

The Assignment of Workers to Jobs In an Economy with Coordination Frictions

Robert Shimer*

Department of Economics

University of Chicago

shimer@uchicago.edu

First Version: September 7, 2001

Current Version: August 31, 2004

Abstract

This paper studies the assignment of heterogeneous workers to heterogeneous jobs in a large anonymous environment. Firms can commit *ex ante* to wage offers and workers can direct their job search accordingly. Anonymity requires that firms must offer identical workers the same wage and that identical workers must use the same mixed strategy in deciding where to apply for a job. The randomness induced by the realization of the mixed strategies generates coordination frictions, since two identical workers may happen to apply for a particular job, while an identical job gets no applications. In equilibrium, firms choose to attract applications from multiple types of workers, earning higher profits when they are able to hire a more productive worker. Identical workers apply for multiple types of jobs and get higher wages when they obtain a more productive job. Despite the resulting mismatch, I show that the model can generate assortative matching, with a positive correlation between matched workers' and firms' types. I also prove a version of the welfare theorems: the decentralized equilibrium maximizes the value of output in the economy given the anonymity restriction and the resulting coordination frictions.

*I am grateful to Jaap Abbring, Daron Acemoglu, James Heckman, Ian King, Shouyong Shi, and Chris Sims for useful discussions, and to Fernando Alvarez and three anonymous referees for their detailed reports. I have also benefited from comments by numerous seminar participants. I thank the National Science Foundation and the Sloan Foundation for financial support, and Sebastian Ludmer for excellent research assistance.

1 Introduction

This paper analyzes a large anonymous labor market in which aggregate output is affected by which worker is assigned to which job. Koopmans and Beckmann (1957) and Shapley and Shubik (1972) first explored assignment models, implicitly assuming that all workers and firms can communicate with each other. They obtained a powerful characterization of the core, i.e. the set of Pareto optimal assignments.¹ Identical workers earn the same wage, even if they take different types of jobs. Identical firms earn the same profits, even if they hire different types of workers. If workers' and firms' characteristics are complements in production, a more productive worker always has a better job than a less productive one (Becker 1973). Unemployment and vacant jobs cannot coexist. If there is unemployment, only the least productive workers are unemployed, while if there are vacancies, only the least productive firms fail to hire a worker.

These predictions are inconsistent with existing empirical evidence. Consider a panel data regression of wages on an individual fixed effect, any time-varying worker characteristics, and on some measure of the quality of the worker's job. The individual fixed effect should soak up observable or unobservable time-invariant individual characteristics, and so the textbook assignment model predicts that the quality of the worker's job should not affect her wage. On the contrary, the data indicate that workers in better quality jobs consistently earn higher wages (Krueger and Summers 1988, Gibbons and Katz 1992, Abowd, Kramarz, and Margolis 1999). Similarly, firms' profits appear to depend on the quality of their employees. The other stark predictions fail as well. Although on average more productive workers have better jobs than less productive ones, that is not true on a case-by-case basis. Unemployment and vacancies coexist, high productivity workers are sometimes unemployed, and high productivity jobs sometimes go unfilled.

This paper argues that the introduction of a small restriction on behavior, motivated by the anonymity of a large market economy, yields an assignment model that is qualitatively consistent with the facts. I consider an economy consisting of a continuum of risk-neutral workers, each described by one of M types, and a continuum of risk-neutral firms, divided into N types. These agents interact in a simple environment. Firms commit *ex ante* to type-contingent wage offers. Workers observe those offers and apply for one job. Each job that receives at least one application hires one worker; those that receive multiple applications

¹Sattinger (1993) and Roth and Sotomayor (1990) provide comprehensive reviews of the literature on matching with transferable and non-transferable utility, respectively.

choose which worker to hire. A type n firm that employs a type m worker pays the contracted wage and produces output $x_{m,n}$. Unmatched workers and firms are unemployed and vacant, respectively, producing nothing.

This model exhibits many equilibria. In a typical one, one type m worker applies to a particular type n firm, while another identical worker applies to a different type n firm. If such an equilibrium exists, there is another equilibrium in which the workers' roles are reversed, with the first worker applying for the second job and vice versa. But how could each worker know which of the multitude of equilibria is being played? While this may be an equilibrium in the sense of Nash—given the behavior of all other workers and firms, no one benefits by behaving differently—the presumed coordination seems implausible in a large market economy.

I use two formal assumptions to model the anonymity that seems inherent to a large market economy. First, firms' wage offers may be conditioned on a worker's type but not on her individual identity. This prevents firms from creating jobs tailored to a particular individual. Second, in equilibrium, identical workers must use identical (mixed) strategies when deciding where to apply for a job. That is, if one type m worker applies with probability 1 to a particular type n job, then all type m workers must apply with probability 1 to that job. That will not be an equilibrium, since only one worker would be hired and the remaining workers would stay unemployed. Instead, type m workers use mixed strategies, applying with probability $p_{m,n}$ to *some* type n firm, and with equal probability to *each* type n firm. The eponymous coordination friction is introduced by the random realization of these mixed strategies. Some type n firms will get multiple applications from type m workers, while others receive none. Given the first restriction that firms treat identical workers identically, the second restriction that identical workers behave identically seems like a plausible limitation on workers' behavior.

Section 2 develops the basic model and notation and discusses the anonymity restrictions at length. Section 3 turns to a social planner's problem in order to explore what is feasible in this environment. More precisely, I consider a hypothetical social planner who wishes to maximize the expected output in the economy. The planner is subject to an anonymity restriction that is similar to the one in the decentralized economy: he cannot tell a particular type m worker to apply to a particular type n firm, but instead must tell all type m workers to apply for some type n job with probability $p_{m,n}$. I provide a set of necessary and sufficient first order conditions that, together with a resource constraint, provide a concise characterization of the unique Social Optimum.

Section 4 shows that the Social Optimum is decentralized if firms can post wages in an effort to attract applicants, as in a Competitive Search Equilibrium (Moen 1997, Shimer 1996). This yields a version of the First and Second Welfare Theorems. Section 5 develops a set of empirical predictions regarding the cross-sectional behavior of wages and profits. I prove that if workers' and firms' types are complements in production, wages are an increasing function of a firm's type after conditioning on the worker's type. The reason workers do not always apply for high wage jobs is that such jobs are more difficult to get. Thus high wages are a compensating differential for high unemployment risk, and the observed correlation between a worker's wage and the quality of her job is an example of sample selection bias, since econometricians do not observe workers who applied for, but failed to get, high wage jobs. Similarly, I prove that a firm's profit is increasing in the quality of its employee after conditioning on the firm's characteristics. On the other hand, more productive workers do not necessarily earn higher wages after conditioning on the firm's characteristics. Instead, firms may compensate high productivity workers primarily through a higher hiring rate, rather than higher compensation conditional on hiring.

Section 6 develops a special case in considerable detail. I restrict attention to production functions in which no worker has a comparative advantage at any firm. That is, if worker m_1 produces twice as much with firm n_2 as with firm n_1 , then so does any other worker n_2 . I derive a closed-form solution for equilibrium and optimal application decisions. Workers apply for all jobs above a type-contingent threshold, an increasing function of the worker's type. Moreover, workers are equally likely to apply for all jobs above this threshold.² This has several significant implications. First, there is considerable overlap between different workers' application decisions. Given the randomness inherent in mixed strategies, some high productivity firms are forced to hire low productivity workers while some low productivity firms are able to hire higher productivity workers, a phenomenon that I call a 'mismatch' of workers and jobs. Second, despite this mismatch, there are patterns in matching behavior. A more productive worker is relatively more likely to match with a high productivity firm than a low productivity firm, compared to a less productive worker. In particular, there is a positive correlation between a worker's type and the type of her job, although that correlation is less than 1. Finally, I show that the model without comparative advantage is sufficiently tractable so as to be amenable to simple comparative statics exercises. Section 7

²To be precise, a worker is equally likely to apply for any job above the threshold of the next most productive worker. She is less likely to apply for jobs that lie between her threshold and the next most productive worker's threshold.

briefly describes the related literature before I conclude in Section 8.

2 Model

2.1 Participants

There are two kinds of risk-neutral agents in the economy, workers and firms. Workers are divided into M different types, $m = 1, \dots, M$. Let $\mu_m > 0$ denote the exogenous measure of type m workers. Each worker has a distinct ‘name’: for $i \in [0, \mu_m]$, (m, i) is the name of a particular type m worker. Similarly, there are N different types of firms, $n = 1, \dots, N$, with $\nu_n > 0$ denoting the exogenous measure of type n firms and (n, j) , $j \in [0, \nu_n]$, the name of a particular type n firm. Each firm has one job opening, and I refer interchangeably to a ‘firm’ and a ‘job’. There is no necessary relationship between the number of worker and firm types M and N , nor between the measures of workers and firms $\sum_{m=1}^M \mu_m$ and $\sum_{n=1}^N \nu_n$.

2.2 Production

Workers and firms match in pairs. Define $x_{m,n}$ to be the output that a type m worker and type n firm produce when matched, hereafter the production function. An unmatched agent produces nothing and a firm cannot match with more than one worker, nor a worker with more than one firm. More generally, one can view $x_{m,n}$ as the output produced by a type m worker and type n firm in excess of what they would get while single. For this reason, I assume throughout that x is nonnegative. I also impose that it is increasing in each argument. Finally, it is convenient to define $x_{0,n} \equiv 0$ for all n .

