

Hirsch, Boris; Jahn, Elke J.; Manning, Alan; Oberfichtner, Michael

Working Paper

The Wage Elasticity of Recruitment

IZA Discussion Papers, No. 15675

Provided in Cooperation with:

IZA – Institute of Labor Economics

Suggested Citation: Hirsch, Boris; Jahn, Elke J.; Manning, Alan; Oberfichtner, Michael (2022) : The Wage Elasticity of Recruitment, IZA Discussion Papers, No. 15675, Institute of Labor Economics (IZA), Bonn

This Version is available at:

<https://hdl.handle.net/10419/267412>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

DISCUSSION PAPER SERIES

IZA DP No. 15675

The Wage Elasticity of Recruitment

Boris Hirsch
Elke J. Jahn
Alan Manning
Michael Oberfichtner

OCTOBER 2022

DISCUSSION PAPER SERIES

IZA DP No. 15675

The Wage Elasticity of Recruitment

Boris Hirsch

*Leuphana University of Lüneburg, IWH
and IZA*

Elke J. Jahn

*Institute for Employment Research,
University of Bayreuth and IZA*

Alan Manning

*London School of Economics and Centre
for Economic Performance*

Michael Oberfichtner

Institute for Employment Research and IZA

OCTOBER 2022

Any opinions expressed in this paper are those of the author(s) and not those of IZA. Research published in this series may include views on policy, but IZA takes no institutional policy positions. The IZA research network is committed to the IZA Guiding Principles of Research Integrity.

The IZA Institute of Labor Economics is an independent economic research institute that conducts research in labor economics and offers evidence-based policy advice on labor market issues. Supported by the Deutsche Post Foundation, IZA runs the world's largest network of economists, whose research aims to provide answers to the global labor market challenges of our time. Our key objective is to build bridges between academic research, policymakers and society.

IZA Discussion Papers often represent preliminary work and are circulated to encourage discussion. Citation of such a paper should account for its provisional character. A revised version may be available directly from the author.

ISSN: 2365-9793

IZA – Institute of Labor Economics

Schaumburg-Lippe-Straße 5–9
53113 Bonn, Germany

Phone: +49-228-3894-0
Email: publications@iza.org

www.iza.org

ABSTRACT

The Wage Elasticity of Recruitment*

One of the factors likely to affect the market power of employers is the sensitivity of the flow of recruits to the offered wage, but there is very little research on this. This paper presents a methodology for estimating the wage elasticity of recruitment and applies it to German data. Our estimates of the wage elasticity of recruitment are about 1.4. We also report evidence that high-wage employers are more selective in hiring, in which case the relevant recruitment elasticity should be higher, about 2.2. Together with prior estimates of the quit elasticity these results imply that wages are 72–77% of the marginal product of labour. Further, we find lower elasticities for recruits hired from non-employment as well as for women, non-German nationals, non-prime-age workers, less skilled workers, and workers with less complex jobs.

JEL Classification: J42, J31

Keywords: monopsony, imperfect labour markets, wage elasticity of recruitment

Corresponding author:

Elke J. Jahn
Bayreuth University
Universitaetsstrasse 30
95447 Bayreuth
Germany
E-mail: Elke.Jahn@uni-bayreuth.de

* We would like to thank Marie Paul for her helpful comments. We further appreciate comments by participants of the EALE 2022, IAAE 2022, SOLE 2022, Vfs 2022 conferences and a Bank of Portugal/Spain conference and by seminar participants at Hebrew University Jerusalem, HKUST, RHUL, Drexel University, in Berlin, Halle, Lüneburg, and Nuremberg.

1 Introduction

There has been a resurgence of interest in recent years in the idea of monopsony as relevant for understanding the labour market (see the surveys by Manning, 2011, 2021). To measure the monopsony power of employers, a natural measure is the wage elasticity of the labour supply curve to the firm (Robinson, 1933).

A common approach to estimating the labour supply elasticity has been to use a simple dynamic model in which a higher wage makes it easier for firms to recruit and retain workers. In a steady state, the supply of labour $N(w)$ to a firm paying wage w can be written as:

$$N(w) = \frac{R(w)}{q(w)} \quad (1)$$

where $R(w)$ denotes the number of recruits with $R' > 0$ and $q(w)$ the quit rate of incumbent workers with $q' < 0$. From equation (1), the wage elasticity of the labour supply curve ε_{Nw} is given by the difference of the recruitment elasticity ε_{Rw} and the quit rate elasticity ε_{qw} :

$$\varepsilon_{Nw} = \varepsilon_{Rw} - \varepsilon_{qw} \quad (2)$$

The elasticity of labour supply to the firm is an attractive measure of monopsony power because, in a simple model of monopsony, it directly determines the wage as a fraction of the marginal product of labour MPL :

$$w = \frac{\varepsilon_{Nw}}{\varepsilon_{Nw} + 1} MPL \quad (3)$$

There is a large literature estimating the quit rate elasticity: 868 out of the 1,320 estimates of the labour supply elasticity considered by Sokolova and Sorensen (2021) rely on estimates of ε_{qw} or related concepts. Most of these studies rely on a theoretical result in Manning (2003) that, on average, the wage elasticity of recruitment for job-to-job moves should be equal in absolute value to the wage elasticity of the quit rate. The intuition is that one firm's quit is

another firm's recruit. The result does not apply to hires from non-employment and depends on the assumption that the probability of moving from one firm to another depends on the relative wage, which while it is plausible is not beyond criticism. Given this it is important to have independent estimates of the recruitment elasticity to see whether recruitment and separations are equally sensitive to the wage.

A few papers provide direct estimates of the recruitment elasticity: Falch (2017) uses quasi-experimental variation in wages to study teacher recruitment in Norway finding a short-run recruitment elasticity of around 1.4. He concludes that for this specific market the recruitment elasticity is similar in magnitude to the quit elasticity estimated in Falch (2010). Dal Bó et al. (2013) provide evidence from a field experiment that randomly assigned wages for public sector positions in Mexico. They find a recruitment elasticity of around 2, which they decompose into an effect of the wage on the likelihood that a job offer will be accepted by the worker (elasticity around 1.1) and an effect on the number of applicants (elasticity around 0.8).

Related to these estimates of the recruitment elasticity are studies of the application elasticity. Marinescu and Wolthoff (2020) find that within job titles the application elasticity is around 0.8 using observational data from a US online job board. A field experiment in the UK by Belot et al. (2022) also yields an application elasticity of this size. While applications and recruitment are related they are not the same, so these estimates are not directly informative on the recruitment elasticity (see also Azar et al., 2022). Finally, Dube et al. (2022) find only a small effect of the wage on the ease of filling tasks in on-demand online labour markets.

One reason for the much larger literature on the quit than recruitment elasticity is methodological. There is only one current job a worker could leave (ignoring the minority who have multiple jobs at once) and the wage in that job is observable. This means it is easy to estimate models of the quit rate. But there are many firms which a worker might join (becoming a recruit for the destination firm) and, at best, we observe the wage in the firm which they do join.

This paper develops a method that can be used to study the recruitment elasticity directly. The key idea is to use a sample of recruits and estimate a model for the probability that the recruit is hired by the firm they joined as opposed to others they might have joined; this is what we call the recruit probability function. We show that the elasticity of the recruit probability function with respect to the wage is informative about the recruitment elasticity.

To overcome the problem that we do not observe the recruits' wages at those firms they might have joined, we use the firm wage effect of a two-way fixed-effects wage decomposition along the lines of Abowd, Kramarz, and Margolis (1999, AKM henceforth) as the wage variable. The AKM decomposition splits up individual workers' wages into permanent worker and firm components and has been found to provide a suitable approximation of the West German wage structure (Card et al., 2013). In the AKM framework, the firm fixed wage effect represents a pure wage premium from working for this rather than another firm adjusted for observed and unobserved worker quality (Card et al., 2018; Hirsch and Müller, 2020) and thus captures the firm's wage policy. Along these lines, a recent paper by Bassier et al. (2022) uses the AKM firm wage effects as the wage variable when estimating the quit rate elasticity.

We apply these ideas to the study of recruitment in the Hamburg area of Germany. We use administrative data on the universe of recruits to the private sector between July 1, 2013 and June 30, 2014. Our baseline estimate of the recruitment elasticity is around 1.4 and thus similar in absolute value to the estimated quit rate elasticity for the Hamburg area in Hirsch et al (2022) which is 1.2. Combining the estimated recruitment and quit rate elasticities, equation (3) implies that wages will be 74% of the marginal product of labour.

We further document substantial heterogeneities in the recruitment elasticities of workers hired from employment and from non-employment and across subgroups of workers. Specifically, we find that the recruitment elasticity from employment is larger than the recruitment elasticity from non-employment, which we expect to find if firms that pay higher wages find it easier to poach workers from other firms. Estimated recruitment elasticities are substantially

lower for women than for men; for non-German than for German nationals; for both younger and older workers than for prime-age workers; for less skilled workers; and for workers with less complex jobs. These findings are in line with previous studies on heterogeneities in quit rate elasticities and suggest that employers have more monopsony power over these less-wage elastic groups of workers.

