Efficient Learning and Job Turnover in the Labor Market*

Fei Li[†] and Xi Weng[‡]

February 2, 2015

Abstract

This paper nests a continuous-time learning model *à la* Jovanovic (1984) into a directed on-the-job search framework. We prove that the socially efficient allocation is separable, i.e., the workers' value functions and optimal controls are independent of both the distribution of workers across their current match qualities and the unemployment rate. We characterize the dynamics of job transitions in the efficient allocation. Furthermore, when the matching technology is linear, our numerical results show that increasing the vacancy creation cost and the speed of learning have ambiguous effects on the unemployment rate and aggregate job transition.

Keywords: Learning, Directed On-the-Job Search, Brownian Motion, Separability

JEL Classification Codes: D83, J31

*We thank the associate editor, Guido Menzio, whose comments have significantly improved the quality of the paper. We also thank Naoki Aizawa, Benjamin Lester, Kenneth Burdett, Jan Eeckhout, Hanming Fang, Chao Fu, Manolis Galenianos, Nils Gornemann, Kyungmin Kim, Ricardo Lagos, George Mailath, Giuseppe Moscarini, Peter Norman, Andrew Postlewaite, Moritz Ritter, Robert Shimer, Can Tian, Gabor Virag, and Pierre-Olivier Weill. Any remaining errors are ours. Weng also acknowledges financial support from the National Natural Science Foundation of China (Grant No. 71303014) and Guanghua Leadership Institute (Grant No. 12-02), as well as support from the Spanish Ministry of the Economy and Competitiveness (Project ECO2012-36200) and the Key Laboratory of Mathematical Economics and Quantitative Finance (Peking University), Ministry of Education, China.

[†]Department of Economics, University of North Carolina Chapel Hill; `lifei@email.unc.edu`

[‡]Guanghua School of Management, Peking University; `wengxi125@gsm.pku.edu.cn`

1 Introduction

This paper nests a continuous-time learning model *à la* Jovanovic (1984) into a directed on-the-job search framework. In our model, each worker-firm pair gradually learns its unknown match quality based on cumulative output. The unknown match quality follows a two-point distribution: it is either high or low. Search is directed in the sense that a worker knows the terms of trade offered by different firms before choosing where to apply for a job, as in Moen (1997) and Acemoglu and Shimer (1999). Based on the unidimensional posterior belief about the match quality, a worker decides if and where to search on-the-job.

Due to the ex post heterogeneous performance of matches, any non-trivial allocation inevitably generates a time-varying distribution of matches over the ex post qualities. In general, each individual's optimal decision may depend on this time-varying distribution. As a result, previous studies mainly analyze the steady state where the distribution is constant over time.[1] In a recent pioneering paper, Menzio and Shi (2011) developed a discrete-time framework with directed on-the-job search and aggregate productivity fluctuation. They showed that the unique socially efficient solution is separable, in the sense that it does not depend on the time-varying distribution. As the efficient allocation can be implemented by a decentralized market equilibrium if firms and workers can sign complete contracts, their technique allows for equilibrium analysis of the dynamics of job-to-job transitions.

The current paper can be viewed as a continuous-time analog of Menzio and Shi (2011) with Gaussian learning. Similar to Menzio and Shi (2011), the efficient allocation depends neither on the distribution of the quality of current matches nor unemployment rate. Consequently, the planner's problem can be decomposed into two parts. The planner decides where to send unemployed workers to search for jobs. For employed workers, the planner gradually learns the quality of their matches and decides when to separate them and where each worker should search on-the-job. Since employed workers are allowed to search on-the-job, the planner is able to replace an unpromising match with a new one without suffering the inefficient delay during job-to-job transition caused by search frictions.

We fully characterize the efficient allocation. When the belief about the quality of a match being high is large enough, the planner prefers to maintain the match, so it is inefficient to send the worker to search on-the-job or separate the match. When the belief is low enough, the planner immediately separates the match by ending the worker's employment. In the case where the belief

[1] See Burdett and Mortensen (1998) as an example.

is intermediate, it is efficient to replace the current match with a new one but inefficient to destroy the current match; thus, the worker is assigned to search on-the-job so that the current match will be destroyed only if a new match is formed. More precisely, the job-finding rate is decreasing in the quality of the employed worker's current match because the benefit of job-to-job transition rises as the quality of the current match becomes less promising. Like Jovanovic (1984) and Moscarini (2005), our model can explain a number of robust empirical observations on individual turnovers such as the hump shaped relationship between tenure and the hazard rate of separation.

We also analyze a parametrized example of the model with linear matching technology. This example has a closed-form solution of the planner's efficient allocation and the corresponding ergodic stationary distribution of match quality and unemployment rate. We numerically study the effects of changing the vacancy creation cost and the individual learning speed. Remarkably, we find that (1) reducing the vacancy creation cost has an ambiguous effect on the unemployment rate in the presence of learning, (2) the rate at which employed workers move into unemployment (the EU rate) changes non-monotonically as the cost of vacancy creation declines, and (3) while improving the speed of learning monotonically enhances the unemployment rate, it has a non-monotonic impact on the the rate at which workers move from one employer to another (the EE rate).

Our main contribution is to investigate the role of learning in an equilibrium economy. To the best of our knowledge, Moscarini (2005) is the first paper that integrates a Jovanovic (1984)-like learning model into an equilibrium search framework.[2] Our model is different from that of Moscarini (2005) in the following aspects. First, Moscarini (2005) assumes that an employer's on-the-job search decision is a yes-or-no choice, which allows him to model the firm's problem as a simple stopping-time problem.[3] However, the model cannot explain the heterogeneous job-finding rate among different employed workers engaged in on-the-job search.[4] In contrast, on-the-job search dynamics naturally appear in our model. Employed workers's job-finding rates may differ

[2] Gonzalez and Shi (2010) also develop an equilibrium learning model with directed search. In their model, all matches are homogeneous, but over time a worker learns his job-finding ability which is production-irrelevant.

[3] Strictly speaking, in Moscarini (2005), workers do not actively search on-the-job. New jobs arrive randomly, and workers passively choose whether to accept the new job. While one can add search intensity into Moscarini (2005), the complications arising from his setting as the problem is no longer a stopping-time problem; the value and policy functions cannot be solved explicitly in general as Moscarini (2005) did; and thus it is worth to consider the setting with non-trivial on-the-job search choices.

[4] For example, Fujita (2012) documents a positive relationship between the employers' job security concerns and their job-finding rates.

due to the differences in the quality of their current matches. Second, in Moscarini (2005), search is random and neither firms nor workers have any commitment power, so the allocation is inefficient in general; in contrast we consider a directed search model and characterize the efficient allocation.[5] To implement the efficient allocation, as in Menzio and Shi (2011), one needs to assume that workers and firms can sign complete contracts, which allows one to focus on the ex ante optimal allocation.

Our second contribution is to establish the separability of the efficient allocation in a continuous-time directed on-the-job search model with Gaussian learning.[6] In Menzio and Shi (2011), the problem is formulated in discrete time, and the separability result is proved by using a contraction mapping argument. However, such an argument does not work in a continuous-time model. We therefore develop a different way to prove the separability result. We believe our approach can also be applied to other similar settings.[7]

The rest of the paper is organized as follows. In Section 2, we present the model. In Section 3, we formulate the planner's problem and characterize the efficient allocation and its implications. In Section 4, we parametrize the model with a linear match technology and present some comparative statics results. Section 5 concludes. All proofs appear in the appendix.

2 Model

Population of Firms and Workers. Time is continuous: $t \in [0, +\infty)$. The economy is populated by a unit measure of workers and a sufficiently large measure of long-lived firms to ensure free entry. Both firms and workers are *ex ante* homogeneous. Workers and firms are risk-neutral and discount future payoffs at a rate $r > 0$. Utility is transferable.

[5] As suggested by much empirical evidence, the assumption of random search may not capture the reality of job search. Hall and Krueger (2008) present a survey of the US labor market showing that a large proportion of workers either "knew exactly" or "had a pretty good idea" about their future job at the beginning of the application process. Also, Holzer, Katz, and Krueger (1991) finds a negative relationship between wage and the number of attracted applicants applicants, which supports the assumption of directed search.

[6] Shi (2009) also considers a continuous-time directed search model and proves that the equilibrium is block recursive (separable). However, the proof in Shi (2009) can ensure neither the uniqueness nor the efficiency of the equilibrium.

[7] For example, Papageorgiou (2014) considers a two-sector model where each worker learns his comparative advantage in each sector, and decides which sector to work in. However, due to technical difficulties, this paper cannot allow directed on-the-job search.

Production Technology and Filtering. A consumption good is produced by pairwise firm-worker matches (jobs). The quality of each match, μ, is ex ante uncertain and idiosyncratic, and is randomly assigned by Nature upon matching. The assigned quality of the match is unknown, and the matched worker and firm share a common prior belief about the quality of the match: $\Pr(\mu = \mu_H) = \alpha_0$ and $\Pr(\mu = \mu_L) = 1 - \alpha_0$ where $\alpha_0 \in (0, 1)$, $\mu_H \in \mathbb{R}$ denotes a high-quality match and $\mu_L \in \mathbb{R}$ denotes a low-quality match such that $\mu_H > \mu_L$.

At any moment, a match is exogenously destroyed at a rate $\rho > 0$. Before it is destroyed, the match can generate a consumption good at any instant. The performance of each match depends on its quality. To save notations, we also use t to denote the duration (or tenure) of a given match in the rest of the paper. The cumulative output of a match of duration t, X_t, follows a Brownian motion with drift μ and known variance σ^2:

$$X_t = \mu t + \sigma Z_t \sim N(\mu t, \sigma^2 t),$$

where $\mu \in \{\mu_H, \mu_L\}$ and Z_t is a Wiener process. The realized performance is public information, so both parties commonly update their beliefs about the quality of the match according to Bayes' rule. Specifically, denote $\alpha_t \equiv \Pr(\mu = \mu_H | X^t)$ as the posterior belief about the match quality with duration t where $X^t = \{X_\tau\}_{\tau=0}^t$ is the realized path of its performance. The standard result by Liptser and Shiryaev (2001) implies that

$$d\alpha_t = \alpha_t(1 - \alpha_t) s \frac{[dX_t - \alpha_t \mu_H dt - (1 - \alpha_t)\mu_L dt]}{\sigma}, \tag{1}$$

where $[dX_t - \alpha_t \mu_H dt - (1 - \alpha_t)\mu_L dt]/\sigma$ is the innovation process and follows the standard Wiener process, and the signal-to-noise ratio $s = (\mu_H - \mu_L)/\sigma$ measures the informativeness of the learning. An unemployed worker enjoys a flow payoff b from leisure. Assume $b \in (\mu_L, (1 - \alpha_0)\mu_L + \alpha_0 \mu_H)$, so that a new match with unknown quality is ex ante desirable, but a low-quality match is not. At any moment, the planner can separate a match immediately without paying any cost.

