University of Toronto
Department of Economics

Working Paper 308

On-the-Job Search and Business Cycles

By Guido Menzio and Shouyong Shi

January 21, 2008

On-the-Job Search and Business Cycles

GUIDO MENZIO
University of Pennsylvania

SHOUYONG SHI*
University of Toronto

October 2007

Abstract

In this paper, we develop a tractable model of the labor market where workers search for jobs both while unemployed and while on the job. Search is directed in the sense that each worker chooses to search for the offer that provides the optimal tradeoff between the probability of obtaining the offer and the increase in the value relative to the worker's current employment. There are both aggregate and match-specific shocks, on which the wage path in an offer can be contingent. We characterize the equilibrium analytically and show that the equilibrium is unique and socially efficient. On the quantitative side, we calibrate the model to the US data to measure the effect of aggregate productivity fluctuations on the labor market. We find that productivity fluctuations account for approximately 64% of the cyclical volatility in US unemployment. Moreover, productivity fluctuations generate the same matrix of correlations between unemployment and other labor market variables as in the US. In particular, the Beveridge curve is negatively sloped over business cycles, and the magnitude of the slope is the same as in the data. In light of these findings, we conclude that productivity shocks are one of the main forces driving labor market fluctuations over business cycles. Furthermore, we find that recessions have a cleansing effect on the economy.

JEL Codes: E24, E32, J64

Keywords: Directed Search, On the Job Search, Unemployment Fluctuations

*Menzio: Department of Economics, University of Pennsylvania, 160 McNeil Building, 3718 Locust Walk, Philadelphia, PA 19104, U.S.A. (email: gmenzio@sas.upenn.edu); Shi: Department of Economics, University of Toronto, 150 St. George Street, Toronto, Ontario, Canada, M5S 3G7 (email: shouyong@chass.utoronto.ca).

1 Introduction

Empirical evidence suggests that a substantial fraction of US workers move from job to job without any intervening spell of unemployment or non-employment. For example, Blanchard and Diamond (1989) find that approximately twenty percent of all the workers hired during a given month directly move from another job. Using NLSY data, Parsons (1991) finds that approximately fifty percent of the workers who voluntarily quit their job have another one lined up. Using CPS data, Nagypál (2004) finds that more than forty percent of the workers who, for whatever reason, have left their job during a given month are employed somewhere else at the end of the month.

This evidence has motivated the development of on-the-job search models of the labor market, i.e. models where workers search for new employment opportunities both while unemployed and while on the job. On-the-job search models have greatly contributed to our understanding of the US labor market. For example, partial equilibrium models with on-the-job search have been successfully used to explain the relationship between the worker's tenure on a job and his probability of leaving it (see Burdett 1978). And general equilibrium models with on-the-job search have been successfully used to explain the extent and shape of the distribution of wages across firms (see Burdett and Mortensen 1998, Mortensen 2003 and Burdett and Coles 2003).

Nevertheless, on-the-job search models are not commonly used to study the dynamics of the labor market over the business cycle. The reason is that—in all the existing general equilibrium models of on-the-job search (such as Pissarides 1994, Burdett and Mortensen 1998 and Barlevy 2002)—the worker's probability of transition from one employment state to another is a non-trivial function of the entire distribution of workers across jobs. Therefore, finding the equilibrium transition probabilities outside the steady-state is a task that is analytically hopeless and numerically daunting. Building on recent contributions by Delacroix and Shi (2006) and Shi (2007), this paper develops the first analytically tractable general equilibrium model with on-the-job search and aggregate fluctuations.

In our model, the economy is populated by a continuum of workers and a continuum of firms. Each worker supplies an indivisible unit of labor when employed, and each firm has a production technology with constant returns to scale. Workers and firms come together through a search-and-matching process. At the beginning of the period, a large number of "submarkets" open, each of which is defined by the terms-of-trade between firms and workers.

These terms can be contingent on both aggregate shocks and productivity shocks specific to the match that will be formed. Based on his current employment state, a worker decides in which submarket he wants to search for a new employment opportunity. Based on the distribution of workers across submarkets, each firm decides whether and where to create vacancies. These choices of the submarkets by workers and firms generate the tightness of each submarket, i.e. the ratio of vacancies to workers in that submarket. Because workers and firms can anticipate the dependence of the tightness of each submarket on the terms of trade in that submarket, their search decisions are *directed* by these terms of trade.

Once inside a submarket, workers and vacancies meet bilaterally. The probability that a worker meets a vacancy positively depends on the tightness of the submarket, while the probability that a vacancy meets a worker negatively depends on the tightness. If a worker and a vacancy meet, they form a production unit whose per-period output is the sum of an idiosyncratic and an aggregate component. If a worker does not meet a vacancy, he returns to his previous employment position. If a vacancy does not meet a worker, it fully depreciates.

In equilibrium, because of free entry of firms, the firm's benefit from creating an additional vacancy in any given submarket is equal to its cost. The cost of creating an additional vacancy is a constant. The benefit is the probability that a vacancy meets a worker, multiplied by the value of a filled vacancy. The meeting probability only depends on the tightness of the submarket, while the value of a filled vacancy only depends on the current realization of the aggregate productivity shock and the terms-of-trade prescribed in the submarket. Therefore, in any given submarket, the equilibrium tightness only depends on the realization of the aggregate productivity shock, and not on the distribution of workers across jobs.

In equilibrium, a worker visits the submarket where the benefit from searching is maximized. This benefit is the probability of meeting a vacancy, multiplied by the increase in the value of working in the firm. The meeting probability only depends on the tightness of the market, and the increase in the value of working in the firm only depends on the worker's previous employment position and on the terms-of-trade prescribed in the submarket. Therefore, the worker's equilibrium search strategy only depends on the tightness of the various labor markets and the worker's initial employment position.

Taken together, the two equilibrium conditions imply that the worker's probability of transition from one employment state to another depends on the current realization of the aggregate productivity shock but not on the distribution of workers across jobs. Therefore, in our model, finding the equilibrium transition probabilities is an analytically feasible and numerically straightforward task.

In the second part of the paper, we use our model to measure the effect of cyclical fluctuations in aggregate productivity on the US labor market. In order to choose parameter values, we require that the model's non-stochastic steady-state matches some features of the US labor market such as the average unemployment rate, the average transition rate from unemployment to employment and the relationship between the worker's tenure on a job and his probability of leaving it. In order to choose the persistence and variance of aggregate productivity shocks, we require that the stochastic version of the model matches the persistence and variance of the cyclical component of the US average productivity of labor.

By comparing model-generated and US data, we find that cyclical fluctuations in aggregate labor productivity account for approximately 64% of the cyclical volatility of US unemployment. Moreover, we find that cyclical fluctuations in aggregate productivity generate the same matrix of correlations between unemployment and other labor market variables as in the US. For example, in both model-generated and US data, unemployment features a strong negative correlation with vacancies (i.e., the Beveridge curve over business cycles), a negative correlation with the unemployment-to-employment transition rate and a positive correlation with the employment-to-unemployment transition rate. In light of these findings, we conclude that aggregate productivity shocks are likely to be one of the central forces behind business cycles.

Finally, we use the model to measure the effect of aggregate productivity shocks on job quality, i.e. the average of the idiosyncratic component of labor productivity over all active matches. We find that a one percent negative shock to aggregate productivity increases job quality by half a percentage point. Conversely, we find that a one percent positive shock to aggregate productivity decreases job quality by approximately one percent. In light of these findings, we conclude that recessions have a cleansing effect on the economy.

Our first result is in sharp contrast with Shimer (2005) who finds that aggregate productivity fluctuations account for approximately ten percent of the cyclical volatility of US unemployment. Our model's quantitative predictions match the data better than Shimer's do because our model, unlike Shimer's, allows for on-the-job search and match heterogeneity. Both additional elements are important for the quantitative performance of our model.

In particular, without on-the-job search, the model produces a positive contemporaneous correlation between unemployment and aggregate productivity.

Our second result is in sharp contrast with Barlevy (2002) who finds that recessions have a sullying effect on the economy. We conjecture that our result is different because our model allows for a continuum of realizations of the idiosyncratic productivity shock.

2 Market Equilibrium

2.1 The Model

The economy is populated by a continuum of workers with measure one and a continuum of firms whose measure is determined by competitive entry. Each worker maximizes the expected sum of discounted income, with a discount factor $\beta \in (0,1)$. Each firm maximizes the expected sum of discounted profits discounted, with the same discount factor β.

At the beginning of each period, the state of the economy is summarized by three variables: the aggregate component of labor productivity $y \in [\underline{y}, \overline{y}]$; the fraction of workers $u \in [0,1]$ who are unemployed; and the distribution of employed workers, represented by the function $G(.)$, where $G(z) \in [0,1]$ is the measure of workers whose jobs have an idiosyncratic component of labor productivity equal to or less than $z \in [\underline{z}, \overline{z}]$. It is useful to denote the state of the economy as $\omega = \{y, u, G(z)\}$.

Each period is divided in three stages: *separation*, *search* and *production*. In the first stage, each employed worker is exogenously displaced from his job and into unemployment with probability $\delta \in (0,1)$. Also, during the separation stage, each employed worker can voluntarily quit his job to become unemployed and each firm can voluntarily dismiss his employee.