Our framework so far assumes, as is common in the monopsony literature, that the flow of recruits to different firms is determined solely by the preferences of workers, ignoring the possibility that employer selection plays any role. We next present a simple model of employer selection and show how our recruitment elasticity needs to be modified to take it into account. Using data from the German Vacancy Survey, we document that high-wage employers are more selective in hiring workers. To take account of employer selection, our baseline recruitment elasticity should perhaps be raised by about 0.8. With the modified recruitment elasticity, equation (3) would now imply wages are 77% of marginal products.

The remainder of this paper is organised as follows: Section 2 derives our estimation method and explains its econometric implementation. Section 3 describes our data and provides descriptive evidence. Section 4 presents and discusses the results from our econometric approach, Section 5 discusses employer selection, and Section 6 concludes.

2 Estimation of the wage elasticity of recruitment

2.1 The recruit probability function

Assume that the rate at which an individual i with characteristics \mathbf{x}_i is recruited to firm f that pays wage w_f and has other characteristics \mathbf{z}_f is given by:

$$R_{if} = R(w_f, \mathbf{z}_f, \mathbf{x}_i) \tag{4}$$

The dependence of the flow of recruits on individual characteristics might be because some

recruits are more likely to move than others.¹ We assume that firms have a wage policy represented by the wage w_f so that high-wage firms pay higher wages to all workers; this does not preclude that firms pay different wages to workers with different characteristics. This latter assumption is consistent with our use in estimation of the AKM firm wage premium as our measure of firms' wage policy.

Define $p(w_f, \mathbf{z}_f, \mathbf{x}_i)$ to be the fraction of recruits with characteristics \mathbf{x}_i who move to firm f , to which we refer as the recruit probability function. The recruit probability function is given by:

$$p(w_f, \mathbf{z}_f, \mathbf{x}_i) = \frac{R(w_f, \mathbf{z}_f, \mathbf{x}_i)}{\sum_{f'} R(w_{f'}, \mathbf{z}_{f'}, \mathbf{x}_i)} = \frac{R(w_f, \mathbf{z}_f, \mathbf{x}_i)}{H(\mathbf{x}_i)} \quad (5)$$

where $H(\mathbf{x}_i)$ is the total flow of recruits with characteristics \mathbf{x}_i , i.e. the sum of the recruit functions $R(w_f, \mathbf{z}_f, \mathbf{x}_i)$ over all firms (including f). Our empirical strategy involves estimating the recruit probability function and then using this estimate to infer the recruit function. If we make the common assumption that the individual firm is small in relation to the market as a whole, so the total flow of recruits $H(\mathbf{x}_i)$ in the denominator of (5) does not depend on the wage offered by any single firm, the wage elasticity of the recruit probability function is exactly equal to the elasticity of the recruit function.²

2.2 Econometric implementation

Our estimation approach is based on the estimation of the recruit probability function (5). This is attractive because we can estimate this function on a sample of recruits. Specifically, we use as a baseline empirical model for the recruit probability function a multinomial logit form, that

¹ (4) is expressed at the individual level while the flow of recruits at firm level is what is required for an assessment of employer market power. The flow of recruits to firm f will be given by $R_f = \sum_i R_{if}$ so that knowledge of how R_{if} varies with the wage allows us to work out the recruitment elasticity at firm level.

² Our empirical method does not depend on assuming individual firms are small as the denominator in (5) is modelled as a fixed effect so that the estimated wage elasticity of the recruit probability function will be that of the recruit function.

is:

$$p(w_f, \mathbf{z}_f, \mathbf{x}_i) = \frac{\exp(\phi \log w_f + \boldsymbol{\beta}_1 \mathbf{z}_f + \boldsymbol{\beta}_2 \mathbf{d}_{if})}{\sum_{f'} \exp(\phi \log w_{f'} + \boldsymbol{\beta}_1 \mathbf{z}_{f'} + \boldsymbol{\beta}_2 \mathbf{d}_{if'})} \quad (6)$$

In this baseline specification, the controls are the firm’s wage policy, other firm characteristics, and some “distance” variables \mathbf{d}_{if} representing how attractive firm f is to individual i .

The individual characteristics only enter (6) through interactions with either the wage or other firm characteristics as a “level” effect would affect numerator and denominator equally and disappear from (6). This does not mean that individual characteristics will not affect the rate at which workers join firms, but does mean a level effect is not relevant for the recruit probability function and its wage elasticity. Using such a multinomial logit form can be given a micro-foundation as in Card et al. (2018) if recruits have idiosyncratic preferences over working for different firms that follow a type 1 extreme value distribution.

Although (6) has a multinomial logit structure, it is hard to estimate because there is a very large number of firms that workers could conceivably move to resulting in a multinomial logit model with a number of alternatives that is too large to estimate directly. To implement the model, we exploit the equivalence of the likelihoods of the multinomial logit model and the Poisson model (Baker, 1994; Guimarães, 2004) that allows estimation of (6) through a fixed-effects Poisson regression at the level of the recruit on an expanded data set with one observation for every possible firm–worker pair. The dependent variable is an indicator taking the value one for the firm that hired the recruit and zero otherwise. An individual worker fixed effect is also included. The estimates are identical to those from the infeasible multinomial logit regression. But, in contrast to the multinomial logit, the fixed-effects Poisson model is feasible, albeit computationally demanding due to the large number of observations of the expanded data set, which has as number of observations the product of the number of recruits and the number of firms.

3 Data and descriptive analysis

3.1 Administrative linked employer–employee data

We use this approach to study recruitment between July 1, 2013 and June 30, 2014 in the labour market region of Hamburg, Germany. This region is depicted in Figures 1 and 2 and comes from an updated classification by Kosfeld and Werner (2012) based on commuting links. In choosing the local labour market for our analysis we followed three criteria. We looked, first, for a region with a large number of recruits and firms that has, second, a core–periphery structure and thus can be considered as a well-defined self-contained local labour market and has, third, no national borders nearby and thus no cross-border commuting, which we could not observe in our data. Hamburg fulfils all these three criteria: It is the second-largest city in Germany, and it is at the centre of a local labour market with a clear core–periphery structure and no national border nearby.³ Examining Hamburg yields a data set with almost 500 million observations that is sufficiently large while still tractable.

Our linked employer–employee data on recruits and employers in the Hamburg area combine two administrative data sets: the Integrated Employment Biographies (IEB) and the Establishment History Panel (BHP; Ganzer et al., 2020a), which are both provided by the Institute for Employment Research (IAB). The data on workers and their characteristics (age, education, occupation, job’s task, gender, and nationality) stem from the IEB (for details on the IEB, see Jacobebbinghaus and Seth, 2007). The IEB includes all wage and salary employees registered with the German social security system, in total about 80% of employment. Information on plants (we observe plants in our data rather than firms) come from the BHP that also consists of data from the German social security system aggregated at the level of the plant at

³ We decided against examining the even bigger local labour market of Berlin because part of post-unification Berlin and all its surrounding districts were part of the former GDR and there is still a marked East–West divide in many labour market outcomes (e.g. Schnabel, 2016) and because of the nearby Polish border.

June 30 of a year (for details on the BHP, see Ganzer et al., 2020b). The information in these administrative data is used to calculate social security payments and it is thus highly reliable.

To these two data sets we merge the AKM plant wage effects calculated by Bellmann et al. (2020) for the IEB based on full-time workers aged 18–60. Following the methodology from Card et al. (2013), Bellmann et al. (2020) provide separate AKM effects for five time intervals since 1985. We make use of the AKM plant effects for the estimation period 2011–2018. In the AKM framework, which has been shown to provide a good approximation of the West German wage structure (Card et al., 2013), the plant wage effect gives the log wage premium enjoyed by every worker holding a job at this plant adjusted for observed and unobserved worker quality. In other words, it represents a pure wage premium adjusted for observed and unobserved worker quality and thus is a suitable measure of a plant’s wage policy (Card et al., 2018; Hirsch and Müller, 2020).

There are two elements to our data set: the selection of employers and the selection of workers. We select for the sample of employers all plants located in the local labour market of Hamburg (i.e. the dark blue shaded districts in Figures 1 and 2) in the private sector, excluding temporary work agencies, that existed on June 30, 2013 and employed more than ten full-time equivalent workers (where we count part-time workers as half workers). We restrict our sample to the private sector because wages in the public sector are centrally determined so the notion of a plant’s wage policy hardly applies. We further exclude temporary work agencies, as we do not have information on temporary workers’ user plants and their location and thus cannot infer these workers’ actual commute, which will be a core regressor in our models. Finally, we do not consider plants with ten or less full-time equivalent workers because these are not subject to German dismissal protection.

Note that given the multinomial logit structure assumed in our empirical model (6) we do not have to consider the full universe of potential employers as a choice set when estimating the model. Conditional on the choice being in a restricted set of options, the probability of

choice within that set also follows a multinomial logit structure (McFadden, 1974). That said, the method is likely to work better with a wide range of employers and our sample selection is designed to balance this.

Apart from the AKM plant wage effects, our regressions also include plant characteristics drawn on June 30, 2013, i.e. before the hiring window starts. Plant controls include information on workforce composition (worker shares in terms of gender, nationality, age, education, and job complexity) as well as dummies for two-digit industry and plant location (the districts depicted in Figures 1 and 2).⁴

Distance to the job will be an important variable. The IEB GEO provides the exact geographic location of worker’s residence and workplace (see Ostermann et al., 2022). We use these geocodes to compute commuting distance; we experiment with a number of measures but our main measure is the linear distance (in km) between the worker’s residence and workplace.⁵

Our final sample of employers, which constitutes the recruited workers’ choice set, includes 10,134 plants. Descriptive statistics are shown in Table 1 for all plants and the subsample of 7,895 plants that hire at least one worker during our period of observation. As is clear from Table 1, the vast majority of plants are recruiting some workers in our observational window and there are little differences in characteristics between the wider set of all plants in the Hamburg area and the narrower set of hiring plants.