2.3 Wage Posting Game

The interaction between workers and firms can be represented as a three stage game. First, each firm (n, j) makes a wage offer to each worker (m, i) . Then each worker observes all the wage offers and applies for one job. Finally, firms that receive at least one application hire exactly one worker, pay the promised wage, and produce. Workers who are not hired are unemployed and jobs that are unfilled are vacant. In equilibrium, firms’ wage offers are optimal given other firms’ offers and workers’ application strategies, workers’ applications strategies in each subgame are optimal given the wage offers and all other workers’ application strategies, and firms hire the most productive applicant.

This game exhibits infinitely many equilibria, including the frictionless assignment.³ All jobs offer all workers their wage in the competitive equilibrium, and each worker looks for a different job, with the same assignment as in the competitive equilibrium. For reasons that I discuss below in Section 2.4, I preclude this possibility by requiring that firms' wage offers and hiring decisions depend only on workers' types, and that in every subgame, workers with the same type use the same payoff-maximizing application strategies.⁴ Since the frictionless assignment requires that identical workers apply for different jobs, it is ruled out by this anonymity restriction.

2.4 Discussion

There are two critical assumptions in this model: identical workers use identical mixed application strategies in every subgame; and although there are many types of workers, there are also many workers of each type.⁵ This section briefly discusses the role and plausibility of each of those assumptions in turn.

To understand the restriction that identical workers use identical mixed application strategies, it helps to consider the following static two player game. Each player simultaneously announces 'heads' or 'tails'. If both players make the same announcement, they both get a payoff of zero. If they make different announcements, they both get a payoff of one. This game has three equilibria. In one, player 1 announces 'heads' and player 2 announces 'tails'; in the second, the roles are reversed; and in the third, both players use identical mixed strategies, announcing 'heads' half the time. If there is pre-play communication, it is likely that the players will coordinate on one of the first two equilibria. But without pre-play communication, i.e. in an anonymous market, neither of those equilibria seems particularly plausible. From the perspective of player 1, it may seem equally likely that player 2 will play 'heads' or 'tails', and so player 1 might as well flip a coin to decide which announcement to make, and similarly for player 2. In other words, the mixed strategy

³For details in a related model, see Burdett, Shi, and Wright (2001).

⁴The requirement that firms' wage offers and hiring decisions depend only on workers' types is a restriction on the strategy space. The requirement that workers with the same type use the same application strategies is a refinement on the set of equilibria. The second restriction would not make sense without the first. Conversely, the first restriction by itself does not substantially reduce the set of equilibria. In particular, the frictionless assignment is an equilibrium even if firms' wage offers and hiring decisions depend only on workers' types.

⁵The assumption that there is a discrete number of worker and firm types is not important for the results. It is possible to redo the analysis in this paper with a continuum of worker and firm types, as long as there is a continuum of workers and firms of each type; none of the main results change, but the analysis with a double continuum is conceptually more complicated.

equilibrium, in which half the time the players make the same announcement, may be a good prediction of how players actually behave. Ochs (1990) and Cason and Noussair (2003) conduct laboratory experiments that support this conclusion. Even with small numbers of subjects, it appears as if the ‘workers’ employ mixed strategies.

The focus on the mixed strategy equilibrium is not novel to the wage posting literature. Montgomery (1991) first made this assumption explicitly in the context of a model with two workers and two jobs, writing “While in the simple 2×2 case presented above a pure-strategy equilibrium may seem more likely, this implies coordination on the part of applicants... But in a large labor market with many openings and many applicants, such coordination becomes nearly impossible” (p. 167). Burdett, Shi, and Wright (2001) explore the full set of equilibria in the case of $m = 2$ buyers (in the context of this paper, workers) and $n = 2$ sellers (jobs) in great detail, but conclude that “all of these pure-strategy equilibria require a lot of coordination, in the sense that every buyer has to somehow know where everyone else is going. This may not be so unreasonable when $n = m = 2$, but it seems hard to imagine for general n and m , which is what we want to consider below” (p. 1066).

In contrast, much less information is needed to play a mixed strategy equilibrium. Each type m worker must know the expected income of all types of workers $m' \geq m$ who are at least as productive as she is, $v_{m'}$ in this paper’s notation. Proposition 2 shows that this is uniquely determined by the economic environment. Using this, she can compute the difficulty of obtaining a job at a particular firm (n, i) conditional on the wage offers that the firm makes to type $m' \geq m$ workers. She then randomly selects one job which gives her the maximum expected income, using probability weights that are consistent with the computed difficulty of obtaining jobs at each of the firms. If all type m workers behave in the same manner, her behavior is indeed a best response. This seems like a reasonable prediction of how this game would in fact be played.

The second important assumption is that there are many workers of each type. If no two workers were identical, the anonymity restriction that “workers with the same type use the same payoff-maximizing application strategies” would be vacuous. Coles and Eeckhout (2000) show in a two-worker, two-firm example that with enough heterogeneity,⁶ there is a unique pure strategy equilibrium. In the real world, workers differ in their genetic characteristics, education, and experiences, so no two workers are exactly the same. This might suggest that the results in this paper are practically irrelevant.

Instead, it seems that the pure strategy equilibrium that emerges from an environment in

⁶The critical assumption is that aggregate output depends on which worker matches with which firm.

which no two workers are identical pushes the information structure of the assignment game too hard. Shimer (2003) shows, in the spirit of Harsanyi's (1973) purification argument, that in the presence of asymmetric information, heterogeneity can generate a unique equilibrium with coordination frictions. I assume that although each worker knows her ability at each job, she does not know the other workers' ability and hence does not know her comparative advantage. I show that if comparative advantage is sufficiently important, there is a unique equilibrium in which all workers use the same application strategy, a function of their comparative advantage but not their identity. The equilibrium is indistinguishable from the mixed strategy equilibrium in the simplest model with homogeneous workers and symmetric information. In other words, for the coordination friction in this paper to disappear, my preliminary research suggests that we require that (i) no two workers are identical and (ii) all workers know the characteristics of all the other workers. Montgomery (1991), Peters (1991), Burdett, Shi, and Wright (2001), and this paper violate condition (i), while Shimer (2003) violates condition (ii).

3 Social Planner's Problem

Before analyzing the symmetric equilibrium of this game, it is useful to consider a related centralized assignment problem, hereafter referred to as the Social Planner's Problem. A social planner wishes to maximize the value of output in this economy, a utilitarian welfare function given the risk-neutrality of workers and firms. To achieve that objective, the planner provides workers with instructions on where to apply for a job and firms with instructions on which applicant to hire. Each worker can make only one job application and each firm may hire only one applicant.

3.1 Job Applications

An unconstrained social planner can assign each worker (m, i) to the desired job (n, j) . I introduce coordination frictions in the planner's problem via an assumption that he must treat identical workers and firms identically. That is, the planner can tell type m workers to apply with probability $p_{m,n} \geq 0$ to some type n firm, where $\sum_{n=1}^N p_{m,n} = 1$, but cannot otherwise distinguish between workers and firms according to their names. More precisely, each type m worker is equally likely to apply to any of the type n firms, so a type n firm on

average receives

$$q_{m,n} \equiv \frac{\mu_m p_{m,n}}{\nu_n} \quad (1)$$

applications from type m workers. Moreover, assuming the realization of these stochastic applications is independent across workers, the actual number of such applications is a Poisson random variable. Each type n firm receives $z \in \{0, 1, 2, \dots\}$ applications from type m workers with probability $\frac{1}{z!} q_{m,n}^z e^{-q_{m,n}}$. In a standard abuse of the law of large numbers, I assume that this in fact represents the proportion of type n firms that receives z applications from type m workers. Note that as long as $q_{m,n}$ is finite, a positive fraction of type n firms receives no type m applications. I refer to $q_{m,n}$ as the (expected) queue of type m workers for a job n . It will also be convenient to define the (expected) queue of at-least-as-productive job applicants:

$$Q_{m,n} \equiv \sum_{m'=m}^M q_{m',n}, \quad (2)$$

with $Q_{M+1,n} \equiv 0$.

3.2 Hiring and Output

Since the production function $x_{m,n}$ is increasing in m , the planner optimally instructs firms to hire the most productive applicant. This implies that a type n firm employs a type m worker if it does not receive any applications from type $m' \in \{m+1, \dots, M\}$ workers and it receives at least one application from a type m worker. Given the expected queue lengths defined above, such an event occurs with probability $e^{-Q_{m+1,n}}(1 - e^{-q_{m,n}})$, in which case the firm produces output $x_{m,n}$. Aggregate output is the product of the hiring probabilities and the output that a matched worker-firm pair produces, summed across worker and firm types:

$$Y(q) = \sum_{n=1}^N \nu_n \sum_{m=1}^M e^{-Q_{m+1,n}} (1 - e^{-q_{m,n}}) x_{m,n} \quad (3)$$

The planner maximizes output by choosing nonnegative queue lengths $q_{m,n}$ that satisfy a resource constraint, combining the requirement that the application probabilities $p_{m,n}$ sum to one with the definition of q in equation (1):

$$\mu_m = \sum_{n=1}^N q_{m,n} \nu_n \text{ for all } m \quad (4)$$

I refer to the solution to this constrained optimization problem as the Social Optimum.

3.3 Characterization

Let v_m be the multiplier on constraint (4) and write the Lagrangian as

$$\mathcal{L}(q, v) = \sum_{n=1}^N \nu_n \sum_{m=1}^M \left(e^{-Q_{m+1,n}} (1 - e^{-q_{m,n}}) x_{m,n} - v_m q_{m,n} \right) + \sum_{m=1}^M v_m \mu_m. \quad (5)$$

Paying attention to the non-negativity constraints on $q_{m,n}$, a typical first order condition is

$$v_m \geq e^{-Q_{m,n}} x_{m,n} - \sum_{m'=1}^{m-1} e^{-Q_{m'+1,n}} (1 - e^{-q_{m',n}}) x_{m',n} \text{ and } q_{m,n} \geq 0, \quad (6)$$

with complementary slackness. In words, type m workers should apply for type n jobs only if their marginal product is highest at these jobs, where their marginal product is defined as the additional output that a particular type n job produces in excess of what it would have produced without the application from a type m worker.

