD) in the firms employing workers that appear in the ASHE-NES sample. As our concentration measure is calculated from sampled data, as opposed to the complete population of employment, it will be noisy and biased to some extent. We discuss this in greater detail in appendix A1. To mitigate this problem, labour markets containing fewer than 10 observations of workers are dropped. This leads to a loss of 75,000 observations, around 2.5% of our data.

Two reasonable alternatives to this baseline measure would be (1) to include the employment of all firms in the BSD when calculating concentration, including those not appearing in the ASHE-NES sample, or (2) to rely exclusively on the random sample of the firm size distribution implied by the ASHE-NES. The first alternative will include the population of firms, a potential advantage, but among these are very many extremely small firms which are often pass-through entities or a disguised form of self-employment. The second alternative will involve noisier measurement of the amount of employment, but will admit (1) an occupation-based measure of labour-market concentration and (2) a longer sample period as it is not constrained by the availability of the BSD. We present robustness checks to these alternatives below.

Finally, we adapt a measure of worker outside options used by Schubert et al. (2020). We employ the panel element of our data to measure the transition rate ω from each labour market in our data to each of the others. We then construct a worker's outside option as the transition-rate-weighted average wage in other labour markets as follows

$$OO_{ind^*,region^*,t} = \sum\nolimits_{ind \neq ind^*, region \neq region^*} \omega_{ind^* \to ind, region^* \to region, t} w_{ind,region,t}$$

where $w_{ind,region,t}$ is the average wage in a labour market at time t and $\omega_{ind^* \to ind, region^* \to region, t} w_{ind,region,t}$ is the transition rate from labour market $\{ind^*, region^*\}$ to labour market $\{ind, region\}$at time t.

4 Descriptive statistics

In this section of the paper we present some key descriptive statistics about labour market concentration and CBA coverage.

Figure 1 shows a time series for the median and mean level of concentration faced by employees over time, where averages are taken across labour markets. Concentration levels at the end of our sample are similar to those in the starting period. Mean and median labour market concentration levels have fallen by only 1% and 3% respectively between 1998 and 2017, between which there was a substantial rise and fall in concentration. Over the first decade of the series, mean and median concentrations increased by 24% and 28% respectively, from 1998-2008, and subsequently declined back to their starting levels.

[Figure 1 about here.]

Behind this rise and fall in the aggregate time series, there is, moreover, substantial variation at the cross sectional level. Figure 2 shows the distribution of concentration measures by industry - where each observation represents an industry-region-year concentration measure. Two observations are readily apparent. First, there is large variation in concentration between industries along lines we would expect - there is high competition for workers in retail and residential care industries for example, while there are relatively few employers in industries such as sewerage, mining and courier services. Secondly there is a strong rightwards skew to the data (it appears to be log-normally distributed) suggesting that, even in relatively competitive industries, some workers may still face very monopsonistic labour markets.

[Figure 2 about here.]

The rightward skew in industry concentrations is not driven by regional variation. One possibility is that the pattern we observe in figure 2 is due to the fact some industries are in more remote locations and so will naturally have less competitors. To examine this, in Figure 3, we show concentration aggregated by NUTS1 region. While there is some regional variation, with the relatively sparely populated regions of South Yorkshire, the Highlands and Cumbria being the most concentrated regions and West London, East Anglia and Oxfordshire being the least concentrated, we see that there is substantially more within than between regions. Even workers in parts of Manchester or London face highly concentrated labour markets, depending on their industry of work.

[Figure 3 about here.]

Figure 4 shows how the fraction of employees which are covered by CBA has declined sharply in the UK over recent decades. This makes the UK an ideal laboratory to test hypotheses about how union coverage affects the labour market. Figure 5 shows that union coverage varies a great deal across countries, and furthermore that coverage and density (the fraction of workers who are members of unions), while often thought to be interchangeable, can be very different in some contexts.

[Figure 4 about here.]

[Figure 5 about here.]

5 Econometrics

This section sets out our econometric methodology and results.

