

DISCUSSION PAPER SERIES

DP13265

MONOPSONY IN THE UK

Will Abel, Silvana Tenreyro and Gregory Thwaites

**LABOUR ECONOMICS AND MONETARY
ECONOMICS AND FLUCTUATIONS**



MONOPSONY IN THE UK

Will Abel, Silvana Tenreyro and Gregory Thwaites

Discussion Paper DP13265
Published 20 October 2018
Submitted 16 October 2018

Centre for Economic Policy Research
33 Great Sutton Street, London EC1V 0DX, UK
Tel: +44 (0)20 7183 8801
www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **LABOUR ECONOMICS AND MONETARY ECONOMICS AND FLUCTUATIONS**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Will Abel, Silvana Tenreyro and Gregory Thwaites

MONOPSONY IN THE UK

Abstract

We study the evolution and effects of monopsony power in the UK private sector labour market from 1998 to 2017. Using linked employee-firm micro-data, we find that: (1) Measures of monopsony have been relatively stable across the time period examined - rising prior to the crisis, before subsequently falling again. (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 is greatly reduced and in most cases disappears. (5) The link between productivity and wage levels is weaker when labour markets are more concentrated.

JEL Classification: N/A

Keywords: monopsony, labour markets, market power, Unionization

Will Abel - William.Abel@bankofengland.co.uk
Bank of England

Silvana Tenreyro - s.tenreyro@lse.ac.uk
LSE and CEPR

Gregory Thwaites - gregorythwaites@yahoo.co.uk
LSE

Monopsony in the UK*

Will Abel[†], Silvana Tenreyro[‡] and Gregory Thwaites[§]

10th October 2018

Abstract

We study the evolution and effects of monopsony power in the UK private sector labour market from 1998 to 2017. Using linked employee-firm micro-data, we find that: (1) Measures of monopsony have been relatively stable across the time period examined - rising prior to the crisis, before subsequently falling again. (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 is greatly reduced and in most cases disappears. (5) The link between productivity and wage levels is weaker when labour markets are more concentrated.

*The views expressed in this paper are those of the authors and not necessarily those of the Bank of England. We would like to thank Nikola Dacic, Patrick Schneider and Michael McLeay for helpful comments. This work contains statistical data from ONS which is Crown Copyright. The use of the ONS statistical data in this work does not imply the endorsement of the ONS in relation to the interpretation or analysis of the statistical data. This work uses research datasets which may not exactly reproduce National Statistics aggregates.

[†]Bank of England (william.abel@bankofengland.co.uk).

[‡]Bank of England, LSE, CfM and CEPR (s.tenreyro@lse.ac.uk).

[§]LSE and CfM (g.d.thwaites@lse.ac.uk)

1 Introduction

The growing prominence of giant companies in advanced economies has raised concerns of increased monopsony power in labour markets. Monopsony power, a phenomenon first studied by Robinson (1933), arises when an employer faces little competition for workers from other firms. This allows the employer to set wages at a level lower than would be the case in a competitive market. Manning (2003) outlines how monopsonistic labour markets are more likely to emerge when employees have insufficient information on work opportunities and mobility costs are high.

In this paper we use employee level micro-data from the United Kingdom to study the evolution and effects of monopsony power in the UK private sector from 1998 to 2017. Using a 1% representative sample of all employees, along with unique firm identifiers, we are able to construct industry-region measures of concentration in employment over time. We find that while there are large variations in labour market concentration across industries and regions, aggregate measures of labour market concentration have been remarkably stable over the period we examine.

We document a complex relationship between wage growth, monopsony power, firm productivity and employee bargaining power. Increasing monopsony power, in general, is associated with lower pay. However this is dependent on the productivity of the firm the worker is employed at and on whether they are covered by a collective bargaining agreement (CBA). At high productivity firms, higher employer concentration is associated with lower pay for workers regardless of CBA coverage. For an employee covered by a CBA and at a firm with productivity at the median of the distribution, increases in concentration have little relationship with pay levels. For the same worker without CBA coverage, increases in employer concentration are again 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%.

Our approach and findings are very close to 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.

Azar et al. (2017) also look at the effects of monopsony power on wage growth in the US - finding a negative relationship between the two. However their results are in the magnitude of 10-15 times larger than those we find. This is potentially driven by many differences in our approaches - they focus on a small number of occupations (26), use a unique dataset of posted vacancies which spans only 4 years, and have fewer controls in their estimation.