In the second stage, workers and firms come together via search. First, each worker finds out whether he can search the labor market in the current period. In particular, if the worker was unemployed at the beginning of the separation stage, he can search the labor market with probability $\lambda_u \in (0,1]$. If the worker was employed at the end of the separation stage, he can search the labor market with probability $\lambda_e \in [0,1]$. If the worker became unemployed during the separation stage, he cannot seek for a new job during the search stage. Secondly, if a worker has the opportunity to search, he chooses which submarket $W \in \mathbb{R}$ to visit. Finally, after having observed the locations of all workers, firms choose how many vacancies to create in each submarket W. The cost of a vacancy is $k \in \mathbb{R}_{++}$.

Let $\theta(W; y) \in \mathbb{R}_+$ denote the ratio between the number of vacancies and the number of workers in the submarket W. Following Pissarides (1985), we refer to $\theta(W; y)$ as the *tightness* of submarket W. We assume that a worker visiting submarket W is hired with probability $p(\theta(W; y))$, where $p : \mathbb{R}_+ \to [0, 1]$ is a twice continuously differentiable function such that $p' > 0$, $p'' < 0$, $p(0) = 0$ and $p'(0) = 1$. Conversely, we assume that a vacancy created in submarket W is filled with probability $q(\theta(W; y)) \equiv \theta(W; y)^{-1} \cdot p(\theta(W; y))$, where $q : \mathbb{R}_+ \to [0, 1]$ is a twice continuously differentiable function such that $q' < 0$, $q(0) = 1$ and $\lim_{\theta \to \infty} q(\theta) = 0$. Perhaps, the reader may find useful to interpret $p(\theta(W; y))$ as $u^{-1} m(u, \theta(W; y)u)$, where $m(u, v)$ is a constant return to scale function that tells how many matches are created when u workers and v vacancies visit submarket W.

If a worker visiting submarket W is not hired, he maintains the employment state he had before the search stage. If a vacancy created in submarket W is not filled, it fully depreciates. And if a match is formed in submarket W, the firm offers to the worker an employment contract that provides him with the lifetime utility W. The contract specifies the probability of employment and the terms-of-trade between the firm and the worker in the current period and—conditional on the survival of the employment relationship—in future periods. As in Burdett and Coles (2003) and Stevens (2004), the terms-of-trade can be contingent on the aggregate state of the economy, the idiosyncratic productivity of the match and the length of the relationship, but they cannot depend on the worker's outside offers. Moreover, either party retains the option to unilaterally terminate the employment relationship during the separation stage. Finally, at the end of the search stage, each new match draws its idiosyncratic component of labor productivity z from the cumulative distribution function $F(z)$. For simplicity, we assume that this match specific productivity will remain the same as long as the worker and the firm keep the match.[1]

In the third stage, production takes place, workers receive their income and firms realize their profits. In particular, if a worker is unemployed, his income is equal to b units of output. If a worker is employed at a job with idiosyncratic quality z, he produces $y + z$ units of output. The division of output between the worker and the firm is determined by the contractual agreement that the two parties have signed upon forming the production unit.

[1] In contrast to our modeling, Mortensen and Pissardies (1994) assume that a new match always starts with the highest match specific productivity, and then experiences new draws of z in future periods.

Finally, at the end of the production stage, nature draws next period's aggregate component of productivity y_+ from the cumulative distribution function $\Phi(y_+|y)$. Figure 1 illustrates the timing of events.

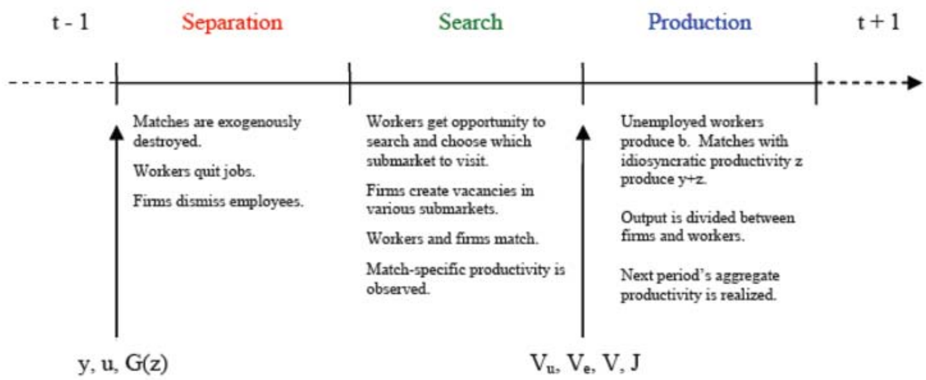

Figure 1: Timing of Events

2.2 Worker's Problem

Consider a worker at the beginning of the search stage. If the worker does not get an opportunity to search in the labor market, he enters the production stage in the same employment position that he held at the end of the separation stage. In this case, the worker's lifetime utility at the beginning of the production stage is denoted as V. We refer to V as the worker's *fall-back position.* If the worker does get an opportunity to search, he chooses the submarket in which to search, W. With probability $p(\theta(W;y))$, the worker finds a new job and his continuation utility is W. With probability $1-p(\theta(W;y))$, the worker does not find a new job, he returns to his previous employment position, and his lifetime utility is V at the beginning of the production stage. Therefore, the worker chooses W in order to solve the following maximization problem:

$$D(V;y) = \max_{W} p(\theta(W;y))\,(W-V)\,. \tag{1}$$

We denote with $W(V;y)$ the solution for W to (1).

Now, consider a worker who enters the production stage without a job and denote with $V_u(y)$ his lifetime utility. During the current stage, the worker obtains b from unemployment. During the next search stage, the worker gets the opportunity to look for a job with probability λ_u. In this case, the worker's expected lifetime utility at the beginning of the next production stage is $V_u(y_+) + D(V_u(y_+); y_+)$. If the worker fails to get the search opportunity, the worker's lifetime utility at the beginning of the next production stage is $V_u(y_+)$. Therefore, $V_u(y)$ is given by:

$$V_u(y) = b + \beta \int [V_u(y_+) + \lambda_u D(V_u(y_+); y_+)] d\Psi(y_+|y). \tag{2}$$

Next, consider a worker who enters the production stage with a job. In this case, the worker's lifetime utility, V_e, is determined by the prescriptions of his employment contract. In particular, let the employment contract prescribe that the worker should receive the wage w in the current period, that the match should be destroyed with probability $d(y_+)$ during the next separation stage, and that the worker's lifetime utility at the next production stage should be $V_+(y_+)$, conditional on the survival of the match. Then, V_e is given by:

$$\begin{aligned} V_e\left(w, d_+, V_+; y\right) \quad = \quad & w + \beta \int d(y_+) V_u(y_+) d\Psi(y_+|y) + \\ & \beta \int (1 - d(y_+))[V_+(z; y_+) + \lambda_e D(V_+(z; y_+); y_+)] d\Psi(y_+|y). \end{aligned} \tag{3}$$

Finally, consider a worker who is employed at the beginning of the production stage and has complete ownership over the output of his job. Denote with $V(z;y)$ the lifetime utility of this worker. During the current stage, the worker produces $y + z$ units of output. During the next separation stage, the worker is exogenously displaced from his job with probability δ and he is not with probability $1 - \delta$. In the first case, the worker's lifetime utility at the beginning of the next production stage is $V_u(y_+)$. In the second case, the worker can choose to voluntarily quit his job or to keep it. If the worker quits, his lifetime utility at the beginning of the next production stage is $V_u(y_+)$. If the worker does not quit, his lifetime utility at the beginning of the next production stage is $V(z; y_+) + \lambda_e D(V(z; y_+); y_+)$. Therefore $V(z;y)$ is given by

$$\begin{aligned} V(z;y) \quad = \quad & y + z + \beta\delta \int V_u(y_+) d\Psi(y_+|y) + \\ & \beta(1-\delta) \int \max\left\{V(z; y_+) + \lambda_e D(V(z; y_+); y_+),\ V_u(y_+)\right\} d\Psi(y_+|y). \end{aligned} \tag{4}$$

Define $z_*(y)$ as the level of idiosyncratic productivity such that the worker is indifferent between keeping and quitting the job to get unemployed, i.e.

$$V(z_*(y);y) + \lambda_e D(V(z_*(y);y);y) = V_u(y). \tag{5}$$

We will show later that a worker will quit the job to get unemployed if and only if $z < z_*(y)$.

2.3 Firm's Problem

At the beginning of the search stage, a firm decides whether and where to create a vacancy. In making this decision, the firm takes as given the set $\mathcal{W}$ of submarkets that are populated by a positive number of workers and the equilibrium tightness function $\theta(W;y)$. If it decides to create a vacancy in submarket $W \in \mathcal{W}$, the firm first pays the cost k. Then, the firm finds a worker to fill its vacancy with probability $q(\theta(W;y))$ and does not with probability $1-q(\theta(W;y))$. In the first case, the firm's expected profits at the production stage are $J_0(W;y)$. In the second case, the firm's expected profits at the production stage are zero. If it decides against creating a vacancy, the firm does not sustain any cost at the search stage and does not make any revenue at the production stage. Overall, the firm's expected profits at the beginning of the search stage are given by:

$$\max\{0, -k + \max_{W\in\mathcal{W}} q(\theta(W;y))J_0(W;y)\}. \tag{6}$$

In equilibrium, (6) is equal to zero because firms are free to enter the labor market.

