IFS

Institute for Fiscal Studies

Abi Adams-Prassl
Maria Balgova
Matthias Qian
Tom Waters

Working paper

23/14

Firm concentration & job design: the case of schedule flexible work arrangements

Economic and Social Research Council

Firm Concentration & Job Design: The Case of Schedule Flexible Work Arrangements

Abi Adams-Prassl, Maria Balgova, Matthias Qian & Tom Waters *

March 22, 2023

Abstract

We build a model of job design under monopsony that yields predictions over the relationship between: (i) the amenity value of non-wage job features; (ii) whether they are costly or profitable to firms; (iii) monopsony power. We analyse the amenity value of schedule flexibility offered in the labour market by combining our model's predictions with a new measure of schedule flexibility, which we construct from job vacancy text using a supervised machine learning approach. We show that the amenity value of schedule flexibility depends crucially on whether it is offered alongside a salaried contract that insures workers from earnings variation.

*Adams-Prassl: University of Oxford (abi.adams-prassl@economics.ox.ac.uk); Balgova: IZA (balgova@iza.org); Qian: University of Oxford (matthias.qian@sbs.ox.ac.uk). Waters: IFS (email: tom_w@ifs.org.uk). We are very grateful for financial support from UKRI (ES/S010424/1). Waters gratefully acknowledges co- funding from the ESRC-funded Centre for the Microeconomic Analysis of Public Policy (ES/M010147/1) at the IFS.

1 Introduction

A job is more than just a wage. From flexibility over when and where one works to the provision of health insurance and tuition assistance programs, there are many dimensions along which jobs vary. Workers and employers can have strong, and potentially conflicting, preferences over the non-wage characteristics of a job. Workers commonly cite non-monetary features as key determinants of their search behaviour (Adams-Prassl et al. 2023; Carrillo-Tudela et al. 2023) and have a significant willingness to pay for a wide range of job characteristics including freedom from harassment and dignity at work (Folke and Rickne 2022; Dube, Naidu, and Reich 2022). Non-wage dimensions of jobs can have a direct effect on firm profit; certain features may be costly to supply or have an impact on workplace productivity (Gibbs, Mengel, and Siemroth 2021; Etheridge, Wang, and Tang 2020).

Schedule flexibility is an increasingly salient dimension of job design. However, the extent to which the job flexibility offered in the labour market is a positive amenity that satisfies workers' demand for work-life balance, as opposed to a cost-saving measure for employers looking to shift risk onto employees, is a matter of fierce policy debate. As described in an independent review of working practises for the UK government, "being able to work when you want is a good thing; not knowing whether you have work from one day to the next when you have bills to pay is not" (Taylor et al. 2017). An experimental literature highlights that the amenity value of flexibility for workers depends crucially on whether workers or employers control the work schedule (Mas and Pallais 2017; Datta 2019). The "type" of flexibility that is realised in particular segments of the labor market is likely to depend on the relative bargaining power of firms and workers, although this has not been examined from a theoretical nor empirical perspective.

In this paper, we build a new model of endogenous job feature provision under monopsony to analyze the amenity value of job flexibility on offer across the wage and occupational distribution. While there is an active literature on the relationship between employer power and wage setting, the impacts of monopsony power on the provision of non-wage job features is less developed. Our model yields clear predictions on the relationship between the amenity value of a job feature to workers, whether it is costly or profitable for firms to supply, and the degree of monopsony power. We apply the insights of our model to assess the amenity value of flexibility by analyzing the impact of exogenous changes in employer concentration on the provision of flexibility in a local labour market.

In the model, we allow jobs to be characterised both by their wage and by their non-monetary features. Non-monetary job features can increase workers' utility, e.g. employee-led flexibility, or decrease workers' utility, e.g. employer-led flexibility. These features can also be profitable or costly to the firm. This gives rise to three types of job features: (i) costly amenities, (ii) profitable disamenities, and (iii) profitable amenities.[1] We assume that higher employer concentration reduces workers' outside option, which in turn changes their relative preferences over monetary and non-monetary features. Depending on whether the feature is an amenity or a disamenity to the employee, and whether it is costly or profitable for the firm to provide, we prove the conditions under which the provision of a given job feature will systematically increase or decrease with firms' monopsony power.

We apply the insights of our model to assess the amenity value of schedule flexibility available to workers in practise. There is little consistent evidence on the "type" of flexibility experienced by different groups of workers in the labour market. Administrative data rarely captures non-wage job characteristics, while survey questions on work arrangements have not been asked consistently over time and rarely ask about who controls the variability in one's work schedule (Mas and Pallais 2020). It is more common for surveys to elicit information on the contractual label associated with a job (e.g. annualized hours, on-call working). However, different groups of workers may have very different experiences of contractually similar schedule flexible jobs due to variation in their relative bargaining position. The implications for worker welfare might be particularly pronounced for those without a fixed salary; in such cases, the number and timing of hours performed each week translates directly into the level and timing of income.

Given the absence of a consistent measure of schedule flexibility over time in survey data,[2] we construct a new data set of flexible jobs from job vacancy postings. We analyse the full text of more than 46 million online job vacancies posted in the UK between 2014 and 2019 collected by Burning Glass Technologies (now 'Lightcast'). We take a supervised machine learning approach to classify the work arrangements described in the vacancy text because relying on a keyword search for "flexibility" can lead to both false positives ("We are looking for a motivated, flexible & committed individual") and false negatives ("Working hours will vary according to the needs

[1] Costly disamenities should not exist – perhaps outside of some legal obligations – because neither the employer nor the worker benefits.

[2] There have been attempts to measure some dimensions of flexibility, e.g. telework potential, by analyzing tasks at the occupational level. However, this approach does not facilitate a within-occupation analysis to analyze, e.g. heterogeneity across wages.

of the business"). We manually label the work arrangements described in ∼7,000 job vacancies and use these as the training data set for a logistic classification model with Lasso regularization, selecting the tuning parameter by a grid search over a 5-fold cross-validated balance F-measure. Our approach achieves high predictive accuracy and yields an occupational distribution of flexible work arrangements that is strongly correlated with a high-quality cross-section measure of job flexibility that the authors designed in a nationally representative survey.[3]

When piloting our approach, we found that it wasn't possible consistently to identify whether workers or employers had control over the variation in schedules from the vacancy text of flexible jobs. For example, does a work arrangement described as a "Casual contract! Allows for flexibility" describe a worker or employer determined schedule?[4] Thus, in addition to flexibility, we also identify whether a job vacancy is salaried or not (i.e. a fixed income versus a wage contract). Conceptually, unsalaried schedule flexible jobs open the worker up to the possibility of an unpredictable, precarious income stream, because flexibility in the timing and number of hours worked can translate directly into earnings variation. These "risky" flexible jobs are in contrast to salaried flexible jobs which come with a fixed income, and where schedule flexibility only impacts the timing of work but not of remuneration. As a result, while remuneration type and schedule control are separate concepts, their impact on worker welfare is intertwined since unsalaried workers are more exposed to the consequences of employer-control over schedule. In the paper, we refer to unsalaried and salaried flexible positions as "risky flexible" and "safe flexible" jobs respectively.

We first provide new stylized facts on the characteristics of flexible jobs. We find that schedule flexibility is relatively widespread, even in the years before the Covid-19 pandemic: 30% of all vacancies in the UK in 2019 came without a fixed schedule. About half of these jobs are risky. We document considerable differences between the characteristics of risky flexible, safe flexible, and fixed-schedule jobs. Safe flexible jobs are more likely to be offered in higher-paying occupations and require significantly more skill, while the opposite is true for risky flexibility. In occupations such as caring and leisure services, elementary jobs, and process operatives, more than 40% of vacancies come with a flexible schedule, and more than 80% of these flexible jobs are risky.

Even within occupations and local labour markets, we find systematic differences in the wage and skill content of different types of flexible jobs. Both risky and safe flexible jobs are associated

[3]The measure was designed by Adams-Prassl et al. (2020) and was also included in the Understanding Society household panel surveys over the summer of 2020.

[4]Furthermore, who controls the job schedule might vary over time within an employment relationship given changes in the relative bargaining power of workers and firms.

with a wage penalty compared to fixed-schedule positions. Risky flexible jobs are the worst remunerated on average, with a 10 p.p wage gap on average with fixed-schedule, salaried positions. This wage gap for risky flexible jobs is not simply a result of variation in the prevalence of salaried contracts; risky flexible jobs are paid less than otherwise similar non-salaried fixed-schedule jobs. We find that flexible jobs are advertised alongside a broader set of skills requirements. However, the type of flexibility again matters. Risky flexible vacancies are associated with the fewest number of skills demanded, less than both safe flexible and non-flexible vacancies.

To test the predictions of our theoretical model and assess the amenity value of flexibility, we estimate the causal relationship between employer concentration and the prevalence of flexibility in a local labour market. We construct the Herfindahl-Hirschman Index (HHI) of employer concentration using job vacancy data at the 3-digit occupation-county level as our measure of monopsony power (Azar, Marinescu, Steinbaum, and Taska 2020). In addition to including a rich set of fixed effects to control for any systematic variation across local labour markets that might be driving the concentration-flexibility relationship, we also instrument for employer concentration with a predicted HHI measure based on nationwide shocks to firm-level hiring. The motivating idea behind this instrumental variable approach is that a firm's local hiring will partly reflect its conditions nationally, which are in turn independent of changes in local conditions.

We find a robust positive relationship between employer concentration and schedule flexibility. As workers' outside options deteriorate, firms are more likely to offer jobs with flexible schedules. However, estimating the relationship separately for risky and safe flexible vacancies shows that different types of flexibility behave differently in response to a change in concentration. While risky flexibility rises with employer concentration, the prevalence of safe flexible jobs changes little with worker outside options. Interpreted through the lens of our model, these results suggest that risky flexibility is a disamenity that increases firms' profits, while safe flexibility is an amenity. We find that the positive relationship between employer concentration and risky flexibility is driven by those occupations in which employer-controlled flexibility is most prevalent in a one-off nationally representative survey. This is reasonable: in occupations where workers can control their schedules, risky flexibility is not such a disamenity.

Our results highlight that schedule flexibility is not the same job feature for all workers. Risky flexibility is much more prevalent at low wages, in less-skilled occupations, and in less competitive labour markets. This feature is a disamenity for workers, although one that is productive for firms. Yet, safe flexibility is an amenity and is the dominant form of flexibility at high wages and in high skilled occupations. Our results suggest that the distribution of welfare-relevant job features accentuates overall inequality in the labour market. Furthermore, our findings suggest different mechanisms lie behind the wage penalty for safe and risky flexible jobs compared to fixed-schedule jobs. Safe flexible vacancies pay less because workers are willing to give up part of their wages for an important amenity, while risky flexible vacancies are paid less because they are more likely to be offered in labour markets with poor outside options for workers (Hamermesh 1999; Bonhomme and Jolivet 2009; Hwang, Mortensen, and Reed 1998).

This paper makes three main contributions to literature. First, we contribute to the emerging literature on alternative work arrangements (Mas and Pallais 2020; Datta, Giupponi, and Machin 2019; Prassl 2018; Boeri et al. 2020). Recent research estimates worker preferences over flexible arrangements using experimental variation, finding that the average worker is not willing to pay for schedule flexibility and prefers the characteristics of "traditional jobs" (Mas and Pallais 2017; Datta 2019). We complement this work by showing that worker preferences for job features can be inferred from their relationship with employer concentration. This facilitates an analysis of the amenity value of flexibility based on observational data beyond a tightly controlled setting. We show that the welfare impact of schedule flexibility depends on its interaction with remuneration type: risky flexibility is a profitable disamenity, while safe flexibility is an amenity. These types of flexibility are unevenly distributed across the wage and occupational distribution in a manner that accentuates overall inequality in worker welfare.

Second, this paper makes a theoretical and empirical contribution to the growing literature on monopsony in the labour market. While there are several models that use firms' provisions of amenities as a microfoundation for monopsony power (Card et al. 2018; Lamadon, Mogstad, and Setzler 2022), our model does the opposite: it takes monopsony power as given, and describes the optimal provision of wages and non-monetary remuneration under these conditions. Following recent work on the impact of employer concentration on workers' wages (Benmelech, Bergman, and Kim 2020; Azar, Marinescu, Steinbaum, and Taska 2020; Rinz 2022; Schubert, Stansbury, and Taska 2022), we show that it reduces non-monetary compensation. An analysis of wages alone is therefore insufficient to capture the welfare impact of monopsony power.

The closest paper to our own in this literature is recent work by Dube, Naidu, and Reich (2022). Their focus is on the relationship between dignity at work and monopsony power, and how this might mediate the impact of changes in the minimum wage. They show that under monopsony, the prevalence of positive amenities could rise or fall depending on the interaction of amenities and wages in the firm's profit function. Our papers nicely complement each other. We emphasise different microfoundations for monopsony power[5] and have different aims: Dube, Naidu, and Reich 2022 focus is on the interaction of amenities and wages with changes in minimum wage rates, while our focus is on the relationship between the extent of monopsony power and the prevalence of different types of job features.

Third, we show how job vacancy text represents a rich source of information about the non-wage characteristics of jobs. Online job vacancy data has been used to analyze a growing range of topics, from firm demand for skill (Hershbein and Kahn 2018; Deming and Kahn 2018; Clemens, Kahn, and Meer 2020), to changes in firm hiring strategies to public policy interventions (Marinescu 2017; Duchini, Simion, and Turrell 2020), and economic shocks (Javorcik et al. 2020; Forsythe et al. 2020). This is one of the first papers to use machine learning tools to exploit the full text of the vacancy, rather than performing a keyword search (Duchini, Simion, and Turrell 2020) or focusing only on job titles (Marinescu and Wolthoff 2020). In a recent paper, Hansen et al. (2023) apply frontier natural language processing methods to extract information on working from home (WFH) from job vacancy text. In the case of WFH, they find that their approach can offer substantial accuracy improvements over alternative methods. The machine learning approach that we take in this paper is based on a logistic classifier and is much simpler and less computationally demanding to implement. However, Hansen et al. (2023) find that our simple approach yields an identical accuracy to their model in the case of flexible scheduling (Footnote 20, Hansen et al. (2023)) validating our measurement strategy.

The rest of this paper proceeds as follows. In Section 2 we develop our model of endogenous job feature provision under monopsony. In Section 3, we introduce the job vacancy data, describe our classification of schedule flexible jobs, and outline the machine learning algorithm used to identify the job features specified in the vacancy text. In Section 4, we summarise the main descriptive patterns of flexible vacancies in our data and describe the differences between safe flexible, risky flexible, and fixed-schedule vacancies. In Section 5 we analyze the relationship between changes in employer concentration and different types of flexibility to identify the amenity value of these job features. Section 6 concludes.

[5] Dube, Naidu, and Reich (2022) consider horizontal job differentiation while we specify a granular search model.

2 Theoretical Framework

We develop a model of the equilibrium provision of non-monetary job features and wages under monopsony. When a worker and firm meet, they bargain over wage-job feature bundles. A worker's outside option in this bargaining problem is a function of employer concentration in the local labour market. The model yields predictions about the relationship between employer concentration and job features that differ according to whether a job feature is utility-enhancing or utility-diminishing for workers, and whether it is profitable or costly for firms to supply. As our focus is on the relationship between wages and non-monetary job features in this bargaining problem, we consider a static framework; a dynamic model of search behaviour would not yield much insight into the core question of interest, while adding significantly more complication.

2.1 Set Up

Jobs are characterized by the wage, w, and a continuous non-monetary job feature, f. Workers and firms have preferences defined over bundles of job characteristics, $\{w, f\}$. For simplicity we assume that firms and workers are homogeneous, and all firms are equally productive. Workers are characterised by the utility function $u(w, f)$. While utility is always strictly increasing in the wage ($u'_w > 0$), a job feature can be either an amenity ($u'_f > 0$) or a disamenity ($u'_f < 0$). A firm's profit function is given by:

$$\pi = y + \delta(f) - w \tag{2.1}$$

where y gives the output of the match and $\delta(f)$ captures the relationship between the job feature and profits. Provision of a feature can be either profitable or costly; it is profitable when $\delta(f)$ is positive and costly when it is negative.

This set-up gives rise to three types of job features:

(i) Costly amenities, which are utility-increasing but costly for the firm to provide, e.g. designer offices and a company car.

(ii) Profitable disamenities, which increase profits but at the expense of lowering workers' utility, e.g. non-compete clauses in employment contract.

(iii) Profitable amenities, which directly benefit both workers and firms. The option to work from home, if it makes workers more productive, is an example of a profitable amenity.

Note that the fourth logical combination, costly disamenities, should not exist in equilibrium (absent legal obligations) because neither the firm nor the worker benefits from their provision.

Outside Options Firms post vacancies to fill empty positions, and unemployed workers search to find these vacancies.[6] The labor market is characterised by a fixed, finite number of firms, each with a potentially different exogenous size, N_i. Let the employment share of firm i be given by $s_i = N_i / \sum_j N_j$. Assuming exogenous job destruction and no on-the-job search, a firm's vacancy share also corresponds to its employment share.[7] Workers match to one random vacancy with probability λ each period. Therefore, the probability that a worker matches with a vacancy from a particular firm i increases in firm size (since bigger firms post more vacancies): λs_i.

We assume that firms exercise a degree of size-based monopsony power.[8] For concreteness, we follow Jarosch, Nimczik, and Sorkin (2019) to ground this power in a credible threat by the firm not to hire a worker in the future if current bargaining breaks down (although in equilibrium this will never happen). This implies that other vacancies advertised by the firm do not form part of a worker's outside option. However, any theory providing a link between a worker's outside option and employer concentration will give the same central predictions as our framework. When bargaining breaks down, or when the worker doesn't receive any job offer, she receives unemployment benefit b.