This interpretation is most easily understood for type $m = 1$ workers. If a more productive worker applies for the same job, a type 1 worker will not be hired and so produces nothing. Even if another type 1 worker applies for the job but this particular worker is hired, the worker's marginal product is still zero, since the firm would have produced just as much output without this application. In other words, the marginal product of a type 1 worker applying for a type n job is just equal to the probability that no other worker applies for the same job, $e^{-Q_{1,n}}$, times the output produced by a type 1 worker in a type n job, $x_{1,n}$.

More generally, the marginal product of a type m worker reflects the possibility that if the worker had not applied for the job, the firm might have employed a less productive worker. If $m = 2$, the right hand side of condition (6) is equal to the probability that the firm does not get an application from a type 2 or better worker, $e^{-Q_{2,n}}$, times the output produced by a type 2 in a type n job, $x_{2,n}$; minus the probability that the firm would have employed a type 1 worker in the absence of this application, $e^{-Q_{1,n}}$, times the output produced by a type 1 worker in a type n job, $x_{1,n}$.

I prove in the appendix that condition (6) is necessary and sufficient for an optimum because the planner's problem is convex. To summarize:

Proposition 1. *Any queue lengths $\{q_{m,n}\}$ and shadow values $\{v_m\}$ satisfying the feasibility*

constraint (4) and the complementary slackness condition (6) is socially optimal. The Social Optimum is unique.

4 Decentralization

This section shows that the Social Optimum can be decentralized as a Competitive Search Equilibrium (Moen 1997, Shimer 1996).

4.1 Competitive Search Equilibrium

In a Competitive Search Equilibrium, each firm (n, j) posts type-contingent wage offers $w_{m,(n,j)}$.⁷ Each worker observes all the wage offers and applies for one job. Firms that receive at least one application hire one worker, pay the promised wage, and produce. Workers who are not hired are unemployed and jobs that are unfilled are vacant.

I look for an equilibrium in which firms always hire the most productive applicant.⁸ Moreover, I impose the anonymity restriction that identical workers use identical mixed strategies in the second stage of the game. Firm (n, j) anticipates that if it offers a wage schedule $\{w_{m,(n,j)}\}$, it will attract on average a queue of $q_{m,(n,j)}$ type m workers, although the realized number will be a Poisson random variable, as in the centralized economy.

The critical question is how those queues are determined. In a Competitive Search Equilibrium, workers adjust their application strategies so that they are indifferent between applying for this job or their best alternative job. More precisely, let v_m denote a type m worker's expected income at her best alternative job. Since firm (n, j) is infinitesimal, it believes that its wage offer does not affect v_m ,⁹ although in equilibrium these values are determined to clear the market for applications. Instead, if $q_{m,(n,j)}$ is positive, type m workers

⁷This notation already embeds part of the anonymity restriction. A firm cannot tailor the wage contract to a particular worker, i.e. offer identity-contingent wage contracts $\tilde{w}_{(m,i),(n,j)}$.

⁸This must be the case if the equilibrium decentralizes the Social Optimum. Proposition 4 confirms that the most productive applicant is in fact the most profitable applicant, i.e. the applicant with the maximal value of $x_{m,n} - w_{m,(n,j)}$, so firms want to hire the most productive applicant. An earlier version of this paper formally establishes that there is no equilibrium in which firms rank applicants differently (Shimer 2001). Because the notation is cumbersome, that proof is omitted from the current version of the paper; however, the text before equation (13) provides some intuition for the result.

⁹Burdett, Shi, and Wright (2001) prove that in a homogeneous agent economy, the 'price-taking' approximation is correct if the number of workers and firms is sufficiently large.

must get the same expected income from applying to firm (n, j) as their next best alternative,

$$v_m = e^{-Q_{m+1,(n,j)}} \frac{1 - e^{-q_{m,(n,j)}}}{q_{m,(n,j)}} w_{m,(n,j)}, \quad (7)$$

where $e^{-Q_{m+1,(n,j)}} \frac{1 - e^{-q_{m,(n,j)}}}{q_{m,(n,j)}}$ is the probability that the worker is hired¹⁰ and $w_{m,(n,j)}$ is her wage if she is hired. Alternatively, if $v_m > e^{-Q_{m+1,(n,j)}} w_{m,(n,j)}$, type m workers do not apply for the job, $q_{m,(n,j)} = 0$, because applying elsewhere gives higher expected utility even if no other type m workers apply for the job. These optimality requirements both imply

$$q_{m,(n,j)} v_m = e^{-Q_{m+1,(n,j)}} (1 - e^{-q_{m,(n,j)}}) w_{m,(n,j)}. \quad (8)$$

For given values of $\{v_m\}$, this equation uniquely determines the queue lengths as a function of firm (n, j) 's wage offers. First solve for $q_{M,(n,j)}$ as a function of $w_{M,(n,j)}$ and then inductively compute $q_{m,(n,j)}$ as a function of $w_{m,(n,j)}$ and $Q_{m+1,(n,j)}$. At each step, look first for a positive solution, $q_{m,(n,j)} > 0$; if none exists, use $q_{m,(n,j)} = 0$ to solve the equation.

Given this determination of the queues, firm (n, j) 's expected profit is the product of the probability that it hires a type m worker, $e^{-Q_{m+1,(n,j)}} (1 - e^{-q_{m,(n,j)}})$, times the profit it gets when it does so, $x_{m,n} - w_{m,(n,j)}$, summed across worker types:

$$\sum_{m=1}^M e^{-Q_{m+1,(n,j)}} (1 - e^{-q_{m,(n,j)}}) (x_{m,n} - w_{m,(n,j)}). \quad (9)$$

A *Competitive Search Equilibrium* is a tuple (w, q, v) , where each firm chooses wages w and queue lengths q so as to maximize profits (9), taking as given workers' expected income $\{w_m\}$ and the constraint (8); and the queue lengths q satisfy a resource constraint, generalizing (4) to allow different type n firms to offer different wages and hence have different queues:

$$\mu_m = \sum_{n=1}^N \int_0^{\nu_n} q_{m,(n,j)} dj. \quad (10)$$

¹⁰With probability $\frac{1}{z!} q_{m,(n,j)}^z e^{-Q_{m,(n,j)}}$, there are $z \in \{0, 1, \dots\}$ type m job applicants and no better applicants, in which event the worker is hired with probability $\frac{1}{z+1}$. Summing across z gives a hiring probability, conditional on applying for the job, of

$$\sum_{z=0}^{\infty} \frac{1}{(z+1)!} q_{m,(n,j)}^z e^{-Q_{m,(n,j)}} = \frac{e^{-Q_{m+1,(n,j)}}}{q_{m,(n,j)}} \sum_{z=1}^{\infty} \frac{1}{z!} q_{m,(n,j)}^z e^{-q_{m,(n,j)}} = e^{-Q_{m+1,(n,j)}} \frac{1 - e^{-q_{m,(n,j)}}}{q_{m,(n,j)}}.$$

4.2 Equilibrium Characterization

Substitute the M constraints (8) into the firm's objective function (9) to eliminate the wages. Firm (n, j) 's profits may be expressed as

$$R_{(n,j)}(q_{1,(n,j)}, \dots, q_{M,(n,j)}) - \sum_{m=1}^M q_{m,(n,j)} v_m, \quad (11)$$

where

$$R_{(n,j)}(q_{1,(n,j)}, \dots, q_{M,(n,j)}) \equiv \sum_{m=1}^M e^{-Q_{m+1,(n,j)}} (1 - e^{-q_{m,(n,j)}}) x_{m,n} \quad (12)$$

is the expected revenue a firm receives as a function of its queue lengths. The revenue function is the sum of the probability that a firm hires a type m worker to fill the job times the resulting output. We subtract from this the expected cost of attracting an applicant queue of $q_{m,(n,j)}$ type m workers, each of whom must be paid v_m on average. Viewed this way, the firm has a particular production function translating 'expected applications' $q_{m,(n,j)}$ into revenue, and it faces a competitive market for applicants with v_m representing the cost of type m 'expected applications'.

This is a textbook profit maximization problem. A Competitive Search Equilibrium creates a competitive market for job applicants. A standard argument establishes the existence of a solution to this problem assuming workers' expected income v_m is positive for all m . Moreover, the proof of Proposition 1 established that the revenue function is strictly concave. It follows that all type n firms choose the same queue lengths, and so from now on I drop the redundant firm name and refer to a firm simply by its type. With this notational simplification, the necessary and sufficient first order conditions for profit maximization are identical to the first order condition of the planner's problem, (6). In addition, since all type n firms behave identically, the resource constraint (10) reduces to the feasibility constraint (4).

The representation of profit maximization as a choice of queue lengths to maximize (11) clarifies why firms always hire the most productive applicant in a Competitive Search Equilibrium. The cost of attracting a given queue $\{q_{m,n}\}$ is $\sum_{m=1}^M q_{m,n} v_m$, regardless of whom the type n firm hires. The only affect of changing the ranking of job applicants would be to reduce the revenue function R_n , which is not optimal.

I can also deduce the equilibrium wage from the preceding analysis. When $q_{m,n} > 0$,

combine the first order condition (6) with the equation for workers' expected income (8):

$$w_{m,n} = \frac{q_{m,n}e^{-q_{m,n}}}{1 - e^{-q_{m,n}}} \left(x_{m,n} - \sum_{m'=1}^{m-1} e^{-(Q_{m'+1,n} - Q_{m,n})} (1 - e^{-q_{m',n}}) x_{m',n} \right). \quad (13)$$

If $q_{m,n} = 0$, a type n firm can offer this wage without attracting any type m applicants, but the model does not pin down the wage uniquely. For example, a zero wage would have the same effect.