Individual and Aggregate States. The posterior belief about the match quality, $\alpha \in [0, 1]$, is a sufficient statistic for performance history and determines its future prospects; it is also a natural state variable of an employed worker. An unemployed worker's state is simply his unemployment status. As a result, the space of individual state can be defined as a set $\Xi = [0, 1] \cup \{-1\}$. Here we abuse the notation by allowing $\alpha = -1$ if the worker is unemployed to define the individual state of a worker at a given time as a random variable. For each worker, denote his expected "output"

$y(\alpha)$ as follows,

$$y(\alpha) = \begin{cases} \alpha\mu_H + (1-\alpha)\mu_L & \text{if } \alpha \in [0,1], \\ b & \text{if } \alpha = -1. \end{cases}$$

When a worker is unemployed, we abuse the notation and interpret his unemployment benefit b as his "output".

Define $\Delta(\Xi)$ as the set of all cross-sectional probability measures of the individual worker states. At any time t, denote $G_t : \Xi \to [0,1]$ as the cumulative density function corresponding to the cross-sectional probability measure of the individual states of workers at time t, which is a natural aggregate state of the economy. Namely, the measure of unemployed workers at time t is represented by $G_t(0) - G_t(-1) = u_t$ for the sake of convenience.

Search and Matching Technology. At any instant, the planner sends workers and firms searching for new matches at different locations. Specifically, the planner chooses how many vacancies to open in each location and and which locations each worker should search. As is standard in models of directed search such as Moen (1997), the planner finds it optimal to send workers in different individual states to search in different locations but has no incentive to send workers in the same individual state to different locations.

At each location, the workers and the vacancies meet (and match) according to a constant-returns-to-scale matching technology that can be described in terms of the tightness of the location θ (i.e., the vacancy-to-worker ratio at the location). Specifically, at any instant, a worker meets a vacancy at a rate $p(\theta)$ where $p : \mathbb{R}_+ \to \mathbb{R}_+$ is a twice continuously differentiable, strictly increasing, and concave function such that $p(0) = 0$, and $\lim_{\theta\to\infty} p(\theta) = \bar{p} < \infty$.[8] The flow cost of maintaining a vacancy is $k > 0$, which implies that the social cost of maintaining a vacancy also constantly returns to scale and can be described in terms of the tightness θ. The per-worker social cost of maintaining a tightness θ is given by $k\theta$.

[8]In our continuous-time model, $p(\theta)$ denotes the rate rather than the measure of matching at a given time, so the usual restrictions in discrete-time models, $p(\theta) \leq \min\{1, \theta\}$, is not binding.

3 Analysis

3.1 Formulation of the Planner's Problem

At time t, the planner observes the aggregate state of the economy G_t. The planner then decides whether to separate each match. Denote $\delta(\alpha) \in \{0, 1\}$ to be the planner's separation decision for a match whose quality is believed to be high with probability $\alpha \in [0, 1]$, i.e., $\delta(\alpha) = 1$ represents separation while $\delta(\alpha) = 0$ represents no separation.[9] In addition, the planner also chooses $\theta(\alpha) \in \mathbb{R}_+$, the tightness at the location where the α-worker looks for new matches for $\alpha \in \Xi$.[10]

For each individual, we allow the planner's allocation to potentially depend not only on the individual's state but also the calendar time and the aggregate state G_t. Formally, an admissible plan for the planner is a measurable function $(\theta, \delta) : \mathbb{R}_+ \times \Xi \times \Delta(\Xi) \to \mathbb{R}_+ \times [0, 1]$ which is right-continuous in time. Namely, $\theta(t, \alpha, G_t)$ denotes the tightness of the location where the α-worker is sent to search at time t when the aggregate state is G_t, and $\delta(t, \alpha, G_t)$ is the separation decision for a match with posterior α at time t when the cross-sectional distribution of the individual states is G_t, and $\delta(t, -1, G_t)$ is irrelevant so is assumed to be 0. Denote $\mathcal{A}$ as the set of all admissible plans.

Fix an admissible plan and the initial distribution of the states is given by G_0, and by the "law of large number," the process of $\{G_t\}_{t\geq 0}$ is deterministic.[11] The planner's payoff is given by

$$\int_0^\infty e^{-rt} \left[\int_{\alpha\in\Xi} y(\alpha) dG_t(\alpha) \right] dt - \int_0^\infty e^{-rt} \left[\int_{\alpha\in\Xi} k\theta(t, \alpha) dG_t(\alpha) \right] dt, \tag{2}$$

where $\int_{\alpha\in\Xi} y(\alpha)dG_t(\alpha)$ is the total social "output" at time t and $k \int_{\alpha\in\Xi} \theta(t, \alpha)dG_t(\alpha)$ is the total social cost of vacancy creation at time t.

Let $g(t, \alpha)$ denote the probability density function corresponding with $G_t(\cdot)$ whenever it is well-defined. We can use a Kolmogorov forward equation to describe the law of motion of $g(t, \alpha), \forall\alpha \in \Xi$. For any $\alpha \in [0, 1]$ such that $\delta(t, \alpha, G_t) = 0$, we have

[9] As is clear from the subsequent analysis, the planner has no incentive to randomize to separate a match almost surely.

[10] Effectively, the choice of θ is non-trivial only if the match is preserved.

[11] We assume the law of large number holds. Duffie and Sun (2012) provide a formal treatment in a model with a continuum of agents in which the set of individual types (states) of agents is finite. Some of their results still hold in the case with a complete separable metric type space (see Duffie and Sun (2007)).

$$\frac{\partial g(t,\alpha)}{\partial t} = \frac{\partial^2}{\partial \alpha^2}[\Sigma(\alpha)g(t,\alpha)] - [\rho + p(\theta(t,\alpha,G_t))]g(t,\alpha) \tag{3}$$

where

$$\Sigma(\alpha) = \frac{1}{2}s^2\alpha^2(1-\alpha)^2.$$

The first term of the right-hand-side of equation (3) captures the density change due to arriving information and the second term reflects the change due to the separation of matches. On the other hand,[12]

$$g(t+,\alpha) = 0, \forall \alpha \text{ s.t. } \delta(t,\alpha,G_t) = 1. \tag{4}$$

If the measure of matches being destroyed at time t according to the planner's plan is zero[13], then

$$\dot{u}_t = -p(\theta(t,u))u_t + \rho \int_{\alpha\in[0,1]} [1-\delta(t,\alpha,G_t)]dG_t(\alpha) + \int_{\alpha\in[0,1]} \delta(t,\alpha)\frac{\partial g(t,\alpha)}{\partial t}d\alpha \tag{5}$$

where the first term of the right-hand side of equation (5) represents the current unemployed workers who find jobs, the second term represents the employed workers whose matches are destroyed exogenously, and the last term represents the workers whose matches will be destroyed by the planner deliberately. On the contrary, if the measure being destroyed by the planner is positive, then

$$u_{t+} = u_t + \int_{\alpha\in[0,1]} \delta(t,\alpha,G_t)dG_t(\alpha), \tag{6}$$

which implies that the measure of unemployed workers jumps at time t.

The planner's problem is to choose an admissible plan to maximize (2) subject to (3), (4), (5), and (6). Denote the planner's value function as $S(t,G_t)$. In principle, the planner's problem depends on the aggregate state of the economy, the cross-sectional distribution G_t; thus the planner's problem has infinitely many dimensions of "state variables." However, thanks to the seminal work by Menzio and Shi (2011), no such difficulty arises in our model. Theorem 1 is a continuous-time analog of the separability result in Menzio and Shi (2011). It shows that the planner's problem can be broken down into a set of individual problems, and the optimal plan is distribution-free.

[12]For any function f, $f(x-) = \lim_{y\nearrow x} f(y)$ and $f(x+) = \lim_{y\searrow x} f(y)$.

[13]That is to say, $\int_{\alpha\in\delta_t^{-1}(1)} dG_t(\alpha) = 0$ where $\delta_t^{-1}(1) \subset [0,1]$ denotes the (Borel) set of employed workers whose matches will be destroyed at time t according to the planner's plan.

Theorem 1 (Separability of the Planner's Problem)**.** *The planner's value function is linear in $G_t(\cdot)$ and it depends on the calendar time only through G_t. That is,*

$$S(t, G_t) = \int_{\alpha \in [0,1]} V(\alpha) dG_t(\alpha) + u_t U, \tag{7}$$

for any $t \geq 0$ where $V(\alpha)$ and U are the component value functions such that

$$rV(\alpha) = \max_{\delta \in \{0,1\}} \{(1-\delta)rU + \delta \max_{\theta \geq 0} \{y(\alpha) + \rho(U - V(\alpha)) + p(\theta)[V(\alpha_0) - V(\alpha)] - k\theta + \Sigma(\alpha)V''(\alpha)\}\}, \tag{8}$$

and

$$rU = \max_{\theta \geq 0} \{b + p(\theta)[V(\alpha_0) - U] - k\theta\}. \tag{9}$$

where the component value function $V(\alpha)$ is increasing and convex in α, i.e., the planner's optimal allocation depends neither on the aggregate state G_t nor the calendar time t.

The planner's problem in equation (8) is associated with an employed worker or match with a different belief; while the value function U in equation (9) is associated with the unemployed worker's problem. The economics behind equations (8) and (9) will be explained in section 3.2 in great detail. However, it is worth noting that although the planner's optimal solution is stationary, his payoff $S(t, G_t)$ critically depends on the aggregate state G_t, as does the total measure of vacancy created.

As in Menzio and Shi (2011), the separability is driven by the assumption of directed search so that workers can be sent to search at different locations having different tightness measures and job-finding rates. Instead, imagine the search is random. The planner ideally may want to assign different tightness levels for different workers. However, as the search is random, the planner has to assign the same tightness for every worker; thus the distribution of the state of workers, G_t, has to enter the planner's trade-off in general.

The assumption of constant return to scale of both the matching technology and the vacancy creation technology is also crucial for the separability result. It ensures that one cannot express the planner's problem per worker in terms of the tightness θ. In the labor search literature, the cost of vacancy creation is almost exclusively assumed to be a linear function of the measure of vacancies, which is constant return to scale.

3.2 Planner's Solution

Although the planner's problem is the solution of the component value functions (8) and (9), it is not straightforward that such a solution exists and is unique. In Moscarini (2005), the agent effectively needs to decide only when to separate the match, so the existence of the optimal solution is proved by solving it explicitly. In the current model, the planner needs to consider the worker's on-the-job search location as well. In addition, the matching function appears in equation (8). Without assuming some special function form, one cannot solve the optimal solution explicitly. In the following, we show that the optimal solution of the planner's problem exists and is unique.

From the perspective of the planner, the component value function U in equation (9) is the social value function associated with an unemployed worker. The "opportunity cost" of unemployment rU equals the sum of the three terms in the right-hand side of equation (9). The first term is the flow benefit of leisure, the second term is the product of the probability of finding a new job, $p(\theta_u)$, and the "capital gain" from a new match $V(\alpha_0) - U$, and the last term is the social cost of maintaining the tightness θ_u for each unemployed worker in the location where he is searching for a job where $\theta_u^* \in \mathbb{R}_+$ is the planner's optimal choice.

