Interviews and the Assignment of Workers to Jobs

Benjamin Lester*

Federal Reserve Bank of Philadelphia

Ronald Wolthoff

University of Toronto

May 30, 2012

Preliminary and Incomplete

Abstract

This paper studies the effect of screening costs on the equilibrium allocation of workers with different productivities to firms with different technologies. In the model, a worker's type is private information, but can be learned by the firm during a costly screening or interviewing process. We characterize the planner's problem in this environment and determine its solution. A firm may receive applications from workers with different productivities, but should in general not interview them all. Once a sufficiently good applicant has been found, the firm should instead make a hiring decision immediately. We show that the planner's solution can be decentralized if workers direct their search to contracts posted by firms. These contracts must include the wage that the firm promises to pay to a worker of a particular type, as well as a hiring policy which indicates which types of workers will be hired immediately, and which types will lead the firm to keep interviewing additional applicants.

1 Introduction

A central challenge to labor economists is to understand how labor markets allocate workers of different abilities to firms with different technologies. This challenge was first addressed within the context of a frictionless environment by Becker (1973), and has recently been re-examined within the context of environments that incorporate various frictions within the matching process (see Shi, 2001, 2002; Shimer, 2005; Peters, 2010 and Eeckhout & Kircher, 2010). However, all of these papers abstract from an important, realistic feature of modern labor markets: the screening or *interviewing* of applicants.

In this paper, we construct a model in which firms with different technologies post vacancies, and workers with different abilities observe the contracts being offered at these firms and apply for a single position. Following the literature on directed search, we assume that workers cannot coordinate their application strategies, so that multiple workers may apply to the same firm, and these workers may be of different abilities. Unlike the existing literature, however, we assume that a firm cannot distinguish one applicant from another *prima facie*. Instead, in order to discover a worker's ability, we assume that a firm must interview the applicant, which is a costly process. Given this cost, a firm may not choose to interview each applicant, but rather it may implement a hiring policy that indicates which types of workers will be hired immediately, and which types will lead the firm to keep interviewing additional applicants.

In this environment, we characterize the optimal wages and hiring policies of each firm, the optimal application behavior of workers, and the equilibrium allocation of workers to firms. We show that the equilibrium allocation is constrained efficient, i.e. a social planner would use the same hiring policy and would choose the same allocation of workers to firms.

The model provides testable implications with respect to various labor market outcomes and can be used to analyze to what extent these implications depend on the magnitude of the screening costs. For example, the model predicts how unemployment rates and wage dispersion vary with worker skill. Further, the model is informative of how an inflow of low-skilled workers into the pool of job seekers affects the equilibrium allocation and wages.

This paper proceeds as follows. Section 2 introduces the environment and explains how we model recruitment. We solve the planner's problem in section 3. In section 4, we derive the market equilibrium and show that it decentralizes the planner's solution. Section 5 concludes.

2   The Environment

We consider an economy populated by a measure 1 of workers and a measure v of firms. Each worker is endowed with a single, indivisible unit of labor, and each firm requires one such unit of labor to produce. Workers and firms are risk-neutral and heterogeneous with respect to productivity. In particular, each worker can be represented by a characteristic x which is distributed according to the cumulative density function $F(x)$, which has support on a closed and connected interval $\mathcal{X} \equiv [\underline{x}, \overline{x}] \subset \mathbb{R}_+$. A worker's type is private information. Like workers, firms differ in a characteristic y distributed according to $G(y)$ over the interval $\mathcal{Y} \equiv [\underline{y}, \overline{y}] \subset \mathbb{R}_+$. However, unlike workers, a firm's type is publicly observable. Both $F(x)$ and $G(y)$ are assumed to differentiable, with $f(x) = F'(x)$ and $g(y) = G'(y)$. Finally, a match between a worker of type x and a firm of type y yields $z(x, y)$ units of output. We assume that $z(x, y) \geq 0$ for all $(x, y) \in \mathcal{X} \times \mathcal{Y}$, $\partial z(x, y)/\partial x > 0$, and $\partial z(x, y)/\partial y > 0$.