Next, consider a firm that has just met a worker in submarket W. At this stage, the firm chooses the function $V_0 : [\underline{z}, \overline{z}] \times Y \to \mathbb{R}$, where $V_0(z)$ is the lifetime utility that the employment contract offers to the worker conditional on the realization z of the idiosyncratic productivity shock. The firm's choice is constrained by the requirement that, in expectation, the employment contract must offer the worker the lifetime utility W, i.e. $\int V_0(z;y)dF(z) = W$. Finally, the firm's choice maximizes the expectation of $J(V_0(z;y), z;y)$ with respect to the realization z of the idiosyncratic productivity shock, where $J(V_0(z;y), z;y)$ is the profit that a firm can expect from having a match with quality z and an obligation to deliver its employee the lifetime utility V_0. Therefore, $J_0(W;y)$ is given by:

$$J_0(W;y) = \max_{V_0(z)} \int J(V_0(z;y), z;y)dF(z), \text{ s.t. } W = \int V_0(z;y)dF(z). \tag{7}$$

Finally, consider a firm that has just entered the production stage with a match of quality z and an obligation to deliver the lifetime utility V to its employee. The firm designs the employment contract—i.e. it chooses the current wage w, next period's separation probability $d(y_+)$ and the worker's continuation value $V_+(z;y_+)$—in order to maximize the discounted sum of profits. The firm's choice is subject to the *promise-keeping* constraint—i.e. the contract has to provide the worker with the lifetime utility V—and the *individual rationality* constraint—i.e. the separation probability $d(y_+)$ has to be consistent with the worker's incentives to quit and the firm's incentives to destroy the match. Therefore, $J(V,z;y)$ is given by

$$\begin{aligned} J(V,z;y) = &\max_{w,d\in[\delta,1],V_+} y+z-w+ \\ &\beta \textstyle\int [1-d(y_+)]\left[1-\lambda_e p(\theta(W(V_+(z;y_+);y_+)))\right] J(V_+(z;y_+),z;y_+)d\Psi(y_+|y), \text{ s.t.} \\ &V = V_e(w,d(y_+),V_+(z;y_+);y), \\ &d(y_+) = 1 \text{ if } V_+(z;y_+)+\lambda_e D(V_+(z;y_+);y_+) < V_u(y_+) \text{ or } J(V_+(z;y_+),z;y_+)<0. \end{aligned} \tag{8}$$

2.4 The Optimal Contract

First, we consider the contracting problem (8). We denote with $M(V,z;y)$ the joint value of the employment relationship between the firm and the worker, i.e. $M(V,z;y) = J(V,z;y)+V$. From (8) and (3), it follows that $M(V,z;y)$ is given by

$$\begin{aligned} M(V,z;y) = &\, y+z+\beta \textstyle\int d(y_+)V_u(y_+)d\Psi(y_+|y)+ \\ &\beta \textstyle\int (1-d(y_+))M(V_+(z;y_+),z;y_+)d\Psi(y_+|y)+ \\ &\beta \textstyle\int (1-d(y_+))\lambda_e p(\theta(W(V_+(z;y_+);y_+))))[W(V_+(z;y_+);y_+)-M(V_+(z;y_+),z;y_+)]d\Psi(y_+|y). \end{aligned} \tag{9}$$

Since the worker does not care about the firm's profits, the search strategy $W(V_+(z;y_+);y_+)$ does not necessarily maximize the joint value of the employment relationship. Similarly, because one partner does not care about the other's share of the match output, the separation strategy $d(y_+)$ does not necessarily maximize the joint value of the employment relationship. Therefore, $M(V,z;y)$ is smaller or equal than

$$\begin{aligned} M(V,z;y) \le &\, y+z+\beta\delta \textstyle\int V_u(y_+)d\Psi(y_+|y)+ \\ &\beta(1-\delta) \textstyle\int \max\left\{V_u(y_+), M(V_+(z;y_+),z;y)+\lambda_e D(M(V_+(z;y_+),z;y_+);y_+)\right\} d\Psi(y_+|y). \end{aligned} \tag{10}$$

From the comparison of (10) and (4), it is immediate to conclude that the joint value of the employment relationship—namely, $M(V,z;y)$—cannot be greater than the worker's value of the production unit—namely, $V(z;y)$. And, consequently, the firm's expected profits $J(V,z;y)$ cannot be greater than $V(z;y)-V$.

Now, consider the employment contract $\{w^*, d^*, V_+^*\}$, where the current wage w^* is equal to $y+z+V-V(z;y)$, the separation probability d^* is equal to one if $V_u(y_+)$ is greater than $V(z;y_+)+\lambda_e D(V(z;y_+);y_+)$ and equal to δ otherwise, and the worker's continuation value $V_+^*(y_+)$ is equal to $V(z;y)$. The contract $\{w^*, d^*, V_+^*\}$ satisfies the promise-keeping constraint because $V_e\left(w^*, d^*, V_+^*; y\right)$ is equal to V. The separation probability $d^*(y_+)$ satisfies the individual rationality constraint because the worker's continuation value is equal to $V(z;y)$ and the firm's continuation profits are equal to zero. Finally, given $\{w^*, d^*, V_+^*\}$, the sum of the firm's expected profits and the worker's lifetime utility is equal to $V(z;y)$. Hence, the employment contract $\{w^*, d^*, V_+^*\}$ is the solution to the firm's problem (8).

Next, we consider the contracting problem (7). The firm's expected profits $J_0(W;y)$ are given by

$$\begin{aligned} J_0(W;y) = \quad & \left[\max_{V_0(z)} \int J(V_0(z;y),z;y)dF(z), \text{ s.t. } W=\int V_0(z;y)dF(z)\right] = \\ & \left[\max_{V_0(z)} \int [V(z;y)-V_0(z;y)]dF(z), \text{ s.t. } W=\int V_0(z;y)dF(z)\right] = \\ & \int V(z;y)dF(z) - W. \end{aligned} \tag{11}$$

The second line in (11) uses the fact that $J(V_0,z;y)$ is equal to $V(z;y)-V_0$. The third line in (11) states that the firm's profits $J_0(W;y)$ are equal to the difference between the expected value of the production unit, $\int V(z;y)dF(z)$, and the lifetime utility, W, that the firm has to provide to a worker hired in submarket W. Finally, notice that the solution to the contracting problem (7) is indeterminate.

2.5 Laws of Motion

Given the properties of the optimal employment contract, we can derive the probability with which a worker transits from one employment state to another. First, consider a worker who is unemployed at the beginning of the period. Define $\theta_u^*(y)=\theta(W(V_u(y);y);y)$. With probability $1-\lambda_u p(\theta_u^*(y))$, the worker does not find an employment opportunity during the search stage and, at the beginning of next period, he will still be unemployed. With probability $\lambda_u p(\theta_u^*(y))F(z)$, the worker finds an employment opportunity during the search stage and, at the beginning of next period, he will be employed at a job with idiosyncratic productivity equal to or less than z.

Next, consider a worker who is employed at a job with productivity z at the beginning of the period. If z is smaller than $z_*(y)$, the worker quits his job during the separation stage and does not have the opportunity to find another job during the search stage. At the beginning of next period, the worker will be unemployed. If z is greater than $z_*(y)$, the worker is exogenously displaced from his job with probability δ. In this case, the worker will be unemployed at the beginning of next period. With probability $(1-\delta)(1-\lambda_e p(\theta^*(z;y)))$, where $\theta^*(z;y)$ is defined as $\theta(W(V(z;y);y);y)$, the worker keeps his job during the separation stage and does not find a new employment opportunity during the search stage. At the beginning of next period, the worker will still be employed at a job with idiosyncratic productivity z. Finally, with probability $(1-\delta)\lambda_e p(\theta^*(z;y))F(z)$, the worker keeps his job during the separation stage and finds a new employment opportunity during the search stage. At the beginning of next period, the worker will be employed at a job with idiosyncratic productivity smaller or equal than z.

Given the state of the economy at the beginning of the period, $\omega = \{y, u, G(z)\}$, and given the worker's transition probabilities, we can derive the state of the economy at the beginning of next period, $\omega_+ = \{y_+, u_+(\omega), G_+(z;\omega)\}$. In particular, the measure of workers who will be unemployed at the beginning of next period is:

$$u_+(\omega) = u + (1-u)\left[\delta + (1-\delta)G(z_*(y))\right] - u\lambda_u p^*(\theta_u^*(y)), \tag{12}$$

For $z < z_*(y)$, the measure of employed workers who will have a job at the beginning of next period with idiosyncratic productivity equal to or less than z is:

$$G_+(z;\omega) = \left[u\lambda_u p(\theta_u^*(y)) + (1-u)(1-\delta)\lambda_e \int_{z_*(y)}^{\overline{z}} p(\theta^*(\tilde{z};y))dG(\tilde{z})\right] \tfrac{F(z)}{1-u_+(\omega)}. \tag{13}$$

For $z \geq z_*(y)$, the measure of employed workers who will have a job at the beginning of next period with idiosyncratic productivity equal to or less than z is:

$$\begin{aligned} G_+(z;\omega) = & \left[u\lambda_u p(\theta_u^*(y)) + (1-u)(1-\delta)\lambda_e \int_{z_*(y)}^{\overline{z}} p(\theta^*(\tilde{z};y))dG(\tilde{z})\right] \tfrac{F(z)}{1-u_+(\omega)} + \\ & \left[\int_{z_*(y)}^{z} \left[1-\lambda_e p(\theta^*(\tilde{z};y))\right] dG(\tilde{z})\right] \tfrac{(1-u)(1-\delta)}{1-u_+(\omega)}. \end{aligned} \tag{14}$$

2.6 Definition of an Equilibrium

Definition 1: A market equilibrium is a tuple $D(V;y)$, $W(V;y)$, $V_u(y)$, $V(z;y)$, $z_*(y)$, $\theta(W;y)$, $u_+(\omega)$, $G_+(z;\omega)$ which satisfies the following properties:

(i) Optimal search: $D(V;y)$ and $W(V;y)$ are, respectively, the value and the policy function associated with the maximization problem (1).

