

Efficient Search on the Job and the Business Cycle*

Guido Menzio

Department of Economics

University of Pennsylvania

and NBER

(email: gmenzio@sas.upenn.edu)

Shouyong Shi

Department of Economics

University of Toronto

(email: shouyong@chass.utoronto.ca)

March 2010

Abstract

We develop a model of search on the job where match-specific productivity induces workers to move between jobs and where workers can direct search toward particular jobs. After characterizing the social planner's allocation, we prove that there exists a unique equilibrium in the market economy and it decentralizes the social planner's allocation. We calibrate the model to the US economy to measure the contribution of aggregate productivity shocks to the cyclical volatility of labor market variables. Aggregate productivity shocks account for a sizable fraction of the cyclical volatility of the labor market if match quality is observed after the match is created, but only for a small fraction of the cyclical volatility if match quality is observed before the match is created. However, we argue that the version of the model in which matches are experience goods provides an unambiguously better description of the US labor market and, hence, deserves more confidence in its predictions about the effect of aggregate productivity shocks.

*We are grateful to an editor and two referees for comments. We have also received comments from participants at the Search Theory conference at the University of Quebec in Montreal (April 2007), the Society for Economic Dynamics Meetings (New Orleans, January 2008, and Cambridge, July 2008), the Econometric Society Meeting in Pittsburgh (June 2008), the Vienna Macroeconomic Workshop (October, 2008), the NBER Summer Institute in Cambridge (July, 2008), the Canadian Macro Study Group Meeting (November, 2009) and at the seminars at Michigan, Stanford, Princeton, Columbia, Wisconsin, Boston University, Bank of Canada and Hoover Institution. Discussions with Rudi Bachmann, Gadi Barlevy, Mike Elsby, Martin Gervais, Marcus Hagedorn, Bob Hall, Nir Jaimovich, Nobuhiro Kiyotaki, Dale Mortensen, Giuseppe Moscarini, Thijs van Rens and Ludo Visschers led to significant improvements in the paper. We thank Frank Diebold, Jason Faberman, Giuseppe Moscarini, Eva Nagypál, David Neumark and Daniel Polsky for generously sharing their data with us. Menzio gratefully acknowledges the financial support and the hospitality of the Hoover Institution. Shi gratefully acknowledges the financial support from the Social Sciences and Humanities Research Council of Canada and from the Bank of Canada Fellowship. The usual disclaimer applies.

1 Introduction

The US labor market exhibits a number of stylized patterns at the business cycle frequency. First, there are large movements of workers between the states of employment, unemployment and across different employers. On average, the rate at which unemployed workers move into employment (henceforth, the UE rate) is 42 percent a month, the rate at which employed workers move into unemployment (the EU rate) is 2.6 percent a month, and the rate at which workers move from one employer to the other (the EE rate) is 2.9 percent a month. Second, these transition rates are very volatile, contributing to the large volatility of the unemployment rate and the vacancy rate. As documented in Table 1, monthly UE, EU and EE rates are five to six times as volatile as labor productivity, while the unemployment rate and the vacancy rate are ten times as volatile as labor productivity. Third, there is a clear pattern of comovement between these labor market variables. The unemployment rate is mildly negatively correlated with labor productivity, and strongly negatively correlated with the vacancy rate. Moreover, the unemployment rate is negatively correlated with the UE and EE rates, and positively correlated with the EU rate. The question that we want to address in this paper is: To what extent can cyclical movements in aggregate labor productivity account for these patterns of comovement and volatility in unemployment, vacancies and workers' transition rates?

To answer this question, we develop a model of search on the job in which workers move between employment, unemployment and across firms because of differences in the quality of different firm-worker matches. We calibrate the model to match the average of workers' transition rates in the data, as well as other features of the US labor market. Then we simulate the model to measure the contribution of aggregate productivity shocks to the cyclical volatility of unemployment, vacancies and workers' transition rates. Our main finding is that the contribution of aggregate productivity shocks to the cyclical volatility of the labor market critically depends on whether the quality of a firm-worker match is observed before or after the match is created. If the quality is observed after the match is created (i.e., if matches are experience goods), then aggregate productivity shocks account for a sizeable fraction of the cyclical volatility of the labor market. If the quality is observed before the match is created (i.e., if matches are inspection goods), the contribution of aggregate productivity

shocks to the cyclical volatility of the labor market is marginal. However, we argue that the version of the model in which matches are experience goods provides an unambiguously better description of the US labor market and, hence, we should have more confidence in its predictions about the effect of aggregate productivity shocks.

In our model, the search process is directed—as in Shimer (1996), Moen (1997), and Burdett et al. (2001)—rather than random—as in Mortensen (1982) and Pissarides (1985). On one side of the market, firms choose how many and what type of vacancies to create. On the other side of the market, workers choose what type of vacancies to search. The type of a vacancy is described by the conditions under which it hires a worker and by the employment contract that it offers to a new hire. Employment contracts are bilaterally efficient, in the sense that they maximize the surplus of a match. Workers and vacancies are brought into contact by a meeting function that has constant returns to scale. Upon meeting, a worker and a firm observe a signal about the idiosyncratic productivity of their match (i.e., the quality). If the signal meets the conditions specified by the vacancy’s type, the worker and the firm begin to produce, and eventually observe the actual quality of their match. If the signal does not meet those conditions, the worker returns to his previous employment position. Depending on the informativeness of the signal, the model captures different views of the matching process. If the signal is uninformative, the match is an experience good. If the signal is fully informative, the match is an inspection good. If the signal contains some but not all information, a match is partly an inspection and partly an experience good.

In the theoretical part of the paper, we first characterize the social planner’s allocation. We then prove that there exists a unique equilibrium in the market economy and that this equilibrium decentralizes the social planner’s allocation. Moreover, we prove that the equilibrium is *block recursive*, in the sense that the agents’ value and policy functions depend on the aggregate state of the economy only through the realization of the exogenous aggregate productivity shocks, and not through the endogenous distribution of workers across employment states (unemployment and employment in different matches). This property of the equilibrium implies that the block of equations that determine the agents’ value and policy functions can be solved independently from the block of equations that determine the evolution of the distribution of workers across employment states. Therefore, even though

in our model matches are ex-post heterogeneous, we can solve the equilibrium as easily as one would solve the equilibrium of a representative agent model. Moreover, because the equilibrium is efficient, we can characterize the equilibrium policy functions by examining the first order conditions of the planner's problem.

In the quantitative part of the paper, we consider two versions of the model that provide, a priori, an equally plausible description of the US labor market. Specifically, we calibrate the version of the model in which matches are experience goods and the version in which matches are inspection goods. Following Shimer (2005), we ask the model to match the empirical average of the workers' transition rates between employment states and the empirical elasticity of the UE rate with respect to the vacancy-to-unemployment ratio. In addition, we use the empirical tenure distribution to calibrate the shape and the variance of the distribution of match-specific productivities.

Given the calibrated parameter values, we simulate the two versions of the model to measure the effect of aggregate productivity shocks on the labor market. When we use the version of the model in which matches are experience goods, we find that aggregate productivity shocks generate the same pattern of comovement between unemployment, vacancies, and workers' transition rates as in the US labor market (see the first paragraph of this Introduction). Moreover, we find that aggregate productivity shocks account for a large fraction of the cyclical volatility of the US labor market. Specifically, these shocks account for 85 percent of the cyclical volatility of unemployment, for 25 percent of the cyclical volatility of vacancies, and for approximately 40, 100 and 80 percent of the cyclical volatility of the UE, EU and EE rates.

When we use the version of the model in which matches are inspection goods, we also find that aggregate productivity shocks generate the same pattern of comovement between unemployment, vacancies and workers' transition rates as in the data. However, these shocks account for very little of the cyclical volatility of the labor market. Specifically, aggregate productivity shocks account for 8 percent of the cyclical volatility of unemployment, for 20 percent of the cyclical volatility of vacancies, and for 15, 0 and 2 percent of the cyclical volatility of the UE, EU and EE rates. As we will elaborate in section 6, some of the differences between the predictions of the inspection and experience versions of the model

are due to the fact that the informativeness of signals affects the way in which the economy responds to aggregate productivity shocks, while other differences are due to the fact that the informativeness of the signals affects the calibrated values of the parameters.

In order to choose between the experience-good and the inspection-good versions of the model, we examine their fit of the calibration targets. We find that both versions of the model are able to match exactly the standard calibration targets (i.e., the average workers' transition rates and the elasticity of the UE rate with respect to the vacancy-to-unemployment ratio). However, the experience-good version of the model fits well the empirical tenure distribution, while the inspection-good version does not. We conclude that, at least within the confines of our dataset, the experience-good version of the model provides an unambiguously better description of the US labor market and, hence, we should have more confidence in its predictions about the cyclical effect of aggregate productivity shocks.

The paper contains two contributions. On the theoretical side, the paper develops a model of directed search on the job in which the transitions of workers between unemployment, employment and across firms are driven by differences in the quality of different firm-worker matches. For this model, the paper establishes that the equilibrium is unique, efficient and block recursive. Block recursivity is the most important property of equilibrium because it allows us to solve the aggregate dynamics of our model with ex-post heterogeneous agents as easily as one would solve the aggregate dynamics of a representative agent model. In earlier work (Shi 2009, Menzio and Shi 2009, 2010), we already established the existence of block recursive equilibria in models of directed search on the job. In this paper, we sharpen those results by showing that, in fact, all equilibria are block recursive.¹

The equilibrium of our model is block recursive because the search process is directed. When the search process is random, models of search on the job are not block recursive, in the sense that the agents' value and policy functions depend on the entire distribution of workers across employment states (unemployment and employment in different jobs). For this reason, models of random search on the job are difficult to solve outside of the steady state. To circumvent this difficulty, the existing literature has had to impose some strong

¹An earlier model of directed search on the job is Delacroix and Shi (2006), where the equilibrium is block recursive but all matches have the same productivity. In the literature of directed search, most models abstract from on-the-job search, e.g. Acemoglu and Shimer (1991), Moen (1997), Burdett et al. (2001), Shi (2001), Menzio (2007) and Gonzalez and Shi (2010).

restrictions on the environment. For example, in order to solve their models outside of the steady state, Moscarini and Postel-Vinay (2009) and Robin (2009) need to assume that the rate at which firms and workers come into contact is exogenous. Mortensen (1994) and Pissarides (1994, 2000) need to assume that an employed worker moves into unemployment before bargaining the wage with his new employer. Moreover, these papers need to assume that all new matches have exactly the same productivity. Directed search allows us to study the aggregate dynamics of the model without such restrictive assumptions.

On the empirical side, the contribution of the paper is to measure the effect of aggregate productivity shocks on unemployment, vacancies and workers' transition rates using a calibrated model of search on the job with heterogeneous matches. When matches are experience goods, we show that aggregate productivity shocks can account for the pattern of comovement between unemployment, vacancies and workers' transition rates and for a large fraction of their volatility. These findings are novel and are quite different from those obtained using models that abstract from search on the job and match heterogeneity (e.g. Shimer 2005), models that abstract from search on the job (e.g., Mortensen and Pissarides 1994 and Merz 1995), and models of random search on the job (e.g. Mortensen 1994, Nagypal 2008, Moscarini and Postel-Vinay 2009 and Robin 2009). We will detail these differences at the end of section 5. When matches are inspection goods, we show that aggregate productivity shocks account for a very small fraction of the cyclical volatility of the labor market. This finding is novel since, as far as we know, a search model in which matches are inspection goods has never been used for studying the cyclical dynamics of the labor market.

The remainder of the paper is organized as follows. In section 2, we present the model and characterize the social planner's problem. In section 3, we describe a market economy and prove that the equilibrium is unique, block recursive and efficient. In section 4, we describe the data and the calibration. In section 5, we measure the effect of aggregate productivity shocks on the labor market using the version of the model in which matches are experience goods. In section 6, we measure the effect of aggregate productivity shocks using the version of the model in which matches are inspection goods. Moreover, we argue that the experience model provides a better fit of the acyclical properties of the labor market. Section 7 concludes. The proofs of all propositions and theorems are in the appendix.

2 Planner's Problem

2.1 Preferences and technologies

The economy is populated by a continuum of workers with measure 1 and a continuum of firms with positive measure. Each worker is endowed with an indivisible unit of labor and maximizes the expected sum of periodical consumption discounted at the factor $\beta \in (0, 1)$. Each firm operates a constant return to scale technology that turns one unit of labor into $y + z$ units of output. The first component of productivity, y , is common to all firms and its value lies in the set $Y = \{y_1, y_2, \dots, y_{N(y)}\}$, where $y_1 < y_2 < \dots < y_{N(y)}$ and $N(y) \geq 2$ is an integer. The second component of productivity, z , is specific to a firm-worker pair, and its value lies in the set $Z = \{z_1, z_2, \dots, z_{N(z)}\}$, where $z_1 < z_2 < \dots < z_{N(z)}$ and $N(z) \geq 2$ is an integer.² Each firm maximizes the expected sum of profits discounted at the factor β .

Time is discrete and continues forever. At the beginning of each period, the state of the economy can be summarized by the triple $\psi = (y, u, g)$. The first element of ψ denotes aggregate productivity, $y \in Y$. The second element denotes the measure of workers who are unemployed, $u \in [0, 1]$. The third element is a function $g : Z \rightarrow [0, 1]$, with $g(z)$ denoting the measure of workers who are employed in matches with the idiosyncratic productivity z . Let Ψ denote the set in which ψ belongs.

Each period is divided into four stages: separation, search, matching and production. At the separation stage, the planner chooses the probability $d \in [\delta, 1]$ with which a match between a firm and a worker is destroyed. The lower bound on d denotes the probability that a match is destroyed for exogenous reasons, $\delta \in (0, 1)$.

At the search stage, the planner sends workers and firms searching for new matches across different locations. In particular, the planner chooses how many vacancies each firm should open at different locations. The cost of maintaining a vacancy for one period is $k > 0$. Moreover, the planner chooses which location each worker should visit if he has the opportunity of searching for a new match. The probability that a worker has the opportunity of

²The assumption that y and z are discrete random variables simplifies the notation but plays no role in the derivation of our theoretical results. In fact, it is straightforward to generalize the proof of uniqueness, efficiency and block recursivity of equilibrium to the case in which y and z are continuous random variables. Moreover, the assumption plays no role in the derivation of our quantitative results, because continuous random variables would eventually have to be discretized in order to simulate the model.

searching depends on his employment status. If the worker was unemployed at the beginning of the period, he can search with probability $\lambda_u \in [0, 1]$. If the worker was employed at the beginning of the period and did not lose his job during the separation stage, he can search with probability $\lambda_e \in [0, 1]$. Finally, if the worker lost his job during the separation stage, he cannot search. As it is standard in models of directed search (e.g. Acemoglu and Shimer 1999, Burdett et al. 2001, and Shi 2001), we assume that the planner must send workers in the same employment state to search at the same location.

At the matching stage, the workers and the vacancies who are searching at the same location are brought into contact by a meeting technology with constant returns to scale that can be described in terms of the vacancy-to-worker ratio θ (i.e., the tightness). Specifically, the probability that a worker meets a vacancy is $p(\theta)$, where $p : \mathbb{R}_+ \rightarrow [0, 1]$ is a twice continuously differentiable, strictly increasing, and strictly concave function which satisfies the boundary conditions $p(0) = 0$ and $p(\bar{\theta}) = 1$, $0 < \bar{\theta} < \infty$. Similarly, the probability that a vacancy meets a worker is $q(\theta)$, where $q : \mathbb{R}_+ \rightarrow [0, 1]$ is a twice continuously differentiable and strictly decreasing function such that $q(\theta) = p(\theta)/\theta$, $q(0) = 1$ and $q(\bar{\theta}) = 0$.

When a firm and a worker meet, Nature draws the idiosyncratic productivity of their match, z , from the probability distribution $f(z)$, $f : Z \rightarrow [0, 1]$. Nature also draws a signal about the idiosyncratic productivity of their match, s . With probability $\alpha \in [0, 1]$, the signal is equal to z ; with probability $1 - \alpha$, the signal is drawn from the distribution f independently of z . After observing s but not z , the planner chooses whether to create the match or not. If the planner chooses to create the match, the worker's previous match is destroyed (if the worker was employed). If the planner chooses not to create the match, the worker returns to his previous status (unemployment or employment in the previous match).

