Labor Market Experience and Worker Flows*

Aspen Gorry

University of Chicago

aspen@uchicago.edu

November 21, 2008

Abstract

This paper presents a model of learning where labor market experience improves the accuracy of information about the quality of future job matches. The model can explain the decline in both job finding and separation rates with age found in the data. Beyond accounting for labor market flows, the model gives predictions that are consistent with observed empirical wage distributions. I evaluate the effects of minimum wages on labor outcomes and find that the endogenous decline in job finding rates is essential to understand high unemployment rates for young workers subjected to minimum wages. Finally, I examine the lasting consequences for workers who start their careers in a bad economy. The model generates sizable wage losses that last eight to ten years without lasting differences in unemployment. These findings are consistent with empirical studies.

*This is a copy of my job market paper. Please send comments to aspen@uchicago.edu. I thank my advisors Robert Shimer, Fernando Alvarez, Derek Neal, and Nancy Stokey for helpful comments. All mistakes are my own.

1 Introduction

During the first ten years of labor market experience, workers transition from high job turnover into stable employment and have rapid wage growth. About two-thirds of lifetime job turnover and wage growth occurs during these early years (see Topel and Ward (1992)). Initial high turnover manifests itself in both high job finding and separation rates for young workers. The mechanism where young workers transition from rapid turnover to stable employment is not well understood. Given the importance of early work experience, a theory is needed to explain labor market outcomes of young workers. A full understanding of these labor market transitions must account for the decline in both job finding and separation rates over the life cycle.

I propose a model where labor market experience allows workers to learn about the quality of future matches. Past work experience allows a jobless worker to differentiate between good and bad matches before accepting a new job. This simple mechanism generates the decline in unemployment, job separation and job finding rates with age, the decline in separations with job tenure, the rise in wages with labor market experience, and fat right tails in the wage distribution.

The novel feature of the model is that unemployed workers receive a signal about the quality of a new match. The precision of the signal increases as the worker gains experience. Like in Jovanovic (1979) workers are uncertain about the quality of their match and learn about the quality while working. In each period when employed the worker receives a signal and updates her belief about the quality of the job.

In this model, workers learn rapidly about the quality of their current job by observing output and this experience is useful in helping them determine the quality of future opportunities when seeking a new job. This result contrasts with classic models of learning that have been used to understand wage growth over workers' careers. These models seek to explain how transferable information learned at a given job is to future employment. On one extreme is the Jovanovic (1979) model where workers only learn about their current job and all human capital investments are job specific[1]. Here all information is lost once the worker decides to change jobs. On the other extreme are models where workers learn about their entire vector of abilities when employed as in Gibbons et al. (2005). In these models, workers learn about their ability to perform at all jobs at a constant rate. This paper provides a mechanism that is between the two extremes. My model allows the worker to learn quickly about the quality of her current job and let only some of that information carry over to future employment.

My model has implications for both job finding and separation rates as workers gain experience. Job finding rates are determined in the model by two factors. First, workers receive job offers at an exogenous rate. Second, workers can choose to accept or reject job offers that they receive. The chance of a job being accepted depends both on the minimum value of accepted opportunities that the worker will take and the distribution of jobs. Experience implies that the worker has more precise information about the quality of a job offer. This model is able to generate a decline in job finding rates as experience allows individuals to reject bad offers that were initially accepted. For inexperienced workers

[1] This model could also be reinterpreted so that each worker firm match denotes an industry or occupation.

jobs are experience goods; they only learn about the quality of the match by trying it out. However, as workers gain experience jobs become inspection goods. Market experience influences decisions by unemployed workers about which jobs to accept. This contrasts with the standard Jovanovic (1979) model where workers accept all or a constant fraction of job offers since their information when unemployed does not change with experience.

Moreover, the model with experience is able to account for the full decline in the job separation rate. In Jovanovic's (1979) model, job durations are identically and independently distributed random variables and hence the turnover generated by the model is a renewal process. Each time the worker becomes unemployed she is in an identical position; job finding rates are constant. In the standard Jovanovic (1979) model only statistical sorting generates a decline in job separation rates. In his model, older workers are more likely to have been in their job longer. While this sorting is able to qualitatively match the decrease in separations with age, it does not quantitatively account for the magnitude of the decline found in the data. Adding experience to the model generates a second force that causes separations to decline. Older workers are selective, so their new jobs are more likely to be good. This additional feature predicts both a decline in job finding rates and quantitatively captures the full decline in job separation rates.