Turning to recruits, we select our worker sample in the following way. First, we restrict to full-time workers aged 18–60 (used in the identification of the AKM wage effects by Bellmann et al., 2020) who live in the local labour market of Hamburg or surrounding districts (i.e. the dark blue and light blue shaded districts depicted in Figures 1 and 2). This is to make sure

⁴ With respect to education, we distinguish low-skilled, medium-skilled, and high-skilled workers where low-skilled workers are workers with neither a vocational nor an academic degree, medium-skilled workers those with a vocational degree, and high-skilled workers those with an academic degree. With respect to job complexity, we distinguish jobs requiring simple tasks, expert tasks, specialist tasks, and complex tasks.

⁵ We use the Open Source Routing Machine (OSRM) provided by Huber and Rust (2016) and an offline version of OpenStreetMap for 2014.

that we have workers who would have a high probability of working for a member of our employer sample. Second, we consider the set of all new recruits into one plant of our employer sample between July 1, 2013 and June 30, 2014. We only include hires where the new job lasts at least 6 months, which is the probation period in Germany, so that it is reasonable to think of the hires as permanent ones.

The choice set for these workers is assumed to be all employers in the labour market region of Hamburg irrespectively of the fact whether they are hiring or not, i.e. the employers that we selected in the first step (located in the dark blue shaded districts depicted in Figures 1 and 2). But we do a number of robustness checks in which we vary the choice set. In the expanded data set with one observation for every recruit–plant pair, the dependent variable is one for the plant where the worker is hired and zero for all other plants in the choice set.

As worker characteristics, we have the usual demographics (gender, age, education, and whether the recruit has German nationality) and, if applicable, information about the task complexity of the recruit’s previous job (see fn. 4 about the different education and task complexity categories). We also know the exact residential location of the worker which we use, combined with geocoded information on employer location, to compute the linear distance to every plant in our employer sample. In some cases, we further use the information on the previous employer of the worker, which is not restricted to have been in our employer sample.

Moreover, we know whether recruits are hired from employment or non-employment. We define a hire as being from employment if we observe the worker in a job with a different plant within a month of starting the current job. Hires from non-employment are those with a previous spell in registered unemployment or no prior spell in our data. This latter possibility may mean the worker has been unemployed without receiving benefits or was employed without being covered by social security (e.g. as a self-employed). Although we cannot disaggregate this category of unknown origin in our data, information from other German data sets suggests

that the vast majority of workers in this category have indeed started new jobs from non-employment (e.g. Hirsch et al., 2018).

3.2 Descriptive evidence

Our final worker sample is made of 48,311 recruits where 26,168 or 54.2% come from employment and the remaining 22,143 or 45.8% from non-employment. Table 2 reports descriptive statistics for the sample of recruits and the subsamples of recruits hired from employment and from non-employment. The last column of Table 2 further reports the share of hires from employment and documents that different subgroups of recruits (in terms of gender, nationality, education, and age) differ substantially in the probability of being hired from employment.

It is noticeable that groups of workers with less favourable labour market outcomes, such as women, non-German nationals, and low-skilled workers, are less likely to start a new job after a job-to-job move. This suggests that these groups of workers find it harder to improve their outcomes through job changes compared to their male, German, or more skilled counterparts and we will therefore check for heterogeneities in the recruitment elasticities across these groups in our later regression analysis.

Turning to the wage sensitivity of recruitment, Figure 3 presents a binned scatterplot of the fraction of recruits hired by all plants within a certain decile of the distribution of the AKM plant wage effects against the average AKM plant wage effect in that decile. In this plot, the fraction of recruits serves as the sample counterpart of the recruit probability function (5), so its relation to the wage provides us with a simple descriptive measure of the wage elasticity of recruitment.

In line with expectations, we see that on average plants that pay higher wages attract more recruits. Turning to the extremes, just 5.6% of recruits move to the 10% lowest-paying plants, whereas 12.1% of recruits and thus roughly twice as much move to the 10% highest-paying plants. The slope of the regression line in Figure 3 is about 0.12 meaning that on average

10% larger wages are accompanied by an increase in the fraction of recruits by about 1.2 percentage points. Since the average fraction of recruits over the deciles is 0.1 by construction and since the AKM plant wage effects are in logs, the slope of the regression line of 0.12 implies a wage elasticity of the fraction of recruits of 1.2 ($= 0.12/0.1$).

Figure 3 groups together all recruits no matter whether they are hired from employment or from non-employment. Yet, there are good reasons to expect that the recruits' responsiveness to wages differs depending on their previous labour market status. As is shown by Manning (2003: 100, Proposition 4.5), the wage elasticity of recruitment from employment will be larger than the wage elasticity of recruitment from non-employment if paying a higher wage raises the share of hires coming from employment. Intuitively, the latter reflects that employers paying higher wages find it easier to poach workers from other employers (Hirsch et al., 2018). Available empirical evidence indeed points in that direction as several studies document a significant positive relationship between wages and the share of hires from employment (e.g. Manning, 2003; Hirsch et al., 2018; Webber, 2022).

Our data are no exception as is obvious from Figures 4 and 5 that show separate binned scatterplots for recruitment from employment and from non-employment, respectively. For both types of recruitment we find a positive relationship between the fraction of recruits and wages, but the relationship is markedly steeper for recruits hired from employment than it is for recruits hired from non-employment. For recruitment from employment the regression line has slope 0.16, whereas for recruitment from non-employment the slope is just 0.06. These slopes, in turn, imply wage elasticities of 1.6 ($= 0.16/0.1$) for the fraction of recruits from employment and 0.6 ($= 0.06/0.1$) for the fraction of recruits from non-employment, respectively.

4 Estimation results

4.1 Baseline estimates of the recruitment elasticity

We now turn to the results of applying our econometric approach detailed in Section 2 to the data on recruits in the local labour market of Hamburg. For each of the 48,311 workers in our recruit sample, we first create 10,134 observations corresponding to each of the plants in our employer sample, so the total number of observations amounts to almost 490 million. The dependent variable takes the value one for the plant that the recruit moves to and zero otherwise. On these expanded data, we then run fixed-effects Poisson regressions where the fixed effect is at the level of the recruit. The core regressor in the models is the AKM plant wage effect, which is in logs and whose coefficient, therefore, informs us on the wage elasticity of recruitment we are looking for.

Table 3 presents some baseline estimates of the wage elasticity of recruitment from successively richer fixed-effects Poisson regressions. The starting point is a simple regression that just includes the AKM plant wage effect (model 1). We then successively add the linear distance (in km) between the worker's residence and plant location (model 2); plant controls capturing the workforce composition, that is the shares of male and non-German workers, the shares of workers aged below 30 and of workers aged 45 or older, the shares of low-skilled and high-skilled workers, and the shares of workers on jobs that require simple, specialist, and complex tasks (model 3); dummies for two-digit industry and plant location at district level (model 4).

The estimated recruitment elasticity for models 1–4 all lie in the range 1.1–1.4 and all estimates are statistically significant at the 1% level. Notably, in the simple fixed-effects Poisson regression that only includes the AKM effect (model 1) the estimate is 1.2, i.e. an increase of the firm AKM effect by 10% increases the probability of recruitment by 12%. This estimate is similar to the wage elasticity of the fraction of recruits from our descriptive analysis in the previous section. In our preferred specification with the richest set of controls (model 4), the

estimated elasticity of 1.4 is somewhat larger. Our estimates of the recruitment elasticity are similar in magnitude to those found in other studies, e.g. Dal Bó et al. (2013) and Falch (2017).

One interesting feature is that these estimates of the recruitment elasticity are very similar to estimates of the quit rate elasticity suggesting that the assumption of equality of these elasticities commonly made in the literature is not unreasonable; Hirsch et al. (2022) estimate a wage elasticity of about -1.2 for job separations in the local labour market of Hamburg.⁶ The implied overall labour supply elasticity is well within the range of estimates reported by Sokolova and Sorensen (2021) and very similar in magnitude to previous estimates for Germany based on quit rate elasticities, e.g. by Hirsch et al. (2010, 2018, 2022). Hirsch et al (2022) estimated the quit rate elasticity for the Hamburg area to be 1.2. Combining the estimated recruitment and quit rate elasticities, equation (3) implies that wages will be 74% of the marginal product of labour.

When included as a control variable, we find a strong negative impact of distance on the recruitment probability that is stable across specifications and statistically significant at the 1% level throughout. An increase in the linear distance between a worker's residence and plant location by 1 km reduces the probability of the worker's recruitment substantially by about 10%.

The models presented in Table 3 are for total recruitment and do not distinguish whether recruits are hired from employment or from non-employment. Yet, as detailed in the previous section and corroborated by descriptive evidence, we might expect recruitment from employment to be more wage-elastic than recruitment from non-employment. Therefore, Table 4 presents estimates of the wage elasticity of recruitment from employment (panel A) and from non-employment (panel B) from successively richer fixed-effects Poisson regressions akin to models 1–4 in Table 3.

