

Directed Search over the Life-Cycle

Guido Menzio
University of Pennsylvania and NBER

Irina A. Telyukova
University of California at San Diego

Ludo Visschers
U. Carlos III, Madrid
Simon Fraser University

July 2010
Preliminary

Abstract

We build a life-cycle model of directed search in the labor market, in which workers move between the states of unemployment, employment and across employers because of differences in labor market experience and in idiosyncratic productivity of different firm-worker matches. We calibrate the model to age-independent facts based on the SIPP, and find that it is able to match key life-cycle facts in the labor market, including the rates of transition between labor market states by age and the average age-earnings profile.

1 Introduction

Workers' earnings and mobility rates vary considerably over the life-cycle. The life-cycle profile of average wages, for example, is hump-shaped: average wages increase rapidly for young workers, they flatten out for middle-aged workers and they decline for older workers. It could be the case that average wages have this hump-shaped profile because the workers individual productivity is hump-shaped over the life-cycle. However, it could also be the case that average wages are hump-shaped because young workers have had little time to find high paying jobs, and older workers do not have strong incentives to invest in searching for high paying jobs. Similarly, empirical changes in the variance of the wage distribution over the life cycle could be caused by movements in the variance of the workers individual productivity (talent), but they could also be caused by changes in the variance of the outcome of workers search (luck). Most likely, both changes in worker-embodied productivity and the success of search play a role in the life-cycle movements in the average wages and in the variance of the wage distribution. However, for many issues, among them the design of the majority of labor market policies, it is necessary to understand *how much* of the life-cycle experiences of workers are shaped by inherent talent or by (mis)fortune.

To understand the role and interplay of luck and talent, we build a life-cycle model of directed search in the labor market, in which workers move between the states of unemployment, employ-

ment and across employers because of differences in labor market experience and in idiosyncratic productivity of different firm-worker matches. On one side of the market, firms choose how many and what type of vacancies to create. The type of a vacancy is given by the value that it offers to a worker. On the other side of the market, workers choose which type of vacancies to search. Workers are different because they have different individual productivity, because they have different age, and because they are in different employment states (unemployment, and employment in different matches). Different workers have a different trade-off between the probability of finding a new job and the value offered by the job. For this reason, different workers choose to search for different types of vacancies and display different job finding rates and different wage dynamics. In turn, the heterogeneity in job finding rates and wage dynamics creates a wealth of predictions about the life-cycle profile of workers transition rates across employment states, about the life-cycle profile of the wage distribution, and about the life-cycle profile of labor productivity. To quantify these predictions, we calibrate the parameters of the model to match some basic features of the US labor market (as measured in the SIPP). In particular, we calibrate the model to match the average transition rate from unemployment to employment (UE rate), the average transition rate from employment to unemployment (the EU rate), and the average transition rate from one employer to another (the EE rate). We use data on the tenure and wage distribution to calibrate the extent of worker and match-specific heterogeneity. Likewise, the data on the unconditional distribution of job tenure and the average wage difference of young and old workers imply the extent of worker- and match-specific heterogeneity in the model.

We find, first, that the calibrated model is able to generate the same patterns of workers transition rates that we observe in the data. Like in the data, the model generates an unemployment-to-employment transition rate (UE rate) that is approximately constant for workers of age between 25 and 55, and sharply declining for older workers. Like in the data, the model generates an employment-to-unemployment transition rate (EU rate) that declines throughout the life-cycle. Like in the data, the model generates a transition rate from employer to employer (EE rate) that declines with age. The mechanisms for these results are intuitive. The UE rate declines late in life when the value of new matches begins to fall relative to the cost of creating vacancies. The EU and the EE rates fall throughout the life-cycle because workers select into better and better matches as they become older.

Second, we use the model to compute average labor productivity and average wages conditional on age. We find that the earnings have a concave profile in age. In the early part of the work life, average earnings are steeply increasing due to the large returns to search. As the worker ages,

the growth rate of average earnings falls because of the fall in the return to search, until earnings become nearly flat. This last effect is created because as the worker ages, although he is becoming more productive, his value to the firm begins to decline as he approaches retirement and the firm expects him to retire.

We will next use the model to study the variance of (log) labor productivity and (log) wages conditional on age, to understand how much of the wage inequality of workers is due to differences in productivity, and how much is due to search frictions. This decomposition will help answer the question we posed above. In addition, our work can be seen from a different perspective. In the literature devoted to studying household saving and borrowing, and in particular the extent of precautionary saving and risk-sharing, a crucial input for analyzing working-age households is idiosyncratic uncertainty in income. There is extensive effort that has been and continues to be devoted to measuring the extent of this idiosyncratic risk in the data based on observed income fluctuations. This paper can be seen as a *model-based* approach to answering the question of how much of the lifecycle fluctuations in earnings are due to the workers' endogenous decisions, as opposed to exogenous events that are unpredictable to the worker.

The paper is organized as follows. Section 2 discusses related literature. We present the model in section 3, show key data facts in section 4, and discuss our calibration strategy in section 5. Section 6 presents the results of the model, section 7 discusses some comparative statics, and section 8 concludes.

2 Literature

There are relatively few papers which study the impact of search frictions in a labor market *equilibrium* when working lives are finite. Most related to this paper is Hairault et al. (2007), who show that in a Pissarides model with forced retirement in finite time, employment rates and hiring rates drop before the end of the working life. Interestingly, they document that in a cross-country comparison, the drop in employment rates occurs sooner the earlier is the retirement time – supporting an important role for search frictions, since the finite horizon appears to affect directly the asset value of a filled vacancy. In their paper, in contrast to ours, all workers of a given age with a job are essentially identical (apart from an i.i.d. productivity shock triggering a match breakup). Moreover, they do not quantitatively investigate the joint process of wages and transition rates over the life cycle, and abstract from human capital accumulation in their quantitative analysis, focussing instead on the role of firing taxes and hiring subsidies over the life cycle. In a number of

follow-up papers, Cheron et al. (2008a) and Cheron et al. (2008b) investigate, respectively, the case where the matching function is the same for all ages, and the effects of age-dependent employment protection.

In a partial-equilibrium or decision-theoretic framework, the effects of the finiteness of working lives when the labor market is frictional are studied more extensively, going back at least to Saeter (1977). In this category, a recent paper is Liu (2009), who estimates a decision-theoretic model of the life cycle of workers in a frictional market with job-to-job transitions and endogenous human capital accumulation. The wage distribution in this model is taken exogenously: the focus is on the search and human capital investment choices of workers.¹

There are many papers that focus on the experience of young workers in the labor market, of which we mention only a few. The role of job mobility in income improvement at young age, a channel that occurs in our model as well, has been studied, for example, in well-known papers by Topel and Ward (1992) and Keane and Wolpin (1997). This paper extends the study of the importance of job mobility for individual earning gains to the entire life cycle.² Understanding labor market behavior at different ages is both important at the individual level and at the aggregate level: changes in the age structure of the population have important aggregate implications, e.g. for the unemployment rate, as shown e.g. in Shimer (1998) and Abraham and Shimer (2002).

This paper relates also to a recent literature which studies the interaction of human capital and search frictions. Again, this is mostly done in models with infinitely-lived agents. Barlevy (2008), Burdett et al. (2010), Burdett and Coles (2010), Carrillo-Tudela (2010) and Fu (2009) study the interaction between on-the-job search with wage posting and human capital accumulation. In this context one can decompose wage growth into human capital growth and improvements from job mobility. In the equilibrium both forces have important and interesting interactions, for example increasing the amount of wage dispersion. Bagger et al. (2007) study human capital in the setting of Postel-Vinay and Robin (2002), where on-the-job search triggers Bertrand wage competition whenever the worker has a possibility to move to a new firm.

Since one of the goals of our paper is to get more insight on the role of search frictions in the (stochastic) evolution of earnings at the level of *individual* workers, the paper relates to the extensive literature that attempts to estimate the stochastic processes underlying individual-level

¹Another recent example is Hoffman (2009) who estimates a dynamic Roy model, enriched with a partial-equilibrium search setting, over the life cycle.

²Theoretical models of early life experience often have employed infinitely-lived agents, as the impact of the finite horizon goes to insignificance at early age. The experiences of young workers are shaped then by their initial conditions, while for (somewhat) older workers these conditions gradually wash out.

earnings histories. Recent examples of papers in this vast literature are Meghir and Pistaferri (2004) and Low, Meghir and Pistaferri (2009). Like this paper, there are a few papers which try to bridge the gap between the search literature and the empirical earning dynamics literature, notably Bagger et al. (2007) and Postel-Vinay and Turon (2010), who add productivity shocks within existing employment relations.

Finally, we solve for the equilibrium on the labor market, which helps us infer how much of the innovations to wages are expected. Thus, this paper relates broadly to a set of papers that attempt to infer how much unpredictability there is in wages. Besides using restrictions on behavior from the equilibrium, one can also employ restrictions from joint optimization to identify unexpected income shocks. An example of this is Guvenen and Smith (2009), which uses consumption decisions in the face of volatile income to identify such shocks.

3 Environment and Equilibrium

3.1 Market economy

The labor market is populated by a continuum of workers with measure one. Workers are finitely lived. A worker enters the labor market at age $t = 1$ with productivity y_1 , where y_1 is a random variable drawn from the distribution $h(y)$. At age $t = 2, 3, \dots, T$, the worker participates to the labor market with productivity $y_t = y_1 \rho^{e_t}$, where e_t is the worker's experience (i.e., the number of periods during which the worker was employed) and $\rho \geq 1$ is a parameter that determines the return to experience. At age $t = T + 1$, the worker exits the labor market and is replaced by a worker of age $t = 1$. Each worker is endowed with one indivisible unit of labor and maximizes the expected sum of periodical consumption discounted at the factor $\beta \in (0, 1)$.