Notice that the information structure above encompasses a number of interesting special cases. If $\alpha = 0$, the planner has no information about the quality of a match when choosing whether to create it or not, in which case a match is a pure *experience good*. If $\alpha = 1$, the planner has perfect information about the quality of a match before choosing whether to create it or not, in which case a match is a pure *inspection good*. If $\alpha \in (0, 1)$, a match is partly an experience good and partly an inspection good.

At the production stage, an unemployed worker produces $b > 0$ units of output. A worker

employed in a match with idiosyncratic productivity z produces $y + z$ units of output, and z is observed. At the end of this stage, Nature draws next period's aggregate component of productivity, \hat{y} , from the probability distribution $\phi(\hat{y}|y)$, $\phi : Y \times Y \rightarrow [0, 1]$. Throughout the paper, the caret indicates variables or functions in the next period.

2.2 Formulation of the planner's problem

At the beginning of a period, the social planner observes the aggregate state of the economy $\psi = (y, u, g)$. At the separation stage, the planner chooses the probability $d(z)$ of destroying a match with idiosyncratic productivity z , $d : Z \rightarrow [\delta, 1]$. At the search stage, the planner chooses $\theta_u \in [0, \bar{\theta}]$, the ratio of vacancies to workers at the location where unemployed workers search. Moreover, the planner chooses $\theta_e(z) : Z \rightarrow [0, \bar{\theta}]$, the ratio of vacancies to workers at the location where the workers employed in matches of type z look for new matches. At the matching stage, the planner chooses the probability $c_u(s)$ of creating a match between an unemployed worker and a firm given the signal s , $c_u : Z \rightarrow [0, 1]$. Moreover, the planner chooses the probability $c_e(s, z)$ of creating a match between a worker who is employed in a type- z match and another firm given that the signal of the new match is s , $c_e : Z \times Z \rightarrow [0, 1]$. Given the choices $(d, \theta_u, \theta_e, c_u, c_e)$, aggregate consumption is given by

$$F(d, \theta_u, \theta_e, c_u, c_e | \psi) = -k \{ \lambda_u \theta_u u + \sum_z [(1 - d(z)) \lambda_e \theta_e(z) g(z)] \} + b \hat{u} + \sum_z [(y + z) \hat{g}(z)], \quad (1)$$

where (\hat{u}, \hat{g}) denotes the distribution of workers across employment states at the production stage (and, hence, at the beginning of next period).

In order to derive detailed expressions for \hat{u} and \hat{g} in terms of the aggregate state of the economy and the planner's choices, it is useful to derive the transition probabilities for an individual worker. First, consider a worker who enters the period unemployed. With probability $1 - \lambda_u p(\theta_u)$, the worker does not meet any firm at the matching stage. In this case, the worker remains unemployed. With probability $\lambda_u p(\theta_u)$, the worker meets a firm during the matching stage. In this case, the worker and the firm receive a signal s about the idiosyncratic productivity of their match. With probability $1 - c_u(s)$, the match is not created and the worker remains unemployed. With probability $c_u(s) [\alpha + (1 - \alpha) f(s)]$, the match is created and its idiosyncratic productivity is $z = s$. With probability $c_u(s) (1 - \alpha) f(z)$, the

match is created and its idiosyncratic productivity is $z \neq s$. Overall, at the production stage, the worker remains unemployed with probability $1 - \lambda_u p(\theta_u) m_u$, where $m_u = \sum_s [c_u(s) f(s)]$, and he is employed in a match of type z with probability $\lambda_u p(\theta_u) [\alpha c_u(z) + (1 - \alpha) m_u] f(z)$.

Next, consider a worker who enters the period in a match of type z . With probability $d(z)$, the worker leaves his match and becomes unemployed during the separation stage. With probability $(1 - d(z)) (1 - \lambda_e p(\theta_e(z)))$, the worker does not meet any firm during the matching stage. In this case, the worker remains in the same match as at the beginning of the period. With probability $(1 - d(z)) \lambda_e p(\theta_e(z))$, the worker meets a firm at the matching stage. In this case, the worker and the firm receive a signal s about the idiosyncratic productivity of their match. With probability $1 - c_e(s, z)$, the match is not created. With probability $c_e(s, z) [\alpha + (1 - \alpha) f(s)]$, the match is created and its idiosyncratic productivity is $z' = s$. With probability $c_e(s, z) (1 - \alpha) f(z')$, the match is created and its idiosyncratic productivity is $z' \neq s$. Overall, at the production stage, the worker is unemployed with probability $d(z)$, he is employed in the same match as at the beginning of the period with probability $(1 - d(z)) (1 - \lambda_e p(\theta_e(z)) m_e(z))$, where $m_e(z) = \sum_s [c_e(s, z) f(s)]$, and he is employed in a new match of type z' with probability $(1 - d(z)) \lambda_e p(\theta_e(z)) [\alpha c_e(z', z) + (1 - \alpha) m_e(z)] f(z')$.

After aggregating the transition probabilities of individual workers, we find that the measure of workers who are unemployed at the production stage is given by

$$\hat{u} = u [1 - \lambda_u p(\theta_u) m_u] + \sum_z [d(z) g(z)]. \quad (2)$$

Similarly, the measure of workers who are employed in matches of type z' is given by

$$\begin{aligned} \hat{g}(z') = & u \lambda_u p(\theta_u) [\alpha c_u(z') + (1 - \alpha) m_u] f(z') \\ & + g(z') [1 - d(z')] [1 - \lambda_e p(\theta_e(z')) m_e(z')] \\ & + \sum_z g(z) \{ [1 - d(z)] [\lambda_e p(\theta_e(z))] [\alpha c_e(z', z) + (1 - \alpha) m_e(z)] f(z') \}. \end{aligned} \quad (3)$$

The planner maximizes the sum of present and future consumption discounted at the factor β . Hence, the planner's value function, $W(\psi)$, solves the following Bellman equation

$$\begin{aligned} W(\psi) = & \max_{(d, \theta_u, \theta_e, c_u, c_e)} F(d, \theta_u, \theta_e, c_u, c_e | \psi) + \beta \mathbb{E} W(\hat{\psi}) \\ \text{s.t.} & \quad (2) \text{ and } (3), \quad d : Z \rightarrow [\delta, 1], \quad \theta_u \in [0, \bar{\theta}], \\ & \quad \theta_e : Z \rightarrow [0, \bar{\theta}], \quad c_u : Z \rightarrow [0, 1], \quad c_e : Z \times Z \rightarrow [0, 1]. \end{aligned} \quad (4)$$

Throughout this paper, the expectation operator is taken over the future state of the aggre-

gate economy, $\hat{\psi}$, unless it is specified otherwise.

Theorem 1 (*Block recursivity of the planner's problem*): (i) The planner's value function, $W(\psi)$, is the unique solution to (4). (ii) $W(\psi)$ is linear in u and g . That is, $W(\psi) = W_u(y)u + \sum_z [W_e(z, y)g(z)]$, where $W_u(y)$ and $W_e(z, y)$ are called the component value functions. The component value function $W_u(y)$ is given by

$$\begin{aligned} W_u(y) = & \max_{(\theta_u, c_u)} \{ -k\lambda_u\theta_u + [1 - \lambda_u p(\theta_u)m_u] [b + \beta\mathbb{E}W_u(\hat{y})] \\ & + \lambda_u p(\theta_u) \sum_{z'} \{ [\alpha c_u(z') + (1 - \alpha)m_u] [y + z' + \beta\mathbb{E}W_e(z', \hat{y})] f(z') \} \} \\ \text{s.t. } & \theta_u \in [0, \bar{\theta}], \quad c_u : Z \rightarrow [0, 1]. \end{aligned} \tag{5}$$

The component value function $W_e(z, y)$ is given by

$$\begin{aligned} W_e(z, y) = & \max_{(d, \theta_e, c_e)} \{ d [b + \beta\mathbb{E}W_u(\hat{y})] - (1 - d) k\lambda_e\theta_e \\ & + (1 - d) [1 - \lambda_e p(\theta_e)m_e] [y + z + \beta\mathbb{E}W_e(z, \hat{y})] \\ & + (1 - d)\lambda_e p(\theta_e) \sum_{z'} \{ [\alpha c_e(z') + (1 - \alpha)m_e] [y + z' + \beta\mathbb{E}W_e(z', \hat{y})] f(z') \} \} \\ \text{s.t. } & d \in [\delta, 1], \quad \theta_e \in [0, \bar{\theta}], \quad c_e : Z \rightarrow [0, 1]. \end{aligned} \tag{6}$$

(iii) $W_e(z, y)$ is strictly increasing in z . (iv) The policy correspondences $(d^*, \theta_u^*, \theta_e^*, c_u^*, c_e^*)$ associated with (4) depend on ψ only through y and not through (u, g) .

The planner's problem depends on the aggregate productivity, y , the measure of workers who are unemployed, u , and the measure of workers who are employed in each of the $N(z)$ types of matches, $g(z)$. If $N(z)$ is large, as it is needed to properly calibrate and simulate the model, solving for the planner's problem is in general a difficult task both analytically and computationally. However, Theorem 1 shows that the planner's problem in our model can be decomposed into $N(z) + 1$ independent components, one for every worker in a different employment state. In the component of the problem associated with an unemployed worker, (5), the planner chooses θ_u and $c_u(s)$ to maximize the present value of the output generated by this worker, net of the cost of the vacancies that are assigned to him. Similarly, in the component of the problem associated to a worker employed in a match of type z , (6), the planner chooses $d(z)$, $\theta_e(z)$ and $c_e(s, z)$ to maximize the present value of consumption generated by this worker, net of the cost of the vacancies that are assigned to him. Since each

of these problems only depends on the aggregate state of the economy, ψ , through the one-dimensional aggregate productivity, y , and not through the multidimensional distribution of workers, (u, g) , solving the planner's problem in our economy is computationally just as easy as solving the planner's problem in a representative agent economy. We refer to the separability of the planner's problem into components as block recursivity, following Shi (2009), Menzio and Shi (2009b, 2010) and Gonzalez and Shi (2010).

It is important to notice that the planner's problem is block recursive because the search process is directed (i.e., the planner can send each particular type of workers to search at a particular location) rather than random (i.e. the planner has to send all workers to the same location). Under random search, workers in different employment states search in the same location and face the same probability of meeting a firm. Since the value of a meeting depends on the workers' employment state, the optimal choice of the tightness depends on u and g , and the planner's problem cannot be decomposed into different components. Under directed search, the planner can choose a different tightness for the location visited by workers in different employment states. This property, together with the linearity of the production function, is sufficient to guarantee that the planner's problem is block recursive.

2.3 Solution to the planner's problem

The efficient choice for the probability of creating a match between an unemployed worker and a firm is $c_u^*(s, y) = 1$ if

$$b + \beta \mathbb{E}W_u(\hat{y}) \leq \alpha[y + s + \beta \mathbb{E}W_e(s, \hat{y})] + (1 - \alpha)\mathbb{E}_z[y + z + \beta \mathbb{E}W_e(z, \hat{y})], \quad (7)$$

and $c_u^*(s, y) = 0$ otherwise, where s is the signal about the quality of the match. The expression above is intuitive. The left-hand side is the value of keeping the worker unemployed. The right-hand side is the value of matching the worker to the firm. This is equal to the value of a worker employed in a match with idiosyncratic productivity z , where z is equal to s with probability α and to a value drawn randomly from the distribution f with probability $1 - \alpha$. The planner finds it optimal to create the match between the firm and the worker if and only if the left-hand side is smaller than the right-hand side. Notice that the left-hand side is independent of s , while the right-hand side is strictly increasing in s . Hence, the creation

probability $c_u^*(s, y)$ is an increasing function of s , and can be represented by a reservation signal $r_u^*(y)$ such that $c_u^*(s, y) = 0$ if $s < r_u^*(y)$ and $c_u^*(s, y) = 1$ if $s \geq r_u^*(y)$.

The efficient choice for the probability of creating a new match between a firm and an employed worker is $c_e^*(s, z, y) = 1$ if

$$y + z + \beta \mathbb{E} W_e(z, \hat{y}) \leq \alpha [y + s + \beta \mathbb{E} W_e(s, \hat{y})] + (1 - \alpha) \mathbb{E}_{z'} [y + z' + \beta \mathbb{E} W_e(z', \hat{y})], \quad (8)$$

and $c_e^*(s, z, y) = 0$ otherwise, where s is the signal about the quality of the match, and z is the idiosyncratic productivity of the worker's current match. The left-hand side of (8) is the value of keeping the worker in his current match. The right-hand side is the value of moving the worker into the new match. As in (7), this is the value of a worker employed in a match with idiosyncratic productivity z' , where z' is equal to s with probability α and to a value randomly drawn from the distribution f with probability $1 - \alpha$. The planner finds it optimal to create the new match between the firm and the worker if and only if the left-hand side is smaller than the right-hand side. It is easy to verify that the creation probability $c_e^*(s, z, y)$ is an increasing function of s . Hence, the creation policy can be represented by a reservation signal $r_e^*(z, y)$ such that $c_e^*(s, z, y) = 0$ if $s < r_e^*(z, y)$ and $c_e^*(s, z, y) = 1$ if $s \geq r_e^*(z, y)$. It is also easy to verify that $c_e^*(s, z, y)$ is a decreasing function of z . Hence, the cutoff signal $r_e^*(z, y)$ is increasing in z .

The efficient choice for the vacancy-to-worker ratio at the location where unemployed workers search for new matches is $\theta_u^*(y)$ such that

$$k \geq p'(\theta_u^*(y)) \sum_{s \geq r_u^*(y)} \left\{ \begin{array}{l} \alpha [y + s - b + \beta \mathbb{E} (W_e(s, \hat{y}) - W_u(\hat{y}))] \\ + (1 - \alpha) \mathbb{E}_{z'} [y + z' - b + \beta \mathbb{E} (W_e(z', \hat{y}) - W_u(\hat{y}))] \end{array} \right\} f(s), \quad (9)$$

and $\theta_u^*(y) \geq 0$, where the two inequalities hold with complementary slackness. The left-hand side of (9) is the marginal cost of increasing the vacancy-to-worker ratio. The right-hand side is the marginal benefit of increasing the vacancy-to-worker ratio, which is given by the product of two terms. The first term is the marginal increase in the probability with which an unemployed worker meets a firm. The second term is the expected value of a meeting between an unemployed worker and a firm. If $\theta_u^*(y)$ is positive, the marginal cost and the marginal benefit of increasing the vacancy-to-worker ratio must be equal. Otherwise, the marginal cost must be greater than the marginal benefit.

Similarly, the efficient choice for the vacancy-to-worker ratio at the location visited by workers who are employed in a match with idiosyncratic productivity z is $\theta_e^*(z, y)$ such that

$$k \geq p'(\theta_e^*(z, y)) \sum_{s \geq r_e^*(z, y)} \left\{ \begin{array}{l} \alpha [s - z + \beta \mathbb{E}(W_e(s, \hat{y}) - W_e(z, \hat{y}))] \\ + (1 - \alpha) \mathbb{E}_{z'} [z' - z + \beta \mathbb{E}(W_e(z', \hat{y}) - W_e(z, \hat{y}))] \end{array} \right\} f(s), \quad (10)$$

and $\theta_e^*(z, y) \geq 0$, where the two inequalities hold with complementary slackness. The interpretation of (10) is similar to that of (9), except that the increase in the vacancy-to-worker ratio takes place at the location that is visited by workers employed in matches of type z rather than at the location visited by unemployed workers. Notice that the left-hand side of (10) does not depend on z , while the right-hand side strictly decreases in z . Hence, as long as $\theta_e^*(z, y) > 0$, the vacancy-to-worker ratio $\theta_e^*(z, y)$ is a strictly decreasing function of z .

Finally, the efficient choice for the probability of destroying a match between a worker and a firm is $d^*(z, y) = 1$ if

$$\begin{aligned} b + \beta \mathbb{E}W_u(\hat{y}) > k \lambda_e \theta_e^*(z, y) + (1 - \lambda_e p(\theta_e^*(z, y)) m_e^*(z, y)) [y + z + \beta \mathbb{E}W_e(z, \hat{y})] \\ + \lambda_e p(\theta_e^*(z, y)) \mathbb{E}_{z'} \{ [\alpha c_e^*(z', z, y) + (1 - \alpha) m_e^*(z, y)] [y + z' + \beta \mathbb{E}W_e(z', \hat{y})] \}, \end{aligned} \quad (11)$$

and $d^*(z, y) = \delta$ otherwise, where z is the idiosyncratic productivity of the match. The left-hand side of (11) is the value of a worker who is unemployed and does not have the opportunity to search for a new match in the current period. This is the value of destroying the match. The right-hand side is the value of a worker who is employed in a match of type z and has the opportunity to search for a new match with probability λ_e . This is the value of keeping the match alive. When the left-hand side is greater than the right-hand side, the planner destroys the match with probability 1. Otherwise, Nature destroys the match with probability δ . Notice that the left-hand side does not depend on z , while the right-hand side is strictly increasing in z . Hence, the destruction probability $d^*(z, y)$ is a decreasing function of z , and can be represented by a reservation productivity $r_d^*(y)$ such that $d^*(z, y) = 1$ if $z < r_d^*(y)$ and $d^*(z, y) = \delta$ if $z \geq r_d^*(y)$.