The wage given in equation (13) is the marginal value of an application from a type m worker conditional on hiring such a worker,¹¹ where again the marginal value of an application reflects the expected increase in output from receiving a type m application in excess of hiring the firm's next best applicant. The first term in equation (13), $\frac{q_{m,n}e^{-q_{m,n}}}{1 - e^{-q_{m,n}}}$, is the probability that a job receives exactly one type m application conditional on receiving at least one application.¹² If a type m worker is hired but the firm receives another type m application, the marginal value of her application is zero. Otherwise, the marginal value of the application is the output produced, $x_{m,n}$, in excess of the next best possibility. With probability $e^{-(Q_{m'+1,n} - Q_{m,n})} (1 - e^{-q_{m',n}})$, the firm does not get any applications that are worse than m and better than m' , but it does get at least one type m' application. In this event, the foregone output is $x_{m',n}$.

I summarize these results in a Proposition (proof in the preceding text):

Proposition 2. *A Competitive Search Equilibrium is described by queue lengths $\{q_{m,n}\}$, wages $\{w_{m,n}\}$, and expected incomes $\{v_m\}$ satisfying the resource constraint (4), the complementary slackness condition (6), and the wage equation (13). It is unique and the queue lengths are identical to those in the Social Optimum.*

If the planner has access to lump-sum transfers, the equivalence between the Competitive Search Equilibrium and Social Optimum is a version of the First and Second Welfare Theorems.

¹¹It follows that the wage firms offer equals the expected wage that the worker would earn if the firm used a sealed-bid second-price auction to sell the job to one of the applicants (Julien, Kennes, and King 2000, Shimer 1999). Existing results on the equivalence between auctions and *ex ante* wage commitments (Kultti 1999, Julien, Kennes, and King 2001) extend to environments with heterogeneous workers and firms.

¹²The probability of receiving exactly one such application is $q_{m,n}e^{-q_{m,n}}$, while the probability of receiving at least one such application is $1 - e^{-q_{m,n}}$. The result follows from Bayes rule.

5 Empirical Predictions

According to the textbook assignment model (Sattinger 1993), a worker's wage should be determined by her characteristics, not by her job; however, in a regression of wages on a worker's characteristics, much of the residual can be explained through the characteristics of her job (Krueger and Summers 1988, Groshen 1991, Gibbons and Katz 1992, Abowd, Kramarz, and Margolis 1999). One possible explanation is unobserved worker heterogeneity (Murphy and Topel 1987). Since more productive workers get better jobs, the job reveals something about the worker's productivity that is observable to firms but unobservable to the econometrician. But that does not seem to be the whole story. Krueger and Summers (1988) and Gibbons and Katz (1992) find that workers who move from a high to a low wage firm lose approximately the wage differential between the two firms. Another explanation is that workers in some industries receive a compensating differential. Again, this explanation appears to be incomplete, since Krueger and Summers (1988) find that industry fixed effects have little explanatory power. Most wage dispersion appears to be at the level of individual firms.

This paper provides a concise explanation that is consistent with this evidence. A single type of worker typically opts to search over a range of different types of jobs. A worker who earns a high wage relative to her characteristics sought and found a high wage, high productivity job. The presence of firm effects in a wage regression is thus a classic sample selection problem: the data set does not include the workers who seek but fail to find high wage jobs. This interpretation is also consistent with Holzer, Katz, and Krueger's (1991) observation that high wage firms attract more applicants. In equilibrium, high wage firms are more productive, and more productive firms expect more applicants, i.e. $Q_{1,n}$ is increasing.

The key to this logic is that more productive firms pay higher wages. This is true if the production function is supermodular: for all $m_1 < m_2$ and $n_1 < n_2$, $x_{m_2,n_2} - x_{m_1,n_2} > x_{m_2,n_1} - x_{m_1,n_1}$. This is a familiar condition in the assignment literature, since it ensures positively assortative matching in a frictionless environment.

Proposition 3. *Assume x is supermodular. Then $Q_{m,n}$ is strictly increasing in n when it is positive. In addition, a more productive job is more likely to be filled, a worker is less likely to obtain a more productive job conditional on applying for it, and a worker's wage is increasing in her employer's productivity.*

The appendicized proof uses a variational argument to establish that $Q_{m,n}$ is strictly increasing in n . The remaining results follow immediately. It is harder to get a more productive

job, since it attracts more good applicants. Workers then require compensation—a higher wage—in return for the lower hiring probability.

A natural followup question is whether in equilibrium more productive workers receive higher wages at a given firm, i.e. is $w_{m,n}$ increasing not only in n but also in m ? The answer in general is no. Workers are compensated not only through their wage but through their hiring probability. In fact, it is easy to construct examples in which less productive workers earn a higher wage. Suppose type m and $m + 1$ workers are almost equally productive, and there are many type $m + 1$ workers. When a type m worker is hired, there are no type $m + 1$ applicants, and so her marginal product is high. In contrast, when a type $m + 1$ worker is hired, it is likely that she is not the only such applicant, and so she receives a low wage. In the limit, if m and $m + 1$ are equally productive and there is a positive measure of either type of applicant, type m workers will always receive a strictly higher wage.¹³ One can even construct nongeneric examples in which the wage $w_{m,n}$ is a function of the firm's type n but not the worker's type m , a finding that, if not viewed through the lens of this model, would appear to be at odds with a competitive labor market.

Similarly, one might expect that more productive workers are more likely to find jobs. Since firms rank workers according to their productivity, this is true conditional on the type of job that a worker applies for. But a countervailing force is that less productive workers may apply for less productive jobs which are easier to obtain. In fact, it is easy to construct examples in which the latter force is dominant, so the employment probability is increasing in a worker's type.¹⁴

On the other hand, it is generally true that after conditioning on firm characteristics, jobs filled by better workers should earn higher profits:

Proposition 4. *$x_{m,n} - w_{m,n}$ is increasing in m whenever $q_{m,n}$ is positive. Thus there is a positive correlation between a firm's profit and the quality of its worker after conditioning on the firm's type n .*

The proof is in the appendix.

An important corollary of this Proposition is that firms always want to hire the most

¹³But in such a case, there would also be another equilibrium in which type m workers are hired in preference to type $m + 1$ workers, as well as an equilibrium in which the two are ranked equally. All these equilibria yield the same output. This example violates an assumption in this paper, strict monotonicity of the production function.

¹⁴Suppose there are two types of workers and two types of jobs, with production function $x_{m,n} = \min\{m, n\}$. Also assume $\log 2 + \frac{\mu_1}{\nu_1} > \frac{\mu_2}{\nu_2} > \frac{\mu_1}{\nu_1}$. Then all type i workers apply for type i jobs and the relative scarcity of type 2 jobs ensures that type 2 workers are unemployed more frequently.

productive job applicant, so there is no tension between the *ex ante* efficiency of hiring the most productive applicant and the firm’s desire to maximize *ex post* profits by choosing the most profitable applicant.

The intuition for this Proposition comes from the nature of the wage in this model. It reflects the marginal value to the firm of receiving an application from a type m worker conditional on hiring that worker. Now compare the marginal value of applications from a type m_1 and a type $m_2 > m_1$ worker. A type m_2 worker produces more output, which tends to raise her wage in proportion to her productivity. On the other hand, when a type m_2 worker is hired, there is a possibility that the firm received other type m_2 applications, a possibility that the firm received type $m \in \{m_1 + 1, \dots, m_2 - 1\}$ applications, and a possibility that the firm received type m_1 applications. All of these reduce the marginal value of hiring a type m_2 worker, and hence reduce her wage. When the firm hires a type m_1 worker, only the last possibility remains, and even that probability is lower. This tends to raise the relative wage of a low productivity worker and hence increase the *ex post* profit from hiring a high productivity worker.

There is less evidence in support of Proposition 4, since an empirical investigation must utilize a matched worker-firm data set to measure both firm profits and worker characteristics. To my knowledge, Abowd, Kramarz, and Margolis (1999) provide the only direct test of this hypothesis, using French data. Their Table X on page 298 shows that a firm’s profits, measured as the ratio of operating income divided by capital stock, is increasing in its workers’ observable characteristics (‘Average Predicted Effect of x Variables ($x\nu$)’). Moreover, Abowd, Kramarz, and Margolis follow workers over time, and so can include individual fixed effects in their regression to control for unobserved heterogeneity. They find that workers’ unobserved characteristics have a small but statistically insignificant effect on firm profits. An important caveat in interpreting these results is that my model does not capture the institutional structure of the French labor market, e.g. high minimum wage levels and centralized bargaining. Nonetheless, existing empirical evidence on the correlation between firm profits and worker characteristics is consistent with the model.

6 Example: No Comparative Advantage

This section analyzes a special case in detail in an effort to clarify the more general results presented in the rest of the paper. Assume $x_{m,n}$ can be expressed as $h_m k_n$ for some increasing vectors $\{h_1, \dots, h_M\}$ and $\{k_1, \dots, k_N\}$. In other words, the ratio of output produced by type

m_1 and type m_2 workers in type n_1 and n_2 firms is identical,

$$\frac{x_{m_1, n_1}}{x_{m_2, n_1}} = \frac{x_{m_1, n_2}}{x_{m_2, n_2}}.$$

Following Sattinger (1975), I refer to this as the case without comparative advantage.