5.1 Methodology

The starting point for our analysis is the following reduced-form equation:

$$w_{i,t} = \alpha + \beta_1 HHI_{ind,t,region} + \beta_2 X_{i,t} + e_{i,t} \tag{1}$$

where $w_{i,t}$ is the log of an individual i's gross weekly wage in year t. $HHI_{ind,t,region}$ is the labour market concentration for a given industry-year-region combination. $X_{i,t}$ is a vector of individual controls that may include firm-level turnover per head, age, age squared, gender, CBA coverage, size of firm the individual is employed at, whether a worker is full or part time and whether they are on a temporary contract, plus dummies for industry, occupation, region and year. We first present results in which a small subset of controls are included, and then successively add controls and/or interactions in order to assess the stability of our core results to their inclusion, and in some cases to measure the impact of these variables in their own right.

5.2 Results

Table 1 shows the results from our baseline regression. The first column shows that being covered by a CBA raises an employee's wages by about 1 per cent when we do not control for concentration. This parameter is well-determined and highly statistically significant but is nonetheless a smaller union premium than many other researchers have found.

If a local productivity shock affects both wages and the extent of local labour-market concentration, the partial correlation coefficient between these two variables may not reflect the causal impact of concentration on wages, but rather their joint dependence on an unobserved productivity shock. Other authors (insert refs) have accordingly instrumented for the concentration variable in their regressions. The inclusion of firm-level productivity in our specification obviates the need to employ an instrumental variable strategy. The results in Table 1 this column also show a 1 per cent rise in productivity at the firm level raises a worker's wages by about 5 per cent, suggesting that firms typically share a small fraction of the rents they generate with their workers. In order to guard against the possibility that firm- or market-level productivity shocks may in turn be endogenous with respect to wages, we have experimented with using lagged rather than contemporaneous productivity. This variant has essentially no effect on our results.

The second column shows that concentration has a small negative effect on wages. Moving from the 25th to the 75th percentile of concentration would be associated with a decline in pay of around 1.6 per cent.

The third column contains the first instance of our key result. When we include CBA coverage, concentration and their interaction, we find that the coefficients on the latter two variables are approximately equal and opposite in sign. This means that an increase in labour-market concentration reduces the wages of those who are not covered by a CBA, but has no effect on workers who are covered. Furthermore, the coefficients on both the CBA and concentration variables increase somewhat. The union wage premium increases by a factor of 4-5, depending on the degree of concentration in the local labour market, while the effect of labour market concentration increases by around 30% for employees who are not covered by a CBA.

Our baseline measure of labour-market concentration defines a labour market by industry rather than by occupation. This is for two reasons. First, labour-market transitions between industries are lower than between occupations (see Figure 6). Secondly, classifying labour markets by industry allows us to use the BSD to measure employment within a given labour market. One potential concern with our results is that our measure of labour-market concentration, based on the local concentration of employment by industry, could be measuring a combination of concentration in both the labour and product markets, such that our key RHS variable would be mismeasured to some extent. To guard against this, we estimate two variants of our model.

First, we partition our sample into tradable and non-tradable industries, using the classification employed in Broadbent et al. (2019). Local product market concentration should matter much less with tradable goods, such that our measure of labour-market concentration will have less noise for this sector of the economy. Consistent with this, the results shown in columns 4 and 5 of Table 1 show larger effects of concentration, and a larger offsetting impact of CBA coverage, for tradable than for nontradable goods.

Secondly, we estimate a variant of our model in which our HHI index is defined by occupation-region rather than industry-region. Table 2 shows the results, alongside the baseline regression from Table 1. The results show that the impact of all the variables is qualitatively the same. The negative impact of concentration on wages is not statistically significant at standard levels, but the interaction with union coverage is both larger in absolute value and statistically more significant, suggesting that union members in labour markets with concentrated occupations actually enjoy higher wages.

Following Hershbein et al. (2018), Table 3 shows the results for our baseline regression, in which the data have been disaggregated by skill level[1]. The results suggest that union membership boosts pay the most among low-skill workers, and that this result is stronger in more concentrated labour markets. Union membership seems to be associated with lower pay among high-skilled workers.

