Wages, Job Queues, and Skills

Ioana Marinescu and Ronald Wolthoff*

Preliminary and Incomplete

February 14, 2012

Abstract

We study the relationship between wages and the number and quality of applicants that a vacancy attracts. Using data from a large US employment website, we show that higher wages attract better applicants. Surprisingly, higher wages are associated with fewer applications, and this is robust to controlling for industry and occupation fixed effects. Only within specific job titles are higher wages associated with more applications. Our theoretical model shows that such a pattern is consistent with skills demanded by firms being highly job specific. The model has additional testable implications about rent sharing and unemployment rates by skill.

1 Introduction

The rise of employment websites over the past decade has made it relatively easy for job seekers to find and compare vacancies. Whereas in the past job ads were typically spread out over many different news papers, these days most vacancies can be found in a few mouse clicks and at zero marginal monetary cost. Background information on employers is also much easier to obtain than before. It seems intuitive that these developments must have reduced information frictions and have increased the level of competition in this market. In a homogeneous world, this competition would cause a firm that offers more attractive terms of employment than other firms to attract more applicants and to fill its vacancy more easily. Of course, reality is more complex. Workers are heterogeneous in skills and jobs differ in their skill requirements, so the relationship between the terms of employment that a firm offers and the likelihood that it fills its vacancy is not straightforward . For example, hospitals looking for brain surgeons may very well need longer to fill their vacancy than a local school searching for a janitor, even though they pay considerably more.

*Ioana Marinescu: Harris School of Public Policy, University of Chicago, ioana.marinescu@gmail.com; Ronald Wolthoff: Department of Economics, University of Toronto, ronald.p.wolthoff@gmail.com.

Understanding the relationship between skills, wages, and job queues (or number of applications) is important because it provides insights into various questions that are of interest to labor economists. For example, the relation between skills and wages determines the returns to investment in human capital. Further, the link between skills and job queues is informative about how unemployment varies with skill levels. Finally, the relation between wages and job queues is related to the level of competition within a market, with a strong positive impact of the wage on the number of applications implying a more competitive market.

Despite the relevance of these questions, the literature that studies the relationship between skills, wages, and job queues is limited in size. The main cause of this appears to be a lack of data containing all the necessary variables. This paper overcomes this problem by using data from CareerBuilder.com, the largest employment website in the US. This data set contains detailed information on available vacancies in two US cities in the beginning of 2011. It includes the wage that these vacancies pay, the number of applicants that each vacancy attracts, as well as various firm and applicant characteristics that can be used as measures of skills.

We use this data to document a number of new empirical facts. We first consider the economy-wide relation between skills, wages, and job queues and show that higher wage jobs get fewer applicants, but these applicants are higher quality. The same patterns emerge when we look within detailed industries and occupation categories. It is only when we consider a more stringent definition of a labor market that the results change. In particular, a positive relationship between the wage and the job queues is found when a market is defined by the job title chosen by firms when posting their vacancies. Specifically, a 10% increase in the wage is associated with a 7.5% increase in applicants. In the second part of the paper, we show that these facts are consistent with a simple directed search model where both firms and workers are heterogenous, and in which the degree of heterogeneity is large across labor markets but modest within a labor market. The model shows that the empirical patterns we uncover are consistent with a limited degree of skill transferability across labor markets with different skill levels. Given that a labor market is best defined empirically by a very specific job title, our empirical and theoretical results taken together imply that most labor markets are fairly thin. Additionally, the elasticity of applications with respect to wages within a labor market is less than one, suggesting that wage competition within a market is limited. The model generates additional empirically testable implications. In particular, the cost of creating vacancies increases less than proportionally with the skill level of the vacancy. Furthermore, high skilled workers are less likely to remain unemployed and capture a higher share of the surplus than low skilled workers.

Our findings add to the literature in various ways. First, our results on the relation between the wage and the number of applicants that a firm attracts are related to work by Holzer et al. (1991) and Faberman & Menzio (2010). Compared to these papers, we employ a data set that is larger and more representative of the entire labor market, instead of mostly focusing on low-skilled jobs. While there are some elements in previous work suggesting that high wage jobs attract fewer applicants or have shorter queues, we are to our knowledge the first to show that this relationship is very robust and is only reversed when controlling for the specific job title. Second, our paper is the first to our knowledge to document the relationship between wages and the quality of the applicant pool in a large and fairly representative sample of jobs. Third, our results shed light on the specificity of human capital, i.e. how transferable skills are from one market to another, adding to a line of research by e.g. Kambourov & Manovskii (2009). Our findings suggest that skills are not very transferable and labor markets are therefore relatively thin. Fourth, the theoretical model that we adopt extends the model introduced by Lang et al. (2005) to a setting with multiple submarkets. Interestingly, this extension allows us to better understand an empirically puzzling aspect of the results inLang et al. (2005), namely that less productive workers have lower unemployment durations. Indeed, while it is also the case in our model that less productive workers within a job title have lower unemployment duration, we find that workers in low skill labor markets have longer unemployment duration, consistent with the empirical evidence on this issue.

This paper proceeds as follows. Section 2 describes the data set and documents the key empirical facts. In section 3, we discuss a directed search model with two-sided heterogeneity that captures these facts. Section 4 concludes.

2 Empirical Analysis

In this section, we discuss the empirical part of our study. We start by describing the data in section 2.1, before presenting the results in section 2.2.

2.1 Data

We use proprietary data provided by CareerBuilder.com, the largest US employment website. Some background work was done to compare job vacancies in CareerBuilder.com with data on job vacancies in the representative JOLTS (Job Openings and Labor Turnover Survey). CareerBuilder.com vacancies represent a 35% of all vacancies in the US in January 2011. Some industries are overrepresented, in particular information technology, finance and insurance, and real estate, rental and leasing. The most underrepresented industries are state and local government, accommodation and food services, other services, and construction.