(ii) Optimal quitting: $V_u(y)$ and $V(z;y)$ are the solution to the system of functional equations (2) and (4). Moreover, $z_*(y)$ is the solution to the equation (5).

(iii) Profit maximization: For all W,

$$-k + q(\theta(W;y)) \left[\int V(z;y)dF(z) - W\right] \leq 0$$

and $\theta(W;y) \geq 0$, with complementary slackness.

(iv) Consistent aggregation: $u_+(\omega)$ and $G_+(z;\omega)$ are given by (12) and (13)–(14).

Condition (i) guarantees that the worker's search strategy is consistent with utility maximization given the equilibrium market tightness $\theta(W;y)$. Condition (ii) guarantees that the worker's quitting strategy is consistent with utility maximization given the properties of the optimal employment contract. Condition (iii) guarantees that the equilibrium market tightness $\theta(W;y)$ is consistent with the firms' optimal decisions about the creation and location of vacancies. Finally, condition (iv) guarantees that the law of motion for the state of the economy is consistent with the decisions made by individual agents.

3 Social Planner's Problem

3.1 Formulation of the Problem

At the beginning of each period, the physical state of the economy can be summarized by two state variables: the aggregate productivity $y \in [\underline{y}, \overline{y}]$ and the derivative $h : [\underline{z}, \overline{z}] \to \mathbb{R}_+$ of the measure of workers who are employed at jobs with idiosyncratic productivity smaller than z. The measure of workers u who begin the period without a job is recovered from the *adding-up constraint*:

$$u + \int h(z)dz = 1. \tag{15}$$

In each period, the social planner makes three choices. First, it chooses the tightness $\hat{\theta}_u \in \mathbb{R}_+$ in the labor market visited by the workers whose fall-back position is unemployment. Secondly, it chooses the tightness $\hat{\theta} : [\underline{z}, \overline{z}] \times Y \to \mathbb{R}_+$ in the market visited by the workers whose fall-back position is being employed at a job with idiosyncratic productivity z. Finally, the planner chooses the productivity level $\hat{z}_*$ below which workers are instructed to quit their job.

The amount of output C available for consumption in the current period is determined by the initial state of the economy $\{h, y\}$ and by the planner's choices $\{\hat{\theta}_u, \hat{\theta}, \hat{z}_*\}$. More specifically, C is given by

$$C = -k\left[\lambda_u u\hat{\theta}_u + \lambda_e(1-\delta)\int_{\hat{z}_*}^{\hat{\bar{z}}} \hat{\theta}(z)h(z)dz\right] + u_+ b + \int_{\underline{z}}^{\bar{z}} (z+y)h_+(z)dz. \tag{16}$$

The first term on the right-hand-side of (16) measures the output that is invested in the creation of new vacancies. The second term measures the amount of output that is home-produced by unemployed workers. And the third term measures the amount of output that is produced by employed workers.

In the next period, the state of the economy $\{h_+, y_+\}$ is determined by $\{h, y\}$ and by the planner's choices $\{\hat{\theta}_u, \hat{\theta}, \hat{z}_*\}$. More specifically, next period's aggregate productivity y_{+1} is a random variable distributed according to the distribution function $\Psi(y_+|y)$. And at the beginning of the next period, the density of workers $h_+ : [\underline{z}, \overline{z}] \to \mathbb{R}_+$ across jobs with different productivity is given by

$$\begin{aligned} h_+(z) &= \left[\lambda_u up(\hat{\theta}_u) + \lambda_e(1-\delta)\textstyle\int_{\hat{z}_*}^{\hat{\bar{z}}} p(\hat{\theta}(z))h(z)dz\right]F'(z) + \Delta(z), \\ \Delta(z) &= \begin{cases} 0 & \text{if } z \leq \hat{z}_*, \\ (1-\delta)(1-\lambda_e p(\hat{\theta}(z)))h(z) & \text{if } z \geq \hat{z}_*. \end{cases} \end{aligned} \tag{17}$$

Finally, the measure u_{+1} of workers who are unemployed at the beginning of the next period can be recovered from the following adding-up constraint:

$$u_+ + \textstyle\int h_+(z)dz = 1. \tag{18}$$

The planner 's objective function is the sum of current and future consumption discounted with the factor, β. Therefore, the planner's value in state $\{h, y\}$ is given by

$$S(h, y) = \max_{\{\hat{\theta}_u, \hat{\theta}, \hat{z}_*\}} C + \beta \textstyle\int S(h_+, y_+)d\Psi(y_+|y), \text{ s.t. (15)–(18).} \tag{19}$$

We conjecture that the Volterra derivative of $S(h, y)$ with respect to $h(z)$ is increasing in the productivity z and independent from the function h, i.e.,

$$\frac{\partial S(h, y)}{\partial(h(z))} = \hat{s}(z, y).$$

This is derivative is the option value to the society of an employed worker in a match (z, y), as opposed to keeping the worker unemployed.

3.2 Solution to the Problem

In any submarket, the optimal tightness (if positive) is such that the cost and benefit of creating an additional vacancy are equalized. The cost of creating an additional vacancy is k. The benefit of creating an additional vacancy is the product between the increase in the number of workers who find a successful match and the increase in the output produced by each one of them. Therefore, in the submarket visited by unemployed workers, the optimality condition for $\hat{\theta}_u(y)$ is:

$$k = p'(\hat{\theta}_u(y)) \left[\int_{\underline{z}}^{\overline{z}} \hat{\mu}(z; y) dF(z) - \hat{\mu}_u(y) \right], \tag{20}$$

where $\hat{\mu}_u(y)$ is the social value per period for an unemployed worker, and $\hat{\mu}(z; y)$ is the social value per period generated by an employed worker in a match with productivity (z, y). These social values are given, respectively, as

$$\hat{\mu}_u(y) = b,\ \hat{\mu}(z; y) = z + y + \beta \int \hat{s}(z; y_+) d\Psi(y_+|y). \tag{21}$$

Note that the social value of employing a worker in a match (z, y) contains not only the level of output in the match, but also the option value of the employed worker in the future relative to the social value of an unemployed worker. In the submarkets visited by workers employed at jobs with productivity z, the optimality condition for $\hat{\theta}(z; y)$ is:

$$\begin{aligned} & k = p'(\hat{\theta}(z; y)) \left[\int_{\underline{z}}^{\overline{z}} \hat{\mu}(z; y) dF(z) - \hat{\mu}(z; y) \right] && \text{if } z \leq \hat{z}^*(y), \\ & \hat{\theta}(z; y) = 0 && \text{if } z > \hat{z}^*(y). \end{aligned} \tag{22}$$

The optimal cutoff $\hat{z}_*(y)$ equates the output of a worker who is employed at the beginning of the search stage at a job with idiosyncratic productivity $\hat{z}_*(y)$ to the output of a worker who is unemployed at the production stage. Formally, the optimality condition for $\hat{z}_*(y)$ is:

$$\hat{\mu}(\hat{z}_*(y); y) + \lambda_e k \left[\frac{p(\hat{\theta}(\hat{z}_*(y); y))}{p'(\hat{\theta}(\hat{z}_*(y); y))} - \hat{\theta}(\hat{z}_*(y); y) \right] - \mu_u(y) = 0. \tag{23}$$

Conversely, the optimal cutoff $\hat{z}^*(y)$ equates the output of a worker who is employed at the production stage at a job with productivity $\hat{z}^*(y)$ and the expected output of a worker who has just been hired at a new job (net of the vacancy cost). Formally, the optimality condition for $\hat{z}^*(y)$ is:

$$\int_{\underline{z}}^{\overline{z}} \hat{\mu}(z;y)dF(z) - \hat{\mu}(\hat{z}^*(y);y) - k = 0. \tag{24}$$

Finally, from the envelope condition for the value function $S(h,y)$, it follows that the derivative $\hat{s}(z;y)$ satisfies the following condition:

$$\hat{s}(z;y) = -\lambda_u k\left(\frac{p(\hat{\theta}_u(y))}{p'(\hat{\theta}_u(y))} - \hat{\theta}_u(y)\right) + \Delta(z;y),$$

$$\Delta(z;y) = \begin{cases} 0 & \text{if } z \leq \hat{z}_*(y), \\ (1-\delta)\left[\hat{\mu}(z;y) - \hat{\mu}_u(y) + \lambda_e k\left(\frac{p(\hat{\theta}(z;y))}{p'(\hat{\theta}(z;y))} - \hat{\theta}(z;y)\right)\right] & \text{if } z \in (\hat{z}_*(y), \hat{z}^*(y)), \\ (1-\delta)\left[\hat{\mu}(z;y) - \hat{\mu}_u(y)\right] & \text{if } z \geq \hat{z}^*(y). \end{cases} \tag{25}$$

The derivative $\hat{s}(z;y)$ is the difference between the output of a worker who is employed at the beginning of the period at a job with productivity z and the output of a worker who is unemployed at the beginning of the period. According to condition (25), when z is smaller than $\hat{z}_*(y)$, $\hat{s}(z;y)$ is the negative of the option value of searching while unemployed. When z is greater than $\hat{z}_*(y)$, $\hat{s}(z;y)$ is the sum of two terms. The first term is the difference between the option value of searching while employed at a job z and searching while unemployed. The second term is the difference between the output of a worker who is employed at a job z at the production stage and a worker who is unemployed at the production stage.