Schubert et al. (2020) include a measure of workers' outside options in their analysis of the impact of labour market oligopsony in the US, arguing that an outside option can improve the bargaining power of a worker that would otherwise have to accept a low wage offered. The outside option may be correlated with the degree of labour-market concentration, such that the coefficient on the latter will be biased if the former

[1] We group the SOC 2010 codes 1-3 as high skill, 4-6 as medium skill and 7-9 as low skill.

is excluded.

Column 2 of Table 4 shows that the outside option, defined as the transition-rate weighted average of wages in other labour markets, exerts a strong effect on wages: a 1% rise in the outside option translates into a 0.3% rise in wages. This effect is especially large in light of the inclusion of occupation, region, industry and time fixed effects in the regression. The inclusion of the outside option attenuates somewhat the coefficients on CBA coverage, concentration and their interaction. Interestingly, we also find that CBA coverage reduces the impact of the outside option somewhat. The interpretation is that being covered by a CBA affords workers another means of capturing the rents created by the employment relationship.

[Figure 6 about here.]

[Table 1 about here.]

[Table 2 about here.]

[Table 3 about here.]

[Table 4 about here.]

6 Conclusion

In this paper we describe what is, to our knowledge, the first multi-decade, economy-wide time series measurement of monopsony power in the labour market of a major industrialised country. We document that monopsony power increased from 1998-2008, before declining from 2008-2018, and then subsequently returning to levels broadly in line with those seen at the begining of our sample. We also document substantial variation across industries and regions.

We have shown how higher levels of concentration are associated with lower levels of pay for workers not covered by a collective bargaining agreement, and that for those who *are* covered by a CBA that this negative correlation between pay and monopsony mostly disappears. These effects are stronger in tradable goods sectors, and for lower-paid workers. They remain when we control for a worker's outside options, and CBAs attenuate the effects of these options.

These results emphasise the importance that labour-market institutions can play in counterbalancing the loss of workers' power (Stansbury and Summers (2020)), and potentially in helping workers bargain for rents created in concentrated product markets or by 'superstar' firms (De Loecker et al. (2020), Autor et al. (2020)),

References

Autor, D., Dorn, D., Katz, L. F., Patterson, C., and 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., and Steinbaum, M. (2020). Labor market concentration. *Journal of Human Resources*, pages 1218–9914R1.

Benmelech, E., Bergman, N., and Kim, H. (2018). Strong employers and weak employees: How does employer concentration affect wages? (NBER 24307).

Blanchflower, D. G. and Bryson, A. (2004). What effect do unions have on wages now and would freeman and medoff be surprised? *Journal of Labor Research*, 25(3):383–414.

Boal, W. M. and Ransom, M. R. (1997). Monopsony in the labor market. *Journal of economic literature*, 35(1):86–112.

Broadbent, B., Pace, F. D., Drechsel, T., Harrison, R., and Tenreyro, S. (2019). The Brexit Vote, Productivity Growth and Macroeconomic Adjustments in the United Kingdom. Discussion Papers 1916, Centre for Macroeconomics (CFM).

Card, D. (1992). The effect of unions on the distribution of wages: redistribution or relabelling? Technical report, National Bureau of Economic Research.

Card, D., Lemieux, T., and Riddell, W. C. (2004). Unions and wage inequality. *Journal of Labor Research*, 25(4):519–559.

De Loecker, J., Eeckhout, J., and Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2):561–644.

Forth, J. and Millward, N. (2002). Union effects on pay levels in britain. *Labour Economics*, 9(4):547–561.

Freeman, R. B. and Medoff, J. L. (1979). The two faces of unionism. Technical report, National Bureau of Economic Research.

Hendricks, W. (1975). Labor market structure and union wage levels. *Economic Inquiry*, 13(3):401–416.

Hershbein, B., Macaluso, C., and Yeh, C. (2018). Concentration in us local labor markets: evidence from vacancy and employment data. Technical report, Working paper.

Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press.

MARTINS, P. (2018). Making their own weather? estimating employer labour-market power and its wage effects. *CGR*.

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

Robinson, J. (1933). *The economics of imperfect competition.* London Macmillan.