Our main data set contains all job vacancies posted on CareerBuilder.com in the Chicago and Washington DC Designated Market Areas (DMA) in January and February 2011. A DMA is a geographical region set up by the A.C. Nielsen Company. The counties that make up a city's television viewing area define a DMA. DMAs are slightly larger in size than Metropolitan Statistical Areas, and they include rural zones. For each job, we observe the following characteristics: the job title, the salary if specified, whether the salary is by hour or by year, the education required, the experience required, the name of the firm, the number of days the vacancy has been posted for. We normalize all salaries to be expressed in yearly amounts, assuming a full-time work schedule. When a salary range is provided, we take the middle of the interval. The job title is the title of the job posting, as freely chosen by the firm: this is something like "senior accountant". Because job titles are not normalized, there are many unique job titles. We did some basic cleaning to make job titles more comparable, the most important of which was to put every word in lower case and get rid of punctuation signs. We also determined that the first three words were the crucial ones in most cases, so we define a job title variable based on the first three words. Based on the full content of the job posting, an internal CareerBuilder algorithm assigns an SOC (Standard Occupational Classification) code to the job posting. Additionally, based on the firm's name, CareerBuilder uses external data sets like Dun & Bradstreet to retrieve the NAICS (North American Industry Classification System) industry code and the number of employees of the firm.

In terms of outcomes, we observe the number of times that a job was viewed as part of a list of jobs. When a worker searches for jobs, he is presented with a list of jobs. For the jobs that appear in the list, the job seeker can see the job title, salary, DMA and firm. To get more details, the job seeker must click on the job snippet, and we observe the number of times a job was clicked on. Finally, we observe the number of applications to each job. An application is defined as a person clicking on the "Apply Now" button in a job ad. Our key outcome of interest is applications per view, but we will also examine clicks per view. We have a second data set containing a random sample of jobs from the Chicago and Washington DC DMAs in January and March 2011. This data contains the same information as above, but additionally we have measures of applicant quality. Specifically, we have the share of applicants with a masters' degree, and the share of applicants with more than 10 years of experience.

Table 1 shows summary statistics. The average job ad receives almost 6 clicks and a bit more than one application per hundred views. The average yearly salary is $57,323; this number is somewhat higher than the US average wage in 2010 (see BLS Occupational Employment Statistics), which is consistent with the higher than average education of applicants on the website (see below). This number is obtained after we cleaned the data by removing the bottom and top 0.5% of salaries to eliminate outliers and errors. The average posting firm has about 19000 employees. Finally, on average 25% of applicants have a masters degree, and 50% of them have 11 years of experience or more.

2.2 Empirical Results

We start with examining the association between log wages and the number of applications per 100 views (table 2). In column I, without any controls, we find that there is a significant negative association between the wage and the number of applicants a vacancy gets: a 10% increase in the wage is associated with a 6.6% decline in applications per view. Subsequent columns add a number of job characteristics controls, including the required education and experience for the job. In column II, we also add industry and detailed occupation fixed effects (5 digits SOC codes). In column II, there are 595 occupation fixed effects. In principle, this should allow us to compare jobs that are very similar. Yet we still get a negative and significant association between the wage and the number of applicants. In column III, we add firm fixed effects instead of SOC fixed effects. The coefficient on the wage remains unchanged. The key lesson from columns I-III is that there is a strong negative correlation between the posted wage and the number of applications, even when controlling for large sets of observables. Remarkably, the magnitude of the coefficient on wages is fairly insensitive to the addition of controls, suggesting that the negative association is very robust. In column IV we control for job title fixed effects, where the job title is limited to the first three words. Interestingly, the coefficient on the wage is positive and significant and almost as large in absolute value as in column II when we controlled for SOC fixed effects.

The expected positive relationship between applications and wages only appears within job title (table 2, col. IV). The fact that the relationship between wages and the number of applicants switches sign when using job title fixed effects is not due to sample selection: indeed, as shown in table 7, a significant and negative relationship between wages and the number of applicants is also observed within the sample of job titles with at least two job postings, and when not controlling for job title fixed effects. By exploring job titles, we were able to determine two key reasons why job titles are more precise than detailed SOC codes. The first is that job titles indicate specialties. For example, registered nurse is SOC 29-1140, but the job title specifies if it's oncology or ICU. The second reason why job titles are more precise is that SOC codes do not do a very good job at representing hierarchy or experience level: for example, a job title may say "registered nurse supervisor" or "senior accountant", even though the first will still fall within SOC 29-1140 while the second will fall within the general SOC for accountants and auditors (13-2010). Finally, in column V, we add both job title and firm fixed effects, which should essentially absorb all of the firm-side heterogeneity in the data. Potential omitted variable bias could occur from a variable that is contained within the text of the job posting and that is systematically correlated with the wage but not absorbed by the job characteristic controls, nor the job title and firm fixed effects. Column V shows that within essentially identical jobs, higher wages indeed are associated with more applicants: the point estimate implies that a 10% increase in the wage is associated with a 7.5% increase in applicants. We summarize our findings as follows:

Empirical Result. *Across job titles, vacancies that offer higher wages receive fewer applications.*

Empirical Result. *Within a job title, vacancies that offer higher wages receive more applications.*

In table 3, we investigate the relationship between wages and the quality of applicants that a vacancy receives. We measure quality by the share of applicants with 11 years of experience or more, and the share of applicants with a master's degree. We find that higher wages are associated with a significantly higher share of high experience applicants, and the relationship is not sensitive to the addition of controls. A 10% increase in the wage is associated with a 1.3 percentage point increase in the share of high experience applicants, which represents a 2.6% increase in high experience applicants. Similarly, a higher wage increase the share of applicants with a master's degree (col. III). In this case, controlling for job title makes the coefficient insignificantly different from 0 (col. IV), but the coefficient is still positive, and not statistically significantly different from the estimate in column III. Overall, we conclude that higher wages attract higher quality applicants, and this relationship is robust to the addition of a large set of controls. We summarize our findings about the relationship between wages and applicant quality below:

Empirical Result. *Across job titles, vacancies that offer higher wages get higher quality applicants.*

Empirical Result. *Within a job title, vacancies that offer higher wages get higher quality applicants.*

One caveat to our analysis is that many jobs do not post wages, and so the relationship we estimate is based on a selected sample of jobs that do post a wage. One important reason for the wage not being posted is the use by many companies of Applicant Tracking Systems (ATS) software that keeps track of job postings and applications. These ATS also send out the job posting to online job boards such as CareerBuilder. Before sending out the job posting, ATS software typically strips down the wage information, even if it was available to start with. We believe the use of ATS is an important explanation for the absence of a posted wage because about two thirds of jobs are posted by ATS software and this proportion is similar to the proportion of jobs without a posted wage. To assess the extent of a potential selection bias, we examine whether jobs with a posted wage get more or better applicants than jobs without a posted wage. In table4, we examine the relationship between posting a wage and the number of applicants. We find that jobs with a posted wage get a larger number of applicants, but this relationship becomes insignificant when controlling for both job title and firm fixed effects (col. III). The point estimate in column III is also quantitatively small since it implies that jobs with a posted wage get 7.8% more applications per view. Since ATS use is typically determined at the firm level, and seems responsible for the non-posting of the wage, it makes sense that the impact of posting a wage is wiped out when firm fixed effects are controlled for.

In table 5, we examine the relationship between posting a wage and the quality of applicants in terms of education and experience. We do not find any relationship between posting a wage and the quality of applicants. Overall, we conclude that, as a group, jobs without a posted wage are not different from jobs with a posted wage once we condition on observables. We speculate that jobs without a posted wage are probably able to signal their salary through other means. Another possibility consistent with this pattern is that jobs without a wage offer a roughly average pay, and job candidates correctly infer this when they do not see a posted wage. In both cases, we think that the existence of many jobs without a posted wage is unlikely to bias our results about the relationship between the level of posted wages and the number and quality of applicants.

One remaining open question is whether the relationship between wages and the number of applications found in table 2 is contaminated by omitted variable bias. To better assess this, we turn to an examination of the impact of the wage on clicks per 100 views (table 6). Indeed, when a job is listed as a snippet, only the salary, job title, firm and DMA are listed. The applicant must click to see more details. As such, we have all the variables that can drive the applicant's click decision. When no controls are used (col. I), we see a significant and negative association between the wage and clicks per 100 views. When controlling for basic job characteristics (vacancy duration, dummy for salary expressed by hour, DMA and calendar month), industry dummies, and job title fixed effects, the coefficient on the wage becomes positive and highly significant, implying that a 10% increase in the wage is associated with a 2.8% increase in clicks per 100 views. The fact that the qualitative results in table 2 can be reproduced for clicks per view, an outcome whose job-level determinants are fully known, makes us more confident about our basic results. A higher wage is generally associated with fewer clicks and applications per view. It is only within job title that a higher wage results in more clicks and more applications per view.

3 Theory

In this section, we show that the patterns that we find in the data are qualitatively consistent with a simple directed search model with multiple submarkets, in which there is little worker heterogeneity within a submarket but a large amount of worker heterogeneity across submarkets. In the comparison with the empirical findings, we will assume that each submarket corresponds to a job title. Submarkets or job titles are defined by skill sets: workers differ in their skills and vacancies differ in the skills that they require. We will initially assume that a worker can only match with a firm that demands his exact skill set, i.e. a worker trained to be a nurse can only work in nursing jobs. In section 3.4, we will relax this assumption and discuss what happens if skills are (partially) transferable, i.e. a doctor might also be able to work as e.g. a nurse, although he will not necessarily produce the same amount of output as someone trained to do that job since the required skill sets are different.

3.1 Setting

Consider a static model of the labor market. The labor market is populated with a mass 1 of unemployed workers and a mass v of firms, determined by free entry. Each worker supplies one unit of indivisible labor and each firm has one job vacancy, which can be filled by at most one worker.[1]

Each workers can be characterized by a two-dimensional human capital vector $z_{ij} \equiv (x_i, y_j)$, which is exogenously determined when he enters the market place. The first component of this vector indicates his skill set, governing the job title in which he can work. To keep the model as simple as possible, we assume that x_i takes one of two values, x_1 and x_2, with $x_1 < x_2$.[2] The second component determines the worker's productivity within his job title: a worker with human capital z_{ij} working in job title x_i produces $x_i y_j$ units of output. Hence, output differences exist both within (determined by y_j) and across (determined by x_i) job titles. We distinguish two different productivity levels, y_1 and y_2 with $y_1 < y_2$, yielding a total of four different types of workers (e.g. bad nurses, good nurses, bad doctors, and good doctors).

The fraction of the workers with human capital z_{ij} is given by α_{ij}, which is assumed to be positive for all i and j. We assume that worker heterogeneity is smaller within a particular job title than across job titles, i.e. $1 < \frac{y_2}{y_1} < \frac{x_2}{x_1}$: it seems reasonable to assume that the difference in productivity between good and bad nurses is lower than between nurses and doctors. In most cases, we will focus on the limit $\frac{y_2}{y_1} \to 1$, i.e. on the case where the difference between high and low productivity workers within an job title is very small. Note that y_1, the productivity shifter for the low productivity worker, can be normalized to 1 without loss of generality.

Firms choose their vacancy's job title upon entry and this decision is irreversible. We index the firm's choice by k, i.e. x_k. Creating a vacancy in job title $x_k \in X$ costs the firm $c(x_k)$. We assume that $c(0) = 0$, $c(x) < xy_1$, $c'(x) > 0$ and $c'' \leq 0$. Hence, the cost of creating a vacancy with high skill requirements is higher than the cost of creating a vacancy with low skill requirements in absolute terms, but smaller as a fraction of the output being produced, i.e. $\frac{c(x)}{x}$ is decreasing in x.