Given this set-up, the outside option of a worker bargaining with firm j is:

$$U_j = \lambda \sum_{i \neq j} s_i u(w_i, f_i) + (1 - \lambda + \lambda s_j) u(b, 0) \tag{2.2}$$

$$= \lambda \bar{u} - \lambda s_j u(w_j, f_j) + (1 - \lambda + \lambda s_j) u(b, 0) \tag{2.3}$$

where $\bar{u}$ is the expected utility of jobs available on the market. The average outside option (or

[6]There is no on-the-job search, and no firm growth, i.e. firms only advertise vacancies to find replacement for destroyed matches.

[7]We will often refer to s_i as firm size.

[8]This is in contrast to approaches that ground monospony power in the provision of non-wage amenities (Lamadon, Mogstad, and Setzler 2022; Dube, Naidu, and Reich 2022).

equivalently the ex ante outside option before the worker receives a particular job offer) is:

$$\begin{aligned}\bar{U} &= \sum_i s_i U_i \\ &= \lambda \bar{u} + (1-\lambda)u(b,0) - \lambda \sum_i s_i^2 \left[u(w_i, f_i) - u(b,0)\right] \end{aligned} \tag{2.4}$$

This expression shows a direct link between the degree of employer concentration in the market, $\sum_i s_i^2$, and the average outside option of a worker. A useful feature of the model is that $\sum_i s_i^2$, the measure of monopsony power, corresponds to the standard measure of employer concentration, the Herfindahl-Hirschman Index.

Theorem 1 *A worker's average outside option falls in employer concentration, $\sum_i s_i^2$.*

Proof. Theorem 1 follows directly from the expression of average outside option of a worker. The relationship between $\bar{U}$ and $\sum_i s_i^2$ is negative as long as $u(w_i, f_i) - u(b, 0) > 0$ for all firms i. Since the worker has no incentive to accept a job leaving her worse of than remaining unemployed, this inequality will always hold, and workers' outside option will always decline in employers' monopsony power. ■

Bargaining In standard models of job search, the meeting of a worker and a firm is followed by bargaining over the optimal split of the match surplus – the optimal wage. In our model, the worker and firm also bargain over the optimal level of the job feature. As job features have a direct impact on utility and profits, the bargaining process also determines the *size* of the match surplus. When the optimal amount of a job feature varies with wage, the size and the split of the match surplus are determined jointly, creating a link between the provision of a job feature and a firm's monopsony power.

In the main text, we solve the model for the representative firm and worker as this provides the most direct link to our market-level empirical analysis. In Appendix B, we solve the model conditional on firm size: in this case, the optimal $\{w, f\}$ depends on s_i, since workers bargaining with larger firms have weaker outside options. Given the assumption that firms are homogeneous up to their (exogenous) size, abstracting from heterogeneity in s_i does not change the main mechanism between monopsony power and amenity provision at the heart of the model.

Once matched with a vacancy, a firm and worker bargain over the $\{w, f\}$ bundle that maximises the match surplus and split it optimally, given their bargaining power and outside options. Modelling this process as standard Nash bargaining, the problem can be represented as:

$$\max_{w,f} \quad [u(w,f) - \bar{U}]^{\beta}[y + \delta(f) - w]^{1-\beta} \tag{2.5}$$

where β denotes a worker's bargaining power, and $\bar{U}$ is their average outside option (Equation 2.4) given our focus on the representative scenario. The outside option of the employer, as well as the cost of posting the vacancy, is set to 0.[9]

2.2 Solution

We derive the conditions under which there exists a unique, interior solution to this bargaining problem. We are specifically looking for interior solutions as, empirically, the proportion of jobs offering the job features we study in this paper is very rarely 0 or 1. Let the optimal wage-feature bundle be represented as $\{w^*, f^*\}$. The worker and firm will optimally decide to set $f^* > 0$ as long as its net impact on the match value is positive, i.e.

$$u(w^*, f^*) - u(w^{*\prime}, 0) + \delta(f^*) + w^* - w^{*\prime} > 0 \tag{2.6}$$

where $w^{*\prime}$ is what the optimal wage would be if f were constrained to be zero. This implies that each side benefits from a job feature even if its direct impact on utility or profits is negative. For example, a disamenity could lower workers' utility directly, but if it is profitable, then its provision could still leave them better off through the wage bargain. In the bargaining process, the increase in match surplus due to f is split the same way as match output y.

The optimal wage-job feature bundle is characterised by a familiar optimality condition. The worker's marginal rate of substitution between the job feature and wage must equal the marginal rate of transformation on the part of the firm:

$$\frac{u'_f(w^*, f^*)}{u'_w(w^*, f^*)} = -\delta'_f(f^*) \tag{2.7}$$

[9] This is a simplifying assumption, in line with our earlier assumption that firms are fixed in size and only hire to replace destroyed matches. For more details, see Schubert, Stansbury, and Taska 2022.

We present this optimality condition for each type of job feature graphically in Figure 1. Figure 1 also demonstrates the necessary conditions that the utility and profit functions must satisfy for the problem to have a unique interior solution. We formalise these conditions in Theorem 2.[10]

Theorem 2 *A unique interior solution to the bargaining problem exists for any type of job feature f if*

(i) there is some interval for $\delta(f)$ on which $\mathrm{sgn}(\delta'_f) \neq \mathrm{sgn}(u'_f)$

(ii) the second derivative of the job feature profit function $\delta(f)$ is negative: $\delta''_f < 0$

(iii) the second-order derivatives of the utility function $u(w,f)$ are weakly negative: $u''_w \leq 0, u''_f \leq 0$

(iv) the sign of the second-order cross-derivative of the utility function is the same as the sign of marginal utility of job feature[a]*:* $\mathrm{sgn}(u''_{wf}) = \mathrm{sgn}(u'_f)$

Proof. See Appendix B. ■

[a]This includes the possibility that $u''_{wf} = 0$.

The conditions underlying Theorem 2 are satisfied for any standard utility function, and for any profit function with enough curvature. Starting with Condition (ii), it implies that marginal profit must be decreasing with respect to f: i.e. $\delta(f)$ is strictly concave when f is profitable, and convex when f is costly. Condition (iii) imposes a similar requirement on the utility function: it must be concave in wages and amenities, and convex for disamenities. Conditions (ii) and (iii) are consistent with standard utility and profit functions: costly (or utility-decreasing) goods become more costly (or more unpleasant) as their quantity increases, while profitable and utility-increasing goods are subject to decreasing marginal returns. The implication of condition (iv) is similarly intuitive: it states that the job feature is a complement for wage when it is an amenity, and a substitute when it is a disamenity; in other words, as the wage increases, workers will demand more amenities and fewer disamenities.[11] Finally, condition (i) ensures that the unique solution

[10]Note that the existence of a unique equilibrium, including corner solutions, may not require all the conditions of Proposition 1 to hold. The existence conditions are also more stringent for profitable amenities than for other types of job features. For further discussion, see the Proof of Proposition 1.

[11]We define "complements" and "substitutes" using the q-concept: two goods are complements if the marginal utility of one increases with the quantity of the other good, and substitutes if the marginal utility of one decreases in the quantity of the other. Importantly, q-complementarity is uncompensated. See also Seidman 1989.

Figure 1: Solutions to the bargaining problem for different types of job features

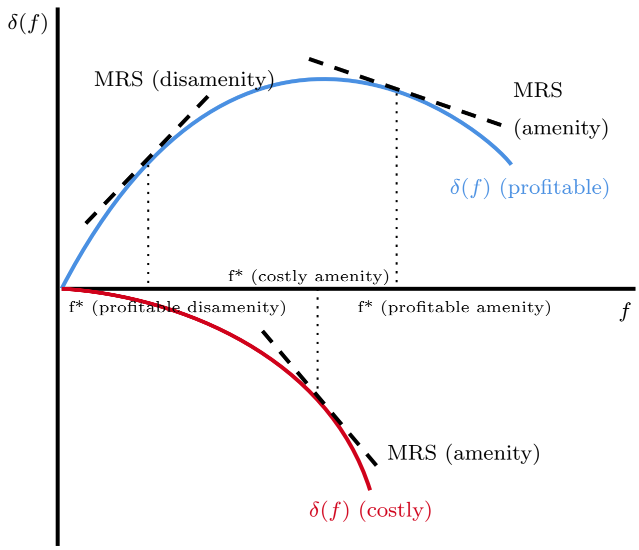

Notes: The figure presents possible solutions to the bargaining problem for the three types of job features. The solution is always characterised by points where the worker's MRS equals the (negative of the) marginal impact of a job feature on firm profit. The solution for profitable job features is found in the upper half of the diagram; the solution for costly features is found in the lower half where $\delta(f)$ is negative. The MRS for profitable disamenities is positive (workers must be compensated for increases in a disamenity with a higher wage) while for profitable amenities, the MRS is negative.

is an interior one. It follows from Equation 2.7 (i.e. $MRS = MRT$) which requires that there is some f at which marginal profit and marginal utility are of opposite signs. This is satisfied automatically for costly amenities and profitable disamenities; for the case of profitable amenities, it means that the marginal profitability turns negative at some f.

2.3 Job Features & Employer Concentration

The model implies a systematic relationship between the properties of a job feature, f, and a worker's outside option, $\bar{U}$. Specifically, the provision of profitable disamenities will rise when outside options are lower, while the provision of amenities (especially costly ones) increases with $\bar{U}$. As we show below, the sufficient condition for this to hold is that utility is not quasi-linear in the wage, and job features and the wage are not perfectly substitutable.

Theorem 3 *The optimal quantity of a job feature, f, increases in a worker's outside option, $\bar{U}$, when f is an amenity, and decreases when f is a disamenity if:*

$$\frac{\partial^2 u(\cdot)}{\partial w^2} \neq 0 \quad or \quad \frac{\partial^2 u(\cdot)}{\partial w \partial f} \neq 0 \tag{2.8}$$

and

$$\frac{\partial u^2 / \partial w \partial f}{\partial u / \partial f} < \frac{2}{\beta \pi(w^*, f^*)} \tag{2.9}$$

The size of this effect is falling in $\delta(f)$, the profitability of supplying the job feature. Therefore, the impact of changes in $\bar{U}$ on the provision of costly amenities is greater than its impact on the provision of profitable amenities.

Proof. See Appendix B. ■

Theorem 3 specifies the conditions under which firms' monopsony power affects the provision of job features. Condition 2.8 rules out quasi-linearity of utility in w. If the utility function is linear with respect to w, workers can be compensated for the increase in their outside option by increasing their wage, and this increase in utility will not impact the workers' preferences over the optimal quantity of the job feature.[12] Condition 2.9 places a restriction on the substitutability of

[12] In other words, utility that is quasi-linear in flexibility would imply that workers' preferences over flexibility (and stability of income) don't vary with their wages. While this is a theoretical possibility, the fact that schedule flexibility (and variability of income) tend to be an important dimension of the job makes this unlikely. When

wages and a job feature. For a worker to demand more of an amenity f when her outside option improves, the substitutability between wage and f must not be "too big" — otherwise the worker would demand a higher wage before the job amenity.

If the conditions underlying Theorem 3 hold, it possible to identify the type of a job feature from its relationship with employer concentration. If a job feature increases in employer concentration, then it is a profitable disamenity. Where the opposite is true, the job feature is an amenity. When the relationship between the provision of a job feature and employer concentration is close to 0, the job feature is a profitable amenity. This will be true in alternative search models that also give rise to a strictly decreasing relationship between $\bar{U}$ and employer concentration.

Theorem 4 captures that monopsonistic firms under-supply amenities and over-supply disamenities relative to the no-monopsony power benchmark in which firms are price-takers and workers find jobs immediately. Intuitively, workers' remuneration is higher under perfect competition. Given our assumptions on preferences, this means that workers in a competitive labour market enjoy more amenities, and fewer disamenities, than workers in monopsonistic labour markets.

Theorem 4 *If Theorems 2 and 3 hold, firms provide a lower level of amenities and a higher level of disamenities than under a setting where firms have no monopsony power over workers.*

Proof. See Appendix B. ■

3 Measuring Flexible Work Arrangements

We harness our model to provide new insights into the amenity value of flexible work arrangements in the labour market. Existing experimental work has shown that workers' preference for schedule flexibility depends crucially on whether they or their employer has control of the schedule (Mas and Pallais 2017; Datta 2019). When workers have control of the schedule, flexibility is typically found to be an amenity. However, if it is the employer who has the flexibility to vary schedules, the job feature is a disamenity. Yet, little is known about the characteristics and amenity value of the flexible work arrangements that are actually on offer in the labour market, and how this varies

utility function is not quasi-linear, the optimal wage and optimal f are directly linked and a worker's marginal rate of substitution between w and f will vary as their outside option, and overall compensation, changes.

across the wage distribution and by occupation. There is limited consistent time series evidence on the prevalence of schedule flexibility, let alone whether variation in schedules is controlled by the worker or employer (Adams, Prassl, et al. 2018; Abraham and Amaya 2019; Mas and Pallais 2020).

To make progress, we use two novel sources of data on flexible work arrangements. Specifically, we analyze the full text of the near-population of job vacancies advertised online in the UK labour market from Burning Glass Technologies and identify the characteristics of the advertised work arrangements using supervised machine learning methods. This yields a measure of schedule flexibility that is consistent across occupations, local labour markets, and time. In Section 5, we will use variation in these work arrangements with employer concentration to identify the amenity value of the flexibility offered across different occupations. We complement our vacancy-based flexibility measure with (one-off) high quality, representative data on job flexibility from Understanding Society, a large UK household survey. This allows us to check the quality of our vacancy based measure and to provide additional insights into the characteristics of schedule flexible jobs that are not captured by vacancy data.

3.1 Data

Burning Glass Technologies Data We use the UK NOVA data feed from Burning Glass Technologies (BGT) as our corpus of job vacancies, a widely used database that aims to cover the universe of online job postings. Since 2012, BGT has collected the near-universe of electronically posted job vacancies in the UK by scraping 7,500 online job boards and company web pages. We restrict our analysis to data from 2014; between 2012 and 2014, the number of webpages that BGT scraped for job adverts increased rapidly. However, the source of vacancy information has been stable since this date. Our analysis is based on the full vacancy text of over 46 million unique online job postings between 2014 and 2019. Remuneration information is extracted from job postings using text and number patterns. If wages are specified in the vacancy text, then standardized language is used to specify the salary band. For example, the wage can be stated in a section titled "Salary", or "Remuneration Structure". Skills are identified with keyword patterns, which represent inclusion and exclusion criteria. For example, the keyword *python* would appear as a computer programming skill if it is followed by *programming*, but not if it is followed by *hunter*.

As with all studies that make use of job vacancy data, it is important to note that the data comes with some caveats. The source only captures jobs advertised online. However, some jobs are not advertised at all and others may only be advertised on offline or closed portals. Table A.2 gives the headline summary statistics for the BGT data and their closest counterparts from representative surveys. BGT captured 83% of all vacancies posted in the UK over this time period (as given by the Office for National Statistics' official vacancy count). The BGT wage distribution closely aligns with the distribution of wages of employees (i.e. a stock measure) in the Annual Survey of Hours and Earnings (ASHE), a 1% sample of employee jobs in the UK.[13] 37% of BGT vacancies omit wage information but we demonstrate there are no statistically significant observable differences between vacancies with and without posted wages (Table A.3). Our main empirical specification will rely on variation in firm concentration and schedule flexibility within occupation-county cells over time. Figure A.8 analyzes whether there has been any change in the representativeness of BGT data over time within these cells. It shows that there has been no change in the relationship between the number of vacancies in each occupation-county cell (from BGT data) and the number of employees in the same cell as given by ASHE.[14]

Understanding Society To check the quality of our measure of flexibility, we use one-off survey data from Understanding Society. Understanding Society is one of the largest household panel studies in the world. The main panel is known to provide a strong basis for population inferences (Crossley, Fisher, and Low 2021). In response to the Covid-19 pandemic, the survey introduced a special web-based survey in April 2020. All individual members of households enrolled in the existing panel who were aged sixteen or over in April 2020 were invited to participate in the COVID-19 Study.

We make use of responses of individuals to the April to September 2020 waves of the Covid-19 study who reported that they were employed in January or February 2020 (i.e. pre-pandemic). These waves included the following retrospective question on job flexibility and schedule control for a respondent's main job pre-pandemic:[15]

[13] The share of jobs that are permanent in the BGT data is considerably less than in ASHE, but this likely reflects the fact that by their very nature, temporary jobs have a higher churn rate and so are overrepresented in the flow of vacancies than the stock of employees.

[14] Specifically, we regress the number of vacancies in each occupation-county cell on the number of employees in the same cell as given by ASHE. We run this regression separately for each year of BGT data, and plot the resulting R-squared, which allows us to see whether the predictiveness of ASHE for BGT changes over time. We also run the analysis for median wage in each occupation-county cell. Both panels of Figure A.8 show that this R-squared stays flat over time.

[15] This question was piloted in the real-time surveys collected by Adams-Prassl et al. 2020.

"*How were your hours set during January and February?*

- *Fixed weekly hours;*
- *I chose my hours;*
- *My employer chose my hours, with a minimum guaranteed number of hours;*
- *My employer chose my hours, with no minimum guaranteed number of hours.*"

We make use of answers to this question, in addition to survey responses concerning other job characteristics such as occupation, remuneration type (salaried versus non-salaried), and wage to validate our vacancy-based approach. We cannot use this data for our main analysis as the question was only collected over summer 2020 and there is insufficient sample size for analysis at the local labour market level.