2 Data Overview

In this section, we first introduce the data which we use in the analysis and then report descriptive statistics.

2.1 Data

Our data comes from 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 to 2017, where we know the pay, industry, occupation, union coverage, region and size of each individual's place of work. However 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 to 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.

Our data on union coverage only indicates whether an individual's pay makes reference to a CBA. For example this would cover non-union members whose employers have agreed to engage with unions (to any degree) during annual pay rounds. As such, the share of employees we observe as covered by a CBA is far in excess of those who are actually members of unions. 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 this with firm data from the Business Survey Database (BSD), which has annual turnover and employment for the universe of UK firms from 1997 to 2017. From this we construct measures of turnover per head which we use 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 to 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 we use NUTS2 regions. 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:

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

where $s_{j,ind,t,region}$ is the employment share of firm j in a given industry-year-region cell. Cells constructed with less than 10 individual observations are dropped leading to a loss of 75,000 observations; around 2.5% of our data.

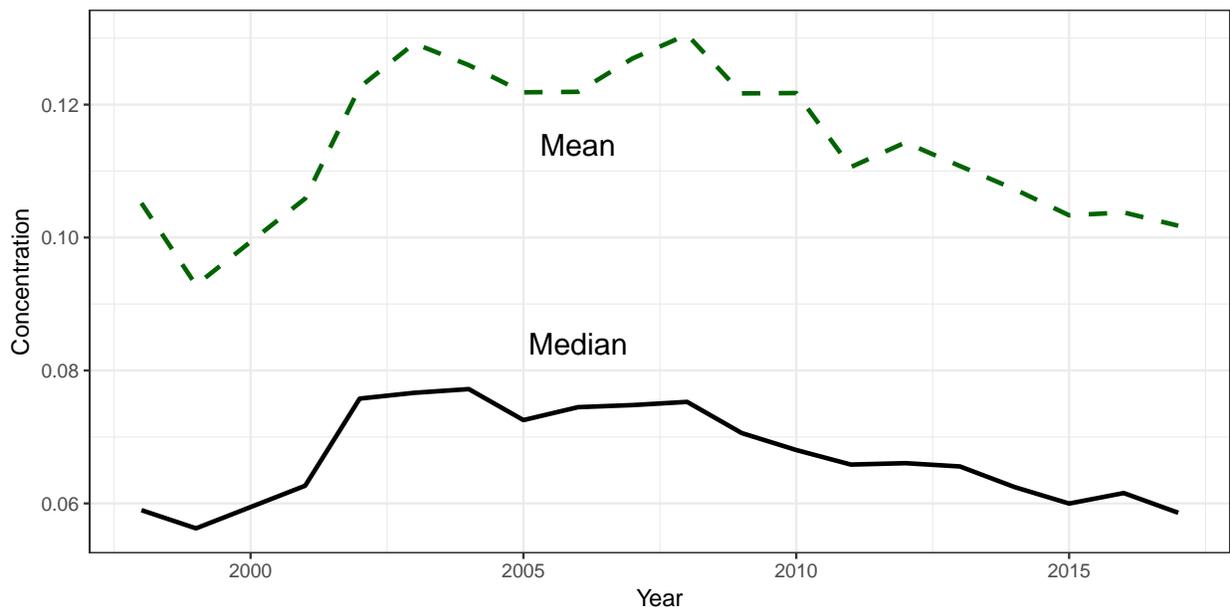
As our concentration measure is calculated from sampled data, as opposed to the complete population of employment, it will be noisy and biased. We discuss this in greater detail in appendix A1.

2.2 Descriptive statistics

Figure 1 shows a time series for the median and mean level of concentration faced by employees over time. 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. This masks a substantial rise and fall which occurred over this period though. Over the first decade of the series, mean and median concentrations increased by 24% and 28% respectively, from 1998 to 2008, and subsequently declined back to their starting levels.

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 1: Monopsony over time



The graph shows the mean and median of HHI across industries and regions over time.

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 A2, we show these same concentration figures sorted by region. While there is some regional variation, with the relatively sparsely populated regions of South Yorkshire, the Highlands and Cumbria being the most concentrated regional markets and West London, East Anglia and Oxfordshire being the least concentrated, we see that there is substantially more within rather than between-region variation. Even workers in parts of Manchester or London face highly concentrated labour markets, depending on their industry of work.

