

Who Gets Hired? The Importance of Finding an Open Slot

Edward P. Lazear
Kathryn L. Shaw
Christopher T. Stanton

Working Paper 16-128



Who Gets Hired? The Importance of Finding an Open Slot

Edward P. Lazear
Stanford University

Kathryn L. Shaw
Stanford University

Christopher T. Stanton
Harvard Business School

Working Paper 16-128

Copyright © 2016 by Edward P. Lazear, Kathryn L. Shaw, and Christopher T. Stanton

Working papers are in draft form. This working paper is distributed for purposes of comment and discussion only. It may not be reproduced without permission of the copyright holder. Copies of working papers are available from the author.

Who Gets Hired? The Importance of Finding an Open Slot

Edward P. Lazear*

Kathryn L. Shaw**

and

Christopher T. Stanton***

April 14, 2016

*Graduate School of Business and Hoover Institution, Stanford University

** Graduate School of Business, Stanford University

***Harvard Business School, Harvard University

The authors thank Peter Kuhn for his insightful comments, and participants at the University of Chicago, Copenhagen Business School, Stanford GSB.

Abstract

Despite seeming to be an important requirement for hiring, the concept of a slot is absent from virtually all of economics. Macroeconomic studies of vacancies and search come closest, but the implications of slot-based hiring for individual worker outcomes has not been analyzed in a market context. A model of hiring into slots is presented in which job assignment is based on comparative advantage. Crucially, and consistent with almost all realistic hiring contexts, being hired and assigned to a job depends not only on one's own skill but also on the skill of other applicants. The model has many implications, the most important of which are: First, bumping of applicants occurs when one job-seeker is slotted into a lower paying job or pushed into unemployment by another applicant who is more skilled. Second, less able workers are more likely to be unemployed because high ability workers are more flexible in what they can do. Third, vacancies are higher for difficult jobs because easy jobs can be filled by more workers. Fourth, some workers are overqualified for their jobs whereas others are underqualified. Miss-assigned workers earn less than they would have had they found an open slot in a job that more appropriately matches their skills. Despite that, overqualified workers earn more than the typical worker in that job. These implications are borne out using four different data sets that match the data requirements to test these points and others implied by the model.

Two workers, seemingly identical in qualifications, apply for a job. There is one opening so the employer must choose only one to hire. One worker is employed; the other continues looking. Luck has played a role in determining outcomes, presumably good luck for the one hired and bad for the one turned down.

Consider another scenario. Two similar, although not identical workers apply. The better of the two is offered the job. The slightly inferior worker is told that the position has been filled, but that another, lower paying job is still available. The worker accepts the position, fearing that the alternative in a poor labor market might be long-term unemployment. In this case, luck takes the form of job assignment, but one worker enjoys good luck while the other's luck is less favorable.

Neither of these situations is well-described by standard theory. Most production technologies are assumed to be smooth, with substitution across worker types and numbers being permitted. But at least some situations in the real world may be closer to a type of technology with some complementarities, where the notion of a job slot makes more sense. Job slots give rise to stochastic outcomes, where luck plays an important role. In markets, luck may take many forms, but the luck that is the focus of attention here is that which affects job offers. Central to the analysis is that a given applicant's luck depends on the others who apply for a job at the same time. A worker has no control over what others do, but the outcome of a job search process likely depends crucially on the applicants with whom the worker competes. Differentiating this view from more standard ones is that whether a worker obtains a job depends not only on a worker's own skills but on the skills of others. For example, in human capital theory, a worker's wage depends only on his own stock of human capital and its market price. Here, even workers with high levels of human capital may be underemployed and may receive low wages because an even better worker took the scarce job slot.

Whether luck in hiring is important in affecting lifetime worker wealth depends on the cost of mobility and on the thickness of markets. If bad luck in job assignment can be undone rapidly by subsequent job search, it may not be of major consequence. But in some markets, like that for academics, where hiring occurs only at scheduled times, luck in finding a job may have long term effects. The documented importance of cohort effects on wealth at least suggests that hiring luck may be of some consequence (Oyer 2008, Kahn 2010). At a minimum, it is important to understand the way in which the existence of slots and competition for jobs affects outcomes.

In what follows, a model of slots and within-firm job assignment is presented that yields specific, testable implications. Some help reconcile puzzles in empirical findings that are not understood in the context of standard economic theory. Specifically, the analysis below produces the following results, which are borne out in the empirical section.

1. Less able workers are more likely to be unemployed than the more able. Although this finding is well-established empirically, it is hardly obvious. Usually, markets for more homogeneous products are thicker than those for less homogeneous ones. Low ability workers are likely to be more similar to one another than are the high ability ones. This puzzle is explained by the fact that high ability workers are more flexible and can do a larger number of jobs.

2. Analogously, vacancy rates are highest in high-paying, high-quality jobs. Most workers can fill lower quality positions, but higher quality positions require high ability applicants who may not be present. This implication, coupled with the first, provides a cross-sectional Beveridge-curve. Friction in job search implies mismatch and here it takes a specific form that has been observed in data. Firms complain that they cannot find workers while at the same time, workers cannot find jobs. But the workers who have a difficult time finding jobs are not the ones suited to the jobs that are vacant.
3. Consistent with earlier search literature, unemployment, at least of the frictional variety, is a consequence of bad luck. Because firms are slot constrained, a worker remains unemployed when other applicants for a job have superior qualifications. But for those others applicants, the worker would have been employed.
4. It is common to speak about a person being over-qualified for a position, but what does that mean formally? A clear definition of over- and under-qualification is provided and that definition yields implications about observed wages for workers who find themselves over- or under-qualified for a job. Those who are in the wrong job receive wages below that which would be expected had they been lucky enough to find the appropriate job for their skills.
5. “Bumping” creates over-qualification. A worker’s probability of employment itself and the job to which he is assigned depends not only on that worker’s skill, but on the skill of competing applicants. Workers who are better suited to high level jobs are bumped into lower level ones or unemployment by even better qualified workers. Conversely, under-qualification results when an unfilled high-level job is filled by a worker who must take that job because there is no higher quality worker available to do it.

Model

The goal of the model is to capture the idea that luck is important in getting hired. A worker must encounter a firm that can make use of the worker’s skills, which depends on the qualifications of others who are employed by the firm. Key here is the notion of “slots” that is central and absent from most prior standard analyses.¹ A firm is not free to simply add workers to increase output. For example, a school might have a given number of classrooms and, if there is already one teacher per classroom, it may not be cost-effective to add another teacher to that room.²

¹Some of the earliest related work is Diamond (1982), which recognized that a worker’s ability to find a job depends on the number of others who are searching. Here, a slot is skill specific and a particular substitution technology is modeled.

²The matching literature that goes back to Gale and Shapley (1962) has the feature that slots are crucial and matching may not be perfect. The algorithms and results of this literature tend to abstract from market

The use of slots in the model is an innovation that will be shown to be crucial and helpful in understanding the existence of unemployment. Absent the concept of slots, it is difficult to generate unemployment in equilibrium without reverting to some kind of rigidity, the most obvious of which is sticky wages. Search theory uses a weak notion of slots implicitly because whether a worker locates a firm that wants that particular worker's skills is stochastic and based on the idiosyncratic aspects of both the firm and worker. But the level of abstraction in search theory is generally too high to generate the implications that are required for analyzing the detailed micro-data that will be used in this study. As a consequence, the notion of slots and how those slots relate to others already employed or also applying to a firm is explored.

A formal search theoretic literature by Shimer and Smith (2000) and Shimer (2005) comes closest to generating the implications that are provided by this model. In particular, implicit in Shimer and Smith is a matching model. Imagine heterogeneity among firms and workers such that each worker, in equilibrium, is optimally sorted to a particular firm. The equilibrium wage function generates wages conditional on worker types, where in a frictionless world, the firm that is the best match bids the most for a worker of every given type. However, if search is costly, then workers will accept jobs that are not their perfect match and firms will hire workers who are not their most preferred type. Implicit in this structure is a definition of over- and under-qualification. Workers who end up working at a firm that prefers lower quality, cheaper workers than themselves are overqualified and are less valuable there than they would be were they to find the optimal match. Those who end up working at a firm that prefers better, more expensive workers than themselves are underqualified and are less valuable there than they would be were they to find the optimal match.

Analogously, the vacancy literature essentially treats the job opening as a firm. In most structures, it is the overall flow of vacancies that determines hiring rates and if, because of stochastic search, more than one worker arrives at a vacancy, then some of the applicants are necessarily unemployed. This (typified by Shimer (2005)) is a close predecessor to the slot structure used below, but the one described here is more specialized and therefore more concrete in its set of implications. The strength of the model is that the implications are borne out in all the data examined. The search technology modeled, however, is in the spirit of Shimer (2005), but is modified to allow two workers to be employed by any given firm.

Explicitly modeling the structure of slots contrasts with the vast literature on supply-side unemployment, where business cycles induce workers to stay home because their reservation values exceed their productivity. These prior models are consistent with some real-world observations, but not all. For example, one of two otherwise identical workers, both in terms of productivity and alternative uses of time, may be offered a job while the other is forced into unemployment. This cannot be explained in standard technologies with smooth production. The notion of slots is particularly helpful here and conforms to standard intuition about job finding. Once a job is already filled, another equally qualified applicant will not be offered employment.

considerations and do not provide the specific implications that are the focus of this paper. Closest to imbedding markets into a matching framework is Becker (1993), which considers matching in marriage markets.

It is also desirable that the model does not create unemployment by assumption. Thus, no worker is inherently unemployable in the sense of having ability so low that he can never add positive value. If a worker does not obtain a job, it is because he has encountered bad luck that precludes productive employment because of the workforce composition, not because he is so unproductive that no firm will employ him.

Another goal of the model is that the worker's wage and standard of living be affected by luck that takes the form of being hired or not and on the job assignment if hired. A worker may be overqualified in the sense of being more productive in another job, were it available, but may take the one offered because the better job is already filled.

Production

The production function has both slot features and complementarity. Smooth production functions give no role to slots. In standard theory, labor utilization is a continuous variable and, despite diminishing marginal productivity, everything occurs smoothly. This is at odds with what is frequently observed in the real world where positions are discrete and having an open slot is necessary before a worker can be hired.

The assumption of complementarity creates slots, but also a reason for having more than one worker in a firm. For example, there might be an advantage to having one firm serve the same client. Customers may prefer that the same outsourcer provide both programming and administrative services so that if there is a problem, one outsourcer cannot blame the other for the difficulties encountered.

A useful benchmark against which to compare results of this model is a world where slots do not exist in the sense that a firm can always hire a worker and the resulting effect on output is smooth. In the standard production function, output is

$$(1) \quad Q = Q(q_1, q_2, \dots, q_n),$$

with $Q_i > 0$, $Q_{ii} < 0$, $\forall i$. Each i can be thought as a worker type and the q reflects the amount of labor of type i that is used. Production is smooth in the sense that increasing the input of labor of type i increases the output of the firm, albeit at declining rates.