Posted Contracts, Applications, and Queue Lengths

We consider a one-shot game that proceeds in stages. First, each firm posts and commits to a contract, which consists of two elements. The first element is a hiring policy $\mu \in \mathcal{X}$, which we describe in more detail below. The second element of the contract is a wage schedule $w(x)$, indicating the wage that the firm will pay if it hires a worker of type x. We focus our attention on equilibria in which firms play symmetric, pure strategies. Therefore, the contract of a type y firm can be summarized by a pair $\{\mu(y), w(x, y)\}$.

Workers observe the contracts posted by all firms and apply to a single firm. As is customary in models of directed search, we restrict attention to equilibria in which workers follow symmetric strategies: all workers of type x apply to (one of the) type y firms with the

same probability (and *each* type y firm with equal probability). As a result, the number of type x workers to apply to a particular type y firm is a random variable distributed according to the Poisson distribution.[1] We denote by $\lambda(x, y)$ the expected number, or *queue length*, of type x workers to arrive at a firm of type y, so that the probability that a firm of type y receives n applications from workers of type x is given by the usual $\left[\lambda(x, y)e^{-\lambda(x,y)}\right]/n!$.

When workers apply across different types of firms with positive probability, a firm will typically receive applications from workers of different productivity levels. However, since a worker's type is private information, the firm will not be able to distinguish these different types of workers *prima facie*. In order to overcome this informational asymmetry, we assume that a firm can learn a worker's type through a job interview. Job interviews are costly and must be conducted sequentially. In particular, we assume that firms interview candidates one at a time in a randomly assigned order and incur a cost k for each interview, with the exception of the first interview, which we assume is costless.[2] Since interviews are costly, firms will typically not want to interview all of their applicants. Instead, given the sequential nature of the interviewing process, firms will optimally choose a cutoff or hiring policy $\mu(y)$, which specifies that the firm will hire the first applicant whose type exceeds $\mu(y)$ (and not interview any remaining candidates). If none of the applicants are of type $\mu(y)$ or higher, the firm will interview all applicants and hire the most productive.[3]

Once the interviewing process is over and matches are formed, production occurs and wages are paid. A firm of type y who hires a worker of type x receives profit $\pi(x, y) = z(x, y) - w(x, y)$, while the worker receives wage $w(x, y)$. Firms that do not match produce

---

[1]This can be derived explicitly by analyzing the game with a finite number of workers and firms, and taking the limit as the number of agents tends to infinity; see, for example, Burdett et al. (2001) and Peters (2010).

[2]The assumption that the first interview is costless is convenient here. For example, if screening the first applicant is costly as well, firms with only one applicant would prefer to hire him without an interview. However, in that case the firm cannot pay a wage conditioned on the worker's type, as the type is unknown. More generally, our analysis would have to allow firms to post contracts that were conditional on the number of applications they received, which complicates the analysis considerably.

[3]Notice that this contracting space allows the firm considerable flexibility in shaping the distribution of workers that apply and the type of the worker that is ultimately hired. For example, a firm of type y that wants to discourage workers of type x' from applying (so as to avoid congestion at the interviewing stage) can specify a wage $w(x', y) = 0$. In addition, the choice of $\mu(y)$ allows the firm a variety of policies ranging from hiring an applicant at random ($\mu(y) = \underline{x}$) to hiring the best applicant ($\mu(y) = \overline{x}$).

zero units of output, and workers that are unmatched receive payoff zero.

Interviewing Costs

We now derive the expected interviewing costs incurred by a firm of type y with queue lengths $\{\lambda(x, y)\}$ who has chosen cutoff $\mu(y)$. In order to do so, it will be convenient to define $L(x, y) = \int_x^{\overline{x}} \lambda(x', y)dx'$ and $\Lambda(x, y) = \int_{\underline{x}}^x \lambda(x', y)dx'$ as the queue of applicants who are more or less productive, respectively, than a worker of type x. Given this notation, $\Lambda(\overline{x}, y) = L(\underline{x}, y)$ represents the total queue length of workers applying to a type y firm, while $\Lambda(\underline{x}, y) = L(\overline{x}, y) = 0$. In what follows, we will suppress the argument of $\mu(y)$ whenever possible for notational simplicity, though it should be understood that the hiring policy will, in general, depend on the firm's type.