Schubert, G., Stansbury, A., and Taska, B. (2020). Monopsony and outside options. *Available at SSRN.*

Stansbury, A. and Summers, L. H. (2020). Declining worker power and american economic performance. *Brookings Papers on Economic Activity.*

7 Appendix: Measuring HHIs with sample data

In contrast to estimates of simple population means of quantities, estimates of market concentration derived from random samples will be biased. This bias comes from two distinct sources. The first is simply that when a sample of N workers is drawn from a given market (defined in our case by occupation and region), there is a lower bound on estimated shares of N^{-1} and hence on concentration of N^{-2} attained if each worker in the sample works for a different firm. For low values of N, the true value of oligopsony could be lower than this. The second source of bias comes from Jensen""""'s inquality - i.e. the fact that, given a set of unbiased estimates s_i of true market shares s_i such that $E\left[s_i\right] = \sigma_i$, in general $E\left[s_i^2\right] > \sigma_i^2$.

To investigate the size of this bias empirically, we ran Monte Carlo trials and generated random populations of labour markets with different degrees of concentration and then calculated the observed HHI for different sample sizes and compared them to the true value. Figure 7 below shows that, as expected, estimates of concentration are upwardly biased but that this bias is fairly constant for different true values of concentration, and declines fairly quickly for moderate sample sizes. So while our coefficient estimates are likely to be affected to some degree by the sample data we use to calculate monopsony, the effect seems likely to be small in practice.

[Figure 7 about here.]

List of Figures

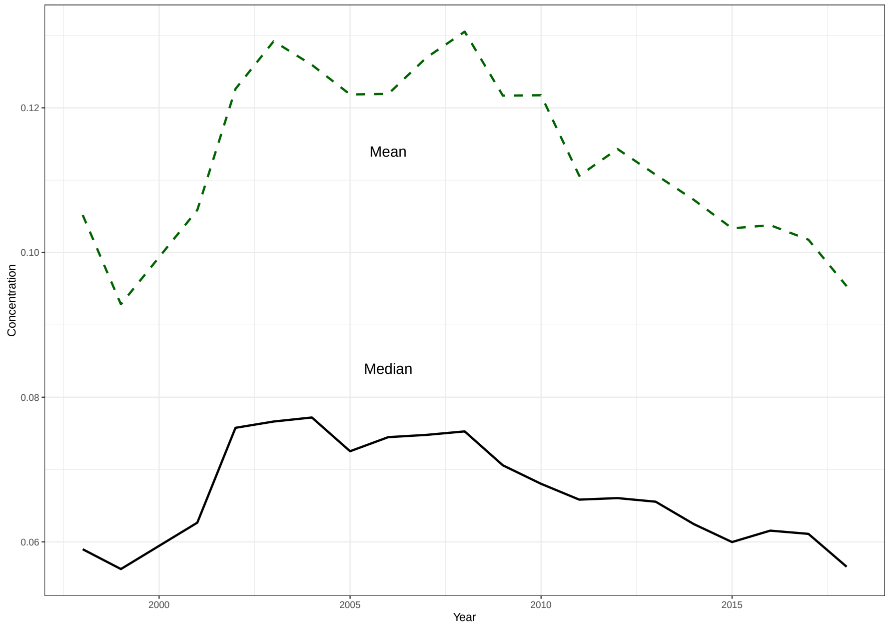

Figure 1: Trends in monopsony over time

This chart shows the mean and median Herfindahl-Hirschman concentration index across region-industry labour markets in the United Kingdom