Here, a different technology is assumed. Each firm can use at most one worker for each job, which for simplicity is assumed to equal two. Again, this is closer to the vacancy literature discussed above, where a vacancy defines a unique slot and the arrival of more than one worker does not result in more output than can be produced by filling that vacancy.

The comparison, then, is not against what would occur in a frictionless world, but rather in a world where any number of workers can be accommodated by the firm, just at decreasing marginal product. Unlike smooth production in (1), in a two-slot case, output at the firm takes the form

$$Q(A_1, A_2) = Z(A_1, A_2) [q_1 + q_2] + [1 - Z(A_1, A_2)] q_1$$

where A_1 is the ability of the worker assigned to job 1, A_2 is the ability of the worker assigned to job 2, and $Z(A_1, A_2)$ equals one if it pays to hire workers into both slots and zero if it pays to hire a worker into only one of the two slots. $Z(A_1, A_2)$ is derived optimally below (in lemma 2).

The point here is that even if multiple workers of identical ability apply to the firm, only one of them can be placed in each job. Unlike the production function in (1), more labor cannot be applied to job 1. The number of heads, as opposed to hours, are the determining factor. A slot structure may arise when capital is discrete as in the case of a single position at the control panel in an automated steel plant. If the slot is unfilled, then q_i is defined to be equal to zero.

To make things concrete, suppose that $Q[q_1(A_1), q_2(A_2)]$ is derived from underlying production

$$(2) \quad \begin{array}{l} \text{a.} \quad q_1 = (\gamma + \delta A_1) \\ \text{b.} \quad q_2 = (\alpha + \beta A_2) \end{array}$$

Think of job 1 as the difficult job and job 2 as the easy job. In the oDesk data examined below, job 1, the difficult job, is programming and job 2, the easy job, is administrative support. High ability workers are better in every job, but they have a comparative advantage in job 1. Thus, let $\alpha > \gamma$ and $\delta > \beta$ as shown in figure 1. High ability workers produce more than low ability workers in each of the two jobs because both β and δ are positive, but ability has a greater effect in augmenting output in the difficult job than in the easy job.

Let the worker's reservation value be K , thought of as the value of leisure. The value of a worker's output is then Rq_i where R is the price of the product. Define A_0 as that ability such that a worker would have a value of output in the difficult job that just equals the value of the alternative, K . Using (2a),

$$A_0 = (K/R - \gamma) / \delta.$$

Further, define A^* as the ability such that the worker is equally productive in both jobs (as shown in figure 1), so that $\gamma + \delta A^* = \alpha + \beta A^*$. Thus,

$$A^* = (\alpha - \gamma) / (\delta - \beta).$$

Then,

$$(3) \quad R(\gamma + \delta A) > K \text{ for } A > A_0.$$

Any worker with $A < A_0$ would never be hired into the difficult job because his ability is so low that his output would be below the value of not working.

Also, in keeping with the desire to avoid assuming that no worker is inherently unemployable, at least in normal times, the minimum ability in the population A_{\min} is assumed to be sufficiently high so that

$$(4) \quad R(\alpha + \beta A_{\min}) > K.$$

Finally, to complete the model, assume a single period and that costly search takes the form of allowing each worker to search at one and only one firm. If there are N workers and M firms, each of which is equally likely to be searched by any worker, then the number of workers who arrive at any given firm follows a binomial distribution. The probability of a given worker arriving at any firm is then $1/M$. If h is the number of workers that arrive at a firm, then the density of arrivals is

$$(5) \quad p(h) = \frac{N!}{h!(N-h)!} \left(\frac{1}{M}\right)^h \left(1 - \frac{1}{M}\right)^{N-h}$$

where $p(h)$ is the proportion of firms that have h applicants, $h=0,1,\dots,N$. Although N is exogenous (perfectly inelastic labor supply is assumed for simplicity), M is derived as part of a competitive equilibrium, described in a later section.

Firms that have zero applicants are uninteresting for the analysis. Those that have one applicant hire that applicant and place him in the job to which his output is highest based on the production function in (2a,b). The action all comes from firms that receive two or more applicants because it is only those firms that ever leave workers unemployed.

Because (2a,b) implies that output and profit are both (weakly) increasing in worker ability, the firm will always choose to hire at most the two highest ability workers from among the pool of applicants. Thus, define A_1, A_2 as the ability of the best and second best applicant, respectively. If more than two workers apply to a firm, then some applicants necessarily are unemployed in the same way that unemployment occurs in the Shimer model. More interesting, however, is the assignment problem among the two best workers and the situations where firms choose optimally to use only one worker, even though at least two apply. This is the essence of the slot-based assignment model and the two lemmas and proposition derived below describe the full array of outcomes for the economy.

Costly search creates bilateral monopoly (the worker has at most one offer and the firm sees two and only two workers) so there is a need to allocate the rents. Although the structure is competitive in the sense that there are many firms and many workers, once pairing occurs, bargaining opportunities exist. The bargaining game is not crucial to the model as long as it results in some rent splitting parameter that is common across firms. Since ex post wage determination is not the focus of this analysis, firms are assumed to commit to a wage schedule

when they advertise a job. Firms advertise wage schedules that will later be shown to be consistent with competitive equilibrium.⁷ Thus, wages in the two jobs are given

$$(6) \quad \begin{aligned} \text{a.} \quad & w_1(A) = [\lambda (\gamma + \delta A)] R \\ \text{b.} \quad & w_2(A) = [\lambda (\alpha + \beta A)] R \end{aligned}$$

with $0 \leq \lambda \leq 1$.

It is now possible to derive the implications of the model. First, firms that receive two applicants are no different from those that receive more than two, because the firm discards all but the two highest ability applicants. Thus, the following lemmas relate to firms with two or more applicants.

Lemma 1: If a firm assigns any worker to the difficult job, it will always be the highest ability worker, that is, the worker with ability A_1 .

Proof:

The better worker should be assigned to the difficult job at a given firm and the poorer worker to the easy job if

$$\gamma + \delta A_1 + \alpha + \beta A_2 > \gamma + \delta A_2 + \alpha + \beta A_1$$

or if

$$(\delta - \beta) (A_1 - A_2) > 0$$

which must hold because $\delta > \beta$ and $A_1 > A_2$.

Lemma 1 merely says that workers should be assigned according to comparative advantage. The worker of ability A_1 is better in every job, but has a comparative advantage in the difficult job. Another result is useful.

Lemma 2: Both slots are filled if and only if

$$A_2 > [(\beta - \delta) A_1 - \gamma] / \beta$$

Proof: The choice is between hiring two workers or hiring only the best worker and assigning her to the easy job. If the highest ability worker is best assigned to the difficult job, then there is always gain to hiring the low ability worker into the easy job by (4). Thus, two slots are filled if and only if

$$\alpha + \beta A_1 < \gamma + \delta A_1 + \alpha + \beta A_2$$

or iff

$$A_2 > [(\beta - \delta) A_1 - \gamma] / \beta .$$

|||

Recall that $Z(A_1, A_2)$ is an indicator variable that is 1 when a firm fills both slots and zero when the firm fills only the low ability slot. Because all workers have positive output in the easy

⁷ Reneging on the wage offer is possible, but is thought to be analogous to refusing to pay a worker at the end of a pay period after the work has already been done.

job, a firm that has at least two applicants will surely fill the easy job if it fills the difficult job. This is merely a direct outcome of Lemma 2. Thus, directly from lemma 2,

$$Z(A_1, A_2) = 1 \text{ if } A_2 > [(\beta - \delta) A_1 - \gamma] / \beta$$

$$= 0 \text{ otherwise .}$$

It is now possible to state a proposition that provides necessary and sufficient conditions for full employment in the economy.

Proposition: Unemployment occurs if for at least one firm, $A_2 < [(\beta - \delta) A_1 - \gamma] / \beta$ or if at least one firm sees more than two applicants.

Proof:

The second part of the statement is trivial. If a firm can only employ two workers and it receives more than two applications, some of the applicants are denied employment.

The first part of the statement follows directly from lemma 2, which states that one and only one worker is employed when $A_2 < [(\beta - \delta) A_1 - \gamma] / \beta$. The other worker is then unemployed. |||

There are a number of points that come out of this slot-based structure. First, “bumping” occurs. If two high ability workers, defined as having $A > A^*$, show up at the firm, then the lowest ability of the two is bumped down to the easy job, even though he is inherently more productive in the difficult job than in the easy job. The worker must do the easy job, not because he is low ability, but because the difficult job is best assigned to the even-higher ability worker. The worker who is bumped into the easy job earns less than he would have had he been able to secure a difficult job. Similarly, when a firm’s two best applicants are low ability workers, defined as having $A < A_0$, then the lowest ability worker is bumped out of a job altogether and ends up being unemployed despite the fact that the firm has two job openings. In this situation, because one job is too difficult for either of the workers to do successfully, it goes unfilled.

Second, low ability workers are more likely to be unemployed than high ability ones. Again, part of this is trivial. Given that profits are increasing in worker ability in (2a,b), a firm that receives more than two applicants discards all but the two best workers. But even firms that receive only two applicants may turn one down; lemma 2 and the proposition describe the conditions under which this occurs. Bad luck for low ability workers takes the form of applying to a firm where at least one applicant is of higher ability but no applicant has sufficiently high ability that the firm wants to employ multiple workers. Because no applicant can do the difficult job, only one worker is hired and other applicants are unemployed. For very low ability workers, good luck means either that other applicants are of even lower ability or that there is only one higher ability applicant and that worker has sufficiently high ability to induce the firm to employ both workers (meaning that the high ability worker has ability greater than $\gamma/(\beta - \delta) + [\beta/(\beta - \delta)]A_2$, which is the condition in lemma 2, rewritten).

The second implication of this model, that low ability employees suffer more unemployment, is not an obvious one, and is in some respects counterintuitive. The market for

high ability workers might be thought to be thinner than that for low ability workers, just as the market for mansions is thinner than for low priced development houses.⁴ The time on the market for more idiosyncratic goods and services is generally expected to be longer, not shorter, than those for homogenous ones. But high ability workers are not idiosyncratic. The ability to work in both jobs, as opposed to only one, makes them more employable. The empirical implication is that workers with low levels of education have higher unemployment rates. This is Hypothesis 1 below. While it is generally well-known that well-educated workers suffer less unemployment, it is useful to see the magnitudes of these rates and to understand a theoretical logic that is consistent with this observation.

Third, it is possible, although highly unlikely, that the application process is such that no unemployment results. There is nothing in the model that guarantees unemployment. Unemployment is not assumed; it occurs only when there is some bad luck in the world. Under the right distribution of applicants across firms, no unemployment occurs. The unemployment describes by this model is of the “frictional” variety, which can result even in very tight labor markets.

Fourth, the jobs dominated by low ability workers have the lowest vacancy rates. There are never unfilled easy jobs; only difficult jobs sometimes go unfilled. This Beveridge-curve type result (low vacancies with high unemployment) is testable and consistent with occupational difference in mismatch between vacancies and unemployment found in Lazear and Spletzer (2012).⁶ Using education as an observable measure of ability, the empirical implication in Hypothesis 2 below is that vacancy rates rise with education.