The efficient choice for the tightness at the location visited by unemployed workers $\theta_u^* \in \mathbb{R}_+$ satisfies

$$k \geq p'(\theta_u^*)[V(\alpha_0) - U] \tag{10}$$

The left-hand side is the marginal social cost of increasing the tightness at the location where unemployed workers search for jobs; while the right-hand side is the marginal social benefit, 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 vacancy, and the second term is the capital gain of forming a match. If the planner's optimal choice of tightness is an interior solution ($\theta_u^* > 0$), then the marginal cost equals the marginal benefit. The validity of the interior solution can be ensured by assuming k to be sufficiently small (see Theorem 3 below).

Similarly, for an employed worker, the planner chooses a control variable $\theta(\alpha)$ which is the tightness of the location where the α-worker is sent to search on-the-job, and a stopping rule $\delta(\alpha)$ to separate the match. As the value function $V(\alpha)$ is increasing in α, the stopping rule is characterized by a cutoff belief $\underline{\alpha} \in [0, 1]$:

$$\delta(\alpha) = \begin{cases} 0 & \text{if } \alpha > \underline{\alpha} \\ 1 & \text{otherwise.} \end{cases} \tag{11}$$

If the planner finds it optimal to separate a match at α, its associated value function $V(\alpha) = U$; otherwise, it can be rewritten as the following Hamilton-Jacobi-Bellman (HJB) equation:

$$rV(\alpha) = \max_{\theta \geq 0}\{y(\alpha) + \Sigma(\alpha)V''(\alpha) + p(\theta)[V(\alpha_0) - V(\alpha)] + \rho(U - V(\alpha)) - k\theta\} \tag{12}$$

The "opportunity cost" of a match that is good with posterior α equals the sum of the expected flow payoff $y(\alpha)$, the value of diffusion-learning $\Sigma(\alpha)V''(\alpha)$, the potential expected capital gain of forming a new match, the social cost of maintaining tightness $\theta^*(\alpha)$ in the location where α-workers are sent to search on-the-job, and the expected capital loss following exogenous separation at rate ρ where $\theta^*(\alpha) \in \mathbb{R}_+$ is the optimal choice of the planner.

The efficient choice of the tightness of the location visited by an employed worker is $\theta^*(\alpha) \in \mathbb{R}_+$ such that

$$k \geq p'(\theta^*(\alpha))[V(\alpha_0) - V(\alpha)] \tag{13}$$

where α is the posterior belief about his current match quality. As the condition is similar to (10), we do not discuss it again. However, it is noteworthy that the optimal tightness choice tightness $\theta^*(\alpha)$ is strictly positive only if $V(\alpha_0) - V(\alpha)$ is sufficiently large, i.e., the current match is believed to be bad with high probability. Unlike the case of the unemployed worker, one may not obtain interior solutions for all α by assuming k to be sufficiently small. Fix a $k > 0$, if $p'(\cdot)$ is bounded, as the value function $V(\alpha)$ is continuous, for α smaller than but close to α_0, one must have $\theta^*(\alpha) = 0$.

Last, the planner's optimal separation choice is an optimal stopping-time decision. Because $b > 0$, $U > 0$ (the planner can always choose $\theta_u = 0$). In addition, because $b \in (\mu_L, \mu_H)$, one must have $\underline{\alpha} \in (0, 1)$. By the standard result (see Dixit (1993)), the cutoff belief $\underline{\alpha}$ must satisfy the following two boundary conditions:

$$V(\underline{\alpha}) = U \text{ (value matching) and } V'(\underline{\alpha}) = 0 \text{ (smooth pasting)}, \tag{14}$$

Roughly speaking, the planner maintains a match only if its expected social value, $V(\alpha)$ is greater than its opportunity cost, U. At the cutoff belief, the social value does not have a kink; otherwise, there is an additional value associated with choosing to "wait and see," which contradicts the optimality of separation. We summarize the properties of the planner's solution into the following theorem.

Theorem 2. *There is a unique triple* $(V(\alpha), U, \underline{\alpha})$ *such that* $V : [0, 1] \to \mathbb{R}$ *is continuous and twice differentiable,* $U \in \mathbb{R}_+$, $\underline{\alpha} \in (0, \alpha_0)$*, and they solve (8), (9), and (14) and characterize the planner's component value function in Theorem 1. In the efficient allocation,* *1. a match is separated by the planner only if its quality $\alpha \le \underline{\alpha}$, and*

2. $\theta^*(\alpha)$ is decreasing in α.

The idea behind Theorem 2 is as follows. Given any U, the initial boundary conditions (14) and the second-order differential equation (12) has a unique solution, which pins down the value of $V(\alpha_0)$. Next, given $V(\alpha_0|U)$, the planner further pins down the optimal search plan for the unemployed worker and obtains a value $\hat{U}$, which has to be equal to U for consistency. Because $V(\alpha|U)$ is increasing, to maintain the equality of the marginal cost and the marginal benefit of vacancy creation in condition (13), $\theta^*(\alpha)$ has to increase to response the decline of α. To put it differently, as the worker's current match becomes worse, the planner will create more vacancy for him. Because $p(\cdot)$ is increasing, the job-finding rate of an employed worker is also decreasing in the quality of his current match α if he is engaged in on-the-job search. Furthermore, since $V(\alpha)$ is increasing and greater than U as long as $\alpha \ge \underline{\alpha}$, the job-finding rate of any employed worker who is searching on-the-job is lower than that of an unemployed worker. As the statement holds for any $\alpha \in (\underline{\alpha}, 1]$, the average job-finding rate of employed workers is low regardless of the cross-sectional distribution of the quality of matches.

As mentioned earlier, there may be no interior solution for matches with high α. If k is so small that unemployed worker's solution is an interior one, $\theta_u^* > 0$, then by the continuity of $V(\alpha)$ and the boundary conditions (14), there must exist a $\bar{\alpha} \in [\underline{\alpha}, \alpha_0]$ such that $\theta^*(\alpha) > 0$ for $\alpha \in (\underline{\alpha}, \bar{\alpha}]$.[14] In other words, if k is so large that it is inefficient for unemployed workers to search, then it is also inefficient to allow employed workers to search; eventually all workers become unemployed. To avoid such a trivial case, we assume that the vacancy creation cost is sufficiently small in the rest of this paper. Formally,

Theorem 3. *Let $\hat{k} = \frac{\mu_H - \mu_L}{r+\rho} \frac{z(1-z)}{\beta_2 - (1-z)} z^{\beta_2 - 1}(1-z)^{-\beta_2}$ where $z = \frac{(\beta_2 - 1)(b - \mu_L)}{\beta_2(\mu_H - \mu_L) + \mu_L - b}$. Suppose that $k \in (0, \hat{k})$. In the efficient allocation,*

1. unemployed workers search, $\theta_u^* > 0$, and

2. employed workers search on-the-job, $\theta^*(\alpha) > 0$, if and only if $\alpha \in (\underline{\alpha}, \bar{\alpha}]$ where $0 < \underline{\alpha} < \bar{\alpha} \le \alpha_0$.

Similar to Menzio and Shi (2011), in the current context, the planner's efficient plan can be implemented by a decentralized Block recursive market equilibrium if firms and workers can

[14] We prove that $\bar{\alpha} < \alpha_0$ when $p'(0)$ is bounded. If $p'(0)$ is unbounded, then naturally $\bar{\alpha} = \alpha_0$.

sign bilaterally efficient contracts.[15] Specifically, in the Block recursive equilibrium, there is a continuum of submarkets opening. Submarkets differ from each other in terms of the market tightness and the worker's expected utility from joining a firm. From the perspective of workers, the equilibrium expected utility from joining a firm is decreasing in the market tightness and therefore in the job-finding rate. Given a bilaterally efficient contract, the two parties in the firm-worker pair behave in a way that maximizes their joint surplus. As the market is complete, the market mechanism internalizes the social cost of maintaining vacancy; thus, the firm terminates his current match efficiently, and an employed worker searches on-the-job efficiently. Remarkably, when the quality of his current match is low, the worker believes that he will soon become jobless, so he has a strong incentive to start a new job. As a result, he will search in a submarket with low expected utility but high market tightness (or job-finding rate). On the other hand, when the quality of his current match is high, the worker does not worry about job security, so he will either not search on-the-job or search in a submarket with high expected utility but low job-finding rate. In sum, the motivation for a worker's on-the-job search comes mainly from his concern about his long-term career path (job security and expected discounted payoff) rather than his interest in the short-term payoff. In fact, as both workers and firms are risk-neutral, the optimal flow wage is indeterminate.

3.3 Individual Turnover

Because the efficient allocation does not depend on the aggregate state, a worker's efficient career path is governed by the stationary search rule specified by Theorem 2. For each job, the belief about the match quality starts from α_0 and evolves according to Bayes' rule (1) given the realized performances. A match is terminated in one of three ways: (1) an exogenous separation shock arrives, (2) a new (better) job is found and the worker switches to the new position, or (3) the planner deliberately destroys the current match when α reaches $\underline{\alpha}$.

A match can be viewed as a Bayesian experimentation: from the perspective of the planner, an on-going match is a "risky arm" as the true match quality is unknown, while making the worker unemployed is a "safe arm".[16] As is standard in the Bayesian experimentation literature, the belief

[15] Menzio and Shi (2009) show that, if the contracting space is complete, the firm's profit maximization contract is bilaterally efficient.

[16] Rigorously speaking, the social payoff of an unemployed worker rU is also uncertain due to the uncertain outcome of the job search.

process $\{\alpha_t\}$ of each specific match is a martingale, i.e., $\mathbb{E}[\alpha_{t+\Delta}|\alpha_t] = \alpha_t$ for any $\Delta \geq 0$. However, as a match will be separated if $\alpha \leq \underline{\alpha}$, the belief sequence $\{\alpha_t\}$ of an ongoing project is a strict sub-martingale. In other words, $\mathbb{E}[\alpha_{t+\Delta}|\alpha_\tau > \underline{\alpha}, \forall\tau \in [t, t+\Delta], \alpha_t] > \alpha_t$, which is the standard selection effect of learning: bad matches' performances are low in expectation so that there is a higher probability that they will soon be separated; thus the average quality of an ongoing match becomes better as the tenure of the worker rises. This means that, in the long run, the probability that the match is deliberately destroyed by the planner shrinks over the worker's tenure. This selection effect implies that the average employment-to-unemployment (AEU) transition rate should decline over a worker's tenure. As the efficient on-the-job search policy $p(\theta^*(\alpha))$ is decreasing in α and equals zero when α is sufficiently large, the average employment-to-employment (AEE) transition rate of a worker is also decreasing in his tenure in the long run due to the selection effect.