3.3 Conditions for a Market Equilibrium

In this subsection, we derive a set of conditions that are necessarily satisfied by any market equilibrium. First, let $s^*(z;y)$ denote the the difference between the lifetime utility of a worker who is employed at the beginning of the period at a job with idiosyncratic productivity z and the utility of a worker who does not have a job at the beginning of the period. Formally, let $s^*(z;y)$ denote:

$$s^*(z;y) = (1-\delta)\max\{V(z;y) - V_u(y) + \lambda_e D(V(z;y);y), 0\} - \lambda_u D(V_u(y);y). \tag{26}$$

Secondly, let $\mu_u^*(y)$ and $\mu^*(z;y)$ be defined as

$$\mu_u^*(y) = b,\ \mu^*(z;y) = z + y + \beta\int s^*(z;y_+)d\Psi(y_+|y). \tag{27}$$

From condition (iii) in the definition of an equilibrium, it follows that the market tightness $\theta(W;y)$ is equal to zero whenever W is greater than $\int V(z;y)dF(z)-k$. And for all W smaller than $\int V(z;y)dF(z)-k$, the market tightness $\theta(W;y)$ equates the cost and benefit of creating an additional vacancy, i.e.

$$k = q(\theta(W;y))\left[\int V(z;y)dF(z) - W\right]. \tag{28}$$

Rearranging (28) and using the fact that $q(\theta)\cdot\theta = p(\theta)$, we obtain that

$$W = \int V(z;y)dF(z) - k\frac{\theta(W;y)}{p(\theta(W;y))}. \tag{29}$$

Differentiating (29) with respect to W, we obtain that

$$\theta'(W;y) = \frac{p(\theta(W;y))}{p'(\theta(W;y))\left[\int V(z)dF(z) - W\right] - k}. \tag{30}$$

From condition (i) in the definition of an equilibrium, it follows that the worker's search strategy $W(V;y)$ equates the cost and benefit of seeking a marginally better job. The cost of seeking a marginally better job is the product between the lower probability of being hired, i.e., $-p'(\theta(W;y))\cdot\theta'(W;y)$, and the value of being hired, i.e., $W(V;y)-V$. The benefit of seeking a marginally better job is the higher value of the job, i.e., $p(\theta(W;y))$. Therefore, the worker's search strategy $W(V;y)$ is such that

$$\theta'(W;y)\,(W-V) = -\frac{p(\theta(W;y))}{p'(\theta(W;y))}. \tag{31}$$

If the worker is unemployed, the optimality condition (31) can be written as

$$\begin{aligned} k = p'(\theta_u^*(y))\left[\int V(z;y)dF(z) - V_u(y)\right] = \\ p'(\theta_u^*(y))\left[\int \mu^*(z;y)dF(z) - \mu_u^*(y)\right], \end{aligned} \tag{32}$$

where the first line uses (29) and (30), and the second line uses (26) and (27). If the worker is employed at a job with idiosyncratic productivity $z \leq z^*(y)$, where $z^*(y)$ is such that $V(z^*(y);y) = \int V(z;y)dF(z) - k$, the optimality condition (31) can be written as

$$\begin{aligned} k = p'(\theta^*(z;y))\left[\int V(z;y)dF(z) - V(z;y)\right] = \\ p'(\theta^*(z;y))\left[\int \mu^*(z;y)dF(z) - \mu_u^*(z;y)\right]. \end{aligned} \tag{33}$$

Finally, if the worker is employed at a job with idiosyncratic productivity $z \geq z^*(y)$, the optimality condition (31) implies that $W(V(z;y);y)$ is greater than $\int V(z;y)dF(z) - k$.

From the optimality conditions (32) and (33), it follows that the expected gain from searching while unemployed (i.e., $D(V_u(y);y)$) and the expected gain from searching while employed at a job with idiosyncratic productivity $z \leq z^*(y)$ (i.e., $D(V(z;y);y)$) can be written as

$$\begin{aligned} D(V_u(y);y) &= k\left(\frac{p(\theta_u^*(y))}{p'(\theta_u^*(y))} - \theta_u^*(y)\right), \\ D(V(z;y);y) &= k\left(\frac{p(\theta^*(z;y))}{p'(\theta^*(z;y))} - \theta^*(z;y)\right). \end{aligned} \tag{34}$$

Finally, from the optimality condition (31), it follows that the option value of searching while employed at a job with idiosyncratic productivity $z \geq z^*(y)$ is zero.

From condition (ii) in the definition of an equilibrium, it follows that a worker employed at a job with idiosyncratic productivity z voluntarily moves into unemployment whenever z is smaller or equal than $z_*(y)$, where $z_*(y)$ is implicitly defined by equation (**??**). Using (26)–(27) and (34), we can rewrite equation (**??**) as

$$\mu^*(z_*(y);y) + \lambda_e k\left[\frac{p(\theta^*(z_*(y);y))}{p'(\theta^*(z_*(y);y))} - \theta^*(z_*(y);y)\right] - \mu_u(y) = 0. \tag{35}$$

From conditions (i) and (iii) in the definition of an equilibrium, it follows that a worker employed at a job with idiosyncratic productivity $z \geq z^*(y)$ does not move to another job during the search stage. Using (26)–(27) and (29), we can rewrite the equation that defines $z^*(y)$ as

$$\int \mu^*(z;y)dF(z) - \mu^*(z^*(y);y) - k = 0. \tag{36}$$

Finally, using (2), (4), (27) and (34), we can rewrite (26) as

$$s^*(z;y) = -\lambda_u k\left(\tfrac{p(\theta_u^*(y))}{p'(\theta_u^*(y))} - \theta_u^*(y)\right) + \Delta(z;y),$$

$$\Delta(z;y) = \begin{cases} 0 & \text{if } z \leq z_*(y), \\ (1-\delta)\left[\mu^*(z;y) - \mu_u^*(y) + \lambda_e k\left(\tfrac{p(\theta^*(z;y))}{p'(\theta^*(z;y))} - \theta^*(z;y)\right)\right] & \text{if } z \in (z_*(y), z^*(y)), \\ (1-\delta)\left[\mu^*(z;y) - \mu_u^*(y)\right] & \text{if } z \geq z^*(y). \end{cases} \tag{37}$$

3.4 Efficiency and Uniqueness of the Market Equilibrium

Taken together, (20)–(25) are a system of seven equations in the seven unknowns $\hat{\theta}_u(y)$, $\hat{\theta}(z;y)$, $\hat{z}_*(y)$, $\hat{z}^*(y)$, $\hat{s}(z;y)$, $\hat{\mu}(z;y)$ and $\hat{\mu}_u(y)$. Using the concavity of the job-finding probability $p(\theta)$ and the uniqueness of the solution of the Bellman Equation (19), we verify that this system of equations has a unique solution. Then, since the necessary conditions of an equilibrium (27), (32)–(33) and (35)–(37) constitute the same system as (20)–(25), we conclude that the market equilibrium is unique and efficient. Proposition 2 contains the details of this argument.

Proposition 1: (Uniqueness and Efficiency of the Equilibrium) *Any market equilibrium generates a unique and efficient allocation of labor* $\theta_u^*(y)$, $\theta^*(z;y)$, $z_*(y)$.

Proof: In the Appendix. $\|$

4 Quantitative Analysis

4.1 Calibration of the Model

In order to simulate the model, we need to choose values and functional forms for: (i) households' preferences $\{\beta, b\}$, (ii) matching process $\{\lambda_u, \lambda_e, m(u,v), \delta\}$, (iii) production technology $\{k, F(z), \Psi(y'|y)\}$. Following most of the labor search literature, we choose $m(u,v)$ to be a Cobb-Douglas matching function $Mu^\gamma v^{1-\gamma}$. Following most of the real business cycle literature, we choose the distribution $\Psi(y'|y)$ of the current level of aggregate productivity y' conditional on the past realization y to be normal with mean $\overline{y} + \rho(y - \overline{y})$ and standard deviation σ_y. Finally, we choose the distribution $F(z)$ of match-specific productivity to be normal with mean zero and standard deviation σ_z. Given these functional specifications, we have twelve parameters to identify.

Three of these twelve parameters can be normalized. First, the average aggregate productivity level, $\overline{y}$, affects only the units of $\{\alpha k, \alpha^2\sigma_z, \alpha\overline{y}, \alpha^2\sigma_y, b\}$, and so it can be normalized to 1. Second, notice that the model's outcomes are identical for all parameterization $\{\beta, b\}$, $\{\alpha^{-1}\lambda_u, \alpha^{-1}\lambda_e, \alpha M, \gamma, \delta\}$, $\{\alpha k, \sigma_z, \sigma_y, \rho\}$, where $\alpha > 0$. Therefore, the coefficient M in front of the matching function can be normalized to 1. Finally, notice that the model's outcomes are identical for all parameterization $\{\beta, b\}$, $\{\alpha\lambda_u, \alpha\lambda_e, \gamma, \delta\}$, $\{\alpha^{(1-\gamma)^{-1}\gamma}k, \sigma_z, \sigma_y, \rho\}$, where $\alpha > 0$. Therefore, the probability λ_u that an unemployed worker receives a job application opportunity can be normalized to 1.