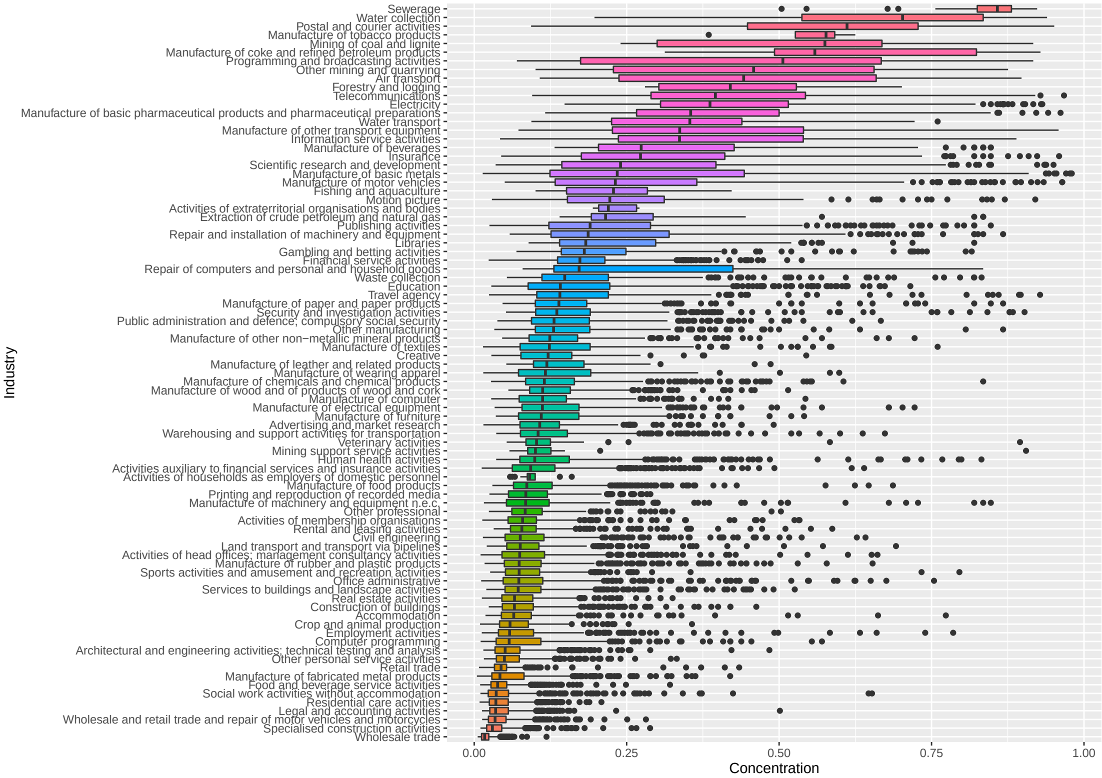

Figure 2: Cross-section of labour-market concentration by industry

This chart show the distribution of HHI concentration indices across regions within different industries in the United Kingdom

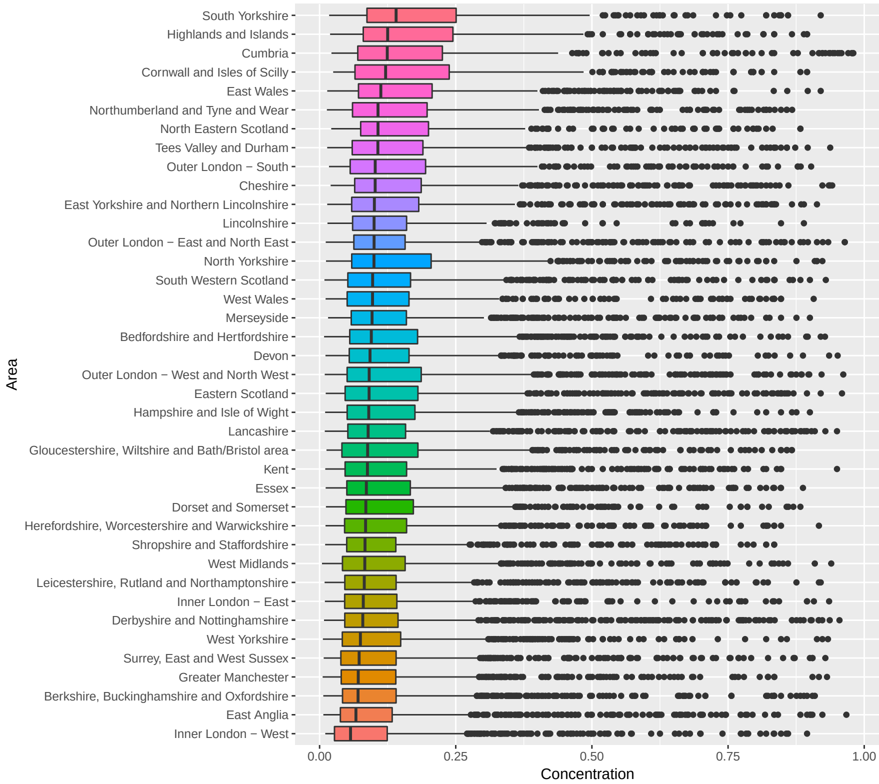

Figure 3: Cross-section of labour-market concentration by region

This chart show the distribution of HHI concentration indices across different industries within different regions in the United Kingdom