6.1 Characterization

The assignment of workers to jobs can be characterized as follows: for each worker type m , there is a threshold productivity level $\bar{\kappa}_m$ such that workers only apply for jobs with $k_n \geq \bar{\kappa}_m$. These thresholds are positive and increasing. Moreover a type m worker is equally likely to apply for a job at any type n firm with $k_n \geq \bar{\kappa}_{m+1}$ and strictly less likely to apply for a job at a firm with $\bar{\kappa}_{m+1} > k_n \geq \bar{\kappa}_m$.¹⁵ More precisely:

Proposition 5. *Assume there is no comparative advantage. Define $\bar{\kappa}_m$ inductively by $\bar{\kappa}_{M+1} \equiv k_N$ and*

$$\mu_m = \sum_{n=1}^N \nu_n (\min\{\log \bar{\kappa}_{m+1}, \log k_n\} - \min\{\log \bar{\kappa}_m, \log k_n\}).$$

Then $0 < \bar{\kappa}_1 < \dots < \bar{\kappa}_M < k_N$. The queue of type m workers for a type n job satisfies

$$q_{m,n} = \min\{\log \bar{\kappa}_{m+1}, \log k_n\} - \min\{\log \bar{\kappa}_m, \log k_n\},$$

which is positive for $k_n > \bar{\kappa}_m$ and independent of n for $k_n > \bar{\kappa}_{m+1}$. A type m worker's shadow value or expected income satisfies

$$v_m = \sum_{m'=1}^m \bar{\kappa}_{m'} (h_{m'} - h_{m'-1}).$$

The proof is in the appendix.

Figure 1 depicts the threshold function $\bar{\kappa}_m$ graphically. Because workers use threshold rules, if type m workers apply for type n jobs with positive probability, then there is also a positive probability that they apply for any better (type $n+1, \dots, N$) job. And because the thresholds are increasing in the worker's type, if type m workers apply for type n jobs with

¹⁵That is, $q_{m,n}$ is independent of n if $k_n \geq \bar{\kappa}_{m+1}$, and $q_{m,n} > q_{m,n'} > 0$ if $k_n \geq \bar{\kappa}_{m+1} > k_{n'} > \bar{\kappa}_m$.

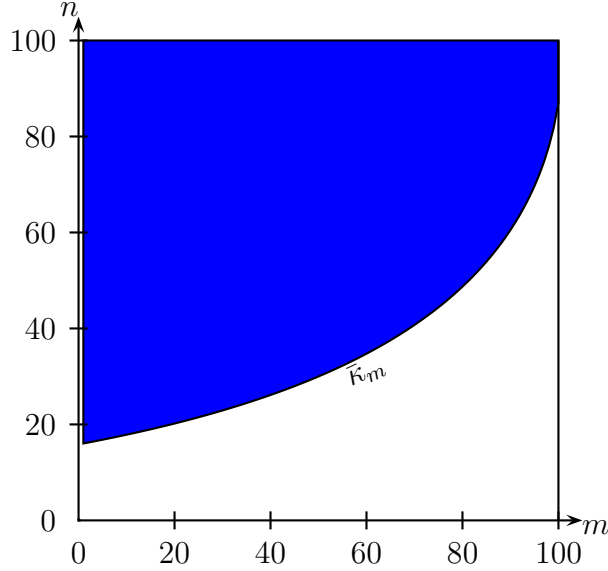


Figure 1: Matching sets: $M = N = 100$, $x_{m,n} = mn$, $\mu_m = \nu_n = 1$. The shaded region indicates the set of job types that a type m worker will apply for. The threshold \bar{k}_m is increasing, so more productive workers are more likely to apply for more productive jobs.

positive probability, there is also a positive probability that any worse (type $1, \dots, m - 1$) worker applies for a type n job.

A corollary is that ‘mismatch’ is normal. That is, it is possible to find both a more productive firm that hires a less productive worker and a less productive firm that hires a more productive worker. For example, if type M workers apply for type $N - 1$ jobs,¹⁶ it follows that all workers apply for both types of jobs. Some type N jobs will fail to attract a type M worker and so will instead hire a less productive worker, while some type $N - 1$ jobs will succeed in attracting type M workers. More generally, there will be mismatch in this model unless type all type $m \in \{2, \dots, M\}$ workers apply exclusively for type N jobs.¹⁷ The shaded region in Figure 1 indicates that there may be a substantial amount of mismatch.

The social planner’s decision to mismatch workers and firms is intuitive. On the one hand, the planner has an incentive to raise the employment rate of the most productive workers, which he does by spreading these workers out across firms. This effect is limited by the planner’s desire to take advantage of the complementarity between worker and firm characteristics. On the other hand, the planner also wants to ensure that high productivity

¹⁶This occurs if and only if $\mu_M > (\log k_N - \log k_{N-1})\nu_N$, so type M workers are plentiful, type N firms are scarce, or the productivity difference between type N and $N - 1$ jobs is small.

¹⁷The necessary and sufficient condition for some mismatch is $\sum_{m=2}^M \mu_m > (\log \bar{k}_N - \log \bar{k}_{N-1})\nu_N$.

jobs are filled, which he does by instructing less productive workers to apply for those jobs as well. These workers effectively serve as cheap insurance against the possibility that no high productivity worker applies for the job.¹⁸ In the decentralized equilibrium, firms' ability to rank applicants makes it more profitable to get one high productivity application and one low productivity application on average, rather than getting two medium applications.

The pervasiveness of mismatch is important for at least two reasons. First, it may help to explain why one does not observe a perfect rank correlation between worker and firm characteristics in the data. Second, it implies that one may meaningfully ask questions like those posed in Propositions 3 and 4: how does a worker's unemployment risk and wage depend on the quality of the job she applies for, conditional on the worker's type? How does a firm's profit depend on the quality of its employees, conditional on the firm's type? Such questions would be nonsensical in an economy without mismatch.

Introducing a small amount of comparative advantage, either for good workers in good jobs or good workers in bad jobs, would not eliminate mismatch. In fact, numerical examples indicate that, unless high productivity workers have a very strong comparative advantage in high productivity jobs (e.g. $x_{m,n} = \min\{h_m, k_n\}$) or there is a severe imbalance in the ratio workers to firms, there will be mismatch.¹⁹ Of course, the functional forms in Proposition 5 depend on the absence of comparative advantage.

6.2 Assortative Matching

Despite the mismatch between workers and firms, the model still makes strong predictions about the relationship between matched worker and firm characteristics in the absence of comparative advantage. Consider an econometrician who has a data set containing a matched sample of workers and jobs. An observation consists of a worker's type m and her employer's type n (or alternatively h_m and k_n). In a frictionless version of the assignment model, the rank correlation coefficient should be equal to one.²⁰ In an economy with coordination frictions, I find that the rank correlation coefficient is still positive, a weaker notion of assortative matching. In fact, I prove a significantly stronger result. Let $\epsilon_{m,n}$ denote the

¹⁸A critical assumption is that firms can rank the job applicants and hire only the most productive one, so less productive workers do not crowd out more productive ones. Without such ranking, it is optimal to avoid mismatch by segregating different workers to different firms (Shi 2001).

¹⁹A previous version of this paper (Shi 2001) allowed for the possibility that there is a continuum of firms. I proved that if there are at least two types of workers, the support of the distribution of firm characteristics is convex, and the production function is strictly increasing, then there will always be some mismatch.

²⁰This is exactly correct if there are no atoms in the type distribution or if $M = N$ and $\mu_m = \nu_m$ for all $m \in \{1, \dots, M\}$. The latter assumption holds in the example in Figure 1.

fraction of type m workers who are employed by type n firms. I prove that high productivity workers are relatively more likely to be employed in high productivity jobs than in low productivity jobs. Allowing for the possibility that $\epsilon_{m,n} = 0$ for some worker-firm pairs, this is formally a statement of log-supermodularity of $\epsilon_{m,n}$: for $m_1 < m_2$ and $n_1 < n_2$, $\epsilon_{m_1,n_1}\epsilon_{m_2,n_2} > \epsilon_{m_1,n_2}\epsilon_{m_2,n_1}$. Similarly, the fraction of type n firms that employ type m workers, $\epsilon_{m,n}\frac{\mu_m}{\nu_n}$, is also log-supermodular:

Proposition 6. *Assume there is no comparative advantage. Then the fraction of type m workers who are employed by type n firms, $\epsilon_{m,n}$, and the fraction of type n firms that employ type m workers, $\epsilon_{m,n}\frac{\mu_m}{\nu_n}$, are log-supermodular.*

The proof is in the appendix.

In principle, it is possible to test this prediction, but in practice the data demands may be unrealistic, so the power of such a test may be minimal. It is therefore useful to note increasingly weak but more easily testable implications of log-supermodularity:

1. Log-supermodularity of $\epsilon_{m,n}$ implies that the distribution of employers for type m workers first order stochastically dominates the distribution of employers for type $m - 1$ workers. Similarly, the distribution of employees for type n firms first order stochastically dominates the distribution for type $n - 1$ firms.
2. First order stochastic dominance implies that expected employer's type (either the expected value of n or of k_n) is increasing in a worker's type, and that the expected employee's type (m or h_m) is increasing in a firm's type.
3. Monotonicity of the expected partner implies a positive correlation between matched worker and firm types.

I omit the proof of these standard results.

One can also prove that $p_{m,n}$, the probability that a type m worker applies for a type n job, satisfies log-supermodularity. If an econometrician had access to a data set consisting of an unemployed worker's characteristics and the quality of the job that she applies to, she should find that more productive workers are relatively more likely to apply for more productive jobs, with the analogous subsidiary implications.