3.2 Classifying Work Arrangements

Our goal is to retrieve all vacancies that describe flexible work arrangements from the set of BGT job adverts. Duchini, Simion, and Turrell (2020) is the only other paper we are aware of that studies schedule flexibility from vacancy text; they classify a vacancy as flexible if a particular set of words are present in the advert. A limitation of this approach is that it is subject to both false positives and false negatives. For example, the word "flexible" can be used to describe a personality trait rather than a work arrangement ("We are looking for a motivated, flexible & committed individual"). A large number of words are used to describe flexible arrangements making it hard to specify the full dictionary a priori ("Be your own boss! Set your own schedule."). Given these difficulties, we take a supervised machine learning approach that relied on manual annotations rather than rely on a keyword search - and, as we show shortly, the machine learning approach produces a much more accurate classification of vacancies. Our method proceeds as follows:

1. Manually label a set of job vacancies for the dimensions of work arrangements of interest. This serves as a training dataset for the machine learning model and creates a "ground truth" for the correct classification of the work arrangements described in job vacancies;
2. Define the vocabulary and represent each job vacancy in a matrix format;

3. Train a machine learning model to classify work arrangements on the basis of vacancy text;

4. Apply the machine learning model to all 46 million job vacancies.

Defining Schedule Flexibility To manually label the training set of job vacancies, we require a working definition of schedule flexibility. We define a schedule flexible work arrangement as one in which the timing of work is not fixed in the contract and is left open to future (possibly repeated) negotiations between the employer and the employee. Given the impossibility of fully enumerating all the possible ways in which schedule flexibility might be described, a subset of 1,700 vacancies in the training set were annotated by two readers independently to ensure consistency of our definition.[16] In practice, we categorize a job to be schedule flexible if it mentions shift or rota work without a fixed pattern, specifies that it offers flexible working, or specifies that work will be organised according to the needs of the business.

Control over schedule When piloting our approach to develop instructions for annotators, it was not possible for us consistently to identify whether workers or employers had control over the variation in schedules from the vacancy text. For example, does a work arrangement described as a "Casual contract! Allows for flexibility" describe a worker or employer determined schedule? Even promises of "work as much or as little as you need" may be subject to managerial approval, and demands of "working according to the needs of the business" will eventually need to be agreed to by the worker.

Responses to the Understanding Society survey also provide grounds for doubting a clean distinction between employer versus employee control in vacancy text exists. Appendix Figure A.1 (a) gives the prevalence of different forms of schedule flexibility in the survey data. 24% of employed respondents reported having variable hours. Of those, 47% report having control over their schedule, and a further 36% report that their employer has control over their schedule but *with* guaranteed hours. Only around one in six of those employed on flexible contracts report having employer control over the schedule *without* minimum guaranteed hours. Thus, in practice, which party has control over the schedule is likely to lie on a continuum, rather than being represented by a simple binary.

[16] At times, the vacancy text can contain ambiguous or contradictory statements about whether the vacancy is of a particular work arrangement. Having two researchers label each vacancy independently and then reconcile any disagreements in a mutual discussion, potentially involving the senior author on this paper, also provided a mechanism for addressing this ambiguity.

Remuneration type Concern about employer-controlled flexible work arrangements rarely relates solely to the variation in the timing and number of hours that they generate. Rather, misgivings about these contracts typically concern the earnings variation that they can expose workers to. For schedule flexibility to translate into earnings variation, a worker must be employed on an unsalaried contract, i.e. rather than being paid a fixed amount each month, a worker is paid according to a wage or daily contract. For example, an academic economist is usually paid a constant, fixed amount each period even when the number of hours spent teaching and researching varies by month. The earnings of a delivery driver on an hourly wage contract, however, will vary according to the realised hours of work.

The Understanding Society survey data demonstrates that there is a strong correlation between remuneration type and who controls the schedule. Appendix Figure A.1 (b) shows that more than 65% of workers who can choose their own hours receive a salary, compared to 40% of those whose hours are chosen for them by their employer with some minimum guaranteed hours. Conversely, as one might expect, only a minority of workers without guaranteed hours are paid a salary. While we wouldn't expect there to be a perfect mapping between remuneration type and control over one's schedule, the data shows that these two concepts are strongly related.

Motivated by this relationship, we also extract information on remuneration type from the job vacancy, which is easily identifiable from the 63% of vacancies that include earnings information. This information is usually given either at an annual level (e.g. "Salary of £24,300 per annum") in which case the position is salaried, or at the wage level (e.g. "Competitive wage of £10 per hour"), in which case it is unsalaried. We also annotate whether a vacancy describes a permanent arrangement to provide another example of a positive job amenity.

We continued to annotate vacancies as schedule flexible, salaried, and permanent until we achieved sufficient accuracy in identifying these features from the vacancy text using our machine learning model. For schedule flexibility and permanent job features, our training set comprises 6,650 vacancies; for salary type there are 1,192 vacancies, as this feature is relatively easy to identify.

3.3 Supervised Machine Learning Methodology

To apply the machine learning classifier to the training data, we represent the vacancy text in matrix format. To do so, we first define our 'vocabulary', the set of relevant language components.

The terms of the vocabulary are determined by the tokenization of the text of our manually annotated vacancies. Tokenization is a way of separating a piece of text into smaller units. We use the tokenisation of 1-grams (single words) to build a parsimonious and interpretable machine learning model. We add to this engineered 2-grams and 3-grams on the basis of key sentences highlighted when annotating the training set; that is, we hand-pick phrases to build a vocabulary which can better differentiate among employment arrangements. We further tune the vocabulary through the use of a stop word list which filters out common words (e.g. and, the, it) from the vocabulary before further processing. Stop words have limited lexical content and their presence adds little signal to help us distinguish between the dimensions of employment arrangements.

With our vocabulary in hand, we represent vacancies with the binary term-document incidence matrix. Each row of the incidence matrix is a term of our vocabulary and each column is a vacancy in the training set. Equation 3.1 gives an example of the structure of this matrix. Our incidence matrix is binary as opposed to continuous, as we do not count the number a words that appear in the text, but instead only assign a one if a word appears and a zero if it does not.[17] As the incidence matrix does not take into account the order of the words as they appear in the sentence, it resembles a bag-of-words approach. Finally, we limit the size of our vocabulary, and therefore the number of rows of the incidence matrix, to 5000, to take into account the limited number of annotations.

$$X = \begin{array}{c l} \begin{bmatrix} 1 & 0 & \dots & 1 \\ 1 & 1 & \dots & 0 \\ 0 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \vdots \\ 0 & 0 & \dots & 1 \end{bmatrix} & \begin{array}{l} \text{Flexible} \\ \text{Maternity} \\ \text{Fixed Term} \\ \vdots \\ \text{Permanent} \end{array} \end{array} \tag{3.1}$$

Logistic Prediction Model with Lasso Regularization. The incidence matrix corresponds to our matrix of regressors for predicting the work arrangements, where each token is a regressor. As the number of regressors is large, and can be larger than the number of observations, we use a logistic regression model with lasso regularisation. This algorithm chooses a subset of the original 5000-token vocabulary to predict work arrangements from the vacancy text to avoid over-fiiting.

[17]We experimented with the term frequency inverse document frequency representation of text, however, here was no noticeable improvement in accuracy over the binary term-document incidence matrix.

Let $\mathbb{C}$ be the set of C classes that we want to predict (in our case the different employment arrangements) and $\mathbb{D}$ represent the set of D vacancies in the training set. Let $c_i = 1$ if vacancy i has been labelled with work arrangement $c \in \mathbb{C}$ and $c_i = 0$ otherwise. x_i represents the vector representation of document $i \in \mathbb{D}$ using our vocabulary. With the logistic specification, the probability of a vacancy i being labelled with class c given the representation of its text, x_i, is:

$$\Pr_c(c_i = 1|x_i) = \frac{\exp(x_i'\beta_c)}{1+\exp(x_i'\beta_c)}. \tag{3.2}$$

Under lasso regularization, the log-likelihood of the logistic regression model is:

$$\sum_{i=1}^{D} \log \Pr(c_i|x_i) - \frac{\lambda}{2}||\beta_c||_1 \tag{3.3}$$

where λ is the tuning parameter which governs the magnitude of the penalty term. Intuitively, with sufficiently high λ, the lasso regularization penalizes the inclusion of non-zero coefficients on vocabulary elements.

Algorithm 1 in Appendix D lists the steps taken to train the logistic regression model. In summary, we determine the tuning parameter by a grid search over a 5-fold cross-validated balance F-measure. Consider a work arrangement c. For a given value of the tuning parameter λ, we randomly sample a balanced set of vacancies[18] from the training set, $\mathbb{D}_S$. We split $\mathbb{D}_S$ into five equally sized groups $\mathcal{S}_j$ for $j = 1, ..., 5$. For each group $\mathcal{S}_\ell$, we select the parameter, $\hat{\beta}_{c,\lambda,\ell}$ that maximises the objective function at Equation 3.3 on the subset of vacancies excluding group $\mathcal{S}_\ell$, i.e. $\mathbb{D}_S \setminus \mathcal{S}_\ell$. $\hat{\beta}_{c,\lambda,\ell}$ is then used to predict the presence of work vacancy c in the excluded group $\mathcal{S}_\ell$. The model's predictions are compared to the true vacancy labels for vacancies in $\mathcal{S}_\ell$ and the accuracy of the model evaluated (see below). We repeat this procedure 500 times for each value of λ for each work arrangement. We select the value of λ that maximises the average of the cross-validated balanced F-measure over the 500 sampling draws.

We use the balanced F-measure to evaluate the accuracy of our classification algorithm. This equally weights precision and recall using the harmonic mean. The precision score penalises false positives (irrelevant items retrieved) and the recall score penalises false negatives (relevant items

[18] in ap an equal number of zeros and ones for that class.

that are not retrieved).

$$\text{Balanced F-measure} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}. \tag{3.4}$$

where

$$\text{Precision} = \frac{\sum_i (\hat{c}_i = 1 \ \&\ c_i = 1)}{\sum_i (\hat{c}_i = 1 \ \&\ c_i = 1) + \sum_i (\hat{c}_i = 1 \ \&\ c_i = 0)} \tag{3.5}$$

$$\text{Recall} = \frac{\sum_i (\hat{c}_i = 1 \ \&\ c_i = 1)}{\sum_i (\hat{c}_i = 1 \ \&\ c_i = 1) + \sum_i (\hat{c}_i = 0 \ \&\ c_i = 1)} \tag{3.6}$$

with $\hat{c}_i$ giving the predicted label of vacancy i and c_i giving the true label.

Table 1 shows the precision, recall and balanced F-measure for our trained model. We have a consistently strong performance across all dimensions of work arrangements. This is both in absolute terms and also the five most positive keywords identified by the logistic regression model. This is the most charitable keyword approach we could have taken as it is not clear that these are the words we would have selected ex-ante.

Table 1: Prediction Accuracy of Trained Logistic Regression Model

	Logistic Regression Model			Δ Keyword
Contract type	**Precision**	**Recall**	**F**	**F**
Schedule flexible	0.8540	0.8083	0.8303	0.1855
Permanent	0.9294	0.9736	0.9510	0.0223
Full-time	0.9162	0.8881	0.9019	0.1314
Salaried	0.8604	0.8415	0.8503	0.2070

Notes: Measures as defined in Equations 3.4 to 3.6 applied to our training and hold-out data sets. Improvement to keywords, Δ Keyword. is defined as the absolute difference in the balanced F-score when our ML approach is used over a keyword search.

Algorithm 2 in Appendix D lists the steps to apply our trained machine learning model to predict the employment arrangement given a previously unseen vacancy text. The trained machine learning model is used to evaluate the probability that a work arrangement is described in a given vacancy, $\Pr_c(1|x_i)$. If this probability is greater than .5, the vacancy is labelled with that particular work arrangement.

Relationship with Survey Data We finally validate our measure by comparing the distribution of flexible work arrangements that it yields to the a high quality measure of flexible work arrangements that is available in the Understanding Society survey data. We take the prevalence of flexible work arrangements in 3-digit occupation codes in 2019 in BGT and compare these to the prevalence recorded in Understanding Society for January 2020. Appendix Figure A.2 (a) shows that there is a strong positive relationship between the share of flexible work arrangements in a three-digit occupation code in the BGT vacancy data and in our representative survey data. Regressing the BGT flexible vacancy share on the USoc vacancy share yields a significant positive point estimate of 0.55 (p-value$<$0.0001). We also observe a strong, signficant relationship between the prevalence of salaried positions in the BGT vacancy data and survey data (Appendix Figure A.2 (b)).

4 Schedule Flexibility: Descriptive Facts

Before analyzing the relationship between schedule flexibility and employer concentration, we provide new stylized facts on the prevalence and characteristics of flexible work arrangements. In what follows, we restrict the sample to the set of vacancies that are flagged by BGT as having wage information. This allows us to keep the sample stable as we analyse the interaction of schedule flexibility with remuneration type. In Table A.3, we show that there is no statistically significant difference in the non-wage characteristics between vacancies that do and do not post a wage. For simplicity, we will refer to vacancies that are schedule flexible and unsalaried as "risky flexible" jobs and those that are schedule flexible and salaried as "safe flexible" jobs.

Figure 2 gives the overall share of flexible vacancies across occupations, the wage distribution and by the number of skills a vacancy requires. Approximately a third (29%) of all vacancies in the UK are flexible, with an equal split between safe and risky flexibility. This closely corresponds to the prevalence in representative survey data reported in Section 3.2. Figure 2 (a) demonstrates that the prevalence and type of flexibility varies significantly across occupations. The share of flexible vacancies varies from 12% amongst managers to 44% in caring and leisure services. Flexibility in lower-skill occupations is much more likely to be risky rather than safe: in elementary occupations (which include cleaning, low-skilled farm work, and security guards), only about 20% of flexible jobs come with a salary.

Figures 2 (b) and (c) make use of measures that BGT themselves have extracted from the vacancy text. Figures 2 (b) shows the prevalence of flexible vacancies by BGT's measure of the number of skills required by an advert. We see that job adverts which mention a larger number of skills, are less likely to be flexible. However, the prevalence of risky and safe flexibility varies significantly across the skill distribution. Risky flexibility is much more common amongst adverts that specify only a small number of skills. The opposite is true for safe flexibility. BGT also extracts and standardizes information on hourly wages from job adverts that include remuneration information. Risky flexibility is much more prevalent amongst low wage jobs (Figure 2 (c)). Approximately 40% of vacancies paid less than £8 per hour are flexible, and the majority of these positions are risky flexible jobs. For jobs paying more than £20 per hour, only 10% of the positions are described as flexible. This gradient is driven by changes in the prevalence of risky flexibility across the wage distribution.

Table 2 gives the results of descriptive regression models where we regress vacancy characteristics on dummy variables capturing whether a job is flexible, unsalaried, and their interaction. We include three-digit occupation, time, and county fixed effects in all specifications. This allows us to probe whether the patterns in Figure 2 reflect differences in the occupational distribution of flexible and non-flexible jobs or the distribution of salaried and non-salaried jobs. This exercise shows that, even among jobs of the same occupation, in the same area, at the same time, those that are flexible are associated with a significant wage penalty, and risky flexible jobs (flexible non-salaried) a *further* penalty. Table 2 also shows (again, conditional on occupation, location and time) strong relationships between skills and flexibility. Compared to salaried jobs with a fixed schedule, safe flexible jobs are associated with a broader set of skills, but risky flexible a smaller set.

Columns (4) and (5) analyze the interaction between flexibility and whether the job is permanent. Flexible salaried jobs are more likely to be temporary than non-flexible salaried ones, while risky flexible jobs are *less* likely to be temporary. This raises the question of what it means for a risky flexible job to be permanent if an employer is able to simply offer a non-salaried worker zero hours in perpetuity without formally dismissing them. Indeed, the policy and legal commentary in this area has questioned whether the terms "permanent" and "temporary" are really meaningful for low-wage flexible jobs (Adams, Prassl, et al. 2018).

Figure 2: Prevalence of Flexible Vacancies

(a) Occupation

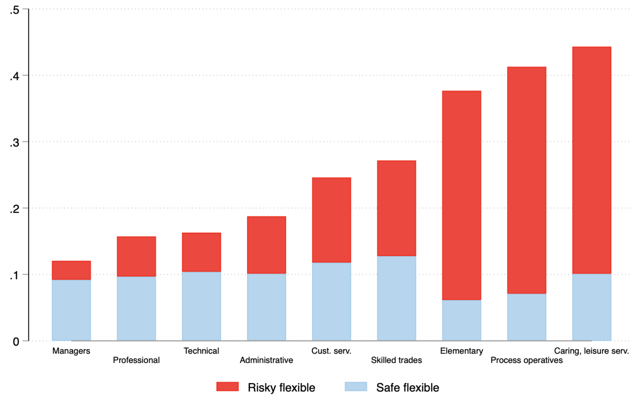

(b) Number of Skills

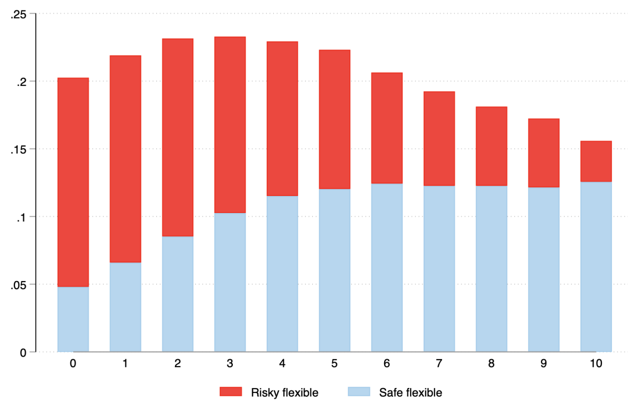

(c) Wage

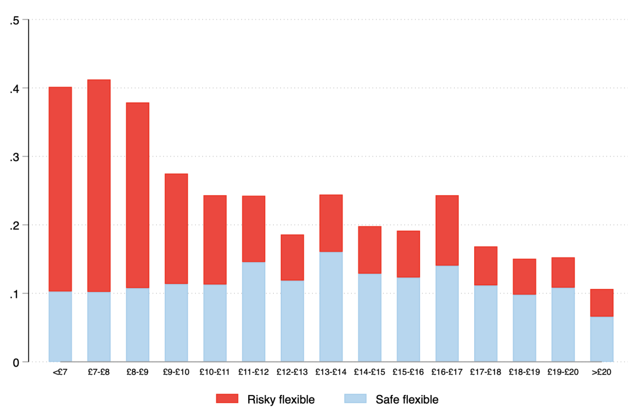

Notes: BGT data from 2014-2019. The sample is restricted to job adverts with an occupaiton code, wage, and county. Wages deflated to the 2019 price level.