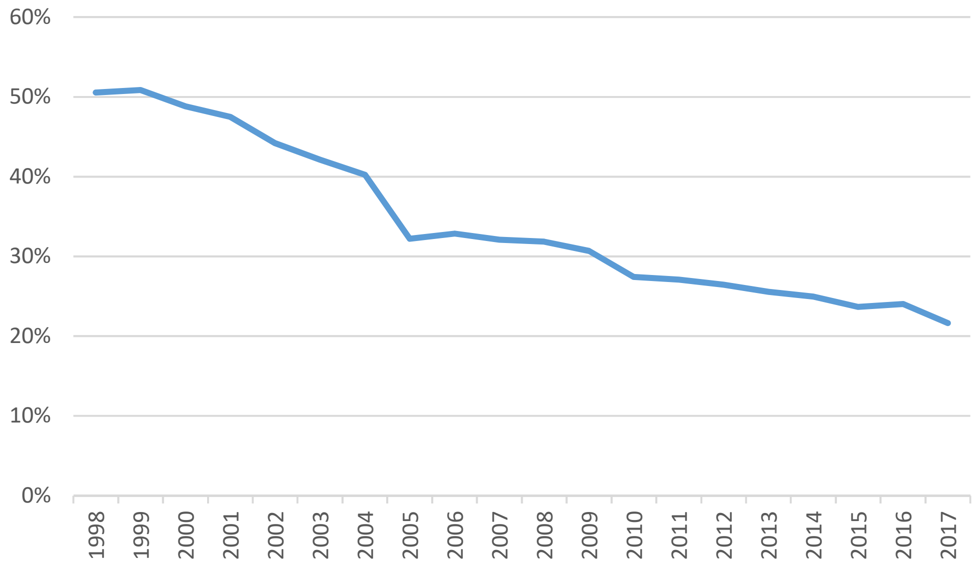

Figure 4: CBA coverage in the UK

This chart show the fraction of UK employees who were covered by a CBA. Source: ASHE-NES

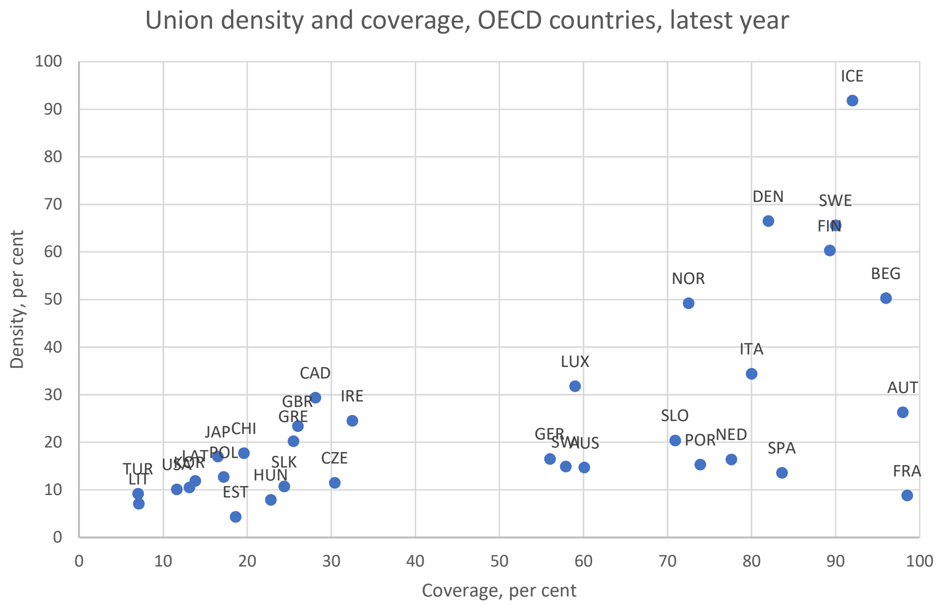

Figure 5: Cross-section of CBA coverage and union density across the OECD

This chart show the fraction of employees who were covered by a union agreement, and the fraction that were members of trade unions. Source: OECD