We summarize the properties of the efficient choices in the proposition below.

Proposition 2 (*Planner's policy functions*): (i) The policy correspondences $(d^*, \theta_u^*, \theta_e^*, c_u^*, c_e^*)$ are single valued. (ii) There is $r_d^*(y)$ such that $d^*(z, y) = 1$ if $z < r_d^*(y)$ and $d^*(z, y) = \delta$ else. (iii) There is $r_u^*(y)$ such that $c_u^*(s, y) = 0$ if $s < r_u^*(y)$ and $c_u^*(s, y) = 1$ else. Similarly, there

is $r_e^*(z, y)$ such that $c_e^*(s, z, y) = 0$ if $s < r_e^*(z, y)$ and $c_e^*(s, z, y) = 1$ else. Moreover, $r_e^*(z, y)$ is increasing in z . (iv) $\theta_e^*(z, y)$ is decreasing in z .

With respect to a standard search model (e.g. Pissarides 1985), our model identifies a number of additional channels through which an aggregate productivity shock may affect the transitions of workers across employment states. First, by affecting not only θ_u^* and θ_e^* but also r_u^* and r_e^* , an aggregate productivity shock may affect not only the probability that a worker meets a firm but also the probability that a meeting between a firm and a worker turns into a match. Clearly, both channels may contribute to the response of the UE and EE rates to an aggregate productivity shock. Second, by affecting r_d^* , an aggregate productivity shock may affect the probability that the match between a firm and a worker is destroyed and, hence, it may affect the EU rate. As we shall see in sections 5 and 6, the quantitative importance of these additional channels depends on the informativeness of the signals, and on the shape of the distribution of match-specific productivity.

3 Decentralization

In this section we describe a market economy which decentralizes the efficient allocation. First, we describe the structure of the labor market and the nature of the employment contracts. Second, we formulate the conditions on the individual agent's value and policy functions that need to be satisfied in any equilibrium of the market economy. Third, we prove that an equilibrium of the economy exists, is unique and is block recursive. Specifically, because workers who are in matches with different levels of idiosyncratic productivity choose to search in different markets, the agents' value and policy functions depend on the aggregate state of the economy, ψ , only through the aggregate productivity, y , and not through the entire distribution of workers across employment states, (u, g) . Finally, we prove that the equilibrium decentralizes the efficient allocation. The reader who is mostly interested in the quantitative analysis can skip this section.

3.1 Market economy

In section 2, we only needed to describe the physical environment of the economy as we were concerned with the planner's problem. Here, we have to describe the structure of markets and the nature of the employment contracts as we are interested in the characterization of equilibrium. We assume that the labor market is organized in a continuum of submarkets indexed by (x, r) , $(x, r) \in \mathbb{R} \times Z$, where x is the value offered by a firm to a worker and r is a selection criterion based on the signal s . Specifically, when a firm meets a worker in submarket (x, r) , it hires the worker if and only if the signal s about the quality of their match is greater than or equal to r . If the firm hires the worker, it offers him an employment contract worth x in lifetime utility. The vacancy-to-worker ratio of submarket (x, r) is denoted as $\theta(x, r, \psi)$. In equilibrium, $\theta(x, r, \psi)$ will be consistent with the firms' and workers' search decisions.

At the separation stage, an employed worker moves into unemployment with probability $d \in [\delta, 1]$. At the search stage, each firm chooses how many vacancies to create and in which submarkets to locate them. On the other side of the market, each worker who has the opportunity to search chooses which submarket to visit. At the matching stage, each worker searching in submarket (x, r) meets a vacancy with probability $p(\theta(x, r, \psi))$. Similarly, each vacancy located in submarket (x, r) meets a worker with probability $q(\theta(x, r, \psi))$. When a worker and a vacancy meet in submarket (x, r) , the hiring process follows the rule specified for that submarket (i.e., the worker is hired if and only if the signal is higher than r and, conditional on being hired, he receives the lifetime utility x). At the production stage, an unemployed worker produces b units of output, and a worker employed in a match of type z produces $y + z$ units of output.

We assume that the contracts offered by firms to workers are bilaterally efficient, in the sense that they maximize the joint value of the match, i.e., the sum of the worker's lifetime utility and the firm's lifetime profits from the match. We focus on bilaterally efficient contracts not only because we are interested in decentralizing the solution to the planner's problem, but also because these contracts emerge as an equilibrium outcome under many different specifications of the contract space. For example, in a previous version of this paper (Menzio and Shi 2009a), we prove that the contracts offered by firms in equilibrium

are bilaterally efficient when the contract space is complete, in the sense that the employment contract can specify the wage w , the separation probability d , and the submarket where the worker searches while on the job, (x_e, r_e) , as functions of the history of the aggregate state of the economy, ψ , and the quality of the match, z . This result is intuitive. The firm maximizes its profits if it chooses the contingencies for d , x_e and r_e that maximize the joint value of the match, and the contingencies for w that provide the worker with the promised lifetime utility x . Moreover, we can prove that the contracts offered by firms in equilibrium are bilaterally efficient even when the employment contract can only specify the wage as a function of tenure and productivity. This result is also intuitive. The firm maximizes its profits if it chooses the wage in the first period of the employment relationship to provide the worker with the promised lifetime utility, x , and if it chooses wages in the subsequent periods to induce the worker to maximize the joint value of the match (e.g., by setting the wage equal to the worker's marginal product). Alternatively, equilibrium contracts are bilaterally efficient if they can specify severance transfers that induce the worker to internalize the effect of his separation and search decisions on the profits of the firm.

3.2 The problem of the worker and the firm

First, consider an unemployed worker at the beginning of the production stage, and let $V_u(\psi)$ denote his lifetime utility. In the current period, the worker produces and consumes b units of output. In next period, the worker matches with a vacancy with probability $\lambda_u p(\theta(x, r, \psi))m(r)$, where (x, r) is the submarket where the worker searches and $m(r) = \sum_{s \geq r} f(s)$ is the probability that the signal about the quality of the match is above the selection cutoff r . If the worker matches with a vacancy, his continuation utility is x . If the worker does not match with a vacancy, his continuation utility is $V_u(\hat{\psi})$. Thus, $V_u(\psi)$ is

$$V_u(\psi) = b + \beta \mathbb{E} \max_{(x, r)} \left\{ V_u(\hat{\psi}) + \lambda_u D(x, r, V_u(\hat{\psi}), \hat{\psi}) \right\}, \quad (12)$$

where D is defined as

$$D(x, r, V, \psi) = p(\theta(x, r, \psi))m(r)(x - V), \quad (13)$$

We denote as $(x_u(\hat{\psi}), r_u(\hat{\psi}))$ the policy functions for the optimal choices in (12).

Second, consider a worker and a firm who are matched at the beginning of the production stage. Let $V_e(z, \psi)$ denote the sum of the worker's lifetime utility and the firm's lifetime profits, given that the employment contract is bilaterally efficient. In the current period, the sum of the worker's utility and the firm's profit is equal to the output of the match, $y + z$. In the next period, the worker and the firm separate at the matching stage with probability d . In this case, the worker's continuation utility is $V_u(\hat{\psi})$ and the firm's continuation profit is zero. The worker and the firm separate at the next matching stage with probability $(1 - d)[\lambda_e p(\theta(x, r, \psi))m(r)]$, where (x, r) is the submarket where the worker searches for a new match. In this case, the continuation utility of the worker is x and the firm's continuation profit is zero. Finally, the worker and the firm remain together until the next production stage with probability $(1 - d)[1 - \lambda_e p(\theta(x, r, \psi))m(r)]$. In this case, the sum of the worker's continuation utility and the firm's continuation profit is $V_e(z, \hat{\psi})$. Thus, the joint value of the match $V_e(z, \psi)$ is

$$V_e(z, \psi) = y + z + \beta \mathbb{E} \max_{(d, x, r)} \left\{ d V_u(\hat{\psi}) + (1 - d) \left[V_e(z, \hat{\psi}) + \lambda_e D(x, r, V_e(z, \hat{\psi}), \hat{\psi}) \right] \right\} \quad (14)$$

where D is the function defined in (13). We denote as $d(z, \hat{\psi})$ and $(x_e(z, \hat{\psi}), r_e(z, \hat{\psi}))$ the policy functions for the optimal choices in (14).

At the search stage, a firm chooses how many vacancies to create and where to locate them. The firm's cost of creating a vacancy in submarket (x, r) is k . The firm's benefit from creating a vacancy in submarket (x, r) is

$$q(\theta(x, r, \psi)) \sum_{s \geq r} \{ [\alpha V_e(s, \psi) + (1 - \alpha) \mathbb{E}_z V_e(z, \psi) - x] f(s) \}, \quad (15)$$

where $q(\theta(x, r, \psi))$ is the probability of meeting a worker, $V_e(s, \psi)$ is the joint value of the match if the signal is correct, $\mathbb{E}_z V_e(z, \psi)$ is the joint value of the match if the signal is not correct, and x is the part of the joint value of the match that the firm delivers to the worker. When the cost is strictly greater than the benefit, the firm does not create any vacancy in submarket (x, r) . When the cost is strictly smaller than the benefit, the firm creates infinitely many vacancies in submarket (x, r) . And when the cost and the benefit are equal, the firm's profit is independent of the number of vacancies it creates in submarket (x, r) .

In any submarket that is visited by a positive number of workers, the tightness $\theta(x, r, \psi)$

is consistent with the firm's incentives to create vacancies if and only if

$$k \geq q(\theta(x, r, \psi)) \sum_{s \geq r} \{[\alpha V_e(s, \psi) + (1 - \alpha) \mathbb{E}_z V_e(z, \psi) - x] f(s)\}, \quad (16)$$

and $\theta(x, r, \psi) \geq 0$ with complementary slackness. In any submarket that workers do not visit, the tightness $\theta(x, r, \psi)$ is consistent with the firm's incentives to create vacancies if and only if k is greater or equal than (15). However, following the literature on directed search on the job with heterogeneous workers (i.e. Shi 2009, Menzio and Shi 2009, 2010, and Gonzalez and Shi, 2010), we restrict attention to equilibria in which $\theta(x, r, \psi)$ satisfies the above complementary slackness condition in every submarket.³

3.3 Equilibrium, block recursivity and efficiency

Definition 3 *A Block Recursive Equilibrium (BRE) consists of a market tightness function $\theta : \mathbb{R} \times Z \times Y \rightarrow \mathbb{R}_+$, a value function for the unemployed worker $V_u : Y \rightarrow \mathbb{R}$, a policy function for the unemployed worker $(x_u, r_u) : Y \rightarrow \mathbb{R} \times Z$, a joint value function for the firm-worker match $V_e : Z \times Y \rightarrow \mathbb{R}$, and policy functions for the firm-worker match $d : Z \times Y \rightarrow [\delta, 1]$ and $(x_e, r_e) : Z \times Y \rightarrow \mathbb{R} \times Z$. These functions satisfy the following conditions:*

- (i) $\theta(x, r, y)$ satisfies (16) for all $(x, r, \psi) \in \mathbb{R} \times Z \times \Psi$;
- (ii) $V_u(y)$ satisfies (12) for all $\psi \in \Psi$, and $(x_u(y), r_u(y))$ are the associated policy functions;
- (iii) $V_e(z, y)$ satisfies (14) for all $(z, y) \in Z \times \Psi$, and $d(z, y)$ and $(x_e(z, y), r_e(z, y))$ are the associated policy functions.

Condition (i) guarantees that the search strategy of an unemployed worker maximizes his lifetime utility, given the market tightness function θ . Condition (ii) guarantees that the employment contract maximizes the sum of the worker's lifetime utility and the firm's lifetime profits, given the market tightness function θ . Condition (iii) guarantees that the market tightness function θ is consistent with the firm's incentives to create vacancies. Taken together, conditions (i)-(iii) insure that in a BRE, just like in a recursive equilibrium, the

³This assumption pins down the tightness of an inactive submarket by a firm's indifference condition. That is, the tightness is such that a firm's expected profit from visiting any inactive submarket is equal to the firm's expected profit from visiting one of the active submarkets. A justification for this assumption comes from the following thought experiment. Imagine a sequential game in which unemployed workers choose (with a tremble) where to look for vacancies and, then, firms choose where to create their vacancies. Because of the tremble, the tightness is well defined everywhere. As the probability of the tremble goes to zero, the tightness of every submarket remains well defined and converges to the one given by (16).

strategies of each agent are optimal given the strategies of the other agents. However, unlike in a recursive equilibrium, in a BRE, the agent's value and policy functions depend on the aggregate state of the economy, ψ , only through the aggregate productivity, y , and not through the distribution of workers across different employment states, (u, g) . For this reason, a BRE is much easier to solve than a recursive equilibrium. But does a BRE exist? And why should we focus on a BRE rather than on a recursive equilibrium?

The following theorem answers these questions. Specifically, the theorem establishes that a BRE exists, that a BRE is unique and that it decentralizes the solution to the social planner's problem. Moreover, the theorem establishes that there is no loss in generality in focusing on the BRE because all equilibria are block recursive.

Theorem 4 (*Block recursivity, uniqueness and efficiency of equilibrium*):(i) All equilibria are block recursive. (ii) There exists a unique BRE. (iii) The BRE is socially efficient in the sense that: (a) $\theta(x_u(y), r_u(y), y) = \theta_u^*(y)$, and $r_u(y) = r_u^*(y)$; (b) $d(z, y) = d^*(z, y)$; (c) $\theta(x_e(z, y), r_e(z, y), y) = \theta_e^*(z, y)$, and $r_e(z, y) = r_e^*(z, y)$.

The equilibrium is block recursive because searching workers are endogenously separated in different markets and, as in the social planner's problem, such separation is possible only when search is directed. To explain why directed search induces workers to separate endogenously, note that workers choose in which submarket to search in order to maximize the product between the probability of finding a new match and the value of moving from their current employment position to the new match. For a worker in a low-value employment position (unemployment or employment in a low quality match), it is optimal to search in a submarket where the probability of finding a new match is relatively high and the value of the match is relatively low. For a worker in a high-value employment position (i.e., employment in a high quality match), it is optimal to search in a submarket where the probability of finding a new match is relatively low and the value of the match is relatively high. Overall, workers in different employment positions choose to search in different submarkets. As a result of the self-selection of workers, a firm that opens a vacancy in submarket (x, r) knows that it will only meet one type of worker. For this reason, the expected value to the firm from meeting a worker in submarket (x, r) does not depend on the entire distribution of workers across employment states and, because of the free entry condition (16), the probability

that a firm meets a worker in submarket (x, r) has the same property. Since the meeting probability across different submarkets is independent from the distribution of workers across employment states, it is easy to see from (12) and (14) that the value of unemployment and the joint value of a match will also be independent from the distribution.

If we replaced the assumption of directed search with random search, the equilibrium could not be block recursive. Under random search, workers in high and low-value employment positions all have to search in the same market. When this is the case, the firm's expected value from meeting a worker depends on how workers are distributed across different employment positions, as this distribution determines the probability that the employment contract offered by the firm will be accepted by a randomly selected worker. In turn, the free-entry condition implies that the probability that a firm meets a worker must also depend on the distribution of workers. Since the meeting probability between firms and workers depends on the distribution, so do all of the agents' value and policy functions.⁴

It is important to clarify that the assumption of bilaterally efficient contracts is not necessary for establishing the existence of a block recursive equilibrium. In fact, in some of our work (Shi 2009, Menzio and Shi 2009b, 2010), we have shown that block recursive equilibria exist also in economies where the contract space is so limited that bilateral efficiency cannot be attained (e.g., economies in which contracts can only specify a wage that remains constant over the entire duration of an employment relationship).⁵

However, we make use of the assumption of bilaterally efficient contracts in order to establish the equivalence between the block recursive equilibrium and the social plan, and to rule out equilibria that are not block recursive. When contracts are bilaterally efficient, the joint value of a match to the firm and the worker satisfies the equilibrium condition (14). After solving the free-entry condition (16) for x and substituting the solution into (14), we

⁴One should clearly distinguish block recursivity from the property that the market tightness is independent of unemployment in simple models of random search (e.g. Pissarides 1985, Mortensen and Pissarides 1994). The latter feature arises only when searching workers are identical, so that a vacancy knows exactly the type of worker it will meet. In fact, when there is on-the-job search or when searching workers are heterogeneous ex ante, random search will cause the market tightness to depend on their distribution.