The labor market is also populated by a continuum of firms with positive measure. Each firm operates a constant return to scale technology that turns a unit of labor into yz units of output, where y is the component of productivity that is specific to the worker, and z is the component of productivity that is specific to the firm-worker match. In the remainder of the paper, we shall refer to y as the worker-specific component of productivity and to z as the match-specific component of productivity. Each firm maximizes the expected sum of periodical profits discounted at the factor β .

The labor market is organized in a continuum of submarkets indexed by (x, y, t) . When a firm meets a worker with productivity y and age t in submarket (x, y, t) , it offers him an employment contract that is worth x in lifetime utility. When a firm meets a worker with productivity and age

$(y', t') \neq (y, t)$ in submarket (x, y, t) , it does not offer him an employment contract. In other words, a submarket is indexed by the productivity and age of the workers that the firm is willing to hire and by the value offered to them conditional on being hired. In submarket (x, y, t) , the ratio of the number of vacancies created by firms to the number of workers who are looking for jobs is given by the tightness $\theta_t(x, y) \in \mathbb{R}_+$. In equilibrium, $\theta_t(x, y)$ will be consistent with the firms' and workers' search decisions.

At the beginning of each period, the state of the economy can be summarized by the couple $\{u_t, g_t\}_{t=1}^T \equiv \psi$. The first element of ψ is a list of functions $\{u_t\}_{t=1}^T$, $u_t : Y \rightarrow [0, 1]$, with $u_t(y)$ denoting the measure of unemployed workers of age t and productivity y . The second element of ψ is a list of functions $\{g_t\}_{t=1}^T$, $g_t : Y \times Z \rightarrow [0, 1]$, with $g_t(y, z)$ denoting the measure of workers of age t and productivity y who are employed in jobs with match-specific productivity z .

Each period is divided into four stages: separation, search, matching and production. During the separation stage, an employed worker becomes unemployed with probability $d \in [\delta, 1]$, where $\delta \in (0, 1)$ is the probability that the worker has to leave his job for exogenous reasons.

During the search stage, a worker gets the opportunity to search for a job with a probability that depends on his recent employment history. In particular, if the worker was unemployed at the beginning of the period, he gets the opportunity to 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 gets the opportunity to search with probability $\lambda_e \in [0, 1]$. If the worker lost his job during the separation stage, he cannot search. Conditional on being able to search, the worker chooses which submarket to visit. In this sense, search is directed. Also, during the search stage, a firm chooses how many vacancies to create and where to locate them. The cost of maintaining a vacancy for one period is $k > 0$. Both workers and firms take the tightness $\theta_t(x, y)$ parametrically.

During the matching stage, the workers and the firms in submarket (x, y, t) come together through a frictional meeting process. In particular, a worker meets a vacancy with probability $p(\theta_t(x, y))$, where $p : \mathbb{R}_+ \rightarrow [0, 1]$ is a twice-continuously differentiable, strictly increasing, strictly concave function such that $p(0) = 0$ and $p(\infty) = 1$. Similarly, a vacancy meets a worker with probability $q(\theta_t(x, y))$, where $q : \mathbb{R}_+ \rightarrow [0, 1]$ is a twice-continuously differentiable, strictly decreasing function such that $q(\theta) = p(\theta)/\theta$, $q(0) = 1$ and $q(\infty) = 0$. When a vacancy meets a worker of age t and productivity y , it offers him an employment contract worth x in lifetime utility. If the worker rejects the offer, he returns to his previous employment position (unemployment, or employment at some other firm). If the worker accepts the offer, the two parties form a new match. Only after the match is formed and the worker has forgone the option of returning to his previous employment

position, the match-specific component of productivity z is realized. The productivity z is drawn from the distribution $f(z)$ and remains constant throughout the duration of the match.

During the last stage, an unemployed worker produces and consumes $b > 0$ units of output, independently of his productivity y . A worker employed in a match of type z , produces yz units of output and consumes w of them, where y is the worker's productivity and w is the wage prescribed by the worker's employment contract.

We assume that employment contracts are complete, in the sense that they specify the wage paid to the worker, w , the probability of separation, d , and the submarket where the worker applies for a new job, x_e , as a function of the worker's age, t , the worker-specific component of productivity, y , and the match-specific component of productivity, z . Given that contracts are complete, it is easy to show that any optimal employment contract is bilaterally efficient, in the sense that it maximizes the sum of the worker's lifetime utility and the firm's profits from the match. Given that contracts are complete and utility is perfectly transferrable, it is also easy to show that the wage w is not uniquely pinned down. In fact, the firm and the worker only care about the present value of wages, not about the particular profile that delivers such a value. In the quantitative section of the paper, we assume that the wage paid by the firm to the worker is a constant fraction of the output of the match.

3.2 The problem of the worker and the firm

First, consider an unemployed worker at the beginning of the production stage. Let $U_t(y, \psi)$ denote the worker's lifetime utility, given that his age is t and his productivity is y . In the current period, the worker produces and consumes b units of output. In the next period, the worker exits the labor market if $t = T$. Otherwise, the worker finds a new job with probability $\lambda_u \theta_{t+1}(x, y, \hat{\psi})$ and he remains unemployed with probability $1 - \lambda_u \theta_{t+1}(x, y, \hat{\psi})$, where $(x, y, t + 1)$ is the submarket where the worker chooses to search and $\hat{\psi}$ is next period's aggregate state. In the first case, the worker's continuation utility is x . In the second case, the worker's continuation utility is $U_{t+1}(y, \hat{\psi})$. Thus, $U_t(y, \psi)$ is

$$U_t(y, \psi) = b + \beta \mathbb{I}[t < T] \left[U_{t+1}(y, \hat{\psi}) + \lambda_u \max_x D_{t+1}(x, y, U_{t+1}(y, \hat{\psi}), \hat{\psi}) \right], \quad (1)$$

where D_t is defined as

$$D_t(x, y, v, \psi) = p(\theta_t(x, y, \psi))(x - v). \quad (2)$$

We denote as $x_{u,t+1}(y, \hat{\psi})$ the policy function associated to (1).

Second, consider a worker and a firm who are matched at the beginning of the production stage. Let $V_t(y, z, \psi)$ denote the sum of the worker's lifetime utility and the firm's lifetime profits, given that the worker's age is t , the worker-specific component of productivity is y , the match-specific component of productivity is z , and 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, yz . In the next period, the worker exits the labor market if $t = T$. Otherwise, the worker moves into unemployment with probability d , he moves to another firm with probability $(1 - d)\lambda_e p(\theta_{t+1}(x, \rho y, \hat{\psi}))$, and he remains with the firm with probability $(1 - d) \cdot (1 - \lambda_e p(\theta_{t+1}(x, \rho y, \hat{\psi})))$, where $(x, \rho y, t + 1)$ is the submarket where the worker searches for a new job. In the first case, the worker's continuation utility is $U_{t+1}(\rho y, \hat{\psi})$ and the firm's continuation profit is zero. In the second case, the worker's continuation utility is x and the firm's continuation profit is zero. In the third case, the sum of the worker's continuation utility and the firm's continuation profit is $V_{t+1}(\rho y, z, \hat{\psi})$. Thus, $V_t(y, z, \psi)$ is

$$V_t(y, z, \psi) = yz + \beta \mathbb{I}[t < T] \max_{1 \geq d \geq \delta} \left\{ \begin{aligned} & dU_{t+1}(\rho y, \hat{\psi}) + (1 - d)V_{t+1}(\rho y, z, \hat{\psi}) \\ & + (1 - d)\lambda_e \max_x D_{t+1}(x, \rho y, V_{t+1}(\rho y, z, \hat{\psi}), \hat{\psi}) \end{aligned} \right\}, \quad (3)$$

where D_t is the function defined in (2). We denote as $d_{t+1}(\rho y, z, \hat{\psi})$ and $x_{e,t+1}(\rho y, z, \hat{\psi})$ the policy functions associated with (3).

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, y, t) is k . The firm's benefit from creating a vacancy in submarket (x, y, t) is

$$q(\theta_t(x, y, \psi)) \{E_{z'} [V_t(y, z', \psi)] - x\}, \quad (4)$$

where $q(\theta_t(x, y, \psi))$ is the probability of meeting a worker, $E_{z'} [V_t(y, z', \psi)]$ is the joint value of the match, 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 vacancies in submarket (x, y, t) . When the cost is strictly smaller than the benefit, the firm creates infinitely many vacancies in submarket (x, y, t) . And when the cost and the benefit are equal, the firm's profit is independent from the number of vacancies it creates in submarket (x, y, t) .

In any submarket that is visited by a positive number of workers, the tightness $\theta_t(x, y, \psi)$ is consistent with the firm's incentives to create vacancies if and only if

$$k \geq q(\theta_t(x, y, \psi)) \{E_{z'} [V_t(y, z', \psi)] - x\}, \quad (5)$$

and $\theta_t(x, y, \psi) \geq 0$ with complementary slackness. In any submarket that workers do not visit, the tightness $\theta_t(x, y, \psi)$ is consistent with the firm's incentives to create vacancies if and only if k is greater or equal than (4). However, following the literature on directed search on the job (i.e., Shi 2009 and Menzio and Shi 2008, 2010 a), we restrict attention to equilibria in which $\theta_t(x, y, \psi)$ satisfies the above complementary slackness condition (5) in every submarket.