Fifth, workers may be over-qualified or under-qualified for jobs, but are still profitably employed in those jobs. Recall that A^* is defined as that ability level such that the worker is equally productive in both jobs, given before as $A^* = (\alpha - \gamma) / (\delta - \beta)$. Workers for whom $A > A^*$ prefer to be assigned to the difficult job, which happens if the worker in question is the highest ability worker of the two who apply. But it is possible, even for a worker whose $A > A^*$, to be the lower ability worker of the two best to apply to the firm, in which case he will be forced to do the easy job. He is “over-qualified,” but still successful in the sense that he is more valuable in that task than in taking leisure. Good luck for a high ability worker (whose $A > A^*$) consists of being paired with a workmate who is of lower ability because the difficult job, which yields higher wages for those whose $A > A^*$, goes to the highest ability worker.

In the over-qualified job, the worker will be underpaid relative to his expected pay on the higher skilled job. To see this, use the wage model of equation (6), but simplify to compare wages across ability levels by setting $R=1$ and $\lambda=1$. When $A > A^*$ and the worker is in the over-qualified job, then the actual wage will be w_2 but the worker would have earned a predicted wage of w_1 if in the job for which he is best qualified. The wage gap defined as the predicted minus the actual wage, $\hat{W} - W$, is then equal to $(\gamma + \delta A) - (\alpha + \beta A)$. As is evident in Figure 1, for $A > A^*$, the value of the wage gap is positive for the overqualified who work in job 2.

⁴See Lazear (1986) for an analysis of pricing, time on the market and inventories in thick and thin markets.

⁶There are chronic “shortages” of workers to fill professional jobs.

Although the over-qualified worker earns less than she would were she in the difficult job, she earns more than the typical worker in the easy job because her ability is high relative to those workers and because productivity increases in ability in both jobs. This is a testable implication once an empirical definition of over- and under-qualified is established.

Consider next workers who are under-qualified for jobs, but profitably employed in those jobs. If both applicants to a firm have ability below A^* , but the condition of Lemma 2, namely, $A_1 > \gamma/(\beta-\delta) + [\beta/(\beta-\delta)]A_2$, is satisfied, the highest ability worker is assigned to the difficult job. But because his ability is below A^* , his absolute output in the easy job would be higher. In that sense, he is under-qualified for the job, producing low output there relative to what he would have produced in the easy occupation. Good luck in that case consists of being paired with a workmate whose ability is even greater because the higher ability worker is assigned to the difficult job where his output and wage is lower than were he assigned to the easy job. In this case, being the higher ability worker is bad luck.

The wage gap for under-qualified workers who are in the difficult job but have ability $A < A^*$ remains $\hat{W} - W$, but now is given by $(\alpha + \beta A) - (\gamma + \delta A)$. The appropriate job for under-qualified workers is the easy job, whereas the actual job is the difficult job. As before and, in this case, perhaps counterintuitively, the wage gap is predicted to be positive because these workers would be earning more in the easy job for which they are better suited than in the difficult one into which they are thrust. Were wages attached to jobs, rather than to workers, a worker who was under-qualified would receive more in that job than in his appropriate job. Thus, the model's prediction, if borne out (as it is), is certainly not an obvious one. Thus, for both over-qualified and under-qualified workers, the wage gap is predicted to be positive, which is stated as Hypothesis 3 below.

There is one difference in wage predictions between under- and over-qualified workers. Unlike the over-qualified worker, the under-qualified worker earns less than the typical worker in the difficult job. Again, because productivity increases in ability in both jobs and because his ability is lower than that of the typical worker in the difficult job, his wages are expected to be lower than average for that job. This is also testable.

A general statement is that good luck consists of applying to a firm where the other applicant's ability permits the worker to be assigned to the job in which he has a comparative advantage were there open slots in each job. The assignment to jobs, existence of unemployment, wages and profits all depend on the distribution of talents in the population, the number of slots of each type and on luck that takes the form of the distribution of applicants across firms.

Sixth, it is quite possible that firms will complain about not being able to find qualified workers, while workers simultaneously complain about not being able to find a job. This is a cross-market Beveridge-like result. When vacancies are high among the difficult jobs, then unemployment is more prevalent among the low ability workers. Vacancies occur in the difficult jobs only when neither of the workers who arrives at the firm has sufficiently high ability to satisfy the employ-two conditions of lemma 2. When that occurs, the highest ability applicant bumps the other applicants out of their easy job and into unemployment, creating higher

unemployment rates among the less skilled workers. This is the issue of mismatch. Programmer jobs go unfilled when all of a firm's applicants have ability that is too low, again as described by Lemma 2⁸

Seventh, over time, the variance of income rises with education. One well known stylized fact is that there has been a rising return to education over time. Such time series implications can be captured by shifts of skill gradients displayed in Figure 1. In Figure 2, the skill gradients shift upward, but the upward shift is greater for the difficult job than for the easy job. Output q_1 in the difficult job is today higher at higher ability levels than it was in the past – δ has shifted up to δ' . Output q_2 in the easy job has also shifted up as a function of ability, with the slope of the line rising from β to β' . By assumption, the rising return to skills over time has resulted in a greater increase in δ than in β . The assumption is consistent with the notion that technological progress is more complementary with skill level (skill biased) in difficult jobs than it is in easy jobs. The idea is that technology has increased the difference between the output of the more able farmer and the less able one, but it has increased the difference between the output of the more able engineer and the less able one by even more.

The implication is that the skilled worker who is in the easy job will get a bigger pay reduction today than he would have in the past. This outcome is not due merely to the rising return to human capital. It results because the gap between productivity in the job for which a worker is appropriately qualified and the one for which he is over-qualified has grown over time. It is the interaction between skill-biased technical change and the slot allocation that comes out of this personnel economics model that generates the result. This is also in keeping with a rising variance of wages over time. Among the highly able, there will be some workers in the difficult job and there will be some in the easy job. The pay gap between these two jobs has risen over time, so the variance of pay has risen over time.

Competitive Equilibrium and Endogenous Slots

It is important to show that the number of firms and the rent distribution parameter, λ , can be derived from a competitive equilibrium and that the number of firms is consistent with the search technology assumed, namely that each worker searches at one and only one firm. Allowing free entry among firms, given perfectly inelastic labor supply, produces a zero-profit competitive equilibrium. Intuitively, the margin of adjustment is the stochastic arrival of workers. Were positive profits available, more firms would enter, which would lower the probability of an applicant showing up at any given firm. This reduces the highest and second-highest order statistic among those firms that obtain at least two applicants and also increases the proportion of firms that obtain zero or one applicant.⁹

Given the technology, adjustments in product markets come about through variations in the number of firms, rather than in output per firm. Since firms cannot adjust on the intensive

⁸ Again, see Shimer and Smith (2000) for earlier work that derives this implication.

⁹ This is a direct property of the binomial distribution that results from this search technology.

margin by hiring more labor, the marginal cost of output is the cost per unit that results from adding another firm, i.e., the average cost.

Because each of N workers can search one and only one firm, the distribution of arrivals at the various M firms is given by a binomial distribution with $1/M$ being the probability of any given applicant arriving at a specific firm. Define $P(k; 1/M, N)$ as the binomial density that any given firm sees k applicants from the population of N total applicants.¹⁰

Since the firm keeps at most two applicants and only the best two are relevant, the distributions of the highest and second-highest order statistics are inputs into the firm's choice. In general, the density of the r^{th} order statistic from k arrivals is given by

$$(7) \quad f_{r,k}(A) = \frac{k!}{(r-1)!(k-r)!} F(A)^{r-1} (1-F(A))^{k-r} f(A) \quad k=0\dots N$$

As before, A_1 and A_2 are defined as the best and second best applicant to come to the firm. Then the density of A_1 is given by

$$(8) \quad f_{k,k}(A_1) = \frac{k!}{(k-1)!} F(A_1)^{k-1} f(A_1)$$

and the density of A_2 is given by

$$(9) \quad f_{k-1,k}(A_2) = \frac{k!}{(k-2)!} F(A_2)^{k-2} (1-F(A_2)) f(A_2)$$

The expected output for the firm is

$$(10) \quad E(Q) =$$

$$P(1; \frac{1}{M}, N) \left[\int_{A^*}^{\infty} (\gamma + \delta A) f(A) dA + \int_0^{A^*} (\alpha + \beta A) f(A) dA \right] \\ + \sum_{k=2}^N P(k; \frac{1}{M}, N) \left[\int_0^{\infty} \int_{[(\beta-\delta)A_1 - \gamma]/\beta}^{\infty} (\gamma + \delta A_1 + \alpha + \beta A_2) f_k^k(A_1) f_k^{k-1}(A_2) dA_2 dA_1 \right. \\ \left. + \int_0^{\infty} \int_0^{(\beta-\delta)A_1 - \gamma/\beta} (\alpha + \beta A_1) f_k^k(A_1) f_k^{k-1}(A_2) dA_2 dA_1 \right]$$

¹⁰From the binomial, $P(k; 1/M, N) = \frac{N!}{k!(N-k)!} (\frac{1}{M})^k (1-\frac{1}{M})^{N-k}$

Next, cost is given by

(11) Average Cost=

$$P(1; \frac{1}{M}, N) \left[\int_{A^*}^{\infty} \frac{W_1(A)f(A)dA}{Q(A,0)} + \int_0^{A^*} \frac{W_2(A)f(A)}{Q(0,A)} dA \right] \\ + \sum_{k=2}^N P(k; \frac{1}{M}, N) \left[\int_0^{\infty} \int_{[(\beta-\delta)A_1-\gamma]^{1/\beta}}^{\infty} \frac{[W_1(A_1) + W_2(A_2)]}{Q(A_1, A_2)} f_k^k(A_1) f_k^{k-1}(A_2) dA_2 dA_1 \right. \\ \left. + \int_0^{\infty} \int_0^{(\beta-\delta)A_1-\gamma]^{1/\beta}} \frac{W_2(A_1)}{Q(0, A_1)} f_k^k(A_1) f_k^{k-1}(A_2) dA_2 dA_1 \right]$$

There is free entry of firms, which guarantees zero profits:

(12) Average Cost = E(R)

Average cost must equal the expected price.

Next, ex ante supply must equal ex ante demand in the product market so

(13) $M E(Q) = D(E(R))$

where D(R) is product demand.

Finally the actual price, R, is determined by spot market supply and demand so

(14) $\sum_m Q_m = D(R)$

where Q_m is defined as the actual output of firm m . Note that M is determined ex ante, but both the actual Q and R are determined ex post.

Because there are no fixed costs, the zero profit condition requires that $\lambda = 1$. If λ were less than 1, then firms would earn ex post profits that would not be offset by any costs. This nails down the wage functions, given in (6).

The system of equations (10)-(14) is five equations, which uniquely determine the five unknowns, $E(Q)$, AC, $E(R)$, actual R , and M.