⁶ More specifically, this estimate comes from a stratified Cox regression at the worker level for the separation rate of jobs held by male workers in the Hamburg area (akin to the elasticity estimates used in model I from Table 2 in Hirsch et al., 2022).

In all four specifications of panel A, we obtain estimates of the recruitment elasticity from employment around 1.6 that are all statistically significant at the 1% level. Again, this number coincides with the wage elasticity of the fraction of recruits from employment from the descriptive analysis above.

As expected, the estimates for the recruitment elasticity from non-employment in panel B are lower than that number. They are all statistically significant at the 1% level and lie in the range 0.6–1.1 where the estimate in the simple model without controls again coincides with the descriptive wage elasticity of the fraction of recruits from non-employment. Notably, estimates get bigger in richer specifications, which is behind the somewhat larger overall recruitment elasticity in Table 3 in the richer specifications with extensive sets of control variables (models 3 and 4). In our preferred specification that controls for workforce composition, industry, and plant location (model 4), the estimated recruitment elasticity from non-employment is 1.1 and thus only about two-thirds of the estimated recruitment elasticity from employment of 1.6.

Turning to linear distance as the other key regressor in models 2–4, the estimated coefficient of -0.1 is statistically significant at the 1% level throughout and is the same for recruitment from employment and non-employment and across specifications. As before, an increase in the linear distance between a worker’s residence and plant location by 1 km reduces the probability of recruitment markedly by about 10%.

4.2 Heterogeneity in the recruitment elasticity across workers

In the fixed-effects Poisson regressions, worker characteristics are absorbed in the recruit fixed effect and thus they are controlled for. Still, worker characteristics may be important as interactions to the AKM plant wage effect because groups of workers may differ in their responsiveness to wages in the hiring process. In particular, such interactions may reflect that groups of workers face a varying intensity of search frictions as suggested by the descriptive analysis of group differences in the share of hires from employment in Section 3.2.

What is more, such interactions would square with a growing body of evidence that employers possess more monopsony power over groups of workers who have lower quit rate elasticities, which, in turn, explains part of the wage differences between these and more wage-elastic groups (see Hirsch et al., 2010; Ransom and Oaxaca, 2010; Hirsch, 2016 for women vs. men and Hirsch and Jahn, 2015, for immigrants vs. natives).

Against this background, we next run several fixed-effects Poisson regressions including interactions of the AKM plant wage effect with the recruit's gender, nationality, education, and age. If part of wage differences across groups documented in previous research indeed reflected their wage responsiveness in the hiring process, we would expect groups of workers with more favourable labour market outcomes to have a higher recruitment elasticity, such as men compared to women, German compared to non-German nationals, more skilled workers, and prime-age workers.

Table 5 reports substantial differences in the recruitment elasticity along all these four dimensions that are in line with expectations and that are all statistically significant at the 1% level. To start with, the estimated recruitment elasticity is 1.5 for male but just 1.2 for female workers, and it is 1.5 for Germans and drops to zero for non-German nationals, where the latter suggests that non-German nationals take up any job available to them irrespectively of the wage offered. Turning to education, we obtain the lowest recruit elasticity of 0.6 for low-skilled workers, followed by 0.9 for medium-skilled workers and 2.9 for high-skilled workers. Finally, we find the largest elasticity estimate of 1.7 for prime-age workers aged 30–44 compared to 1.2 for workers aged under 30 and 1.3 for workers aged 45 or older.⁷

Interestingly, we obtain corresponding interaction effects for the distance between a worker's residence and plant location, with the only exception of workers aged 45 or older. Those groups of workers who are less responsive to wages are at the same time more responsive

⁷ Note that the same patterns also show up in fully interacted models that add all these interactions simultaneously. Yet, as these are harder to read, we only report them in the appendix (see Appendix Table A1).

to distance. For example, an increase in the distance between the worker's place of residence and plant location by 1 km reduces the probability of recruitment by 10.3% if the worker has German nationality and by even 11.6% for a non-German national. These differences in the distance responsiveness strongly suggest that less wage-elastic recruits like women or non-German nationals are less willing to commute, which arguably limits the geographic scope of their job search and thus contributes to less favourable outcomes. The gender differences are in line with Le Barbanchon et al. (2020), who document a similar pattern for jobseekers in France, and with Caldwell and Daniele (2022), who conclude that women face a larger cost of distance in the German labour market.

With a few exceptions, the group differences in recruitment elasticities also show up when estimating separate models for recruits hired from employment and for recruits hired from non-employment (see Appendix Tables A2 and A3). The two exceptions are male and young recruits aged under 30 hired from non-employment who do not differ significantly in their recruitment elasticity from their female and prime-age counterparts, respectively.

These estimates have assumed that the plant wage fixed effect is the same for all groups of workers. But it could be that the plant wage fixed effect is driven by plant-level productivity and the pass-through of productivity to wages is lower for groups of workers over which employers have some market power. Under this interpretation the lower wage elasticity for women could be not because women are less responsive to wage differentials but because the plant wage effect is flatter for women than men. It is hard to provide precise estimates of AKM plant wage effects for sub-groups because of small sample sizes but one study for Portugal (Card et al., 2016) found that women receive only 90% of the firm-specific pay premium earned by men. An effect of this magnitude is insufficient to overturn our conclusion as we find that women's recruitment elasticity is 82% of the male level.

For recruits from employment we further have information on their previous job, which allows us to include interactions of the AKM plant wage effect with job's task complexity. We

find that recruits are much more wage-elastic when their previous job was more complex (see Appendix Table A3). The recruitment elasticity from employment is 0.5 for simple tasks, 1.1 for expert tasks, and 2.5 (2.9) for specialist tasks (complex tasks), with these differences in the elasticity being statistically significant at the 1% level. At the same time, the responsiveness to distance is more pronounced for the less wage-elastic recruits with less complex previous jobs.

In summary, we find clear evidence of heterogeneities in recruitment elasticities that strongly suggest that workers with less favourable labour market outcomes, such as females, non-German nationals, the less educated, non-prime-agers, and those working in less complex jobs, are less wage-elastic. Notably, the same groups are also more responsive to the distance to prospective employers, which suggests that they have a lower willingness to commute that, in turn, is likely to limit the geographic scope of their job search.

So far, we have considered how the recruitment elasticity varies with demographic characteristics. Next, we investigate how it varies with the AKM worker wage effect, which measures the worker's permanent skills and other factors that are rewarded equally across all employers and which thus can be interpreted as a summary measure of how well a worker fares in the labour market (Card et al., 2013). The first column of Table 6 includes an interaction of the AKM plant wage effect with the AKM worker wage effect centred around its mean so that the main effect of the AKM plant wage variable is an estimate of the recruitment elasticity for the average worker. The estimate implies that workers who have higher wage effects have a significantly larger recruitment elasticity. The elasticity is about 1.4 for a recruit with an AKM worker wage effect at the mean, and it rises significantly by about 0.13 for every additional 10 log points in the AKM worker wage effect. This finding squares with the group heterogeneities in the previous subsection where we found that groups of recruits who tend to achieve better labour market outcomes are also more wage-elastic.

All the heterogeneities investigated so far have related to permanent characteristics of

the worker. Next, we investigate whether the position of recruits' previous job in the distribution of the AKM plant wage effects matters for their recruitment elasticity with respect to the AKM wage effect of potential employers. For brevity, we refer to these two plant effects as the source and destination AKM plant effects, respectively. Specifically, we interact the destination AKM plant effect with dummy variables for the quartiles of the distribution of the source AKM plant effects. As the information on the source AKM plant effect is only available for job-to-job movers, this specification is restricted to recruitment from employment.

There are several reasons why source AKM plant effects might matter though not all reasons predict the same sign of effect. On the one hand, recruits coming from high-wage employers face fewer attractive options in their choice set (i.e. fewer employers offering wage increases) than recruits coming from low-wage employers. For this reason, we would expect recruits to be less responsive to wages the better was their position in the distribution of the source AKM plant effects. On the other hand, job-to-job movers with a low source AKM plant effect are recruits who had ended up in a job with a low-wage employer in the first place, this group may be disproportionately made of workers with low recruitment elasticities.

Because we have already shown pronounced heterogeneities in the recruitment elasticity across groups of workers, and the disadvantaged groups (females or non-German nationals) are disproportionately more often hired from low-wage than from high-wage employers, it is important to control for these factors when investigating the impact of source plant AKM effects. To prevent this channel driving our results, we include all the group interactions with the AKM plant wage effect explored in the previous section (that is, we run a fully interacted model for recruitment from employment akin to the one for overall recruitment reported in Appendix Table A1).

Figure 6 plots the coefficients of the destination AKM plant effect and linear distance for each quartile of the distribution of the source AKM plant effects, showing that the recruit-

ment elasticity is significantly higher for recruits hired from high-wage employers.⁸ The estimated recruitment elasticity is 4.3 for recruits hired from employers in the top quartile of the distribution of the source AKM plant effects and 2.1 for recruits from the third quartile. In stark contrast, recruits coming from the bottom half of the distribution of the source AKM plant effects are hardly responsive to wages, with estimated elasticities of 0.2 and -0.5 for recruits from the second and first quartile, respectively. Our findings suggest that recruits hired from low-wage employers are disproportionately made of workers with low recruitment elasticities, who may end up in low-wage employers for that reason.