3.3 Equilibrium, block recursivity and efficiency

Definition 1: A Block Recursive Equilibrium (BRE) consists of a list of market tightness functions $\{\theta_t\}_{t=1}^T$, $\theta_t : \mathbb{R} \times Y \rightarrow \mathbb{R}_+$, value functions for unemployed workers $\{U_t\}_{t=1}^T$, $U_t : Y \rightarrow \mathbb{R}$, policy functions for unemployed workers $\{x_{u,t}\}_{t=1}^T$, $x_{u,t} : Y \rightarrow \mathbb{R}$, joint value functions for the firm-worker match $\{V_t\}_{t=1}^T$, $V_t : Y \times Z \rightarrow \mathbb{R}$, and policy functions for the firm-worker match $\{d_t, x_{e,t}\}_{t=1}^T$, $d_t : Y \times Z \rightarrow [\delta, 1]$ and $x_{e,t} : Y \times Z \rightarrow \mathbb{R}$. For $t = 1, 2, \dots, T$, these functions satisfy the following conditions:

- (i) $U_t(y)$ satisfies (1) for all $(y, \psi) \in Y \times \Psi$, and $x_{u,t}(y)$ is the associated policy function;
- (ii) $V_t(z, y)$ satisfies (3) for all $(z, y, \psi) \in Z \times Y \times \Psi$, and $\{d_t, x_{e,t}\}$ are the associated policy functions;
- (iii) $\theta_t(x, y)$ satisfies (5) for all $(x, y, \psi) \in Z \times Y \times \Psi$.

Condition (i) guarantees that the search strategy of an unemployed worker maximizes his life-time utility, given the 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. Overall, conditions (i)-(iii) guarantee that in a BRE, just like in a recursive equilibrium, the strategies of each agent are optimal given the strategies of others. However, unlike in a recursive equilibrium, the agent's value and policy functions do not depend on the distribution of workers across employment states, ψ . For this reason, a BRE describes the equilibrium of the economy not only in the steady state, but also along the transition towards the steady-state.

However, why should a BRE exist? And why should we focus on a BRE rather than on other recursive equilibria? The next theorem answers these questions. Specifically, the theorem establishes that a BRE exists and that there is no loss in generality in focusing on the BRE because all equilibria are block recursive.

Theorem 1: The unique equilibrium is block recursive.

Proof. In Appendix A.

By inspecting equations (1), (3) and (5), one can immediately establish that the equilibrium is block recursive. In fact, from equations (1) and (3), it follows that the value of unemployment U_T and the joint value of a match V_T do not depend on the distribution of workers across employment states. This implies that the value to the firm from filling a vacancy in submarket (x, y, T) does not depend on the distribution of workers and, because of the free-entry condition (5), neither does the equilibrium market tightness θ_T . In turn, since θ_T is independent of the distribution of workers, it follows from equations (1) and (3) that the value of unemployment U_{T-1} and the joint value of a match V_{T-1} are also independent of the distribution of workers. The same argument can then be used to establish that θ_{T-1} , U_{T-2} , V_{T-2} ..., θ_1 , U_1 , V_1 are all independent of the distribution of workers across employment states.

The equilibrium is block recursive because the search process is directed. When search is directed, workers only apply for jobs that they intend to accept. Therefore, a firm that opens a vacancy in submarket (x, y, t) knows that it will exclusively meet workers of type (y, t) who are willing to fill the vacancy. For this reason, the distribution of workers across different employment states does not affect the firm's value from meeting a worker in submarket (x, y, t) . In turn, because of the free entry condition (5), the distribution of workers across employment states does not affect the probability that a firm meets a worker in submarket (x, y, t) . Finally, since the meeting probability across different submarkets is independent from the distribution of workers, it follows immediately from (1) and (3) that the value of unemployment and the joint value of a match are also independent from the distribution of workers.

The next theorem analyzes the welfare properties of the block recursive equilibrium. Specifically, the theorem establishes that the block recursive equilibrium is efficient, in the sense that it decentralizes the solution to the problem of a utilitarian social planner.

Theorem 2: The Block Recursive Equilibrium is efficient.

Proof. In Appendix B.

It is important to clarify that the assumption of complete contracts is not necessary for establishing the existence of a block recursive equilibrium. In fact, Menzio and Shi (2010a,b) show that a block recursive equilibrium exists also in economies where the contract space is incomplete (e.g. employment contracts can only specify a wage that remains constant throughout the entire duration of the employment relationship). However, we make use of the assumption of complete contracts in order to establish the efficiency of the equilibrium.

3.4 Characterization of equilibrium

From the equilibrium condition (iii), it follows that the market tightness $\theta_t(x, y)$ is

$$\theta_t(x, y) = \begin{cases} q^{-1} [k / (E_{z'} [V_t(y, z')] - x)] & \text{if } x \leq E_{z'} [V_t(y, z')] - k, \\ 0 & \text{else.} \end{cases} \quad (6)$$

The market tightness $\theta_t(x, y)$ is decreasing in x . This property is intuitive. The profits of the firm from filling a vacancy, $E_{z'} [V_t(y, z')] - x$, are decreasing in the value offered by the vacancy to the workers, x . Since firms must be indifferent between opening vacancies that offer different values to the workers, the probability that the firm fills a vacancy, $q(\theta_t(x, y))$, must be increasing in x . Hence, $\theta_t(x, y)$ must be decreasing in x . Moreover, given our calibration of the model, the market tightness $\theta_t(x, y)$ is increasing in y and decreasing in t . These properties are also intuitive. The profits of the firm from filling a vacancy, $E_{z'} [V_t(y, z')] - x$, are increasing in the productivity, y , and decreasing in the age, t , of the workers who apply for it. Since firms must be indifferent between opening vacancies for different types of applicants, the probability that a firm fills a vacancy, $q(\theta_t(x, y))$, must be decreasing in y and increasing in t . In turn, this implies that $\theta_t(x, y)$ must be increasing in y and decreasing in t .

From the equilibrium conditions (i) and (ii), it follows that the search strategy of the worker is the solution to

$$\max_x p(\theta_t(x, y)) (x - v), \quad (7)$$

where t is the worker's age, y is the worker's idiosyncratic component of productivity, and v is the value of the worker's current employment position (i.e. $U_t(y)$ if the worker is unemployed, and $V_t(y, z)$ if the worker is employed in a match of quality z). After solving (6) with respect to x and substituting the solution into (7), we can rewrite the search problem of the worker as

$$\max_{\theta \geq 0} p(\theta) \{E_{z'} [V_t(y, z')] - v\} - k\theta. \quad (8)$$

The worker chooses the tightness, θ , of the submarket in which he searches for a new match in order to maximize the sum of two terms. The first term is the product between the probability that the worker finds a new match, $p(\theta)$, and the difference between the joint value of a new match to the worker and a new employer and the value of the worker's current employment position, $E_{z'} [V_t(y, z')] - v$. The second term is the cost of creating θ vacancies. The optimality condition for θ is

$$k \geq p'(\theta) \{E_{z'} [V_t(y, z')] - v\}, \quad (9)$$

and $\theta \geq 0$ with complementary slackness. Notice that, since p is a strictly concave function, the optimal market tightness is increasing in the difference between the joint value of a new match and the value of the worker's current employment position, $E_{z'} [V_t(y, z')] - v$.

The difference between the value of a new match and the value of unemployment is $E_{z'} [V_t(y, z')] - U_t(y)$. For our calibration of the model, we find that this difference is increasing in the worker's idiosyncratic component of productivity, y , and decreasing in the worker's age, t . Hence, the tightness of the submarket visited by an unemployed worker, $\theta_{u,t}(y)$, and the job-finding probability of an unemployed worker, $p(\theta_{u,t}(y))$, are increasing in y and decreasing in t .

The difference between the value of a new match and the value of an existing match is $E_{z'} [V_t(y, z')] - V_t(y, z)$. First, we find that this difference is decreasing in the quality of the worker's existing match, z . Hence, the tightness of the submarket visited by an employed worker, $\theta_{e,t}(y, z)$, and the job-finding probability of an employed worker, $p(\theta_{e,t}(y, z))$, are decreasing in z . Second, we find that the difference $E_{z'} [V_t(y, z')] - V_t(y, z)$ is increasing in the worker's specific component of productivity y for low values of z , and decreasing for high values of z . Hence, $\theta_{e,t}(y, z)$, and $p(\theta_{e,t}(y, z))$ are increasing in y for low values of z , and decreasing for high values of z . Similarly, we find that the difference $E_{z'} [V_t(y, z')] - V_t(y, z)$ is decreasing in the worker's age t for low values of z , and increasing for high values of z . Hence, $\theta_{e,t}(y, z)$, and $p(\theta_{e,t}(y, z))$ are decreasing in t for low values of z , and increasing for high values of z .

From the equilibrium condition (iii), it follows that the separation strategy is $d_t(y, z) = 1$ if

$$U_t(y) > V_t(y, z) + \lambda_e \max_x D_t(x, y, V_t(y, z)), \quad (10)$$

and $d_t(y, z) = \delta$ otherwise. The left-hand side of (10) is the value of unemployment to the worker, given that the worker does not have the option to look for a new match in the current period. This is the joint value to the firm and the worker from destroying the match. The right-hand side of (10) is the value of the match, given that the worker has the opportunity to search for a new match with probability λ_e . This is the joint value to the firm and the worker from maintaining the match. If the left-hand side is greater than the right-hand side, the worker and the firm find it optimal to destroy the match with probability 1. Otherwise, Nature destroys the match with probability δ . Note that the left-hand side does not depend on z , while the right hand side is strictly increasing in z . Hence, the optimal separation strategy can be represented by a reservation productivity $z_{d,t}(y)$ such that $d_t(y, z) = 1$ for all $z < z_{d,t}(y)$, and $d_t(y, z) = \delta$ for all $z \geq z_{d,t}(y)$. In our calibration, we find that $z_{d,t}(y)$ is decreasing in y and increasing in t . Intuitively, the higher

is the worker-specific component of productivity, y , the higher is the value of the match relative to the value of unemployment and, hence, the lower is the realization of z for which the firm and the worker find it optimal to separate. Similarly, the older is the worker, the lower is the value of the match relative to the value of unemployment and, hence, the higher is the realization of z for which the firm and the worker find it optimal to separate.