In order to find a worker for their vacancy, firms post job ads when they enter the market. These job ads contain the vacancy's job title and the wage w that the firm promises to pay to the worker that it will hire. Firms commit to the posted wage (and the job title) for the remainder of the period. All job ads are posted in a central location (the employment website), where they can be observed by workers at zero cost. Hence, search in this economy is directed. After observing all job ads, workers decide to which job they wish to apply.

[1] Throughout this paper, the terms *firm*, *job* and *vacancy* are used interchangeably.

[2] It is straightforward to extend the model to a larger number of worker types.

Applying is a costly activity, which we capture here by allowing workers to apply only once.[3] As standard in the literature, workers and firms are assumed to use symmetric and anonymous strategies, introducing coordination frictions. If a firm receives applications from different types of workers, it will be able to identify and hire the most productive worker.

It is well known that in urn-ball models like the one presented here, the number of applications that a particular firm receives is a random variable following a Poisson distribution with mean equal to the queue length $\lambda(\cdot)$, the ratio of applications to vacancies at the combination of the job title and the wage chosen by the firm.[4] Hence, a firm facing a queue length λ matches with probability $m(\lambda) \equiv 1 - e^{-\lambda}$. Workers and firms who fail to match obtain a zero payoff.[5] A worker that matches with a firm earns the posted wage w, while the firm keeps the difference between the amount of output that is created and the wage that it pays.

3.2 Equilibrium

Consider a firm posting a vacancy $x_k \in \{x_1, x_2\}$.[6] Since we assumed that only workers with the required skill (i.e. $x_i = x_k$) can produce output in this job, workers of the wrong type ($x_i \neq x_k$) will not apply. Hence, each job title x_k operates as an independent labor market in which workers of type z_{k1} and z_{k2} compete for a job. If both types of workers apply to the same firm, the firm will give priority in hiring to a worker of type z_{k2}, since they generate a larger amount of output. Hence, this environment is very similar to the one discussed in Lang et al. (2005). Analogous to the analysis in their paper, we can show that in equilibrium these two worker types cannot be applying to the same vacancy. The following lemma formalizes this.

Lemma 1. *There exists no pooling equilibrium in which a firm posts a wage w in job title x_k and receives applications from both z_{k2}-workers and z_{k1}-workers.*

Instead, the two types must be separated. Some firms in job title x_k post high wages and only attract z_{k2}-workers, while other firms post low wages and only receive applications from workers of type z_{k1}. Let w_{kj} denote the wage attracting workers of type z_{kj}. In order to sustain such an equilibrium, two incentive compatibility constraints must be satisfied: the high productivity z_{k2}-workers must not want to apply to the jobs aimed at the low productivity z_{k1}-workers, and vice versa. One can show that the former condition binds, reducing the wage of low productivity workers to a level below what it would have been in a world without

[3] See Albrecht et al. (2006), Galenianos & Kircher (2009), Kircher (2009) and Wolthoff (2010) for models with multiple applications.

[4] See e.g. Burdett et al. (2001).

[5] A zero outside option for firms follows from the free entry assumption. Workers may have a positive outside option in a richer model due to e.g. unemployment benefits or the option value of search. The appendix presents such a model and shows that it generates the same predictions as the simpler model presented here.

[6] We will discuss firms' choice of which type of job to open below.

the high productivity z_{k2}-workers. The intuition is that a high productivity worker may be tempted to deviate and apply to a low wage job because they are more likely to be hired. Indeed, the firm always prefers to hire a high productivity applicant, so applying to the low productivity job is a sure way of getting hired. To make such a deviation unprofitable for high skilled workers, wages in jobs aimed at low skill workers must be low enough.

Lemma 2. *A unique equilibrium exists. In each job title x_k, two wages w_{k1} and w_{k2}, with $w_{k1} < w_{k2}$, are posted:*

$$w_{k1} = m'(\lambda_{k2})\, x_k y_2 \tag{1}$$

$$w_{k2} = \frac{\lambda_{k2} m'(\lambda_{k2})}{m(\lambda_{k2})} x_k y_2, \tag{2}$$

A firm posting w_{kj} attracts applications from workers of type z_{kj}. The queue lengths λ_{k1} and λ_{k2} are determined by the condition that a firm offering w_{kj} makes a profit π_{kj} equal to zero, i.e.

$$\pi_{k1} = m(\lambda_{k1})\, x_k \big(y_1 - m'(\lambda_{k2})\, y_2\big) - c(x_k) = 0$$
$$\pi_{k2} = \big(m(\lambda_{k2}) - \lambda_{k2} m'(\lambda_{k2})\big)\, x_k y_2 - c(x_k) = 0.$$

A worker of type z_{kj} obtains an expected payoff u_{kj} equal to

$$u_{k1} = \frac{m(\lambda_{k1})}{\lambda_{k1}} m'(\lambda_{k2})\, x_k y_2 \tag{3}$$

$$u_{k2} = m'(\lambda_{k2})\, x_k y_2. \tag{4}$$

3.3 Empirical Content

The simple model presented above provides several testable predictions regarding the relationship between productivity, wages, queue lengths and payoffs. In this subsection, we discuss a few of these predictions and show that they match the empirical facts obtained in the data. Proofs are again relegated to the appendix.

Focus first on a particular job title x_k. Two wages are being posted, a low wage w_{k1} attracting applications from z_{k1}-workers and a high wage w_{k2} attracting applications from z_{k2}-workers. This immediately yields the first prediction of the model.

Model Prediction. *Within an job title, vacancies that offer higher wages get higher quality applicants.*

Even though this prediction seems fairly intuitive, it is good to realize that not all models generate a positive relationship between wages and applicant quality. In particular, such a relationship is at odds with random search models or models in which there is only little correlation between the applicant's type and the probability that he gets hired. The latter situation can arise if unobserved or match-specific productivity components are the main determinant in hiring decisions, or if it is very costly for firms to screen workers.