4.3 Robustness

We will now scrutinise the robustness of our recruitment elasticity estimates to changes in the model specification, such as workers' choice set of employers, changing the sample of recruits, investigating the functional form of the wage effect, and alternative distance measures.

Changing the choice set

Our first check of robustness varies the choice set of workers to address the concern that firms that are not looking to recruit workers should arguably not be in the choice set. First, we alter the choice set to plants that hired at least one worker in our period of observation. As is seen from the second column of Table 6, restricting to hiring plants hardly changes the estimate of the recruitment elasticity that is now 1.5 and thus very close to the baseline estimate of 1.4 (from model 4 of Table 3).

However, it may be the case that not all firms are looking to recruit at the same time that a worker is looking for a job. To address this concern, we take the start date of a recruit and construct as a choice set all the plants that recruited a worker in a window around this start date and thus can be presumed to be engaging in some hiring activity around the time the worker

⁸ Note that we decided against adding confidence bands to the coefficient plot presented in Figure 6, which would be overly narrow because of the smallish standard errors.

was searching. We use two windows – three and five weeks, chosen in part because a large proportion of job starts (68.9%) are on the first of the month. For a recruit who starts on the first of a month the three-week window will exclude start dates on the first of the neighbouring months while the five-week window will include them. The choice sets constructed in this way are recruit-specific, something that can be handled easily using the Poisson transformation. The results are reported in columns 3 and 4 of Table 6. The estimated elasticities are lower when we restrict choice sets in this way, implying employers have greater market power.

The AKM wage effects are measured with some error which may lead to attenuation bias in the estimated recruitment elasticity. To address this concern, we use as a choice set employers who had at least 30 workers as of June 30, 2013 following the result of Bonhomme et al. (2022), which suggest that such a bias should be small for plants of this size. McFadden (1974) showed that restricting the choice set to a subset of the true choice set should lead to the same estimates if the model is correctly specified and column 5 of Table 6 shows this is the case.

Changing the worker sample

We also conduct further robustness checks that explore whether our results change with the definition of our worker sample. The first regression reported in the sixth column of Table 6 only includes workers living in the local labour market of Hamburg (i.e. the light blue area of Figures 1 and 2), and the second regression reported in the seventh column restricts to workers living in the city of Hamburg (i.e. the core of that local labour market, the dark blue area). Reassuringly, both models yield recruit elasticities that are almost identical to the baseline estimate.

Functional form

Our baseline specification assumes an iso-elastic recruitment function with a wage elasticity that is constant along the distribution of the AKM plant wage effects. One plausible alternative

to this case is that recruitment is more responsive to wages in the middle of the AKM distribution where employer density is much larger than in the tails of the distribution where employer density is low. This case is plausible because a given increase in the wage offer (induced by being offered a higher AKM plant wage effect) is accompanied by a much larger rise in the ranking of potential employers in the middle of the distribution than in its tails.

To explore this possibility, we run a check of robustness that adds interactions between the AKM plant wage effect and dummies for the quartiles of the distribution of the AKM plant wage effects to the fixed-effects Poisson regression. To rule out that outliers in the AKM effects affect our findings, we symmetrically trim our data and drop both the top 2% and bottom 2% of observations from the regression. The core results from this robustness check are shown in Figure 7 that plots the log predicted values from the fixed-effects Poisson regression (evaluated at the means of all other regressors) along with 95% confidence bands against the AKM plant wage effect. Figure 7 is clearly in line with the notion that the probability of recruitment is particularly responsive in the middle of the distribution of the AKM plant wage effects and less so in the tails. The wage elasticity of recruitment is 1.4 in the first, 3.4 in the second, 1.0 in the third, and zero in the fourth quartile of the AKM distribution. This finding of more market power at the extremes of the wage distribution is in line with the estimates reported by Langella and Manning (2021) for separation elasticities.

Different measures of distance

In a final set of robustness checks reported in Table 7, we check whether our results change when using different distance measures in geographic space and in other dimensions. Since not all measures are available for recruits hired from non-employment, we focus on recruitment from employment. To ease comparison, the first column of Table 7 repeats the baseline estimate that we obtain when using the linear distance between workers' place of residence and plant

location. The second column instead uses linear distance based on the recruit's place of residence at the end of the previous job and thus prior to hiring. In consequence, this check of robustness allows us to assess whether our results could be due to recruits who move near their new employer. This does not seem to be the case as the estimated recruitment elasticity and the impact of distance on the recruit probability are left unaltered by this change in the distance measure.

The third and fourth columns use the street distance (in km) and the commute time (in minutes) between the recruit's place of residence and plant location rather than linear distance. We obtain these two distance measures using OpenStreetMap 2014 and the user-written command by Huber and Rust (2016). Again, using these alternative distance measures has hardly any impact on the estimate of the recruitment elasticity from employment; our results seem robust to the use of different distance measures.

So far, our regressions have controlled for the distance between the recruit and possible employers in geographic space. Employers and workers could be nearer or farther away from each other in other dimensions as well. As a case in point, taking up a job with a certain employer could require a switch in occupation or in industry meaning that this employer is much farther away from the recruit in these dimensions than is an alternative employer from the same industry offering a job in the same occupation. We therefore include a dummy for a plant from the same two-digit industry as the previous employer and a dummy for a plant employing workers in the same three-digit occupation as in the previous job. As is clear from the fifth column of Table 7, the two dummy variables have a substantial, positive influence on the recruit probability that is orthogonal to the influence of distance in geographic space. That said, controlling for employer distance along these dimensions has only little impact on the estimate of recruitment elasticity from employment that is now 1.4 compared to the baseline estimate of 1.6.

5 Employer selection: Evidence from the German Vacancy Survey

Our approach so far has assumed that the destination of recruits represents only worker choice over employers; this is in line with the approach of Card et al. (2018). But it is also possible that employer selection plays a role in determining the destination of recruits. This section sketches a simple way to think about employer selection and how one might assess its importance in practice.

Suppose the firm only hires a fraction θ of those willing to join the firm. A firm will only turn some of these workers away if the average productivity declines in θ : denote this by $y(\theta)$. Given this, total employment is $N(w) = R(w)/q(w)$ where $R(w) = \theta P(w)$ is the flow of the recruits selected by the firm, that is the fraction θ of potential recruits $P(w)$ actually joining the firm. Our approach so far implicitly assumed that all recruits willing to join the firm are also hired by the firm (and is nested as the case where $\theta = 1$) so that potential and actual recruitment coincide.

With employer selection, profits are given by:

$$\pi = [y(\theta) - w]N(w) = [y(\theta) - w]\theta \frac{P(w)}{q(w)} \quad (7)$$

The first-order condition for choice of θ can be written as:

$$[y(\theta) - w] + \theta y'(\theta) = 0 \quad (8)$$

(8) could be solved for a solution $\theta(w)$. From (7) the first-order condition for the wage can be written as:

$$w = \frac{\varepsilon_{Nw}}{\varepsilon_{Nw} + 1} y[\theta(w)] \quad (9)$$

where $\varepsilon_{Nw} = \varepsilon_{Pw} - \varepsilon_{qw}$. (9) tells us that the wage elasticity of the total flow of potential recruits ε_{Pw} is relevant for assessing the market power of employers as measured by the gap

between the marginal product and the wage. But the actual flow of recruits will be $R(w) = \theta(w)P(w)$, so the elasticity we have estimated so far will be $\varepsilon_{Rw} = \varepsilon_{\theta w} + \varepsilon_{Pw}$; we will refer to this as the estimated recruitment elasticity. If higher-wage employers are more selective (i.e. $\varepsilon_{\theta w} < 0$) this is an under-estimate of the elasticity needed to measure employer power.

The data we have used so far cannot be used to investigate employer selection; all we observe are workers who are actual recruits, not those who wanted to be but were not hired by the firm. To investigate employer selection, we turn to the German Vacancy Survey (GVS; Bossler et al., 2021a, 2021b). The GVS is a repeated cross-sectional survey of plants focusing on vacancies and hires. Sample plants are interviewed for four consecutive quarters starting in the fourth quarter of a year and a new sample of plants is drawn every year. The interviewed plants are a stratified random sample drawn from all plants who employ at least one worker subject to social security. The survey can be merged with administrative data which allows us to link plants in the GVS with their estimated AKM wage effect.

In the fourth quarter of each year, the questionnaire includes several items that refer to hires in the previous twelve months. The questionnaire asks for the total number of hires and, for plants that hired a worker in the previous twelve months, the number of applicants for the latest filled position in total and by gender. We follow our earlier sample restrictions regarding sector and type of employment, but include plants from all of Germany to obtain a sufficiently large sample. To cover approximately the same period as in our earlier analysis, we use the survey waves from the fourth quarters of 2013 and 2014.

As a check that results from the GVS are broadly similar to those reported earlier, we first estimate negative binomial regressions for the number of recruits over the past year as a function of the AKM wage plant effect. These take account of the count data nature of the dependent variable and the zeros in the data in particular. Estimates are reported in Table 8. The starting point is a simple regression that just includes the AKM plant wage effect and survey year effects (model 1). We then successively add control variables for the plant's workforce

composition (model 2) and for industry and plant location (model 3). Reassuringly, for total recruits we obtain a recruitment elasticity similar to what we found earlier.