4 Data

We use the Survey of Income and Program Participation³ (SIPP) to investigate the patterns of wages and job mobility over the life cycle. The SIPP is a longitudinal survey that follows the same individuals for periods up to four year, collecting data on e.g. income, employment status, assets, and health and health insurance. In addition we use Panel Survey of Income Dynamics (PSID) to verify these patterns on data that follows subjects over their entire working life.

The SIPP consists of multiple panels: here, we use the 1996 and 2001 panel. Households are assigned to one of four ‘rotation groups’ and interviewed once every four months, for the duration of the panel. Different rotation groups are interviewed in different months, in a staggered pattern. When all rotation groups are interviewed, a ‘wave’ of four months is completed, and a new wave is started with a new interview of the first rotation group. In an interview household members are asked questions about their labor market status in each of the previous four months. For some questions the retrospective precision is at weekly level (see below).

The SIPP is not a simple random sample of the US population, and as result, we have to take the stratified sampling design into account when calculating moments like means and variances. As an important part of the sampling design, for example, the SIPP oversamples lower-income households. We use the appropriate weights and variance calculation methods to make the statistics below relevant for the US population as a whole.⁴

Moreover, the US population, and the sampled individuals are not spread uniformly across ages 18-66. In fact, the population pyramid for the 1996 SIPP, for example, even using the sample weights, shows a sharp decline with age in the number of men above age 40. In our calculations, we correct for this by computing statistics conditional on each age, and taking the unconditional average of these statistics. We also adjust the weighting scheme accordingly. This, among other

³www.census.gov/sipp

⁴To calculate the transition probabilities between labor market states, we use the SIPP as a set of repeated cross-sections. However, to get a measure of income variability, we exploit the panel structure of the data. These two different approaches require the use of different weights, as not every household or individual is in the panel for the whole time.

things, allows us to map the data numbers into the model, whose population structure does not have the pyramid shape. The data appendix has more details on our adjustment to the weights.

In the analysis below, we concentrate on men between 18 years old and 66 years old, who do not have their own business. This leaves about 30,000 individuals in a panel. In each interview, the SIPP asks, for each week in the relevant 4-month period, whether the worker was working in a job, absent without pay, or looking for a job. Any worker who did not have a job in all weeks is asked, in addition, if he was available to work, thus giving a clear idea whom to classify as unemployed as opposed to a nonparticipant in the labor force. Additional questions are asked to get an idea about the reasons behind nonparticipation - e.g. illness, temporary layoff, discouragement, retirement. These reasons are important to know, as they allow us to separate exogenous causes of nonparticipation (e.g. illness) from endogenous causes (like a very low arrival rate of offers).

For the flow measures that we report below, we first need to define the states of unemployment and employment. Our model is monthly, and thus out of the weekly data, we have to construct a monthly measure. The way we approach this is to take the observations of labor market status in the first week of each month. We do not use SIPP's monthly labor market status recode, because it can classify a worker in multiple states in a given month, which our model does not allow. In the benchmark calibration, we just use the first week's observation and are agnostic about what happens in the intervening weeks. This is a close mapping of the data to the model. We conduct robustness analysis by looking at the intervening weeks more closely and even by constructing weekly transition measures.

The lifecycle flow patterns reported below, as well as the benchmark calibration, are based on the following definitions of labor market states. We classify a worker as employed in a given month, if in the first week of that month he reports having a job, and being either present and working, or absent without pay, either on layoff or not. We classify a worker as unemployed if in the first week he reports having no job, but looking for work actively or being on layoff. We classify a worker out of the labor force if in the first week he reports having no job, and not actively looking nor being on layoff.

In robustness analysis, we have also tried the following alternative measures. First, we re-classified workers who report being employed in the current first week, and the first week of the following month, but unemployed in weeks 2, 3, or 4 of the current month, as unemployed in the current month. This includes in the unemployment measure people with short unemployment spells. Second, we re-classified workers who reported having a job but being on layoff as unemployed




Figure 1: Unemployment-Employment Transition Rates by Age, Data (SIPP)

in the current month. We report some of the results of this robustness analysis below.⁵

4.1 Flows Within the Labor Force

Figure 1 shows the monthly transition rates, by age, between unemployment and employment (UE). The overall shape of the transition profile is downward-sloping, from around 20% for young workers to approximately 12% for older ones, but for most of the lifecycle, it is fairly flat, declining only slightly over time. Starting at age 50, however, there is a noticeable decline in the transition rates: older workers, as they approach retirement, are significantly less likely to find jobs out of unemployment.

The transition profile for flows between jobs and unemployment (EU) has a pronounced downward slope, as pictured in figure 2. Young workers are the most likely to transition into unemployment from a job, at the rate of 2.5% per month, but as they age into their mid-30's, that rate drops to 0.08% and the profile flattens out, and the transition rates continue to decline through retirement age, where the rate of transition is close to zero.

Figure 3 documents the average monthly transition rates between jobs (EE), by age. Similarly to the EU flow, young workers transition relatively frequently (at the rate of 4.5% per month), but as they age, their likelihood of switching between jobs declines rapidly. The transition rate continues to decline for the remainder of the working life, reaching just above 2% monthly around

⁵We have also measured weekly transition rates between employment states, based on the weekly data. We do not use these in the model, as our model is monthly, but ascertained that the lifecycle profiles we report below are robust to looking at weekly data.




Figure 2: Employment-Unemployment Transition Rates by Age, Data (SIPP)

age 30, and dropping below 1% by age 65.

4.2 Robustness: Flows Within the Labor Force, Alternative Definitions

Figure 4 shows the smoothed UE profiles using three different definitions of unemployment (and accordingly, employment). The solid line refers to the benchmark case, same as above, while “robust1” (dashed line) refers to the measure of unemployment capturing short unemployment spells, and “robust2” (dot-dashed line) labels those who report being with job but on layoff as unemployed as well, as described above. The last measure is the broadest definition of unemployment that we try. With the first robustness measure, the UE transition rate increases for all ages, from around 20 to 24% in at age 23, and from 17 to 18.7% at age 65. The profile becomes more downward-sloping as well; however, the spirit of the lifecycle profile remains very similar. The second robustness measure increases the rate of transition further: to nearly 26% at age 23, 18% at age 61, and 25% at age 65. What is remarkable in this case is that the lifecycle picture is overall much flatter, thanks to a dramatic U-shape from age 50 onward. These two measures suggest that one can count many UE transitions by capturing unemployment spells that are shorter than a month (possibly only a week), and that among oldest ages, there is relatively pronounced movement in and out of layoffs.

The robustness exercise on the EU profile is in figure 5. Once again, we see that the two alternative definitions of unemployment lead to many more counts of EU transitions. In this case, the shape of the lifecycle profile is not as dramatically affected, but it does become steeper with the first alternative definition, and flatter with the second. The levels are again affected: at age




Figure 3: Employment-employment Transition Rates by Age, Data (SIPP)




Figure 4: UE Transitions, Alternative Definitions, Data (SIPP)

23, the transition rate increases from 1.3% monthly to nearly 1.8%, and to nearly 2% with the second definition. Toward the end of the lifecycle, the first two measures produce nearly identical EU rates, while the third diverges again with a mild U-shape. The interpretation is similar to above: we record more EU transitions if we count more people as unemployed by more loose definitions, involving the short unemployment spells and layoffs.

Finally, we computed the EE profiles corresponding to the alternative definitions of employment,




Figure 5: EU Transitions, Alternative Definitions, Data (SIPP)

shown in figure 6. This profile is least sensitive to the changes, but the measure that does not count as continually employed those who have an unemployment spell after the first week of the current month produces slightly lower transition rates for younger ages (dashed line).




Figure 6: EE Transitions, Alternative Definitions, Data (SIPP)

Because layoffs do not explicitly exist in our model, and because our model is monthly – so that we are agnostic about what happens between two consecutive reference points in a month –

the baseline measures of the labor force states are the closest to the model in spirit, and thus our calibration is to this measure.

4.3 Wages

The age-earnings profile from SIPP is plotted in figure 7. The profile has a slight hump shape, which peaks around age of 60. For the calibration, we average the profiles across cohorts, and target the ratio of the average wage at age 55 and the average wage at age 30 for this average profile.




Figure 7: Average Cohort-Specific Hourly Earnings Profile by Age, Data (SIPP)

5 Calibration

The first goal of our exercise is to see whether our model can reproduce the life-cycle facts described above, based on a calibration that does not target the life cycle explicitly. We set the model period to a month. We estimate the key parameters of the model - λ_e , k , b , dispersion of the log-normal distribution of z , δ , ρ , and the spread of the log-normal distribution of idiosyncratic productivity σ_{y_0} - by the Simulated Method of Moments. We set the discount factor β to a value consistent with the real interest rate of 4% on an annual basis.

From the data already described, we target the rates of UE, EE and EU transitions *averaged across age*. We also target the ratio of average earnings of 55-year-olds to average earnings of 30-year-olds. The additional targets are described below.

5.1 Tenure Distribution

Figure 8 plots the tenure distribution from the SIPP, that is, the distribution of workers by the length of their current job, as well as the log-distribution. We choose six points from this distribution to give us six additional calibration targets: those who have been in their current job under 3 months, 6 months, 1 year, 2 years, 5 years and 10 years. This distribution is an unconditional average of all tenure distributions calculated for each age, and as such is ‘age-free’.