The second prediction concerns the relationship between the wage and the number of applicants within an job title. The number of applicants that a firm attracts is ultimately determined by the free entry condition

$$\pi_{kj} = m\left(\lambda_{kj}\right)\left(x_k y_j - w_{kj}\right) - c\left(x_k\right) = 0.$$

Since the entry cost is the same for all firms in an job title, the equilibrium must satisfy $m\left(\lambda_{k1}\right)\left(x_k y_1 - w_{k1}\right) = m\left(\lambda_{k2}\right)\left(x_k y_2 - w_{k2}\right)$, i.e. a negative relationship between the firm's matching probability (or its number of applicants) and its payoff from a match. Firms attracting z_{k2}-workers create more output but also pay higher wages than firms with z_{k1}-workers. In the appendix, we show that the latter effect dominates if the difference between y_2 and y_1 is not too large, i.e. wages increase faster than output, such that $\lambda_{k1} < \lambda_{k2}$. In other words, since the difference between wages and output is lower in high productivity jobs, firms are making lower profit after hiring high productivity workers. So, to make firms indifferent between creating low and high productivity vacancies, it must be easier to fill high productivity vacancies, i.e. there are more applicants to the high productivity vacancy.

Model Prediction. *Within an job title, vacancies that offer higher wages receive more applications, provided that productivity difference between workers within an job title is sufficiently small.*

The assumption that the difference between y_2 and y_1 is small is important here. Together with the assumption that firms can perfectly screen the workers, this implies that is attractive for the high productivity z_{k2}-workers to consider a deviation to a w_{k1}-job. In order to sustain a separating equilibrium, w_{k1} needs to be reduced to a level which is lower than what it would have been in a world without z_{k2}-workers. Mathematically speaking, w_{k1} is determined by an incentive compatibility constraint rather than by first order conditions. If y_1 is much smaller than y_2, then the large wage difference between firms would prevent a z_{k2}-worker from considering a deviation in the first place. Consequently, w_{k1} would be determined by the first order conditions and the firm's match payoff would be proportional to y_j. Getting a higher match payoff, high-wage firms would need to fill their job with a lower probability than low-wage firms to insure the same expected profit across firms. Hence, in that case high-wage firms would receive fewer applications than low-wage firms, which is at odds with our empirical results.

Note that the literature provides several alternative explanations for a positive relationship between the wage and the queue length. Two examples are heterogeneity in workers' outside options (see e.g. Albrecht & Axell, 1984), and on-the-job search (Delacroix & Shi, 2006). Both mechanisms imply a smaller match payoff for firms paying higher wages, requiring those firms to match with larger probability in order to obtain equal profit. However, the positive relationship between wages and worker quality that we found in our data and model cannot be captured by the basic versions of those models since they assume homogeneity in worker productivity.

Next, we consider the relationship between wages, the number of applicants and applicant quality across job titles. When we compare wages w_{kj} across $k \in \{1, 2\}$ while focussing on a particular productivity type $j \in \{1, 2\}$, it follows immediately that higher-ranked job titles pay higher wages and attract applicants of higher quality. If we allow both k and j to vary, the result is less obvious, since it is possible that the low-type worker in the high job title earn a lower wage but produce more output than the high-type workers in the low job title, i.e. $w_{21} < w_{12}$ while $x_2 y_1 > x_1 y_2$. In the appendix, we show that generally the positive relationship survives.

Model Prediction. *Across job titles, vacancies that offer higher wages get higher quality applicants.*

Similar arguments play a role when considering the relationship between the wage and the number of applicants. From the equal-profit condition $\pi_{12} = \pi_{22} = 0$ and the concavity of $c(x)$, it is straightforward to show that w_{k2} and λ_{k2} are negatively related, i.e. that jobs in high skill job titles aimed at high productivity workers receive fewer applicants than jobs in low skill job titles aimed at high productivity workers. The main intuition is as follows: equation**??** reveals that the wage that the firm pays is proportional to the output being created in the match. If the fraction of output being paid to the worker were constant across job titles, the firm posting w_{22} would make a positive profit since entry costs increase less than proportionally to the output generated by the worker. To ensure equal profits across job titles, firms must pay a higher fraction of output to the workers in high-skill job titles, which is the case if the competition for workers is high and there are few applications in high skill job titles, i.e. λ_{22} low. With some algebra, a similar negative relationship can be established for w_{k1} and λ_{k1}. The overall relationship is not trivial because of the positive relationship between wages and applicants within an job title, but this effect is typically dominated, such that, across job titles, higher wages are associated with fewer applicants.

Model Prediction. *Across job titles, vacancies that offer higher wages receive fewer applications.*

Note that this result depends on the concavity of $c(x)$. If entry costs were to be convexly increasing, a counterfactual positive relationship between w_{kj} and λ_{kj} would be obtained.