We next consider whether the probability of being selected varies with the wage. For the last worker hired, we have information on the number of applicants. We estimate models for the number of applicants with different controls for the hire's occupation. Our preferred specification includes three-digit occupation dummies yielding an applicant elasticity of 0.8, see third column of Table 9. We don't know how many workers were hired but if the number of workers hired is uncorrelated with the wage (a prominent special case satisfying this condition is if only one worker is hired to fill each vacancy), the selection elasticity will be minus the applicant elasticity, so 0.8 should be added to our earlier estimates of the recruitment elasticity to obtain the best measure of employer market power. We obtain similar estimates of the applicant elasticity if we restrict the sample to employers where only one worker was hired in the past year where we know the vacancy led to only one hire, see column 4. With the modified recruitment elasticity, equation (3) would now imply wages are 77% of marginal products. In summary, allowing for employer selection reduces the estimated market power of employers, but it remains substantial.

Earlier in the paper we showed that the estimated recruitment elasticity differs across groups of workers. This difference might be the result of differences in the selection elasticity across groups (perhaps because of discrimination) rather than differences in labour supply behaviour. In the GVS, the only dimension we can consider is gender; we know the fraction of women in the applicant pool and can see whether a woman is hired. If there is no gender difference in employer selection, the probability of a woman being hired should reflect the share of female applicants and should not be related to the plant wage. We investigate this in Table 10. Our estimates do not suggest differential employer selection as the explanation for our earlier results.

6 Conclusions

This paper provides a method to estimate the wage elasticity of recruitment. The core idea is to estimate for a sample of recruits the wage elasticity of the probability that a recruit joins a specific employer rather than other employers the recruit could have moved to. Given the large number of potential employers, computational issues are important and we propose using the equivalence between the multinomial logit and a fixed-effects Poisson regression. To deal with the problem that we do not observe recruits' wages at those employers they did not move to, we use the employer wage effect from an AKM decomposition as a measure of an employer's wage policy.

We applied this estimation approach to administrative data on the universe of recruits to the private sector in the labour market region of Hamburg, Germany who were hired between July 1, 2013 and June 30, 2014. Our baseline estimate of the recruitment elasticity is around 1.4. In line with theory, we found that recruitment from employment is more wage-elastic than recruitment from non-employment. We further present some evidence from the German Vacancy Survey that high-wage employers are more selective in hiring applicants; correcting for this implies that the relevant recruitment elasticity for measuring market power is at most 2.2. Combined with prior estimates of the quit elasticity, these estimates imply that wages amount to 72–77% of the marginal product of labour.

We further documented marked differences in the recruitment elasticity across groups of workers where groups who typically achieve more favourable labour market outcomes also show a larger recruitment elasticity. Specifically, we found that men are more elastic than women, German nationals more elastic than non-German nationals, and prime-age workers more elastic than younger or older workers as are more skilled workers and workers on jobs involving more complex tasks. These findings square with the previous literature that found

lower quit rate elasticities for most of these groups as well and argued that these elasticity differences provide employers with additional monopsony power over the less elastic groups. Notably, we found that groups of workers who are less responsive to wages are at the same time more responsive to the distance to their prospective employer. This finding suggests that these workers are less willing to commute, which, in turn, limits the scope of their job search and is thus one candidate explanation for employers' additional monopsony power over them.

References

- Abowd, J. M., Kramarz, F., and Margolis, D. N. (1999), “High wage workers and high wage firms”, *Econometrica*, 67(2): 251–333
- Azar, J., Berry, S., and Marinescu, I. (2022), Estimating labor market power, NBER Working Paper No. 30365, National Bureau of Economic Research, Cambridge, MA
- Baker, S. G. (1994), “The multinomial–Poisson transformation”, *Journal of the Royal Statistical Association Series D (The Statistician)*, 43(4): 495–504
- Bassier, I., Dube, A., and Naidu, S. (2022), “Monopsony in movers: The elasticity of labor supply to firm wage policies”, *Journal of Human Resources*, 57(S): S50–S86
- Bellmann, L., Lochner, B., Seth, S., and Wolter, S. (2020), AKM Effects for German Labour Market Data, FDZ-Methodenreport 01/2020, Institute for Employment Research, Nuremberg
- Belot, M., Kircher, P., and Muller, P. (2022), “How wage announcements affect job search – a field experiment”, *American Economic Journal: Macroeconomics*, 14(4): 1–67
- Bonhomme, S., Holzheu, K., Lamadon, T., Manresa, E., Mogstad, M., and Setzler, B. (2022), “How much should we trust estimates of firm effects and worker sorting?”, *Journal of Labor Economics*, forthcoming
- Bosler, M., Gartner, H., Kubis, A., and Küfner, B. (2021a), The IAB Job Vacancy Survey: Establishment survey on labour demand and recruitment processes, waves 2000 to 2018 and subsequent quarters 2006 to 2019, FDZ-Datenreport 09/2021, Institute for Employment Research, Nuremberg
- Bosler, M., Gürtzgen, N., Kubis, A., Küfner, B., and Popp, M. (2021b), “German Job Vacancy Survey of the IAB – Version 0018 v1”, Research Data Centre of the Federal Employment Agency (BA) at the Institute for Employment Research (IAB), doi: 10.5164/iab.iabse0018.de.en.v1
- Caldwell, S. and Danieli, O. (2022), Outside options in the labor market, UC Berkeley, mimeo
- Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018), “Firms and labor market inequality: Evidence and some theory”, *Journal of Labor Economics*, 36(S1): S13–S70
- Card, D., Cardoso, A. R., and Kline, P. (2016), “Bargaining, sorting, and the gender wage gap: Quantifying the impact of firms on the relative pay of women”, *Quarterly Journal of Economics*, 131(2), 633–686
- Card, D., Heining, J., and Kline, P. (2013), “Workplace heterogeneity and the rise of West German wage inequality”, *Quarterly Journal of Economics*, 128(3): 967–1015
- Dal Bó, E., Finan, F., and Rossi, M. A. (2013), “Strengthening state capabilities: The role of financial incentives in the call to public service”, *Quarterly Journal of Economics*, 128(3): 1169–1218
- Dube, A., Jacobs, J., Naidu, S., and Suri, S. (2020), “Monopsony in online labor markets”, *American Economic Review: Insights*, 2(1): 33–46
- Falch, T. (2011), “Teacher mobility responses to wage changes: Evidence from a quasi-natural experiment”, *American Economic Review (Papers and Proceedings)*, 101(2): 460–465
- Falch, T. (2017), “Wages and recruitment: Evidence from external wage changes”, *ILR Review*, 70(2): 483–518

- Ganzer, A., Schmidlein, L., Stegmaier, J., and Wolter, S. (2020a), “Establishment History Panel (BHP) – Version 7518 v1”, Research Data Centre of the Federal Employment Agency (BA) at the Institute for Employment Research (IAB), doi: 10.5164/iab.bhp7518.de.en.v1
- Ganzer, A., Schmidlein, L., Stegmaier, J., and Wolter, S. (2020b), Establishment History Panel 1975–2018, FDZ-Datenreport 01/2020, Institute for Employment Research, Nuremberg
- Guimarães, P. (2004), “Understanding the multinomial–Poisson transformation”, *Stata Journal*, 4(3): 265–273
- Hirsch, B. (2016), “Gender wage discrimination: Does the extent of competition in labor markets explain why female workers are paid less than men?”, *IZA World of Labor*, 310
- Hirsch, B., and Jahn, E. J. (2015), “Is there monopsonistic discrimination against immigrants?”, *ILR Review*, 68(3): 501–528
- Hirsch, B., Jahn, E. J., Manning, A., and Oberfichtner, M. (2022), “The urban wage premium in imperfect labor markets”, *Journal of Human Resources*, 57(S): S111–S136
- Hirsch, B., Jahn, E. J., and Schnabel, C. (2018), “Do employers have more monopsony power in slack labor markets?”, *ILR Review*, 71(3): 676–704
- Hirsch, B., and Müller, S. (2020), “Firm wage premia, industrial relations, and rent sharing in Germany”, *ILR Review*, 73(5): 1119–1146
- Hirsch, B., Schank, T., and Schnabel, C. (2010), “Differences in labor supply to monopsonistic firms and the gender pay gap: An empirical analysis using linked employer–employee data from Germany”, *Journal of Labor Economics*, 28(2): 291–330
- Huber, S., and Rust, C. (2016), “Calculate travel time and distance with OpenStreetMap data using the Open Source Routing Machine (OSRM)”, *Stata Journal*, 16(2): 416–423
- IEB, Integrated Employment Biographies, Version V14.04.00-190927, Nuremberg, 2019
- IEB GEO, Geo-coded Address Histories of Persons, Establishments and Households in the IEB, Version V02.00.00-201804, Nuremberg, 2020
- Jacobebbinghaus, P., and Seth, S. (2007), “The German Integrated Employment Biographies Sample IEBS”, *Schmollers Jahrbuch*, 127(2): 335–342
- Kosfeld, R., and Werner, A. (2012), “Deutsche Arbeitsmarktregionen – Neuabgrenzung nach den Kreisgebietsreformen 2007–2012”, *Raumforschung und Raumordnung*, 70(1): 49–64
- Langella, M., and Manning, A. (2021), “Marshall lecture 2020: The measure of monopsony”, *Journal of the European Economic Association*, 19(6): 2929–2957
- Le Barbanchon, T., Rathelot, R., and Roulet, A. (2021), “Gender differences in job search: Trading off commute against wage”, *Quarterly Journal of Economics*, 136(1): 381–426
- Manning, A. (2003), *Monopsony in Motion: Imperfect Competition in Labor Markets*, Princeton, NJ: Princeton University Press
- Manning, A. (2011), “Imperfect competition in the labor market”, in O. Ashenfelter and D. Card (eds.), *Handbook of Labor Economics*, vol. 4B, pp. 973–1041, Amsterdam: Elsevier
- Manning, A. (2021), “Monopsony in labor markets: A review”, *ILR Review*, 74(1): 3–26
- Marinescu, I., and Wolthoff, R. (2020), “Opening the black box of the matching function: The power of words”, *Journal of Labor Economics*, 38(2): 535–568
- McFadden, D. L. (1974), “Conditional logit analysis of qualitative choice analysis”, in P. Zarembka (ed.), *Frontiers in Econometrics*, pp. 105–142, New York: Academic Press