Figure 8: Job Tenure Distributions: Density and Log-Density (*100), Data (SIPP)

5.2 Flows In and Out of the Labor Force

In the benchmark version of our model, all agents enter the labor force at age 1, corresponding to data age 18, and exit at age 564 months, corresponding to data age 65. However, entry and exit ages are actually heterogeneous in the data. Table 1 captures the flows in and out of the *labor force*, averaged across ages. We see that these flows are significant.

Further, we decompose these flows by age. Since the need and scope to find better jobs, and the resulting application behavior, is very strong just after entering the labor market, overall wage and transition patterns are likely to be shaped by the distribution over times of entry. The entry distribution, based on the SIPP, is in figure 9. It is apparent that while the probability of entering the labor force is largest (near 3% per month) at the youngest ages, the age of entry is not uniform.

flow	mean
NE	0.0443
EN	0.0117
NU	0.0222
UN	0.0792
N→NP	0.0145
NP→P	0.0665

Table 1: Participation Flows, Data (SIPP)




Figure 9: Monthly Probability of Labor Market Entry by Age, Data (SIPP)




Figure 10: Retirement Hazard, Data (SIPP)

Target	Data	Model
UE	0.176	0.173
EU	0.0067	0.0053
EE	0.017	0.017
Tenure < 3 months	0.044	0.04
Tenure < 6 months	0.1001	0.095
Tenure < 1 year	0.1831	0.16
Tenure < 2 years	0.2985	0.25
Tenure < 5 years	0.4892	0.42
Tenure < 10 years	0.6719	0.61
wage 55/ wage 30	1.37	1.41

Table 2: Calibration Targets

Similarly, at the end of the lifetime, some workers will retire before age 65. From the perspective of our model, in their valuation of the match, firms have to take into account the increasing probability that a match will be broken up due to retirement. In figure 10 we plot the monthly probability of retiring, decomposed by age, in the data. It is apparent that the probability of exiting the labor force is flat and near zero until age 50, begins increasing noticeably after age 50 from zero to nearly 1% per month by age 55, and spikes up thereafter, reaching close to 2% per month by age 60, and above 5% per month by mid-60's.

While the model we described does not generate endogenous flows in and out of the labor force, the empirical heterogeneity of entry and exit ages that happen for reasons exogenous to the model (e.g. due to time spent in college early on, or disability later on in life) will be relevant for the shapes of transition profiles in the data. Thus, we also calibrate a version of the model where we allow for exogenously heterogeneous entry and exit. It is for this version that we present the results below.

5.3 Calibration Target and Parameter Summary

The targets and the parameters are summarized in tables 2 and 3. Once again we emphasize that all of our targets are averaged across ages, so that we are not predisposing the model to get the life cycle facts right. In the version of the model presented below, we make the workers nearly homogeneous in terms of their initial productivity y_0 . We will also investigate the version with the initial productivity calibrated to match the cross-sectional dispersion of earnings at entry into the labor force. The computation procedure is described in the appendix.

Parameter	Value
β	0.9966
δ	0.005
λ_u	1.0
λ_e	0.78
k	25
b	0.73
ρ^*	1.01
σ_z	0.40
σ_{y_0}	0.10

*: Annualized value

Table 3: Parameters of the Model

6 Results

In this section, we show the life-cycle results from our benchmark model with exogenous heterogeneity of entry and exit times.

6.1 UE Transition Profile

The UE rate in the data is slightly downward-sloping/constant for most of the early part of the life-cycle, and becomes increasingly declining in the last part of the working life, after age 50. The calibrated model mimics the pattern well, as can be seen in figure 11: initially the UE rate stays close to 20% per month until age 50, after which the transition rate from unemployment to employment drops dramatically. The shape of the profile is driven by two opposing effects in the model. On the one hand, as workers age, they become more productive, and thus more valuable to the firm, which makes it easier for them to find a job out of unemployment. On the other hand, as workers approach the end of the life cycle, they also approach retirement - that is, the firm is certain to lose the worker rapidly. This makes the oldest workers much less valuable to the firm, despite the fact that they are the most productive. The two effects roughly balance each other out, with the first slightly dominating, through most of the life cycle, but after age 50 (due also to the fact that some workers will begin exiting into retirement exogenously early), the second effect begins to dominate; workers rapidly become less valuable, which creates a rapid decline in the probability of finding a job out of unemployment.

6.2 EE and EU Transition Profiles

Figures 12 and 13 show that the model successfully captures, by age, the transition rates between jobs, and out of jobs into unemployment. The EU profile is explained as follows. Young workers




Figure 11: UE Transition Rates, Model with Heterogeneous Entry and Exit

in the model have the highest value of searching for new jobs. They are also the most likely to be in low-productivity and low-paid jobs, with joint value below the value of unemployment. Thus it is the youngest workers who are likely to quit into unemployment, as they can search with higher probability when unemployed, thus further increasing their chances of getting a better match in the future. In contrast, the oldest workers sort themselves into high-value matches already, so that the chances of finding a better-paid job are small, while the joint value of the current match is high. Thus they never voluntarily quit into unemployment, and so the UE rate at the end of the lifecycle converges to the exogenous separation rate δ . Similarly, as the youngest workers have the highest returns to search and are the most likely to have low-value current matches, they are more likely to switch between jobs, and that likelihood declines dramatically as workers age, which creates the EE profile we see in the figure.

6.3 Wages

In figure 14, we plot the wage-age profiles in the data and the model. The model successfully reproduces the concave shape of the wage-age profile and the magnitude of the increase in wages that occurs through the lifecycle. While the scale of the increase is targeted in calibration, the shape of the profile is produced by the search frictions in the model. Absent these frictions, increases in wages with age would be driven only by the exogenous productivity process, which is estimated to be convex ($\rho = 1.01$). Thus, it is the search process that creates the concave shape of the profile. The mechanism for this is similar to that described above: initial rapid increase in wages is driven




Figure 12: EE Transition Rates, Model




Figure 13: EU Transition Rates, Model and Data

by the increases in worker productivity that happen as the worker climbs the ladder rapidly due to high returns to search. As the worker ages, returns to search decline, as the worker becomes less valuable to the firm when he approaches retirement. This slows down wage growth.

7 Discussion: Comparative Statics

In this section, we vary parameters of the model numerically to show the role of human capital accumulation and the role of heterogeneous entry and exit in the model.




Figure 14: Wage-Age Profiles, Model and Data

7.1 Role of Human Capital Accumulation

As workers accumulate experience, they become more attractive on the labor market. On the other hand, the flip side of becoming more experienced is becoming older, and therefore closer to the end of one's working life. These two forces work against each other. In figure ?? we see that a higher rate of human capital accumulation leads to a more upward-sloping UE age profile (blue line), whereas the absence of human capital accumulation leads to a decreasing UE age profile across the board (green line).

In terms of the EU rate, figure 15 shows that keeping all other parameters constant, the attractiveness of a job is also affected by the rate of human capital growth. If human capital growth




Figure 15: UE and EU Profiles in Model with Different Rates of Human Capital Accumulation ρ

with experience is higher, then workers become less picky early on in their life, and the result is a flatter EU profile (figure 15, blue line).

7.2 Role of Heterogeneous Entry and Exit

Incorporating the additional risk of match break-ups due to retirement makes the employment of older workers less attractive, and works to offset any increased productivity. Hence, shutting this channel down leads to a (more) increasing UE transition rate with age, until a later moment in life, when the shrinking horizon finally kicks in. Figure 16 demonstrates this.

8 Conclusion

In this paper, we construct a lifecycle model of directed search where workers accumulate human capital and search on and off the job based on their age, human capital, and current wage. This simple model is able to account for key lifecycle facts in the labor market, including transition




Figure 16: UE transition rates in model with and without heterogeneous exit dynamics

profiles between labor market states by age, and age-earnings profiles.

We are now using the model to study the variance of (log) labor productivity and (log) wages conditional on age, to answer the question of how much of the lifecycle experience in the labor market is shaped by differences in productivity versus search frictions. Our future research agenda includes introducing wealth accumulation into the model and studying the properties of the model with aggregate uncertainty.

References

- [1] Acemoglu, D., and R. Shimer. 1999. “Efficient Unemployment Insurance.” *J.P.E.* 107: 893-927.
- [2] Burdett, K., S. Shi, and R. Wright. 2001. “Pricing and Matching with Frictions.” *J.P.E.* 109: 1060–85.
- [3] Menzio, G. and S. Shi. 2009a. “Efficient Search on the Job and the Business Cycle.” *NBER Working Paper* 14905.
- [4] Menzio, G. and S. Shi. 2009b. “Block Recursive Equilibria for Stochastic Models of Search on the Job.” *J. Econ. Theory*.
- [5] Menzio, G. and S. Shi. 2010. “Directed Search on the Job, Heterogeneity, and Aggregate Fluctuations.” *A.E.R. Papers and Proc.*
- [6] Moen, E. 1997. “Competitive Search Equilibrium.” *J.P.E.* 105: 694–723.
- [7] Mortensen, D. 1982. “Property Rights and Efficiency in Mating, Racing and Related Games.” *A.E.R.* 72: 968–79.
- [8] Shi, S. 2009. “Directed Search for Equilibrium Wage-Tenure Contracts.” *Econometrica*, 77: 561-584.
- [9] Shimer, R. 1996. “Contracts in a Frictional Labor Market.” Manuscript M.I.T.