3.4 Transferable Skills

In this section, we relax the assumption that a worker with a particular skill set x can only produce output in one type of job. Instead, we allow him to be productive in other jobs as well, although not necessarily equally productive as workers trained to do those jobs. We show that the equilibrium described above survives as long as the transferability of skills is not too high. Consider the following asymmetric specification of skill transferability. If a worker possesses less skill than required by the job ($x_i < x_k$), he will produce no output. On the other hand, if he possesses more skill than required by the job ($x_i > x_k$), he will be able to produce a positive amount of output, although the exact amount depends on how transferable skills are. We formalize these ideas in a parsimonious way by imposing the following structure:

$$f(z_{ij}, x_k) = \begin{cases} 0 & \text{if } x_i = x_1, x_k = x_2 \\ x_i y_j & \text{if } x_i = x_k \\ \tau x_i y_j & \text{if } x_i = x_2, x_k = x_1 \end{cases}$$

where the parameter $\tau \geq 0$ captures the extent to which a worker can transfer his skill to a different job title. If τ is small, skills are very job title-specific, while a large value of τ implies that skills are relatively more general. In words, this production function means that a worker trained to be a doctor is most productive performing medicine ($x_i = x_k = x_2$). A nurse will generally not be able to do his job ($x_i = x_1 < x_k = x_2$). Conversely however, a doctor might be able to work as a nurse ($x_i = x_2 > x_k = x_1$), although he will generally produce less output in that case than in his own job ($\tau < 1$). Given this structure, it is straightforward to show that the equilibrium described above survives when $\tau \leq \phi \equiv \frac{x_1}{x_2}$, i.e. the output of a doctor working as a nurse will not exceed the amount of output created by someone trained to do the nursing job.

Lemma 3. *Suppose $\tau \leq \phi$, i.e. skills are not very transferable. Then in equilibrium, workers of type z_{ij} maximize their payoff by applying to vacancies in job title x_i offering a wage w_{ij}.*

It is clear that if the transferability of skills is sufficiently large ($\tau > \phi$) workers would have an incentive to apply to a lower job title since they will be placed first in the queue there, and the characterized equilibrium breaks down. In that case, we either get a pooling equilibrium or an equilibrium determined by incentive compatibility constrains, similar to what we have within an job title. Such an equilibrium would not yield the negative relationship between wages and queue lengths that we observe in reality.

3.5 Discussion

Our model is able to generate the stylized facts we have uncovered in our empirical analysis. Namely, higher wage jobs attract better quality applicants. Across job titles, higher wage jobs attract fewer applicants, while within job title higher wage jobs attract more applicants. In order to generate these results, our search model relies on three key assumptions. First, the heterogeneity in worker productivity within job titles is smaller than the heterogeneity across job titles. Second, skills are not very transferable across jobs with different skill levels. Third, the cost of creating a higher skill vacancy increases less than proportionally with skill.

The fact that there is less heterogeneity in worker productivity within labor markets than across labor markets sounds like a tautology. And it is for good reason since we can define a labor market as a grouping such that heterogeneity in worker productivity within labor markets is lower than across labor markets. Taking this definition as a given, there is new information in our empirical analysis about what a labor market is. Indeed the labor market could have been a 5-digit SOC occupation, but we find that the labor market is more likely to be something like a job title, i.e. a labor market is typically very specific and smaller than a 5-digit SOC occupation. The assumption that skills are not very transferable across markets with different skill levels is a way to further flesh out the definition of a labor market by taking into account worker's freely chosen search strategies. Indeed, even if there is less heterogeneity in worker productivity within a job title than across job titles, workers may be willing to apply to multiple job titles. Low transferability insures that workers only apply to their own job title, so that job titles indeed constitute closed labor markets in the model. The fact that it is necessary to make these two assumptions about skill heterogeneity and skill transferability to yield correct empirical predictions has important implications for labor mobility and the cost of job loss. Our results suggest that switching to an even slightly different job likely has high costs, and that search frictions are likely to be very important in these relatively thin markets.

The assumption that the cost of creating a higher skill vacancy increases less than proportionally with skill is crucial in yielding a negative relationship between wages and job queues across job titles. This assumption implies that if it were equally easy to fill low and high skilled jobs, firms would find it more profitable to create high skilled jobs. Therefore, to make firms indifferent between creating low and high skilled jobs, it has to be the case that high skilled jobs are harder to fill, i.e. queues are shorter in high skilled jobs. This feature implies that high skilled workers have shorter unemployment duration and capture a higher share of the surplus. Both these implications are empirically testable, and it is in fact well known that unemployment duration decreases with education.

Our model does not explicitly model the demand for different skills. In reality, one may think that the nature of production is such that the demand for high skilled workers is high relative to supply. This would explain why jobs with high wages have shorter queues. However, we can reinterpret our model to be consistent with this explanation. In our model, each firm has a single job and high skilled workers are assumed to produce more physical output than low skilled workers. In a firm with multiple jobs, we could interpret xy as the marginal revenue product of a worker, i.e. the marginal contribution of this worker to the revenue of the firm. As the demand for high skilled workers increases relative to supply, the marginal revenue product of a high skilled worker should increase. This will incentivize firms to open more high skill vacancies, which will further diminish the number of applications received in equilibrium by high wage vacancies relative to low wage vacancies. This reinterpretation also allows for the possibility that the vacancy creation cost is not concavely increasing in physical output, but instead in the value of output.

4 Conclusion

In this paper, we analyze the relationship between skills, wages and job queues. We first document a number of new empirical facts. Using data from CareerBuilder.com, the largest US employment website, we show that firms that pay higher wages attract applicants with higher skill. The relationship between the wage that a firm posts and the number of applicants that it attracts crucially depends on the definition of a market. Economy-wide, the relation is negative and this continues to be the case if one controls for sector and/or occupation. However, when controlling for the job title as specified by the firm in the job ad, the sign of the relationship reverses: within a job title, higher wage jobs get more applicants.

We explain these patterns with a directed search model with two-sided heterogeneity and submarkets by skill. We argue that the combination of data and theory suggests that skills are very specific in the sense that they cannot easily be transferred from one job title to another. We further show that the empirical findings can be explained by a small degree of heterogeneity within a job title, but a large degree of heterogeneity across job titles.