- Ostermann, K., Eppelsheimer, J., Gläser, N., Haller, P., and Oertel, M. (2022), “Geodata in labor market research: Trends, potentials and perspectives”, *Journal for Labour Market Research*, 56: 5
- Ransom, M. R., and Oaxaca, R. L. (2010), “New market power models and sex differences in pay”, *Journal of Labor Economics*, 28(2): 267–289
- Robinson, J. (1933), *The Economics of Imperfect Competition*, London: Macmillan
- Schnabel, C. (2016), “United, yet apart? A note on persistent labour market differences between western and eastern Germany”, *Journal of Economics and Statistics*, 236(2): 157–179
- Sokolova, A., and Sorensen, T. (2021), “Monopsony in labor markets: A meta-analysis”, *ILR Review*, 74(1): 27–55
- Webber, D. A. (2022), “Labor market competition and employment adjustment over the business cycle”, *Journal of Human Resources*, 57(S): S87–S110

Figures




Figure 1: Labour market region of Hamburg (dark blue shaded) and adjacent districts (light blue shaded)




Figure 2: Detail of the labour market region of Hamburg (dark blue shaded) and adjacent districts (light blue shaded)




Figure 3: Fraction of recruits and the AKM plant wage effect

Note: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). The binned scatterplot plots the fraction of recruits hired by all plants with wages within a certain decile of the distribution of the AKM plant wage effects against the average AKM plant wage effect in that decile together with the regression line from a simple OLS regression.




Figure 4: Fraction of recruits hired from employment and the AKM plant wage effect

Note: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). The binned scatterplot plots the fraction of recruits from employment hired by all plants with wages within a certain decile of the distribution of the AKM plant wage effects against the average AKM plant wage effect in that decile together with the regression line from a simple OLS regression.




Figure 5: Fraction of recruits from non-employment and the AKM plant wage effect

Note: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). The binned scatterplot plots the fraction of recruits from non-employment hired by all plants with wages within a certain decile of the distribution of the AKM plant wage effects against the average AKM plant wage effect in that decile together with the regression line from a simple OLS regression.




Figure 6: Coefficients of the destination AKM plant effect and linear distance for each quartile of the distribution of the source AKM plant effects

Note: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from a fixed-effects Poisson regression on expanded data containing an observation for every recruit-plant pair where the dependent variable is an indicator for the plant where the recruit is hired. The regression includes all plant, industry, and plant location controls as detailed in the notes to Table 3 and all the interactions investigated in Table 5 where all dummy variables are centred around their mean.




Figure 7: Log predicted value for each quartile of the distribution of the AKM plant wage effects

Note: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from a fixed-effects Poisson regression on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired. The regression includes all plant, industry, and plant location controls as detailed in the notes to Table 3 and dummies for the quartiles of the distribution of the AKM plant wage effect as well as interactions between these and the AKM plant wage effect.

Tables

Table 1: Descriptive statistics on plant sample

	All plants		Hiring plants	
	Mean	SD	Mean	SD
Hiring plant	0.779	0.415	1.000	0.000
AKM plant wage effect	0.320	0.189	0.324	0.171
Share of male workers	0.603	0.255	0.600	0.252
Share of non-German workers	0.080	0.130	0.075	0.118
Share of workers aged under 30	0.227	0.161	0.238	0.160
Share of workers aged 30–44	0.363	0.139	0.365	0.134
Share of workers aged 45 or older	0.410	0.198	0.396	0.192
Share of low-skilled workers	0.115	0.107	0.115	0.104
Share of medium-skilled workers	0.658	0.218	0.658	0.215
Share of high-skilled workers	0.187	0.215	0.189	0.214
Share of workers on simple tasks	0.133	0.205	0.127	0.194
Share of workers on expert tasks	0.578	0.287	0.581	0.280
Share of workers on specialist tasks	0.162	0.197	0.164	0.194
Share of workers on complex tasks	0.127	0.181	0.128	0.178
Plants	10,134		7,895	

Notes: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Hiring plants are plants that hire at least one worker during our period of observation.

Table 2: Descriptive statistics on recruit sample (means)

	All recruits	Recruits from employment	Recruits from non-employment	Share of recruits from employment (%)
Male	0.605	0.639	0.564	57.3
Female	0.395	0.361	0.436	49.4
German national	0.926	0.939	0.912	54.9
Non-German national	0.074	0.061	0.088	45.1
Aged under 30	0.432	0.371	0.504	46.5
Aged 30–44	0.377	0.421	0.324	60.6
Aged 45 or older	0.192	0.208	0.172	58.8
Low-skilled	0.159	0.071	0.263	24.1
Medium-skilled	0.567	0.639	0.483	61.0
High-skilled	0.274	0.290	0.254	57.5
Simple tasks		0.095		
Expert tasks		0.553		
Specialist tasks		0.185		
Complex tasks		0.166		
Recruits	48,311	26,168	22,143	54.2

Notes: IEB and BHP, 2013/2014. Average linear distance refers to the average linear distance (in km) between the residence of the recruit and plant location (for all plants in the sample of plants summarised in Table 1). Information on task levels refers to the previous job and is thus only available for recruits from employment.

Table 3: Estimates of the wage elasticity of recruitment

Model	(1)	(2)	(3)	(4)
AKM plant wage effect	1.193 (0.023)	1.125 (0.023)	1.276 (0.027)	1.396 (0.029)
Linear distance		-0.102 (0.001)	-0.103 (0.001)	-0.103 (0.001)
Plant controls			✓	✓
Industry and plant location controls				✓
Recruits		48,311		
Plants		10,134		
Observations		489,583,674		

Notes: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from fixed-effects Poisson regressions on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired. Plant controls are the shares of male and non-German workers, the shares of workers aged under 30 and of workers aged 45 or older, the shares of low-skilled and high-skilled workers, as well as the shares of workers on simple, specialist, and complex tasks. Industry and plant location controls comprise 57 dummies for two-digit industry and four dummies for the districts where plants are located. Standard errors clustered at recruit level in parentheses.

Table 4: Estimates of the wage elasticity of recruitment from employment and recruitment from non-employment

Model	(1)	(2)	(3)	(4)
<i>Panel A: Recruits from employment</i>				
AKM plant wage effect	1.638 (0.031)	1.565 (0.031)	1.584 (0.036)	1.631 (0.038)
Linear distance		-0.099 (0.001)	-0.100 (0.001)	-0.100 (0.001)
Plant controls			✓	✓
Industry and plant location controls				✓
Recruits		26,168		
Plants		10,134		
Observations		265,186,512		
<i>Panel B: Recruits from non-employment</i>				
AKM plant wage effect	0.682 (0.034)	0.624 (0.034)	0.926 (0.040)	1.128 (0.044)
Linear distance		-0.106 (0.001)	-0.107 (0.001)	-0.107 (0.001)
Plant controls			✓	✓
Industry and plant location controls				✓
Recruits		22,143		
Plants		10,134		
Observations		224,397,162		

Notes: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from fixed-effects Poisson regressions on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired. Plant controls are the shares of male and non-German workers, the shares of workers aged under 30 and of workers aged 45 or older, the shares of low-skilled and high-skilled workers, as well as the shares of workers on simple, specialist, and complex tasks. Industry and plant location controls comprise 57 dummies for two-digit industry and four dummies for the districts where plants are located. Standard errors clustered at recruit level in parentheses.

Table 5: Heterogeneity in the wage elasticity of recruitment

	Gender	Nationality	Education	Age
AKM plant wage effect	1.237 (0.039)	1.524 (0.030)	0.889 (0.034)	1.739 (0.042)
× male	0.264 (0.047)			
× non-German national		-1.606 (0.083)		
× low-skilled			-0.291 (0.066)	
× high-skilled			1.983 (0.048)	
× aged under 30				-0.586 (0.051)
× aged 45 or older				-0.465 (0.067)
× non-employment				
Linear distance	-0.108 (0.001)	-0.103 (0.001)	-0.109 (0.001)	-0.100 (0.001)
× male	0.008 (0.001)			
× non-German		-0.013 (0.003)		
× low-skilled			-0.008 (0.002)	
× high-skilled			0.031 (0.002)	
× aged under 30				-0.008 (0.002)
× aged 45 or older				0.001 (0.002)
Recruits		48,311		
Plants		10,134		
Observations		489,583,674		

Notes: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from fixed-effects Poisson regressions on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired. All regressions include plant, industry, and plant location controls as detailed in the notes to Table 3. Standard errors clustered at recruit level in parentheses.