However, in the short run, a wait-and-see effect emerges. Because the path of a Brownian motion is continuous, the posterior belief cannot jump, so some workers' beliefs may reside in the area such that $\alpha_t \geq \underline{\alpha}$ and $\theta^*(\alpha_t) > 0$ with positive probability when t is small. These workers start to actively search on-the-job, resulting in a positive AEE transition rate when t is small but positive. In addition, the posterior belief α_t will reach $\underline{\alpha}$ with positive probability, resulting in a positive AEU transition rate when t is small but positive. As a result, both the AEU and AEE rates rise initially but eventually decline over tenure; thus the job separation rate is also non-monotone in tenure as it is in both Jovanovic (1979) and Moscarini (2005).[17] We summerize the implications above in the following theorem.

Theorem 4. *Conditional on match continuation, the expected match quality is increasing in tenure. Both the average employment-to-employment transition rate and the average employment-to-unemployment transition rate are initially increasing and then decreasing in tenure.*

4 Linear Matching Technology

In this section, we examine a special case of our model with a linear matching function. Doing so allows us to obtain the closed-form solution of the efficient allocation and to conduct some comparative statics exercises numerically. Specifically, suppose that $p(\theta) = \min\{\theta, \lambda\}$ where λ is a finite upper bound of the matching rate that ensures the existence of the efficient solution. What

[17]See Farber (1999) for evidence of the hump-shaped average separation rate. See Menzio, Telyukova, and Visschers (2012) for evidence of the hump-shaped AEE rate.

is remarkable is that when both the matching function and the cost function of vacancy creation are linear, the efficient allocation is a bang-bang solution generically; and thus the efficient search allocation is effectively random: If k is sufficiently small, there exists $\bar{\alpha} \in [\underline{\alpha}, \alpha_0]$ such that

$$\theta^*(\alpha) = \begin{cases} 0 & \text{if } \alpha > \bar{\alpha} \\ \lambda & \text{if } \alpha \in [\underline{\alpha}, \bar{\alpha}]. \end{cases}$$

Without loss of any generality, we normalize $\mu_H = 1, \mu_L = 0$, so $y(\alpha) = \alpha$. When the match quality is believed to be sufficiently high, the worker does not search on-the-job, so for $\alpha > \bar{\alpha}$, the value function $V(\alpha)$ satisfies

$$rV(\alpha) = \alpha + \rho(U - V(\alpha)) + \Sigma(\alpha)V''(\alpha).$$

On the other hand, when $\alpha \in (\underline{\alpha}, \bar{\alpha})$, we have

$$rV(\alpha) = \alpha + \rho(U - V(\alpha)) + \lambda[V(\bar{\alpha}) - V(\alpha) - k] + \Sigma(\alpha)V''(\alpha).$$

At $\underline{\alpha}$ and $\bar{\alpha}$, we have the boundary conditions including the value matching and smooth pasting conditions:

$$V_1(\underline{\alpha}) = U, \quad V_1'(\underline{\alpha}) = 0, \quad V_1(\bar{\alpha}) = V_0(\bar{\alpha}), \quad V_1'(\bar{\alpha}) = V_0'(\bar{\alpha}). \tag{15}$$

In addition, at $\bar{\alpha}$, the marginal social benefit of on-the-job search is $\lambda[V_0(\alpha_0) - V_0(\bar{\alpha})]$, which is the product of the probability that the worker finds a new job and the social capital gain of the job-to-job transition; while the corresponding minimal social marginal cost to maintain a matching rate $p(\theta) = \lambda$ is $k\lambda$. The optimality implies that the marginal social benefit equals the marginal social cost. Canceling λ yields

$$V_0(\alpha_0) - V_0(\bar{\alpha}) = k, \tag{16}$$

which constitutes the last boundary condition.

We can solve for the value function $V(\alpha)$ explicitly as:

$$V(\alpha) = \begin{cases} \frac{\alpha+\rho U}{r+\rho} + \kappa_0 \alpha^{1-\beta_2}(1-\alpha)^{\beta_2}, & \text{if } \alpha > \bar{\alpha} \\ \frac{\alpha+\rho U+\lambda V(\bar{\alpha})}{r+\rho+\lambda} + \kappa_1 \alpha^{\beta_1}(1-\alpha)^{1-\beta_1} + \kappa_2 \alpha^{1-\beta_1}(1-\alpha)^{\beta_1}, & \text{if } \alpha \in [\underline{\alpha}, \bar{\alpha}] \\ U & \text{if } \alpha < \underline{\alpha} \end{cases}$$

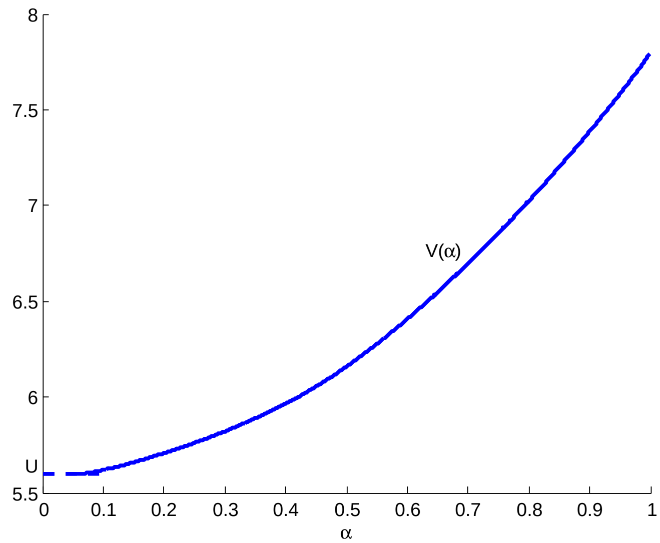

Figure 1: Planner's component value function for an employed worker.

where $\beta_1 = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2(r+\rho+\lambda)}{s^2}}$, $\beta_2 = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2(r+\rho)}{s^2}}$, and κ_0, κ_1, κ_2 are coefficients satisfying the boundary conditions at $\underline{\alpha}$ and $\bar{\alpha}$, (15, 16).

To further address the properties of the efficient plan, we parametrize the model and numerically examine some comparative statics with respect to the learning speed s and the vacancy creation cost k. Suppose that the parameter values are given in Table 1. The corresponding numerical results are $\underline{\alpha} = 0.055$, $\bar{\alpha} = 0.4522 < \alpha_0$, $U = 5.5982$, and the value function of employed workers is given in Figure 1. Consistent with our theoretical analysis, the value function $V(\alpha)$ is increasing and convex.

λ	r	ρ	α_0	b	k
1	0.1	0.1	0.5	0.1	0.1

Table 1: Parameter Values

4.1 The Ergodic Distribution

The closed-form solution of the efficient allocation allows us to explicitly characterize the stationary distribution of workers' employment status.[18] By the Kolmogorov forward equation (3), the stationary and ergodic density $g(\alpha)$ should satisfy the following differential equation:

$$0 = \begin{cases} \frac{d^2}{d\alpha^2}[\Sigma(\alpha)g(\alpha)] - \rho g(\alpha), & \text{if } \alpha > \bar{\alpha} \text{ and } \alpha \neq \alpha_0 \\ \frac{d^2}{d\alpha^2}[\Sigma(\alpha)g(\alpha)] - (\rho+\lambda)g(\alpha) & \text{if } \alpha \in [\underline{\alpha}, \bar{\alpha}] \end{cases}$$

since an employed worker searches on-the-job only if $\alpha \in [\underline{\alpha}, \bar{\alpha}]$, and the job-finding rate is λ. The above equation does not hold at α_0 where the inflow from unemployment and from other jobs (quits) produces a kink in the density, as shown by Moscarini (2005).

The support of the ergodic density contains three intervals: $[\underline{\alpha}, \bar{\alpha}]$, $(\bar{\alpha}, \alpha_0]$ and $(\alpha_0, 1]$. The general solution on each interval is[19]

$$g(\alpha) = \begin{cases} \eta_0 \alpha^{-1-\gamma_1}(1-\alpha)^{\gamma_1-2}, & \text{if } \alpha > \alpha_0 \\ \eta_1 \alpha^{\gamma_1-2}(1-\alpha)^{-1-\gamma_1} + \eta_2 \alpha^{-1-\gamma_1}(1-\alpha)^{\gamma_1-2}, & \text{if } \alpha \in (\bar{\alpha}, \alpha_0] \\ \eta_3 \alpha^{\gamma_2-2}(1-\alpha)^{-1-\gamma_2} + \eta_4 \alpha^{-1-\gamma_2}(1-\alpha)^{\gamma_2-2}, & \text{if } \alpha \in [\underline{\alpha}, \bar{\alpha}] \end{cases}$$

where $\gamma_1 = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2\rho}{s^2}}$ and $\gamma_2 = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2(\rho+\lambda)}{s^2}}$.

Following Moscarini (2005), we impose the following boundary conditions:

$$g(\underline{\alpha}+) = 0 \tag{17}$$

$$g(\bar{\alpha}+) = g(\bar{\alpha}-) \tag{18}$$

$$g(\alpha_0+) = g(\alpha_0-) \tag{19}$$

$$\Sigma(\underline{\alpha})g'(\underline{\alpha}+) + \rho \int_{\underline{\alpha}}^{1} g(\alpha)d\alpha = p(\theta^*(\underline{\alpha}))(1 - \int_{\underline{\alpha}}^{1} g(\alpha)d\alpha) \tag{20}$$

$$\Sigma(\alpha_0)[g'(\alpha_0-) - g'(\alpha_0+)] = p(\theta^*(\underline{\alpha}))(1 - \int_{\underline{\alpha}}^{1} g(\alpha)d\alpha) + \int_{\underline{\alpha}}^{\bar{\alpha}} p(\theta^*(\alpha))g(\alpha)d\alpha. \tag{21}$$

[18] Notice that we need the separability result even if we focus on the socially efficient allocation in the steady state. In general, the planner's problem is difficult to solve because she has to take into account the fact that her decision changes the stationary distribution. If the planner's problem is not separable, it is unclear whether the socially efficient allocation is stationary.

[19] For $\alpha > \alpha_0$, the general solution is actually $\eta_0 \alpha^{-1-\gamma_1}(1-\alpha)^{\gamma_1-2} + \eta_0' \alpha^{\gamma_1-2}(1-\alpha)^{-1-\gamma_1}$, but since $\alpha = 1$ is included in the domain, η_0' has to be zero to ensure the boundedness of the distribution function.