⁵Menzio and Moen (2008) establish the existence of a block recursive equilibrium in models of directed search off the job in which contracts are not bilaterally efficient.

can rewrite this equilibrium condition as

$$\begin{aligned}
V_e(z, \psi) = & y + z \\
& + \beta \mathbb{E} \max_{(d, \theta, r)} \left\{ dV_u(\hat{\psi}) - (1-d)\lambda_e k \theta + (1-d)\lambda_e V_e(z, \hat{\psi}) \right. \\
& \left. + (1-d)\lambda_e p(\theta) \sum_{s \geq r} \left[\alpha V_e(s, \hat{\psi}) + (1-\alpha) \mathbb{E}_z V_e(z, \hat{\psi}) - V_e(z, \hat{\psi}) \right] f(s) \right\}.
\end{aligned} \tag{17}$$

One can easily verify that the functional equation (17) is satisfied not only by the joint value of a match to the firm and the worker, $V_e(z, \psi)$, but also by the value of an employed worker to the planner, $y + z + \beta \mathbb{E} W_e(z, \hat{y})$. Moreover, one can easily verify that the functional equation (17) is a contraction mapping and, hence, it admits a unique solution. Therefore, the joint value of a match to the firm and the worker must be equal to the value of an employed worker to the planner. Similarly, one can establish the equivalence between the value of unemployment to a worker, $V_u(\psi)$, and the value of an unemployed worker to the planner, $b + \beta \mathbb{E} W_u(\hat{y})$. The equivalence between the value functions of individual agents and the component value functions of the planner is sufficient for establishing that any equilibrium is efficient and block recursive.

4 Data and Calibration

In the remainder of the paper, we want to use the model to measure the contribution of aggregate productivity shocks to the cyclical volatility of the US labor market. We accomplish this task in three steps. First, we calibrate the parameters of the model to match as many of the acyclical features of the US labor market as possible. Second, we feed into the calibrated model aggregate productivity shocks of the same magnitude and persistence as those observed in the US economy. This step allows us to uncover how the US labor market would behave if aggregate productivity shocks were, counterfactually, the only shocks affecting the economy. Third, we compare the actual and the counterfactual behavior of the US business cycle. This last step allows us to gauge the contribution of aggregate productivity shocks to the cyclical volatility of the US labor market. The three-step procedure described above is the same procedure typically followed in the real business cycle literature to measure the importance of TFP shocks (e.g. Cooley 1995). We carry out this procedure for two versions of the model which are, a priori, equally plausible, i.e., the version of the model in

which matches are pure experience goods (i.e., $\alpha = 0$) and the version of the model in which matches are inspection goods (i.e., $\alpha = 1$). In this section, we calibrate these two versions of the model and review the key features of the US labor market over the business cycle. In section 5, we simulate the version of the model in which matches are experience goods. In section 6, we simulate the version of the model in which matches are inspection goods.

4.1 Calibration

The preference parameters are the discount factor β and the value of leisure b . The firm's technology is described by the vacancy cost k , the distribution of the match-specific component of productivity f , the stochastic process for aggregate productivity ϕ , and the exogenous match-destruction probability δ . We restrict f to be a 200 point approximation of a Weibull distribution with mean μ_z , shape ν_z , and scale σ_z .⁶ We also restrict the stochastic process for aggregate productivity to be a 3-state Markov process with unconditional mean μ_y , autocorrelation ρ_z , and standard deviation σ_y . The matching process is described by the search probabilities λ_u and λ_e , the meeting probability p , and the precision of the signal about the quality of a new match, α . As in most of the literature (e.g., Shimer 2005, Mortensen and Nagypál 2007), we restrict p to be of the form $p(\theta) = \min\{\theta^\gamma, 1\}$, $\gamma \in (0, 1)$.

We choose the model period to be one month. We set β so that the real interest rate in the model is 5 percent per year. We choose k and δ so that the average UE and EU transition rates are the same in the model and in the data. In the model, the UE rate is given by $h^{ue} = p(\theta_u)m_u$, and the EU rate is given by $h^{eu} = [\sum d(z)g(z)]/(1-u)$. In the data, we measure these transition rates following the methodology developed by Shimer (2005).⁷

⁶The Weibull density function is:

$$f(z) = \frac{\nu_z}{\sigma_z} \left(\frac{z - \mu_z}{\sigma_z} + \Gamma \left(\frac{1}{\nu_z} + 1 \right) \right)^{\nu_z - 1} \exp \left[- \left(\frac{z - \mu_z}{\sigma_z} + \Gamma \left(\frac{1}{\nu_z} + 1 \right) \right)^{\nu_z} \right],$$

where Γ is the gamma function. The parameters ν_z and σ_z control respectively the shape and the variance of the distribution. In particular, the shape of the Weibull distribution is similar to the shape of the exponential distribution for $\nu_z = 1$, to the lognormal distribution for $\nu_z = 2$, to the normal distribution for $\nu_z = 4$, and to a left-skewed version of a normal distribution for $\nu_z = 10$. To keep the calibration manageable, we restrict attention to these four values of ν_z .

⁷There are two differences between the measures of the UE and EU rates constructed by Shimer (2005) and ours. First, Shimer multiplies the short-term unemployment rate by 1.1 in every month after February 1994 in order to correct for the fact that the 1994 redesign of the CPS changed the way in which unemployment duration is measured. In this paper, we follow Elsby et al. (2009) who argue that the short-term unemployment rate should be multiplied by 1.15 not 1.1. Second, Shimer corrects the measures of the UE

Specifically, we measure the UE rate in month t as $h_t^{ue} = u_{t+1}^s / (1 - u_t)$, where u_t is the CPS unemployment rate in month t , and u_{t+1}^s is the CPS short-term unemployment rate in month $t + 1$. Similarly, we measure the EU rate in month t as $h_t^{eu} = 1 - (u_{t+1} - u_{t+1}^s) / u_t$.

We normalize λ_u to 1, and we choose λ_e so that the average EE transition rate is the same in the model as in the data. The EE rate in the model is given by $h^{ee} = [\sum (1 - d(z)) \lambda_e p(\theta_e(z)) m_e(z) g(z)] / (1 - u)$. The EE rate in the data has been measured by Nagypal (2008) using the CPS microdata. Specifically, Nagypal measures the EE rate in month t as $h_t^{ee} = s_t / e_t$, where s_t is the number of workers who are employed at different firms in months t and $t + 1$, and e_t is the number of workers who are employed in month t .

We choose γ so that the elasticity of the UE rate with respect to the vacancy-to-unemployment ratio is the same in the model as in the data. In the model, the vacancy-to-unemployment ratio is given by v/u , where the aggregate measure of vacancies v is given by the sum of $\lambda_u \theta_u u$ and $\sum (1 - d(z)) \lambda_e \theta_e(z) g(z)$. In the data, the vacancy-to-unemployment ratio is measured as the ratio of the Conference Board Help-Wanted Index and the CPS unemployment rate. It is important to note that γ will generally be different from the elasticity of the UE rate with respect to the vacancy-to-unemployment ratio, because in the model $\log h^{ue} = \gamma \log \theta_u + \log m_u$.

Normalize μ_z to 0. We choose the scale ν_z and shape σ_z parameters in the distribution of the idiosyncratic productivity to minimize the distance between the tenure distribution generated by the model and its empirical counterpart.⁸ In the model, the tenure distribution is defined as the fraction of workers who are employed and have been in the same match for t years. In the data, the analogous distribution is measured by Diebold, Neumark and Polsky (1997) using the 1987 CPS tenure supplement.⁹ It is easy to understand why the

and EU rates for time aggregation bias. This correction is appropriate for a continuous-time model in which the rate at which a worker moves from one employment state to another does not depend on how long the worker has been in that state (i.e., a model in which workers' transition rates are duration independent). Unfortunately, the workers' transition rates are duration dependent in our model and, hence, Shimer's correction is not appropriate. Moreover, in order to build the correction that is right for our model, we would have to develop a continuous time version of it. In the face of these difficulties, we simply decided not to correct the UE and EU rates for time aggregation bias.

⁸This identification strategy has a precedent in Moscarini (2003), who considers a model of random search on the job in which workers and firms learn over time the quality of their match by observing their output. He uses the empirical tenure distribution to identify the precision of output as a signal of match quality.

⁹Diebold, Neumark and Polsky (1997) also show that the empirical tenure distribution is stable over time. For this reason, it is appropriate to compare the empirical tenure distribution observed in 1987 with the tenure distribution generated by the steady-state of the model.

tenure distribution helps identifying σ_z and ν_z . For example, if $\sigma_z = 0$, all matches have the same idiosyncratic productivity and the same, constant probability of surviving from one year to the next. Hence, if $\sigma_z = 0$, the average survival probability of a match that has reached tenure t is independent of t . If $\sigma_z > 0$, on the other hand, different matches have a different idiosyncratic productivity and a different probability of surviving from one year to the next. For this reason, if $\sigma_z > 0$, the average survival probability of a match that has reached tenure t is increasing in t . Clearly, both the scale σ_z and the shape ν_z of the match distribution affect the relationship between survival probability and tenure. In turn, since the slope of the tenure distribution is the average survival probability of a match with tenure t , σ_z and ν_z can be recovered from the tenure distribution.

We normalize μ_y to 1, and choose ρ_y and σ_y so that the average productivity of labor in the model has the same autocorrelation and standard deviation as in the data. In the model, the average productivity of labor is measured as $\pi = [\sum(y + z)g(z)] / (1 - u)$. In the data, average labor productivity is measured as the CPS output per worker in the non-farm business sector. Note that, because the distribution of workers across matches with different idiosyncratic productivity may vary over time, the autocorrelation and standard deviation of average labor productivity need not be the same as ρ_y and σ_y . Finally, we choose b so that the ratio of the value of leisure to the average productivity of labor is 0.71, the value recently estimated by Hall and Milgrom (2008).

Table 2 presents the outcomes of the calibration for the version of the model in which matches are experience goods, and for the version in which matches are inspection goods. For these two versions, Figures 1 and 2 illustrate the calibrated distribution of the idiosyncratic productivity of new matches, f , as well as the steady-state distribution of employed workers across matches with different idiosyncratic productivity, $g/(1 - u)$. As one can immediately see by inspecting Table 2 and Figures 1 and 2, the calibration outcomes for the two versions of the model are very different even though the calibration targets are exactly the same. We will discuss the causes and consequences of some of these differences in section 6.

4.2 Business cycle facts

We want to use the calibrated model to understand labor market fluctuations over the business cycle. Table 1 presents summary statistics of the US labor market over the period 1951(I)-2006(II), where the cyclical component of each variable is computed as the difference between the log of the variable and an HP-trend with smoothing parameter 1600. Several features of the data are worth noting: (i) Unemployment and labor productivity are negatively correlated, but unemployment is much more volatile than labor productivity. Specifically, the standard deviation of unemployment is 9.5 times larger than that of labor productivity. (ii) Vacancies and unemployment are nearly perfectly negatively correlated and have similar volatility. Specifically, the correlation between vacancies and unemployment is -0.9 and the standard deviation of the former relative to the latter is 1.1. (iii) The workers' transition rates between unemployment, employment and across employers are approximately half as volatile as unemployment and are strongly correlated with it. Specifically, unemployment displays a strong negative correlation with the UE and EE rates (respectively, -0.92 and -0.63), and a strong positive correlation with the EU rate (0.78).

5 Experience Model

In this section, we use the experience model to measure the contribution of aggregate productivity shocks (henceforth, y -shocks) to the cyclical volatility of the US labor market. We find that, according to this model, y -shocks generate the same pattern of comovement between labor market variables as in the data. Moreover, y -shocks can account for a substantial fraction of the empirical volatility of unemployment and workers' transition rates, as well as for a smaller (yet non-negligible) fraction of the empirical volatility of vacancies.

5.1 Transitional dynamics

To illustrate the mechanics of the experience model, we begin this section by characterizing the response of the labor market to a 1 percent permanent increase to aggregate productivity. Using the first order conditions of the planner's problem, we characterize the responses to the shock of: (i) θ_u and θ_e , the vacancy-to-worker ratio at the locations where unemployed and employed workers search for new matches, and (ii) r_d , the cutoff on the idiosyncratic

productivity below which a match is destroyed. Then, using the laws of motion for the distribution of workers across employment states, we characterize the responses of unemployment, vacancies and workers' transition rates. We carry out this analysis under the assumption that the economy is at its non stochastic steady-state when it is hit by the aggregate productivity shock. Note that, since in the experience model signals are uninformative, the creation cutoffs r_u and r_e play no role in the analysis.

In the experience model, $\alpha = 0$ and the vacancy-to-worker ratio θ_u is such that

$$k \geq p'(\theta_u) \{ \mathbb{E}_{z'} [y + z' + \beta \mathbb{E}W_e(z', \hat{y})] - [b + \beta \mathbb{E}W_u(\hat{y})] \}, \quad (18)$$

and $\theta_u \geq 0$ with complementary slackness. The term in brackets on the right-hand side of (18) is the difference between the value of a worker employed in a new match and the value of an unemployed worker. The y -shock increases this difference and, hence, it increases the vacancy-to-worker ratio θ_u .

The vacancy-to-worker ratio $\theta_e(z)$ is such that

$$k \geq p'(\theta_e(z)) \{ \mathbb{E}_{z'} [y + z' + \beta \mathbb{E}W_e(z', \hat{y})] - [z + \beta \mathbb{E}W_e(z, \hat{y})] \}, \quad (19)$$

and $\theta_e(z) \geq 0$ with complementary slackness. The term in brackets on the right-hand side of (19) is the difference between the value of a worker employed in a new match and the value of a worker employed in match of type z . The y -shock increases the value of a worker employed in a new match. Moreover, the y -shock increases the value of a worker employed in a type- z match, and the increase is larger the higher is z . This property is intuitive because workers who are employed in better matches have a higher probability of being employed in the future and, hence, a higher probability of taking advantage of the increase in y . Overall, the effect of the y -shock on the difference between the value of a worker employed in a new match and in a type- z match is positive for relatively low values of z , and negative otherwise. Consequently, the effect of the y -shock on the vacancy-to-worker ratio $\theta_e(z)$ is positive for relatively low values of z , and negative otherwise.

The endogenous match destruction cutoff r_d is the solution for z to the following equation

$$b + \beta \mathbb{E}W_u(\hat{y}) = -k\lambda_e\theta_e(z, y) + [1 - \lambda_e p(\theta_e(z, y))] [y + z + \beta \mathbb{E}W_e(z, \hat{y})] + \lambda_e p(\theta_e(z, y)) \mathbb{E}_{z'} \{ [y + z' + \beta \mathbb{E}W_e(z', \hat{y})] \}, \quad (20)$$

The term on the right-hand side of (20) is the value of a worker employed in a type- z match who has the option of searching for a better match with probability λ_e . The term on the left-hand side of (20) is the value of an unemployed worker who does not have the option to look for a new match. The y -shock causes the term on the right to increase more than the term on the left and, hence, it causes the endogenous destruction cutoff r_d to fall.

Figure 3 shows how the y -shock affects the UE rate, $h^{ue} = \theta_u^\gamma$, the EU rate, $h^{eu} = [\sum d(z)g(z)] / (1 - u)$, and the EE rate, $h^{ee} = [\sum (1 - d(z))\lambda_e\theta_e(z)^\gamma] / (1 - u)$. The UE rate increases because the y -shock raises θ_u , the vacancy-to-worker ratio in the location where unemployed workers look for new matches. The EU rate falls because the y -shock lowers r_d and, consequently, it lowers the fraction of new matches that are endogenously destroyed after their idiosyncratic productivity is observed. On impact, the EE rate increases because the y -shock raises the average vacancy-to-worker ratio in the location where employed workers look for new matches. Over time, the EE rate continues to grow because the y -shock shifts the distribution of employed workers towards matches with lower productivity, which have a higher probability of terminating with an EE transition. Quantitatively, the 1 percent increase in y leads to a 2 percent increase in the steady state UE rate, a 4 percent decrease in the steady state EU rate, and a 3.5 percent increase in the steady state EE rate. As a result of both the increase in the UE rate and the decline in the EU rate, the steady state unemployment rate falls by 6 percent.