Table 6: Further heterogeneities in the recruitment elasticity and robustness checks

	Interaction with AKM worker wage effect	Hiring plants only	Recruit-specific choice set (three-week window)	Recruit-specific choice set (five-week window)	Plants with at least 30 workers only	Only workers from the local labour market of Hamburg	Only workers from the city of Hamburg
AKM plant wage effect	1.365 (0.026)	1.476 (0.035)	0.944 (0.041)	1.089 (0.045)	1.352 (0.040)	1.356 (0.032)	1.347 (0.039)
× AKM worker wage effect (centred)	1.281 (0.018)						
Linear distance	−0.102 (0.001)	−0.104 (0.001)	−0.104 (0.001)	−0.103 (0.001)	−0.101 (0.001)	−0.117 (0.001)	−0.097 (0.002)
× AKM worker wage effect (centred)	0.020 (0.001)						
Recruits	46,326	48,311	43,405	35,444	39,435	41,625	28,523
Plants	10,134	7,895	7,895	7,895	5,025	10,134	10,134
Observations	469,467,684	381,415,345	60,793,794	140,072,557	198,160,875	421,827,750	289,052,082

Notes: IEB and BHP, 2013/2014, and the AKM plant and worker wage effects provided by Bellmann et al. (2020). Estimates come from fixed-effects Poisson regressions on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired. In the specifications with recruit-specific choice sets, the expanded data set contains only those plants with some hiring activity in a three-week of five-week window around the start of the recruit’s job. All regressions include plant, industry, and plant location controls as detailed in the notes to Table 3. In the model that includes an interaction with the AKM worker wage effect, the latter is centred around its mean. Standard errors clustered at recruit level in parentheses.

Table 7: Robustness to the use of different distance measures

	Linear distance at hiring	Linear distance before hiring	Street distance	Commuting time (in minutes)	Controlling for staying in the same industry and occupation
AKM plant wage effect	1.631 (0.038)	1.631 (0.038)	1.647 (0.038)	1.647 (0.038)	1.424 (0.042)
Distance measure	-0.100 (0.001)	-0.100 (0.001)	-0.079 (0.001)	-0.094 (0.001)	-0.097 (0.001)
Plant belongs to the same industry as previous employer					2.311 (0.018)
Plant employs workers in the same occupation as previous job					2.923 (0.021)
Recruits			26,168		
Plants			10,134		
Observations			265,186,512		

Notes: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from fixed-effects Poisson regressions on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired. All regressions include plant, industry, and plant location controls as detailed in the notes to Table 3. Standard errors clustered at recruit level in parentheses.

Table 8: Estimates of the wage elasticity of recruitment using the German Vacancy Survey

Model	(1)	(2)	(3)
AKM plant wage effect	1.685 (0.208)	1.832 (0.214)	1.966 (0.156)
Plant controls		✓	✓
Industry and plant location controls			✓
Observations	9,269		

Notes: GVS, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from negative binomial regressions where the dependent variable is the number of hires in the last 12 months. Plant controls as detailed in the notes to Table 3. Industry and plant location controls comprise 62 dummies for two-digit industry and 401 dummies for the districts where plants are located. All models control for survey year effects. Robust standard errors in parentheses.

Table 9: Estimates of the wage elasticity of applications

Model	(1)	(2)	(3)	(4)
AKM plant wage effect	1.242 (0.136)	1.023 (0.133)	0.829 (0.128)	0.565 (0.357)
Occupation dummies		two-digit	three-digit	three-digit
Employer only hired one worker				✓
Observations	5,655			824

Notes: GVS, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from negative binomial regressions where the dependent variable is the number of applicants for the latest filled vacancy for a full-time position conditional on having hired at least one worker in the last twelve months. All models include the plant, industry, plant location, and survey year controls as detailed in the notes to Table 8. Column (4) uses only observations with exactly one hiring in the last twelve months. Robust standard errors in parentheses.

Table 10: Probability of hiring a woman

Model	(1)	(2)	(3)
Share of women among applicants	1.070 (0.028)	0.947 (0.051)	0.950 (0.051)
AKM plant wage effect			0.094 (0.077)
Three-digit occupation dummies		✓	✓
Observations	2,436		

Notes: GVS, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from OLS regressions where the dependent variable is an indicator whether a woman was hired. We use only observations where the share of female applicants is neither 0 nor 1. All models include the plant, industry, plant location, and survey year controls as detailed in the notes to Table 8. Robust standard errors in parentheses.

Appendix

Table A1: Heterogeneity in the recruitment elasticity from fully interacted models

	All recruits	Recruits from employment	Recruits from non-employment
AKM plant wage effect	1.394 (0.027)	1.612 (0.036)	1.163 (0.041)
× male	0.361 (0.044)	0.366 (0.058)	0.238 (0.065)
× non-German national	−1.552 (0.078)	−0.893 (0.122)	−1.737 (0.099)
× low-skilled	−0.138 (0.069)	−1.343 (0.108)	0.750 (0.091)
× high-skilled	1.974 (0.048)	1.483 (0.074)	2.160 (0.076)
× aged under 30	−0.171 (0.051)	−0.222 (0.067)	−0.034 (0.079)
× aged 45 or older	−0.251 (0.062)	−0.135 (0.077)	−0.304 (0.097)
× simple tasks		−0.199 (0.101)	
× specialist tasks		0.666 (0.079)	
× complex tasks		0.667 (0.091)	
Linear distance	−0.102 (0.001)	−0.099 (0.001)	−0.106 (0.001)
× male	0.008 (0.001)	0.008 (0.002)	0.009 (0.002)
× non-German	−0.014 (0.003)	−0.008 (0.005)	−0.011 (0.004)
× low-skilled	−0.006 (0.002)	−0.003 (0.004)	−0.007 (0.003)
× high-skilled	0.032 (0.002)	0.014 (0.002)	0.035 (0.003)
× aged under 30	−0.001 (0.002)	−0.005 (0.002)	0.006 (0.003)
× aged 45 or older	0.003 (0.002)	0.007 (0.002)	−0.005 (0.003)
× simple tasks		−0.033 (0.003)	
× specialist tasks		0.022 (0.002)	
× complex tasks		0.022 (0.003)	
Recruits	48,311	26,168	22,143
Plants	10,135	10,134	10,134
Observations	489,583,674	265,186,512	224,397,162

Notes: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from fixed-effects Poisson regressions on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired and all dummies are centred around their means. All regressions include plant, industry, and plant location controls as detailed in the notes to Table 3. Standard errors clustered at recruit level in parentheses.

Table A2: Heterogeneity in the wage elasticity of recruitment from employment

	Gender	Nationality	Education	Age	Tasks
AKM plant wage effect	1.403 (0.053)	1.707 (0.039)	1.178 (0.044)	2.023 (0.052)	1.132 (0.048)
× male	0.356 (0.062)				
× non-German national		-1.216 (0.132)			
× low-skilled			-1.602 (0.104)		
× high-skilled			1.894 (0.060)		
× aged under 30				-0.878 (0.067)	
× aged 45 or older				-0.327 (0.084)	
× simple tasks					-0.602 (0.100)
× specialist tasks					1.325 (0.076)
× complex tasks					1.788 (0.077)
Linear distance	-0.105 (0.002)	-0.099 (0.001)	-0.106 (0.001)	-0.097 (0.001)	-0.106 (0.001)
× male	0.008 (0.002)				
× non-German		-0.015 (0.005)			
× low-skilled			-0.015 (0.004)		
× high-skilled			0.029 (0.002)		
× aged under 30				-0.014 (0.002)	
× aged 45 or older				0.007 (0.002)	
× simple tasks					-0.033 (0.003)
× specialist tasks					0.028 (0.002)
× complex tasks					0.033 (0.003)
Recruits			26,168		
Plants			10,134		
Observations			265,186,512		

Notes: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from fixed-effects Poisson regressions on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired. All regressions include plant, industry, and plant location controls as detailed in the notes to Table 3. Standard errors clustered at recruit level in parentheses.

Table A3: Heterogeneity in the wage elasticity of recruitment from non-employment

	Gender	Nationality	Education	Age
AKM plant wage effect	1.119 (0.058)	1.302 (0.045)	0.392 (0.053)	1.244 (0.067)
× male	0.015 (0.069)			
× non-German national		-1.732 (0.105)		
× low-skilled			0.737 (0.085)	
× high-skilled			2.184 (0.076)	
× aged under 30				-0.024 (0.078)
× aged 45 or older				-0.588 (0.103)
Linear distance	-0.111 (0.002)	-0.106 (0.001)	-0.114 (0.001)	
× male	0.006 (0.002)			
× non-German		-0.010 (0.004)		
× low-skilled			-0.003 (0.002)	
× high-skilled			0.035 (0.003)	
× aged under 30				-0.001 (0.002)
× aged 45 or older				-0.007 (0.003)
Recruits		22,143		
Plants		10,134		
Observations		224,397,162		

Notes: IEB and BHP, 2013/2014, and the AKM plant wage effects provided by Bellmann et al. (2020). Estimates come from fixed-effects Poisson regressions on expanded data containing an observation for every recruit–plant pair where the dependent variable is an indicator for the plant where the recruit is hired. All regressions include plant, industry, and plant location controls as detailed in the notes to Table 3. Standard errors clustered at recruit level in parentheses.