3 Econometric Analysis

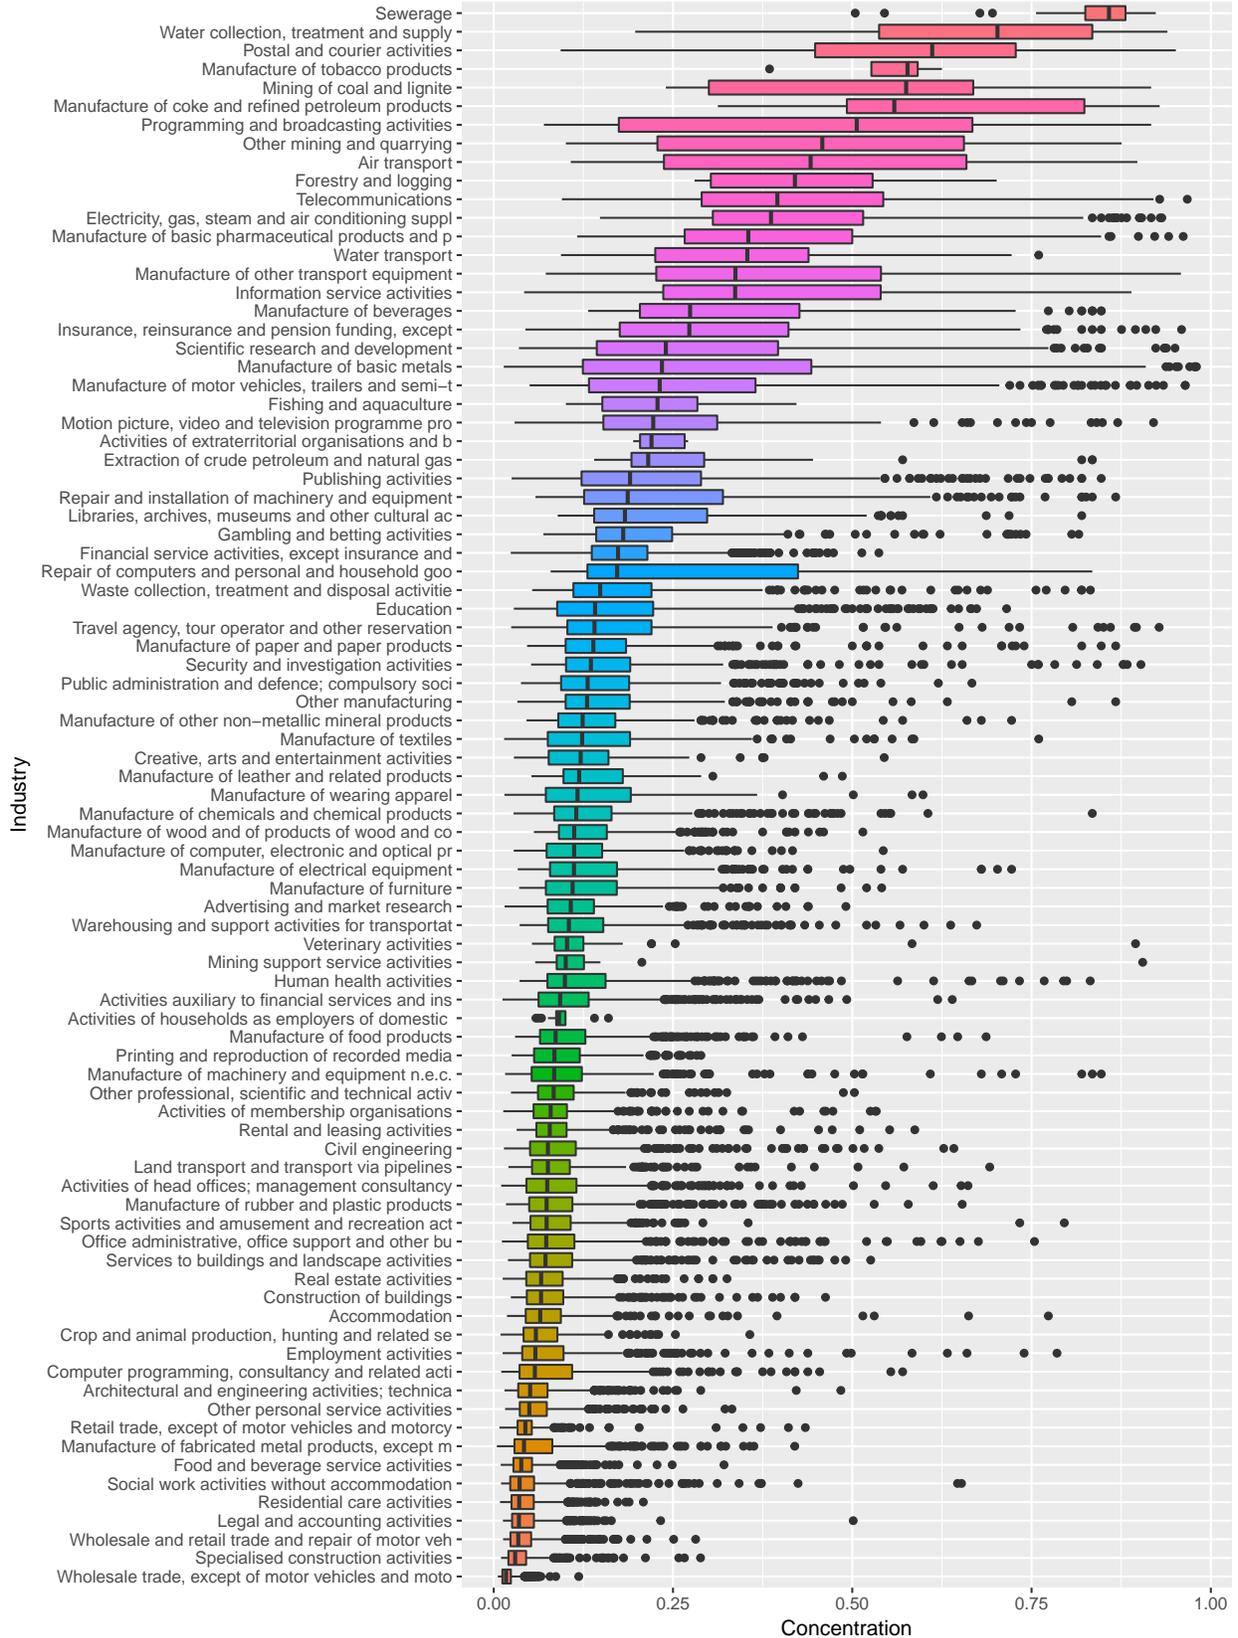
In this section, we first describe the specification of the main empirical model before presenting our results.