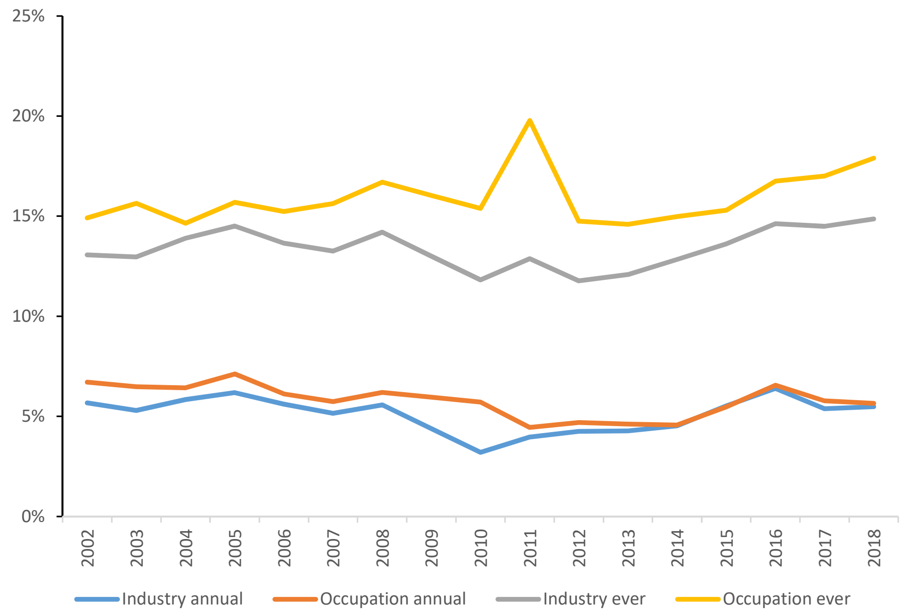

Figure 6: Transition rates of private-sector employees between occupations and industries

This chart shows the fraction of workers that move from one industry or occupation (both at the 2-digit level) to another in each year and in the entire sample

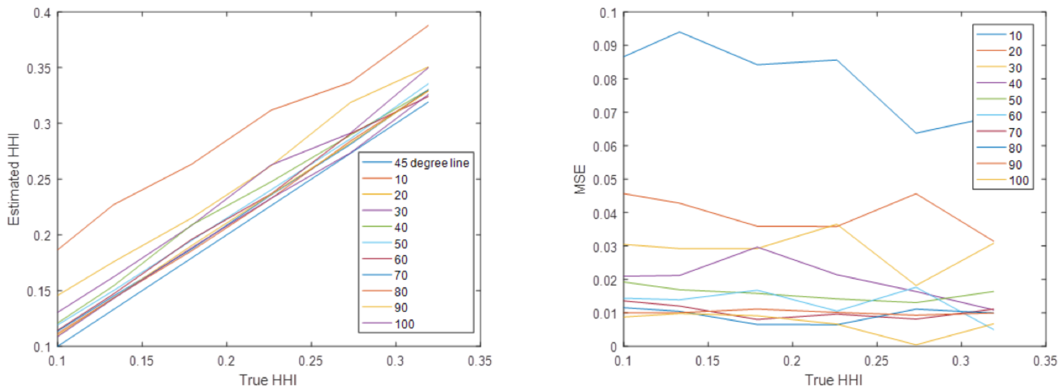

Figure 7: Monte Carlo estimates of HHIs

The left-hand panel shows the relationship between average estimated and true levels of market concentration for different sample sizes in a Monte Carlo trial. The right-hand panel shows the how the mean square error of the estimated HHI depends on sample size and the true HHI..