The remaining parameters are identified as follows, and the outcome is reported in Table 1, column (a). Four of the remaining nine parameters, i.e., (k, λ_e, σ_z, δ), are identified by requiring that the non-stochastic steady state of the model matches some statistical characteristics of the US labor market. In particular, the steady-state unemployment rate

$\hat{u}$ is required to match the average US unemployment rate over the period 1951-2003. The steady state hazard rate $\hat{h}_{ue}$ from unemployment to employment is required to match the average US monthly job-finding rate over the period 1951-2003. The steady state job-to-job transition rate $\hat{h}_{ee}$ is required to be equal to the steady state hazard rate from employment to unemployment as it is (approximately) the case in the US (see Nagypál 2004). Finally, notice that the empirical probability that a worker leaves his job during a given month is decreasing with tenure and converging to 1.6 percent (see Topel and Ward 1992). In the non-stochastic steady state version of the model, the probability that a worker leaves his job during a given month is decreasing with tenure and converging to the exogenous destruction rate δ. Therefore, we set δ to be equal to 1.6.

Three of the remaining nine parameters, γ, ρ and σ_y, are identified by requiring that the data generated by the stochastic version of the model matches some statistical properties of the US labor market. The elasticity γ of the matching function with respect to applicants is such that the regression coefficient of the job-finding rate $\hat{h}_{ue}$ over the market tightness v/u in the data generated by the model is the same as in the US. The parameters ρ and σ_y are such that the quarterly autocorrelation and standard deviation of labor productivity in the data generated by the model are the same as in the US.

The two remaining parameters, β and b, are identified as follows. The parameter β is such that the real interest rate in the model is equal to the average of the real monthly interest rate on AAA-bonds over the period 1951-2003. The parameter b is such that the flow value of non-market to market activity is equal to 68 percent of the average labor productivity.[2]

4.2 The Cyclical Behavior of Unemployment and Vacancies

Table 2 presents a statistical summary of the cyclical behavior of the US labor market over the period 1951–2003. The first two rows in Table 2 report the standard deviation and the quarterly autocorrelation for the log-deviation from trend of the unemployment rate u, the aggregate vacancy rate v, the tightness of the labor market v/u, the hazard rate from unemployment to employment h_{ue}, the hazard rate from employment to unemployment h_{eu} and the average labor productivity p. The last six rows in Table 1 report the correlation matrix between these labor market variables.

[2] Note that his value of b is higher than the one used by Shimer (2005) but significantly lower than the one used by Hagedorn and Manovskii (2006).

Four facts about the behavior of the US labor market clearly emerge from Table 2. First, the unemployment rate displays large fluctuations over the business cycle. Specifically, the standard deviation of the cyclical component of the unemployment rate is almost ten times as large as the standard deviation of the average labor productivity. Second, over the business cycle, the fluctuations in the vacancy rate have the opposite sign and the same magnitude as the fluctuations in the unemployment rate. Specifically, the standard deviation of the cyclical component of the aggregate vacancy rate is 1.1 times the standard deviation of unemployment, and the contemporaneous correlation between the two variables is -.894. Consequently, the cyclical fluctuations in labor market tightness are twice as large as the fluctuations in unemployment and vacancies. Third, when the cyclical component of the unemployment rate is positive, the hazard rate from employment to unemployment is typically above its long-run trend while the hazard rate from unemployment to employment is typically below trend. Specifically, the contemporaneous correlation of the unemployment rate with the probability of transition from employment to unemployment is .709, and the contemporaneous correlation of the unemployment rate with the probability of moving from unemployment to employment is -.949. Finally, the quarterly autocorrelation of all six variables is at least as large as 0.7.

Table 3 presents the statistical summary of the cyclical behavior of the labor market in our model. Specifically, using the parameter values from Table 1 and an arbitrary initial condition for the distribution of workers across employment states, we have solved for the equilibrium of the model and simulated 9,000 months of economic activity. Then, we have thrown away the first 3,000 months of data to make our findings independent from the initial conditions. Finally, we have used the last 6,000 months of data to construct Table 3.

Four facts about the effect of aggregate productivity shocks on the labor market emerge from Table 3. First, productivity shocks generate large cyclical fluctuations in the unemployment rate. Specifically, in the data generated by the model, the standard deviation of the unemployment rate is more than seven times greater than the standard deviation of the average labor productivity. Secondly, vacancies and unemployment fluctuate in opposite directions in response to productivity shocks. Specifically, in the data generated by the model, the contemporaneous correlation between the aggregate vacancy rate and the unemployment rate is -.819, which indicates a negatively sloped Beveridge curve over business cycles. Third, in response to a positive productivity shock, the hazard rate from employment to unemployment falls and the hazard from unemployment to employment increases. Specifically, in the simulation, the contemporaneous correlation of the unemployment rate with the probability of moving from employment to unemployment is .949 and its correlation with the probability of moving from unemployment to employment is -.990. Finally, productivity shocks generate high autocorrelation in all six variables. The reader can find an illustration of these findings in Figures 1–3, where we have plotted the path of labor productivity, unemployment, vacancies and hazard rates (expressed as percent deviations from trend) over one hundred quarters of model-generated data.

By comparing Tables 2 and 3, we conclude that aggregate productivity shocks are the fundamental cause of cyclical fluctuations in the US unemployment rate during the postwar years. First, productivity shocks alone account for 64% of the historical volatility in the unemployment rate. Second, productivity shocks generate a matrix of contemporaneous correlations between unemployment, vacancies, market tightness and hazard rates that have exactly the same signs and approximately the same magnitudes as in the historical data. The only discrepancies between the US and the model-generated data regard the magnitude of the cyclical fluctuations of hazard rates and vacancies. In particular, while the cyclical component of the hazard rate from unemployment to employment is thirty percent more volatile than the hazard rate from employment to unemployment in the US data, it is only half as volatile in the model-generated data. Moreover, while the cyclical component of the vacancy rate has approximately the same volatility as the cyclical component of the unemployment rate in the US data, it is only 20 percent as volatile in the model-generated data.

Figure 1: Unemployment Fluctuations

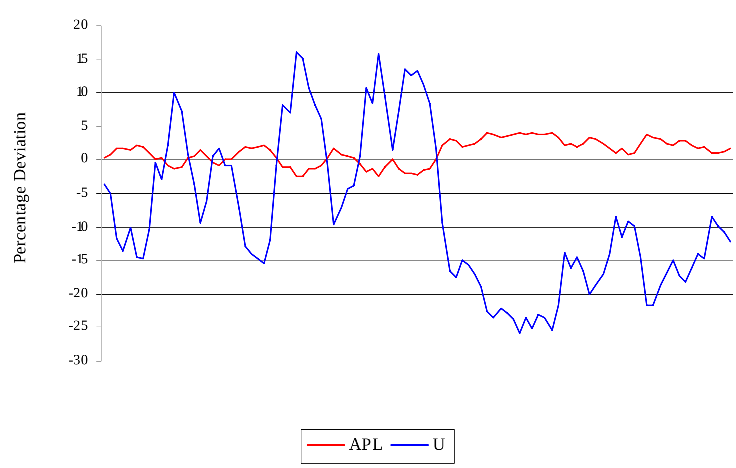

Figure 2: Transition Rates over the Cycle

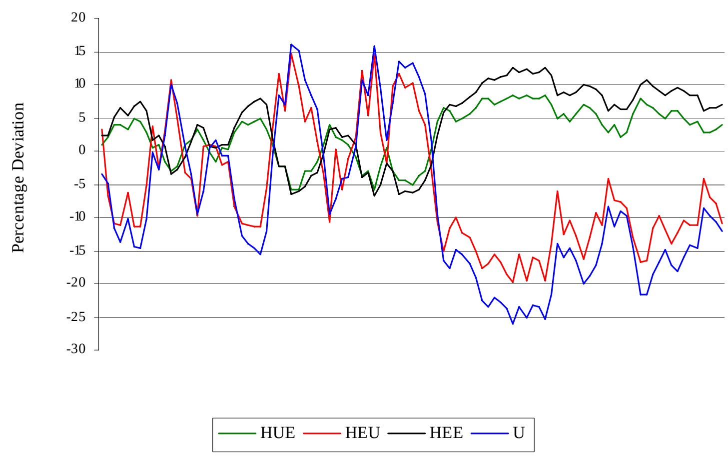

4.3 The Role of On-the-Job Search and Match Heterogeneity

How different is the effect of aggregate productivity shocks on labor market fluctuations in our model from the effect in the canonical search model by Pissarides (1985)? In order to answer this question precisely, we first reduce our model to Pissarides' by eliminating on-the-job search and match heterogeneity, i.e. we impose the constraint $\lambda_e = \sigma_z = 0$.[3] Then, we calibrate the constrained model using all the targets described in Section 4.1, with the exception of $h_{ee} = h_{eu}$ and $\delta = 0.16$. Finally, we use the calibrated model to generate a long time-series for unemployment, vacancies, hazard rates and labor productivity. Table 4 presents a statistical summary of the data generated by Pissarides' model (labeled P85).