I consider two experiments with the model to demonstrate the importance of finding and separation rates in understanding the employment experiences of young workers. First, I consider the effects of minimum wage restrictions on worker outcomes. In the model, minimum wages restrict the jobs that young workers are able to accept. Minimum wage restrictions

vary dramatically between the U.S. and Europe. Along with much higher minimum wages, European employment is characterized by having lower levels of job finding and separation rates. I show that high minimum wages drive down job finding and separation rates early in workers' careers leading to the high levels of youth unemployment observed in many European countries. To correctly predict the effect of policies on labor market outcomes, a model that can generate changes in both job finding and separation rates is needed.

Finally, I examine the lasting consequences for a worker who enters the labor force in a bad economy. I assume that for the first two years of labor market experience workers are subject to lower than normal job finding rates. After facing poor job prospects for two years, workers are shown to experience lower wages for 6-8 more years. These lasting declines in wages do not correspond to greater amounts of future unemployment as workers quickly revert to normal levels of employment after job finding rates return to their standard value. The model's predictions of persistent wage losses with no lasting employment effects are consistent with the Kahn (2006) who examines the effects of graduating during recessions.

Related to the literature on learning is an empirical literature that examines the transferability of human capital across jobs. Altonji and Shakotko (1987), Topel (1991), and Altonji and Williams (2005) examine the extent to which wages rise with tenure in a given job rather than through total job market experience. While they come to slightly different conclusions, Altonji and Williams (2005) reconcile the methods to find that tenure has a modest effect on wage growth taking into account the effects of labor market experience. Learning in my model has both a firm specific and a general effect. While some experience

transfers to allow individuals to better identify the quality of future matches, workers learn about the quality of their current job at a faster rate. Although much of wage growth can be accounted for by career experience, there is still a premium for job tenure. Similarly, Mincer and Jovanovic (1982) and Bartel and Borjas (1982) explore the relationship between turnover and wage growth. They find that much of wage growth is due to general experience while smaller portions can be attributed to firm experience and mobility choices. Job changes early in the career are correlated with positive wage gains where changes later in life have negative effects. These findings are all consistent with model predictions.

Moscarini (2003), Moscarini (2005), and Papageorgiou (2007) present models that are closely related. Moscarini (2005) and Moscarini (2003) assume that jobs are drawn from a distribution of only two types to allow Jovanovic's (1979) model to be embedded into a general equilibrium matching framework. Moscarini (2005) shows that this model can generate a wage distribution of the same shape as the empirical distribution. Moscarini (2003) applies the model to think about the empirical tenure distribution. My paper adapts assumptions about the distribution of jobs in these papers to allow the framework to be embedded into a general equilibrium setting. In both of these models the value of unemployment is static during the worker's career as in the standard Jovanovic (1979) framework, so they are unable to explain changes in the job finding rate.

Papageorgiou (2007) extends Moscarini (2005) by giving workers a vector of abilities on possible jobs that they learn about as they work, similar to the Gibbons et al. (2005) framework. This paper seeks to explain worker flows across occupations. Since workers

direct their search to the job that is best for them in terms of both productivity and learning about their abilities, the model does not account for differences in job finding rates over the career. Separations and occupational changes decline as workers discover their type over time. While Papageorgiou (2007) focuses on the worker's optimal decision regarding occupation selection, he does not explore how learning in his model affects job finding and separation rates with age.

An empirical literature related to Papageorgiou (2007) on career and job specific matches seeks to explain the decline in turnover during the life cycle. Neal (1999) presents a model where workers search for both a career and job specific match. Once they have a career match they can draw a new job match, but if they decide to get a new career match they must also draw a job match. Pavan (2007) argues that a model of this type is better able to match separation behavior than standard models while Pavan (2006) explores this model's predictions for the behavior of wages. My model is able to generate observed declines in job finding and separation rates without adding the complexity of a second type of career match.

The paper proceeds as follows. Section 2 presents the model. Section 3 describes how the parameters of the model are chosen. Section 4 presents the results from the calibrated model and compares them to the standard Jovanovic (1979) model. The effects of minimum wages are explored in Section 5, and the lasting consequences of starting a worker's career in a bad economy are explored in Section 6. Section 7 concludes.

2 Model

2.1 Economic Environment

This section describes the optimal decision problem for a risk neutral worker who maximizes the present discounted value of consumption. The worker's problem will be described in a general matching framework. When functional forms are specified the classic Jovanovic (1979) model will be a special case. There is no storage technology; in each period the worker must consume what she produces. The worker's preferences are given by:

$$U = \sum_{t=0}^{\infty} \beta^t c_t$$

Workers make two decisions. When employed they decide between quitting to search for a new job and continuing employment. When unemployed, if they receive an opportunity, they choose to accept or reject the job offer.