Related Literature

Research on hiring has recently increased due to the availability of new data, but little has been done that resembles the approach taken here. In the past, those studying hiring needed to work with firms to obtain personnel records. But even that would not guarantee success because

most firms kept data only on those hired, not on all applicants.¹⁰ The advent of online job boards and online contracting firms has changed this. Much more is known about the types of workers firms seek and who is hired.¹¹

There is a very large literature on labor demand that could be broadly included in research on hiring.¹² Prominent in this literature is the empirical research on skill-biased technical change. Overall, the rising introduction of information technologies in the workplace has resulted in rising returns to education and greater demand for workers who do non-routine cognitive tasks.¹³ This literature would be consistent with the model here, in which more educated workers are demanded across a variety of jobs because the more educated can do a range of jobs that require cognitive skills. But that could always have been true. In the literature on technical change, the absolute advantage of more skilled workers rises over time.

Some results from search theory indirectly support the bumping implication. Kudlyak, Lkhagvasuren, and Sysuyev (2012) study a job applications website and show that when applicants begin their search, their education level explains their sorting to job postings: some job vacancies attract applications from highly educated and some don't. However, as search proceeds over time, education becomes a weaker predictor of what jobs applicants seek. Over time, the average job seeker begins to apply to jobs that only attracted the less educated in the first week of search. Thus, the highly educated have a broader range of job choices than the less educated and over time will bump the less educated from jobs as they search. This is in keeping with the search result that wage demands decline with the duration of unemployment.¹⁴

Other related literature involves vacancy dynamics. Andrews, Bradley, Stott, and Upward (2008) show, using data from the UK, that vacancies for non-manual work are less likely to be filled. Van Ours and Ridder (1991) examine data from the Dutch Bureau of Statistics. They find that jobs that require more education fill more slowly and that vacancy flows are more sensitive to the business cycle for low-education openings. Van Ours and Ridder (1992) find that higher education requirements are associated with longer vacancy durations.

One issue is whether the assumption that two jobs are available at the firm to which the worker applies is even broadly consistent with typical job posting behavior. In particular, the

¹⁰ Burks, Cowgill, Hoffman, and Housman (2015) show that firms source job applicants using current employee referrals. Hoffman, Kahn, and Li (2015) work with one large job applicant testing firm to identify optimal hiring processes.

¹¹ Agarwal, Horton, Lacetera, and Lyons (2013), Agrawal, Lacetera and Lyons (2012), Autor (2001), Brencic, (2009), Gee (2015), Ghani, Kerr, and Stanton (2014), Kuhn and Mansour (2004), Kuhn and Shen, (2013a, b), Kuhn, and Skuterud (2004), Marinescu and Wolthoff (2015), Nakamura, et. al. (2009), Stanton and Thomas (2015a,b), Pallais (2014). One relevant result is that employers have increased their minimal skill demand in response to the increase in job seekers during the Great Recession (Modestino, Shoag, and Balance (2015)).

¹² Education effects on hiring arise in the job market signaling literature in which workers invest in education to signal worker quality for the hiring decision (Spence (1973)). Altonji and Pierret (2001) show that education as a sorting device for workers diminishes with experience.

¹³ For early papers, see Abowd, Haltiwanger, Lane, McKinsey, and Sandusky (2007), Autor, Katz, and Krueger (1998), Bresnahan, Brynjolfsson, and Hitt (2002), Katz, and Murphy (1992).

¹⁴ See Rogerson, Shimer, and Wright (2005) for a review of the search theory literature.

notion is that a worker applies for multiple jobs at a firm and that a firm receives more than one applicant per job. The latter is surely correct, but Kuhn (2015)¹⁵ has shown, using Chinese Internet job posting data, that the typical (modal) number of jobs per ad is one. His data do reveal, however, that the *average* number of jobs available per advertised job is well above one and appears to be closer to two.

In this model, the worker applies to the firm and accepts the best job that is offered. It is useful to ask whether a worker's application is general or whether it is directed at one specific job. It is difficult to answer this question directly with administrative data, but some introspection, along with audit studies of the job application process, may be helpful in assessing the assumption of a general application. There are some jobs that are quite specific and it seems reasonable to assume that the worker applies to only that job. An example is an assistant professorship at a given department. But even here, an applicant might consider a post-doctoral fellowship if the assistant professorship has already been filled. In more typical jobs, applicants often do not even know what the job entails until going to the job interview. Consider, for example, an individual who applies for a managerial position at a major retailer. Without knowing the details of the firm's hierarchy and the duties of each position, it is unlikely that the worker is applying for the advertised job and that only. More likely is that the individual is applying for a managerial position at the firm and is willing to consider any potentially suitable job offered.

The sociology literature offers some support for the notion that individuals apply to firms rather than for a specific job at the firm, despite the vacancy being advertised as specific to a particular position. In a field experiment studying discrimination, Pager et al. (2009) document that minority applicants are more likely to be steered or channeled to jobs that are not customer facing even though they apply for the same positions as non-minority candidates. For example, black applicants who applied for server jobs at restaurants were more likely than whites to have busser or runner jobs suggested instead. Also consistent with the model, Pager et al. document cases of applicants being channeled into better jobs, presumably because no other suitable applicants for those jobs have arrived.

Data Sets

The predictions detailed below are tested using data from four sources. The Conference Board provides data on vacancies, the Current Population Survey (CPS) and the Panel Study of Income Dynamics (PSID) provide data on wages, education and occupation over time. oDesk provides data with hiring and job applicants to assess how the pool of applicants influences who is hired.

The Conference Board does a monthly survey of online job postings. The Conference Board Help Wanted OnLine (HWOL) is from jobs posted on 16,000 Internet job boards, corporate job boards and smaller job sites that serve niche markets. The Conference Board has two measures of job postings. One is "new ads" that are ads posted for the first time in the

¹⁵ Peter Kuhn, discussion slides, NBER conference, Stanford University, November, 2015.

previous month. The second is “total ads” that are new ads plus ads reposted from the previous month. The data is available by occupation by year, from 2006-2014, where yearly data is the average of the monthly data. There are 846 observations for 9 years times 94 occupations. The goal in using these Conference Board data is to create a variable that is similar to the vacancy rate, but is measured by occupation so that vacancy rates can be related to occupational skill levels. Using these Conference Board series, there are more vacancies when jobs go unfilled more than one month. Therefore, the measure of vacancies used below is the “unfilled jobs ratio,” which is the ratio of unfilled ads to total ads, where unfilled ads is the difference between new ads and total ads. Table 1 shows the average unfilled jobs ratio across occupations and years is .47.

The second data set used below is the March CPS, which provides information on wages and personal characteristics of CPS respondents. Data are obtained from 1975 to 2013. The sample is restricted to men who work full-time (defined as more than 35 hours per week) and are between the ages of 25 and 54, with a total sample size of 866,432 observations across the years. Wages are defined as real annual earnings, expressed in 2013 dollars. Education is defined as years of education and this measure varies with the survey year because different surveys categorize education groups differently over time. Consequently, all educational groupings are converted into a years of education to make them comparable. Another variable used extensively is occupation. The CPS changes its occupational definitions over time. The one used here is the 1990 occupational code, which is largely carried forward to 2013 and backward to 1975. This occupation variable is at a relatively fine level, with an average of 343 occupations delineated by this variable.¹⁶ Table 1 shows that average earnings are \$66,849 in 2013 dollars.

The third data set is the PSID, which is an unbalanced panel from 1968 through 2010. The dataset follows 5,382 men between the ages of 25 and 65, with an average of 5.25 years of data per person. The PSID originates with a sample in 1968 and then introduces new respondents into the sample as parents have children who become respondents. Table 1 shows that the average wage is \$25/hour in 2010 dollars. Note that the definition of occupation is coarse in the PSID as compared with the other data sets with only 25 occupations defined.

To detect more direct evidence of luck in hiring, data on job postings, applications, and new hires are available from online labor market oDesk. oDesk.com (recently rebranded as upwork.com after a merger with their largest competitor, elance.com) is an online labor market for outsourced services. As of the beginning of 2014, oDesk had processed over \$1.3 billion in contracts (Zhu, et. al. 2015). The oDesk platform allows employers to post jobs, hire from the online applicants, make payments to these globally distributed remote workers, and monitor workers with proprietary project management software.

¹⁶ There are a total of 384 unique 1990 codes throughout the dataset. However, not all occupations are present for all years. In earlier years, occupational codes were less precise. For example, no one falls into the “Legislator” occupational code until 1982 because occupational definitions before 1982 were too imprecise to place individuals into this category. There are 293 occupational codes in the dataset for 1975, and a high of 373 occupational codes in 1994.

The oDesk data provides a unique opportunity to study the role of luck in finding jobs because the transactions data contain records of employers' hiring along with the entire set of applications that employers receive. It is also possible to study how luck varies based on underlying skill requirements for a given job. This makes use of data on job requirements that come from job postings. When an employer posts a job opening, the task category is selected, along with a job title and a description of the work to be done remotely. Applicants then submit a short cover letter, their electronic resume as displayed on a profile maintained on oDesk, and, importantly, they also bid an hourly wage.¹⁷ For workers who have worked on oDesk before, there is a public evaluation (1 to 5 score) of past performance done by employers. Some workers are affiliated with agencies and others are unaffiliated. Stanton and Thomas (2015b), Ghani, Kerr, and Stanton (2014), Horton (2010), Pallais (2013) and Agrawal et al (2012) and Agrawal et al (2013) describe oDesk marketplace institutions in more detail.

Table 2 provides summary information about jobs posted on oDesk by different task category. There are 9 main categories of jobs posted on oDesk over the 2006 – 2010 data used here. The two largest job categories are Administrative Support and Web Programming. Taking wages as a proxy for skills, variation in average wages paid across categories is evident over the 2006-2010 period. The same patterns hold at quantiles of the wage distribution, and the 90th percentile is also displayed in the table.

Empirical Results

The main points from the theoretical model are used in sequence to develop a number of implications that are empirical.

Hypothesis 1: Bumping occurs. If two high ability workers show up at the firm, then the lowest ability of the two is bumped down to the easy job.

The oDesk data are used to test for bumping. In this online job site, the employer posts a job opening, applicants respond to that opening with their resumes and wage bids. The employer hires one of these applicants or none at all.

To test for bumping, the first step is to build a model of the probability that each individual worker is hired within a given slot. Stanton and Thomas (2015a) provide this model for the oDesk data by writing a probability model that maps the employer's objective function to data. Because oDesk contracts are for hourly work, the employer on job opening i seeks to hire the individual worker who produces the most output per hourly wage subject to this being greater than the payoff from not doing the job. For each worker j , the employer observes that part of worker quality that is commonly valued in all production tasks, q_j , and chooses the applicant in the set of applicants, denoted C_i , that has the highest expected wage-adjusted quality. This is given by

¹⁷ Employers may also search for worker profiles directly and invite individual workers to apply.