Figure 4 shows how the y -shock affects the number of vacancies for unemployed workers, $v_u = u\theta_u$, the number of vacancies created for employed workers, $v_e = \sum (1 - d(z))\lambda_e\theta_e(z)g(z)$, and the total number of vacancies in the economy, $v = v_u + v_e$. On impact, v_u increases because the y -shock raises θ_u , the number of vacancies that are created for each unemployed worker. Over time, as the number of unemployed workers falls toward its new steady state level, v_u returns to its initial steady state value and then falls below it. The response of v_e to the shock is different. On impact, v_e increases because the y -shock raises the average θ_e , the number of vacancies that are created for each employed workers. Over time, as the number of employed workers grows towards its new steady state level, v_e continues to increase. Quantitatively, the y -shock leads to a 3 percent decline in the steady state value of v_u and to a 5 percent increase in the steady state value of v_e . Since $v = v_u + v_e$ and $v_u \sim v_e$, the

y -shock leads to a 2 percent increase in the steady state value of v .

Figure 5 shows how the y -shock affects the average idiosyncratic productivity, $\bar{z} = [\sum zg(z)] / (1 - u)$, and the average labor productivity, $\pi = y + \bar{z}$. The y -shock has two opposing effects on z . On the one hand, the y -shock tends to lower z because it lowers the endogenous destruction cutoff r_d . On the other hand, the y -shock tends to increase z because it increases the probability that a worker employed in a match with a relatively low z will find another match. In practice, the first effect dominates, and the y -shock leads to a 0.3 percent decline in the steady state value of the average idiosyncratic productivity. Since $\pi = y + \bar{z}$ and $\bar{z} \sim y/3$, the y -shock leads to a 0.7 percent increase in the steady state value of average labor productivity.

5.2 Productivity shocks and labor market fluctuations

Now, we simulate the model to find out what the cyclical behavior of the labor market would be like if, counterfactually, y -shocks were the only source of cyclical fluctuations in the US economy. To this aim, we draw a long time-series for aggregate productivity from the calibrated stochastic process of y . Then, we feed this time-series into the model to generate the time-series for unemployment, vacancies, workers' transition rates and labor productivity. Finally, we isolate the cyclical component of these model-generated data by taking the difference between logs and an HP-trend with smoothing parameter 1600.

Table 3 summarizes the results of the simulation. That is, Table 3 presents the standard deviation and the correlation between labor market variables that we would observe if y -shocks were the only source of cyclical fluctuations in the US economy. A comparison between this table and the table that summarizes the actual behavior of the US labor market (i.e., Table 1) reveals a number of interesting facts. First, y -shocks alone can generate the same pattern of comovement between unemployment, vacancies and workers' transition rates that we observe in the US economy. Specifically, y -shocks generate fluctuations in unemployment that are nearly perfectly negatively correlated with the fluctuations in the UE rate, the EE rate and vacancies, and that are nearly perfectly positively correlated with the fluctuations in the EU rate. Second, y -shocks alone can generate a substantial fraction of the actual volatility of unemployment and workers' transition rates. Specifically, y -shocks generate

85 percent of the actual volatility of unemployment, and 42, 110 and 90 percent of the actual volatility of the UE, EU and EE rates. Third, y -shocks do not generate much of the empirical volatility of vacancies (approximately 25 percent). Fourth, y -shocks cannot account for the empirical correlation between unemployment and labor productivity. In the US economy, this correlation is -0.4. If y -shocks were the only source of cyclical fluctuations, this correlation would be close to -1.

The results above suggest that aggregate productivity shocks are an important source of cyclical fluctuations in the US labor market, but definitely not the only one. There must be additional shocks (e.g., monetary, fiscal or sectoral shocks) that are imperfectly correlated with y -shocks and that account for the difference between the volatility of unemployment, vacancies and workers' transition rates that is generated by y -shocks and their actual volatility. Hence, these additional shocks must be able to increase substantially the volatility of vacancies and to somewhat dampen the volatility of the EU rate.

Naturally, the statements above are valid only if our model is correctly specified. For example, it could be the case that our model overpredicts the effects of y -shocks on the EU rate because it incorrectly abstracts from firing costs. Similarly, it could be the case that our model underpredicts the effect of y -shocks on vacancies because it incorrectly specifies that vacancies depreciate instantaneously.

5.3 Role of match heterogeneity and search on the job

Two classic models are nested into ours. The model by Pissarides 1985 (henceforth, P85) is a version of our model without match heterogeneity or search on the job. That is, P85 is a version of our model in which the parameters λ_e and σ_z are constrained to be zero. The model by Mortensen and Pissarides 1994 (henceforth, MP94) is a version of our model without search on the job. That is, MP94 is a version of our model in which λ_e is constrained to be zero. It is evident from the outcome of our calibration that the constraints $\lambda_e = 0$ and $\sigma_z = 0$ are rejected by the data. However, for the purposes of this paper, a more relevant question is whether these constraints lead to distortions in the measurement of the effect that y -shocks have on the US labor market. To answer this question, we calibrate P85 and MP94 using the same targets described in the previous section, and then we simulate the effect of

y -shocks on unemployment, vacancies and other labor market variables. Table 2 reports the outcomes of the calibrations. Tables 4 and 5 summarize the outcome of the simulations.

First, the contribution of y -shocks to the empirical volatility of the UE rate is 15 percent when it is measured with P85 and MP94, and 42 percent when it is measured with our model. This difference is easy to explain. For all three models, the value of the parameter γ is chosen so that the elasticity of the UE rate with respect to the vacancy-unemployment ratio is the same as in the data (0.27). That is, γ is chosen so that

$$\begin{aligned} d \log h^{ue} / d \log(v/u) &= 0.27, \\ \iff \gamma &= 0.27 d \log(v/u) / d \log(v_u/u), \end{aligned} \tag{21}$$

where the second line makes use of the fact that $h^{ue} = \theta_u^\gamma$. In our model, $\gamma = 0.6$ because the vacancy-to-unemployment ratio, v/u , is 2.2 times as volatile as the vacancy-to-worker ratio in the location where unemployed workers look for matches, v_u/u . In P85 and MP94, the absence of on-the-job search implies $\gamma = 0.27$ because v/u is always equal to v_u/u . Therefore, $\gamma = 0.27$. In turn, a lower value of γ implies a lower elasticity of the UE rate with respect to y because

$$\frac{d \log h^{ue}}{d \log y} = \gamma \frac{d \log(v_u/u)}{d \log y}. \tag{22}$$

Second, the contribution of y -shocks to the empirical volatility of the EU rate is zero when measured with P85, and approximately 100 percent when it is measured with either MP94 or our model. These differences are easy to understand. In our model and in MP94, aggregate productivity shocks affect the endogenous destruction cutoff r_d and, through this channel, they affect the fraction of new matches that are destroyed after their idiosyncratic productivity is observed. Quantitatively, the effect of y -shocks on the EU rate is large because, according to the calibration, the productivity distribution f has a high density around the steady-state value of r_d . In P85, aggregate productivity shocks have no effect on the EU rate because all matches are assumed to be identical.

Third, the contribution of y -shocks to the empirical volatility of unemployment is 10 percent when measured with P85, 60 percent when measured with MP94, and 85 percent when measured with our model. The difference between our model and P85 is due to the fact that, in P85, y -shocks generate smaller fluctuations in both the UE and the EU rates. The difference between our model and MP94 is due to the fact that, in MP94, y -shocks generate

smaller fluctuations in the UE rate.

Finally, we compare the effect of y -shocks on vacancies as measured by the three models. In P85, y -shocks generate the same type of fluctuations in the total number of vacancies as they do in our model. That is, y -shocks generate fluctuations in the total number of vacancies that are approximately 2.5 times as large as the fluctuations in average productivity and are strongly negatively correlated with the fluctuations in unemployment. However, the similarity in the behavior of the total number of vacancies masks important differences in the behavior of the vacancies that are created for unemployed workers, v_u , and in the number of vacancies for employed workers, v_e . First, the correlation between v_u and u is negative in P85 and positive in our model. Second, the volatility of v_e is zero in P85 and large in our model. It is easy to understand these two differences. In both models, an increase in y leads to an increase in θ_u and to a decline in u . In our model, the decline in u is larger (in percentage terms) than the increase in θ_u and, hence, v_u falls. In P85, the decline in u is smaller and v_u increases. Moreover, in our model, an increase in y leads to an increase in both the employment rate and the average tightness of the location visited by employed workers. Hence, it leads to an increase in v_e . In P85, an increase in y has no effect on v_e because the model rules out search on the job.

In MP94, y -shocks generate very different fluctuations in the total number of vacancies than they do in our model. Specifically, in MP94, y -shocks generate fluctuations in the total number of vacancies that are positively correlated with the fluctuations in unemployment. In our model, y -shocks generate fluctuations in total vacancies that are strongly negatively correlated with the fluctuations in u . This difference is easy to understand. In both models, an increase in y leads to a decline in v_u . In our model, an increase in y also leads to an increase in v_e which is sufficiently large to drive v up. In MP94, an increase in y has no effect on v_e and, hence, v falls.

Let us summarize our findings. The constraints in P85 (i.e., $\sigma_z = 0$ and $\lambda_e = 0$) lead to a downward bias in the measurement of the volatility of the UE and EU rates caused by aggregate productivity shocks and, consequently, lead to a downward bias in the measurement of the volatility of unemployment. Moreover, the constraints in P85 distort the measurement of the effect of y -shocks on the number of vacancies that are created for unemployed and

employed workers. However, since these distortions have opposite signs, the measurement of the effect of y -shocks on the total number of vacancies is close to the one obtained with our model. The above observations explain why, when researchers use models that abstract from both search on the job and match heterogeneity, they either find that y -shocks account for a very small fraction of the empirical volatility of unemployment (e.g. Shimer 2005), or that y -shocks can only generate large unemployment fluctuations if the targets of the calibration are modified (e.g. Mortensen and Nagypal 2006, Hagedorn and Manovskii 2008) or if additional amplification mechanisms are introduced (e.g. hiring costs in Mortensen and Nagypál 2007, countercyclical vacancy costs in Shao and Silos 2009, wage rigidity in Hall 2005, Menzio 2005, Menzio and Moen 2008, Kennan 2009, Gertler and Trigari 2009).

The constraint in MP94 (i.e. $\lambda_e = 0$) leads to a downward bias in the measurement of the volatility of the UE rate that is caused by aggregate productivity shocks and, consequently, to a downward bias in the measurement of the volatility of unemployment. Moreover, the constraint in MP94 leads to a distortion in the measurement of the effect that y -shocks have on the total number of vacancies. This last observation explains why, when researchers use models that do not allow from search on the job, they either find that y -shocks generate a positive (or weakly negative) comovement between unemployment and vacancies (e.g. Mortensen and Pissarides 1994, Merz 1995), or they completely abstract from vacancies (e.g. Gomes, Greenwood and Rebelo 2001).

5.4 Role of directed search

Models of random search on the job are difficult to solve outside of the steady state because the equilibrium is not block recursive. That is, the agents' value and policy functions depend on the aggregate state of the economy, ψ , not only on the aggregate productivity, y , but also on the entire distribution of workers across different employment states, (u, g) . To get around this technical difficulty, the existing literature has had to impose some strong assumptions on the environment which limit the scope of their analysis.

Moscarini and Postel-Vinay (2009) consider a model of random search on the job in which firms are heterogeneous with respect to their productivity. In order to solve this model outside of the steady-state, Moscarini and Postel-Vinay need to assume that the rate at

which firms and workers come into contact is exogenous. For this reason, their model cannot be used to measure the effect of aggregate productivity shocks on the workers' transition rates, which is the main object of interest in this paper. On the other hand, their model can be used to measure the effect of aggregate productivity shocks on the distribution of workers across wages, an object that is indeterminate in our model of bilaterally efficient contracts and perfectly transferable utility. Overall, we see the paper by Moscarini and Postel-Vinay (2009) more as a complement to our paper than as a substitute.

Robin (2009) considers a model of random search on the job in which workers are ex-ante heterogeneous with respect to their productivity. In order to solve his model outside of the steady-state, Robin also needs to assume that the rate at which firms and workers come into contact is exogenous. However, his model yields predictions on the effect of productivity shocks on the workers' transition rates because the fraction of meetings between firms and workers that turns into matches is endogenous. As we have discussed above, this channel is not active in the version of our model in which matches are experience goods. As we shall see in the next section, this channel is quantitatively unimportant in the version of our model in which matches are inspection goods.

6 Inspection Model

In this section, we use the inspection model to measure the contribution of aggregate productivity shocks (henceforth, y -shocks) to the cyclical volatility of the US labor market. We find that, according to this model, y -shocks generate the same pattern of comovement between unemployment, vacancies and workers' transition rates that is observed in the data. However, y -shocks account for very little of the empirical volatility of unemployment, vacancies and workers' transition rates.

6.1 Transitional dynamics

To illustrate the mechanics of the inspection model, we begin this section by characterizing the response of the labor market to a 1 percent permanent increase to the aggregate productivity. Using the first order conditions of the planner's problem, we characterize the response to the shock of: (i) r_d , the cutoff on the idiosyncratic productivity below which a match is

endogenously destroyed; (ii) r_u and r_e , the cutoffs on the signal about the idiosyncratic productivity above which a meeting between a firm and a worker is created; (iii) θ_u and θ_e , the vacancy-to-worker ratio in the locations where unemployed and employed workers look for new matches. Then, using the laws of motion for the distribution of workers across employment states, we characterize the response of unemployment, vacancies and workers' transition rates. We carry out the analysis under the assumption that the economy is at its non stochastic steady-state when it is hit by the aggregate productivity shock.

The endogenous match destruction cutoff r_d is the solution for z to the following equation

$$b + \beta \mathbb{E}W_u(\hat{y}) = -k\lambda_e\theta_e(z, y) + [1 - \lambda_e p(\theta_e(z, y))m_e(z, y)] [y + z + \beta \mathbb{E}W_e(z, \hat{y})] + \lambda_e p(\theta_e(z, y)) \sum_{s \geq r_e(z)} \{[y + s + \beta \mathbb{E}W_e(s, \hat{y})]f(s)\}. \quad (23)$$

The term on the left-hand side of (23) is the value of a worker employed in a type- z match who has the option of searching for a better match with probability λ_e . The term on the right-hand side of (23) is the value of an unemployed worker who does not have the option to look for a new match. The y -shock causes the term on the left to increase more than the term on the right and, hence, it causes the endogenous destruction cutoff r_d to fall.

The endogenous match creation cutoff r_u is the solution for s to the following equation

$$y + s + \beta \mathbb{E}W_e(s, \hat{y}) = b + \beta \mathbb{E}W_u(\hat{y}) \quad (24)$$

The term on the left-hand side of (24) is the expected value of creating a match given the signal s . In the inspection model, signals are fully informative and, hence, this value is equal to the value of a worker employed in a match of type s . The term on the right-hand side of (24) is the value of an unemployed worker. The y -shock causes the term on the left to increase more than the term on the right and, hence, it causes the endogenous creation cutoff r_u to fall. Similarly, the endogenous match creation cutoff $r_e(z)$ is the solution for s to the following equation

$$y + s + \beta \mathbb{E}W_e(s, \hat{y}) = y + z + \beta \mathbb{E}W_e(z, \hat{y}) \quad (25)$$

Clearly, $r_e(z)$ is equal to z and is unaffected by the y -shock.

The vacancy-to-worker ratio θ_u is such that

$$k \geq p'(\theta_u) \left\{ \sum_{s \geq r_u} [y + s - b + \beta \mathbb{E}(W_e(s, \hat{y}) - W_u(\hat{y}))] f(s) \right\}, \quad (26)$$

and $\theta_u \geq 0$ with complementary slackness. The term in curly brackets on the right-hand side of (26) is the expected value of a meeting between an unemployed worker and a firm, given that the signal about the quality of the match is fully informative. The y -shock increases this expected value and, hence, it increases the vacancy-to-worker ratio θ_u .