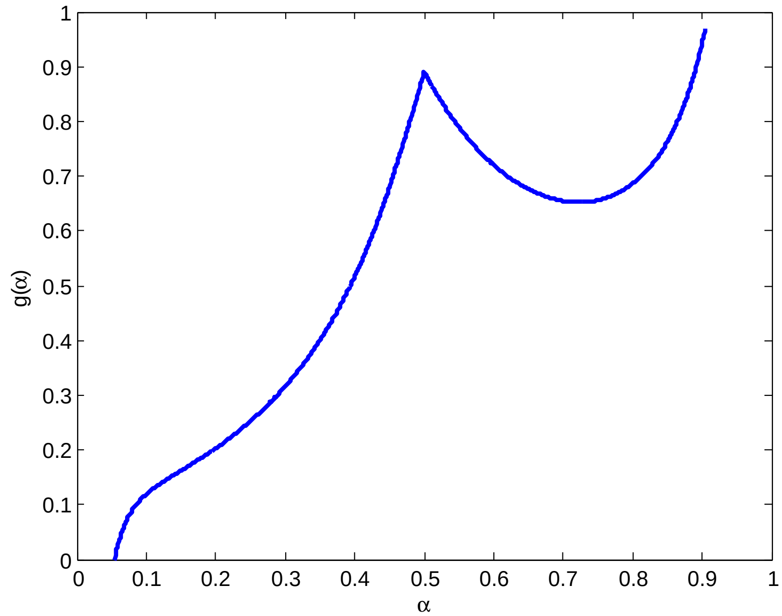

Figure 2: Ergodic density of the quality of matches.

Equation (17) is the standard condition for "attainable" boundaries; equations (18) and (19) require the density function to be continuous at $\bar{\alpha}$ and α_0; equation (20) requires that total flows (respectively) in and out of unemployment must balance; and finally, equation (21) requires that total flows (respectively) in and out of new employment must balance.[20]

Because equations (17)-(21) constitute a system of linear equations about five unknowns η_0-η_4, we can directly solve the coefficients η_i once we know $\underline{\alpha}$ and $\bar{\alpha}$. We parametrize the model by using the same parameter values used for Figure 1. The ergodic density of the quality of existing matches is depicted in Figure 2, and the unemployment rate is about 10%.

The shape of the ergodic density is remarkable. Because the initial quality of every new match is α_0, the ergodic density function $g(\cdot)$ has a kink at α_0. Because on-the-job search is efficient only if α is low, the ergodic density for small α is low. A low quality match (small α) will be separated with high probability according to the efficient allocation. The smaller the α of a match, the lower the chance that it can survive. As a result, when $\alpha < \alpha_0$, the ergodic density declines as α decreases. When $\alpha \geq \alpha_0$, the worker does not search on-the-job. A match is separated if either an exogenous separation shock arrives or the match's performance is poor for sufficient time that so that α declines to some point lower than α_0. As α increases, two effects govern the shape of

[20]Equation (20) is derived from equation (5) and the Kolmogorov forward equation.

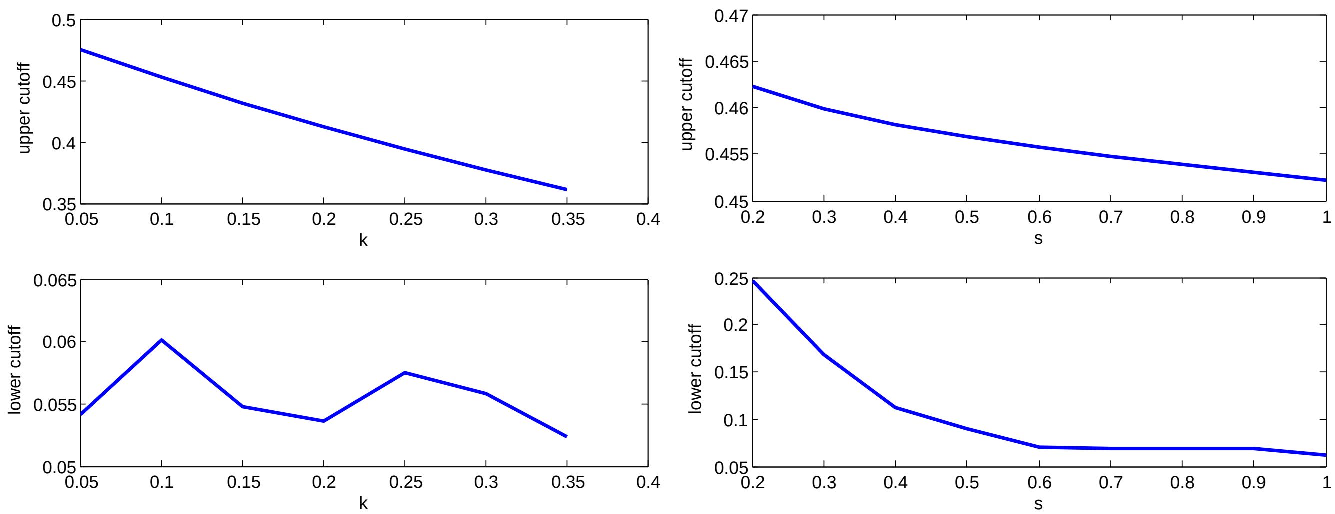

(a) $\bar{\alpha}$ and $\underline{\alpha}$ as functions of k

(b) $\bar{\alpha}$ and $\underline{\alpha}$ as functions of s

Figure 3: Comparative statics on cutoffs.

the ergodic density $g(\alpha)$ for $\alpha \in [\alpha_0, 1]$. First, every new match starts at α_0. However, only some of them are high-quality matches whose α will go beyond α_0 with high probability. The larger α becomes, the less likely that a new (bad) match "ends up" at α due to the selection effect. This suggests that $g(\cdot)$ should be decreasing in α when $\alpha \geq \alpha_0$. Second, by Bayes' rule, as α goes to 1, $d\alpha_t$ goes to zero according to Equation (1); thus once a match is believed to be good with a probability α close to 1, it is more likely that α will escape a neighborhood of 1. This second effect suggests that the ergodic density $g(\cdot)$ should be increasing in α when α is close to 1. The two effects together imply an inverse U-shaped ergodic density for $\alpha \geq \alpha_0$.

4.2 Comparative Statics

In this subsection, we present the results of some numerical comparative statics exercises with respect to the vacancy creation cost k and the learning speed s. First, we show how the efficient learning and on-the-job search plan responses to the changes in k and s. In Figure 3(a), we depict how the cutoffs change in the vacancy creation cost k. It is not surprising that $\bar{\alpha}$ is decreasing in k. On-the-job search becomes unlikely as the search cost increases. However, it is novel that $\underline{\alpha}$ is non-monotonic in k. Economically, a higher cost of vacancy creation leads to less job searching by employed and unemployed workers; and thus the value of being employed and the value of being unemployed both decline. The optimal stopping belief $\underline{\alpha}$ depends on the relative changes of $V(\cdot)$

and U; thus it may not be monotonic in k.

In Figure 3(b), we plot how the cutoffs change in the learning speed s. Both $\bar{\alpha}$ and $\underline{\alpha}$ are decreasing in s. As the speed of learning goes up, the benefit of learning (which is measured by the value of newly arriving information) increases while the cost of learning (which is captured by discounting and the vacancy creation cost) remains. Consequently, when s is higher, it is more desirable to maintain an existing match for a given α; and thus the on-the-job search and separation occur only if the match is believed to be good with lower probability. From Figure 3(b), $\underline{\alpha}$ decreases faster than $\bar{\alpha}$ when s is low; the opposite occurs when s is high.

Next, we turn to the comparative statics of the ergodic unemployment rate and the aggregate (cross-sectional) employment status transition rate with respect to k and s. The results are presented in Figure 4.

Unemployment Rate. Remarkably, the ergodic unemployment rate, u, is non-monotone in k. The non-monotonicity results from the non-monotonic relationship between k and $\underline{\alpha}$. As k increases, the outflow of the pool of unemployed workers decreases. However, as $\underline{\alpha}$ may also decline, a match may be maintained even if its quality is low, which decreases the inflow of the pool of the unemployed workers as well. Hence, the total effect of increasing k on the unemployment rate is ambiguous in the presence of learning.[21]

On the other hand, when the learning speed s increases, more inefficient matches are destroyed sooner. Although the stopping cutoff $\underline{\alpha}$ declines, the unemployment rate goes up. Consistent with the changes in $\bar{\alpha}$ and $\underline{\alpha}$, the unemployment rate does not increase too much in s when s is low, but increases rapidly in s when s is high.

EE Transition Rate. The ergodic cross-sectional employment-to-employment (EE) transition rate is defined as the ratio between the stationary measure of workers who switch jobs in a unit length of time and that of employed workers, i.e.,

$$EE \triangleq \frac{\int_{\underline{\alpha}}^{1} p(\theta^*(\alpha)g(\alpha)d\alpha}{1-u} = \frac{\lambda[G(\bar{\alpha}) - G(\underline{\alpha})]}{1-u}.$$

In the graph, the EE rate is decreasing in k. The result is intuitive. As k increases, it becomes less efficient to create a large amount of vacancy in order to ensure a high job-finding rate of employed workers whose current matches are not sufficiently promising.

When it comes to the change in s, the total effect is unclear. When s is low, $\underline{\alpha}$ decreases faster than $\bar{\alpha}$, and this implies that the total measure of workers who switch jobs increases as s goes up.

[21] Notice that the change in k has a negligible effect on the unemployment rate (it stays around 10% in the example).

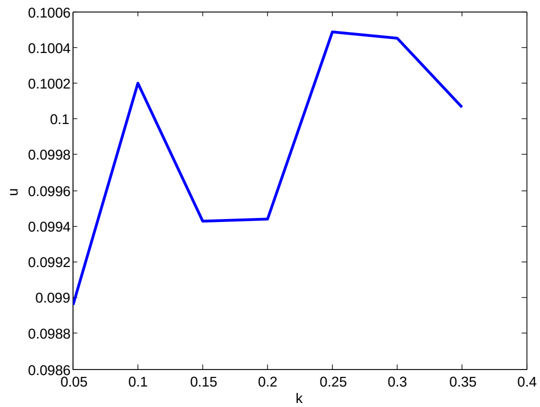

(a) Unemployment rate u as a function of k

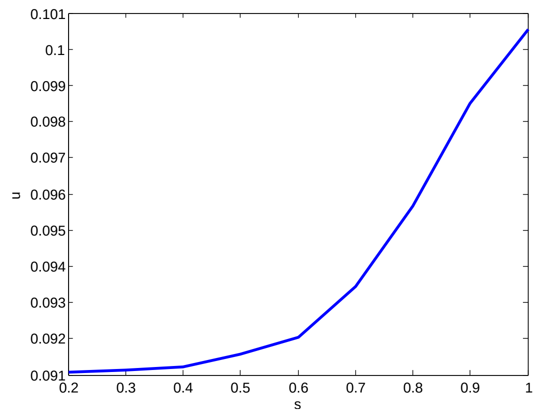

(b) Unemployment rate u as a function of s

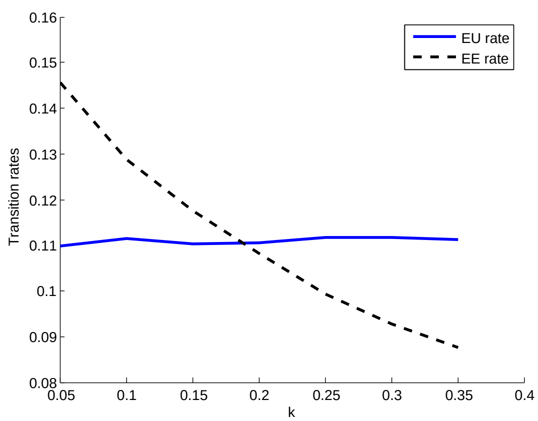

(c) Aggregate EU/EE rate as a function of k

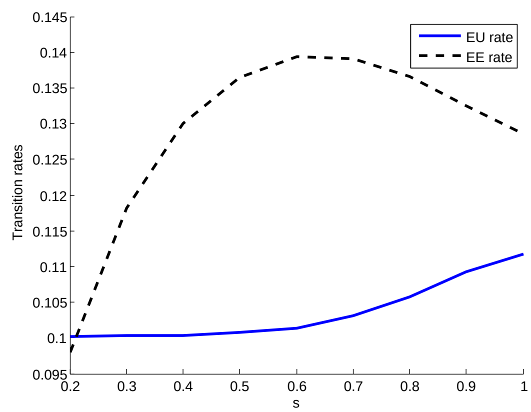

(d) Aggregate EU/EE rate as a function of s

Figure 4: Comparative statics on unemployment rate, employment-to-employment (EE) transition rate and employment-to-unemployment (EU) transition rate.