Appendix

A Proof of Theorem 1

(i) We want to prove that there exists a unique list of value and policy functions $\{U_t^*, V_t^*, \theta_t^*, x_{u,t}^*, x_{e,t}^*, d_t^*\}_{t=1}^T$ that satisfies the equilibrium conditions (i), (ii) and (iii), and that these value and policy functions do not depend on the distribution of workers across employment states.

The unemployment value function U_T satisfies the equilibrium condition (i) if and only if

$$U_T(y, \psi) = b \quad (11)$$

for all $(y, \psi) \in Y \times \Psi$. Clearly, there exists a unique value function U_T^* that satisfies (11). Moreover, U_T^* does not depend on the distribution of workers across employment states ψ , i.e. $U_T^*(y, \psi) = U_T^*(y)$.

The employment value function V_T satisfies the equilibrium condition (ii) if and only if

$$V_T(y, z, \psi) = yz \quad (12)$$

for all $(y, z, \psi) \in Y \times Z \times \Psi$. Clearly, there exists a unique value function V_T^* that satisfies (12). Moreover, V_T^* does not depend on the distribution of workers across employment states, i.e. $V_T^*(y, z, \psi) = V_T^*(y, z)$.

The market tightness function θ_T satisfies the equilibrium condition (iii) if and only if

$$q(\theta_T(x, y, \psi)) \{E_{z'} [V_T^*(y, z')] - x\}, \quad (13)$$

and $\theta_T(x, y, \psi) \geq 0$ with complementary slackness for all $(x, y, \psi) \in \mathbb{R} \times Y \times \Psi$. The unique solution to (13) is

$$\theta_T^*(x, y, \psi) = \begin{cases} q^{-1} [k/(E_{z'} [V_T^*(y, z')] - x)] & \text{if } x \leq E_{z'} [V_T^*(y, z')] - k, \\ 0 & \text{else.} \end{cases} \quad (14)$$

Notice that the RHS of (14) depends on the value promised to the workers, x , the worker-specific productivity, y , but not on the distribution of workers across employment states, ψ . Hence, $\theta_T^*(x, y, \psi) = \theta_T^*(x, y)$.

The search policy function for the unemployed, $x_{u,T}$, satisfies the equilibrium condition (i) if and only if

$$x_{u,T}(y, \psi) \in \arg \max_x p(\theta_T^*(x, y))(x - U_T^*(y)) \quad (15)$$

for all $(y, \psi) \in Y \times \Psi$. Fix an arbitrary $(y, \psi) \in Y \times \Psi$. Consider the case in which $x \leq E_{z'} [V_T^*(y, z')] - k$. In this case, (14) implies that x is equal to $E_{z'} [V_T^*(y, z')] - k/q(\theta_T^*(x, y))$ and, hence, the objective function in (15) can be rewritten as

$$\begin{aligned} & p(\theta_T^*(x, y))(x - U_T^*(y)) \\ &= -k\theta_T^*(x, y) + p(\theta_T^*(x, y)) \{E_{z'} [V_T^*(y, z')] - U_T^*(y)\}. \end{aligned}$$

Next, consider the case in which $x > E_{z'} [V_T^*(y, z')] - k$. In this case, (14) implies that $\theta_T^*(x, y)$ is equal to $E_{z'} [V_T^*(y, z')] - k/q(\theta_T^*(x, y))$ and, hence, the objective function in (15) can be rewritten as

$$\begin{aligned} & p(\theta_T^*(x, y))(x - U_T^*(y)) = 0 \\ &= -k\theta_T^*(x, y) + p(\theta_T^*(x, y)) \{E_{z'} [V_T^*(y, z')] - U_T^*(y)\}. \end{aligned}$$

Therefore, for all $x \in \mathbb{R}$, the objective function in (15) is

$$\begin{aligned} & p(\theta_T^*(x, y))(x - U_T^*(y)) \\ & = -k\theta_T^*(x, y) + p(\theta_T^*(x, y)) \{E_{z'} [V_T^*(y, z')] - U_T^*(y)\}. \end{aligned} \quad (16)$$

Notice that x enters (16) only through θ_T^* . For this reason, the search problem of the unemployed can be expressed as a choice over θ , i.e.

$$\max_{\theta \geq 0} -k\theta + p(\theta) \{E_{z'} [V_T^*(y, z')] - U_T^*(y)\}. \quad (17)$$

The objective function in (17) is strictly concave in θ and does not depend on the distribution of workers across employment states, ψ . Hence, (16) admits a unique solution $\theta_{u,T}^*$ which does not depend on ψ , i.e. $\theta_{u,T}^*(y, \psi) = \theta_{u,T}^*(y)$. Given the optimal choice of market tightness $\theta_{u,T}^*$, one can recover the optimal choice of a submarket as

$$\begin{aligned} x_{u,T}(y, \psi) &= E_{z'} [V_T^*(y, z')] - k/q(\theta_{u,T}^*(y)) & \text{if } \theta_{u,T}^*(y) > 0, \\ x_{u,T}(y, \psi) &\geq E_{z'} [V_T^*(y, z')] - k & \text{if } \theta_{u,T}^*(y) = 0. \end{aligned} \quad (18)$$

There are many policy function $x_{u,T}$ that solve (18). However, these policy functions only differ with respect to the submarket that the worker chooses when $\theta_{u,T}^*(y) = 0$. Since this difference is immaterial, we select without loss in generality the following solution to (18)

$$x_{u,T}^*(y, \psi) = E_{z'} [V_T^*(y, z')] - k/q(\theta_{u,T}^*(y)). \quad (19)$$

Notice that $x_{u,T}^*$ does not depend on the distribution of workers across employment states because none of the terms on the RHS of (19) does. Hence, $x_{u,T}^*(y, \psi) = x_{u,T}^*(y)$.

The search policy function for the unemployed, $x_{e,T}$, satisfies the equilibrium condition (ii) if and only if

$$x_{e,T}(y, z, \psi) \in \arg \max_x p(\theta_T^*(x, y))(x - V_T^*(y, z)) \quad (20)$$

for all $(y, z, \psi) \in Y \times Z \times \Psi$. Again, one can show that the unique solution to (20) is

$$x_{e,T}^*(y, z, \psi) = E_{z'} [V_T^*(y, z')] - k/q(\theta_{e,T}^*(y, z)), \quad (21)$$

where $\theta_{e,T}^*(y, z)$ is defined as

$$\theta_{e,T}^*(y, z) = \arg \max_{\theta \geq 0} -k\theta + p(\theta) \{E_{z'} [V_T^*(y, z')] - V_T^*(y, z)\}.$$

Notice that $x_{e,T}^*$ depends on the worker-specific productivity, y , the match-specific productivity, z , but not on the distribution of workers across employment states, ψ . Hence, $x_{e,T}^*(y, z, \psi) = x_{e,T}^*(y, z)$.

The separation policy function for the unemployed, d_T , satisfies the equilibrium condition (ii) if and only if

$$d_T(y, z, \psi) \in \arg \max_{\delta \leq d \leq 1} \left\{ \begin{aligned} & dU_T^*(y) + (1-d)V_T^*(y, z) \\ & + (1-d)\lambda_e p(\theta_T^*(x_{e,T}^*(y, z), y))(x - V_T^*(y, z)) \end{aligned} \right\} \quad (22)$$

for all $(y, z, \psi) \in Y \times Z \times \Psi$. Clearly, there exists a unique policy function, d_T^* , that satisfies (22). Moreover, d_T^* depends on the worker-specific productivity, y , the match-specific productivity, z , but not on the distribution of workers across employment states, ψ . Hence, $d_T^*(y, z, \psi) = d_T^*(y, z)$.

The unemployment value function U_{T-1} satisfies the equilibrium condition (i) if and only if

$$U_{T-1}(y, \psi) = b + \beta [U_T^*(y) + \lambda_u p(\theta_T^*(x_{u,T}^*(y), y))(x - U_T^*(y))] \quad (23)$$

for all $(y, \psi) \in Y \times \Psi$. Clearly, there exists a unique solution U_{T-1}^* to (23). Moreover, U_{T-1}^* depends on the worker-specific productivity, y , but not on the distribution of workers across employment states, ψ . Hence, $U_{T-1}^*(y, \psi) = U_{T-1}^*(y)$.

The employment value function V_{T-1} satisfies the equilibrium condition (ii) if and only if

$$V_{T-1}(y, z, \psi) = yz + \beta \left\{ \begin{aligned} & d^*(\rho y, z) U_T^*(\rho y) + (1 - d^*(\rho y, z)) V_T^*(\rho y, z) \\ & + (1 - d^*(\rho y, z)) \lambda_e p(\theta_T^*(x_{e,T}^*(\rho y, z), \rho y))(x - V_T^*(\rho y, z)) \end{aligned} \right\} \quad (24)$$

for all $(y, z, \psi) \in Y \times Z \times \Psi$. Clearly, there exists a unique solution V_{T-1}^* to (24). Moreover, V_{T-1}^* depends on the worker-specific productivity, y , the match-specific productivity z , but not on the distribution of workers across employment states, ψ . Hence, $V_{T-1}^*(y, z, \psi) = V_{T-1}^*(y, z)$.