The vacancy-to-worker ratio $\theta_e(z)$ is such that

$$k \geq p'(\theta_e(z)) \left\{ \sum_{s \geq r_e(z)} [s - z + \beta \mathbb{E}(W_e(s, \hat{y}) - W_e(z, \hat{y}))] f(s) \right\}, \quad (27)$$

and $\theta_e(z) \geq 0$ with complementary slackness. The term in curly brackets on the right-hand side of (27) is the expected value of a meeting between a firm and a worker employed in a type- z match, given that the signal about the quality of the new match is perfect. For $z \geq r_u$, the y -shock has no effect on this difference and, hence, it has no effect on the vacancy-to-worker ratio $\theta_e(z)$. Let us give some intuition for this result. The y -shock increases the value of an employed worker. However, the increase is exactly the same for all workers who are employed in matches of type $z \geq r_u$ because all of these workers have exactly the same probability of being employed in the future and, hence, the same probability of taking advantage of the increase in y . Therefore, the y -shock does not increase the return of moving a worker from a match of type $z \geq r_u$ to a better match.

Figure 6 illustrates the effect of the y -shocks on the UE, EE and EU rates. The UE rate is given by $h^{ue} = \theta_u^\gamma m_u$. The y -shock increases the UE rate because it raises the probability that an unemployed worker meets a vacancy, θ_u^γ , and the probability that such a meeting turns into a match, $m_u = \sum_{s \geq r_u} f(s)$. The EE rate is given by $h^{ee} = [\sum (1 - d(z)) \lambda_e \theta_e(z)^\gamma g(z)] / (1 - u)$. Even though the y -shock does not affect $\theta_e(z)$, it does increase the EE rate through a composition effect. Specifically, the y -shock shifts the distribution of employed workers g towards matches with lower quality, which have a higher probability of terminating with an EE transition. The EU rate is given by $h^{eu} = \sum_{z < r_d} g(z) + \delta \sum_{z \geq r_d} g(z)$. The y -shock has no effect on the EU rate because the distribution of employed workers g has no density around the destruction cutoff r_d . Let us explain why this is the case. When signals are perfectly informative, unemployed workers only move to matches with idiosyncratic productivity greater than r_u , and workers employed in type- z matches only move to matches with idiosyncratic productivity greater than z . This implies that the distribution of employed workers g has

no density to the left of r_u . Moreover, since r_u is strictly greater than r_d (as one can see by comparing (23) and (24)), the distribution of employed workers g has no density around the destruction cutoff r_d .

Figure 7 illustrates the response to the y -shock of the number of vacancies that are created for unemployed workers, v_u , the number of vacancies that are created for employed workers, v_e , and the total number of vacancies, v . Similarly, Figure 8 illustrates the response to the y -shock of the average of the idiosyncratic productivity, \bar{z} , and the average productivity, $\pi = y + \bar{z}$. For the sake of brevity, we will not comment on these figures.

6.2 Productivity shocks and labor market fluctuations

Now, we simulate the inspection model to find out what the standard deviation and the correlation between unemployment, vacancies and workers' transition rates would be if y -shocks were the only source of cyclical fluctuations in the economy. The results of the simulation are reported in Table 6. A comparison between this table and Table 1 (which summarizes the actual behavior of the US labor market) leads to two important findings. First, y -shocks generate the same pattern of comovement between unemployment, vacancies and workers' transition rates that we observe in the US labor market. Second, y -shocks are not a quantitatively important source of cyclical fluctuations in the US labor market. Indeed, y -shocks can only account for 8 percent of the empirical volatility of the unemployment rate, for 22 percent of the empirical volatility of vacancies, and for 15, 0 and 1.2 percent of the empirical volatility of the UE, EU and EE rates.

These measures of the contribution of y -shocks to the cyclical volatility of the US labor market are dramatically different from those that we obtained in section 5 using the version of the model in which matches are experience goods. It is useful to review and explain some of these differences in detail.

First, the contribution of y -shocks to the empirical volatility of the UE rate is 45 percent when measured with the experience model, and 15 percent when measured with the inspection model. This difference is surprising, yet easy to explain. In both models, the elasticity

of the UE rate with respect to aggregate productivity is given by

$$\frac{d \log h^{ue}}{d \log y} = \frac{d \log m_u}{d \log y} + \gamma \frac{d \log(v_u/u)}{d \log y}. \quad (28)$$

In the experience model, the first term on the right-hand side of (28) is equal to zero because a meeting between an unemployed worker and a firm turns into a match with probability 1, independently of y . In the inspection model, the first term can be positive and large for some parameter values. However, in our calibration, this term is approximately zero because the productivity distribution f has almost no density around the creation cutoff r_u . Hence, the first term on the right-hand side of (28) is approximately the same in the two models. In contrast, the second term is much smaller in the inspection model than in the experience model, because the calibrated value of γ is much lower. To see why this is the case, remember that we calibrated the value of γ so that the elasticity of the UE rate with respect to the vacancy-to-unemployment ratio would be the same in the model as in the data (0.27). That is, we calibrated the value of γ so that

$$\begin{aligned} d \log h^{ue} / d \log(v/u) &= 0.27 \\ \iff \gamma &= [0.27 d \log(v/u) - d \log m_u] / (d \log(v_u/u)). \end{aligned} \quad (29)$$

In the experience model, $\gamma = 0.6$ because $d \log m_u = 0$ and $d \log(v/u) / d \log(v_u/u) \sim 2.2$. In the inspection model, $\gamma = 0.25$ because $d \log m_u \sim 0$ and $d \log(v/u) / d \log(v_u/u) \sim 1$. In turn, $d \log(v/u) / d \log(v_u/u)$ is larger in the experience model because the elasticity of θ_e with respect to y is larger.¹⁰

Second, the contribution of y -shocks to the empirical volatility of the EU rate is zero when measured with the inspection model, and approximately 100 percent when it is measured with the experience model. This difference is caused by the difference in the value of α , the informativeness of the signal about the quality of a match. In both models, a shock to aggregate productivity y has an effect on the lowest realization of the idiosyncratic productivity z above which an employed worker is more valuable than an unemployed worker. Specifically, a positive (negative) shock to y leads to a decline (increase) in this cutoff realization of z . In the experience model, the idiosyncratic productivity of a match is observed

¹⁰Here is an interesting observation about the inspection model. After substituting (29) into (28), one obtains $d \log h^{ue} / d \log y = 0.27 d \log(v/u) / d \log y$. That is, after the parameter γ is calibrated, the elasticity of the UE rate with respect to y does not depend on the elasticity of m_u .

only after the match is created. Hence, movements in the lowest z above which a match is valuable translate into movements in the EU rate. In the inspection model, the idiosyncratic productivity of a match is perfectly observed before the match is created. Hence, movements in the lowest z above which a match is valuable translate into movements in the UE rate, not into movements of the EU rate.

Third, the contribution of y -shocks to the empirical volatility of the unemployment rate is less than 10 percent when measured with the inspection model, and more than 80 percent when measured with the experience model. The difference is due to the fact that, in the inspection model, y -shocks generate smaller fluctuations in both the UE and the EU rates. The two models also make different predictions about the effect of y -shocks on the EE rate, on the number of vacancies for employed workers, and on the total number of vacancies. However, for the sake of brevity, we will omit the discussion of these ulterior differences.

6.3 Model selection

According to the experience model, y -shocks alone can account for much of the volatility of the US labor market at the business cycle frequency. According to the inspection model, y -shocks cannot even account for 10 percent of the cyclical volatility of the US unemployment rate. So which one of these two models should we trust?

To answer this question, we begin by noticing that the experience model provides a much better fit to the empirical tenure distribution than the inspection model does, as one can recognize instantly by looking at Figure 9. To be more precise, the sum of the absolute deviations between the tenure distribution generated by the experience model and the tenure distribution observed in the data is 0.14. The sum of the absolute deviations between the distribution generated by the inspection model and the one observed in the data is 0.71. That is, the experience model fits the empirical tenure distribution 5 times better than the inspection model. Moreover, we notice that the two models fit equally well (perfectly) all the other calibration targets (the average of the empirical transition rates, the empirical elasticity of the UE rate with respect to the vacancy-to-worker ratio, etc.). Overall, the experience model provides an unambiguously better fit of the data and we should take its predictions with more confidence.

Let us explain why the empirical tenure distribution “prefers” the experience model to the inspection model. First, notice that from the empirical tenure distribution we can recover the empirical hazard/tenure function, i.e. the probability that a match between a firm and a worker will dissolve over the next year as a function of the age of the match. The empirical hazard/tenure function is sharply decreasing for matches with tenure between 1 and 5 years, and it displays no trend for matches with tenure longer than 5 years. More precisely, the hazard probability falls from 42 percent for matches with less than 1 year of tenure to 14 percent for matches with matches with a tenure between 4 and 5 years. For matches with tenure longer than 5 years, the hazard probability does not have a trend as it fluctuates around 14 percent per year or, equivalently, around 1.1 percent per month.

Next, recall that we calibrated the value of the exogenous destruction probability, δ , to target the average EU rate in the model to the value in the data. In the inspection model, the idiosyncratic productivity of a match is perfectly observed before the match is created. For this reason, the endogenous destruction probability is zero and the calibrated value of the exogenous destruction probability δ must be equal to the empirical average of the EU rate (i.e. 2.6 percent monthly). In the experience model, the idiosyncratic productivity of a match is only observed after the match is created. For this reason, the endogenous destruction probability need not be zero and, depending on the distribution f , the calibrated value of the exogenous destruction probability δ may take on any value between 0 and 2.6 percent.

Clearly, the hazard/tenure function generated by either model is bounded below by the exogenous match destruction probability δ . In the case of the inspection model, this implies that the hazard probability has to be greater than 2.6 percent per month at all tenure lengths. However, as we discussed above, the empirical hazard probability is on average 1.1 percent per month for matches with tenure longer than 5 years. Hence, the inspection model cannot possibly fit the empirical hazard/tenure function and, consequently, it cannot fit the empirical tenure distribution. The experience model can (and does) fit much better the empirical hazard/tenure function and the empirical tenure distribution because the exogenous match destruction probability need not be as large as 2.6 percent per month.

7 Conclusion

In this paper, we developed a model of directed search on the job in which workers move between employment, unemployment and across different employers because of heterogeneity in the idiosyncratic productivity of different matches between workers and firms. In the theoretical part of the paper, we proved that the unique equilibrium of the model is such that the agents' value and policy functions depend on the aggregate state of the economy only through the realization of the exogenous shocks and not through the distribution of workers across employment states (unemployment and employment in different matches). This property of the equilibrium (which we refer to as block recursivity) enables us to solve the model outside of steady-state as easily as one would solve the equilibrium of a model in which all workers were identical. In the quantitative part of the paper, we first calibrated the model using data on the acyclical properties of the US labor market. Then we simulated the model to measure the contribution of aggregate productivity shocks to the cyclical volatility of the US labor market. We found that, if matches are experience goods, aggregate productivity shocks generate fluctuations in unemployment, vacancies and workers' transition rates that are quite similar to those observed in the US economy. In contrast, if matches are inspection goods, these shocks contribute very little to the cyclical volatility of the labor market. Finally, we showed that the version of the model in which matches are experience goods provides an unambiguously better description of the acyclical features of the US labor market and, hence, we should take its predictions about the effect of aggregate productivity shocks more seriously.

One of the main implications of our model is that recessions are cleansing. That is, during recessions, workers and firms need a higher idiosyncratic productivity to be willing to produce rather than remain idle and, for this reason, the average idiosyncratic productivity of existing matches increases. In the early nineties, the cleansing view of recession became popular as it provided a natural explanation for the strongly countercyclical behavior of the number of jobs destroyed at the establishment level (a behavior best documented by Davis and Haltiwanger 1992). A few years later, though, the popularity of the cleansing view faded as researchers found evidence that appeared to be at odds with the notion that the quality of matches is higher in downturns. For example, Bowlus (1995) found a negative correlation

between the median duration of a firm-worker match and the unemployment rate at the time the match was formed. Moreover, she found a negative correlation between the starting wage paid by a firm to a worker and the unemployment rate at the time the match between the firm and the worker was formed. This evidence led Barley (2002) to dismiss the cleansing view of recessions in favor of the opposite, sullying view. By the same token, should we dismiss the theory of business cycles advanced in our paper?

Fortunately, the version of our model in which matches are experience goods (our preferred version) is consistent with the evidence uncovered by Bowlus (1995). First, in the experience version of our model, there is a strong negative correlation between the median duration of a match and the unemployment rate at the time the match is formed. Intuitively, a negative shock to aggregate productivity tends to increase the unemployment rate and, at the same time, to reduce the probability that a match survives after its idiosyncratic productivity is revealed. In turn, the latter effect tends to lower the median duration of a match. Second, in our model, it is not difficult to identify a wage setting rule such that the correlation between wages and unemployment is negative because, in response to a negative shock to aggregate productivity, workers find it optimal to look for vacancies that offer them a lower lifetime utility.

Appendix

A Proof of Theorem 1

(i) Let $C(\Psi)$ be the set of bounded continuous functions $R : \Psi \rightarrow \mathbb{R}$ with the sup norm, $\|R\| = \sup_{\psi \in \Psi} R(\psi)$. Define the operator T on $C(\Psi)$ by

$$\begin{aligned}
 (TR)(\psi) &= \max_{(d, \theta_u, \theta_e, c_u, c_e)} F(d, \theta_u, \theta_e, c_u, c_e | \psi) + \beta \mathbb{E}R(\hat{\psi}) \\
 \text{s.t. } \hat{u} &= u [1 - \lambda_u p(\theta_u) m_u] + \sum_z [d(z)g(z)], \\
 \hat{g}(z') &= u \lambda_u p(\theta_u) [\alpha c_u(z') + (1 - \alpha) m_u] f(z') \\
 &\quad + g(z') [1 - d(z')] [1 - \lambda_e p(\theta_e(z')) m_e(z')] \\
 &\quad + \sum_z g(z) \{ [1 - d(z)] [\lambda_e p(\theta_e(z))] [\alpha c_e(z', z) + (1 - \alpha) m_e(z)] f(z') \} \\
 d : Z &\rightarrow [\delta, 1], \theta_u \in [0, \bar{\theta}], \theta_e : Z \rightarrow [0, \bar{\theta}], c_u : Z \rightarrow [0, 1], c_e : Z \times Z \rightarrow [0, 1].
 \end{aligned} \tag{A1}$$

The return function F is defined as

$$F(d, \theta_u, \theta_e, c_u, c_e | \psi) = -k \{ \lambda_u \theta_u u + \sum_z [(1 - d(z)) \lambda_e \theta_e(z) g(z)] \} + b \hat{u} + \sum_z [(y + z) \hat{g}(z)]. \tag{A2}$$

For each $R \in C(\Psi)$ and $\psi \in \Psi$, the maximand in (A1) is continuous in $(d, \theta_u, \theta_e, c_u, c_e)$ and the set of feasible choices for $(d, \theta_u, \theta_e, c_u, c_e)$ is compact. Hence, the maximum is attained. Since the maximand is bounded, TR is bounded; and since the maximand is continuous, it follows from the Theorem of the Maximum (Theorem 3.6. in Stokey, Lucas and Prescott, 1989) that TR is continuous. Hence, $T : C(\Psi) \rightarrow C(\Psi)$.

The operator T maps the set of bounded continuous function $C(\Psi)$ into itself, and one can easily verify that it satisfies the monotonicity and discounting hypotheses in Blackwell's sufficient conditions for a contraction (Theorem 3.3 in Stokey, Lucas and Prescott, 1989). Hence, the operator T is a contraction mapping and it admits one and only one fixed point $R^* \in C(\Psi)$. Since $\lim_{t \rightarrow \infty} \beta^t R^*(\psi) = 0$ for all $\psi \in \Psi$, it follows from Theorem 4.3 in Stokey, Lucas and Prescott (1989) that R^* is equal to the planner's value function W .