Given a hiring policy μ, the workers that apply to a firm can be divided into two groups: those with type $x < \mu$ and those with type $x \geq \mu$. For lack of a better label, we refer to the first type of worker as a "bad" applicant and the latter as a "good" applicant. The queue lengths of the respective groups are $\Lambda(\mu, y)$ and $L(\mu, y)$.

To derive the expected interviewing costs incurred by the firm, consider the event that n total applications are received. Then the number of good applicants, n_G, conditional on n total applicants, is distributed according to the binomial distribution $B\left(n, \frac{L(\mu,y)}{L(\underline{x},y)}\right)$.[4] Similarly, conditional on n, the number of bad applicants, n_B, is distributed according to the binomial distribution $B\left(n, \frac{\Lambda(\mu,y)}{L(\underline{x},y)}\right)$. Therefore, given n, the expected interviewing costs can be written

$$\left\{\sum_{i=1}^n \left(\frac{\Lambda(\mu, y)}{L(\underline{x}, y)}\right)^{i-1} \frac{L(\mu, y)}{L(\underline{x}, y)}(i - 1) + (n - 1)\left(\frac{\Lambda(\mu, y)}{L(\underline{x}, y)}\right)^n\right\} k. \tag{1}$$

The first term in brackets in equation (1) captures the possibility that the first good applicant is discovered in interview $i \in \{1, \ldots, n\}$, in which case the firm will incur cost $(i - 1)k$ (since the first interview is costless). The second term in brackets captures the event that all n applicants are bad, in which case the firm incurs cost $(n-1)k$. Since n is distributed according to the Poisson distribution with parameter $L(\underline{x}, y)$, one can take expectations across all

---

[4]The proof of this result is standard and therefore omitted.

possible realizations of n. Doing so reveals that the total expected costs of interviewing for a firm of type y with queue lengths $\{\lambda(x, y)\}$ who has chosen μ is

$$K(\mu, y) = \left\{ \left[1 - e^{-L(\mu,y)} \right] \frac{L(\underline{x}, y)}{\Lambda(\mu, y)} - \left[1 - e^{-L(\underline{x},y)} \right] \right\} k. \tag{2}$$

Match Output and Social Surplus

In order to derive the total social surplus, we begin by constructing the expected output at a firm of type y with queue lengths $\{\lambda(x, y)\}$ and a hiring policy $\mu(y)$. Again, we suppress the argument of μ whenever possible for notational convenience. There are three relevant events that can occur for any firm after choosing a hiring policy μ: (1) at least one good applicant arrives; (2) no good applicants arrive but at least one bad applicant arrives; and (3) no applicants arrive.

Since $e^{-L(\mu,y)}$ is the probability that the firm receives no applicants with productivity $x \geq \mu$, it follows that the first event occurs with probability $1 - e^{-L(\mu,y)}$. In this event, recall that the firm continues to interview workers until it finds the first good applicant. In particular, if there are multiple good applicants, each of them is equally likely to be hired since the applicants are interviewed in a random order. Hence, the productivity of the worker that is hired is a random draw from the distribution of good workers. Given the properties of the Poisson distribution, this conditional distribution has density

$$\phi(x, y) = \frac{\lambda(x, y)}{\int_\mu^{\overline{x}} \lambda(x', y) dx'} = \frac{\lambda(x, y)}{L(\mu, y)} \tag{3}$$

defined over the range $[\mu, \overline{x}]$. As a result, the output of a firm of type y, conditional on receiving at least one good applicant, can be written

$$\int_\mu^{\overline{x}} \phi(x', y) z(x', y) dx' = \int_\mu^{\overline{x}} \frac{\lambda(x', y)}{L(\mu, y)} z(x', y) dx'. \tag{4}$$