Table 2: Flexibility and vacancy characteristics

	Log Wage	No. of Skills		Permanent	
	(1)	(2)	(3)	(4)	(5)
Flexible	-0.072***	0.551***	0.572***	-0.027***	-0.030***
	(0.0013)	(0.0155)	(0.0154)	(0.0012)	(0.0011)
Non-Salaried	-0.022***	-1.200***	-1.213***	-0.453***	-0.454***
	(0.0043)	(0.0165)	(0.0165)	(0.0033)	(0.0032)
Flexible × Non-Salaried	-0.018***	-0.174***	-0.169***	0.089***	0.089***
	(0.0040)	(0.0178)	(0.0177)	(0.0024)	(0.0024)
Log Wage			0.401***		-0.031***
			(0.0193)		(0.0020)
Constant	2.593***	5.005***	3.964***	0.138***	0.802***
	(0.0014)	(0.00987)	(0.0482)	(0.0011)	(0.0055)
Observations	16,134,476	16,134,476	16,134,476	16,134,476	16,134,476
R^2	0.3590	0.2086	0.2093	0.2732	0.2752
Occupation FE	✓	✓	✓	✓	✓
County FE	✓	✓	✓	✓	✓
Time FE	✓	✓	✓	✓	✓

Notes: Standard errors clustered at county-occupation level in parentheses. $*p < 0.1, **p < 0.05, ***p < 0.01$ Sample restricted to vacancies with information on wage, county, and occupation. Wages are deflated and in 2019 prices. All specifications control for county, year-quarter, and occupation fixed effects.

Variation Over Time A key advantage of our data is that we have access to a consistent time-series on schedule flexible jobs.[19] Panel (a) of Appendix Figure A.3 plots the the share of flexible vacancies between 2014 and 2019. The proportion of flexible vacancies increased significantly over the period from 16% to 29%. Safe and risky flexible vacancies contributed equally to the observed growth.

This rise is partly driven by an increase in the proportion of firms that offer at least one flexible vacancy. Panel (b) of Appendix Figure A.3 gives the proportion of firms offering at least one vacancy with a given flexibility feature in 2014 and 2019. The proportion of firms offering at least one flexible vacancy rose from 46% in 2014 to 57% in 2019. Both the prevalence of "risky" and "safe" flexible employers grew over the period.

5 Schedule Flexibility and Employer Concentration

We now apply the predictions from our theoretical model to identify the amenity value of schedule flexibility offered in the labor market. Interpreted through the lens of our model, the relationship between employer concentration and the prevalence of flexible contracts allows us to identify whether these job features are costly or profitable, amenities or disamenities.

5.1 Defining Employer Concentration

We use a vacancy-based Herfindahl-Hirschman Index (HHI) to measure employer concentration (Azar, Marinescu, Steinbaum, and Taska 2020; Benmelech, Bergman, and Kim 2020; Rinz 2022). Specifically, we define a local labour market as the intersection of a county (of which there are 206) and occupation (we use the SOC 3 digit classification, comprised of 89 occupations). Employer concentration is then measured as:

$$HHI_{c,o,t} = \sum_{i}^{I} \left(\frac{v_{i,o,c,t}}{\sum_j v_{j,o,c,t}} \right)^2 \tag{5.1}$$

$$= \sum_{i}^{I} s_{i,c,o,t}^2 \tag{5.2}$$

[19] In the next section, we harness variation within occupations and counties over time to analyze the relationship between employer concentration and flexible work arrangements.

where $HHI_{c,o,t}$ is the Herfindahl-Hirschman Index in county c and occupation o at time t, $v_{i,o,c,t}$ gives the number of vacancies that firm i posts in that county-occupation-time cell and $s_{i,o,c,t}$ is firm i's percentage share of total vacancies in that cell. Time is measured at a quarterly frequency.[20] We follow this approach for two reasons. First, we can construct HHI at a more granular level using our vacancy data than if we used standard data on employment.[21] Second, this measure directly corresponds to the structural construct in our theoretical model, where the size and number of employers who are hiring determine workers' outside option (Equation 2.4).

However, the index is not without problems. Not all jobs are advertised online and workers may search beyond the 3-digit SOC-county boundary. In principle, our measure could overestimate concentration if there are many employers who do not post vacancies online or underestimate concentration if there are significant differences in the hiring to job posting rate amongst small and large employers (Azar, Marinescu, Steinbaum, and Taska 2020; Schubert, Stansbury, and Taska 2022). As discussed in Section 3, BGT covers 83% of UK vacancies and there has been no change in the representativeness of the source across occupation-county cells over the period considered. As a direct robustness check, we compare our BGT-based measure with HHI estimated from the Annual Survey of Hours and Earnings, an employer survey of a random 1 per cent of employee jobs. Due to sample size limitations, we can only calculate HHI at the county-SOC-1 level for each year of our BGT sample (2014-2019).[22] Appendix Figure A.5 shows a strong and statistically significant relationship between the two measures across; the ASHE HHI measure can explain 72% of the variation in the BGT HHI. Therefore, we consider the potential bias from using job vacancy data is likely to be small. Schubert, Stansbury, and Taska (2022) test the importance of workers searching beyond the boundaries of narrow local labour market definitions by constructing an innovative measure of outside options using occupational transitions recorded in US resumes. In their main IV specification, controlling for outside options only modestly reduces the elasticity of HHI with respect to wages, suggesting this is not likely to be a first order issue.

The most fundamental criticism of employer concentration as a measure of firm monopsony

[20] For this analysis, we use vacancies with information on firm name, county, and occupations. We drop vacancies posted by the National Health Service, and treat Greater London as a single location since not all London vacancies come with an exact borough location.

[21] There is insufficient sample size to calculate HHI within detailed occupation codes and geographic areas in the Annual Survey of Hours and Earnings (ASHE), a 1% sample of employee jobs in the UK.

[22] ASHE is not a perfect source of HHI data. While representative, it is only a 1% sample, which means that as a measure of concentration it will tend to be both noisy and biased upwards (a point discussed in further detail in Abel, Tenreyro, and Thwaites 2018).

power is that it is itself a result of this power rather than a market primitive in its own right.[23] Existing evidence has demonstrated a positive relationship between employer concentration and other measures of monospony power. Yeh, Macaluso, and Hershbein (2022) estimate plant-level wage markdowns for US manufacturing firms. Aggregating these measures up to market level, they find a positive, if weak, correlation with employer concentration. However, Azar, Marinescu, and Steinbaum (2019) find that employer concentration is a good predictor of inelastic labour supply using job application data from an online job platform.

5.2 Empirical approach

We seek the causal relationship between employer concentration and the prevalence of job features in a local labour market. The key concern in our setting is omitted variable bias. For example, if more productive occupations have lower market entry costs and can also provide flexibility more cheaply, this will generate a positive relationship between HHI and flexibility that does not derive from any causal impact of monopsony power on flexibility provision.

To address this concern, our baseline specification includes for a rich set of fixed effects to control for any systematic variation across occupations, regions, and individual labour markets that might be driving the concentration-job feature relationship. Specifically, we control for local labour market (occupation x county) fixed effects (α_{oc}) and occupation-specific and county-specific time trends (δ_{ot} and γ_{ct} respectively). Our baseline estimating equation is:

$$f_{cot} = \beta \log HHI_{oct} + \alpha_{oc} + \delta_{ot} + \gamma_{ct} + \epsilon_{oct} \tag{5.3}$$

where f_{cot} is the share of vacancies in county c and occupation o in quarter t that have job feature f and β is our parameter of interest. Standard errors are clustered at county-occupation level throughout.

Our fixed effect strategy will not be sufficient if there are within-market shocks that simultaneously affect employer concentration and the provision of a job feature. For example, an idiosyncratic shock to productivity in a particular occupation and county could, in theory, affect both concentration and the prevalence of the feature. To address this concern and assess the robustness of our FE specification, we instrument our HHI measure using firm-level shocks to nationwide

[23] In the case of product markets, Syverson (2019) find that the relationship between market concentration and monopoly power depends, for example, on the type of market competition and product substitutability.

hiring. Intuitively, the instrument is motivated by the idea is that a firm's local hiring is driven to some extent by nationwide shocks to firm demand, but local shocks are unlikely to determine its nationwide hiring.

To construct the instrument, we use information on past hiring and nationwide shocks to predict the number of vacancies posted by each establishment x occupation today. Let $-c$ indicate all counties other than county c. For multi-establishment firms - those who, last period, posted vacancies for the occupation in question in multiple counties[24] - we take their local hiring in the previous period for the occupation, $v_{i,o,c,t-1}$. We then scale it by the growth in the leave-one-out sum of vacancies posted by the firm for that occupation, $\frac{v_{i,o,-c,t}}{v_{i,o,-c,t-1}}$. This expression gives us a plausibly exogenous prediction of firm's local hiring. For single-establishment firms, we use the firm's past period's hiring to predict hiring today.

$$\hat{v}_{i,o,c,t} = \begin{cases} v_{i,o,c,t-1} \times \frac{v_{i,o,-c,t}}{v_{i,o,-c,t-1}} & \text{for multi-establishment firms} \\ v_{i,o,c,t-1} & \text{for single-establishment firms} \end{cases} \tag{5.4}$$

In the second step, we use our predicted number of vacancies at firm x occupation level to calculate a new predicted index of employer concentration. We use this predicted HHI measure as our instrument:

$$IV_{c,o,t} = \log\left(\sum_i \left(\frac{\hat{v}_{i,o,c,t}}{\sum_j \hat{v}_{j,o,c,t}}\right)^2\right) \tag{5.5}$$

Our identification approach is based on a large number of quasi-exogenous shocks to local firm hiring. In this respect, the IV resembles a shift-share instrument, with the proportional change in nationwide hiring as the shifts (shocks) and the past local hiring as the shares. As we rely on shocks to the nationwide hiring of thousands of firms, our identification strategy is in the vein of Borusyak, Hull, and Jaravel (2021), who show that shift-share instruments can be valid even if the shares used are endogenous.[25]

[24]Thus, a multi-establishment firm, in a given occupation, is one for whom $v_{i,o,-c,t-1} > 0$.
[25]However, we do not actually use a shift-share IV – instead, we use this structure to help us predict employer concentration at the market level.

5.3 Results

Appendix Table A.1 gives descriptive statistics on firm hiring and employer concentration. On average, 36.7 vacancies are posted in each local labour market (county x occupation) per quarter. The average HHI is 3331, which is equivalent to 3 equal-size employers per local labour market. The US Department of Justice defines a highly concentrated market as one with HHI above 2,500. In the UK, this definition applies to 41% of market-time units, somewhat below the 60% in Azar, Marinescu, Steinbaum, and Taska (2020). There is a single employer in 16% of market-time observations. Appendix Figure A.4 shows that while the local market, county-time, and occupation-time fixed effects account for a large share of the variation in log(HHI), significant variation in the index remains even after including the full set of fixed effects.

Figure 3 summarises our baseline fixed effect estimates of the impact of log HHI on the prevalence of different job features. It is clear that changes in employer concentration do not affect the provision of all job features in the same way. A rise in concentration leads to an increase in the provision of schedule flexible jobs. However, this masks very different variation across safe and risky flexible jobs. As workers' outside options deteriorate, risky flexibility rises, while the provision of safe flexibility does not significantly change. As a robustness check, we also estimate the relationship between permanent vacancies and HHI. We find that permanent vacancies are less likely to be offered in more concentrated labor markets, in line with the strong existing evidence that they are a costly amenity (Datta 2019; Mahmud et al. 2021; Bassanini et al. 2022).

Table 3 further probes the robustness of our main results. Column (1) gives our preferred fixed effects specification but adding controls for wages of the vacancies. Existing research has shown that wages fall in more concentrated labor markers (e.g. Azar, Marinescu, Steinbaum, and Taska (2020)). As job features and wages jointly enter both firms' profits and workers' utility, it is likely that the prevailing market wage will influence the level of job features provided. This suggests that we should include wages on the right hand side of our regressions. However, low wages are also the outcome of strong market power of firms, so wages are likely a "bad control". We therefore include these results as a robustness check.

Columns (2) - (6) in Table 3 give the results of our IV specifications. In all five of our IV specifications, we instrument HHI with HHI based on predicted firm-level hiring (5.5) in addition to including the full set of fixed effects as in our baseline fixed effect specification. Column (2) presents the baseline IV regression, and in column (3) we also control for average market wages

Figure 3: The relationship between employer concentration and flexibility

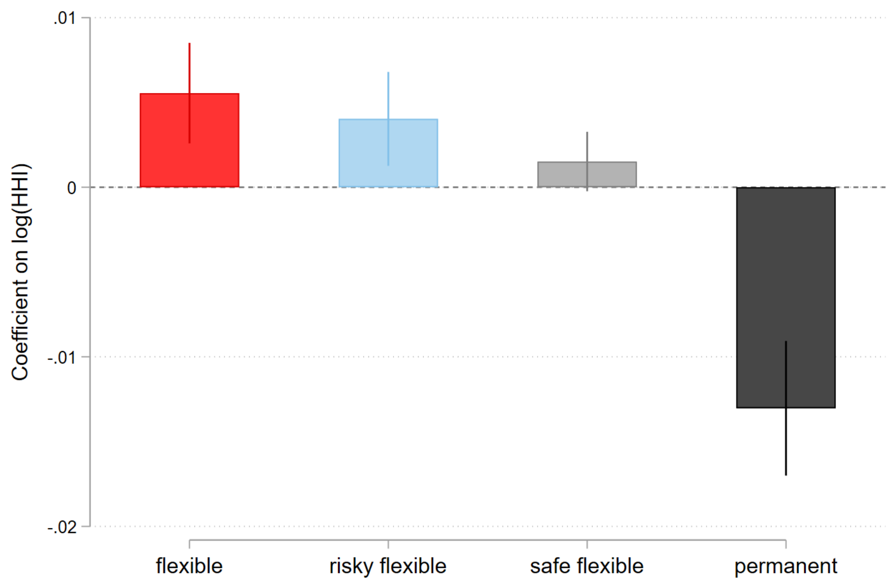

Notes: Each bar plots the size of the coefficient on employer concentration from our baseline regression of job feature provision on log(HHI). The regression includes fixed effects at county x SOC-3 level, county x year-quarter level, and SOC-3 x year-quarter level. The 95% confidence intervals are calculated from standard errors clustered at county-occupation level.

on the righthand side. In column (4), we weigh the regression using market size from the previous period instead of the contemporaneous size we use in the baseline. This robustness check is to address the fact that market-level HHI and market size (i.e. the number of vacancies posted in a given market) are often positively correlated; using past market size thus adds another layer of exogeneity to the estimation procedure. Column (5) re-estimates our baseline IV regression at the SOC-2, rather than the finer SOC-3 level used in the baseline. While the finer occupation categories are probably a better representation of local labour markets in practice, using SOC-3 hiring at firm level might result in many zeros, thus excluding smaller firms from our analysis. Using the coarser SOC-2 levels allows us to include these smaller firms while maintaining much of the high-level of detail in our data. Finally, in Column (6) we keep the definition of local labour market as SOC-3 x county, but we use firm-level hiring, rather than hiring at firm-SOC-3 level, to construct the IV. The rationale for doing this is similar to using SOC-2 occupation categories: since many firms do not hire in each occupation every time period, specifying the nationwide hiring shocks at firm level allows us to keep more firms in the data set, and reduces the noise in our IV.

Our results remain robust throughout. While the IV regressions suffer from larger standard errors, our baseline qualitative results hold across all specifications: risky flexibility is more prevalent in more concentrated labor markets, while the prevalence of safe flexible vacancies is largely unaffected by changes in employer concentration.

Discussion Our results show that the interaction of salary type and flexible scheduling pick out very different types of jobs. In addition to the wage and skill content of risky and safe flexible positions being systematically different (in ways that go beyond the differences in salaried versus non-salaried positions - see Table 2), they are also differentially affected by changes in outside options. Interpreted through the lens of our model, our estimates suggest that safe flexibility is a typical costly amenity, while risky flexibility is a disamenity that increases firms' profits. The zero coefficient on safe flexibility indicates that it is an amenity but our results are not conclusive as to whether it is profitable or costly. Theorem 3 has implications for interpretations of the difference in the magnitude of coefficients on risky vs. safe flexible vacancies; it states that job features that generate greater net improvements in utility and profits respond relatively less to changes in workers' outside option. This result implies that safe flexibility is overall more productive than

Table 3: The relationship between employer concentration and flexibility: alternative specifications

	FE	IV				
	(1)	(2) baseline	(3) wage controls	(4) past weights	(5) SOC2 level	(6) firm level
Panel A: Flexible Vacancies						
log(HHI)	0.00599***	0.0276***	0.0262**	0.0237***	0.0241**	0.0132
	(0.00153)	(0.0101)	(0.0114)	(0.00912)	(0.0112)	(0.00826)
log(wage)	-0.150***					
	(0.00597)					
Panel B: Risky Flexible Vacancies						
log(HHI)	0.00449***	0.0186*	0.0173	0.0160*	0.0259**	0.0170**
	(0.00144)	(0.00954)	(0.0107)	(0.00868)	(0.0112)	(0.00797)
log(wage)	-0.149***					
	(0.00552)					
Panel C: Safe Flexible Vacancies						
log(HHI)	0.00150*	0.00899	0.00890	0.00776	-0.00172	-0.00382
	(0.000892)	(0.00827)	(0.00864)	(0.00746)	(0.00686)	(0.00615)
log(wage)	-0.000428					
	(0.00421)					
N	253056	245170	245170	245170	92404	247600
1st-stage F		61.429	61.892	71.260	57.392	115.931

Notes: Estimating the relationship between employer concentration (log(HHI)) and the provision of different job features at the market level. Column (1) replicates the baseline FE specification (county x SOC3, year-quarter x SOC3, and year-quarter x county fixed effects) with added controls for average market wage and the share of vacancies with posted wage information. Columns (2) - (6) instrument for log(HHI) using an IV based on predicted hiring at establishment level, in addition to including the full set of fixed effects. Column (2) estimates the baseline IV specification. In Column (3), we also control for wage variables. in Column (4), we weigh the regressions using previous year's market size rather than the current size as in the baseline in column (2). In Column (5), we replicate the baseline IV at SOC2, rather than SOC3 level. In Column (6), we use firm-level hiring, rather than hiring at firm x SOC-3 level, to construct the IV. Standard errors clustered at county-occupation level in parentheses. $*p < 0.1, **p < 0.05, ***p < 0.01$ Sample restricted to vacancies with information on wage, county, and occupation.

risky flexibility, because the magnitude of its estimated coefficient is smaller.