3.1 Main specification

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

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

Figure 2: Distribution of Concentration Measures by Industry



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 and fixed effects, including 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. The remaining controls are industry, occupation, region and year fixed effects. In later specifications of the model we also control for firm productivity, using turnover per head as a proxy - this is important as there is a likely inverse correlation between firm density and labour market concentration, such that in specifications where we do not control for this, our concentration measure may be picking up agglomeration affects.

We estimate a wage equation from the canonical search and matching model in Pissarides (2000):

$$w = (1 - \beta)z + \beta\rho(1 + c\theta) \quad (2)$$

where wages w are equal to the weighted average of z , the worker's period utility from unemployment, and their marginal productivity ρ plus the average cost of hiring an unemployed worker $\rho c\theta$. The weight on these two factors is determined by β , the worker's bargaining power. We assume that increasing monopsony acts through lowering workers' bargaining power β and thus allows employers to pay lower wages to employees. Were this the case, we would expect a negative coefficient on β_1 when we estimate equation 1. However, other factors may influence the degree to which employers are able to exercise this power. In particular, workers covered by a CBA may be able to resist employers monopsony power. We control for this in our central specifications by adding an interaction between CBA coverage and our measure of labour market concentration. We assume that both z and $c\theta$ are captured by our fixed effects.

Another implication of 2 is that the marginal effects of increased productivity ρ on wages should depend on β . We can test this by adding an interaction between labour market concentration and firm productivity. Again, a negative coefficient on this term is predicted by the model.

Our results for these specifications are displayed in Table 1. Column (1) shows the results of our main specification with no interaction effects. We find that concentration has no association with individual pay once we control for a range of fixed effects. However this masks a substantial degree of heterogeneity. Column (2) shows our main specification where we allow the effect of concentration to vary by whether a worker is covered by a CBA. For workers not covered by a CBA, we find a negative relationship between concentration and the level of pay. The size of the coefficient on concentration (-0.009) is such that, for a non-protected worker, moving from the 25th percentile to the 75th percentile of concentration would be associated with a decline in pay of 1.1%. The coefficient on the interaction between union membership and