By comparing Tables 3 and 4, we find that the main difference in the results on unemployment between our model and Pissarides' is the magnitude of the unemployment fluctuations caused by aggregate shocks to labor productivity. Specifically, according to our model, productivity shocks alone account for more than *64%* of the volatility in the US unemployment rate over the postwar years. According to Pissarides' model, they account for approximately *13%* of the historical unemployment volatility.

This difference between the two models is easy to explain. In our model, a negative productivity shock increases the unemployment rate through two channels. First, it lowers the hazard rate from unemployment to employment—because it reduces the value of filling a vacancy and the tightness of the labor market—and, secondly, it increases the hazard rate from employment to unemployment—because it raises the minimum level of match-specific productivity that the worker is willing to accept. In Pissarides' model, a negative productivity shock increases the unemployment rate only through the first channel. In fact, because all matches are assumed to be equally productive, the hazard rate from employment to unemployment is given by the exogenous separation rate σ and hence constant over the business cycle. The reader can find an illustration of this explanation in Figures 5 and 6, where we have plotted the impulse response functions of various labor market outcomes in our model and in Pissarides'.

So, if we were to introduce match heterogeneity in the Pissarides' model (as in Mortensen and Pissarides 1994), would aggregate productivity shocks have the same effect on labor market outcomes as in our model? In order to answer the question, we reduce our model to Mortensen and Pissarides' by eliminating on-the-job search. Then, we calibrate the constrained model using all the targets described in Section 4.1 with the exception of $h_{ee} = h_{eu}$. Finally, we use the calibrated model to generate a long time-series for unemployment, vacancies, hazard rates and labor productivity. Table 5 presents a statistical summary of the data generated by the model of Mortensen and Pissarides (labeled MP94).

[3] This amounts to the same calibration strategy adopted by Shimer (2005).

By comparing Tables 3 and 5, we find that the main difference in the behavior of unemployment between our model and Mortensen-Pissarides' is the sign of the contemporaneous correlation between the fluctuations in unemployment and vacancy generated by productivity shocks. According to our model, productivity shocks generate the same strong *negative* correlation between vacancies and unemployment as that observed in the US data. In contrast, in Mortensen and Pissarides' model, productivity shocks generate a strong *positive* correlation between the two variables. This positively sloped Beveridge curve is "an event that has essentially never been observed in the United States at business cycle frequencies" (Shimer 2005, p40).

This difference between the two models is easy to explain. In Mortensen and Pissarides' model, when a negative productivity shock hits the economy, firms create fewer vacancies per each worker searching the market—because the value of filling a vacancy is lower—but more workers choose to search—because the reservation quality of a match is higher. For reasonable calibrations, the second effect dominates the first one and more vacancies are created in response to a negative productivity shock. In our model, when a negative productivity shock hits the economy, firms create fewer vacancies per each worker searching the market, more workers search the market while unemployed but fewer of them search the market while employed. For reasonable calibrations, the third effect offsets the second and fewer vacancies are created. The reader can find an illustration of this explanation in Figures 5 and 7, where we have plotted the impulse response functions of various labor market outcomes in our model and in Mortensen-Pissarides'.

4.4 Are Negative Productivity Shocks Cleansing or Sullying?

In this subsection, we use our model to find out whether a negative shock to the aggregate productivity y has a cleansing effect on the allocation of workers across jobs, in the sense that it increases the average quality of existing matches, or a sullying effect in the sense that it lowers the average match quality.

Theoretically, a negative shock to aggregate productivity affects the average quality of existing matches through two different channels. On the one hand, the shock lowers the relative value of market to non-market activity and, consequently, increases the reservation quality $z_*(y)$. All else equal, this effect increases the average quality of existing matches. On the other hand, when the economy is hit by a negative shock to aggregate productivity, the tightness of all submarkets falls and, consequently, the probability that a worker employed on a job with quality $z \in [z_*(y), z^*(y)]$ finds a new employment opportunity decreases. All else equal, this second effect decreases the average quality of existing matches.

Which of these two effects dominates is a quantitative matter. The blue line in Figure 4 represents the path of aggregate productivity y_t (expressed as percent deviations from trend) over one hundred quarters of data generated from the calibrated version of our model. The green line represents the path of average labor productivity $y_t + \int zdG_t(z)$ over the same interval of time. Finally, the red line represents the path of the average match-specific component of productivity $\int zdG_t(z)$. From this picture, it is evident that negative productivity shocks have a cleansing effect on the economy. In fact, whenever the aggregate productivity lies below its trend, the average of the match-specific component of productivity is above trend (and vice versa). Consequently, average and aggregate productivity fluctuate in the same direction but the fluctuations in average productivity are not as large.

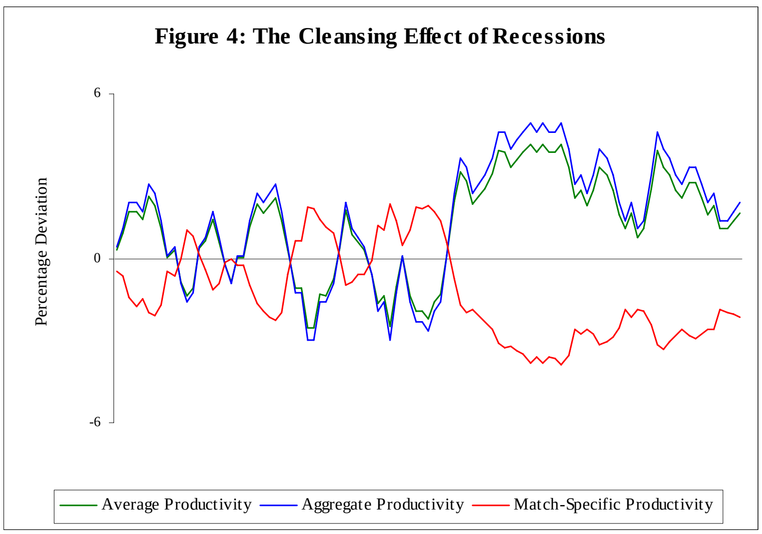

5 Conclusion

In this paper, we develop a tractable model of the labor market where workers search for jobs both while unemployed and while on the job. Search is directed in the sense that each worker chooses to search for the offer that provides the optimal tradeoff between the probability of obtaining the offer and the increase in the value relative to the worker's current employment. There are both aggregate and match-specific shocks, on which the wage path in an offer can be contingent. We characterize the equilibrium analytically and show that the equilibrium is unique and socially efficient. On the quantitative side, we calibrate the model to the US data to measure the effect of aggregate productivity fluctuations on the labor market. We find that productivity fluctuations account for approximately 64% of the cyclical volatility in US unemployment. Moreover, productivity fluctuations generate the same matrix of correlations between unemployment and other labor market variables as in the US. In particular, the Beveridge curve is negatively sloped over business cycles, and the magnitude of the slope is the same as in the data. In light of these findings, we conclude that productivity shocks are one of the main forces driving labor market fluctuations over business cycles. Furthermore, we find that recessions have a cleansing effect on the economy.

Table 1: Calibration Outcomes

	Description	(a) Baseline	(b) P85	(c) MP94
β	discount rate	.996	.996	.996
b	non-market activity value	.750	.680	.703
λ_e	on-the-job application rate	.433	—	—
γ	match elasticity wrt applicants	.500	.715	.715
δ	exogenous destruction rate	.016	.023	.016
k	vacancy cost	$2.1\ 10^{-3}$	$3.7\ 10^{-7}$	$3.5\ 10^{-7}$
σ_z	std match productivity	.117	—	.034
σ_y	std aggregate productivity	.007	.006	.007
ρ	autocor aggregate productivity	.954	.932	.939

Table 2: U.S. Quarterly Data, 1951–2003

	u	v	v/u	h_{ue}	h_{eu}	h_{ee}	p
Standard deviation	19.0	20.2	38.2	11.8	7.50	5.54	2.00
Autocorrelation	.936	.940	.941	.908	.733	.599	.878
u	1	-.894	-.971	-.949	.709	-.755	-.408
v	—	1	.975	.897	-.684	.785	.364
v/u	—	—	1	.948	-.715	.792	.396
h_{ue}	—	—	—	1	-.574	.735	.396
h_{eu}	—	—	—	—	1	-.595	-.524
h_{ee}	—	—	—	—	—	1	-.009
p	—	—	—	—	—	—	1

Source: Shimer (2005)

TABLE 3: PRODUCTIVITY SHOCKS

		u	v	v/u	h_{ue}	h_{eu}	p
Standard deviation		12.2	2.29	14.1	4.11	9.70	2.00
Autocorrelation		.907	.601	.897	.893	.779	.870
	u	1	-.819	-.995	-.990	.949	-.981
	v	—	1	.867	.873	-.922	.876
	v/u	—	—	1	.996	-.967	.992
	h_{ue}	—	—	—	1	-.959	.996
	h_{eu}	—	—	—	—	1	-.957
	p	—	—	—	—	—	1

Table 4: Productivity Shocks in P85

		u	v	v/u	h_{ue}	h_{eu}	p
Standard deviation		2.55	7.09	9.53	2.82	—	2.00
Autocorrelation		.923	.819	.859	.873	—	.870
	u	1	-.939	-.967	-.975	—	-.973
	v	—	1	.995	.991	—	.988
	v/u	—	—	1	.999	—	.996
	h_{ue}	—	—	—	1	—	.997
	h_{eu}	—	—	—	—	—	—
	p	—	—	—	—	—	1

Table 5: Productivity Shocks in MP94

	u	v	v/u	h_{ue}	h_{eu}	p
Standard deviation	10.3	5.66	7.35	2.20	10.6	2.00
Autocorrelation	.788	.475	.877	.877	.511	.870
u	1	.633	-.870	-.870	.896	-.868
v	—	1	-.254	-.254	.437	-.249
v/u	—	—	1	.999	-.719	.997
h_{ue}	—	—	—	1	-.720	.997
h_{eu}	—	—	—	—	1	-.729
p	—	—	—	—	—	1

References

[1] Barlevy, G. 2002. The Sullying Effect of Recessions. *Review of Economic Studies* 69: 65–96.