$\max_{\{j \in J_i \cup 0\}} \frac{q_j \exp(\varepsilon_{ij})}{w_{ij}}$, where $C_i = \{\{q_j, \varepsilon_{ij}, w_{ij}\}_{j=1}^{J_i}, \{0, \varepsilon_{i0}, 0\}\}$ includes a no hire option, $\{0, \varepsilon_{i0}, 0\}$. With a purely idiosyncratic Type-1 extreme value distributed error, ε_{ij} , the probability of being hired takes a conditional logit form within job opening. Taking logs, the probability that worker j is hired is the probability that worker j is the highest order statistic in the choice set, or $\exp(\log(q_j) - \alpha \log(w_{ij})) / [1 + \sum_j \exp(\log(q_j) - \alpha \log(w_{ij}))]$.

Taking this to data, the commonly valued components of productivity are modeled as $\log(q_j) = X_{jt}\beta + v_{jt}$, where v_{jt} is an error component that is possibly correlated with worker quality, may be observable to employers, but is not observed in the data.¹⁸ This creates the possibility that the bids submitted to jobs are endogenous.

The ability for bids to adjust to job characteristics, individual characteristics, or, most importantly, the extent of competition does not mean that luck is absent from labor market allocations. However, if wages adjust to compensate for different probabilities of being hired, endogenous bids are likely to make detection of luck more difficult. Therefore, a technique to incorporate endogeneity of bids is used.

Variation in exchange rates provides an instrument to control for wage bid endogeneity. Workers' local labor market opportunities are denominated in their own currency, while contracts on oDesk are all denominated in dollars. With any friction that limits immediate adjustment of local prices to exchange-rate parity, appreciation of the local currency relative to the dollar (one dollar provides more local currency units) results in a reduction in the dollar-denominated value of workers' outside option. This shift in exchange rates is expected to change equilibrium bidding behavior online, but it is independent of unobserved worker quality, v_{jt} .

With this instrument, the control function approach of Petrin and Train (2009) is used to estimate the parameters governing the probability that an individual worker is hired. The idea behind this empirical strategy is to control for the endogenous portion of wage bids directly in the choice model. The endogenous portion of the wage bid is that part of wages that is correlated with unobserved productivity of the worker. This endogenous component is orthogonal to the fitted values of the first stage regression of bids on the instruments. The first stage regression is

$$\log(w_{jt}) = a_1 + Z_{jt}\gamma_1 + X_{jt}\gamma_2 + u_{jt}$$

where Z_{jt} is the z-score of the local currency-to-US dollar exchange rate in month t for worker j .¹⁹ The first stage results are in Appendix Table 1.

¹⁸ In the richer model in Stanton and Thomas (2015a), employers also have horizontal measures of preferences, but because the horizontal measures are orthogonal to worker skills, they are less important for quantifying the extent of luck due to slots.

¹⁹ One may worry about worker sorting on the instrument at the extensive margin of applying for jobs. See Stanton and Thomas (2015b) for detail about robustness checks regarding extensive-margin selection on the instrument.

Let $CF_{jt} = \hat{u}_{jt}$ denote the fitted residuals from this first stage regression for worker j on job opening i during month t . Petrin and Train’s approach includes CF_{jt} in the conditional logit model. Variation in wage bids caused by variation in the instruments is “good variation” that is useful for fitting the model; the remaining variation is captured by the control function and includes unobserved applicant quality that is correlated with wage bids. Separating these sources of variation in bids allows for consistent estimates of α , the parameter that governs price-related substitution patterns. The residuals are uncorrelated with that part of log wages that vary with the instrument, Z . By putting in the residuals, the model allows different loading on the residuals and on the log wage. The resulting choice probabilities used to form the likelihood are

$$(15) p_{j|i} = \exp(X_{jt}\beta - \alpha \log(w_{ij}) + CF_{jt}\psi) / (1 + \sum_j^i \exp(X_{jt}\beta - \alpha \log(w_{ij}) + CF_{jt}\psi)),$$

and the log likelihood to be maximized is $\sum_i \sum_j d_{ij} \log(p_{j|i})$ where d_{ij} is an indicator for the option chosen. The sum over j in the log likelihood includes the no-hire (outside) option. Table 3 displays the parameter estimates for the two largest types of jobs that span low skill and high skill, administrative support and web programming, respectively. When the model uses generated regressors for the control function, block-bootstrapped standard errors are reported with each block corresponding to a job opening. Otherwise, robust standard errors are reported.

The parameters themselves are not directly the objects of interest, so they are reported in Appendix Table 2. However, the basic results are sensible: the estimated parameter values show that employers value applicants with better feedback scores and more past experience. In most cases, the parameters on worker characteristics are larger when log wage bids and the control function enter the model, suggesting that wage bids are positively correlated with characteristics that employers value.²⁰

With estimates of the parameters in hand, it is possible to assess how sensitive the hiring probability for one worker is with respect to the characteristics of other workers who apply for a job. To capture differences in quality as implied by employers’ revealed preferences, the parameters $\beta_{JobCategory}$ are used as weights on worker characteristics to construct an index of worker quality. An individual worker’s quality index within each job category is then just the fitted value $X_{jt}\beta_{JobCategory}$. Using this index allows applicants to be ranked. Because workers cannot coordinate, it is possible to use the stochastic arrival of other workers to assess how arrival of different workers changes the hiring probability of an individual worker. This allows an assessment of luck. To do so, the actual arrival of workers to jobs is used. The sample is limited to the first six applicants to a job to abstract away from different numbers of job applicants to different vacancies and to balance employers who hire quickly with those who wait to receive more applications.²¹ In some settings, it may also be possible to examine how hiring probabilities vary with the number of applicants to a job. This is difficult on oDesk, however,

²⁰ Also consistent with Stanton and Thomas (2015b), employers more highly value novice applicants with agency affiliation, but this effect fades out for more experienced applicants.

²¹ A spline with applicant order is included in the characteristics, X , but to be included in the sample, job openings must have at least 6 applicants, making the assessment balanced across job openings. Only worker-initiated applicants are considered for the purposes of these calculations. 92 percent of the first 6 applications to web programming jobs are initiated by workers.

because most jobs have many more applicants than available slots. There is no variation that we have discovered that convincingly moves the number of applicants to jobs independently. A cleaner exercise examines the stochastic composition of applicants rather than the number of applicants.

Given these results, it is possible to test Hypothesis 1, namely, that bumping occurs. Using the index of applicant quality, a worker is said to be lucky if the next applicant to arrive has a lower index value. A worker is unlucky if the next applicant to arrive has a higher index value. Additionally, workers themselves can be ranked overall relative to the distribution of quality, and a worker is said to be of good quality if he or she is above the median quality index for a job category; otherwise the worker is classified as bad quality.

Panel A of Table 3 displays the results. Good applicants are uniformly more likely to get a job when the next applicant to arrive is of lower quality. For both administrative support and web programming, the change in hiring probability for good applicants due to luck is about 30% (Columns 2 and 4). In levels, however, individual applicants in web programming are less likely to be hired regardless of luck, consistent with the predictions of the model that not all workers are qualified for more skill-intensive tasks. For bad applicants, those below the median of the quality index, luck also plays a role but the change in hiring probability due to luck is much smaller; these workers are very unlikely to be hired, and with multiple applicants to a job, conditioning only on the identity of the next applicant when computing luck has very little effect on hiring probabilities.

The results in Table 3 corroborate the basic assumptions on which the model is constructed. Hiring depends not only on a given applicant's characteristics, but also on the characteristics of the others with whom he or she competes for the job. The job to which an applicant is assigned, if any, depends on the competition. Although this seems obvious at the most intuitive level, it is, as far as we are aware, the first evidence of its kind that establishes the relative nature of the hiring process. If a better applicant is present, the worker is given a lower quality job or none at all. It is also consistent with the view that slots are a fundamental part of the hiring process.

There is a potential concern with many of these estimates due to the applicant seeing the job description, which is unobserved by the econometrician. Any sorting based on the job description is not a problem for the conditional logit estimates behind the quality index because pairwise comparisons between workers identify the parameters, but sorting may be problematic for assessing an employers' comparison of a worker and the no-hire option. Sorting may also muddle the interpretation of luck: are some workers unlucky due to other applicants' arrival or are workers inappropriately applying to jobs for which they are likely to be unqualified, taking a chance on themselves. For example, when it is a tough job, mostly good applicants apply so the first applicant is less likely to get the job because it is only for highly skilled people. This is like low quality students who take a shot at applying to Stanford. They have a small chance of getting in, not because they were unlucky to be followed by a high quality applicant, but rather because the quality of the applicant pool reflects the overall difficulty of getting into the program in the first place.

To assess the extent to which sorting affects the estimates, employer-by-applicant position fixed effects are included to remove much of the influence of unobserved differences in job or employer quality. These results are presented in Panel B of Table 3. The idea behind these results is to measure whether a high quality first applicant to employer i on job 1 is more likely to get the job if the second applicant is low quality than a similar high quality first applicant to employer i on job 2 when an even higher quality second applicant appears. Because fixed effects are required, the results are presented as regression output, with block-bootstrapped standard errors by employers. The inclusion of employer-by-applicant fixed effects minimally changes the results for good quality applicants. For example, in column 2, a good applicant to an administrative support position has a hiring probability of about 0.021 if the next applicant is ranked lower, but this probability falls to 0.015 if the next applicant is better. This decline is on the order of about 37%. The decline is about 32% in column 4 for web programming. Overall, these results suggest tests for the presence of luck are not being confounded by unobserved attributes about individual jobs or employers.

Hypothesis 2: Low ability workers are more likely to be unemployed than high ability ones. High ability workers can work both the easy and difficult jobs, which makes them more employable. The model predicts that sufficiently high ability workers can never be unemployed.²²

Figure 3 shows the unemployment rate by education using the BLS information. As expected, there is a considerable increase in unemployment as education falls. To reiterate, although the fact is not a new one, the pattern requires explanation. Low skill does not mean unemployable in the same sense that low-priced, lower quality goods are not more likely to stay on the shelves of a store longer than are high-priced, higher quality goods. Indeed, in most cases, the reverse is true. The reason that the low skilled are more likely to suffer unemployment than the highly skilled is that the highly skilled are capable of doing a larger variety of jobs, whereas the less skilled can do many fewer. This means that if an applicant encounters a situation where other applicants are better suited to the job than he, the high ability applicant may be offered another job, when the low ability applicant is not.

Hypothesis 3: Vacancy rates are highest for the high-skilled jobs. The reason is that easy jobs are never unfilled: a high or low ability worker can take them. The difficult jobs may go unfilled if no sufficiently high ability worker arrives at the firm.

As described above, the Conference Board “vacancy rate” in these data is the unfilled jobs ratio, equal to the unfilled job postings divided by total job postings. The hypothesis is that this rate rises with skill level. The Conference Board data include the SOC 3-digit occupation, of which there are 94. An Occupational Skill level is attached to each occupation by going to the

²² Lemma 2 yields a sufficient condition for never suffering unemployment. The most stringent form of the condition occurs when the lowest ability worker has ability equal to zero. Then, as long as A_1 exceeds $-\gamma/(\delta-\beta)$, it is certain that both workers are employed. Thus, any worker with ability greater than $-\gamma/(\delta-\beta)$ can be certain that he will be at a firm that employs both applicants and can never suffer unemployment.