Since the first event occurs with probability $1 - e^{-L(\mu,y)}$ and the third event occurs with probability $e^{-L(\underline{x},y)}$, the second event occurs with probability $1 - \left[1 - e^{-L(\mu,y)} + e^{-L(\underline{x},y)} \right] = e^{-L(\mu,y)} - e^{-L(\underline{x},y)}$. In this event, the firm interviews all applicants and hires the worker with the highest productivity. Thus, the productivity of the selected worker is drawn from the

density function[5]

$$\varphi(x, y) = \frac{\lambda(x, y)e^{-L(x,y)}}{e^{-L(\mu,y)} - e^{-L(\underline{x},y)}},\tag{5}$$

which is defined over the range $[\underline{x}, \mu]$. Therefore, conditional on receiving no good applicants and at least one bad applicant, the output of a firm of type y is equal to

$$\int_{\underline{x}}^{\mu} \varphi(x', y)z(x', y)dx' = \int_{\underline{x}}^{\mu} \frac{\lambda(x', y)e^{-L(x',y)}}{e^{-L(\mu,y)} - e^{-L(\underline{x},y)}} z(x', y)dx'.\tag{6}$$

By assumption, no output is produced when a firm receives zero applicants. Therefore, given the results above, it is straightforward to construct the total expected output at a firm with productivity y and hiring policy μ:

$$Z(\mu, y) = \left[1 - e^{-L(\mu,y)}\right] \int_{\mu}^{\overline{x}} \frac{\lambda(x', y)}{L(\mu, y)} z(x', y)dx' + \int_{\underline{x}}^{\mu} \lambda(x', y)e^{-L(x',y)} z(x', y)dx'.\tag{7}$$

Hence, the total social surplus, integrating across all firm types, can be written

$$S = \int_{\underline{y}}^{\overline{y}} g(y) \left[Z\left(\mu(y), y\right) - K\left(\mu(y), y\right)\right] dy,\tag{8}$$

where we've written μ explicitly as a function of y to remind the reader that hiring policies will, in general, depend on a firm's type.

3   The Social Planner's Problem

The social planner maximizes S by choosing the queue lengths of each type of worker at each type of firm, $\lambda(x, y)$, along with hiring policies $\mu(y)$ at each firm, subject to the constraint that, for all $x \in \mathcal{X}$,

$$f(x) = \int_{\underline{y}}^{\overline{y}} \lambda(x, y)g(y)dy.\tag{9}$$

Equation (9) is a typical resource constraint: it requires that the measure of workers of type x allocated across all firms must be equal to the measure of workers of type x available in the

---

[5]To understand this density function, let x_{max} denote the maximum value of x amongst the applicants. The probability that $x_{max} < x$, conditioning on the occurrence of the second event, is equal to $\frac{e^{-L(x,y)} - e^{-L(\underline{x},y)}}{e^{-L(\mu,y)} - e^{-L(\underline{x},y)}}$. Differentiating with respect to x yields the density function in equation (5).

pool of unemployed workers. Given (8) and (9), the Lagrangian of the optimization problem can be written

$$\mathcal{L} = \int_{\underline{y}}^{\overline{y}} g(y) \left[Z\left(\mu(y), y\right) - K\left(\mu(y), y\right) - \int_{\underline{x}}^{\overline{x}} \xi(x)\lambda(x,y)dx \right] dy + \int_{\underline{x}}^{\overline{x}} \xi(x)f(x)dx, \quad (10)$$

where $\xi(x)$ denotes the Lagrange multiplier for type x workers.