Tables

	obs.	mean	s.d.	min	max
Master degree	2,282	0.273	0.255	0	1
High experience	2,379	0.506	0.265	0	1
Yearly wage	11,900	57,323	31,690	13,500	185,000
Applications per 100 views	61,051	1.168	2.570	0	100
Clicks per 100 views	60,979	5.640	5.578	0	100
Employees	61,135	18,824	59,280	1	2,100,000

Table 1: Summary statistics

	I	II	III	IV	V
Log yearly wage	-0.770***	-0.642***	-0.710***	0.593**	0.876***
	(0.052)	(0.075)	(0.087)	(0.302)	(0.283)
Controls	No	Yes	Yes	Yes	Yes
NAICS (2 digits)	No	Yes	No	Yes	Yes
SOC (5 digits)	No	Yes	Yes	No	No
Firm effects	No	No	Yes	No	Yes
Job title	No	No	No	Yes	Yes
Observations	11,708	11,708	11,708	11,708	11,708
R^2	0.017	0.133	0.363	0.480	

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Controls include a constant, vacancy duration, dummy for salary expressed per hour, required education and experience, log number of employees of posting firm, designated market area, and calendar month.

Table 2: Effect of log wage on the number of applications per 100 views.

	High Experience		Master Degree	
	I	II	I	II
Log yearly wage	0.136***	0.130*	0.173***	0.096
	(0.010)	(0.068)	(0.010)	(0.077)
Controls	No	Yes	No	Yes
NAICS (2 digits)	No	Yes	No	Yes
Job title	No	Yes	No	Yes
Observations	1,755	1,300	1,696	1,257
R^2	0.091	0.061	0.154	0.821

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. High experience = fraction of applicants with at least 11 years of experience; master degree = fraction of applicants with a master degree. Controls include a constant, vacancy duration, dummy for salary expressed per hour, required education and experience, log number of employees of posting firm, designated market area, and calendar month.

Table 3: Effect of log wage on the quality of applicants

	I	II	III
Posted a wage	0.478***	0.172**	0.091
	(0.031)	(0.081)	(0.071)
Controls	No	Yes	Yes
NAICS (2 digits)	No	Yes	Yes
Firm effects	No	No	Yes
Job title	No	Yes	Yes
Observations	61,051	61,050	61,050
R^2	0.006	0.498	

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Controls include a constant, vacancy duration, dummy for salary expressed per hour, required education and experience, log number of employees of posting firm, designated market area, and calendar month.

Table 4: Effect of wage posting on the number of applications per 100 views.

	High Experience		Master Degree	
	I	II	I	II
Posted a wage	0.006	-0.037	0.005	-0.016
	(0.013)	(0.072)	(0.013)	(0.072)
Controls	No	Yes	No	Yes
NAICS (2 digits)	No	Yes	No	Yes
Job title	No	Yes	No	Yes
Observations	2,379	1,774	2,282	1,704
R^2	0.000	0.790	0.000	0.830

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. High experience = fraction of applicants with at least 11 years of experience; master degree = fraction of applicants with a master degree. Controls include a constant, vacancy duration, dummy for salary expressed per hour, required education and experience, log number of employees of posting firm, designated market area, and calendar month.

Table 5: Effect of wage posting on the quality of applicants

	I	II
Log yearly wage	-1.045***	1.582***
	(0.089)	(0.375)
Controls	No	Yes
NAICS (2 digits)	No	Yes
Job title	No	Yes
Observations	11,694	11,694
R^2	0.01	0.568

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Controls include a constant, vacancy duration, dummy for salary expressed per hour, designated market area, and calendar month.

Table 6: Effect of log wage on the number of clicks per 100 views

	I	II	III
Log yearly wage	-0.702***	-0.579***	-0.530***
	(0.067)	(0.115)	(0.104)
Controls	No	Yes	Yes
NAICS (2 digits)	No	Yes	No
SOC (5 digits)	No	Yes	Yes
Firm effects	No	No	Yes
Job title	No	No	No
Observations	5,609	5,609	5,609
R^2	0.012	0.169	0.400

Robust standard errors in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Controls include a constant, vacancy duration, dummy for salary expressed per hour, required education and experience, log number of employees of posting firm, designated market area, and calendar month.

Table 7: Effect of log wage on the number of applications per 100 views.

A Proofs

A.1 Proof of Lemma 1

Proof. The proof is identical to the proof of proposition 3 in Lang et al. (2004).

□

A.2 Proof of Lemma 2

Proof. The proof is identical to the proof of proposition 5 and 6 in Lang et al. (2004).

□

A.3 Proof of Prediction 3.3

Proof. See the proof of lemma 2. □

A.4 Proof of Prediction 3.3

Proof. The proof is identical to the proof of proposition 7 in Lang et al. (2004). □

A.5 Proof of Prediction 3.3

Proof. A sufficient condition to have a positive correlation between wages and applicant quality is $w_{21} > w_{12}$. Substituting (1) and (2) and simplifying the result, this condition can be rewritten as

$$m'(\lambda_{22})\, x_2 > \frac{\lambda_{12} m'(\lambda_{12})}{m(\lambda_{12})} x_1 \tag{5}$$

where the queue lengths λ_{12} and λ_{22} follow from the zero-profit condition.

$$\pi_{k2} = \big(m(\lambda_{k2}) - \lambda_{k2} m'(\lambda_{k2})\big)\, x_k y_2 - c(x_k) = 0.$$

Take the right-hand side of (5) as given and consider how the left-hand side varies with x_2. Clearly, the condition is violated if $x_2 \to x_1$. Note further that the left-hand side is increasing in x_2.

$$\frac{d}{dx_2} m'(\lambda_{22})\, x_2 = m''(\lambda_{22}) \frac{d\lambda_{22}}{dx_2} x_2 + m'(\lambda_{22}) > 0$$

since both $m''(\lambda_{22}) < 0$ and $\frac{d\lambda_{22}}{dx_2} < 0$. If $x_2 \to \infty$, then $\lambda_{22} \to 0$, as follows from rewriting the zero-profit condition as $(m(\lambda_{k2}) - \lambda_{k2} m'(\lambda_{k2}))\, y_2 - \frac{c(x_k)}{x_k} = 0$ and recognizing that $\frac{c(x_2)}{x_2} \to 0$. In that case the left-hand side of (5) converges to x_2, implying that the condition holds. Hence for each x_1, there exists a $\widehat{x}_2(x_1)$ such that there is a positive correlation between wages and applicant quality for all $x_2 > \widehat{x}_2(x_1)$. □