Again, it is interesting to ask the extent to which these results depend on the assumption that there is no comparative advantage, $x_{m,n} = h_mk_n$. Now this assumption is more important. It is easy to construct examples in which assortative matching fails. Suppose there are

two types of workers and $x_{m,n} = h_m + k_n$, so low productivity workers have a comparative advantage in high productivity jobs. Then one can show that more productive workers are relatively more likely to be employed in less productive jobs, making use of their comparative advantage. It is clear why the social planner desires this assignment. He would like to ensure employment for all the high productivity workers by spreading them out across jobs, while he would like to ensure that the highest productivity jobs are filled, which he does by sending low productivity workers exclusively to such jobs. One can prove that if there are only two types of workers and high productivity workers have a comparative advantage in high productivity jobs, then the employment probabilities $\epsilon_{m,n}$ are log-supermodular.²¹ I conjecture that this result carries over to an arbitrary number of worker types.

6.3 Comparative Statics

The model is sufficiently tractable so as to be amenable to comparative statics. Consider an improvement in the composition of the population of workers. A measure η of type m workers suddenly become type $m + 1$ workers. One can verify that $\bar{\kappa}_{m'}$ is unchanged for $m' \neq m + 1$, while $\bar{\kappa}_{m+1}$ falls.²² On the other hand, the shadow value or expected income $v_{m'}$ falls for all type $m' \geq m + 1$ workers, but is unchanged for less productive workers.

Conversely, an improvement in the firm distribution, say taking a measure η of type n firms and raising their type to $n + 1$, raises the threshold $\bar{\kappa}_m$ for all workers who apply to type $n + 1$ or worse jobs. It follows that the shadow value v_m rises not only for low productivity workers but also for high productivity workers who do not apply for type n or $n + 1$ jobs.

7 Related Literature

Montgomery (1991) and Peters (1991) explore the implications of symmetry restrictions in wage or price posting games similar to the one analyzed here. Burdett, Shi, and Wright (2001) refine these analyses, showing that the equilibrium of an economy with a finite number of buyers and sellers converges to the equilibrium of an economy with infinitely many buyers

²¹The proof is contained in Shimer (2001).

²²In particular, note that $\bar{\kappa}_m$ is unchanged since $\mu_m + \mu_{m+1}$ is unchanged and, according to the definition of $\bar{\kappa}_m$ in Proposition 5,

$$\mu_m + \mu_{m+1} = \sum_{n=1}^N \nu_n \left(\min \{ \log \bar{\kappa}_{m+2}, \log k_n \} - \min \{ \log \bar{\kappa}_m, \log k_n \} \right).$$

and sellers. They also extend the earlier papers by allowing firms to create more than one vacancy, a possibility that I do not admit.

Two recent papers by Shouyong Shi have extended the wage posting framework to environments with heterogeneous workers and jobs. Shi (2001) looks at a similar economy to the one in this paper, but assumes that before search begins, each firm must commit to hire a particular type of worker. Thus firms cannot use applications from bad workers as insurance against not getting an application from a good worker. Although it is still possible that identical firms choose to gather applications from different types of workers, Shi proves that this does not happen in equilibrium. My analysis gives firms the option of committing to hire only one type of worker, for example by offering other types a zero wage, but I prove that in general they choose not to exercise that option. Since in equilibrium firms attract applications from different types of workers, my model generates endogenous mismatch, differential unemployment rates for different workers applying for the same type of job, and a correlation between a firm's profit and a worker's characteristics after controlling for the firm's characteristics. None of those results make sense in Shi's equilibrium.

Shi (2002) does not impose the commitment restriction, making the model fairly similar to the one in this paper. The main technical difference between that paper and this one is that Shi analyzes a model with two types of workers and firms and a particular production function, where output is dictated only by the worst partner's type. I consider an arbitrary number of worker and firm types and a much broader class of production functions. This generates the possibility of mismatch which is absent from Shi (2002). Without mismatch, Shi is unable to discuss how wages and profits vary with firms' and workers' productivity (Propositions 3 and 4). On the other hand, Shi endogenizes firms' entry decisions and considers a number of comparative statics results related to skill-biased technical change that go beyond the scope of this paper.

This paper is also related to random search models with heterogeneous agents (Sattinger 1995, Lu and McAfee 1996, Burdett and Coles 1999, Shimer and Smith 2000, Davis 2001). These papers assume workers have no information about jobs and so must randomly look for them. There are several advantages to the present model over the random search framework. First, in the random search framework, wage setting is determined outside the model, typically through a Nash bargaining solution. Equilibrium matching patterns depend on the exact specification of the bargaining game, including the threat points while bargaining. There is no theoretical reason to prefer one specification of the bargaining game over another. Second, the random search framework assumes there is mismatch. With a continuum

of job productivity levels, it would take infinitely long for a worker to find a particular type of job, and so necessarily optimizing agents must compromise on their matching pattern. That is not the case in this paper; the planner could always get rid of mismatch by assigning different workers types to different job types. Even in a decentralized economy, different workers could apply for different job types. They choose not to do so.

Third, in the random search model, low productivity workers impose a congestion externality on the search process, making it harder for jobs to meet high productivity workers. This generates inefficiencies in a decentralized search equilibrium (Davis 2001, Shimer and Smith 2001a) and may imply that limit cycles in which some types of matches are repeatedly created and then destroyed are more efficient than steady state equilibria (Shimer and Smith 2001b). In the assignment model with coordination frictions, firms can (and do) choose not to hire bad workers when good ones are available, eliminating the congestion externality. The decentralized equilibrium is unique and efficient, and even in dynamic extensions to the model, there is nothing to be gained by pursuing nonstationary policies. Fourth, the random search framework is not very tractable, while this paper demonstrates the possibility of performing some simple cross-sectional comparisons and comparative statics in the assignment model with coordination frictions.

Finally, an assignment model with informational frictions is consistent with many of the facts discussed in the introduction. Gibbons and Katz (1992) explain inter-industry wage differentials by developing a model in which workers gradually learn about their productivity and are reassigned to more appropriate industries as the learning process proceeds. Davis (1997) shows how workers can move up a career ladder as the market receives good news about their ability. This can generate an empirical correlation between a worker's place in the corporate hierarchy and her earnings when in fact both are determined by the evolution of the market's perception of the worker's ability. There are features of the coordination frictions model that are absent from the information frictions model, such as equilibrium unemployment and vacancies. Nevertheless, I do not want to claim that learning is unimportant for wage dynamics, only that the assignment model with coordination frictions may provide part of the explanation for the observed data.

8 Conclusion

The assignment model with coordination frictions explains a rich set of interactions between heterogeneous workers and firms. It is also tractable, particularly in the special case with-

out comparative advantage. It should therefore lend itself to a number of extensions. In concluding, I mention only two.

First, I have assumed that workers can only apply for one job. There are conceptual difficulties in allowing workers to apply for multiple jobs simultaneously: can firms make ‘second-round’ offers in the event their first offer is turned down? Albrecht, Gautier, and Vroman (2003) analyze a version of this model with homogeneous workers and firms, showing that the basic properties of the model carry over to an environment without second round offers.

Second, the model must be extended to a dynamic framework if it is to be taken quantitatively seriously. In a dynamic model, the extent of the coordination frictions is governed by the time lag before an unemployed worker can apply for another job. If jobs are geographically disperse, this might be quite long, while in a compact labor market, it is likely that the coordination frictions will quickly resolve themselves. The extension should also have some qualitative effects on the results. For example, I showed that in the absence of comparative advantage, the most productive firms attract applications from all workers. In a dynamic model, these firms would refuse to hire the least productive workers because they could wait until the following period to hire a better worker. This would likely strengthen the assortative matching results discussed in Section 6.

A Appendicized Proofs

Proposition 1. *Any queue lengths $\{q_{m,n}\}$ and shadow values $\{v_m\}$ satisfying the feasibility constraint (4) and the complementary slackness condition (6) is socially optimal. The Social Optimum is unique.*

Proof. I have demonstrated in the text the necessity of conditions (4) and (6). This proof establishes uniqueness by proving that aggregate output is a strictly concave function of q ; the constraint set is obviously linear. It is useful to rewrite output as

$$Y(q) = \sum_{n=1}^N \nu_n R_n(q_{1,n}, \dots, q_{M,n}),$$

where $R_n(q_{1,n}, \dots, q_{M,n})$ is the firm’s revenue function, defined in equation (12). First com-

pute the Hessian of R_n :

$$D^2 R_n = \begin{bmatrix} a_{1,n} & a_{1,n} & a_{1,n} & \cdots & a_{1,n} \\ a_{1,n} & \sum_{m=1}^2 a_{m,n} & \sum_{m=1}^2 a_{m,n} & \cdots & \sum_{m=1}^2 a_{m,n} \\ a_{1,n} & \sum_{m=1}^2 a_{m,n} & \sum_{m=1}^3 a_{m,n} & \cdots & \sum_{m=1}^3 a_{m,n} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_{1,n} & \sum_{m=1}^2 a_{m,n} & \sum_{m=1}^3 a_{m,n} & \cdots & \sum_{m=1}^M a_{m,n} \end{bmatrix},$$

where $a_{m,n} \equiv -e^{-Q_{m,n}}(x_{m,n} - x_{m-1,n}) < 0$. Pre- and post-multiplying this matrix by a non-zero vector $u \equiv \{u_1, \dots, u_M\}$ gives

$$u' D^2 R_n u = \sum_{m=1}^M a_{m,n} \left(\sum_{m'=m}^M u_{m'} \right)^2.$$

Let u_m be the last non-zero element of the vector u . Then $a_{m,n} \left(\sum_{m'=m}^M u_{m'} \right)^2 = a_{m,n} u_m^2$ is negative and the remaining terms are non-positive, proving that $D^2 R_n$ is negative definite.