Our results imply that the value of schedule flexibility varies across different workers according to the other dimensions of job design that it is bundled with. Risky flexibility combines schedule flexibility with no commitment to a level of timing of remuneration. This job feature is profitable for firms but a disamenity for workers. This feature is also much more prevalent in positions at low-wages and in less skilled occupations. Safe flexibility is the dominant form of flexibility at high wages and in high skilled occupations. Our results, therefore, suggest that the distribution of welfare-relevant job features accentuates overall inequality in the labour market.

As we discuss in Section 3.2, our concept of risky flexibility is conceptually distinct from that of employer-controlled schedule flexibility.[26] However, the two concepts are clearly related and interact in determining worker utility: a worker's exposure to earnings risk will be greater in jobs where they do not receive a fixed salary *and* do not control their schedule. To analyze the role that employer control plays in driving the relationship between risky flexible jobs and employer concentration, we combine the Understanding Society survey data on control of flexible schedule jobs (see Section 3.1) with our vacancy data at the 1-digit SOC level. We then interact HHI with the proportion of employer-controlled schedule flexible jobs at the 1-digit occupation level in our baseline IV specification:

$$f_{cot} = \beta \log HHI_{oct} + \rho \log HHI_{oct} \times Employer_O + \alpha_{oc} + \delta_{ot} + \gamma_{ct} + \epsilon_{cot} \tag{5.6}$$

where $Employer_O$ is the proportion of employer controlled schedule flexible contracts (both with an without a minimum number of guaranteed hours) out of all flexible contracts in the 1-digit SOC code encompassing occupation o.

Figure 4 shows the marginal effect of a change in employer concentration on risky flexibility by occupation group given the prevalence of employer-controlled flexibility in that occupation.[27] The positive relationship between employer concentration and risky flexibility is driven by those occupations in which employer-controlled flexibility is most prevalent. This is reasonable: in those

[26]We cannot measure employer control of schedules from the vacancy text.
[27]Figure A.6 gives the same result for all flexible jobs.

Figure 4: Occupation-Specific results: Risky Flexibility & Employer Schedule Control

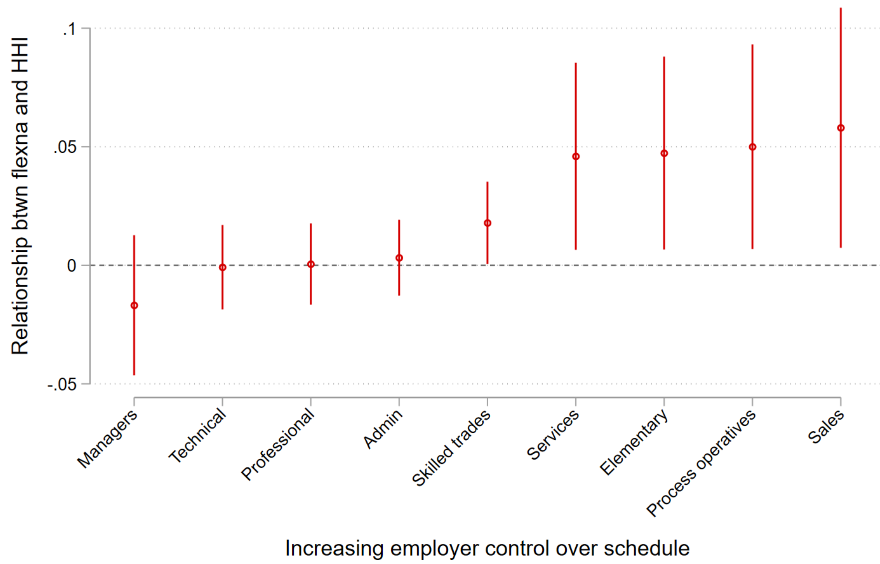

Notes: Figure plots the marginal effect of log(HHI) on risky flexibility by occupation given the observed prevalence of employer-controlled flexibility in that occupation as measured in USoc. These marginal effects are estimated using regression (5.6), with $\hat{\beta} = -0.056\,(0.033)$ and $\hat{\rho} = 0.134\,(0.067)$. The regression includes fixed effects at county x SOC-3 level, county x year-quarter level, and SOC-3 x year-quarter level. The 95% confidence intervals are calculated from standard errors clustered at county-occupation level.

occupations where employees have little control over their schedule and their earnings fluctuate with variation in the number and timing of hours, flexibility is a clear disamenity for workers. In occupations where workers can control their schedules, risky flexibility is not such a disamenity. These results suggest that our approach could be harnessed to identify whether the flexibility experienced by different groups of workers is an amenity or disamenity in data sets that include information on flexibility and salary type, even when measures of schedule control are missing.

6 Conclusion

We have developed a new model of job design under monopsony to analyze the amenity value of job flexibility on offer across the wage and occupational distribution. We applied the insights of our model to a new measure of schedule flexibility derived from a machine learning analysis of job vacancy text. Schedule flexibility is not the same job feature for all workers. We find that risky flexibility, under which schedule flexibility is offered alongside a non-salaried contract, is a profitable disamenity, while safe flexibility, that insures workers from earnings risk, is an amenity. Risky flexibility is much more prevalent at low wages, in less-skilled occupations, and in less competitive labour markets. Safe flexibility, on the other hand, is the dominant form of flexibility at high wages and in high skilled occupations.

These results suggest that the distribution of welfare-relevant job features accentuates overall inequality in the labour market. Risky flexibility, a disamenity, is most prevalent in low-wage occupations and increases in response to a rise in employer power. More broadly, our results imply that to fully capture the impact of monopsonistic power on labour market outcomes, we should also consider the impact of market power on non-monetary job features. For example, a zero employment effect of a rise in the minimum wage could partly be driven by firms increasing profitable disamenities or decreasing costly amenities that are offered to workers. Previous attempts to analyse non-wage job features have often been hampered by data and measurement issues. We show that job vacancy text can provide a rich source of information on non-wage job features. Our machine learning measure of schedule flexibility correlates well with one-off, representative survey evidence on flexibility and achieves high accuracy relative to our annotated "ground truth" without huge computational burden.

Our results suggest several further avenues for research. First, in our data we cannot observe realised outcomes once a worker is matched to a vacancy. Combining the measurement approach developed in this paper, with matched employer-employee data would facilitate a detailed analysis of career trajectories in flexible positions, something we know little about. Second, we found systematic differences in the skill content of risky flexible and safe flexible jobs; risky flexible vacancies are associated with the fewest number of skills demanded, less than both safe flexible and non-flexible vacancies. Understanding the reason for these skill differences is important for understanding the long term consequences for employment progression, worker welfare, and productivity of work in flexible jobs, especially if these positions continue to grow in popularity. Finally, the machinery developed in this paper can be applied to other important non-wage amenities, including health insurance provision and on-the-job training.

References

Abel, Will, Silvana Tenreyro, and Gregory Thwaites (Oct. 2018). *Monopsony in the UK*. Discussion Papers 1827. Centre for Macroeconomics (CFM).

Abraham, Katharine G and Ashley Amaya (2019). “Probing for informal work activity”. In: *Journal of Official Statistics* 35.3, pp. 487–508.

Adams, Abi, Jeremias Prassl, et al. (2018). “Zero-hours work in the United Kingdom”. In: *Conditions of Work and Employment Series* 101.

Adams-Prassl, Abi et al. (2020). “Inequality in the impact of the coronavirus shock: Evidence from real time surveys”. In: *Journal of Public economics* 189, p. 104245.

— (2023). “Perceived returns to job search”. In: *Labour Economics* 80, p. 102307.

Azar, José, Ioana Marinescu, and Marshall Steinbaum (2019). “Measuring Labor Market Power Two Ways”. In: *AEA Papers and Proceedings* 109, pp. 317–321.

Azar, José, Ioana Marinescu, Marshall Steinbaum, and Bledi Taska (2020). “Concentration in US labor markets: Evidence from online vacancy data”. In: *Labour Economics* 66, p. 101886.

Bassanini, Andrea et al. (2022). *Labour Market Concentration, Wages and Job Security in Europe.* Tech. rep. IZA Discussion Paper No. 15231.

Benmelech, Efraim, Nittai K. Bergman, and Hyunseob Kim (2020). “Strong Employers and Weak Employees: How Does Employer Concentration Affect Wages?” In: *Journal of Human Resources.*

Boeri, Tito et al. (2020). “Solo self-employment and alternative work arrangements: A cross-country perspective on the changing composition of jobs”. In: *Journal of Economic Perspectives* 34.1, pp. 170–195.

Bonhomme, Stéphane and Grégory Jolivet (2009). “The pervasive absence of compensating differentials”. In: *Journal of Applied Econometrics* 24.5, pp. 763–795.

Borusyak, Kirill, Peter Hull, and Xavier Jaravel (2021). “Quasi-Experimental Shift-Share Research Designs”. In: *The Review of Economic Studies* 89.1, pp. 181–213.

Card, David et al. (2018). "Firms and Labor Market Inequality: Evidence and Some Theory". In: *Journal of Labor Economics* 36.S1, S13–S70.

Carrillo-Tudela, Carlos et al. (2023). "Search and Reallocation in the COVID-19 Pandemic: Evidence from the UK". In: *Labour Economics*, p. 102328.

Clemens, Jeffrey, Lisa B Kahn, and Jonathan Meer (2020). *Dropouts Need Not Apply? The Minimum Wage and Skill Upgrading*. Tech. rep. National Bureau of Economic Research.

Crossley, Thomas F, Paul Fisher, and Hamish Low (2021). "The heterogeneous and regressive consequences of COVID-19: Evidence from high quality panel data". In: *Journal of public economics* 193, p. 104334.

Datta, Nikhil (2019). *Willing to Pay for Security: A Discrete Choice Experiment to Analyse Labour Supply Preferences*. CEP Discussion Paper No 1632. Centre for Economic Performance.

Datta, Nikhil, Giulia Giupponi, and Stephen Machin (2019). "Zero-hours contracts and labour market policy". In: *Economic Policy* 34.99, pp. 369–427.

Deming, David and Lisa B Kahn (2018). "Skill requirements across firms and labor markets: Evidence from job postings for professionals". In: *Journal of Labor Economics* 36.S1, S337–S369.

Dube, Arindrajit, Suresh Naidu, and Adam D Reich (2022). *Power and Dignity in the Low-Wage Labor Market: Theory and Evidence from Wal-Mart Workers*. Tech. rep. National Bureau of Economic Research.

Duchini, Emma, Stefania Simion, and Arthur Turrell (2020). "Pay Transparency and Cracks in the Glass Ceiling". In: *arXiv preprint arXiv:2006.16099*.

Etheridge, Ben, Yikai Wang, and Li Tang (2020). *Worker productivity during lockdown and working from home: Evidence from self-reports*. Tech. rep. ISER Working Paper Series.

Folke, Olle and Johanna Rickne (2022). "Sexual harassment and gender inequality in the labor market". In: *The Quarterly Journal of Economics* 137.4, pp. 2163–2212.

Forsythe, Eliza et al. (2020). "Labor Demand in the time of COVID-19: Evidence from vacancy postings and UI claims". In: *Journal of Public Economics*, p. 104238.

Gibbs, Michael, Friederike Mengel, and Christoph Siemroth (2021). "Work from home & productivity: Evidence from personnel & analytics data on IT professionals". In: *University of Chicago, Becker Friedman Institute for Economics Working Paper* 2021-56.

Hamermesh, Daniel S. (1999). "Changing Inequality in Markets for Workplace Amenities". In: *The Quarterly Journal of Economics* 114.4, pp. 1085–1123.

Hansen, Stephen et al. (2023). *Remote Work across Jobs, Companies, and Space*. Tech. rep. National Bureau of Economic Research.

Hershbein, Brad and Lisa B Kahn (2018). "Do recessions accelerate routine-biased technological change? Evidence from vacancy postings". In: *American Economic Review* 108.7, pp. 1737–72.

Hwang, Hae-shin, Dale T. Mortensen, and W. Robert Reed (1998). "Hedonic Wages and Labor Market Search". In: *Journal of Labor Economics* 16.4, pp. 815–847.

Jarosch, Gregor, Jan Sebastian Nimczik, and Isaac Sorkin (2019). *Granular search, market structure, and wages*. Tech. rep. National Bureau of Economic Research.

Javorcik, Beata et al. (2020). "Unravelling deep integration: Local labour market effects of the brexit vote". In: *Oxford Discussion Paper*.

Lamadon, Thibaut, Magne Mogstad, and Bradley Setzler (2022). "Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market". In: *American Economic Review* 112.1, pp. 169–212.

Mahmud, Minhaj et al. (2021). "What aspects of formality do Workers value? Evidence from a choice experiment in Bangladesh". In: *The World Bank Economic Review* 35.2, pp. 303–327.

Marinescu, Ioana (2017). "The general equilibrium impacts of unemployment insurance: Evidence from a large online job board". In: *Journal of Public Economics* 150, pp. 14–29.

Marinescu, Ioana and Ronald Wolthoff (2020). "Opening the Black Box of the Matching Function: The Power of Words". In: *Journal of Labor Economics* 38.2, pp. 535–568.

Mas, Alexandre and Amanda Pallais (2017). "Valuing alternative work arrangements". In: *American Economic Review* 107.12, pp. 3722–59.

— (2020). "Alternative work arrangements". In: *Annual Review of Economics* 12, pp. 631–658.

Prassl, Jeremias (2018). *Humans as a service: The promise and perils of work in the gig economy.* Oxford University Press.

Rinz, Kevin (2022). “Labor market concentration, earnings, and inequality”. In: *Journal of Human Resources* 57.S, S251–S283.

Schubert, Gregor, Anna Stansbury, and Bledi Taska (2022). *Employer Concentration and Outside Options.* CEPR Discussion Paper No. DP12919. Centre for Economic Policy Research.

Seidman, Laurence S. (1989). “Complements and Substitutes: The Importance of Minding p’s and q’s”. In: *Southern Economic Journal* 56.1, pp. 183–190.

Syverson, Chad (2019). “Macroeconomics and market power: Context, implications, and open questions”. In: *Journal of Economic Perspectives* 33.3, pp. 23–43.

Taylor, Matthew et al. (2017). *Good work: The Taylor review of modern working practices.* Department for Business, Energy & Industrial Strategy London.

Yeh, Chen, Claudia Macaluso, and Brad Hershbein (2022). “Monopsony in the US labor market”. In: *American Economic Review* 112.7, pp. 2099–2138.

A Additional figures and tables

Figure A.1: Schedule Control & Remuneration Type: Understanding Society

(a) Schedule control among workers without fixed hours

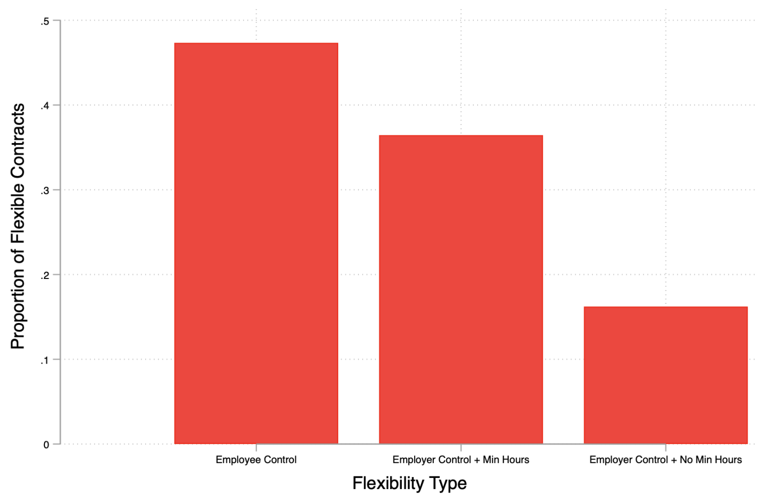

(b) Remuneration type conditional on schedule control

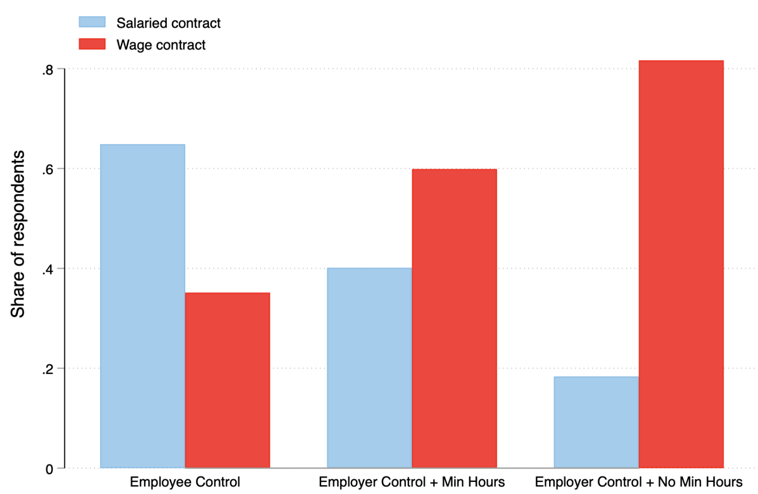

Notes: Understanding Society Covid-19 panel data. Sample restricted to those in employment in January 2020 in a job without a fixed number of hours and without missing values on answers to job flexibility and remuneration type of job in January 2020.