March CPS data and calculating the average education for each of these occupations by year: highly skilled occupations are those occupied by highly educated people.

The first test of Hypothesis 3 is in Table 4. Regression results show that an increase in the Occupational Education results in a higher unfilled jobs ratio. Job postings stay unfilled longer when the jobs are more skilled. The unit of analysis is an occupation-year. The dependent variable is the number of unfilled postings in a month divided by total postings averaged over all the months for that year. Similarly, the independent variable is a year dummy interacted with average education in the occupation. Thus, there are 94 occupations times 9 years, or 846 observations. Each year-occupation variable is an independent test of the hypothesis so the fact that all nine years produce significant results is strong confirmation that unfilled vacancies are higher in the occupations with the highest levels of education. The coefficients are sizable: moving from a high school educated occupation to a college educated increases the percent unfilled by about .15, over a mean of .47.

The second test of Hypothesis 3 is in Table 5 using oDesk online vacancy data. The dependent variable is whether a posted job has been filled. It is regressed on pay as a measure of job skill, as well as employer fixed effects and time fixed effects. Employer effects help to remove unobserved employer differences in familiarity with the platform or differences in unobserved employer attractiveness. The two measures of skill used are the mean wage in the job category and the 90th percentile of wages in the job category. The probability of filling a job is negatively related to both measures. Table 5 also displays point estimates for fill probabilities for each of the main job categories in the data.

By way of background, note that in Table 2 applications per-vacancy are highest in low-skill tasks (as measured by mean the mean or the 90th percentile of wages). Hiring rates are also higher in jobs with low skill requirements, as shown by a comparison of Column 3 for administrative support with web programming or software development.

In sum, using measures of job vacancies for jobs posted online, vacancies rise with skill level. Although this may not be a completely novel finding, it is consistent with the model presented here.

The combination of the results for Hypotheses 2 and 3 produces a cross-sectional analogue of the Beveridge curve. The Beveridge curve is usually applied to the economy as a whole over time and reveals that periods of high unemployment are also periods with low vacancy rates. The cross-sectional version of that point is that occupations that have high vacancies tend to have low unemployment rates and vice versa. This does not follow directly, however, because the vacancy data are for jobs whereas the unemployment data are for workers. Still, the education levels relate to the workers who are occupants of the jobs even in the Conference Board data set so it is reasonable to conclude that highly educated workers not only have low unemployment rates, but that they are also found in jobs with high vacancy rates.

Hypothesis 4: The over-qualified and under-qualified will have a positive observed wage gap, $\hat{W} - W$ where \hat{W} is the wage that the worker would receive were he placed in the job in which he has an absolute advantage.

The PSID is used to test this hypothesis. To do so, it is necessary to define over and under-qualified in the PSID data along with an individual's usual occupation. Consider, for example, an individual who is in one occupation, the usual occupation, say physician, for most of her life and then switches to a less skilled occupation, say, retailing. This could reflect life cycle choice, where the highly skilled person decides to take an easier job as she moves gradually into retirement. It could also reflect an involuntary move that results from a primary job loss that forces the worker to accept another job. Either case is consistent with the formal specification in the model where the worker's productivity, $\gamma + \delta A$, is higher in the usual occupation than is productivity, $\alpha + \beta A$, in the unusual one. The prediction is that her wage should be higher in her usual job than in the unusual one for which she is over-qualified.

An opposite example is also possible. Consider a journeyman machinist who has spent almost his entire career in that job. Now suppose that his plant closes and he is forced to find another job. Unable to find another machinist job, he locates a clerical position in a start-up. He is not well-suited to that position, but because the start-up can find no one better to fill the job, they hire the former machinist. His productivity as machinist, $\alpha + \beta A$, exceeds his productivity in the clerical position, $\gamma + \delta A$, so he is under-qualified for the clerical job. He has an absolute advantage as a machinist, but works as a clerk because he can find nothing that suits his skill set and the firm that hires him can find no one better. This is bad luck. The worker is under-qualified for the clerical position and should earn less there than he did as a machinist. Once again, the predicted wage, based on his skills and assignment to his appropriate job, in this case machinist, should exceed what he earns as a clerk in the start-up.

Related to these examples, over and under-qualification and the usual occupation were defined in the following way. For each individual, the modal occupation was determined, defined as the occupation in which the worker spent the most years. A worker was deemed to be in an "unusual" occupation if the occupation held during that year differs from the modal occupation.¹⁴ About one-fifth of the observations fit this definition of being unusual. The worker in an unusual occupation was defined as over-qualified if the mean wage of her usual occupation exceeded the mean wage of the unusual occupation. Conversely, the worker in the unusual occupation was defined as under-qualified if the mean wage of her unusual occupation exceeded the mean wage of her usual occupation. This resulted in 9.7% of the observations being classified as under-qualified and 10.3% being classified as over-qualified.

There are two implications for the panel data tests. First, as before, the wage in the usual occupation should exceed that received in the unusual occupation for which the worker is either over- or under-qualified: the wage gap, $\hat{W} - W$, should be positive for both subsamples. In Figure 4, the worker with ability A' who is assigned to the easy job is over-qualified. She would have earned the amount that corresponds to point 1 were she in the appropriate job (the difficult one), but earns only the amount that corresponds to point 2. Second, and important, even though over-qualified workers receive less in the job for which they are over-qualified than in their usual job, they should still earn more than the typical worker in the easy job. In Figure 4, although point 2 lies below point 1, it lies above point 3, which yields the wage of the typical worker in

¹⁴Observations for which there were two or more modal occupations were dropped.

the easy job. Even in the easy job, output increases in ability so her wage should be higher than the median for that occupation.

Conversely, under-qualified workers not only receive less in the job for which they are under-qualified than in their usual job, they receive less than the typical worker in the difficult job. In Figure 4, the under-qualified worker is one who has ability level A", but works in the difficult job. Instead of receiving the wage that corresponds to point 3, he receives the lower wage that corresponds to point 4. An under-qualified worker also receives less than the typical worker in the job for which he is under-qualified because output increases in ability and his ability is low for that occupation. Therefore, his wage should be lower than the median for that occupation, where the wage that he receives at point 4 is lower than the wage that the typical worker in the difficult job receives at point 1.

Taking this to data, it is helpful to use differences in log wages because of skewness in levels. The estimating equation for counterfactual log wages is

$$(16) \log \hat{W}_{it} = b_{0i} + b_1 f(\text{Age}_{it}) + b_2 \text{Tenure}_{it} + t + e_{it}$$

The predicted log wage is a function of age and age-squared, weeks of tenure, as well as individual fixed effects, b_{0i} and year fixed effects, t . The coefficients in (16) are obtained by estimating the regression on the sample of qualified workers (who are neither over- nor under-qualified for their jobs as defined above). The predicted log wage for over- and under-qualified workers uses these coefficients to construct the counterfactual.

The wage regression is estimated using only observations that are in the usual occupation, which implies that no under- or over-qualified job spells influence the results. The predicted log wage can be computed for the unusual job spells, using the coefficients and person fixed effects that are estimated in the usual occupation regression. This provides an estimate of what the worker would receive were he in his normal job.²¹⁵

The resulting wage gap, $\log \hat{W}_{it} - \log W$, is predicted to be positive for both over- and under-qualified workers. For example, a worker who is over-qualified has $A > A^*$ in Figure 1, but is working in the easy job, earns only $\alpha + \beta A$ instead of the appropriate and higher $\gamma + \delta A$, as predicted by (16).

Summarizing, there are four patterns that should appear in the results. First, those who are over-qualified for their jobs should earn less in that job than they would be estimated to earn were they properly placed in the job in which they are more productive. Second, those who are

¹⁵The regressions were done two ways. The first specification included age, age-squared, education, quadratic terms and interactions as well as person fixed effects. The second included only year and person fixed effects. The R-squared was almost identical and above .8 in both specifications because once fixed effects are included, only time varying education and aging contributes to the regression, the latter being captured mostly by year effects. Results here are based on the simple specification.

under-qualified for their jobs should also earn less in that job than they would be estimated to earn were they properly placed in the job in which they are more productive. Third, although the over-qualified earn less than in they would in their proper jobs, they should earn more than the typical worker in the job for which they are over-qualified. Fourth, not only do the under-qualified earn less than they would in their proper jobs, but they also should earn less than the typical worker in the job for which they are under-qualified.

The results are reported in Table 6. Column 1 reports the results for those job spells that correspond to under-qualification and Column 2 for those that correspond to over-qualification.²³ The first row reports that the average log wage gap between the predicted wage and the actual wage is positive as predicted for both groups. Because individuals who are incorrectly assigned have lower productivity than they would have were they in their appropriate jobs, their wages are below that predicted. The number reported is the average log wage gap across all person-years that fit the definition of under- or over-qualification. The average log wage gap is positive and statistically significant for the over-qualified, but, although positive, not statistically significant for the under-qualified.

The second set of results reports the average difference between the received log wage and the mean log wage of workers who are in the occupation, most of whom are there appropriately. The prediction is that this gap should be negative for the under-qualified and positive for the over-qualified. Both predictions are borne out, and the differences are statistically significant. The panel data approach is consistent with the theory and treats the issue of unobservables.

Finally, the PSID data allows an assessment of which workers are in jobs for which they are over and under-qualified. Figure 5 presents the results of local polynomial regressions that flexibly characterize the probability of over- and under-qualification as a function of the log wage in the usual occupation. Consistent with the theory, those in the middle of the distribution are those who are likely to be in unusual jobs. The probability of over qualification is increasing with skill up to a point and then declines. Under-qualification is declining with the usual log wage. The previous results say nothing about who is in an unusual job; these results suggest that those who have high wages in their usual occupation are most at risk for over-qualification up to a point. Interestingly, those with the highest usual log wages are less likely to be overqualified than those who earn slightly less in their usual occupations. That is, those at the very top are unlikely to be bumped down, and this inverted u-shape is exactly what comes out of the theory. Those with low wages in their usual occupation are most at risk for under-qualification, again consistent with the theory.

Hypothesis 5: When the return to skills rises over time, there are increasingly adverse wage consequences of mismatch. This result depends on the assumption, made earlier, that the nature

²³ Individuals are dropped from the sample if their wage is an outlier, measured as their wage being greater than three standard deviations from the occupational average wage. Comparing Table 8 to Tables 6 and 7, Table 8 follows the rows that drop wages greater than three standard deviations, and there are no columns in Table 8 because the over- and under-qualified are those that are outside their modal occupation.

of skill-biased technical change steepens the relation of wages to ability more in difficult jobs than in easy jobs.

It is well known that over the last thirty years, the return to education has risen. It is natural to expect, as argued earlier, that the return to ability has increased more in high-skilled jobs than in low-skilled ones. As discussed earlier, this implies that the difference between δ' and δ is greater than the difference between β' and β , as shown in Figure 2. The implication is that the variance of pay is greater today for the highly able than it was in the past: the wage loss for taking the easy job in the past was $b-a$; today the wage loss is $d-c$. Thus, as the return to skills has risen, there is a rising variance of pay for the highly able.