Table 1: Effects of Labour Market Concentration on Pay

	Log weekly pay			
	(1)	(2)	(3)	(4)
Concentration	-0.004 (0.005)	-0.009 (0.004)	-0.009 (0.004)	0.092 (0.014)
Log (Turnover/head)			0.052 (0.002)	0.028 (0.005)
CBA coverage	0.016 (0.003)	0.051 (0.009)	0.040 (0.009)	0.045 (0.008)
Concentration*CBA coverage		0.013 (0.003)	0.009 (0.003)	0.011 (0.003)
Concentration*Log (Turnover/head)				-0.009 (0.001)
Constant	5.665 (0.046)	5.655 (0.046)	5.091 (0.032)	5.363 (0.050)
Years	1998-2017	1998-2017	2002-2017	2002-2017
Observations	1,661,647	1,661,647	1,426,850	1,426,850
Adjusted R ²	0.742	0.742	0.746	0.746

Notes: All models include occupation, industry, region and year fixed effects.

The list of individual controls are: gender, age, age squared, firm size, part time indicator, temporary contract indicator.

All models exclude the year 2000

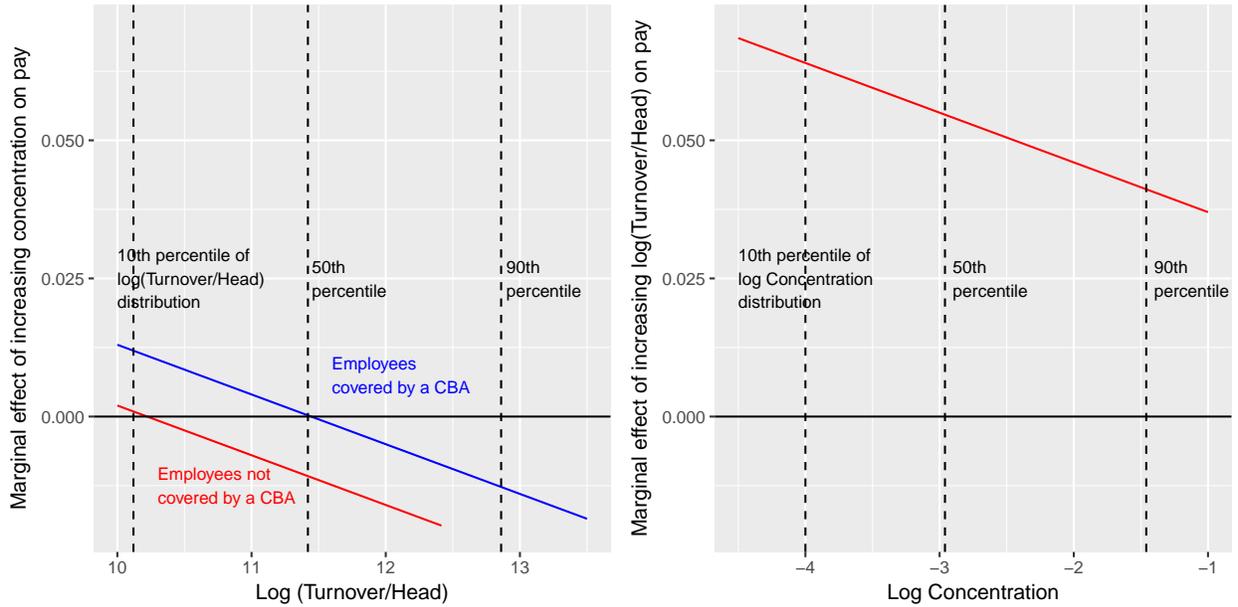
Standard errors are clustered at the region level

concentration almost exactly offsets this effect, suggesting that employer concentration has no relationship with the pay of workers covered by a CBA. This finding holds when we control for firm productivity as shown in column (3).

Column (4) allows the effect of concentration to vary both by CBA coverage and firm productivity. The relationship between concentration and pay is now positive and as per our previous specifications this effect is increased by CBA coverage. The strength of the relationship between concentration and pay declines with increasing productivity. The combination of these factors makes it unintuitive to immediately determine the marginal effects of both concentration and productivity on pay. To best illustrate the triple interaction effect, we plot the marginal effects of increasing concentration and productivity in the two panels of Figure 3. In the first panel we plot the marginal effect of increasing concentration on pay conditional on firm productivity. We have separated the effects by whether the employee is covered by a CBA and marked the 10th, 50th and 90th percentiles of the distribution of log turnover per head. For employees not covered by a CBA (the red line), increases in labour market concentration are correlated with lower pay, except for those in the least productive firms. For an employee not covered by a CBA and at the median of the log (Turnover/Head) distribution, an increase from the 25th percentile of concentration to 75th percentile would again be associated with a decline in pay of 1.1%. For employees who are covered by a CBA (the blue line) increasing labour market concentration is associated with higher pay for employees in firms with productivity below the median of the distribution and lower pay for employees in firms with productivity above this. For employees covered by a CBA and at a firm with median productivity, concentration and pay are uncorrelated.