Figure A.2: Correlation Between Measures in Understanding Society & Burning Glass

(a) Flexibility

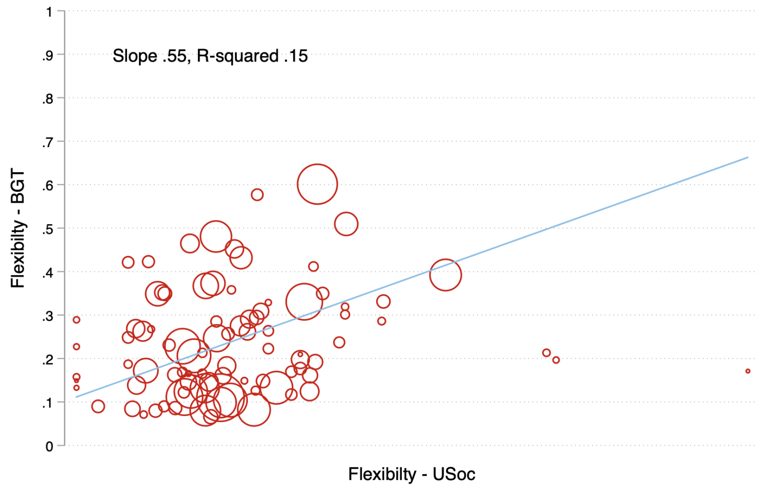

(b) Salaried

Notes: Scatter plot of the average proportion of (a) flexible contracts and (b) salaried contracts across 3-digit occupations codes in Understanding Society Covid-19 panel data (x-axis) and in BGT (y-axis). Sample in USoc restricted to those in employment in January 2020 and without missing values on answers to job flexibility and remuneration type of job in January 2020. Size of the scatter point is proportional to the relative size of a 3-digit SOC code in the USoc data.

Figure A.3: Flexible Vacancies Over Time

(a) Share of flexible vacancies

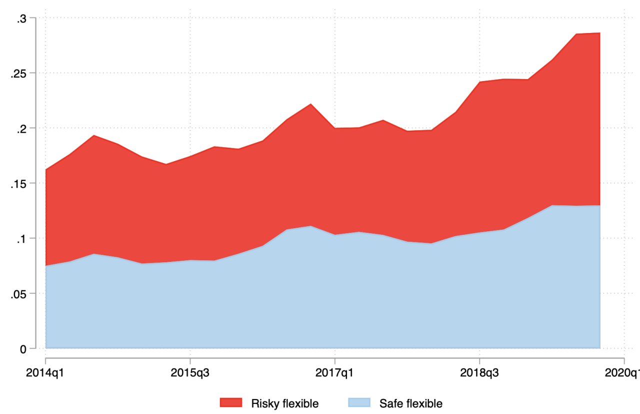

(b) Proportion of firms offering flexible vacancies

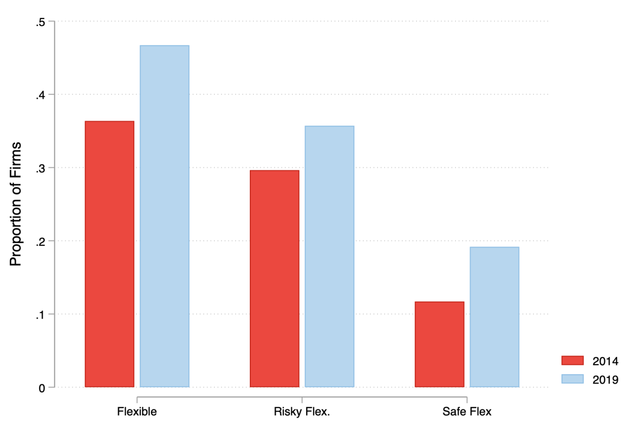

Notes: Panel (a) plots the average prevalence of flexibility in each quarter from 2014Q1 to 2019Q4. Panel (b) plots the share of flexible employers, defined as firms that post at least one flexible vacancy in a given year, for 2014 and 2019. The sample is restricted to vacancies with wage, occupation & county information.

Figure A.4: Variation in log(HHI)

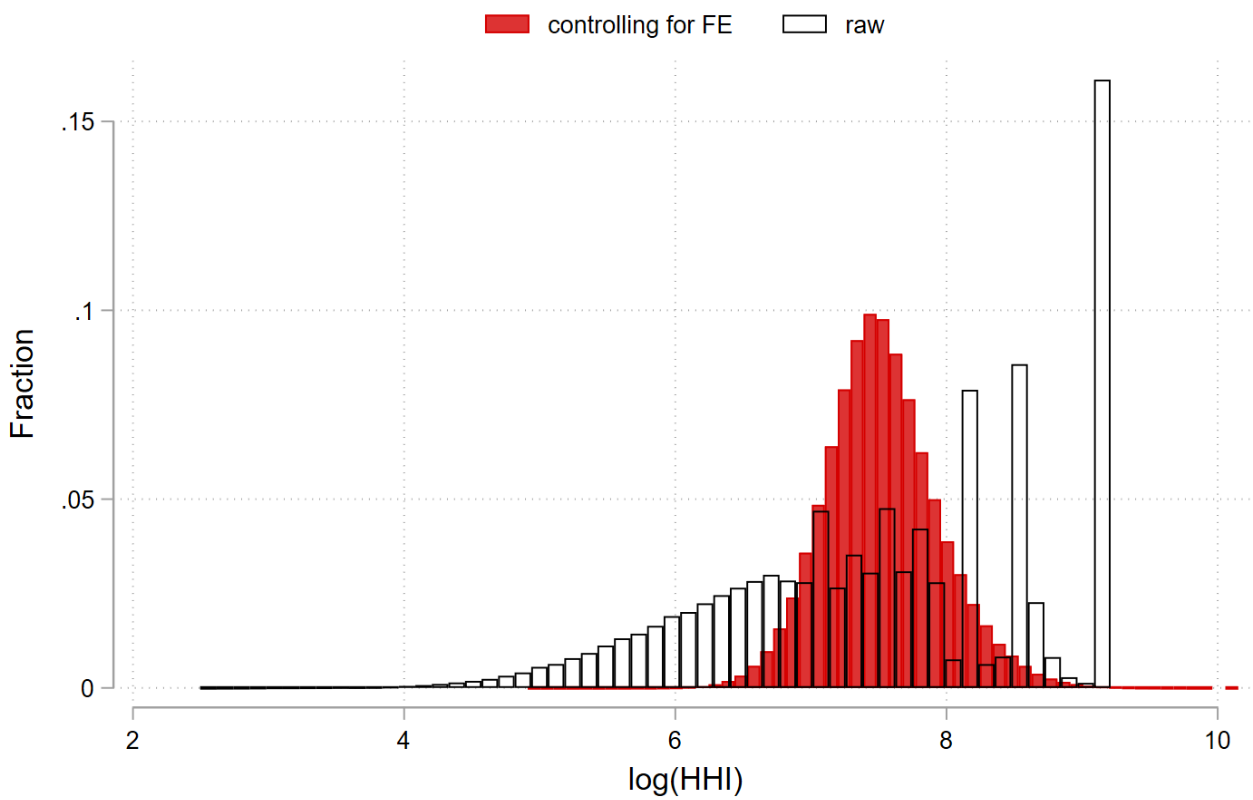

Figure A.5: The relationship between HHI in BGT and ASHE datasets

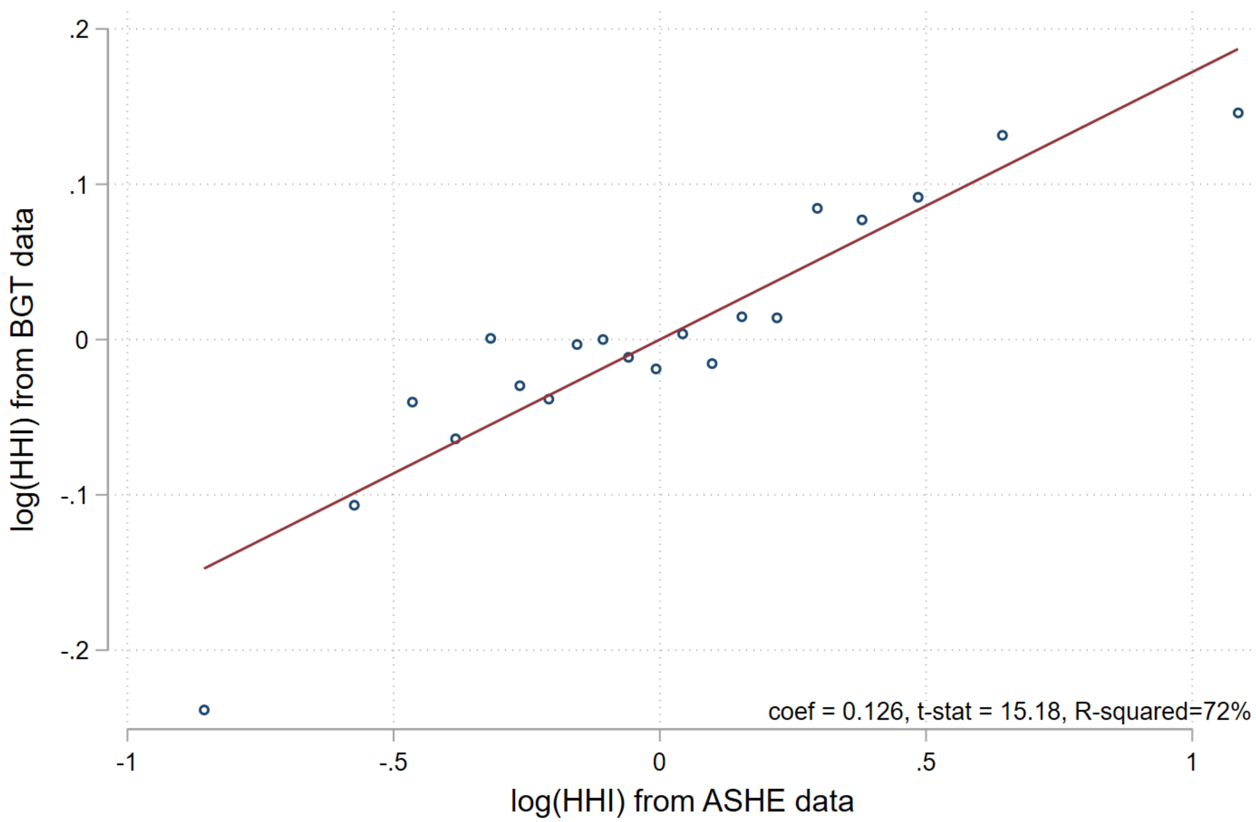

Figure A.6: Occupation-specific results: all flexible vacancies

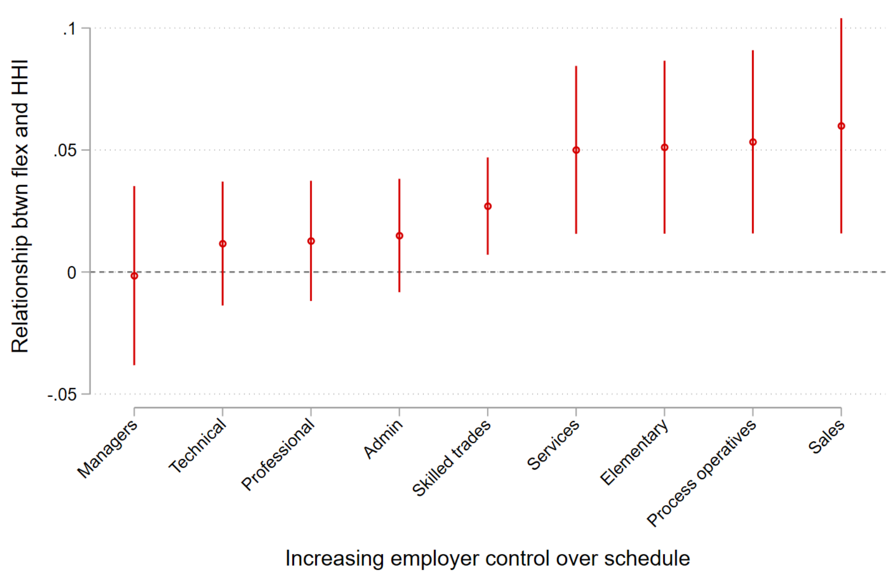

Table A.1: Descriptive statistics: HHI & Local labour market size

	Mean	St. Dev.	10th pct.	90th pct.
No. Vacancies per				
Market	691	4279	8	1371
Firm	26	338	1	16
Firm x Occupation	7.9	81	1	10
Employer concentration				
HHI	3331	3308	365	10000
HHI (vacancy-weighted)	679	1254	35	1605
Predicted HHI	3638	3179	563	10000
Predicted HHI (vacancy-weighted)	1077	1651	76	2653
Firm hiring share	.46	1.5	.015	1
Share of single-establishment firms	.24			
No. Observations				
Market	17933			
Firm	482065			
Firm x Occupation	1571369			

Notes: The first panel gives the number of vacancies within local labour market, firm, and per fim-occupation. The second panel describes employer concentration: the average HHI index of a local labour market and the largest firm-occupation importance weight in the dataset. Note that we define the HHI so that it falls between 1 (perfect competition) and 10,000 (single monopsony employer). The third panel summarises the number of observations as the number of firm-occupations combinations. All statistics are calculated for pooled data with the exception of the middle panel (Employer concentration), where the statistics are summarised over year-quarters.

B Proofs and additional results for the theoretical model

B.1 Proof of Theorem 2

In Theoreom 2, we outline the conditions for the existence of a unique interior solution to the Nash bargaining problem. In general, a maximisation problem will have a unique solution if the function that is being maximised is concave. Below, we derive the second-order derivatives of the joint surplus function $S \equiv [u(w,f) - \bar{U}]^{\beta}[y + \delta(f) - w]^{1-\beta}$ and find the conditions under which the corresponding Hessian matrix is negative semi-definite.

The first-order partial derivatives of S are:

$$\frac{\partial S}{\partial w} = \beta u'_w \frac{S}{u(w,f) - \bar{U}} - (1-\beta)\frac{S}{y + \delta(f) - w} \tag{B.1}$$
$$\frac{\partial S}{\partial f} = \beta u'_f \frac{S}{u(w,f) - \bar{U}} + (1-\beta)\delta'_f \frac{S}{y + \delta(f) - w} \tag{B.2}$$

The second-order partial derivatives of S at the optimal $\{w, f\}$ are:

$$\begin{aligned}
\frac{\partial^2 S}{\partial w^2} &= \beta u''_w \frac{S}{u(w,f) - \bar{U}} - \beta(1-\beta)(u'_w)^2 \frac{S}{(u(w,f) - \bar{U})^2} - u'_w \frac{S}{(u(w,f) - \bar{U})(y + \delta(f) - w)} \\
&= \frac{S}{u(w,f) - \bar{U}} \left(\beta u''_w - \beta(1-\beta)\frac{(u'_w)^2}{u(w,f) - \bar{U}} - \frac{u'_w}{y + \delta(f) - w} \right)
\end{aligned} \tag{B.3}$$

$$\begin{aligned}
\frac{\partial^2 S}{\partial f^2} &= \beta u''_f \frac{S}{u(w,f) - \bar{U}} - \beta(1-\beta)(u'_f)^2 \frac{S}{(u(w,f) - \bar{U})^2} + (1-\beta)\delta''_f \frac{S}{y + \delta(f) - w} \\
&\quad - \beta(1-\beta)(\delta'_f)^2 \frac{S}{(y + \delta(f) - w)^2} + 2\beta(1-\beta)u'_f \delta'_f \frac{S}{(u(w,f) - \bar{U})(y + \delta(f) - w)} \\
&= \beta u''_f \frac{S}{u(w,f) - \bar{U}} + (1-\beta)\delta''_f \frac{S}{y + \delta 9 f) - w} - \frac{\beta}{1-\beta}(u'_f)^2 \frac{S}{(u(w,f) - \bar{U})^2}
\end{aligned} \tag{B.4}$$

$$\begin{aligned}
\frac{\partial^2 S}{\partial f \partial w} &= \beta u''_{wf} \frac{S}{u(w,f) - \bar{U}} - \beta(1-\beta)u'_w u'_f \frac{S}{(u(w,f) - \bar{U})^2} - \beta(1-\beta)u'_f \frac{S}{(u(w,f) - \bar{U})(y + \delta(f) - w)} \\
&\quad + \beta(1-\beta)\delta'_f u'_w \frac{S}{(u(w,f) - \bar{U})(y + \delta(f) - w)} + \beta(1-\beta)\delta'_f \frac{S}{(y + \delta(f) - w)^2} \\
&= \beta \frac{S}{u(w,f) - \bar{U}} \left(u''_{wf} - \frac{u'_f u'_w}{u(w,f) - \bar{U}} - \frac{u'_f}{y + \delta(f) - w} \right)
\end{aligned} \tag{B.5}$$

To determine the sign of the corresponding Hessian, we start by examining the signs of the two second-order derivatives at the optimal $\{w, f\}$. Given the standard assumption that $u'_w > 0$,