Optimal Hiring Standards

Given the queue lengths of type x workers at a firm of type y, for all $x \in \mathcal{X}$, the socially optimal hiring policy $\mu(y)$ satisfies the first order condition

$$0 \geq \left[\int_{\mu}^{\overline{x}} \frac{\lambda\left(x', y\right)}{L(\mu, y)} z\left(x', y\right) dx' - z\left(\mu, y\right) \right] - \frac{L\left(\underline{x}, y\right)}{L(\mu, y)} k, \quad (11)$$

along with the condition $\mu(y) \geq \underline{x}$, with complementary slackness. To understand this expression, consider the effect of a marginal increase from an initial hiring policy μ. First, note that this increase would only affect the outcome if a type μ worker would have been chosen under the initial hiring policy, whereas a type $x' > \mu$ worker will be hired under the new hiring policy. In this case, expected output increases by an amount equal to the first term in brackets on the right hand side of the inequality above: the difference between the expected output when a good worker is randomly drawn from $(\mu, \overline{x}]$ according to the density function ϕ, net of the output that would have been produced when a type μ worker was hired. However, given this increase in the hiring policy, the expected costs of interviewing also increase, which is reflected in the final term of the inequality in equation (11).

Optimal Queue Lengths

We turn now to the optimal queue lengths, taking as given the hiring policies at each firm. For each $(x, y) \in \mathcal{X} \times \mathcal{Y}$, the first order condition with respect to $\lambda(x, y)$ reduces to

$$\xi(x) \geq \left[e^{-L(x,y)} z(x,y) - \int_{\underline{x}}^{x} \lambda(x', y)e^{-L(x',y)} z(x',y)dx' \right] - k \left[\frac{1 - e^{-L(\mu,y)}}{L(\mu, y)} - e^{-L(\underline{x},y)} \right] \quad (12)$$

if $x < \mu(y)$, and

$$\xi(x) \geq e^{-L(\mu,y)} z(x,y) - \int_{\underline{x}}^{\mu} \lambda(x', y) e^{-L(x',y)} z(x',y) dx' - k \left[\frac{1 - e^{-L(\mu,y)}}{L(\mu, y)} - e^{-L(\underline{x},y)} \right] \quad (13)$$

$$+ \frac{1 - e^{-L(\mu,y)} \left[1 + L(\mu, y) \right]}{L(\mu, y)} \left[z(x,y) - \int_{\mu}^{\overline{x}} \frac{\lambda(x', y)}{L(\mu, y)} z(x',y) dx' + \frac{L(\underline{x}, y)}{L(\mu, y)} k \right] \quad (14)$$

8

if $x \geq \mu(y)$. In both cases, we require that $\lambda(x, y) \geq 0$, with complementary slackness.

The first order conditions on $\lambda(x, y)$ are fairly intuitive. For example, for the case of $x < \mu(y)$ and $\lambda(x, y) > 0$, equation (12) implies that the shadow value of adding an additional worker of type x to the applicant pool at a firm of type y is equal to the expected marginal gain in productivity (should the worker be hired) less the expected increase in interviewing costs.[6] To see this more clearly, note that (12) can be re-written in the following form (assuming $\lambda(x, y) > 0$), which decomposes the benefits and costs into four possible cases:

$$
\begin{aligned}
\xi(x) \;=\;& e^{-L(\underline{x},y)} \left[z(x, y)\right] \\
& + \left[e^{-L(x,y)} - e^{-L(\underline{x},y)}\right] \left[z(x, y) - \int_{\underline{x}}^{x} \frac{\lambda(x', y)e^{-L(x',y)}}{e^{-L(x,y)} - e^{-L(\underline{x},y)}} z(x', y)dx' - k\right] \\
& + \left[e^{-L(\mu,y)} - e^{-L(x,y)}\right] \left[-k\right] \\
& + \left[1 - e^{-L(\mu,y)}\right] \left\{\frac{1 - e^{-L(\mu,y)}\left[1 + L(\mu, y)\right]}{\left[1 - e^{-L(\mu,y)}\right] L(\mu, y)}\right\} \left[-k\right].
\end{aligned}
$$