Next, consider to the $MN \times MN$ Hessian of aggregate output $Y(q)$. Since the cross partial derivative with respect to q_{m_1, n_1} and q_{m_2, n_2} is zero for any $n_1 \neq n_2$, the Hessian of Y has a block diagonal structure, with each block corresponding to one of the Hessians of R_n . Since each block is negative definite, the entire matrix is negative definite, which implies that output is a strictly concave function of queue lengths q . \square

Proposition 3. *Assume x is supermodular. Then $Q_{m,n}$ is strictly increasing in n when it is positive. In addition, a more productive job is more likely to be filled, a worker is less likely to obtain a more productive job conditional on applying for it, and a worker's wage is increasing in her employer's productivity.*

Proof. In order to find a contradiction, suppose there is an m_1 and $n_1 < n_2$ with $Q_{m_1, n_1} > 0$ and $Q_{m_1, n_1} \geq Q_{m_1, n_2}$. If $q_{m_1, n_1} > 0$, define $m_2 = m_1$; otherwise, let m_2 be the smallest worker type larger than m_1 with $q_{m_2, n_1} > 0$. In either case, $Q_{m_1, n_1} = Q_{m_2, n_1}$, so $Q_{m_2, n_1} > 0$ and $Q_{m_2, n_1} \geq Q_{m_2, n_2}$ as well. I will prove that reducing Q_{m_2, n_1} by $\varepsilon \nu_{n_2}$ and raising Q_{m_2, n_2} by $\varepsilon \nu_{n_1}$ is feasible for small but positive ε , and it results in higher output.

The proposed deviation is equivalent to reducing q_{m_2, n_1} by $\varepsilon \nu_{n_2}$ and raising q_{m_2, n_2} by $\varepsilon \nu_{n_1}$, which is feasible because $q_{m_2, n_1} > 0$ and because constraint (4) continues to hold. To show

that this perturbation results in higher output, regroup the terms in (3) to express aggregate output as

$$Y(q) = \sum_{n=1}^N \nu_n \left(x_{M,n} - \sum_{m=1}^M e^{-Q_{m,n}} (x_{m,n} - x_{m-1,n}) \right),$$

Then for sufficiently small ε , the change in output is approximately equal to

$$\varepsilon \nu_{n_2} \nu_{n_1} \left(e^{-Q_{m_2, n_2}} (x_{m_2, n_2} - x_{m_2-1, n_2}) - e^{-Q_{m_2, n_1}} (x_{m_2, n_1} - x_{m_2-1, n_1}) \right).$$

Supermodularity implies $x_{m_2, n_2} - x_{m_2-1, n_2} > x_{m_2, n_1} - x_{m_2-1, n_1}$, while $e^{-Q_{m_2, n_2}} \geq e^{-Q_{m_2, n_1}}$ by assumption. This deviation therefore raises output, a contradiction.

Next, the probability that a job is filled is $1 - e^{-Q_{1,n}}$, increasing in $Q_{1,n}$ and hence in n . The probability that a worker obtains a job conditional on applying for it is given in footnote 10,

$$e^{-Q_{m+1,n}} \frac{1 - e^{-q_{m,n}}}{q_{m,n}} = \frac{e^{-Q_{m+1,n}} - e^{-Q_{m,n}}}{Q_{m,n} - Q_{m+1,n}},$$

where the equality uses $q_{m,n} = Q_{m,n} - Q_{m+1,n}$. Simple differentiation shows that this last expression is decreasing in both $Q_{m+1,n}$ and $Q_{m,n}$ and hence decreasing in n . It then follows immediately from equation (7) that workers who apply for more productive jobs are paid higher wages, a compensating differential for the lower employment probability. \square

Proposition 4. *$x_{m,n} - w_{m,n}$ is increasing in m whenever $q_{m,n}$ is positive. Thus there is a positive correlation between a firm's profit and the quality of its worker after conditioning on the firm's type n .*

Proof. Regroup terms in equation (13) to get

$$w_{m,n} = \frac{q_{m,n} e^{-q_{m,n}}}{1 - e^{-q_{m,n}}} \sum_{m'=1}^m e^{-(Q_{m',n} - Q_{m,n})} (x_{m',n} - x_{m'-1,n}). \quad (14)$$

Then take any $m_1 < m_2$ and n with $q_{m_1,n} > 0$, $q_{m_2,n} > 0$, and $q_{m,n} = 0$ for all $m \in \{m_1 + 1, \dots, m_2 - 1\}$. $\frac{q_{m_1,n}}{1 - e^{-q_{m_1,n}}} > 1$, since this is the inverse of the probability that a type m_1 worker is hired when she applies for a job with no more productive applicants. Thus

equation (14) implies

$$w_{m_1,n} > \sum_{m'=1}^{m_1} e^{-(Q_{m',n}-Q_{m_1+1,n})}(x_{m',n} - x_{m'-1,n}).$$

Similarly, $q_{m_2,n} > 0$ implies $\frac{q_{m_2,n}e^{-q_{m_2,n}}}{1-e^{-q_{m_2,n}}} < 1$, since this is the probability that the job receives no identical applications conditional on hiring a type m_2 worker. Thus equation (14) implies

$$w_{m_2,n} < \sum_{m'=1}^{m_2} e^{-(Q_{m',n}-Q_{m_2,n})}(x_{m',n} - x_{m'-1,n}).$$

Subtract the first inequality from the second and simplify by combining terms to get

$$\begin{aligned} w_{m_2,n} - w_{m_1,n} &< (e^{Q_{m_1+1,n}} - e^{Q_{m_2,n}}) \sum_{m'=1}^{m_1} e^{-Q_{m',n}}(x_{m',n} - x_{m'-1,n}) \\ &\quad + \sum_{m'=m_1+1}^{m_2} e^{-(Q_{m',n}-Q_{m_2,n})}(x_{m',n} - x_{m'-1,n}). \end{aligned}$$

Next, use the assumption that $q_{m,n} = 0$ for all $m \in \{m_1 + 1, \dots, m_2 - 1\}$. This implies $Q_{m',n} = Q_{m_1+1} = Q_{m_2,n}$ for all $m' \in \{m_1 + 1, \dots, m_2\}$. In particular, the first line evaluates to zero, as does the exponent in the second line:

$$w_{m_2,n} - w_{m_1,n} < \sum_{m'=m_1+1}^{m_2} (x_{m',n} - x_{m'-1,n}) = x_{m_2,n} - x_{m_1,n}.$$

That is, $x_{m_2,n} - w_{m_2,n} > x_{m_1,n} - w_{m_1,n}$. By transitivity, this holds for arbitrary m_1 and m_2 . \square

Proposition 5. *Assume there is no comparative advantage. Define $\bar{\kappa}_m$ inductively by $\bar{\kappa}_{M+1} \equiv k_N$ and*

$$\mu_m = \sum_{n=1}^N \nu_n (\min\{\log \bar{\kappa}_{m+1}, \log k_n\} - \min\{\log \bar{\kappa}_m, \log k_n\}).$$

Then $0 < \bar{\kappa}_1 < \dots < \bar{\kappa}_M < k_N$. The queue of type m workers for a type n job satisfies

$$q_{m,n} = \min\{\log \bar{\kappa}_{m+1}, \log k_n\} - \min\{\log \bar{\kappa}_m, \log k_n\},$$

which is positive for $k_n > \bar{\kappa}_m$ and independent of n for $k_n > \bar{\kappa}_{m+1}$. A type m worker's shadow value or expected income satisfies

$$v_m = \sum_{m'=1}^m \bar{\kappa}_{m'} (h_{m'} - h_{m'-1}).$$

Proof. The first step is to prove that $0 < \bar{\kappa}_1 < \dots < \bar{\kappa}_M < k_N$. The argument proceeds by induction. Recall that $\bar{\kappa}_{M+1} \equiv k_N$. Taking as given the value of $\bar{\kappa}_{m+1}$, define $\bar{\kappa}_m$ via

$$\mu_m = \sum_{n=1}^N \nu_n \left(\min \{ \log \bar{\kappa}_{m+1}, \log k_n \} - \min \{ \log \bar{\kappa}_m, \log k_n \} \right)$$

The right hand side is continuous in $\bar{\kappa}_m > 0$, is non-positive when $\bar{\kappa}_m \geq \bar{\kappa}_{m+1}$, and is strictly decreasing in $\bar{\kappa}_m$ when $\bar{\kappa}_m < \bar{\kappa}_{m+1}$. Moreover, it approaches infinity when $\bar{\kappa}_m$ gets close to zero. This condition therefore uniquely defines $\bar{\kappa}_m \in (0, \bar{\kappa}_{m+1})$.

Next, sum $\nu_n q_{m,n}$ across n to verify that the feasibility condition (4) is satisfied. To establish the complementary slackness condition (6), first regroup terms and write it as

$$v_m \geq \sum_{m'=1}^m e^{-Q_{m',n}} (x_{m',n} - x_{m'-1,n}) \text{ and } q_{m,n} \geq 0,$$

with complementary slackness. Then note that the proposed queue lengths q imply that the queue of strictly more productive workers satisfies

$$Q_{m,n} = \log k_n - \min \{ \log \bar{\kappa}_m, \log k_n \}. \quad (15)$$

Substitute this into the preceding expression to get

$$v_m \geq \sum_{m'=1}^m \min \{ \bar{\kappa}_{m'}, k_n \} (h_{m'} - h_{m'-1}) \text{ and } q_{m,n} \geq 0,$$

with complementary slackness. Given the proposed expression for the shadow value, the first inequality always obtains, and it holds as an equality if and only if $k_n > \bar{\kappa}_m$, i.e. whenever $q_{m,n} > 0$. Proposition 1 implies (q, v) is optimal since it satisfies the feasibility condition (4) and first order condition (6). Proposition 2 implies it is an equilibrium. \square

Proposition 6. *Assume there is no comparative advantage. Then the fraction of type m workers who are employed by type n firms, $\epsilon_{m,n}$, and the fraction of type n firms that employ type m workers, $\epsilon_{m,n} \frac{\mu_m}{\nu_n}$, are log-supermodular.*

Proof. To get a job, (i) a worker must apply for it, with probability $p_{m,n}$; (ii) there must not be any more productive applicants, with probability $e^{-Q_{m+1,n}}$; and (iii) the worker must be chosen from among the identical applicants, with probability $\frac{1-e^{-q_{m,n}}}{q_{m,n}}$ (see footnote 10). Multiply the probability of these three independent events, using equation (1) to eliminate $p_{m,n}$, to get the fraction of type m workers that obtain type n jobs,

$$\epsilon_{m,n} = e^{-Q_{m+1,n}} (1 - e^{-q_{m,n}}) \frac{\nu_n}{\mu_m}. \quad (16)$$

This holds for any production function.