Hence, the total EE rate increases as well. When s is high, $\bar{\alpha}$ decreases faster than $\underline{\alpha}$, and this implies that the total EE rate decreases as s goes up.

UE and EU Transition Rates. Because the job-finding rate is λ for any non-trivial k, s, the measure of unemployed workers who successfully find jobs in a unit length of time is λu, which implies that the total unemployment-to-employment transition rate (UE) is λ, which depends on neither k nor s.

In the steady state, the measure of unemployed workers is constant over time, so the measure of employed workers who become unemployed in a unit length of time equals λu. Thus the employment-to-unemployment transition rate (EU) is

$$EU = \frac{\lambda u}{1-u} = \lambda(-1 + \frac{1}{1-u}).$$

As a result, the comparative statics of EU with respect to k and s are determined by those of u. As u is non-monotone in k, EU is non-monotone in k as well.

5 Conclusion

Since the seminal work by Jovanovic (1979), the learning theory of job turnover has remained popular among labor economists. A worker-firm match is modeled as an experience good whose quality is initially unknown and is gradually learned over time through a sequence of observations of its performance. This model is useful for studying individual job turnover, and it successfully explains a number of empirical observations (e.g., the initially positive but soon negative relationship between tenure and the hazard rate of separation). On the other hand, a large body of search and matching literature focuses on unemployment and aggregate job turnover.[22] This paper combines two bodies of literature by introducing microeconomic learning into a macroeconomic search model. We prove that the socially efficient allocation is separable, so the socially efficient allocation can be characterized as a simple bandit problem. An employed worker searches on-the-job if the quality of his current match is not sufficiently good, and his job-finding rate is decreasing in his current match quality. When the quality becomes sufficiently low, the match is separated and the worker becomes jobless. The model generates a number of empirical implications about the dynamics of the job turnover rate.

[22] See Mortensen and Pissarides (1994), Burdett and Mortensen (1998) and Menzio and Shi (2011) as examples.

A Appendix: Omitted Proofs

A.1 Proof of Theorem 1

Proof. The proof proceeds in three steps.

1. We show that the planner's value function at time 0 is separable.
2. We show that in the optimal solution, $\hat{U}_t = U$ for all $t \geq 0$, and that the maximization problem has a (stationary) Markovian solution.
3. We show that $V(\alpha)$ is convex and increasing.

Step 1. We want to show that the planner's value function at time 0 can be rewritten as:

$$\int_{\alpha\in[0,1]} \hat{V}_0(\alpha) dG_0(\alpha) + u_0\hat{U}_0,$$

where $\hat{U}_0$ denotes the expected payoff of an unemployed worker ($\alpha = -1$) at time 0:

$$\hat{U}_0 = \sup_{\theta,\sigma\in\mathcal{A}} \mathbb{E}\left[\int_0^\tau e^{-rt}[b - k\theta_t(u)]dt + e^{-r\tau}\hat{V}_\tau(\alpha_0)\right], \tag{22}$$

and $\hat{V}_0(\alpha)$ denotes the expected payoff of an employed worker with prior $\alpha \in [0,1]$ at time 0:

$$\hat{V}_0(\alpha) = \sup_{\theta,\sigma\in\mathcal{A}} \mathbb{E}\left[\int_0^\tau e^{-rt}[y(\alpha_t) - k\theta_t(\alpha_t)]dt + e^{-r\tau}\hat{U}_\tau|\alpha_0 = \alpha\right], \tag{23}$$

and $\hat{U}_\tau$ and $\hat{V}_\tau(\alpha)$ are defined in the same fashion at time $\tau \geq 0$. Specifically, fix any admissible plan $(\theta, \delta) \in \mathcal{A}$. For each worker whose state is α at time 0, the probability measure of his individual state $\hat{\alpha}$ at time t is $\mathbb{P}(\alpha(t) = \hat{\alpha}|\alpha(0) = \alpha)$. Suppose the initial cross-sectional distribution of the individual state is $G_0(\alpha)$, then we have

$$G_t(\hat{\alpha}) = \int_{\alpha\in\Xi} \mathbb{P}(\alpha(t) = \hat{\alpha}|\alpha(0) = \alpha)dG_0(\alpha)$$

for any $\hat{\alpha}, \alpha \in \Xi$, where $G_t(\hat{\alpha})$ is also the cross-sectional distribution of the individual state of the economy at time t.

Define

$$\begin{aligned} h_t(\alpha) &= \mathbb{E}[y(\alpha(t)) - k\theta(t, \alpha(t))|\alpha(0) = \alpha] \\ &= \int_{\hat{\alpha}\in\Xi} [y(\hat{\alpha}) - k\theta(t, \hat{\alpha})]d\Pr(\alpha(t) = \hat{\alpha}|\alpha(0) = \alpha), \end{aligned}$$

which is the expected social flow payoff of a worker at time t conditional on his individual state being α at time 0, and

$$\begin{aligned}\int_{\alpha\in\Xi} h_t(\alpha)dG_0(\alpha) &= \int_{\alpha\in\Xi}\int_{\hat{\alpha}\in\Xi} y(\hat{\alpha}) - k\theta(t,\hat{\alpha})d\Pr(\alpha(t)=\hat{\alpha}|\alpha(0)=\alpha)dG_0(\alpha)\\ &= \int_{\hat{\alpha}\in\Xi}[y(\hat{\alpha}) - k\theta(t,\hat{\alpha})]dG_t(\alpha).\end{aligned}$$

For a given plan $(\theta,\sigma)\in\mathcal{A}$, the planner's payoff

$$\begin{aligned}W^{\theta,\sigma}(0,G_0) &= \int_0^\infty e^{-rt}\int_{\hat{\alpha}\in\Xi}[y(\hat{\alpha}) - k\theta(t,\hat{\alpha})]dG_t(\alpha)dt\\ &= \int_0^\infty e^{-rt}\int_{\alpha\in\Xi} h_t(\alpha)dG_0(\alpha)dt = \int_0^\infty\int_{\alpha\in\Xi} e^{-rt}h_t(\alpha)dG_0(\alpha)dt\\ &= \int_{\alpha\in\Xi}\int_0^\infty e^{-rt}h_t(\alpha)dtdG_0(\alpha).\end{aligned}$$

The planner's goal is to choose $(\theta,\sigma)\in\mathcal{A}$ to maximize $W^{\theta,\sigma}(0,G_0)$. Denote

$$S(0,G_0) = \sup_{(\theta,\sigma)\in\mathcal{A}} W^{\theta,\sigma}(0,G_0).$$

For $\alpha\in[0,1]$, define

$$\sup_{(\theta,\sigma)\in\mathcal{A}}[\int_0^\infty e^{-rt}h_t(\alpha)dt] = \sup_{(\theta,\sigma)\in\mathcal{A}}\mathbb{E}\left[\int_0^\tau e^{-rt}[y(\alpha_t) - k\theta_t(\alpha_t)]dt + e^{-r\tau}\hat{U}_\tau|\alpha_0=\alpha\right] = \hat{V}_0(\alpha)$$

where τ is the random time at which the current match is separated, and if $\alpha=-1$,

$$\sup_{(\theta,\sigma)\in\mathcal{A}}[\int_0^\infty e^{-rt}h_t(\alpha)dt] = \sup_{(\theta,\sigma)\in\mathcal{A}}\mathbb{E}\left[\int_0^\tau e^{-rt}[b - k\theta_t(u)]dt + e^{-r\tau}\hat{V}_\tau(\alpha_0)\right] = \hat{U}_0$$

where τ denotes the random time at which the worker finds a job.

Hence, $\int_{\alpha\in[0,1]}\hat{V}_0(\alpha)dG_0(\alpha) + u_0\hat{U}_0 \geq W^{\theta,\sigma}(0,G_0)$, for any $\theta,\sigma\in\mathcal{A}$, so

$$S(0,G_0) \leq \int_{\alpha\in[0,1]}\hat{V}_0(\alpha)dG_0(\alpha) + u_0\hat{U}_0.$$

Note that as $\int_0^\infty e^{-rt}h_t(\alpha)dt$ does not depend on $G_0(\alpha)$. The optimal plans (θ,σ) of problems (22) and (23) (if they exist) do not depend on $G_0(\alpha)$ either. They depend only on t and $\alpha\in\Xi$.

On the other hand, the solutions of (22) and (23) depend on t and $\alpha\in\Xi$, so they belong to $\mathcal{A}$. By the definition of $S(0,G_0)$, we have

$$S(0,G_0) \geq \int_{\alpha\in[0,1]}\hat{V}_0(\alpha)dG_0(\alpha) + u_0\hat{U}_0,$$

which leads to the desired result.

As the proof works for any G_0, for any $t \geq 0$, one can apply the previous argument to prove that the planner's problem is separable at any time $t \geq 0$ with corresponding aggregate state G_t. In the rest of the proof, we focus on the problem at time 0.

Step 2. First, notice that the planner's objective is to maximize

$$S(0, G_0) = \int_{\alpha \in [0,1]} \hat{V}_0(\alpha) dG_0(\alpha) + u_0 \hat{U}_0.$$

Suppose on the contrary that in the optimal solution, there exists some t such that $\hat{U}_t \neq \hat{U}_0$. This allows for two possibilities. When $\hat{U}_t > \hat{U}_0$, then at time 0, the planner can simply change the unemployed workers' strategies to be the same as the time t unemployed workers' strategies. By doing so, the planner obtains a higher $S(0, G_0)$, which leads to a contradiction. Similarly, when $\hat{U}_t < \hat{U}_0$, the planner can increase $S(0, G_0)$ by letting the time t unemployed workers mimic the time 0 unemployed workers. Hence, it follows that $\hat{U}_t = U$ for all $t \geq 0$ in the optimal solution.