Using the CPS data, the equation to test this is a simple regression:

$$(17) \quad \sigma_{it} = b_1 \text{Year}_{it} + b_2 \text{OccupationalSkill}_{it} + b_3 \text{OccupationalSkill}_{it} * \text{Year}_{it} + e_{it}$$

where the variance of pay is calculated for each occupation i and for each year t , resulting in a data set of 12,733 observations (for 39 years times an average of 326 occupations). The OccupationalSkill is the median education for that occupation each year. The first implication is that $b_1 > 0$ because the rising return to skills increases the variance of earnings over time. The second implication is that $b_3 > 0$ because the variance of earnings rises more for the highly able, as suggested in Figure 2.

Regression results in Table 7 are consistent with both implications. The variance of pay has risen over time, but it has risen most for the highly skilled.

Conclusion

A worker's skills alone do not determine the job in which he or she is hired or indeed, whether the worker is hired at all. The existence of slots or job positions means that even qualified workers may not be hired or may not be assigned to the job for which they are best suited because there is a superior applicant for that position.

Although this idea is intuitive, it has not been modeled in a way that allows its implications to be explored in individual-based data. The model and analysis herein not only provides many specific predictions on what should be observed in hiring and job assignment, but tests and validates those predictions using four different data sets.

First, the job to which an applicant is assigned and whether the applicant is hired at all depends not only on own skills but also on the skills of the competition. This is verified using oDesk data.

Second, bumping occurs where workers take jobs for which they are not well-suited, but receive the offer because their skills are superior to those of other applicants but inferior to those applying to the job that they prefer. The model provides a clear definition of "over-qualification" and of "under-qualification" that has specific empirical meaning. Using these definitions, the PSID data provide evidence that over- and under-qualification occurs and that the wages that are

received in those jobs are exactly as predicted by the model. Namely, both over- and under-qualified workers receive less in those jobs than they would in their appropriate positions, but over-qualified workers receive more than the average worker in that job. Conversely, underqualified workers not only receive less than they would in their appropriate job, but also less than the average worker in the job for which they are underqualified.

Third, less able workers are more likely to be unemployed because the more able workers are capable of doing a wider variety of jobs. This implication is not as obvious as it seems. The model provides this as an implication, and not surprisingly, the implication is found in data provided by the BLS.

Fourth, vacancy rates are higher in jobs that require higher levels of skill. The lower skilled jobs can be filled by almost all workers, but only the smaller group of high ability workers are more able to do the high skilled jobs. The Conference Board data on vacancy rates confirms this prediction.

In sum, a model that explicitly incorporates slots and allows for other applicants to compete with an individual for a job is not only intuitive, but provides many testable implications that are confirmed empirically.

References

- Abowd, J. M., Haltiwanger, J., Lane, J., McKinsey, K. L., and K. Sandusky, 2007, "Technology and the demand for skill: an analysis of within and between firm differences," *NBER Working Paper 13043*.
- Agrawal, Ajay, Lacetera, Nicola, and Elizabeth Lyons, 2012, "How Do Online Platforms Flatten Markets for Contract Labor?" *Working Paper*.
- Agrawal, Ajay, Horton, John, Lacetera, Nicola, and Elizabeth Lyons, 2013, "Digitization and the Contract Labor Market: A Research Agenda," *NBER Working Paper 19525*.
- Altonji, Joseph G., and Charles R. Pierret, 2001, "Employer Learning and Statistical Determination," *Quarterly Journal of Economics*, 116 (1): 313-350.
- Andrews, M. J., Bradley, S., Stott, D., and R. Upward, 2008, "Successful employer search? An empirical analysis of vacancy duration using micro data," *Economica*, 75: 455–480.
- Autor, D. H., 2001, "Wiring the labor market," *Journal of Economic Perspectives*, 15: 25–40.
- Autor, D. H., Katz, L. F., and A.B. Krueger, 1998, "Computing inequality: Have computers changed the labor market?" *Quarterly Journal of Economics*, 113: 1169–1213.
- Autor, Levy, and Murnane, 2003, "The Skill Content of Recent Technological Change: An Empirical Exploration," *Quarterly Journal of Economics*, 118 (4): 1279-1334.
- Becker, Gary S. 1993. *A Treatise on the Family*, Cambridge: Harvard University Press.
- Brencic, V., 2009, "Do employers respond to the costs of continued search?" *Oxford Bulletin of Economics and Statistics*, 71: 1–25.
- Bresnahan, T. F., Brynjolfsson, E., and L. M. Hitt, 2002, "Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence," *Quarterly Journal of Economics*, 117: 339–376.
- Burks, Stephen V., Cowgill, Bo, Hoffman, Mitchell, and Michael Housman, 2015, "The Value of Hiring Through Employee Referrals," *The Quarterly Journal of Economics*, 130 (2): 805-839.
- Diamond, Peter, 1982, "Wage Determination and Efficiency in Search Equilibrium," *The Review of Economic Studies*, 49 (2): 217-227.
- Gale, D. and L.S. Shapley, 1962. "College Admissions and the Stability of Marriage," *The American Mathematical Monthly*, 69,1: 9-15.
- Gee, Laura, 2015, "The More You Know: Information Effects on Job Application Rates by Gender in a Large Field Experiment," *Working Paper*.
- Ghani, Ejaz, Kerr, William, and Christopher Stanton, 2014, "Diasporas and outsourcing: Evidence from oDesk and India." *Management Science*, 60 (7): 1677-1697.

- Gibbons, Robert, Katz, Lawrence, Lemieux, Thomas, and Daniel Parent, 2005, “Comparative Advantage, Learning, and Sectoral Wage Determination,” *Journal of Labor Economics*, 23 (4): 681-724.
- Hoffman, Mitchell, Kahn, Lisa, and Danielle Li, 2015, “Discretion in Hiring,” *Working Paper*.
- Horton, John, 2010, “Online Labor Markets,” *Working Paper*.
- Kahn, Lisa, 2010, “The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy,” *Labour Economics*, 17 (2): 303-316.
- Katz, Lawrence F., and Kevin M. Murphy, 1992, “Changes in Relative Wages, 1963-1987: Supply and Demand Factors,” *The Quarterly Journal of Economics*, 107 (1): 35-78.
- Kudlyak, Marianna, Lkhagvasuren, Damba, and Roman Sysuyev, 2012, “Systematic Job Search: New Evidence from Individual Job Application Data,” *FRB Richmond WP #12-03*.
- Kuhn, Peter, and Hani Mansour, 2014, “Is Internet Job Search Still Ineffective?” *Economic Journal*, 124 (581): 1213–1233.
- Kuhn, Peter, and Kailing Shen, 2013, “Do Chinese Employers Avoid Hiring Overqualified Workers? Evidence from an Internet Job Board,” *Research in Labor Economics*, 37: 1-30.
- Kuhn, Peter, and Kailing Shen, 2013, “Gender Discrimination in Job Ads: Evidence from China,” *Quarterly Journal of Economics*, 128 (1): 287-336.
- Kuhn, Peter, and Mikal Skuterud, 2004, “Internet Job Search and Unemployment Durations,” *American Economic Review*, 94 (1): 218-232.
- Lazear, Edward P., 1986, “Retail Pricing and Clearance Sales,” *American Economic Review*, 76 (6): 14-32.
- Marinescu, Ioana, and Ronald Wolthoff, 2015, “Opening the Black Box of the Matching Function: The Power of Words,” *Working Paper*.
- Modestino, Alicia, Shoag, Daniel, and Joshua Balance, 2015, “Upskilling: Do Employers Demand Greater Skills When Workers are Plentiful?” *Working Paper*.
- Nakamura, Alice, Emi Nakamura, Kathryn Shaw, Richard B. Freeman, and Amanda Pyman, 2009, in David H. Autor (editor), *Studies of Labor Market Intermediation*, University of Chicago Press.
- Oyer, Paul, 2008, “The Making of an Investment Banker: Stock Market Shocks, Career Choice, and Lifetime Income,” *Journal of Finance*, 63 (6): 2601-2628.
- Pager, Devah, Bart Bokowski, and Bruce Western, 2009, “Discrimination in a Low-Wage Labor Market: A Field Experiment,” *American Sociological Review*, 74 (Oct): 777-799.
- Pallais, Amanda, 2014, “Inefficient Hiring in Entry-Level Labor Markets,” *American Economic Review*, 104 (11): 3565-3599.

Petrin, Amil, and Kenneth Train, 2009, “A Control Function Approach to Endogeneity in Consumer Choice Models,” *Journal of Marketing Research*, XLVI.

Rogerson, Richard, Shimer, Robert, and Randall Wright, 2005, “Search-Theoretic Models of the Labor Market: A Survey,” *Journal of Economic Literature*, 43: 959–988.

Shimer, Robert, 2005. “The Cyclical Behavior of Unemployment and Vacancies,” *American Economic Review*, 95,1: 25-49.

Shimer, Robert and Lones Smith, 2000. “Assortative Matching and Search,” *Econometrica*, 68, 2: 343-69.

Spence, Michael, 1973, “Job Market Signaling,” *Quarterly Journal of Economics*, 87 (3): 355-374.

Stanton, Christopher, and Catherine Thomas, 2015a, “Information Frictions and Observable Experience,” *Harvard Business School Working Paper*.


Stanton, Christopher, and Catherine Thomas, 2015b, “Landing the First Job: The Value of Intermediaries in Online Hiring,” *Review of Economic Studies*, Forthcoming.

van Ours, J., and G. Ridder, 1991, “Cyclical variation in vacancy durations and vacancy flows,” *European Economic Review*, 35: 1143–1155.

van Ours, J., and G. Ridder, 1992, “Vacancies and the recruitment of new employees,” *Journal of Labor Economics*, 10: 138–155.

Zhu, Feng, Rory McDonald, Marco Iansiti, and Aaron Smith, 2015, “Upwork: Reimagining the Future of Work,” Harvard Business School Case 616-027.

Figure 1







Figure 3:

Earnings and unemployment rates by educational attainment



Note: Data are for persons age 25 and over. Earnings are for full-time wage and salary workers.
Source: Current Population Survey, U.S. Bureau of Labor Statistics, U.S. Department of Labor

From: Bureau of Labor Statistics, Employment Projections, April 2, 2015:
http://www.bls.gov/emp/ep_chart_001.htm

Figure 4




Figure 5: Probability of Over- and Under-Qualification by Log Wage in Usual Occupation




Table 1
Summary Statistics

Variable	Obs	Mean	Std. Dev.	Min	Max
<u>Conference Board Data:</u>					
Mean Education by Occupation	846	14.02793	1.615229	9.166667	18.13372
Unfilled Jobs Ratio	846	.4704046	.0970282	.0761404	.7067247
<u>CPS Data:</u>					
Age	870665	39.04593	8.267556	25	54
Education	870665	13.61305	2.763408	0	19
Yearly Earnings	866432	66849.35	58876.55	1.085666	1845631
<u>PSID Data:</u>					
Age	28255	40.76602	10.73039	25	65
Education	28106	13.2164	2.697894	0	17
Tenure (weeks)	27010	49.25442	82.5911	0	780
Hourly Wage	28255	25.01306	13.7365	.0004947	154.7392

Note: The sample size for the Conference Board data is the 9 years (2006-2014) for 94 occupations. The sample size for the CPS data is the number of men age 25-54 for March years 1975-2013. The sample size for the PSID data is men age 25-65 from 1968 to 2010.