First, it could be the case that no other worker applies to this firm, in which the introduction of a worker of type x yields a marginal gain of $z(x, y)$ and no cost, since the first interview is free. Second, it could be the case that at least one other worker applies to this firm, but that the worker of type x is the most productive. In this case, he will be interviewed with certainty, so that the interviewing cost k will be incurred with probability one. He will also be hired with certainty, and the expected net gain in output is $z(x, y)$ minus the expected output had this type x worker *not* applied; remember that $\varphi(x', y) = [\lambda(x', y)e^{-L(x',y)}]/\left[e^{-L(x',y)} - e^{-L(\underline{x},y)}\right]$ is the density of the maximum value of x'. The third possibility is that the worker is not the best worker to apply, but that there are no good applicants. In this case, the worker is interviewed with certainty, so that the cost k is incurred, but there is no marginal gain in output. Finally, it may be the case that a good worker applies to this firm. In this case, a worker with type $x < \mu(y)$ will never be hired, but he may or may not be interviewed. Indeed, the probability that a worker with $x < \mu(y)$ is interviewed, conditional on at least one good worker applying, is equal to

$$
\frac{1 - e^{-L(\mu,y)}\left[1 + L(\mu, y)\right]}{\left[1 - e^{-L(\mu,y)}\right] L(\mu, y)}.
$$

---

[6]Note that, at the optimum, the shadow value of adding an additional worker of type x must be equated across all firms with type y such that $\lambda(x, y) > 0$.

Similar reasoning can be used to understand the first order condition with respect to $\lambda(x, y)$ when $x \geq \mu(y)$. The first line of equation (13) reflects the expected marginal increase in output, less the increase in expected interviewing costs, from adding a worker of type $x \geq \mu(y)$ to a firm of type y when no other good workers apply. The second line of equation (13) reflects the expected benefit of this worker being hired when at least one other good worker also applies.[7] In this event, the expected change in output may be positive or negative; since the firm stops interviewing after meeting with the worker of type x, the productivity of the other good worker that would have been hired could be greater or less than x. However, in this event, the addition of the worker of type x to the applicant pool unambiguously decreases the expected costs of interviewing, which explains the final term in (13).

4   Decentralized Equilibrium

In this section, we characterize the decentralized market equilibrium and show that the queue lengths and interviewing policies in this equilibrium coincide constitute a solution to the social planner's problem. As we described earlier, in the decentralized equilibrium a firm of type y posts a contract a contract $\{\mu(y), w(x, y)\}$ to maximize profits, taking as given that this contract will attract a random number of type x workers that is distributed according to the Poisson distribution with mean (or expected queue length) $\lambda(x, y)$.

Optimal Application Behavior and Equilibrium Queue Lengths

The first step is to characterize how $\lambda(x, y)$ is determined. To do so, it is first convenient to define the expected payoff to a worker of type x from applying to a firm of type y, taking as given both the contract posted by the firm and the application behavior of the other workers. In particular, if $x < \mu(y)$, the worker is only hired if no better worker applies, which occurs with probability $e^{-L(x,y)}$. Alternatively, if $x \geq \mu(y)$, the worker is hired if he is selected from the pool of good applicants, which occurs with probability $\left[1 - e^{-L(\mu,y)}\right] / L(\mu, y)$. Thus, for

---

[7]Note that there is neither a cost nor a benefit of adding this worker to the pool of applicants if (i) at least one other good worker applies, and (ii) the worker of type $x \geq \mu(y)$ is not hired; in this case, both the output and the costs of interviewing are unchanged.

a type x worker, the expected payoff from applying to a type y firm can be written

$$V(x,y) = \begin{cases} e^{-L(x,y)}w(x,y) & \text{if } x < \mu \\ \frac{1-e^{-L(\mu,y)}}{L(\mu,y)}w(x,y) & \text{if } x \geq \mu. \end{cases} \tag{15}$$