When $x_{m,n} = h_m k_n$, substitute for q and Q in equation (16) using the expressions in Proposition 5 and equation (15):

$$\epsilon_{m,n} \equiv \left(\min\{\bar{\kappa}_{m+1}, k_n\} - \min\{\bar{\kappa}_m, k_n\} \right) \frac{\nu_n}{\mu_m k_n}.$$

Note that $\epsilon_{m,n} > 0$ if and only if $k_n > \bar{\kappa}_m$. Now suppose for some $m < M$ and $n_1 < n_2$ that $\epsilon_{m,n_1} \epsilon_{m+1,n_2} < \epsilon_{m,n_2} \epsilon_{m+1,n_1}$. Among other things, this implies that $\epsilon_{m+1,n_1} > 0$, so $k_{n_1} > \bar{\kappa}_{m+1}$. Monotonicity of k and $\bar{\kappa}$ then implies $k_{n_2} > k_{n_1} > \bar{\kappa}_{m+1} > \bar{\kappa}_m$, and in particular

$$\frac{\epsilon_{m,n_1} k_{n_1}}{\nu_{n_1}} = \frac{\epsilon_{m,n_2} k_{n_2}}{\nu_{n_2}} = \frac{\bar{\kappa}_{m+1} - \bar{\kappa}_m}{\mu_m}.$$

The supposition that $\epsilon_{m,n_1} \epsilon_{m+1,n_2} < \epsilon_{m,n_2} \epsilon_{m+1,n_1}$ therefore holds if and only if

$$\min\{\bar{\kappa}_{m+2}, k_{n_2}\} - \min\{\bar{\kappa}_{m+1}, k_{n_2}\} < \min\{\bar{\kappa}_{m+2}, k_{n_1}\} - \min\{\bar{\kappa}_{m+1}, k_{n_1}\}.$$

But since $k_{n_2} > k_{n_1} > \bar{\kappa}_{m+1}$, this is impossible. This proves $\epsilon_{m,n_1} \epsilon_{m+1,n_2} \geq \epsilon_{m,n_2} \epsilon_{m+1,n_1}$. Transitivity establishes log-supermodularity of $\epsilon_{m,n}$.

For the case of firms, note that the fraction of type n firms that hire type m workers is $e^{-Q_{m+1,n}} (1 - e^{-q_{m,n}}) \equiv \epsilon_{m,n} \mu_m / \nu_n$. Log-supermodularity of $\epsilon_{m,n}$ immediately implies log-supermodularity of this object. \square

References

- ABOWD, JOHN, FRANCIS KRAMARZ, AND DAVID MARGOLIS (1999): “High Wage Workers and High Wage Firms,” *Econometrica*, 67(2), 251–333.
- ALBRECHT, JAMES, PIETER GAUTIER, AND SUSAN VROMAN (2003): “Matching with Multiple Applications,” *Economics Letters*, 78(1), 67–70.
- BECKER, GARY (1973): “A Theory of Marriage: Part I,” *Journal of Political Economy*, 81(4), 813–846.
- BURDETT, KENNETH, AND MELVYN COLES (1999): “Long-Term Partnership Formation: Marriage and Employment,” *Economic Journal*, 109(456), F307–F334.
- BURDETT, KENNETH, SHOUYONG SHI, AND RANDALL WRIGHT (2001): “Pricing and Matching with Frictions,” *Journal of Political Economy*, 109(5), 1060–1085.
- CASON, TIMOTHY, AND CHARLES NOUSSAIR (2003): “A Market with Frictions in the Matching Process: An Experimental Study,” Mimeo.
- COLES, MELVYN, AND JAN EECKHOUT (2000): “Heterogeneity as a Coordination Device,” Mimeo.
- DAVIS, STEVEN (1997): “Sorting, Learning, and Mobility When Jobs Have Scarcity Value,” *Carnegie-Rochester Conference Series on Public Policy*, 46(1), 327–338.
- (2001): “The Quality Distribution of Jobs and the Structure of Wages in Search Equilibrium,” NBER Working Paper 8434.
- GIBBONS, ROBERT, AND LAWRENCE KATZ (1992): “Does Unmeasured Ability Explain Inter-Industry Wage Differentials,” *Review of Economic Studies*, 59(3), 515–535.
- GROSHEN, ERICA (1991): “Sources of Intra-Industry Wage Dispersion: How Much Do Employers Matter?,” *Quarterly Journal of Economics*, 106(3), 869–884.
- HARSANYI, JOHN (1973): “Games with Randomly Disturbed Payoffs: A New Rationale for Mixed-Strategy Equilibrium Points,” *International Journal of Game Theory*, 2, 1–23.
- HOLZER, HARRY, LAWRENCE KATZ, AND ALAN KRUEGER (1991): “Job Queues and Wages,” *Quarterly Journal of Economics*, 106(3), 739–768.
- JULIEN, BENOIT, JOHN KENNES, AND IAN P. KING (2000): “Bidding for Labor,” *Review of Economic Dynamics*, 3(4), 619–649.
- (2001): “Auctions and Posted Prices in Directed Search Equilibrium,” *Topics in Macroeconomics (BE Journals in Macroeconomics)*, 1(1), Article 1.

- KOOPMANS, TJALLING, AND MARTIN BECKMANN (1957): “Assignment Problems and the Location of Economic Activities,” *Econometrica*, 25(1), 53–76.
- KRUEGER, ALAN, AND LAWRENCE SUMMERS (1988): “Efficiency Wages and the Inter-Industry Wage Structure,” *Econometrica*, 56(2), 259–293.
- KULTTI, KLAUS (1999): “Equivalence of Auctions and Posted Prices,” *Games and Economic Behavior*, 27(1), 106–113.
- LU, XIAOHUA, AND R. PRESTON MCAFEE (1996): “Matching and Expectations in a Market with Heterogeneous Agents,” in *Advances in Applied Microeconomics*, ed. by Michael Baye, vol. 6, pp. 121–156. JAI Press, Greenwich, CT.
- MOEN, ESPEN (1997): “Competitive Search Equilibrium,” *Journal of Political Economy*, 105(2), 385–411.
- MONTGOMERY, JAMES (1991): “Equilibrium Wage Dispersion and Interindustry Wage Differentials,” *Quarterly Journal of Economics*, 106(1), 163–179.
- MURPHY, KEVIN, AND ROBERT TOPEL (1987): “Unemployment, Risk and Earnings: Testing for Equalizing Wage Differences in the Labor Market,” in *Unemployment and the Structure of Labor Markets*, ed. by Kevin Lang, and Jonathan Leonard. Basil Blackwell, Oxford.
- OCHS, JACK (1990): “The Coordination Problem in Decentralized Markets: An Experiment,” *Quarterly Journal of Economics*, 105(2), 545–559.
- PETERS, MICHAEL (1991): “Ex Ante Price Offers in Matching Games: Non-Steady States,” *Econometrica*, 59(5), 1425–1454.
- ROTH, ALVIN, AND MARILDA SOTOMAYOR (1990): *Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis*. Cambridge University Press, Cambridge.
- SATTINGER, MICHAEL (1975): “Comparative Advantage and the Distributions of Earnings and Abilities,” *Econometrica*, 43(3), 455–468.
- (1993): “Assignment Models of the Distribution of Earnings,” *Journal of Economic Literature*, 31(2), 831–880.
- (1995): “Search and the Efficient Assignment of Workers to Jobs,” *International Economic Review*, 36(2), 283–302.
- SHAPLEY, LLOYD, AND MARTIN SHUBIK (1972): “The Assignment Game I: the Core,” *International Journal of Game Theory*, 1(2), 111–130.
- SHI, SHOUYONG (2001): “Frictional Assignment, Part I: Efficiency,” *Journal of Economic Theory*, 98(2), 232–260.

———— (2002): “A Directed Search Model of Inequality with Heterogeneous Skills and Skill-Biased Technology,” *Review of Economic Studies*, 69, 467–491.

SHIMER, ROBERT (1996): “Contracts in a Frictional Labor Market,” MIT Mimeo.

———— (1999): “Job Auctions,” Princeton Mimeo.

———— (2001): “The Assignment of Workers to Jobs in an Economy with Coordination Frictions,” NBER Working Paper 8501.

———— (2003): “Assignment and Unemployment,” Chicago Mimeo.

SHIMER, ROBERT, AND LONES SMITH (2000): “Assortative Matching and Search,” *Econometrica*, 68(2), 343–370.

———— (2001a): “Matching, Search, and Heterogeneity,” *Advances in Macroeconomics (BE Journals in Macroeconomics)*, 1(1), Article 5.

———— (2001b): “Nonstationary Search,” Princeton Mimeo.