By repeating the above steps, it is straightforward to establish the existence, uniqueness and independence of ψ of value and policy functions $\{\theta_{T-t}^*, x_{u,T-t}^*, x_{e,T-t}^*, d_{T-t}^*, U_{T-1-t}^*, V_{T-1-t}^*\}_{t=1}^{T-1}$ that satisfy the equilibrium conditions (i), (ii) and (iii). Hence, an equilibrium exists, is unique and it is block recursive. ■

B Proof of Theorem 2

At the beginning of the period, the social planner observes the aggregate state of the economy $\{u_t, g_t\}_{t=1}^T \equiv \psi$. At the separation stage, the planner chooses the probability $d_t(y, z) \in [\delta, 1]$ of destroying a match between a worker and a firm, given that the worker's age is t , the worker-specific component of productivity is y , and the match-specific component of productivity is z . At the search stage, the planner chooses $\theta_{u,t}(y) \in \mathbb{R}_+$, the ratio of vacancies to workers at the location where unemployed workers of age t and idiosyncratic productivity y look for new matches. Moreover, the planner chooses $\theta_{e,t}(y, z) \in \mathbb{R}_+$, the ratio of vacancies to workers at the location where workers of age t , idiosyncratic productivity y who are employed in matches of quality z look for new jobs. Given the choices $\{d_t, \theta_{u,t}, \theta_{e,t}\}_{t=1}^T$, aggregate consumption is given by

$$\begin{aligned} F(\{d_t, \theta_{u,t}, \theta_{e,t}\}_{t=1}^T | \psi) = & -k \left\{ \sum_{t,y} \lambda_u \theta_{u,t}(y) u_t(y) + \sum_{t,y,z} \lambda_e (1 - d_t(y, z)) \theta_{e,t}(y, z) g_t(y, z) \right\} \\ & + \sum_{t,y} b \tilde{u}_t(y) + \sum_{t,y,z} y z \tilde{g}_t(y, z). \end{aligned} \quad (25)$$

In the above expression, $\{\tilde{u}_t, \tilde{g}_t\}_{t=1}^T$ denotes the distribution of workers at the production stage, which is given by

$$\begin{aligned} \tilde{u}_t(y) &= u_t(y) (1 - \lambda_u p(\theta_{u,t}(y))) + \sum_z g_t(y, z) d_t(y, z), \\ \tilde{g}_t(y, z') &= u_t(y) \lambda_u p(\theta_{u,t}(y)) f(z') + \sum_z g_t(y, z) (1 - d_t(y, z)) \lambda_e p(\theta_{e,t}(y, z)) f(z') \\ &\quad + g_t(y, z') (1 - d_t(y, z')) (1 - \lambda_e p(\theta_{e,t}(y, z'))). \end{aligned} \quad (26)$$

In contrast, the distribution of workers at the beginning of next period is denoted as $\hat{\psi} = \{\hat{u}_t, \hat{g}_t\}_{t=1}^T$ and is given by

$$\hat{u}_1(y) = h(y), \quad \hat{u}_t(y) = \tilde{u}_{t-1}(y), \quad \hat{g}_t(y, z) = g_{t-1}(y/\rho, z). \quad (27)$$

The planner maximizes the sum of present and future consumption discounted at the factor β . Hence, the planner's value function W solves the following functional equation

$$\begin{aligned} W(\psi) = & \max_{\{d_t, \theta_{u,t}, \theta_{e,t}\}_{t=1}^T} F(\{d_t, \theta_{u,t}, \theta_{e,t}\}_{t=1}^T | \psi) + \beta W(\hat{\psi}), \\ \text{s.t. } & d_t : Y \times Z \rightarrow [\delta, 1], \theta_{u,t} : Y \rightarrow [0, \bar{\theta}], \theta_{e,t} : Y \times Z \rightarrow [0, \bar{\theta}]. \end{aligned} \quad (28)$$

Lemma 1 (i) The planner's value function $W(\psi)$ is the unique solution to the functional equation (28) (ii) The planner's value function $W(\psi)$, $\psi = \{u_t, g_t\}_{t=1}^T$, is linear in u_t and g_t . That is, there exist a scalar $W_0 \in \mathbb{R}$, and two lists of functions $\{W_{u,t}, W_{e,t}\}_{t=1}^T$, $W_{u,t} : Y \rightarrow \mathbb{R}$, $W_{e,t} : Y \times Z \rightarrow \mathbb{R}$, such that

$$W(u_t(y), g_t(y, z)) = W_0 + \sum_{t,y} W_{u,t} u_t(y) + \sum_{t,y,z} W_{e,t} g_t(y, z). \quad (29)$$

Proof: (i) Let $B(\Psi)$ be the set of bounded, continuous functions $R : \Psi \rightarrow \mathbb{R}$ with the sup norm. Define the operator $T : B(\Psi) \rightarrow B(\Psi)$ as the mapping associated with the functional equation (28). It is straightforward to verify that $R, R' \in B(\Psi)$ and $R \leq R'$ imply $TR \leq TR'$. Similarly, it is straightforward to verify that $R \in B(\Psi)$ and $a \geq 0$ imply $T(R + a) = TR + \beta a$. Hence, Blackwell's sufficient conditions (Theorem 3.3 in Stokey, Lucas and Prescott, 1989) guarantee that the operator T is a contraction and it admits only one fixed point R^* . 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 $B'(\Psi) \subset B(\Psi)$ be the set of functions $R : \Psi \rightarrow \mathbb{R}$ that are bounded, continuous and linear in u_t and g_t . Clearly, R belongs to the set $B'(\Psi)$ if and only if there exist a scalar $R_0 \in \mathbb{R}$, and two lists of functions $\{R_{u,t}, R_{e,t}\}_{t=1}^T$, $R_{u,t} : Y \rightarrow \mathbb{R}$, $R_{e,t} : Y \times Z \rightarrow \mathbb{R}$, such that

$$R(\psi) = R_0 + \sum_{t,y} R_{u,t}(y) u_t(y) + \sum_{t,y,z} R_{e,t}(y, z) g_t(y, z).$$

Consider an arbitrary function R in $B'(\Psi)$. Then, after substituting the constraints into the maximand of (28), we obtain

$$(TR)(\psi) = \hat{R}_0 + \sum_{t,y} \hat{R}_{u,t}(y) u_t(y) + \sum_{t,y,z} \hat{R}_{e,t}(y, z) g_t(y, z).$$

where $\hat{R}_0$, $\hat{R}_{u,t}(y)$ and $\hat{R}_{e,t}(y, z)$ are given by

$$\begin{aligned} \hat{R}_0 &= \beta E_y[R_{u,1}(y)], \\ \hat{R}_{u,t}(y) &= \max_{\theta \geq 0} \{ -k\lambda_u \theta + [1 - \lambda_u p(\theta)] [b + \beta R_{u,t+1}(y)] \\ &\quad + \lambda_u p(\theta) E_{z'} [yz' + \beta R_{e,t+1}(\rho y, z')] \}, \\ \hat{R}_{e,t}(y, z) &= \max_{1 \geq d \geq \delta, \theta \geq 0} \{ d[b + \beta R_{u,t+1}(y)] + (1-d) [1 - \lambda_e p(\theta)] [yz + \beta R_{e,t+1}(\rho y, z)] \\ &\quad + (1-d) \lambda_e p(\theta) E_{z'} [yz' + \beta R_{e,t+1}(\rho y, z')] \}. \end{aligned}$$

Since R is an arbitrary function in $B'(\Psi)$, (xx) implies that $T : B'(\Psi) \rightarrow B'(\Psi)$. Moreover, since $B'(\Psi)$ is a closed subset of $B(\Psi)$ and $T : B'(\Psi) \rightarrow B'(\Psi)$, Corollary 1 to Theorem 3.2 in Stokey, Lucas and Prescott (1989) implies that the fixed point of the operator T belongs to $B'(\Psi)$. That is, $W \in B'(\Psi)$. ■

From the previous lemma, it follows that the social planner's value function $W(\psi)$ can be

written as in (29), where $\{W_{u,t}\}_{t=1}^T$ and $\{W_{e,t}\}_{t=1}^T$ are given by

$$\begin{aligned} W_{u,t}(y) &= \max_{\theta \geq 0} \{-k\lambda_u\theta + [1 - \lambda_u p(\theta)] [b + \beta W_{u,t+1}(y)] + \lambda_u p(\theta) E_{z'} [yz' + \beta W_{e,t+1}(\rho y, z')]\}, \\ W_{e,t}(y, z) &= \max_{1 \geq d \geq \delta, \theta \geq 0} \{d[b + \beta W_{u,t+1}(y)] + (1 - d) [1 - \lambda_e p(\theta)] [yz + \beta W_{e,t+1}(\rho y, z)] \\ &\quad + (1 - d) \lambda_e p(\theta) E_{z'} [yz' + \beta W_{e,t+1}(\rho y, z')]\}. \end{aligned} \tag{30}$$

Let $\tilde{W}_{u,t}(y)$ be defined as $b + \beta W_{u,t+1}(y)$ for $t = 1, 2, \dots, T - 1$, and b for $t = T$. Similarly, let $\tilde{W}_{e,t}(y, z)$ be defined as $yz + \beta W_{e,t+1}(\rho y, z)$ for $t = 1, 2, \dots, T - 1$, and yz for $t = T$. Using (30), we can rewrite $\tilde{W}_{u,t}$ and $\tilde{W}_{e,t}$ as

$$\begin{aligned} \tilde{W}_{u,t}(y) &= b + \beta \max_{\theta \geq 0} \{-k\lambda_u\theta + [1 - \lambda_u p(\theta)] \tilde{W}_{u,t}(y) + \lambda_u p(\theta) E_{z'} [\tilde{W}_{e,t}(y, z')]\}, \\ \tilde{W}_{e,t}(y, z) &= \max_{1 \geq d \geq \delta, \theta \geq 0} \{d\tilde{W}_{u,t+1}(y) + (1 - d) [1 - \lambda_e p(\theta)] [\tilde{W}_{e,t}(y, z)] \\ &\quad + (1 - d) \lambda_e p(\theta) E_{z'} [\tilde{W}_{e,t+1}(y, z')]\}. \end{aligned} \tag{31}$$

It is straightforward to verify that the social planner's component value functions $\{\tilde{W}_{u,t}, \tilde{W}_{e,t}\}_{t=1}^T$ are equal to the equilibrium value functions $\{U_t^*, V_t^*\}_{t=1}^T$. Similarly, it is immediate to verify that the social planner's policy functions $\{\theta_{u,t}, \theta_{e,t}, d_t\}_{t=1}^T$ (i.e., the policy functions associated with the component value functions (31)) are equal to the equilibrium policy functions $\{\theta_{u,t}^*, \theta_{e,t}^*, d_t^*\}_{t=1}^T$. Therefore, the equilibrium decentralized the solution to the social planner's problem. ■

C SIPP Data: Definitions and Calculations

Sample Design The SIPP uses a two-stage sample design to select the households to take part in the survey. In the first stage, counties (or groups of counties) form PSUs (primary sample units). Large PSUs are selected with certainty, the remaining ones are grouped together in strata based on similarity. Within the PSU, housing units are split into high- and low-income strata (from Census data); the low-income stratum is sampled at a rate of 1.66 of the high-income stratum (Westat 2001).