Let us denote by $\overline{V}(x) = \sup_{y\in\mathcal{Y}} V(x,y)$ the maximal payoff, or *market utility*, of a type x worker. In a competitive search equilibrium, the application behavior of workers—and thus queue lengths—adjust so that workers are exactly indifferent between all firms to which they apply with strictly positive probability. That is, given some market utility $\overline{V}(x)$, the equilibrium queue length of type x workers at a type y firm must satisfy

$$\lambda(x,y) \geq 0 \tag{16}$$

$$\overline{V}(x) \geq V(x,y) \tag{17}$$

$$\lambda(x,y)\left[\overline{V}(x) - V(x,y)\right] = 0. \tag{18}$$

Optimal Posted Contracts and Equilibrium Profits

Consider a firm of type y that has posted a contract $\{\mu(y), w(x,y)\}$ and expects queue lengths $\{\lambda(x,y)\}$. Recall that the firm obtains a payoff $\pi(x,y) = z(x,y) - w(x,y)$ when it matches with a worker of type x, and incurs expected interviewing costs given by $K(\mu,y)$. Hence, the firm's expected profits can be derived using the same logic used in the derivation of social surplus, yielding

$$\begin{aligned} \Pi &= \int_{\underline{x}}^{\mu} \lambda(x,y)e^{-L(x,y)}\pi(x,y)\,dx + \left[1 - e^{-L(\mu,y)}\right]\int_{\mu}^{\overline{x}} \frac{\lambda(x,y)}{L(\mu,y)}\pi(x,y)\,dx - K(\mu,y) \\ &= Z(\mu,y) - K(\mu,y) \\ &\quad - \int_{\underline{x}}^{\mu} \lambda(x,y)e^{-L(x,y)}w(x,y)\,dx - \left[1 - e^{-L(\mu,y)}\right]\int_{\mu}^{\overline{x}} \frac{\lambda(x,y)}{L(\mu,y)}w(x,y)\,dx. \end{aligned} \tag{19}$$

Firms are infinitesimally small; they take as given the behavior of other firms, and thus the market utility of workers. Hence, given $\overline{V}(x)$, equations (16)–(18) define an implicit map between the contracts that firms post and the expected number of applicants of type x that apply to their firm. In particular, a firm of type y with hiring policy $\mu(y)$ that wishes to attract a queue $\lambda(x,y) > 0$ of workers of type x must offer a wage $w(x,y)$ that yields exactly the expected utility $\overline{V}(x)$. If it offers less, it will attract no applicants, $\lambda(x,y) = 0$. Taking

this relationship as given (i.e., as a constraint), we can use (15) and (18) to substitute the wages out of (19) and re-write the firm's objective function as

$$\Pi = Z(\mu, y) - K(\mu, y) - \int_{\underline{x}}^{\overline{x}} \lambda(x, y) \overline{V}(x) dx, \tag{20}$$

where now the firm maximizes with respect to $\mu(y)$ and $\lambda(x, y)$ for all x. Notice that (20) resembles a standard profit maximization problem in which the firm "purchases" queues of type x workers at the "price" $\overline{V}(x)$.

Decentralized Equilibrium and the Planner's Solution

Definition 1. *A (symmetric strategy) directed search equilibrium is characterized by contracts $\{\mu(y), w(x, y)\}$ for each $y \in \mathcal{Y}$, queue lengths $\{\lambda(x, y)\}$ for each $(x, y) \in \mathcal{X} \times \mathcal{Y}$, and market utilities $\{\overline{V}(x)\}$ for each $x \in \mathcal{X}$ such that:*

1. ***Firms behave optimally:*** *for each $y \in \mathcal{Y}$, given $\{\overline{V}(x)\}$ and (16)–(18), the contract $\{\mu(y), w(x, y)\}$ maximizes expected profits Π.*

2. ***Workers behave optimally:*** *for each $(x, y) \in \mathcal{X} \times \mathcal{Y}$, given $\{\overline{V}(x)\}$ and $\{\mu(y), w(x, y)\}$, queue lengths $\{\lambda(x, y)\}$ satisfy (16)–(18).*

3. ***Markets clear:*** *for each $x \in \mathcal{X}$, $\{\overline{V}(x)\}$ is such that $f(x) = \int_{\underline{y}}^{\overline{y}} \lambda(x, y) g(y) dy$.*