As in Moscarini (2005), the economy is composed of good and bad jobs. Let μ denote the average productivity of a worker on a job. I assume $\mu \in \{\mu_h, \mu_l\}$. Let μ_h denote good jobs and μ_l denote bad jobs so that $\mu_h > \mu_l$. The worker is uncertain about the quality of her job. With two job types, the probability that the job is of type μ_h, p, is a sufficient statistic for the worker's current employment. p is a random variable that describes the worker's current belief about the quality of her job.

The worker learns about the quality of the match in two ways. First when employed she gets a signal about the quality of her match in each period and updates her belief according

to a known process G. In general this distribution depends on the value of the current belief, p, so the the distribution of updated beliefs, p', is given by $G(p'|p)$. Two restrictions are made on G. First, I assume G is non-degenerate so that the signal conveys some information about p. Second, G is restricted so that p is a martingale. This is a natural restriction since G is used to update an individual's current beliefs. This implies that the expected value of $G(p'|p)$ is p. That is:

$$\int_0^1 p' G(dp'|p) = p$$

The second way that the worker learns about the quality of a job is through an initial signal that she receives when unemployed. When she gets a job offer, her initial signal depends on past experience. Let $H(p'|\tau)$ be the distribution of beliefs about initial jobs for a worker with experience τ. I assume that H is weakly increasing in τ in terms of second order stochastic dominance. This means that for $\tau_1 > \tau_2$:

$$\int_0^x H(p'|\tau_1) - H(p'|\tau_2) dp' \geq 0 \quad \forall \;\; x \in [0,1]$$

For higher values of τ workers get more initial information about the quality of a job. This increasing information for experienced unemployed workers is the novel feature of the model. A sufficient condition for second order stochastic dominance is that if $\tau_1 > \tau_2$ then $H(p'|\tau_1)$ is a mean preserving spread of $H(p'|\tau_2)$. An example would be for workers to receive a signal from G for each unit of experience τ. In this case, $H(p'|0)$ is degenerate and a prior probability of a job opportunity being good must be specified.

When employed the worker receives a wage, $c_t = w(p)$, that depends on the probability that her job is good. $w(p)$ can be determined in many ways. Since I consider the problem of a worker in isolation, workers receive the output produced on their job in each period. I let $w(p)$ denote average output from a job that is believed to be good with probability p[2].

The worker can separate from the job for two reasons. First, she could receive an unfavorable signal about the job quality and decide to quit. Second, with exogenous probability $\delta > 0$ an employed worker becomes separated from the job in each period. δ captures reasons for job separations not captured by the endogenous quits that arise from learning. Possible examples include plant closures or geographic relocation by the worker.

Let $V(p, \tau)$ be the value function for an employed worker with belief p and experience τ. The value is written as:

$$V(p, \tau) = w(p) + \beta\delta U(\tau + 1) + \beta(1 - \delta) \int_0^1 \max\{U(\tau + 1), V(p', \tau + 1)\} G(dp'|p)$$

A worker with belief p and experience τ gets her expected output $w(p)$ and the continuation value from employment. In the next period, the worker is separated from her job with probability δ, becoming unemployed with experience $\tau + 1$. With probability $1 - \delta$ she is not separated from her job and receives her updated belief from the distribution G. Depending on the realization of her updated belief she can choose to remain employed with belief p' and experience $\tau + 1$ or quit to become unemployed with experience $\tau + 1$.

[2]In an environment with firms the wage can be determined in two different ways. First, I could follow Jovanovic (1979) in assuming that there are competitive firms that compete with workers so that the wage in each period is the expected value of output. Alternately, there could be bargaining over the wage between firms who are matched with workers.

Unemployed workers consume the unemployment value $c_t = b$. I assume that b high enough that if a worker knows for certain that a job is bad it is optimal to quit and low enough so that if the worker knows that the job is good that she will work. These assumptions ensure that the worker's search problem is non-trivial.

When unemployed, the worker with experience τ gets an offer from the distribution of jobs $H(p'|\tau)$ with exogenous probability λ . She must choose between remaining unemployed and becoming employed with belief p'. If she does not receive a job offer she remains unemployed with the same experience.

Let $U(\tau)$ be the value function for an unemployed worker with experience τ. The value function is given by:

$$U(\tau) = b + \beta(1-\lambda)U(\tau) + \beta\lambda \int_0^1 \max\{U(\tau), V(p', \tau)\} H(dp'|\tau)$$

An unemployed worker with experience τ gets to consume b. With probability λ she receives a job offer from $H(p|\tau)$ and must choose to remain unemployed or begin work in the next period. If she does not get a job offer she remains unemployed.