Table 2
Summary Statistics by Job Category, oDesk Data

Job Category	Vacancies	Applications Per Vacancy	Mean Probability that Vacancy Is Filled	90th Percentile of Log Hourly Wages (Filled Vacancies)	Mean Log Hourly Wage	St. Dev. Of Log Hourly Wage
Administrative Support	32,175	26.15	0.49	2.08	0.97	0.91
Business Services	3,129	14.79	0.33	3.00	1.64	1.01
Customer Service	1,565	28.85	0.29	2.57	1.41	1.07
Design & Multimedia	18,737	15.14	0.40	3.00	2.35	0.70
Networking & Information Systems	4,905	11.35	0.34	3.51	2.67	0.77
Sales & Marketing	14,349	11.32	0.37	2.66	1.56	0.87
Software Development	21,818	10.21	0.33	3.22	2.57	0.96
Web Programming	100,051	14.33	0.40	3.09	2.49	0.65
Writing & Translation	20,766	13.02	0.43	2.81	1.75	0.88

Table 3: Hiring Probabilities Differ Based on Who Else Applies for a Job

	(1)	(2)	(3)	(4)
Sample:	Administrative Support	Administrative Support, With Control Function	Web Programming	Web Programming, With Control Function
<i>Panel A: Hiring Probabilities by Own and Future Applicant Type for the First 6 Applicants</i>				
Good Applicant and Better Next Applicant	0.0144	0.0155	0.0117	0.0129
Good Applicant and Worse Next Applicant (Standard Error on Difference)	0.0229 (0.0013)	0.0200 (0.0015)	0.0187 (0.00039)	0.0169 (0.00039)
Bad Applicant and Better Next Applicant	0.00426	0.00984	0.00424	0.00814
Bad Applicant and Worse Next Applicant (Standard Error on Difference)	0.00523 (0.00014)	0.0115 (0.00065)	0.00506 (0.0001)	0.00864 (0.00021)
<i>Panel B: Hiring Probabilities Regressions with Employer and Applicant Position Fixed Effects. Baseline is a bad applicant with worse next applicant</i>				
Good Applicant and Better Next Applicant	-0.00835 (0.0009)	-0.00557 (0.0013)	-0.00698 (0.0003)	-0.00412 (0.0003)
Good Applicant and Worse Next Applicant	0.0168 (0.0017)	0.0106 (0.0023)	0.0133 (0.0007)	0.00862 (0.0007)
Bad Applicant and Better Next Applicant	-0.00158 (0.0004)	-0.000465 (0.0007)	-0.00104 (0.0002)	-0.000342 (0.0002)
Constant	0.00589 (0.0006)	0.00992 (0.0016)	0.00539 (0.0003)	0.00841 (0.0005)
Observations	9,870	9,871	36,954	36,954
R-squared	0.734	0.547	0.657	0.531

The sample is described in the text. In both panels, the dependent variable is the hiring probability from the conditional logit model. A good or bad applicant is defined as one above or below the median for the set of worker characteristics, X times coefficients, β , excluding the control function and the log bid in the conditional logit model as described in the text. Whether the next applicant is better or worse, ordering applicants by application time, is coded using the same underlying measure of applicant quality. All standard errors are calculated using 50 block bootstrap replications to account for parameter uncertainty in forming $X\beta$. In Panel B, regression results are presented from models regressing the hiring probability on the measures of paired applicant quality along with employer-by-applicant fixed effects. Coefficients are additive relative to the constant baseline, which captures the hiring probability for a bad applicant paired with a worse subsequent applicant.

Table 4
Conference Board Online Job Postings Regression Results

	(1) Unfilled Jobs Ratio
Occupational Education * Year 2006	0.040*** (0.0070)
Occupational Education * Year 2007	0.038*** (0.0067)
Occupational Education * Year 2008	0.037*** (0.0065)
Occupational Education * Year 2009	0.036*** (0.0065)
Occupational Education * Year 2010	0.035*** (0.0063)
Occupational Education * Year 2011	0.034*** (0.0062)
Occupational Education * Year 2012	0.036*** (0.0063)
Occupational Education * Year 2013	0.037*** (0.0062)
Occupational Education * Year 2014	0.036*** (0.0062)
Constant	-0.044 (0.094)
N	846
R ²	0.4144

The dependent variable is the percent of online job postings that are unfilled in an average month by year. The Occupational Education level is the mean education level from CPS data for the 94 occupations. The education average for 2014 is imputed from 2013 data. Regression is weighted by the number of observations in each occupation, standard errors clustered by occupation. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5
oDesk Probability of Filling a Job

Dependent Variable: Indicator for Filling Vacancy

	All Job Categories			2 Largest
Mean Wage in Job Category	-0.0683*** (0.00538)			
90th Percentile of Wages in Category		-0.105*** (0.00631)		
Admin Support (Baseline)				
Design and Multimedia			-0.101*** (0.00805)	
Networking and IS			-0.145*** (0.0116)	
Sales and Marketing			-0.127*** (0.00846)	
Software Development			-0.171*** (0.0132)	
Web Programming			-0.100*** (0.00675)	-0.104*** (0.00818)
Writing and Translation			-0.0804*** (0.00772)	
Firm Effects	Y	Y	Y	Y
Time Effects	Y	Y	Y	Y
N Job Openings	217,753	217,753	217,496	132,226
N Firms	60199	60199	60130	41402

Standard errors clustered by employer. Customer and Business Service estimates not displayed. 64 percent of employers post vacancies in different categories in the same month.

Table 6
PSID Average Hourly Log Wage Gap and Average Difference Between Actual Log Wage and Mean Occupation Log Wage for Workers In Job Spells for Which they are Over- and Under-Qualified for Their Job

	Work Spells Corresponding to a Worker Being in a Job for Which He is Under-Qualified	Work Spells Corresponding to a Worker Being in a Job for Which He is Over-Qualified
Predicted Log Wage - Actual Log Wage		
Mean	0.008	0.078
(Std Error)	(0.008)	(0.008)
Fraction Above Zero	56.4	64.5
Number of Observations	2530	2657
Actual Log Wage - Mean Log Wage in Inappropriate Occupation		
Mean	-0.175	0.058
(Std Error)	(0.010)	(0.010)
Number of Observations	2732	2907

Data from the PSID. Sample sizes differ due to missing data for some regressions in the first panel. Predicted log wages come from a model with individual fixed effects estimated on 19,397 person-year observations for individuals in their usual occupation. There are 4385 unique individuals in the model.

Table 7
Rising Mismatch Over Time

	(1)	(2)
Year	770.4*** (73.6)	-273.5 (200.0)
OccupationalSkill	0.86*** (0.071)	-34.5*** (7.54)
OccupationalSkill*Year		0.018*** (0.0038)
Constant	-1543801.5*** (146782.0)	541478.8 (399728.7)
N	12733	12733
R ²	0.5209	0.5415

The sample size for the CPS data by Occupation is 39 years (1975-2013) times an average of 326 occupational groups per year. The number of occupations varies between years, as in earlier years the occupational categories were broader. For example, there are 293 occupations in 1975, compared to 373 in 1994. The dependent variable is the variance of income within occupational group by year. OccupationalSkill is the average education in that occupation. Regression observations are weighted by number of observations per occupation. In some years, there is only 1 observation for some occupations; these observations are dropped.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Appendix: First Stage Regression and Conditional Logit Parameter Estimates

Appendix Table 1 contains the results of the first stage regression. The instruments are estimated precisely and are quite strong

Appendix Table 1: First Stage Regression of Log Hourly Bids on Exchange Rate Instruments

Dependent Variable	Log Hourly Wage Bid (Job Opening Level)	
Model	Linear Regression	
Sample	Administrative Support	Web Programming
	(1)	(2)
Log Local Currency to Dollar (monthly; standardized)	-0.0365*** (0.00423)	-0.0134*** (0.00118)
Feedback Score (Out of 5)	0.0665*** (0.00383)	0.0662*** (0.00121)
Agency Affiliate Indicator	0.0327*** (0.00841)	0.0410*** (0.00209)
Agency Affiliate x Inexperienced Worker	0.0609*** (0.0134)	0.00231 (0.00333)
Prior Work Experience	-0.0529*** (0.00741)	0.117*** (0.00321)
Number of Observations	73,056	240,806
R-Squared	0.166	0.211
F Statistic on Excluded Instruments	74.47	129.8

Notes: The sample is experienced employers who have hired 2 or more previous workers from any job category and have posted at least 2 previous jobs in the job category in question. Robust standard errors in parentheses. All models contain a fifth order polynomial in calendar time, fixed effects for 6 country groups, controls for English skills, and an indicator for having zero feedback. The last country group includes many countries with very small application shares. Models in columns 1 - 4 also include a piecewise-linear spline with 4 knots for the application number, an indicator for an employer-initiated application, and an indicator that the worker only applies to this job during the month. The Log Local Currency to Dollar exchange rate is calculated using monthly data and z-scores are used to make the measure comparable across countries.

Appendix Table 2: Conditional Logit Parameter Estimates

	(1)	(2)	(3)	(4)
Sample:	Administrative Support	Administrative Support, With Control Function	Web Programming	Web Programming, With Control Function
Log Hourly Bid		-8.496 (2.508)		-18.29 (5.151)
Control Function		7.748 (2.52)		17.83 (5.15)
Feedback Score Out of 5	0.245 (0.083)	0.805 (0.191)	0.236 (0.038)	1.448 (0.353)
Agency Affiliate Indicator	-0.164 (0.159)	0.103 (0.181)	-0.468 (0.05)	0.293 (0.247)
Agency Affiliate x Inexperienced Worker	0.498 (0.247)	0.978 (0.302)	0.246 (0.111)	0.273 (0.146)
Indicator for Prior Experience on oDesk	1.970 (0.129)	1.483 (0.222)	1.755 (0.099)	3.882 (0.615)
Mean Own-Bid Elasticity		-8.414 (2.493)		-18.12 (5.122)
Number of Job Openings	2042	2042	8230	8230

The sample is described in the text. To be included, job openings must have had more than six applicants, with at least three worker-initiated applicants. Estimates come from a conditional logit model that includes an option not to hire an applicant. Standard errors come from using 50 block bootstrap replications of the entire procedure in which each job opening forms a block. Other controls are included for English score, an indicator for no feedback, an indicator for being experienced without having feedback, as well as the country dummies and applicant order spline as described in the table notes for the first stage regression.