(ii) Let $C'(\Psi) \subset C(\Psi)$ be the set of functions $R : \Psi \rightarrow \mathbb{R}$ that are bounded, continuous and linear in the measure of unemployed workers, u , and in the measure of workers employed in matches of type z , $g(z)$. Clearly, $R \in C'(\Psi)$ if and only if there exist two functions

$R_u : Y \rightarrow \mathbb{R}$ and $R_e : Z \times Y \rightarrow \mathbb{R}$ such that

$$R(\psi) = R_u(y)u + \sum_z R_e(z, y)g(z). \quad (\text{A3})$$

Consider an arbitrary function R in $C'(\Psi)$. Then, after substituting the constraints into the maximand of (A1), we obtain

$$(TR)(\psi) = \hat{R}_u(y)u + \sum_z \hat{R}_e(z, y)g(z), \quad (\text{A4})$$

where $\hat{R}_u(y)$ is given by

$$\begin{aligned} \hat{R}_u(y) = \max_{(\theta_u, c_u)} \{ & -k\lambda_u\theta_u + (1 - \lambda_u p(\theta_u)m_u) [b + \beta\mathbb{E}R_u(\hat{y})] \\ & + \lambda_u p(\theta_u) \sum_{z'} [\alpha c_u(z') + (1 - \alpha)m_u] [y + z' + \beta\mathbb{E}R_e(z', \hat{y})] f(z') \} \\ \text{s.t. } & \theta_u \in [0, \bar{\theta}], \quad c_u : Z \rightarrow [0, 1]. \end{aligned} \quad (\text{A5})$$

and $\hat{R}_e(z, y)$ is given by

$$\begin{aligned} \hat{R}_e(z, y) = \max_{(d, \theta_e, c_e)} \{ & d [b + \beta\mathbb{E}R_u(\hat{y})] - (1 - d) k\lambda_e\theta_e \\ & + (1 - d) (1 - \lambda_e p(\theta_e)m_e) [y + z + \beta\mathbb{E}R_e(z, \hat{y})] \\ & + (1 - d) \lambda_e p(\theta_e) \sum_{z'} [\alpha c_e(z') + (1 - \alpha)m_e] [y + z' + \beta\mathbb{E}R_e(z', \hat{y})] f(z') \} \\ \text{s.t. } & d \in [\delta, 1], \quad \theta_e \in [0, \bar{\theta}], \quad c_e : Z \rightarrow [0, 1]. \end{aligned} \quad (\text{A6})$$

Since R is an arbitrary function in $C'(\Psi)$, (A4) implies that $T : C'(\Psi) \rightarrow C'(\Psi)$. Moreover, since $C'(\Psi)$ is a closed subset of $C(\Psi)$ and $T : C'(\Psi) \rightarrow C'(\Psi)$, Corollary 1 to Theorem 3.2 in Stokey, Lucas and Prescott (1989) implies $W \in C'(\Psi)$.

(iii) Let $C''(\Psi) \subset C'(\Psi)$ be the set of functions $R : \Psi \rightarrow \mathbb{R}$ such that the associated component R_e is nondecreasing in z . Let R be an arbitrary function in $C''(\Psi)$. From part (ii), it follows that $TR \in C'(\Psi)$ and the associated components \hat{R}_u and \hat{R}_e satisfy the equations (A5) and (A6). Since the maximand in (A6) is nondecreasing in z and the feasible set in (A6) is independent of z , \hat{R}_e is nondecreasing in z . Hence, $T : C''(\Psi) \rightarrow C''(\Psi)$. Since $C''(\Psi)$ is a closed subset of $C(\Psi)$, Corollary 1 to Theorem 3.2 in Stokey, Lucas and Prescott (1989) implies that $W \in C''(\Psi)$.

(iv) From part (ii), it follows that the policy correspondences (θ_u^*, c_u^*) solve the maximization problem (A5) for $(R_u, R_e) = (W_u, W_e)$. Since the maximand and the constraints in (A5) do

not depend on (u, g) , (θ_u^*, c_u^*) depend on ψ only through y and not through (u, g) . Similarly, the policy correspondences (d^*, θ_e^*, c_e^*) solve the maximization problem (A6) for $(R_u, R_e) = (W_u, W_e)$. Since the maximand and the constraints in (A6) do not depend on (u, g) , (θ_u^*, c_u^*) depend on ψ only through y and not through (u, g) . ■

B Proof of Proposition 2

For any $y \in Y$, $\theta_u^*(y)$ and $c_u^*(z, y)$ are the solutions to the maximization problem

$$\begin{aligned} & \max_{(\theta_u, c_u)} \{ -k\lambda_u\theta_u + (1 - \lambda_u p(\theta_u)m_u) [b + \beta\mathbb{E}W_u(\hat{y})] \\ & \quad + \lambda_u p(\theta_u) \sum_{z'} [\alpha c_u(z') + (1 - \alpha)m_u] [y + z' + \beta\mathbb{E}W_e(z', \hat{y})] f(z') \} \\ & \text{s.t. } \theta_u \in [0, \bar{\theta}], \quad c_u : Z \rightarrow [0, 1]. \end{aligned} \quad (\text{A7})$$

which can be rewritten as

$$\begin{aligned} & \max_{\theta_u \in [0, \bar{\theta}]} \{ -k\lambda_u\theta_u + b + \beta\mathbb{E}W_u(\hat{y}) + \lambda_u p(\theta_u) \times \\ & \quad \times \max_{c_u : Z \rightarrow [0, 1]} \sum_{z'} \{ [\alpha c_u(z') + (1 - \alpha)m_u] [y + z' - b + \beta\mathbb{E}(W_e(z', \hat{y}) - W_u(\hat{y}))] f(z') \} \}. \end{aligned} \quad (\text{A8})$$

First, consider the inner maximization problem in (A8). The maximand is linear in c_u and its derivative with respect to $c_u(s)$ is given by

$$\alpha [y + s + \beta\mathbb{E}W_e(s, \hat{y})] + (1 - \alpha)\mathbb{E}_{z'} [y + z' + \beta\mathbb{E}W_e(z', \hat{y})] - b - \beta\mathbb{E}W_u(\hat{y}). \quad (\text{A9})$$

Hence, the solution to the maximization problem is $c_u^*(s, y) = 1$ if (A9) is positive, and $c_u^*(s, y) = 0$ if (A9) is strictly negative. Therefore, $c_u^*(s, y)$ is unique. Moreover, since (A9) is increasing in s , $c_u^*(s, y)$ is increasing in s . Therefore, there exists $r_u^*(y)$ such that $c_u^*(s, y) = 1$ if $s \geq r_u^*(y)$, and $c_u^*(s, y) = 0$ else. This completes the proof of parts (i) and (iii) of the proposition for c_u^* .

Next, consider the outer maximization problem in (A8). The derivative of the maximand with respect to θ_u is given by

$$-k + p'(\theta_u) \sum_{s \geq r_u^*(y)} \left\{ \begin{aligned} & \alpha [y + s - b + \beta\mathbb{E}(W_e(s, \hat{y}) - W_u(\hat{y}))] \\ & + (1 - \alpha)\mathbb{E}_{z'} [y + z' - b + \beta\mathbb{E}(W_e(z', \hat{y}) - W_u(\hat{y}))] \end{aligned} \right\} f(s). \quad (\text{A10})$$

The expression above is strictly decreasing in θ_u because $p''(\theta_u) < 0$, and it is strictly negative

at $\theta_u = \bar{\theta}$ because $p'(\bar{\theta}) = 0$. Hence, the solution to the maximization problem, $\theta_u^*(y)$, is unique and such that

$$k \geq p'(\theta_u^*(y)) \sum_{s \geq r_u^*(y)} \left\{ \begin{array}{l} \alpha [y + s - b + \beta \mathbb{E}(W_e(s, \hat{y}) - W_u(\hat{y}))] \\ + (1 - \alpha) \mathbb{E}_{z'} [y + z' - b + \beta \mathbb{E}(W_e(z', \hat{y}) - W_u(\hat{y}))] \end{array} \right\} f(s), \quad (\text{A11})$$

and $\theta_u^*(y) \geq 0$, with complementary slackness. This completes the proof of part (i) of the proposition for θ_u^* . The proofs of parts (i) and (iii) for c_e^* , θ_e^* and d^* , as well as the proofs of parts (ii) and (iv) are omitted for the sake of brevity. ■

C Proof of Theorem 4

(i)-(ii) Let $(\theta, V_u, V_e, x_u, r_u, d, x_e, r_e)$ be an equilibrium. We take five steps to prove that the equilibrium is unique and block recursive.

Step 1. Unify the notation for V_u and V_e . Let the function $V : \{0, 1\} \times Z \times \Psi \rightarrow \mathbb{R}$ be defined as $V(0, z, y) = V_u(\psi)$ for all $(z, \psi) \in Z \times \Psi$, and $V(1, z, y) = V_e(z, \psi)$ for all $(z, \psi) \in Z \times \Psi$. Given the definition of V , we can rewrite the equilibrium conditions (12) and (14) as

$$\begin{aligned} & V(a, z, \psi) \\ = & a \left\{ y + z + \beta \mathbb{E} \max_{(d, x, r)} \left\{ \begin{array}{l} dV(0, z, \hat{\psi}) + (1 - d)V(1, z, \hat{\psi}) \\ + (1 - d)\lambda_e p(\theta(x, r, \psi)) m(r) [x - V(1, z, \hat{\psi})] \end{array} \right\} \right\} \\ + & (1 - a) \left\{ b + \beta \mathbb{E} \max_{(x, r)} \left\{ \begin{array}{l} V(0, z, \hat{\psi}) + \\ \lambda_u p(\theta(x, r, \hat{\psi})) m(r) [x - V(0, z, \hat{\psi})] \end{array} \right\} \right\}. \end{aligned} \quad (\text{A12})$$

Step 2. Express the value offered in submarket x as a function of the tightness θ , the reservation signal r , and the aggregate state of the economy ψ . Let $x(\theta, r, \psi)$ denote the value offered to a worker in a submarket with tightness $\theta(x, r, \psi) = \theta > 0$. From the equilibrium condition (16), it follows that

$$x(\theta, r, \psi) = \frac{1}{m(r)} \left\{ \sum_{s \geq r} \{ [\alpha V(1, s, \psi) + (1 - \alpha) \mathbb{E}_z V(1, z, \psi)] f(s) \} - \frac{k}{q(\theta)} \right\}. \quad (\text{A13})$$

In any submarket with $\theta(x, r, \psi) = 0$, the value offered to a worker cannot be expressed uniquely as a function of (θ, r, ψ) . However, the value offered to a worker in these submarkets

is irrelevant because the worker meets a vacancy with zero probability. Hence, without loss in generality, let $x(\theta, r, \psi) = 0$ in all submarkets with tightness $\theta(x, r, \psi) = \theta = 0$.

Step 3. Reformulate the equilibrium condition for V . Substituting x with $x(\theta, r, \psi)$ and $\theta(x, r, \psi)$ with θ , we can rewrite (A12) as

$$\begin{aligned}
& V(a, z, \psi) \\
= & a \left\{ y + z + \beta \mathbb{E} \max_{(d, \theta, r)} \left\{ \begin{aligned} & dV(0, z, \hat{\psi}) - (1-d)\lambda_e k \theta + (1-d)(1-\lambda_e p(\theta)m(r))V(1, z, \hat{\psi}) \\ & + (1-d)\lambda_e p(\theta) \sum_{s \geq r} \left[\alpha V(1, s, \hat{\psi}) + (1-\alpha) \mathbb{E}_z V(1, z, \hat{\psi}) \right] f(s) \end{aligned} \right\} \right\} \\
+ & (1-a) \left\{ b + \beta \mathbb{E} \max_{(\theta, r)} \left\{ \begin{aligned} & -k\lambda_u \theta + (1-\lambda_u p(\theta)m(r))V(0, z, \hat{\psi}) + \\ & \lambda_u p(\theta) \sum_{s \geq r} \left[\alpha V(1, s, \hat{\psi}) + (1-\alpha) \mathbb{E}_z V(1, z, \hat{\psi}) \right] f(s) \end{aligned} \right\} \right\}. \tag{A14}
\end{aligned}$$

Step 4. Establish the uniqueness of V and its independence from (u, g) . Let $\Omega = \{0, 1\} \times Z \times \Psi$ and let $C(\Omega)$ denote the space of bounded continuous functions $R : \Omega \rightarrow \mathbb{R}$, with the sup norm. Let $T : C(\Omega) \rightarrow C(\Omega)$ denote the operator associated with (A14). It is straightforward to verify that: (i) $R, R' \in C(\Omega)$ and $R \leq R'$ implies $T(R) \leq T(R')$; (ii) $R \in C(\Omega)$ and $\epsilon \geq 0$ implies $T(R + \epsilon) = TR + \beta\epsilon$. Therefore, by Blackwell's sufficient conditions, it follows that the operator T is a contraction and that it admits a unique solution. Hence, V is unique. Next, notice that if R depends on $\hat{\psi}$ only through \hat{y} , then $T(R)$ depends on ψ only through y . Hence, the fixed point of the operator T depends on ψ only through y . That is, $V(a, y, \psi) = V(a, z, y)$.

Step 5. Establish the uniqueness of the policy functions $(\theta, x_u, r_u, d, x_e, r_e)$ and their independence from (u, g) . Since $V(a, z, \psi)$ only depends on ψ through y , we can rewrite the equilibrium condition (16) as

$$k \geq q(\theta(x, r, \psi)) \sum_{s \geq r} \{ [\alpha V_e(s, y) + (1-\alpha) \mathbb{E}_z V_e(z, y) - x] f(s) \} \tag{A15}$$

and $\theta(x, r, \psi) \geq 0$, with complementary slackness. It is easy to verify that $\theta(x, r, \psi)$ is unique and only depends on ψ through y ; that is, $\theta(x, r, \psi) = \theta(x, r, y)$. Since $V(a, z, \psi)$ and $\theta(x, r, y)$ only depend on ψ through y , we can rewrite the equilibrium condition (12) as

$$V_u(y) = b + \beta \mathbb{E} \max_{(x, r)} \{ V_u(\hat{y}) + \lambda_u p(\theta(x, r, \hat{y})) m(r) [x - V_u(\hat{y})] \}. \tag{A16}$$

Since the maximization problem in (A16) only depends on $\hat{\psi}$ through \hat{y} , the associated policy functions $(x_u(\hat{\psi}), r_u(\hat{\psi}))$ only depend on $\hat{\psi}$ through \hat{y} . That is $(x_u(\hat{\psi}), r_u(\hat{\psi})) = (x_u(\hat{y}), r_u(\hat{y}))$. Similarly, we can show that the policy functions $d(\hat{\psi})$ and $(x_e(\hat{\psi}), r_e(\hat{\psi}))$ only depend on $\hat{\psi}$ through \hat{y} . That is, $d(\hat{\psi}) = d(\hat{y})$ and $(x_e(\hat{\psi}), r_e(\hat{\psi})) = (x_e(\hat{y}), r_e(\hat{y}))$. This completes the proof that there exists a unique equilibrium and that this equilibrium is block recursive.