We will now establish that a decentralized equilibrium also constitutes a solution to the planner's problem, so that competitive search equilibria are constrained efficient. To see this, let us denote by $\{\lambda^*(x, y)\}$ and $\{\mu^*(y)\}$ the queue lengths and hiring policies, respectively, in a competitive search equilibrium. Also, let $\{\hat{\lambda}(x, y)\}$ and $\{\hat{\mu}(y)\}$ denote an alternative *feasible* solution to the planner's problem; that is, let $\{\hat{\lambda}(x, y)\}$ be such that

$$f(x) = \int_{\underline{y}}^{\overline{y}} \hat{\lambda}(x, y) g(y) dy$$

for each $x \in \mathcal{X}$. Since each type y firm maximizes (20) in a competitive search equilibrium, it must be that

$$Z\left(\lambda^*(x, y), \mu^*(y), y\right) - K\left(\lambda^*(x, y), \mu^*(y), y\right) - \int_{\underline{x}}^{\overline{x}} \lambda^*(x, y) \overline{V}(x) dx \geq$$
$$Z\left(\hat{\lambda}(x, y), \hat{\mu}(y), y\right) - K\left(\hat{\lambda}(x, y), \hat{\mu}(y), y\right) - \int_{\underline{x}}^{\overline{x}} \hat{\lambda}(x, y) \overline{V}(x) dx.$$

Integrating across all firms and simplifying yields

$$Z\left(\lambda^*(x,y),\mu^*(y),y\right) - K\left(\lambda^*(x,y),\mu^*(y),y\right) \geq Z\left(\hat{\lambda}(x,y),\hat{\mu}(y),y\right) - K\left(\hat{\lambda}(x,y),\hat{\mu}(y),y\right)$$
(21)

since

$$\int_{\underline{x}}^{\overline{x}} \overline{V}(x) \int_{\underline{y}}^{\overline{y}} \lambda^*(x,y)g(y)dydx = \int_{\underline{x}}^{\overline{x}} \overline{V}(x) \int_{\underline{y}}^{\overline{y}} \hat{\lambda}(x,y)g(y)dydx = \int_{\underline{x}}^{\overline{x}} \overline{V}(x)f(x)dx.$$

Hence, a competitive search equilibrium is a solution the planner's problem, with $\xi(x) = \bar{V}(x)$ for each x. We summarize this in the following proposition.

Proposition 1. *A directed search equilibrium exists and it is constrained efficient.*

5   Conclusion

This paper studies the effect of screening costs on the equilibrium allocation of workers with different productivities to firms with different technologies. The novel element of the paper is that a worker's type is private information, but can be learned by the firm during a costly screening or interviewing process. We characterize the planner's problem in this environment and determine its solution. We find that in general workers with different productivities may apply to the same firm. The firm should not interview all its applicants: once a sufficiently good applicant has been found, the expected gain in match surplus from interviewing additional applicants is smaller than the interviewing cost. We show that the planner's solution can be decentralized if workers direct their search to contracts posted by firms. These contracts must include two components: (i) the wage that the firm promises to pay to a worker of a particular type and (ii) the firm's hiring policy which indicates which types of workers will be hired immediately, and which types will lead the firm to keep interviewing additional applicants.

References

Becker, G. S. (1973). A theory of marriage: Part I. *Journal of Political Economy*, 81(4), 813–846.

Eeckhout, J. & Kircher, P. (2010). Sorting and decentralized price competition. *Econometrica*, 78, 539–574.

Peters, M. (2010). Unobservable heterogeneity in directed search. mimeo.

Shi, S. (2001). Frictional assignment. I. Efficiency. *Journal of Economic Theory*, 98, 232–260.

Shi, S. (2002). A directed search model of inequality with heterogeneous skills and skill-biased technology. *Review of Economic Studies*, 69(2), 467–491.

Shimer, R. (2005). The assignment of workers to jobs in an economy with coordination frictions. *Journal of Political Economy*, 113(5), 996–1025.