A.6 Proof of Prediction 3.3

Proof. Sufficient conditions for a negative correlation between wages and queue lengths are $w_{21} > w_{12}$ and $\lambda_{11} > \lambda_{22}$. Note that if $x_2 \to \infty$, then $\lambda_{22} \to 0$ and the latter condition is satisfied. Moreover, $m'(\lambda_{22}) \to 1$, such that $w_{21} \to x_2 y_2$. Realizing that $x_2 y_2 > x_1 y_2 > \frac{\lambda_{12} m'(\lambda_{12})}{m(\lambda_{12})} x_1 y_2$ completes the proof. □

A.7 Proof of Lemma 3

Proof. It is straightforward to show that workers of type z_{1j} will not apply to jobs for type-$z_{2j'}$ workers since that yields them a zero payoff. Further, by incentive compatibility, a deviation from a z_{22}-worker to a z_{11}-job is dominated by a deviation to a z_{12}-job. Hence, we only need to consider deviations by z_{2j}-workers to z_{1j}-jobs. Consider the z_{22}-workers first. In equilibrium, they obtain a payoff $u_{22} = m'(\lambda_{22})\, x_2 y_2$. If the worker deviates and applies to a wage w_{12}, he will be treated in the same way as z_{12}-workers when $\tau = \phi$. Hence, the expected payoff of such an application is $u_{12} = m'(\lambda_{12})\, x_1 y_2$ as defined in (**??**). We prove that $u_{22} > u_{12}$ by contradiction. Note that since $x_2 > x_1$ and $m'(\cdot)$ is decreasing, $u_{12} = u_{22}$ requires that $\lambda_{12} < \lambda_{22}$. By the concavity of the entry cost, we have that $c_1 \geq \tau c_2$, such that $\lambda_{12} < \lambda_{22}$ implies that

$$
\begin{aligned}
\pi_{22} &= \big(m(\lambda_{22}) - \lambda_{22} m'(\lambda_{22})\big)\, x_2 y_2 - c_2 \\
&\geq \frac{1}{\tau}\left[\big(m(\lambda_{22}) - \lambda_{22} m'(\lambda_{22})\big)\, x_1 y_2 - c_1\right] \\
&> \tau\left[\big(m(\lambda_{12}) - \lambda_{12} m'(\lambda_{12})\big)\, x_1 y_2 - c_1\right] \\
&= 0.
\end{aligned}
$$

In words, a wage w_{22} which provides worker z_{22} with the same expected payoff as a wage w_{12}, generates positive profits for the firm offering w_{22}, which contradicts the fact that firm must earn zero profits in equilibrium. Hence, a profitable deviation exists. In equilibrium, a z_{i22}-worker must therefore apply to vacancies in job title x_2.

A similar approach can be used for z_{21}-type workers. Suppose that some z_{21}-workers apply to a vacancy in job title x_1 and earn an expected payoff $u_{11} = \frac{m(\lambda_{11})}{\lambda_{11}} u_{12}$. A deviating firm opening a vacancy w_{21} in job title x_2 geared towards these students would attract a queue λ_{21} defined by $\frac{m(\lambda_{11})}{\lambda_{11}} u_{12} = \frac{m(\lambda_{21})}{\lambda_{21}} u_{22}$. Since $u_{22} > u_{12}$, this implies $\lambda_{21} > \lambda_{11}$. The profit of the deviant would be

$$\begin{aligned}
\pi_{21} &= m(\lambda_{21})\, x_2 \big(y_1 - m'(\lambda_{22})\, y_2\big) - c_2 \\
&\geq \frac{1}{\tau}\big[m(\lambda_{21})\, x_1 \big(y_1 - m'(\lambda_{22})\, y_2\big) - c_1\big] \\
&> \frac{1}{\tau}\big[m(\lambda_{11})\, x_1 \big(y_1 - m'(\lambda_{12})\, y_2\big) - c_1\big] \\
&= 0,
\end{aligned}$$

which implies the desired contradiction. $\square$

References

Albrecht, J. W. & Axell, B. (1984). An equilibrium model of search unemployment. *Journal of Political Economy*, 92(5), 824–840.

Albrecht, J. W., Gautier, P. A., & Vroman, S. B. (2006). Equilibrium directed search with multiple applications. *Review of Economic Studies*, 73, 869–891.

Burdett, K., Shi, S., & Wright, R. (2001). Pricing and matching with frictions. *Journal of Political Economy*, 109, 1060–1085.

Delacroix, A. & Shi, S. (2006). Directed search on the job and the wage ladder. *International Economic Review*, 47(2), 651–699.

Faberman, J. R. & Menzio, G. (2010). Evidence on the relationship between recruiting and the starting wage. mimeo.

Galenianos, M. & Kircher, P. (2009). Directed search with multiple job applications. *Journal of Economic Theory*, 114, 445–471.

Holzer, H., Katz, L., & Krueger, A. (1991). Job queues and wages. *Quarterly Journal of Economics*, 106, 739–768.

Kambourov, G. & Manovskii, I. (2009). Occupational mobility and wage inequality. *Review of Economic Studies*, 76, 731–759.

Kircher, P. (2009). Efficiency of simultaneous search. *Journal of Political Economy*, 117, 861–913.

Lang, K., Manove, M., & Dickens, W. (2004). Racial discrimination in labor markets with posted wage offers. online version.

Lang, K., Manove, M., & Dickens, W. (2005). Racial discrimination in labor markets with posted wage offers. *American Economic Review*, 95(4), 1327–1340.

Wolthoff, R. P. (2010). Applications and interviews: A structural analysis of two-sided simultaneous search. mimeo.