In the second panel of Figure 3 we show the marginal effect of increasing productivity on pay, conditional on labour-market concentration. Productivity increases are associated with higher pay at all levels of observed labour-market concentration. As labour-market concentration increases the marginal effect of productivity on pay declines. This effect is not insubstantial - moving from the 25th percentile of concentration to the 75th percentile reduces the marginal effect by roughly 20%. This finding is consistent with more monopsonistic firms being better able to extract productivity gains from their employees.

Figure 3: Plotting Marginal Effects from Model 4



The left-hand panel shows the estimated marginal effect of labour-market concentration on pay varies with firm-level productivity and whether the employee is covered by a CBA. The right-hand panel shows how the estimated marginal effect of increasing turnover per head on pay varies with firm-level concentration.

4 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 the UK. We document that monopsony power increased from 1998 to 2008, before declining from 2008 to 2017, and then subsequently returning to levels broadly in line with those seen at the beginning 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. Finally we have shown that the link between productivity and pay is weaker for individuals who are employed in more monopsonistic labour markets.

References

- Azar, J., I. Marinescu, and M. I. Steinbaum (2017, December). Labor market concentration. (NBER 24147).
- Benmelech, E., N. Bergman, and H. Kim (2018, February). Strong employers and weak employees: How does employer concentration affect wages? (NBER 24307).
- Manning, A. (2003). *Monopsony in Motion: Imperfect Competition in Labor Markets*. Princeton University Press.
- Pissarides, C. A. (2000, March). *Equilibrium Unemployment Theory, 2nd Edition*, Volume 1 of *MIT Press Books*. The MIT Press.
- Robinson, J. (1933). *The economics of imperfect competition*. London Macmillan.

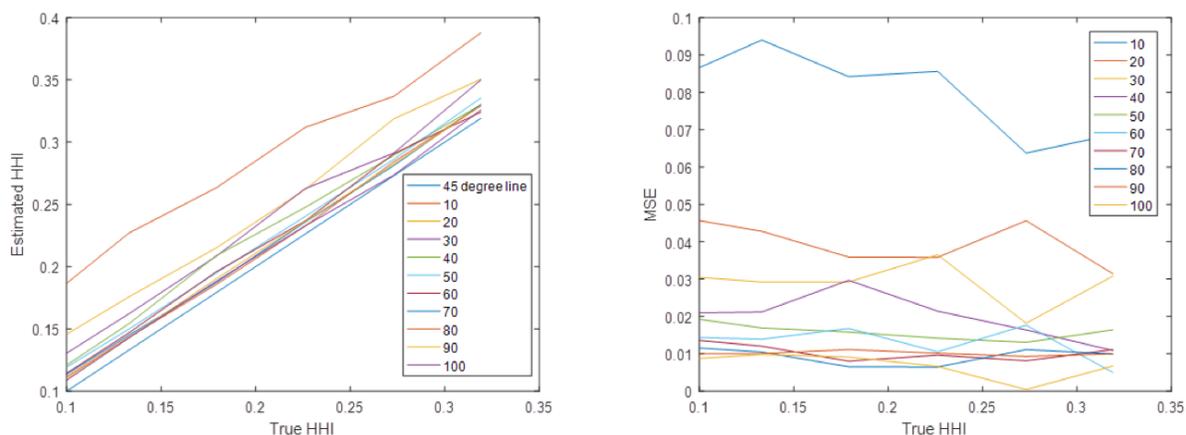
A1 Measuring HHIs with sample data

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 inequality - i.e. the fact that, given a set of unbiased estimates s_i of true market shares σ_i such that $E[s_i] = \sigma_i$, in general $E[s_i^2] > \sigma_i^2$.

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 A1 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.

A2 Distribution of Concentration Measures by Region

Figure A1: 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 how the mean square error of the estimated HHI depends on sample size and the true HHI.

Figure A2: Distribution of Concentration Measures by Region