[2] Burdett, K. 1978. A Theory of Employee Job Search and Quits. *American Economic Review* 68: 212–220.

[3] Burdett, K., and M. Coles. 2003. Equilibrium Wage-Tenure Contracts. *Econometrica* 71: 1377–404.

[4] Burdett, K., and D. Mortensen. 1998. Wage Differentials, Employer Size, and Unemployment. *International Economic Review* 39: 257–73.

[5] Burdett, K., S. Shi, and R. Wright. 2001. Pricing and Matching with Frictions. *Journal of Political Economy* 109: 1060–85.

[6] Delacroix, A. and S. Shi. 2006. Directed Search on the Job and the Wage Ladder. *International Economic Review* 47: 651–99.

[7] Fallick, B. and C. Fleischman. 2004. Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows. Mimeo. Federal Reserve Board.

[8] Gomes, J., J. Greenwood and S. Rebelo. 2001. Equilibrium Unemployment. *Journal of Monetary Economics* 48: 109–52.

[9] Hagedorn, M. and I. Manovskii. 2006. The Cyclical Behavior of Equilibrium Unemployment and Vacancies Revisited. Mimeo. University of Pennsylvania.

[10] Hall, R. 2005. Macroeconomic Fluctuations and the Allocation of Time. *Journal of Labor Economics* 15: S223–50.

[11] Moen, E. 1997. Competitive Search Equilibrium. *Journal of Political Economy* 105: 694–723.

[12] Mortensen, D. 1994. The Cyclical Behavior of Job and Worker Flows. *Journal of Economic Dynamics and Control* 18: 1121–42.

[13] Mortensen, D. 2003. *Why Are Similar Workers Paid Differently?* MIT University Press, Cambridge, MA.

[14] Mortensen, D. and E. Nagypál. 2007. More on Unemployment and Vacancy Fluctuations. Review of Economic Dynamics 10: 327–47.

[15] Mortensen, D. and C. Pissarides. 1994. Job Creation and Job Destruction in the Theory of Unemployment. *Review of Economic Studies* 61: 397–415.

[16] Nagypál, E. 2005. Worker Reallocation Over the Business Cycle: The Importance of Job-to-Job Transitions. Mimeo, Northwestern University.

[17] Pissarides, C. 1985. Short-Run Equilibrium Dynamics of Unemployment, Vacancies and Real Wages. *American Economic Review* 75: 676–90.

[18] Pissarides, C. 1994. Search Unemployment with On-the-job Search. *Review of Economic Studies* 61: 457–75.

[19] Shi, S. 2007. Directed Search for Equilibrium Wage-Tenure Contracts. Mimeo, University of Toronto.

[20] Shimer, R. 2005. The Cyclical Behavior of Unemployment and Vacancies. *American Economic Review* 95: 25–49.

[21] Stevens, M. 2004. Wage-Tenure Contracts in a Frictional Labour Market: Firms' Strategies for Recruitment and Retention. *Review of Economic Studies* 71: 535–51.

[22] Topel, R. and M. Ward. 1992. Job Mobility and the Careers of Young Men. *The Quarterly Journal of Economics* 107: 439–79.

A Appendix

Proof of Proposition 1: The proof is divided in three steps.
Claim 1: The Bellman equation (19) has a unique solution $S(h, y)$.
Proof: Let T denote the mapping, i.e.

$$T(\tilde{S})(h, y) = \max_{\{\hat{\theta}_u, \hat{\theta}, \hat{z}_*\}} C + \beta \int \tilde{S}(h_+, y_+) d\Psi(y_+|y), \text{ s.t. } (15)\text{–}(18).$$

The mapping T satisfies the monotonicity condition—namely, if $\tilde{S}_2(h, y) \geq \tilde{S}_1(h, y)$, then $T(\tilde{S}_2)(h, y)$ is greater or equal than $T(\tilde{S}_1)(h, y)$—and the discounting condition—namely, if $a > 0$, then $T(\tilde{S} + a)(h, y)$ is smaller or equal than $T(\tilde{S})(h, y) + \beta\alpha$. From Blackwell's sufficient conditions, it follows that the mapping T is a contraction and it has a unique fixed point. Since a solution to (19) is a fixed point of T, the Bellman equation (19) has a unique solution $S(h; y)$.

Claim 2: Given that $\hat{s}(z; y)$ is a weakly increasing function of z, the equations (20)-(24) uniquely determine $\hat{\mu}_u(y)$, $\hat{\mu}(z; y)$, $\hat{\theta}_u(y)$, $\hat{z}^*(y)$, $\hat{\theta}(z; y)$ and $\hat{z}_*(y)$.
Proof: Given $\hat{s}(z; y)$, equation (21) uniquely determines $\hat{\mu}_u(y)$ and $\hat{\mu}(z; y)$. Moreover, $\hat{\mu}(z; y)$ is a strictly increasing function of z. Equation (20) uniquely determines $\hat{\theta}_u(y)$ because $p'(\theta)$ is a strictly decreasing function of θ. Equation (24) uniquely determines $\hat{z}^*(y)$ because $\hat{\mu}(z; y)$ is strictly increasing in z. Equation (22) uniquely determines $\hat{\theta}(z; y)$ because $p'(\theta)$ is a strictly decreasing function of θ. Moreover, $\hat{\theta}(z; y)$ is a strictly decreasing function of z. Finally, equation (23) uniquely determines $z_*(y)$ because $\hat{\mu}(z; y)$ and $p'(\hat{\theta}(z; y))^{-1} \cdot p(\hat{\theta}(z; y)) \cdot \hat{\theta}(z; y)$ are both strictly increasing functions of z.

In light of claim 2, we can think of (25) as one equation in the only unknown $\hat{s}(z; y)$. In the next claim, we prove that this equation has a unique solution.

Claim 3: Equations (25) has a unique solution with the property that $\hat{s}(z; y)$ is weakly increasing in z.
Proof: On the way to a contradiction suppose that $\hat{s}^1(z; y)$ and $\hat{s}^2(z; y)$ are two distinct solutions to (25). Denote with $\hat{\mu}_u^i(y)$, $\hat{\mu}^i(z; y)$, $\hat{\theta}_u^i(y)$, $\hat{z}^{*i}(y)$, $\hat{\theta}^i(z; y)$ and $\hat{z}_*^i(y)$, the solution to the equations (20)-(24) given $\hat{s}^i(z; y)$. Denote with $\kappa^i(y)$ the unique solution to the functional equation

$$\kappa^i(y) = -k\lambda_u \hat{\theta}_u^i(y) + (1 - \lambda_u p(\hat{\theta}_u^i(y)))b + \lambda_u p(\hat{\theta}_u^i(y)) \int (z + y) dF(z) +$$
$$\beta \int \left[\kappa^i(y_+) + \lambda_u p(\hat{\theta}_u^i(y)) \int \hat{s}^i(z; y_+) dF(z) \right] d\Psi(y_+|y).$$

Denote with $S^i(h;y)$ the function

$$S^i(h;y) = \kappa^i(y_+) + \int \hat{s}^i(z;y)h(z)dz.$$

Consider the maximization problem

$$T(S^i)(h,y) = \max_{\{\hat{\theta}_u, \hat{\theta}, \hat{z}_*\}} C + \beta \int S^i(h_+, y_+)d\Psi(y_+|y), \text{ s.t. (15)–(18).}$$

It is immediate to verify that $\hat{\theta}_u^i(y)$, $\hat{\theta}^i(z;y)$ and $\hat{z}_*^i(y)$ satisfy the necessary conditions for the optimal choice of $\hat{\theta}_u$, $\hat{\theta}$, $\hat{z}_*$. Also, it is immediate to verify that $\hat{s}^i(z;y)$ is the Volterra derivative of the value function $T(S^i)$ with respect to $h(z)$. Finally, it is immediate to verify that $\kappa^i(y)$ is equal to $T(S^i)(h_0,y)$—where $h_0(z) = 0$ for all $z \in [z, z]$. Overall, the value function $T(S^i)$ is given by

$$T(S^i)(h,y) = \kappa^i(y) + \int \hat{s}^i(z;y)h(z)dz = S^i(h,y).$$

Therefore, the functions $S^1(h_0,y)$ and $S^2(h_0,y)$ are two distinct fixed points of the contraction mapping T. A contradiction.

The market equilibrium is a solution to the system of equations (20)-(25) with the property that $s^*(z;y)$ is weakly increasing in z. The solution to the social planner's problem is a solution to (20)-(25) such that $\hat{s}(z;y)$ is weakly increasing in z. Therefore, claim 3 implies that the labor allocation $\theta_u^*(y)$, $\theta^*(z;y)$, $z_*(y)$ generated by the market equilibrium is unique and efficient. ||