Second, since $\hat{U}_t = U$ for all $t \geq 0$,

$$\hat{V}_0(\alpha) = \sup_{\theta, \sigma \in \mathcal{A}} \mathbb{E} \left[\int_0^\tau e^{-rt} [y(\alpha_t) - k\theta_t(\alpha_t)] dt + e^{-r\tau} U | \alpha_0 = \alpha \right].$$

Each employed worker is essentially facing an optimal stopping problem with stopping value U. As shown by Krylov (1980), this stochastic optimal control problem has a (stationary) Markovian solution.

Therefore, we can further express the value functions $V(\alpha)$ and U as:

$$rV(\alpha) = \max\{U, \max_{\theta \geq 0}\{y(\alpha) + \delta(U - V(\alpha)) + p(\theta)[V(\alpha_0) - V(\alpha)] - k\theta + \Sigma(\alpha) V''(\alpha)\}$$

and

$$rU = \max_{\theta \geq 0}\{b + p(\theta)[V(\alpha_0) - U] - k\theta\}.$$

Step 3. Fix any U. Consider any Markovian strategy $\Gamma = \{\theta(\alpha), \delta(\alpha)\}$. The expected payoff of an employed worker with prior α associated with this strategy can be written as

$$\begin{aligned} \hat{V}^\Gamma(\alpha) &= \alpha \mathbb{E}_{\mu=\mu_H} \left[\int_0^\tau e^{-rt} [y(\alpha_t) - k\theta_t(\alpha_t)] dt + e^{-r\tau} U | \alpha_0 = \alpha \right] \\ &+ (1-\alpha) \mathbb{E}_{\mu=\mu_L} \left[\int_0^\tau e^{-rt} [y(\alpha_t) - k\theta_t(\alpha_t)] dt + e^{-r\tau} U | \alpha_0 = \alpha \right]. \end{aligned}$$

Let $\alpha = \eta\alpha_1 + (1-\eta)\alpha_2$, with $\eta \in [0,1]$. For prior α_1 (resp. α_2), define Γ_1 (resp. Γ_2) to be another strategy in which the planner mistakenly believes the prior to be α and follows strategy Γ.

It is straightforward to see that

$$\begin{aligned}\hat{V}^{\Gamma}(\alpha) &= \eta\hat{V}^{\Gamma_1}(\alpha_1) + (1-\eta)\hat{V}^{\Gamma_2}(\alpha_2) \\ &\leq \eta V(\alpha_1) + (1-\eta)V(\alpha_2).\end{aligned}$$

Taking the supremum of the left-hand side with respect to Γ establishes the convexity of V.

Consider any $\alpha_1 < \alpha_2$, and suppose $\hat{V}^{\Gamma_1}(\alpha_1) = V(\alpha_1)$. For prior α_2, we can similarly define Γ_2 to be another strategy in which the planner mistakenly believes the prior to be α_1 and follows strategy Γ_1. Then, it is straightforward to see that $V(\alpha_1) = \hat{V}^{\Gamma_1}(\alpha_1) \leq \hat{V}^{\Gamma_2}(\alpha_2) \leq V(\alpha_2)$.

□

A.2 Proof of Theorem 2 and 3

Proof. We prove Theorem 2 and 3 together. The proof proceeds in four steps.

1. The first step rewrites equations (9) and (12) such that the control variable θ is in a compact set.

2. The second step constructs candidate solutions to equations (9) and (12) given an arbitrary $\bar{\alpha} < \alpha_0$.

3. The third step shows that there exists a unique $\bar{\alpha}$ satisfying the value matching and smooth pasting conditions at $\underline{\alpha}$.

4. The final step provides conditions guaranteeing that $\bar{\alpha} > \underline{\alpha}$.

Step 1. Notice that the component value function of each individual must take the value in $[\frac{b}{r}, \frac{\mu_H}{r}]$ where the lower bound is the individual rationality payoff since the planner can always leave a worker unemployed, and the upper bound is the expected value if the match is good for sure and it is never separated. Because $p''(\theta) \leq 0$ and $p(\cdot)$ is bounded, $\lim_{\theta\to\infty} p'(\theta) = 0$. Hence, the efficient θ_u^* is bounded above by some finite $\bar{\theta}$ such that $p'(\bar{\theta})[\mu_H - b] = rk$. As $V(\alpha) \geq U$, the efficient $\theta^*(\alpha)$ is also bounded above. In sum, without loss of generality, one can focus on a compact feasible set of tightness $[0, \bar{\theta}]$. Consequently, the value functions $V(\alpha)$ and U satisfy equations

$$rV(\alpha) = \max_{\theta \in [0, \bar{\theta}]} \{y(\alpha) + \rho(U - V(\alpha)) + p(\theta)[V(\alpha_0) - V(\alpha)] - k\theta + \Sigma(\alpha)V''(\alpha)\}, \tag{24}$$

for $\alpha \geq \underline{\alpha}$, and

$$rV(\alpha) = rU = \max_{\theta \in [0, \bar{\theta}]} \{b + p(\theta)[V(\alpha_0) - U] - k\theta\}. \tag{25}$$

if $\alpha < \underline{\alpha}$. If a solution $(V(\alpha), U)$ exists, $V(\alpha)$ must be strictly increasing for $\alpha \geq \underline{\alpha}$ because it cannot be a constant on any interval from equation (24). Since the planner's optimal separation decision is a stopping problem, the "stopping belief" $\underline{\alpha}$ must satisfy $V(\underline{\alpha}) = U$ (value matching) and $V'(\underline{\alpha}) = 0$ (smooth pasting) from Dixit (1993).

Step 2. We construct a candidate solution to the planner's problem. Suppose that $\theta^*(\alpha) > 0$ for some $\alpha \geq \underline{\alpha}$. For any $V(\alpha)$, the corresponding policy $\theta^*(\alpha) > 0$ only if $p'(0)\,[V(\alpha_0) - V(\alpha)] > k$. Define

$$x = \sup\{\alpha | \theta^*(\alpha') = 0, \forall \alpha' > \alpha\}.$$

For any $k > 0$, because p' is bounded and $V(\cdot)$ is continuous, $x < \alpha_0$, and by the definition of x, $p'(0)[V(\alpha_0) - V(x)] = k$.

Hence, the proof of the existence and uniqueness of $V(\alpha)$ is divided into two cases: $\alpha \in [x, 1]$ and $\alpha \in [\underline{\alpha}, x)$. When $\alpha \geq x$, $\theta^* = 0$, so (24) becomes

$$rV(\alpha) = y(\alpha) + \rho(U - V(\alpha)) + \Sigma(\alpha)V''(\alpha), \tag{26}$$

with initial value boundary condition $p'(0)[V(\alpha_0) - V(x)] = k$.

The general solution is given by

$$V_0(\alpha) = \frac{y(\alpha) + \rho U}{r + \rho} + \kappa_0 \alpha^{1-\beta_2}(1-\alpha)^{\beta_2}, \tag{27}$$

where $\beta_2 = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2(r+\rho)}{s^2}} > 0$, and κ_0 is chosen to satisfy the boundary condition

$$V_0(\alpha_0) - V_0(x) = \frac{(\mu_H - \mu_L)(\alpha_0 - x)}{r + \rho} + \kappa_0 \left[\alpha_0^{1-\beta_2}(1-\alpha_0)^{\beta_2} - x^{1-\beta_2}(1-x)^{\beta_2}\right] = \frac{k}{p'(0)}. \tag{28}$$

For a given x, there exists a unique κ_0 satisfying the boundary condition (28). As the value function V_0 is strictly increasing, $V_0'(x) > 0$, the implicit function theorem implies that $\kappa_0(x)$ is decreasing in x: $\frac{\partial \kappa_0}{\partial x} < 0$. Since $\alpha_0 > x$, $V(\alpha_0)$ is uniquely pinned down by

$$V_0(\alpha_0) = \frac{y(\alpha_0) - rU}{r + \rho} + \kappa_0(x)\alpha_0^{1-\beta_2}(1 - \alpha_0)^{\beta_2}.$$

Plugging $V_0(\alpha_0)$ into equation (25) yields

$$rU = T(U) \triangleq \max_{\theta \in [0,\bar{\theta}]} \{b + p(\theta)[\frac{y(\alpha_0) - rU}{r + \rho} + \kappa_0(x)\alpha_0^{1-\beta_2}(1 - \alpha_0)^{\beta_2}] - k\theta\}. \tag{29}$$

Obviously, $T(U)$ is strictly decreasing in U with $T(0) > 0$ and $\lim_{U\to\infty} T(U) < 0$. Therefore, there is a unique $U(x) \in [b/r, V_0(\alpha_0)]$ for each κ_0 (or x). By the envelope theorem, we obtain:

$$\frac{\partial U(x)}{\partial \kappa_0} = \frac{(r + \rho)p(\theta_u^*)}{r(r + \rho + p(\theta_u^*))}\alpha_0^{1-\beta_2}(1 - \alpha_0)^{\beta_2}. \tag{30}$$

Now we turn to the case where $\alpha < x$. The value function satisfies the ODE

$$rV(\alpha) = \max_{\theta \in [0,\bar{\theta}]} \{y(\alpha) + \rho\,(U(x) - V(\alpha)) + p(\theta)[V_0(\alpha_0) - V(\alpha)] - k\theta + \Sigma(\alpha)V''(\alpha)\}$$

with boundary conditions $V(x-) = V_0(x; x)$ and $V'(x-) = V_0'(x; x)$.[23]

The above ODE can be rewritten as:

$$V''(\alpha) = \min_{\theta \in [0,\bar{\theta}]} \frac{rV(\alpha) - y(\alpha) - \rho\,(U(x) - V(\alpha)) - p(\theta)[V_0(\alpha_0) - V(\alpha)] + k\theta}{\Sigma(\alpha)}. \tag{31}$$

Notice that

$$H_\theta(\alpha, V(\alpha)) \triangleq \frac{rV(\alpha) - y(\alpha) - \rho\,(U(x) - V(\alpha)) - p(\theta)[V_0(\alpha_0) - V(\alpha)] + k\theta}{\Sigma(\alpha)}$$

is differentiable in all of its arguments, with uniformly bounded derivatives over all $\theta \in [0, \bar{\theta}]$. Therefore, the right-hand side of equation (31) is Lipschitz-continuous. It follows that the solutions to (31) exist for $\alpha < x$ and are unique and continuous in initial conditions $V(x-)$ and $V'(x-)$. In summary, by fixing $x < \alpha_0$, we can find the unique $V(\alpha; x)$ and $U(x)$ satisfying equations (24) and (25).

Step 3. Define function $\tilde{V}(\alpha; x) = V(\alpha; x) - U(x)$. We want to show that there exists a unique x such that when $\tilde{V}(\underline{\alpha}; x) = 0$, then $\tilde{V}'(\underline{\alpha}; x) = 0$.

Lemma 1. *Consider $x_1 < x_2 < \alpha_0$. Then in the region where $\alpha < x_2$ and $\tilde{V}(\alpha; x_1), \tilde{V}(\alpha; x_2) \geq 0$, we must have $\tilde{V}(\alpha; x_1) > \tilde{V}(\alpha; x_2)$ and $\tilde{V}'(\alpha; x_1) < \tilde{V}'(\alpha; x_2)$.*

[23] The smoothness of the value function has been proven by Strulovici and Szydlowski (2014).