(iii) To establish the equivalence between the equilibrium and the planner's allocation, we rewrite the component value functions (5) and (6). Recall that, in the planner's allocation, a match is formed if and only if the signal s is greater than or equal to the cutoff level $r_u^*(y)$ for an unemployed worker and $r_e^*(z, y)$ for a worker employed in a type- z match (see Proposition 2). Using r_u and r_e as the choices instead of (c_u, c_e) , we can rewrite (5) as

$$W_u(y) = \max_{(\theta_u, r_u)} \left\{ -k\lambda_u\theta_u + (1 - \lambda_u p(\theta_u)m_u) [b + \beta\mathbb{E}W_u(\hat{y})] \right. \\ \left. + \lambda_u p(\theta_u) \sum_{s \geq r_u} \{ [\alpha W_e(s, y) + (1 - \alpha)\mathbb{E}_{z'}(y + z' + \beta\mathbb{E}W_e(z', \hat{y}))] f(s) \} \right\}. \quad (\text{A17})$$

Similarly, we can rewrite (6) as

$$W_e(z, y) \\ = \max_{(d, \theta_e, r_e)} \left\{ d [b + \beta\mathbb{E}W_u(\hat{y})] - (1 - d) k\lambda_e\theta_e \right. \\ \left. + (1 - d) (1 - \lambda_e p(\theta_e)m_e) [y + z + \beta\mathbb{E}W_e(z, \hat{y})] \right. \\ \left. + (1 - d)\lambda_e p(\theta_e) \sum_{s \geq r_e} \{ [\alpha(y + s + \beta\mathbb{E}r_e(s, \hat{y})) + (1 - \alpha)\mathbb{E}_z(y + z + \beta\mathbb{E}r_e(z, \hat{y}))] f(s) \} \right\}. \quad (\text{A18})$$

Using these equations, we can verify that (A14) is satisfied by the function $W'(a, z, y)$ defined as $W'(0, z, y) = b + \beta\mathbb{E}W_u(\hat{y})$ and $W'(1, z, y) = y + z + \beta\mathbb{E}W_e(z, \hat{y})$. Since V is the unique solution to (A14), it follows that $V_u(y) = b + \beta\mathbb{E}W_u(\hat{y})$ and $V_e(z, y) = y + z + \beta\mathbb{E}W_e(z, \hat{y})$. Finally, notice that the equilibrium allocation solves the maximization problems in (A14), while the efficient allocation solves the maximization problems in (A17) and (A18). With the relations $V_u(y) = b + \beta\mathbb{E}W_u(\hat{y})$ and $V_e(z, y) = y + z + \beta\mathbb{E}W_e(z, \hat{y})$, it is not difficult to see that the two sets of allocations coincide. \blacksquare

References

- [1] Acemoglu, D., and R. Shimer. 1999. "Efficient Unemployment Insurance." *J.P.E.* 107: 893-927.
- [2] Barlevy, G. 2002. "The Sullyng Effect of Recessions." *Rev. Econ. Studies* 69: 65–96.
- [3] Bowlus, A. 1995. "Matching Workers and Jobs: Cyclical Fluctuations in Match Quality." *J. Labor Econ.* 13: 335–50.
- [4] Burdett, K., S. Shi, and R. Wright. 2001. "Pricing and Matching with Frictions." *J.P.E.* 109: 1060–85.
- [5] Cooley, F. 1995. *Frontiers of Business Cycle Research*. Princeton University Press, Princeton, NJ.
- [6] Davis, S., and J. Haltiwanger. 1992. "Gross Job Creation, Gross Job Destruction, and Employment Reallocation." *Q.J.E.* 107: 819–63.
- [7] Delacroix, A. and S. Shi. 2006. "Directed Search On the Job and the Wage Ladder," *Int. Econ. Review* 47: 651-699.
- [8] Diebold, F., D. Neumark, and D. Polsky. 1997. "Job Stability in the United States." *J. Labor Econ.* 15: 206–33.
- [9] Elsby, M., R. Michaels, and G. Solon. 2009. "The Ins and Outs of Cyclical Unemployment." *A.E.J. Macroeconomics* 1: 84–110.
- [10] Gertler, M., and A. Trigari. 2009. "Unemployment Fluctuations with Staggered Nash Wage Bargaining." *J.P.E.* 117: 38–86.
- [11] Gomes, J., J. Greenwood, and S. Rebelo. 2001. "Equilibrium Unemployment." *J.M.E.* 48: 109–52.
- [12] Gonzalez, F.M., and S. Shi. 2010. "An Equilibrium Theory of Learning, Search and Wages," *Econometrica* 78: 509-537.
- [13] Hagedorn, M., and I. Manovskii. 2008. "The Cyclical Behavior of Unemployment and Vacancies Revisited." *A.E.R.* 98: 1692–706.
- [14] Hall, R. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *A.E.R.* 95: 50–65.
- [15] Hall, R., and P. Milgrom. 2008. "The Limited Influence of Unemployment of the Wage Bargain." *A.E.R.* 98: 1653–74.
- [16] Kenman, J. 2008. "Private Information, Wage Bargaining and Employment Fluctuations." Manuscript. Univ. of Wisconsin at Madison.
- [17] Menzio, G. 2005. "High Frequency Wage Rigidity." Manuscript. Univ. Pennsylvania.
- [18] Menzio, G. 2007. "A Theory of Partially Directed Search." *J.P.E.* 115: 748–69.

- [19] Menzio, G. and E. Moen. 2008. “Worker Replacement.” *PIER Working Paper* 08-040.
- [20] Menzio, G. and S. Shi. 2009a. “Efficient Search on the Job and the Business Cycle.” *NBER Working Paper* 14905.
- [21] Menzio, G. and S. Shi. 2009b. “Block Recursive Equilibria for Stochastic Models of Search on the Job.” *J. Econ. Theory*.
- [22] Menzio, G. and S. Shi. 2010. “Directed Search on the Job, Heterogeneity, and Aggregate Fluctuations.” *A.E.R. Papers and Proc.*
- [23] Merz, M. 1995. Search in the Labor Market and the Real Business Cycle. *J.M.E.* 36: 269–300.
- [24] Moen, E. 1997. “Competitive Search Equilibrium.” *J.P.E.* 105: 694–723.
- [25] Mortensen, D. 1982. “Property Rights and Efficiency in Mating, Racing and Related Games.” *A.E.R.* 72: 968–79.
- [26] Mortensen, D. 1994. “The Cyclical Behavior of Job and Worker Flows.” *J. Econ. Dynamics and Control* 18: 1121–42.
- [27] Mortensen, D., and E. Nagypál. 2007. “More on Unemployment and Vacancy Fluctuations.” *Rev. Econ. Dynamics* 10: 327–47.
- [28] Mortensen, D., and C. Pissarides. 1994. “Job Creation and Job Destruction in the Theory of Unemployment.” *Rev. Econ. Studies* 61: 397–415.
- [29] Moscarini, G. 2003. “Skill and Luck in the Theory of Turnover.” Manuscript. Yale Univ.
- [30] Moscarini, G., and F. Postel-Vinay. 2009. “Non Stationary Search Equilibrium.” Manuscript. Yale Univ.
- [31] Nagypál, E. 2007. “Labor-Market Fluctuations and On-the-Job Search.” Manuscript, Northwestern Univ.
- [32] Nagypál, E. 2008. “Worker Reallocation Over the Business Cycle: The Importance of Employer-to-Employer Transitions.” Manuscript, Northwestern Univ.
- [33] Pissarides, C. 1985. “Short-Run Equilibrium Dynamics of Unemployment, Vacancies and Real Wages.” *A.E.R.* 75: 676–90.
- [34] Pissarides, C. 1994. “Search Unemployment with On-the-Job Search.” *Rev. Econ. Studies* 61: 457–75.
- [35] Pissarides, C. 2000. *Equilibrium Unemployment Theory*. MIT University Press, Cambridge, MA.
- [36] Pries, M., and R. Rogerson. 2005. “Hiring Policies, Labor Market Institutions, and Labor Market Flows.” *J.P.E.* 113: 811–39.
- [37] Ramey, G. 2008. “Exogenous vs. Endogenous Separation.” Manuscript, U.C. San Diego.

- [38] Shao, E., and P. Silos. 2009. “Firm Entry and Labor Market Dynamics.” Manuscript. F.R.B. Atlanta.
- [39] Shi, S. 2001. “Frictional Assignment I: Efficiency,” *Journal of Economic Theory* 98, 232-260.
- [40] Shi, S. 2009. “Directed Search for Equilibrium Wage-Tenure Contracts.” *Econometrica*, 77: 561-584.
- [41] Shimer, R. 1996. “Contracts in a Frictional Labor Market.” Manuscript M.I.T.
- [42] Shimer, R. 2005. “The Cyclical Behavior of Unemployment and Vacancies.” *A.E.R.* 95: 25–49.
- [43] Stokey, N., R. Lucas, and E. Prescott. 1989. *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, MA.
- [44] Robin. J. 2009. “Labor Market Dynamics with Sequential Auctions.” Manuscript, U.C.L.

TABLE 1: SUMMARY STATISTICS, QUARTERLY U.S. DATA

x	u	v	h^{ue}	h^{eu}	h^{ee}	π	
$std(x)/std(\pi)$	9.56	10.9	5.96	5.48	5.98	1	
$autocorr(x)$.872	.909	.822	.698	.597	.760	
$corr(\cdot, x)$	u	1	-.902	-.916	.778	-.634	-.283
	v	—	1	.902	-.778	.607	.423
	h^{ue}	—	—	1	-.677	.669	.299
	h^{eu}	—	—	—	1	-.301	-.528
	h^{ee}	—	—	—	—	1	.208
	π	—	—	—	—	—	1

Notes: The seasonally adjusted unemployment rate, u , is constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey (CPS). The seasonally adjusted help wanted advertising index, v , is constructed by the Conference Board. The UE and EU rates, h^{ue} and h^{eu} , are constructed from the seasonally adjusted unemployment rate and the short-term unemployment rate as explained in section 5. The EE rate, h^{ee} , is constructed by Nagypal (2007) from the CPS microdata as explained in section 5. The variables u , v , h^{ue} , h^{eu} and h^{ee} are quarterly averages of monthly series. Average labor productivity, π , is seasonally adjusted real average output per worker in the non-farm business sector constructed by the BLS. The series for u , v , h^{ue} , h^{eu} and π cover the period 1951(I)-2006(II). The series for h^{ee} covers the period 1994(I)-2006(II). The standard deviation of h^{ee} is expressed relative to the standard deviation of π over the period 1994(I)-2006(II), and the correlation of h^{ee} with u , v , h^{ue} , h^{eu} and π refers to the period 1994(I)-2006(II). All variables are reported in logs as deviations from an HP trend with smoothing parameter 1600.

TABLE 2: CALIBRATION OUTCOMES

	Description	EXP	INS	P85	MP94
β	discount factor	.996	.996	.996	.996
b	home productivity	.907	.716	.710	.739
λ_u	off the job search	1	1	1	1
λ_e	on the job search	.735	.904	0	0
γ	elasticity of p wrt θ	.600	.250	.270	.270
k	vacancy cost	1.55	2.37	1.85	1.89
δ	exogenous destruction	.012	.026	.026	.012
μ_z	average idiosyncratic prod.	0	0	0	0
σ_z	scale idiosyncratic. prod.	.952	.008	0	.467
α_z	shape idiosyncratic prod.	4	10	—	10

Notes: Calibrated parameters for different versions of the model. The column EXP refers to the version of the model in which matches are experience goods. The column INS refers to the version of the model in which matches are inspection goods. The column P85 refers to a version of the experience model in which the parameters λ_e and σ_z are constrained to be equal to zero. The column MP94 refers to a version of the experience model in which the parameter λ_e is constrained to be equal to zero.

TABLE 3: EXPERIENCE MODEL

x	u	v	v_u	v_e	h^{ue}	h^{eu}	h^{ee}	π
$std(x)/std(\pi)$	7.88	2.54	4.29	8.21	2.51	6.23	5.59	1
$autocorr(x)$.850	.637	.748	.824	.799	.772	.823	.762
$corr(\cdot, x)$	u	1	-.807	.841	-.980	-.976	.972	-.979
	v	—	1	-.380	.855	.897	-.898	.858
	π	—	—	-.729	.984	.999	-.979	.983

Notes: Summary statistics of the last 6,000 month of a 9,000 month long time series for u , v , v_u , v_e , h^{ue} , h^{eu} , h^{ee} , and π generated by the experience model with aggregate productivity shocks. Section 4 provides details on the stochastic process for productivity. All variables are quarterly averages of monthly series. All variables are reported in logs as deviations from an HP trend with smoothing parameter 1600.

TABLE 4: P85 MODEL

x	u	$v = v_u$	h^{ue}	h^{eu}	π	
$std(x)/std(\pi)$	0.82	2.69	0.91	0	1	
$autocorr(x)$.815	.677	.994	1	.745	
$corr(\cdot, x)$	u	1	-.932	-.936	0	-.972
	v	—	1	.990	0	.990
	π	—	—	.999	0	1

Notes: Summary statistics of the last 6,000 month of a 9,000 month long time series for u , v , v_u , h^{ue} , h^{eu} and π generated by a version of the experience model in which the parameters λ_e and σ_z are constrained to be equal to zero. All variables are quarterly averages of monthly series. All variables are reported in logs as deviations from an HP trend with smoothing parameter 1600.

TABLE 5: MP94 MODEL


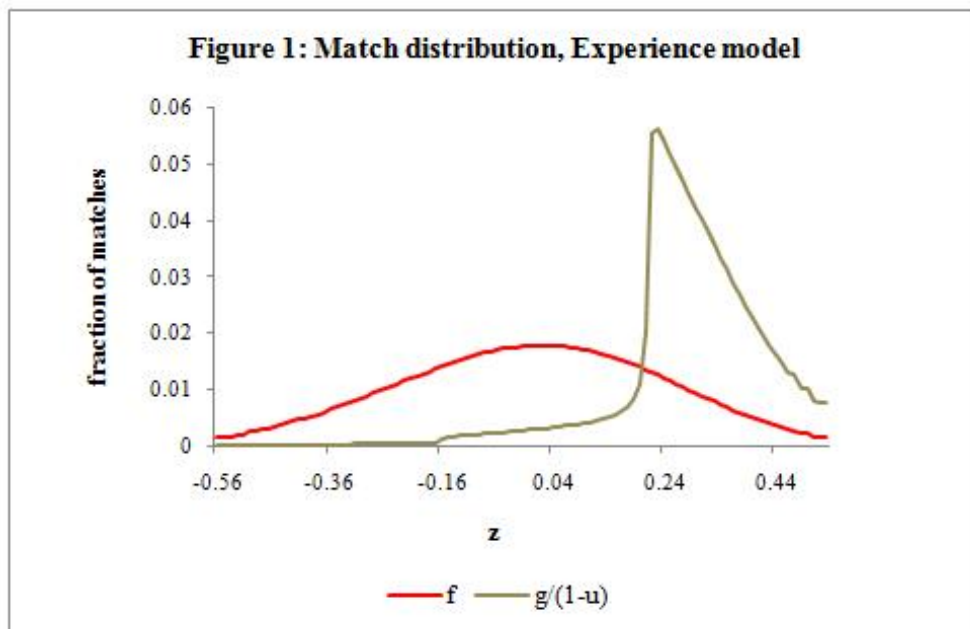
x	u	$v = v_u$	h^{ue}	h^{eu}	π	
$std(x)/std(\pi)$	5.98	4.55	0.83	6.61	1	
$autocorr(x)$.674	.453	.740	.397	.736	
$corr(\cdot, x)$	u	1	.726	-.737	.906	-.732
	v	—	1	-.267	.481	-.259
	π	—	—	.998	-.583	1

Notes: Summary statistics of the last 6,000 month of a 9,000 month long time series for u , v , v_u , h^{ue} , h^{eu} and π generated by a version of the experience model in which the parameter λ_e is constrained to be equal to zero. All variables are quarterly averages of monthly series. All variables are reported in logs as deviations from an HP trend with smoothing parameter 1600.

TABLE 6: INSPECTION MODEL

x	u	v	v_u	v_e	h^{ue}	h^{eu}	h^{ee}	π	
$std(x)/std(\pi)$	0.75	2.40	2.74	0.10	0.84	0	0.06	1	
$autocorr(x)$.829	.686	.684	.753	.749	1	.743	.750	
$corr(\cdot, x)$	u	1	-.935	-.934	-.925	-.971	0	-.817	-.971
	v	—	1	.999	.878	.992	0	.824	.992
	π	—	—	.992	.907	.999	0	.833	1

Notes: Summary statistics of the last 6,000 month of a 9,000 month long time series for u , v , v_u , v_e , h^{ue} , h^{eu} , h^{ee} , and π generated by the inspection model with aggregate productivity shocks. Section 4 provides details on the stochastic process for productivity. All variables are quarterly averages of monthly series. All variables are reported in logs as deviations from an HP trend with smoothing parameter 1600.



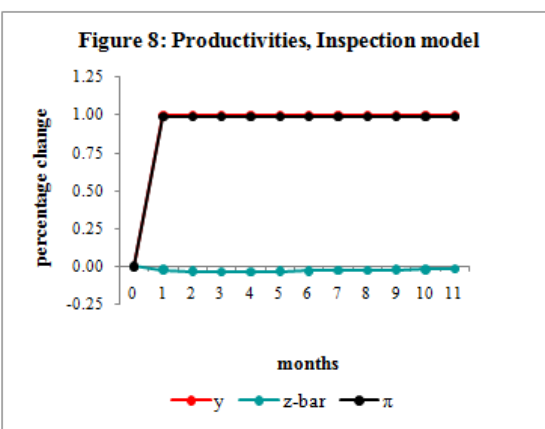
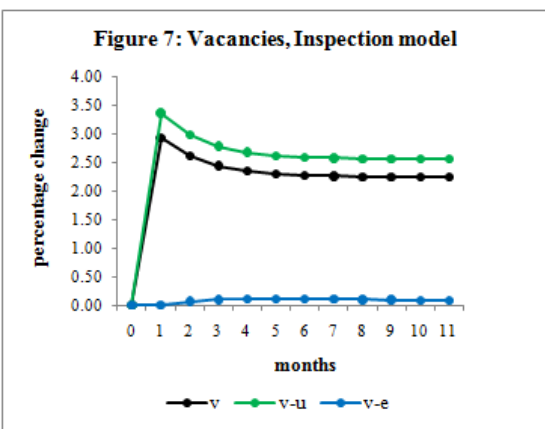






Figure 9: Tenure distribution