$\frac{\partial^2 S}{\partial w^2}$ is negative when the second-order derivative of utilty with respect to wage is weakly negative. Similarly, $\frac{\partial^2 S}{\partial f^2}$ will be unambiguously negative if the second-order derivatives of utility and profit functions with respect to the job feature are weakly negative. We summarise these conditions below:

$$\frac{\partial^2 S}{\partial w^2} < 0 \qquad \text{if} \qquad u'_w > 0, u''_w \leq 0 \tag{B.6}$$
$$\frac{\partial^2 S}{\partial f^2} < 0 \qquad \text{if} \qquad u''_f \leq 0, \delta''_f \leq 0 \tag{B.7}$$

The relationship between worker's and firm's share of the surplus at optimal $\{w^*, f^*\}$ is given by the two FOCs of the bargaining problem:

$$\text{F.O.C.(w)}: \quad \frac{y + \delta(f^*) - w^*}{u(w^*, f^*) - \bar{U}} = \frac{1-\beta}{\beta} \frac{1}{u'_w(w^*, f^*)} \tag{B.8}$$
$$\text{F.O.C.(f)}: \quad \frac{y + \delta(f^*) - w}{u(w^*, f^*) - \bar{U}} = -\frac{1-\beta}{\beta} \frac{\delta'_f(f^*)}{u'_f(w^*, f^*)} \tag{B.9}$$

We use these relationships to find the sign of the determinant of the Hessian:

$$\frac{\partial^2 S}{\partial w^2} = \beta \frac{S}{u(w,f) - \bar{U}} \left(u''_w - \frac{(u'_w)^2}{u(w,f) - \bar{U}} \left((1-\beta) + \frac{1}{1-\beta} \right) \right) \tag{B.10}$$
$$\frac{\partial^2 S}{\partial f^2} = \beta \frac{S}{u(w,f) - \bar{U}} \left(u''_f + \delta''_f u'_w - \frac{(u'_f)^2}{(u(w,f) - \bar{U})} \frac{1}{1-\beta} \right) \tag{B.11}$$
$$\frac{\partial^2 S}{\partial f \partial w} = \beta \frac{S}{u(w,f) - \bar{U}} \left(u''_{wf} - \frac{u'_w u'_f}{u(w,f) - \bar{(U)}} \left(1 - \frac{\beta}{1-\beta} \right) \right) \tag{B.12}$$

Using these expressions, the determinant can be written as

$$\begin{aligned}
\det H &= \frac{\partial^2 S}{\partial w^2} \frac{\partial^2 S}{\partial f^2} - (\frac{\partial^2 S}{\partial f \partial w})^2 \\
&= \left(u''_w - \frac{(u'_w)^2}{u(w,f) - \bar{U}} \left((1-\beta) + \frac{1}{1-\beta} \right) \right) * \left(u''_f + \delta''_f u'_w - \frac{(u'_f)^2}{(u(w,f) - \bar{U})} \frac{1}{1-\beta} \right) \\
&- \left(u''_{wf} - \frac{u'_w u'_f}{u(w,f) - \bar{(U)}} \left(1 - \frac{\beta}{1-\beta} \right) \right)^2
\end{aligned} \tag{B.13}$$

As the next step, we evaluate the above when all second-order derivatives of utility and the

δ function are 0: $u''_w = u''_f = u''_{wf} = \delta''_f = 0$. We can show that, under this assumption, the determinant will be unambiguously positive:

$$\begin{aligned}
\det H &= \left(-\frac{(u'_w)^2}{u(w,f)-\bar{U}}\left((1-\beta)+\frac{1}{1-\beta}\right)\right) * \left(-\frac{(u'_f)^2}{(u(w,f)-\bar{U})}\frac{1}{1-\beta}\right) \\
&\quad - \left(-\frac{u'_w u'_f}{u(w,f)-\bar{(U)}}\left(1-\frac{\beta}{1-\beta}\right)\right)^2 \\
&= \left(\frac{u'_w u'_f}{u(w,f)-\bar{(U)}}\right)^2 \frac{1+3\beta}{1-\beta} > 0
\end{aligned} \tag{B.14}$$

As a consequence, it will be true that $\det H > 0$ when u''_w, u''_f, δ''_f are all weakly negative. Regarding u''_{wf}, a sufficient condition is that it is positive when u'_f is positive, and negative when u'_f is negative: $\text{sgn}(u''_{wf}) = \text{sgn}(u'_f)$.

Finally, because this proof builds on the assumption of an interior solution, this solution will only exist if $u(w,f), \delta(f)$ are such that the first-order conditions can be satisfied. In particular, it follows from the $MRS = MRT$ condition that there must be some f at which the marginal profit of f and its marginal utility are of opposite signs: $\text{sgn}\,\delta'_f \neq \text{sgn}\,u'_f$. This is not an issue for costly amenities and profitable disamenities, where the impact of utility and profit are opposite by definition, but it does call for additional conditions for the case of profitable amenities. It requires $\delta(f)$ to have sufficient curvature so that while marginal profit starts at as positive ($\delta'_f(0) > 0$), it becomes negative at some f. The condition $\delta''_f \leq 0$ thus must hold with strict inequality.

However, a unique (although not necessarily interior) solution exists for some types of job features even if not all the conditions of the Proposition are satisfied. In particular, condition (ii) is not necessary for the existence of a unique solution for costly amenities and profitable disamenities; whether the solution is interior or corner depends on the parameters of the profit and utility functions.[28] On the other hand, condition (ii) is necessary for the existence of an interior solution for profitable amenities. Equation (2.7) states that in an interior solution it must be true that $\text{sgn}(\delta'_f) \neq \text{sgn}(u'_f)$ (condition (i) in Proposition 1). However, for profitable amenities,

[28] For example, when both the utility and profit functions are linear, the solution to the problem is interior if the (constant) marginal utility and marginal profit are the same. If $u'_w \neq \delta'_f$, the firm will provide the maximum or minimum quantity of the job feature, depending whether the net impact of f on the match value is positive or negative (e.g. when the magnitude of marginal profit is greater than the magnitude of marginal disutility, the optimal solution is maximum f).

both the marginal utility and marginal profit start out as positive. In order for an interior solution to exist, the slope of the profit function has to become negative at some f, which is stipulated by condition (ii). If condition (ii) doesn't hold, profitable amenities will be optimally provided at maximum possible amount, a corner solution.

B.2 Proof of Theorem 3

The optimal wage and job feature are implicitly defined in the first-order conditions (equations (B.8) and (B.9)). To see how $\{w^*, f^*\}$ change with $\bar{U}$, we find the total derivatives of both equations around the optimal bundle and solve for $\frac{\mathrm{d}f}{\mathrm{d}\bar{U}}$. The derivatives are:

$$\begin{aligned} \mathrm{d}FOC(w) &= \beta(y + \delta(f^*) - w^*)\left[u''_w \mathrm{d}w + u''_{wf}\mathrm{d}f\right] + \beta\delta'_f u'_w \mathrm{d}f - \beta u'_w \mathrm{d}w \\ &\quad - (1-\beta)[u'_w dw + u'_f \mathrm{d}f - \mathrm{d}\bar{U}] = 0 \qquad \text{(B.15)} \\ \mathrm{d}FOC(f) &= \beta(y + \delta(f^*) - w^*)\left[u''_f \mathrm{d}f + u''_{wf}\mathrm{d}w\right] + \beta\delta'_f u'_f \mathrm{d}f - \beta u'_f \mathrm{d}w \\ &\quad + (1-\beta)\delta'_f[u'_w \mathrm{d}w + u'_f \mathrm{d}f - \mathrm{d}\bar{U}] + (1+\beta)(u(w,f) - \bar{U})\delta''_f \mathrm{d}f = 0 \qquad \text{(B.16)} \end{aligned}$$

Using the optimality condition (2.7), and solving the system for $\frac{\mathrm{d}f}{\mathrm{d}\bar{U}}$ gives us the equilibrium relationship between the quantity of the job feature and workers' outside option:

$$\frac{\mathrm{d}f}{\mathrm{d}\bar{U}} = \frac{(1-\beta)(u''_{wf}u'_w - u''_w u'_f)}{(u'_f)^2\left(\frac{u''_{wf}}{u'_f}\left(2 - \beta[y + \delta(f) - w]\frac{u''_{wf}}{u'_f}\right) - \frac{u''_w}{u'_w}\right) + (u''_f + u'_w\delta''_f)(\beta[y + \delta(f) - w]u''_w - u'_w)} \tag{B.17}$$

The first part of the theorem, that at least one of u''_w, u''_{wf} has to be different from 0, flows directly from the above expression: if both u''_w, u''_{wf} are 0, the numerator of the derivative will be also 0.

This result means that, in order for the optimal f to vary with workers' outside option, the marginal utility of wage u'_w must be a function of wage and/or of the quantity of the job feature. Intuitively, if the marginal utility of wage is independent of wage and the job feature (e.g. the utility function is linear with respect to wage), workers can be compensated for the increase in their outside option by increasing their wage, and this increase in utility will not impact the workers' preferences over the optimal quantity of the job feature. When marginal utility of wage depends on the quantity of the job feature, on the other hand, optimal wage and optimal f are directly linked, and hence will move together. Similarly, when the utility function exhibits diminishing returns to wage, workers' marginal rate of substitution between wage and the job feature will vary as their outside option, and overall compensation, changes. As a result, in either of these cases, optimal f will depend on the outside option $\bar{U}$.

To understand how the sign of the expression depends on the nature of f, we start by noting that the sign of the numerator is the same as sign of u'_f. In the denominator, u''_f and $u'_w\delta''_f$ are both (weakly) negative, and the expression $(\beta(y+\delta(f)-w)u''_w - u'_w)$, is also negative because u''_w is weakly negative. $u''_{wf}u'_f$ is positive since the sings of its two terms are always the same. The only part of the denominator that cannot be easily signed is $(2-\beta''_{wf}u'_f)$. In fact, it will be positive when

$$\frac{u''_{wf}}{u'_f} < \frac{2}{\beta[y+\delta(f^*)-w^*]} \tag{B.18}$$

This condition is the second half of Theorem 3. When it holds, the denominator of $\frac{\mathrm{d}f}{\mathrm{d}\bar{U}}$ is positive. The overall sign of $\frac{\mathrm{d}f}{\mathrm{d}\bar{U}}$ is thus equal to the sign of u'_f.

Inequality B.18 is a sufficient condition. Its interpretation is that $\frac{u''_{wf}}{u'_f}$ is small. Because u''_{wf} has the same sign as u'_f, the marginal utility of wage increases for amenities and decreases for disamenities. If this effect is large – if the curvature of u''_{wf} is large – it means that as the outside option rises, there is some level of f beyond which the worker would demand higher wage, rather than more f even though f is an amenity (and vice versa). For this reason, in order for worker to demand more of amenity f when her outside option improves, the substitutability between wage and f must be relatively small.

The final part of Theorem 3 states that, everything else held constant[29], the response of f to workers' outside option will be greater for costly amenities than for profitable amenities. This result comes from the fact that the sensitivity of the job feature to $\bar{U}$ depends on the *productivity* of the job feature. Productivity denotes the value the job feature adds to the match: a costly

[29] I.e. the preferences for the amenity are fixed, while its impact on the profit function may vary.

amenity when workers value the amenity highly creates greater value than one that is preferred less by workers, or more costly.[30] Intuitively, less productive job features are more responsive to changes in workers' outside option because a greater increase or decrease is needed to compensate the worker for higher $\bar{U}$. For example, when the job feature is a very unpleasant disamenity, a relatively small decrease in f will be sufficient to improve workers' utility enough to compensate for their greater outside option. In contrast, when f is a very costly amenity, it has to increase by a lot to generate enough net value to compensate for higher $\bar{U}$ because some of the utility increase thanks to higher f is offset by its higher costs. It follows from this reasoning that the response of profitable amenities will be relatively muted since they are on average more productive than either costly amenities or profitable disamenities.

B.3 Proof of Theorem 4

Assume a perfectly competitive labour market in which there are no search frictions and firms are price-takers. The worker will receive her full marginal product y. She then decides how much of it to spend on the job feature f: when f is an amenity, she gives up some optimal amount of her wage to cover its cost; when it's a disamenity, she accepts some level of f to increase her monetary remuneration.[31] The optimisation process can be modelled as:

$$u(w_c, f_c) = \max_f \; u(y + \delta(f), f) \tag{B.19}$$

where w_c, f_c denote optimal wage and job feature quantity under perfect competition. The corresponding first order condition is the same as in the bargaining problem (equation 2.7). Since $w_c = y + \delta(f_c)$, firms' profits under perfect competition are 0.

To see how w_c, f_c compare to the solution under monopsony, consider a case where worker's option $\bar{U} = u(w_c, f_c)$. Any bargaining solution must set w_m, f_m such that the worker's utility is at least $u(w_c, f_c)$. This raises two questions: is it possible for bargaining to yield higher utility that perfect competition ($u(w_m, f_m) > u(w_c, f_c)$), and if not, is there another pair of $\{w, f\}$ that

[30]Similarly, a "high productivity" profitable disamenity is one with low distaste for f, or f being very profitable.

[31]When f is a profitable amenity, she similarly chooses some positive quantity of f, stopping before the marginal cost of f doesn't become too large.

delivers $u(w_c, f_c)$?

The answer to the first question is no, because $u(w_c, f_c)$ equals the value of the maximum outside option under which a match is possible – so a worker can never receive more than that. A firm's profit must be at least 0 for the firm to enter the match, so the highest possible wage it would pay is $w_{MAX} = y - \delta(f)$. The maximum utility a worker can earn with w_{MAX} is $u_{MAX} = \max_f u(y - \delta(f), f)$, which coincides with the solution under perfect competition. As a result, $u(w_c, f_c)$ is the highest utility a worker can attain under bargaining.

The above result means that worker's utility when $\bar{U} = u(w_c, f_c)$ is $u(w_c, f_c)$. Is there any other pair of w_m, f_m that could deliver this utility? The answer to this question is also no. To see why, consider again the solution to the worker's maximum utility $u_{MAX} = \max_f u(y - \delta(f), f)$. Under Theorem 2, $u(w, f)$ is well behaved and has a unique solution: $f_{MAX} = f_m = f_c$.

These two results mean that we can treat perfect competition as the upper limit to bargaining: as $\bar{U}$ increases, the bargaining solution is approaching the allocation under perfect competition. As a consequence, we can apply Theorem 3 to understand how other bargaining outcomes differ from $\{w_c, f_c\}$. The Theorem states that f_m increases with $\bar{U}$ when f is an amenity, and falls when it is a disamenity, so f_c is higher than all f_m if f is an amenity, and lower when it is a disamenity. In other words, amenities are under-provided and disamenities and over-supplied when firms have monopsony power.

B.4 Solving the model for heterogeneous firms

In the main text, we solved the model for a case where the worker bargains with a representative firm. In this section, we present the full solution of the model with heterogeneous firm sizes. In this case, the worker faces different outside option depending on the firm she is bargaining with, and the market wage and quantity of job features will be a weighted average of the firm-specific optimal $\{w, f\}$.

B.4.1 Symmetric case

Assume a market in which all firms are of equal size $s_j = s$, and so they offer the same optimal remuneration package $u(w, f)$. For a market with J firms, firm size is $s = 1/J$, and HHI collapses to $1/J$.

The outside option of a worker negotiating with a firm of size $1/J$ is:

$$U_j = \lambda \sum_{i \neq j} s_i u(w_i, f_i) + (1 - \lambda + \lambda s_j) u(b, 0) \tag{B.20}$$

$$= (J-1)\frac{1}{J}\lambda u(w, f) + \left(1 - \lambda + \frac{1}{J}\lambda\right) u(b, 0) \tag{B.21}$$

$$= \lambda u(w, f) + (1-\lambda)u(b, 0) - \frac{1}{J}\lambda[u(w, f) - u(b, 0)] \tag{B.22}$$

The first two terms, $\lambda u(w, f) + (1 - \lambda)u(b, 0)$, are the standard expected value of search. The last term, $\frac{1}{J}\lambda[u(w, f) - u(b, 0)]$, captures the reduction of worker's outside option due to the granularity of search – firm's monopsony power. This last term is negative, and becomes larger when employer concentration $1/J$ increases.

In the next step, we use U_j to model the bargaining process and the optimal $\{w, f\}$. Because firms take their size as given, the key result of the bargaining problem is the same as that derived for the representative firm in section 2, only now the results describe the relationship between outside option and $\{w, f\}$ bundle for an individual firm.

To understand how the market average wage and job feature $\bar{w}$ and $\bar{f}$ respond to employer concentration, we look at the weighted average of firm decisions, using firm size as weights:

$$\bar{w} = \sum_j s_j w_j = J\frac{1}{J}w = w \tag{B.23}$$

$$\bar{f} = \sum_j s_j f_j = J\frac{1}{J}f = f \tag{B.24}$$

Because all firms are the same, the results for $\frac{\mathrm{d}w}{\mathrm{d}U_j}$ and $\frac{\mathrm{d}f}{\mathrm{d}U_j}$ carry through to the market level.

B.4.2 Asymmetric case

The market is defined by its HHI $\sum_j s_j^2$ and a corresponding distribution of firm sizes $g(s_j)$.[32] As defined in section 2.1, worker's outside option when bargaining with firm of size s_j is:

$$U_j = \lambda \sum_{i \neq j} s_i u(w_i, f_i) + (1 - \lambda + \lambda s_j) u(b, 0) \tag{B.25}$$

$$= \lambda \sum_{i} s_i u(w_i, f_i) + (1 - \lambda) u(b, 0) - \lambda s_j (u(w_j, f_j) - u(b, 0)) \tag{B.26}$$

As in the symmetric case, the first two terms correspond to the expected value of search in a standard labour market, while the last term captures the negative impact of the monopsony power of firm j on the worker's outside option. It effectively reduces the general expected value of search by the value of meeting the particular firm, weighted by the probability of such a meeting.