For estimation of means and variances this sample design has to be taken into account. The relevant sample weights (WPFNWGT, LGTPNLWT) give us the information to calculate averages; but we need more information to calculate the variances correctly. The information about the original PSU and stratum is not included in the data set, for privacy concerns. Instead PSUs are combined to form variance strata, with two variance units in each variance stratum (Westat 2001, chapter 8). We use these as (pseudo-)PSUs and (pseudo-)strata, in the variance estimation.⁶

Adjusting Sample Weights to Reflect Age Structure of the Population The age structure in the population, for ages above 35, has a pyramid shape, as we mention in the text. The weights in the SIPP are chosen such that the weighted sample population represents the US population. In the model, however, we are following a cohort, which in the simplest version of the model does not face any attrition. Therefore, when we calculate the average transition rates to match, we calculate the transition probability from state k to l , $f_{kl}(i)$, for each age i (using the weights in the data), and then calculate the simple average transition rate $F_{kl} = 1/(T_{max} - T_0) \sum_{i=T_0}^{T_{max}} f_{kl}(i)$. An equivalent

⁶In particular, using Stata's survey commands (`svy`), we are now able to calculate the variances, using Fay's method (with factor 0.5). In case there is only one observation in a stratum, we use the unconditional mean to calculate the variance.

procedure applies to the calculation of life-time average stocks. This seemingly subtle point is the reason that some of our reported ‘lifetime’ averages are below the standard values: in standard calculations, one commonly treats every individual the same, independent of age. As flows vary with age, and the population in the US has, for ages above 35, a pyramid shape, there is a bias towards the numbers of lower ages. For example, this results in higher average flow numbers (and in the context of a stationary infinitely-lived agent model, it’s often a good approximation choice).

Labor Force States The information in the SIPP allows us to define employment, unemployment and not-in-the-labor force (NLF) in slightly different ways, when trying to approach the equivalent notion in the discrete-time model. One can use the monthly variable RMSEER, the employment status recode for the entire month, constructed from the answers in the interview, we can define those employed who had a job at least part of the month, unemployed those who did not have a job for at least part of the month and were looking in at least one week, and NLF if the worker did not have a job all month, and did not look for one all month. This measure has as an disadvantage that it can assign two labor market states simultaneously to the worker. This makes it harder to get a clear measure of the transitions between months. Therefore, our preferred measure is to use the labor market variable in the first week of the month (which also is also a closer analogue to the measurement in the discrete-time model), RWKESR1: the employment status recode for week 1. In this case we assign the worker to a job if he says he has a job in the first week (even when absent without pay⁷), and make him unemployed if he does not have a job and is currently looking, or has looked in the past month; in the remaining case of reported joblessness, the worker is NLF. As an additional check we create an alternative definition where a worker is assigned to unemployment as opposed to NLF if he has been looking for a job at any point in the wave.⁸

For those individuals who are out of the labor force, we try to uncover the reason for their non-participation: a temporary illness/injury, family circumstances, a chronic illness/disability, retirement, discouragement, or schooling. With regard to the latter, we attempt to specify when a worker enters the labor market for potentially long-term jobs. For those who are without a job for the entire four-month period, we use the ERSNOWRK variable. For those who lost their job, we look at the reason for the loss of job in ERSEND1 and ERSEND2.⁹ As people do not always consistently give the same answer to the question why they are not working, we need to make some assumptions to assign them to one category at each time. We e.g. consider retirement as an absorbing state (in the sense that disabilities occurred while retired are ignored); chronic health conditions that later on do not prevent workers from taking full-time jobs are reclassified as temporary conditions. Similarly, workers who claim to retire but at a future point decide to enter the labor force again, are considered to be temporary out of the labor force, in the ‘other’ category. (In fact, if these individuals re-retire, we assume this is temporary as well – to be on the conservative side.) Individuals who are NLF but do not specify why, are treated as NLF by choice.

An individual who reports to be in school, is below thirty and does not have periods larger than 4 months between schooling periods (e.g. because of summer jobs) is not considered an active participant in the labor market. The entry of a worker in the ‘regular’ labor market is given by

⁷A worker who tells the interviewer that he will begin a new job in 30 days is also in this category, but this group is not large.

⁸Note that the “looking for a job” definition in the SIPP used in the variable ELKWRK presumes that a worker is without a job for at least one week, and available to accept a job, and that the worker took some active measures to find a job. It is not informative about on-the-job search.

⁹ERSNOWRK specifies as reasons for not having a job: temporarily unable to work because of injury/illness; retirement; chronic health condition/disability; pregnancy; taking care of others; going to school; unable to find work; on layoff; not interested. ERSEND variable asks why employment relationship was ended, and includes the same list of reasons, and additionally reasons more specific to the particular company e.g., dismissal, bankruptcy of the company.

the first month in which the worker will not have schooling for the next five months. A person is said to be in school when he answers “going to school” in ERSNOWRK or when he is enrolled full-time according to the school enrollment variables RENROLL, EENRLM. The distribution of entry times is calculated from the switch out of the ‘in-school’ variable, by simple counting when these switches take place. We correct these numbers for the fact that some ages are under- and some are overrepresented.

Tenure, Tenure Distribution We calculate tenures from the variables TSJDATE1, TSJDATE2, but consider the point of measurement at the very beginning of the month. A worker who gets a job during the month is not given a tenure count during this month (but only in the next month). A tenure count of n means that the worker is in the n th month at the current firm. Thus, a tenure < 3 , means that the worker has either passed the first of the month once or twice.

To check how robust the tenure distribution is with respect to changes in our assumptions on counting those who just got a job, we consider a number of ways one can treat jobs that start on the first of the month, among which we mention: to treat a job that was started on the first of a given month, as a job that was created in the given month (and therefore not counted at the beginning of the month); to treat a job started on the first (and possibly some subsequent days) as been created in the previous month, and to count a job created on the first of the month, as a job in the first month of tenure. The implied distribution in terms of tenure groups is given in the table below. The small differences are comforting: changing the assumptions about when the

tenure<...	baseline	first 5 days	first day
< 3months	0.0441	0.0440	0.0451
< 6 months	0.1001	0.1024	0.1036
< 1 year	0.1831	0.1883	0.1895
< 2 years	0.2985	0.3059	0.3051
< 5 years	0.4892	0.4967	0.4943
< 10 years	0.6719	0.6719	0.6719

Table 4: Tenure Moments: robustness

first positive tenure is counted varies the degree of inconclusiveness of tenures of less than a month. These tenures are relatively common, thus depending on the change of definition (and the fact that not all of the jobs with tenures ‘on the margin’ survive long), the change in the size of the first tenure group could be significant. We find that this is not the case.

Employer-to-Employer Movements The SIPP data covers the wage and hours data of jobs at two distinct firms in a four-month period. There are only few cases where a space for two firms’ data entry is insufficient to keep track of the EE movements between ‘dominant firms’, which we consider the firms where most of the hours are worked and (in case of ambiguity) where most income is earned, and (in case of continued ambiguity) the longest tenure. We define an employer-to-employer transition to occur when at the beginning of one month a worker has as his main employer (most hours worked, most income earned) one firm, and in the next month the firm has as his main employer another firm. Transitions from contingent work to a new employer are not counted as an EE change. Likewise, if a worker transitions back to a main firm which had been his main firm before (in the sample), we don’t count this as an EE transition. For robustness, we discard those transitions for which the worker said he was not working and looking for a job in any intermediate weeks (a short period of unemployment, potentially), and those where the individual reported no work in the first week of the month. Neither of these causes any perceptible difference

in the transition rates.

Wages The wages in the SIPP are quite noisily. We employ various definition to minimize the impact of the noise. TPMSUM1 and TPMSUM2 measure the gross earnings (before deductions) received for a given month. TPYRATE1, TPYRATE2 measure the regular hourly payrate for those who are paid on an hourly basis (EPAYHR1). TPMSUM (in dollars) is a monthly variable, while TPYRATE (in cents) is recorded once per wave for each of the two firms recorded. The usual hrs for each of the two firms is recorded (once per wave) in EJBHRS1.

For mean wages, and income dispersion (conditional on having income) we also look at these measures restricted to those who are paid by the hour.

D Computation

We solve the worker's and firm's maximization by backwards induction, searching over a (dynamically adjusting) grid. Then we simulate 10000-20000 realizations of a working life, and calculate the statistics over these, using them either as a repeated cross-section, or using the panel dimension, in close analogue to calculating the same models in the SIPP data. To match the model moments to the data moment, we employ the simplex method (dividing the set of moments to be matched into three, and apply the simplex method hierarchically, in different layers), and as alternative, use one encompassing simplex method, or combinations of grid search and the simplex method.