Proof. Since $\frac{\partial \kappa_0}{\partial x} < 0$, we have $\kappa_0(x_1) > \kappa_0(x_2)$ and hence, $V_0(x_2; x_1) > V_0(x_2; x_2)$ and $V_0'(x_2; x_1) < V_0'(x_2; x_2)$. Moreover, $V_0(\alpha_0; x_1) > V_0(\alpha_0; x_2)$ implies that $U(x_1) > U(x_2)$. Notice that

$$\tilde{V}(x_2; x) = \frac{y(x_2) - rU(x)}{r + \rho} + \kappa_0(x) x_2^{1-\beta_2}(1 - x_2)^{\beta_2},$$

which implies that

$$\frac{\partial \tilde{V}(x_2; x)}{\partial \kappa_0} = x_2^{1-\beta_2}(1 - x_2)^{\beta_2} - \frac{r}{r + \rho}\frac{\partial U(x)}{\partial \kappa_0} > x_2^{1-\beta_2}(1 - x_2)^{\beta_2} - \alpha_0^{1-\beta_2}(1 - \alpha_0)^{\beta_2} > 0.$$

The first inequality comes from equation (30); the second inequality comes from the fact that $x < \alpha_0$. Hence, we obtain $\tilde{V}(x_2; x_1) > \tilde{V}(x_2; x_2)$.

It suffices to show that $\tilde{V}'(\alpha; x_1) < \tilde{V}'(\alpha; x_2)$. Suppose not and let α' be the largest point for which

$$\tilde{V}'(\alpha'; x_1) = V'(\alpha'; x_1) = \tilde{V}'(\alpha'; x_2) = V'(\alpha'; x_2).$$

Obviously, for $\alpha_0 \geq \alpha > \alpha'$,

$$\tilde{V}'(\alpha; x_1) = V'(\alpha; x_1) < \tilde{V}'(\alpha; x_2) = V'(\alpha; x_2),$$

which implies that $V(\alpha_0; x_1) - V(\alpha'; x_1) < V(\alpha_0; x_2) - V(\alpha'; x_2)$ and $\tilde{V}(\alpha'; x_1) > \tilde{V}(\alpha'; x_2)$. From equation (31), it is straightforward to get $V''(\alpha'; x_1) > V''(\alpha'; x_2)$, which leads to a contradiction. □

As $\kappa_0 \geq 0$ (V has to be convex), $x \leq \alpha_0 - \frac{k}{p'(0)}\frac{r+\delta}{\mu_H - \mu_L}$. As we decrease x from $\alpha_0 - \frac{k}{p'(0)}\frac{r+\delta}{\mu_H - \mu_L}$, initially the resulting solutions $\tilde{V}$ must reach 0 at some point $\underline{\alpha}' > 0$, as shown by curve A in Figure 5. By Lemma (1), point $\underline{\alpha}'$ decreases as we decrease x, and $\tilde{V}'(\underline{\alpha}'; x)$ decreases as well. On the other hand, as x becomes sufficiently close to zero, $\tilde{V}(\alpha) > 0$ for all $\alpha > 0$, as shown by curve C in Figure 5. Then we can simply choose the smallest x^* such that $\tilde{V}(\underline{\alpha}^*; x^*) = 0$ for some $\underline{\alpha}^* \in (0, \alpha_0)$ (curve B in Figure 5). Because solutions with a smaller x would never reach 0, it follows that $\tilde{V}(\alpha; x^*) \geq 0$ for all $\alpha \in (0, \alpha_0)$, and hence $\tilde{V}$ has to be tangent to 0 at $\underline{\alpha}^*$: $\tilde{V}'(\underline{\alpha}'; x^*) = 0$. By construction, this is the unique x such that when $\tilde{V}(\underline{\alpha}; x) = 0$, then $\tilde{V}'(\underline{\alpha}; x) = 0$.

It is straightforward to check that the (V, U) constructed is indeed a solution to the optimal control problem following the verification theorems in Fleming and Soner (2006).

When $\alpha > \bar{\alpha}$, $\theta^*(\alpha)$ is constant, so we only show that $\theta^*(\alpha), \forall \alpha \in [\underline{\alpha}, \bar{\alpha}]$ is decreasing. At an interior solution, we have the first-order condition

$$p'(\theta^*(\alpha))[V(\alpha_0) - V(\alpha)] = k.$$

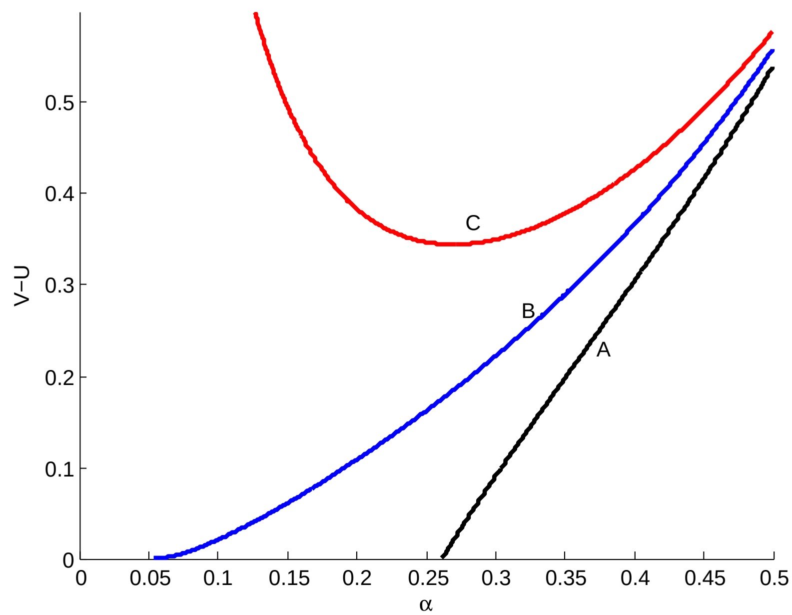

Figure 5: Solutions $\tilde{V}$ for different x.

Because $k, V(\alpha_0)$ is fixed, and $p'' \leq 0$, and $V(\alpha)$ is increasing in this domain, the desired result follows.

Step 4. Notice that the above analysis does not prelude the possibility that $\bar{\alpha} = \underline{\alpha}$ and $\theta^*(\alpha) = 0$ for all α. Suppose that this is the case. Then the value function satisfies:

$$rV(\alpha) = \max\{U, y(\alpha) + \rho(U - V(\alpha)) + \Sigma(\alpha)V''(\alpha)\},$$

and $U = \frac{b}{r}$.

There is an explicit solution to this optimal stopping problem: $\underline{\alpha} = \frac{(\beta_2-1)(b-\mu_L)}{\beta_2(\mu_H-\mu_L)+\mu_L-b}$, where $\beta_2 = \frac{1}{2} + \sqrt{\frac{1}{4} + \frac{2(r+\rho)}{s^2}} > 1$. For $\alpha > \underline{\alpha}$, the value function is

$$\hat{V}(\alpha) = \frac{y(\alpha) + \frac{\rho b}{r}}{r + \rho} + \hat{k}\alpha^{1-\beta_2}(1-\alpha)^{\beta_2},$$

where

$$\hat{k} = \frac{\mu_H - \mu_L}{r + \rho} \frac{\underline{\alpha}(1 - \underline{\alpha})}{\beta_2 - (1 - \underline{\alpha})} \underline{\alpha}^{\beta_2 - 1}(1 - \underline{\alpha})^{-\beta_2}.$$

Therefore, there exists α such that $\theta^*(\alpha) > 0$ if and only if $\hat{V}(\alpha_0) - \frac{b}{r} > \frac{k}{p'(0)}$. □

References

Acemoglu, D., and R. Shimer (1999): "Efficient unemployment insurance," *Journal of Political Economy*, 107, 893–928.

Burdett, K., and D. T. Mortensen (1998): "Wage differentials, employer size, and unemployment," *International Economic Review*, 39, 257–273.

Dixit, A. (1993): *The Art of Smooth Pasting*. Routledge, New York.

Duffie, D., and Y. Sun (2007): "Existence of independent random matching," *The Annals of Applied Probability*, 17, 386–419.

——— (2012): "The exact law of large numbers for independent random matching," *Journal of Economic Theory*, 147(3), 1105–1139.

Farber, H. S. (1999): "Mobility and stability: The dynamics of job change in labor markets," *Handbook of labor economics*, 3, 2439–2483.

Fleming, W. H., and H. M. Soner (2006): *Controlled markov processes and viscosity solutions*, vol. 25. Springer.

Fujita, S. (2012): "An empirical analysis of on-the-job search and job-to-job transitions," Discussion paper, Federal Reserve Bank Philadelphia.

Gonzalez, F. M., and S. Shi (2010): "An equilibrium theory of learning, search, and wages," *Econometrica*, 78(2), 509–537.

Hall, R. E., and A. B. Krueger (2008): "Wage formation between newly hired workers and employers: Survey evidence," Discussion paper, National Bureau of Economic Research.

Holzer, H. J., L. F. Katz, and A. B. Krueger (1991): "Job queues and wages," *The Quarterly Journal of Economics*, 106, 739–768.

Jovanovic, B. (1979): "Job matching and the theory of turnover," *Journal of Political Economy*, 87, 972–990.

——— (1984): "Matching, turnover, and unemployment," *Journal of Political Economy*, 92, 108–122.

KRYLOV, N. V. (1980): *Controlled diffusion processes*. Springer-Verlag.

LIPTSER, R. S., AND A. N. SHIRYAEV (2001): *Statistics of random processes II: Applications*. Springer.

MENZIO, G., AND S. SHI (2009): "Efficient search on the job and the business cycle," Discussion paper, National Bureau of Economic Research.

——— (2011): "Efficient search on the job and the business cycle," *Journal of Political Economy*, 119(3), 468–510.

MENZIO, G., I. A. TELYUKOVA, AND L. VISSCHERS (2012): "Directed search over the life cycle," Discussion paper, National Bureau of Economic Research.

MOEN, E. R. (1997): "Competitive search equilibrium," *Journal of Political Economy*, 105(2), 385–411.

MORTENSEN, D. T., AND C. A. PISSARIDES (1994): "Job creation and job destruction in the theory of unemployment," *Review of Economic Studies*, 61(3), 397–415.

MOSCARINI, G. (2005): "Job matching and the wage distribution," *Econometrica*, 78(2), 481–516.

PAPAGEORGIOU, T. (2014): "Learning your comparative advantages," *Review of Economic Studies*, 81, 1263–1295.

SHI, S. (2009): "Directed search for equilibrium wage–tenure contracts," *Econometrica*, 77(2), 561–584.

STRULOVICI, B., AND M. SZYDLOWSKI (2014): "On the smoothness of value functions and the existence of optimal strategies," *Available at SSRN 1996808*.