The bargaining problem between a worker and firm j is solved in the same way as in section 2, substituting $\bar{U}$ for U_j. The comparative statics for U_j are derived in the same way.

$$\frac{\mathrm{d}f_j}{\mathrm{d}U_j} = \frac{(1-\beta)(u''_{wf}u'_w - u''_w u'_f)}{(u'_f)^2 \left(\frac{u''_{wf}}{u'_f}\left(2 - \beta\pi_j \frac{u''_{wf}}{u'_f}\right) - \frac{u''_w}{u'_w}\right) + (u''_f + u'_w \delta''_f)(\beta\pi_j u''_w - u'_w)} \tag{B.27}$$

$$\frac{\mathrm{d}w_j}{\mathrm{d}U_j} = -\frac{1-\beta}{\beta\pi_j u''_w - u'_w} + \frac{\mathrm{d}f_j}{\mathrm{d}U_j} \frac{\beta\pi_j u''_{wf} - u'_f}{\beta\pi_j u''_w - u'_w} \tag{B.28}$$

We find that firm-specific wage w_j increases in the worker's bargaining power[33], while the quantity of the job feature offered at firm j, f_j, increases or decreases depending on whether it's an amenity or a disamenity, respectively. Overall, then, the utility value of compensation at firm j, $u(w_j, f_j)$, increases in the worker's outside option.

To understand how market-level wages and job features react to monopsony power, we start by deriving the relationship between firm size and the firm-specific outside option. Using the definition of U_j above, consider how it changes as the worker moves from a smaller to a larger firm (s_j increases), keeping the overall distribution of firm sizes (and the corresponding HHI)

[32] There isn't necessarily a one-to-one mapping between these two objects.

[33] Both parts of the expression for $\frac{\mathrm{d}w_j}{\mathrm{d}U_j}$ are positive, and the magnitude of this derivative increases in the magnitude $\frac{\mathrm{d}f_j}{\mathrm{d}U_j}$.

unchanged. The first two terms of U_j, $\lambda \sum_i s_i u(w_i, f_i) + (1-\lambda)u(b,0)$, are the same regardless of firm size. The third term, $-\lambda s_j(u(w_j, f_j) - u(b,0))$, responds to firm size in two ways:

$$\frac{\partial U_j}{\partial s_j} = -\lambda(u(w_j, f_j) - u(b,0)) - \lambda s_j \frac{\partial u(w_j, f_j)}{\partial s_j} \tag{B.29}$$

First, a higher s_j lowers U_j by decreasing the chance of receiving a job offer from elsewhere, making unemployment more likely if bargaining with firm j breaks down. Second, the value of the remuneration package $u(w_j, f_j)$ changes.

To derive the size of $\partial u(w_j, f_j)/\partial s_j$, we decompose it into the change in the utility of firm-specific compensation with respect to the outside option, and the change in the outside option with respect to firm size. Substituting in for the latter using equation B.29 and solving for $\partial u(w_j, f_j)/\partial s_j$ gives us:

$$\begin{aligned}
\frac{\partial u(w_j, f_j)}{\partial s_j} &= \frac{\mathrm{d}u(w_j, f_j)}{\mathrm{d}U_j}\frac{\partial U_j}{\partial s_j} && \text{(B.30)}\\
&= \frac{\mathrm{d}u(w_j, f_j)}{\mathrm{d}U_j}\left(-\lambda(u(w_j, f_j) - u(b,0)) - \lambda s_j \frac{\partial u(w_j, f_j)}{\partial s_j}\right) && \text{(B.31)}\\
&= -\left[\frac{\mathrm{d}u(w_j, f_j)}{\mathrm{d}U_j}\lambda(u(w_j, f_j) - u(b,0))\right]\left[1 + \frac{\mathrm{d}u(w_j, f_j)}{\mathrm{d}U_j}\lambda s_j\right]^{-1} < 0 && \text{(B.32)}
\end{aligned}$$

The sign of this derivative is negative because of our earlier result that utility of compensation increases in the worker's outside option, $\frac{\mathrm{d}u(w_j,f_j)}{\mathrm{d}U_j} > 0$.

Returning to $\frac{\partial U_j}{\partial s_j}$, we can show that it will be negative as long as the elasticity of utility compensation with respect to firm size is smaller than 1:[34]

$$\begin{aligned}
\frac{\partial U_j}{\partial s_j} < 0 \qquad \text{if} \qquad & -\lambda(u(w_j, f_j) - u(b,0)) - \lambda s_j \frac{\partial u(w_j, f_j)}{\partial s_j} < 0 && \text{(B.33)}\\
& u(w_j, f_j) - u(b,0) > -s_j \frac{\partial u(w_j, f_j)}{\partial s_j} && \text{(B.34)}\\
& 1 > \frac{-s_j}{u(w_j, f_j) - u(b,0)}\frac{\partial u(w_j, f_j)}{\partial s_j} && \text{(B.35)}
\end{aligned}$$

To show that this always holds, we return to our solution for $\frac{\partial u(w_j,f_j)}{\partial s_j}$ and use it to derive its

[34] The right-hand side expression collapses to a standard expression for elasticity when $u(b,0) = 0$.

elasticity. (For simplicity, we set $u(b,0)=0$.)

$$-\frac{\partial u(w_j,f_j)}{\partial s_j}\frac{-s_j}{u(w_j,f_j)} = \left[s_j\frac{\mathrm{d}u(w_j,f_j)}{\mathrm{d}U_j}\lambda u(w_j,f_j)\right]\left[u(w_j,f_j)\left(1+\frac{\mathrm{d}u(w_j,f_j)}{\mathrm{d}U_j}\lambda s_j\right)\right]^{-1} \tag{B.36}$$

$$= \left[s_j\lambda\frac{\mathrm{d}u(w_j,f_j)}{\mathrm{d}U_j}\right]\left[1+s_j\lambda\frac{\mathrm{d}u(w_j,f_j)}{\mathrm{d}U_j}\right]^{-1} < 1 \tag{B.37}$$

To summarize, this result shows that worker's outside option is smaller when they meet a larger firm. As a result, larger firms pay less, and offer more disamenities and fewer amenities than smaller firms.

Finally, to analyse the relationship between market wages and job features and market-level employer concentration, we look at how $\bar{w}$ and $\bar{f}$, the weighted average of firm-specific wages and job features, change with HHI. Consider a situation where we increase the size of a (weakly) larger firm s_j and decrease the size of a (weakly) smaller firm s_k by the same amount ($\mathrm{d}s_j=-\mathrm{d}s_k$), so that $HHI=\sum_i s_i^2$ increases. The sizes of all other firms are left the same, so that to derive $\frac{\mathrm{d}\bar{w}}{\mathrm{d}\sum_i s_i^2}$ we only need to derive how $\bar{w}$ changes with respect to $\mathrm{d}s_j$. Using the definition of $\bar{w}=\sum_i s_i w_i$, we get:

$$\frac{\mathrm{d}\bar{w}}{\mathrm{d}\sum_i s_i^2} = \frac{\mathrm{d}\bar{w}}{\mathrm{d}sj} = \frac{\mathrm{d}(s_jw_j+s_kw_k+\sum_{i\neq j,k}s_iw_i)}{\mathrm{d}sj} = \frac{\mathrm{d}(s_jw_j+s_kw_k)}{\mathrm{d}sj} \tag{B.38}$$

$$= w_j+s_j\frac{\mathrm{d}w_j}{\mathrm{d}s_j}-w_k-s_k\frac{\mathrm{d}w_k}{\mathrm{d}s_k} \tag{B.39}$$

Because smaller firms pay more, w_j-w_k is negative, and a shift away from smaller towards larger firms (an increase in HHI) will reduce the market wage. The second effect concerns the change in the firm-specific wages: the wage at the larger firm will fall, while that of the smaller firm will increase, undoing some of the fall in the market wage. The overall impact on market wage will be negative if the firm wage elasticity with respect to firm size is smaller than 1 for all values of s_j:

$$w_j + s_j \frac{\mathrm{d}w_j}{\mathrm{d}s_j} < w_k + s_k \frac{\mathrm{d}w_k}{\mathrm{d}s_k} \tag{B.40}$$

$$\frac{w_j}{w_k} + \frac{w_j}{w_k}\frac{s_j}{w_j}\frac{\mathrm{d}w_j}{\mathrm{d}s_j} < 1 + \frac{s_k}{w_k}\frac{\mathrm{d}w_k}{\mathrm{d}s_k} \tag{B.41}$$

$$\frac{w_j}{w_k}(1-\epsilon_w) < 1-\epsilon_w \tag{B.42}$$

This condition is similar to our earlier result that $-\frac{\partial u(w_j,f_j)}{\partial s_j}\frac{-s_j}{u(w_j,f_j)} < 1$ (equation B.37). When this condition holds, the qualitative result for the relationship between market wage and employer concentration is negative, the same as in the model with representative firm. The result for $\bar{f}$ can be derived analogously.

C Representativeness of BGT data

Table A.2: Summary statistics, 2014-2019

	BGT data	ONS data
number of vacancies (millions)	46	55.4
share salaried (%)	74.8%	60.2%
share permanent (%)	66.3%	92.1%
share with wage info (%)	63.1%	
$< £9$	22.1%	24%
$£9 - £15$	39.9%	36.7%
$£15 - £20$	17.3%	16.1%
$> £20$	20.8%	23.2%

Notes: ONS data on the number of vacancies comes from the Vacancy Survey. The ONS data in the rest of the table comes from the Annual Survey of Hours and Earnings (and therefore measures eht stock of employees). Data is pooled over the 2014-2019 period.

Table A.3: Descriptive statistics for vacancies with and without wage information

	wage missing		difference
	wage missing	wage included	-
N. of skills	5.0558	5.3602	-0.3045
	(4.6999)	(5.5735)	(0.0015)
Permanent contract (%)	0.5799	0.4308	0.1491
	(0.4936)	(0.4952)	(0.0002)
Schedule flexible (%)	0.2244	0.1755	0.0488
	(0.4172)	(0.3804)	(0.0001)
Managers (%)	0.1080	0.1102	-0.0022
	(0.3104)	(0.3131)	(0.0001)
Professionals (%)	0.3414	0.3495	-0.0081
	(0.4742)	(0.4768)	(0.0001)
Technical (%)	0.1747	0.1722	0.0025
	(0.3797)	(0.3776)	(0.0001)
Administrative (%)	0.0893	0.0785	0.0108
	(0.2852)	(0.2689)	(0.0001)
Skilled trades (%)	0.0637	0.0596	0.0040
	(0.2442)	(0.2368)	(0.0001)
Caring, leisure serv. (%)	0.0561	0.0575	-0.0014
	(0.2302)	(0.2328)	(0.0001)
Customer service (%)	0.0910	0.0937	-0.0028
	(0.2876)	(0.2915)	(0.0001)
Process operatives (%)	0.0344	0.0282	0.0062
	(0.1822)	(0.1654)	(0.0001)
Elementary (%)	0.0414	0.0506	-0.0092
	(0.1993)	(0.2192)	(0.0001)

Notes: The number of required skills, the share of permanent and flexible contracts, and the distribution of vacancies over major occupation groups, by vacancies with wage information (column (1)) and without (column(2)). The third column calculates the difference between the two. Standard errors in parentheses.

Figure A.7: Distribution of vacancies and employment in BGT and ASHE data

(a) By occupation

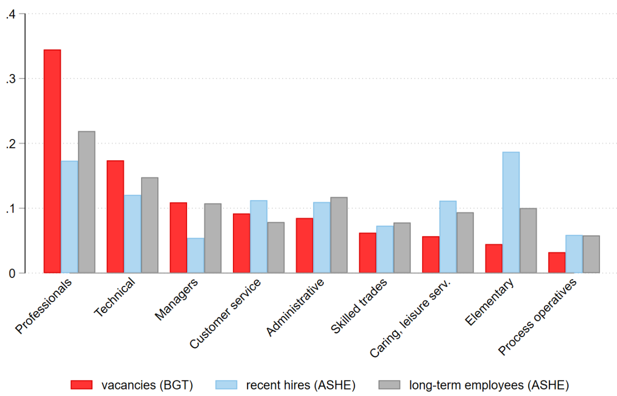

(b) By region

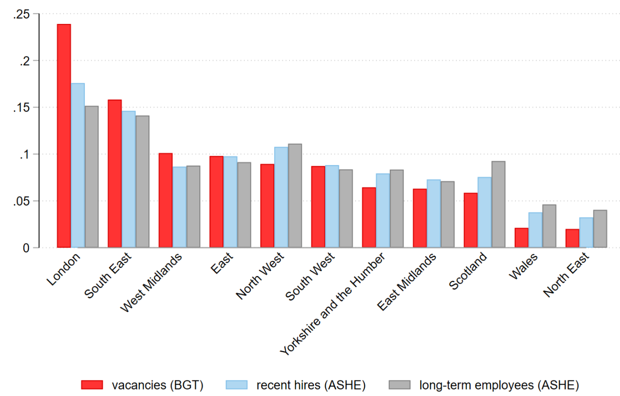

(c) By wage

Notes: Pooled data over 2014-2019. Occupations represent occupation groups at 1-digit SOC level.

Figure A.8: Representativeness of BGT data over time

(a) The number of vacancies

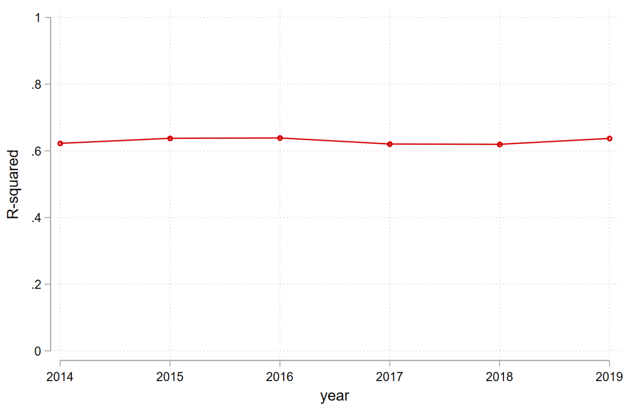

(b) Median wage

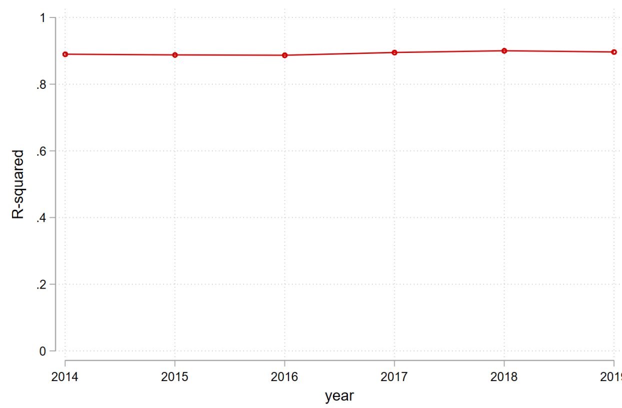

Notes: Each data point is derived by regressing the number of vacancies (panel (a)) or median wage (panel (b)) in occupation-county cells in the BGT data on the corresponding number of employees in the ASHE data, within a given year, and collecting the R-squared. The flat profile of both figures suggests that the represenativeness of BGT relative to ASHE stays constant over time.

D Natural Language Processing Details

Algorithm 1: Training the Logistic Regression model $(\mathbb{C}, \mathbb{D})$

Result: Return the tuple $W = (\hat{\beta}_1^*, \hat{\beta}_2^*, ..., \hat{\beta}_C^*)$ containing the vocabulary weights for each class $c \in \mathbb{C}$

extract the vocabulary in $\mathbb{D}$

for $c \in \mathbb{C}$ **do**

  for $\lambda \in [0.01, 0.05]$ **do**

    let $j = 0$

    repeat

      compute as N_c the number of vacancies in class c which are labelled with 1

      sample N_c vacancy texts labelled with 0 and the same number of text labelled with 1 and call the sample $\mathbb{D}_S$

      randomly divide documents in $\mathbb{D}_S$ into 5 equally sized blocks $\mathcal{S}_j$, for $j = 1, ..., 5$

      for $j = 1, ..., 5$, with $\Pr_c(c_i|x_i)$ defined as in Equation (3.2), find $\hat{\beta}_{c,\lambda,\ell}$ which minimises

$$\sum_{i \in \mathbb{D}_S \setminus \mathcal{S}_\ell} \log \Pr_c(c_i|x_i) - \frac{\lambda}{2}||\beta_{c,\lambda}||_1$$

      where $|| \cdot ||_1$ is the 1-norm while excluding the observations in block ℓ

      compute the cross-validated balanced F-measure, as defined in Equation (3.4), using observations in block ℓ and $\hat{\beta}_{c,\lambda,\ell}$

      increase j by 1

    until $j = 500$;

  end

  select the λ with the lowest average cross-validated balanced F-measure and call the corresponding estimate $\hat{\beta}_c^*$

end

Algorithm 2: Apply the Logistic Regression model $(\mathbb{C}, V, W, d)$

Result: Assign class $c \in \mathbb{C}$ to vacancy i

construct a vector representation x_i according to the vocabulary

for $c \in \mathbb{C}$ **do**

compute the probability $\Pr_c(0|x_i)$ and $\Pr_c(1|x_i)$ using the coefficients $\hat{\beta}_c^*$ in weight vector W as

$$\Pr_c(1|x_i) = \frac{\exp(x_i'\hat{\beta}_c^*)}{1 + \exp(x_i'\hat{\beta}_c^*)}, \quad \Pr_c(0|x_i) = 1 - \Pr_c(1|x_i)$$

if $\Pr_c(1|x_i) > \Pr_c(0|x_i)$ **then**

assign class c to vacancy i

end

end