List of Tables

Table 1: Baseline results

Dependent variable: log weekly wages	(1)	(2)	(3)	(4) Tradables	(5) Nontradables
union = 1	0.0117***		0.0442***	0.0415***	0.0453***
	(0.00396)		(0.00632)	(0.00806)	(0.00802)
	2.949		6.987	5.143	5.649
	(0.00542)		(2.56e-08)	(8.50e-06)	(1.73e-06)
log(concentration)		-0.0131**	-0.0176***	-0.0199**	-0.00569
		(0.00625)	(0.00620)	(0.00775)	(0.00386)
		-2.096	-2.841	-2.561	-1.474
		(0.0428)	(0.00719)	(0.0145)	(0.149)
union x log(concentration)			0.0172***	0.0237***	0.0114***
			(0.00276)	(0.00403)	(0.00358)
			6.245	5.878	3.179
			(2.62e-07)	(8.37e-07)	(0.00294)
lprod	0.0559***	0.0560***	0.0558***	0.0537***	0.0596***
	(0.00264)	(0.00272)	(0.00267)	(0.00391)	(0.00254)
	21.18	20.63	20.87	13.74	23.51
	(0)	(0)	(0)	(0)	(0)
Observations	1,354,448	1,354,448	1,354,448	776,888	577,560
R-squared	0.733	0.733	0.733	0.723	0.713
Region FE	YES	YES	YES	YES	YES
Time FE	YES	YES	YES	YES	YES
Area FE	YES	YES	YES	YES	YES

Robust se pval in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Table 2: Industrial vs occupational concentration

VARIABLES	(1) Baseline	(2) Occupational concentration
union = 1	0.0442***	0.0106***
	(0.00632)	(0.00327)
	6.987	3.234
	(2.56e-08)	(0.00253)
log(concentration)	-0.0176***	-0.0555
	(0.00620)	(0.0800)
	-2.841	-0.694
	(0.00719)	(0.492)
union x log(concentration)	0.0172***	0.221***
	(0.00276)	(0.0371)
	6.245	5.964
	(2.62e-07)	(6.38e-07)
lprod	0.0558***	0.0386***
	(0.00267)	(0.00164)
	20.87	23.57
	(0)	(0)
Observations	1,354,448	1,570,835
R-squared	0.733	0.729
Region FE	YES	YES
Time FE	YES	YES
Area FE	YES	YES
Robust se pval in parentheses		
*** p¡0.01, ** p¡0.05, * p¡0.1		

Table 3: Sample disaggregation by skill level

VARIABLES	(1) Baseline	(2) Low skill	(3) Medium skill	(4) High skill
union = 1	0.0442***	0.081***	0.068***	-0.023**
	(0.00632)	(.011)	(0.007)	(0.011)
log(concentration)	-0.0176***	-0.016***	-0.008***	-0.012***
	(0.00620)	(0.04)	(0.06)	(0.006)
union x log(concentration)	0.0172***	0.021***	0.025***	0.001
	(0.00276)	(0.005)	(0.004)	(0.005)
Observations	1,354,448	589,115	427,442	410,293
R-squared	0.733	0.728	0.644	0.518
Region FE	YES	YES	YES	YES
Time FE	YES	YES	YES	YES
Area FE	YES	YES	YES	YES
Robust se in parentheses				
*** p¡0.01, ** p¡0.05, * p¡0.1				

Table 4: Inclusion of worker outside option

Dependent variable: log weekly wages	(1)	(2)
union = 1	0.0442***	0.406*
	(0.00632)	(0.212)
	6.987	1.919
	(2.56e-08)	(0.0626)
log(concentration)	-0.0176***	-0.00978***
	(0.00620)	(0.00301)
	-2.841	-3.252
	(0.00719)	(0.00241)
union x log(concentration)	0.0172***	0.0149***
	(0.00276)	(0.00267)
	6.245	5.594
	(2.62e-07)	(2.05e-06)
log(outsideoption)		0.343***
		(0.0227)
		15.13
		(0)
union x log(outsideoption)		-0.0214*
		(0.0122)
		-1.748
		(0.0885)
lprod	0.0558***	0.0524***
	(0.00267)	(0.00285)
	20.87	18.39
	(0)	(0)
Observations	1,354,448	1,354,448
R-squared	0.733	0.737
Region FE	YES	YES
Time FE	YES	YES
Area FE	YES	YES
Robust se pval in parentheses		
*** p¡0.01, ** p¡0.05, * p¡0.1		