This setup is similar to that found in Jovanovic (1979) and Moscarini (2005). To get to this well known framework, the worker's signal while employed G can be defined to depend only on observed output in each period and the initial signal of job quality H can be defined so that it does not depend on τ. In these models, the worker's unemployed value function is independent of past experience.

In the model with experience the worker's job finding rate can fluctuate over the life cycle

for two reasons. When workers receive more information about a job offer they can change the reservation probability of jobs that they are willing to accept and the distribution of offers changes as experience accumulates. A higher reservation probability will imply that a worker will accept fewer jobs all else equal. In equilibrium, the reservation probability will depend on the worker's experience τ so it changes over her life. The job finding rate can also depend on changes in the distribution of job offers as experience increases. To understand this effect suppose that the reservation probability is p^* for all τ. That is workers accept job offers with $p \geq p^*$ and reject them otherwise. Moreover, restrict $H(p'|0)$ to be symmetrically distributed in $[0, 1]$ and let $H(p'|\tau)$ be a mean preserving spread of $H(p'|0)$ for all τ. Here the job finding rate can either be increasing or decreasing with τ depending on whether p^* is above or below the mean of the distribution of job offers. As τ increases, the worker gains more information about the quality of the job, so there is more weight of the distribution near 0 and 1. If p^* is above the mean, there is an increasing probability mass concentrated above p^* so the job finding rate will increase with τ. Alternately, if p^* is below the mean the job finding rate will be decreasing.

3 Parameterization

To compute the model, the wage process and information processes are specified and parameters are chosen to match key features of job flows in the United States.

3.1 Wages

Output in each period is composed of the average productivity from the job plus a noise term. Output in each period x_t for a job of quality μ is given by:

$$x_t = \mu + z_t$$

where $z_t \sim N(0, \sigma^2)$ is independently and identically distributed noise on the output process. Therefore, $x_t \sim N(\mu, \sigma^2)$. Given this output process, the average wage received from a worker is given by:

$$w(p) = p\mu_h + (1-p)\mu_l$$

This wage process can arise either from the model with no firms where workers get what they produce each period. Alternately, the assumption of competitive firms in Jovanovic (1979) can be invoked so that workers receive exactly the average expected output each period.

3.2 Information

Given the wage process, I assume workers observe the output from their job in each period and update their beliefs based on this observed output. Given the normality of output noise, for any belief, p, the expected distribution of output is given by:

$$\psi(x|p) = p\frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{1}{2}\left(\frac{x-\mu_h}{\sigma}\right)^2} + (1-p)\frac{1}{\sigma\sqrt{2\pi}}e^{-\frac{1}{2}\left(\frac{x-\mu_l}{\sigma}\right)^2}$$

With probability p output is drawn from a normal distribution with mean μ_h and variance σ, while with probability $1-p$ it is drawn from a normal with mean μ_l and the same variance.

Using this known distribution of output, the worker observers her realization of output and uses it to update her belief about the probability that she has a good match using Bayes' rule. Given any current belief, p, and observed output for a given period, x, the updated belief, p', is is formed by conducting a probability ratio test:

$$f(p,x) \equiv p' = Prob(\mu = \mu_h | p, x) = \frac{pe^{-\frac{1}{2}\left(\frac{x-\mu_h}{\sigma}\right)^2}}{pe^{-\frac{1}{2}\left(\frac{x-\mu_h}{\sigma}\right)^2} + (1-p)e^{-\frac{1}{2}\left(\frac{x-\mu_l}{\sigma}\right)^2}}$$

Here the numerator is proportional to the joint probability of observing output x and the match being good where the denominator is the total probability of observing output x.

With this updating function, define the inverse function $f^{-1}(p';p)$ to be the x required to have posterior p' given prior p. Then the p.d.f. of the G distribution, g, is given by:

$$g(p'|p) = \psi(f^{-1}(p'|p)|p)$$

Finally, the distribution of p for new job offers must be described. When unemployed, I assume that the worker's signal is equivalent to observing $\alpha\tau$ signals from the output process. The normality assumption on output noise is important so that a non-integer number of signals is well defined. Moreover, normality implies that to update beliefs after viewing t observations the worker only needs to know her prior belief p, the average value of the observation $\bar{x}$, and the number of observations observed t not the entire list of observations

$x_1, x_2, \ldots, x_t$. For a worker who observes t periods of output, the distribution of the average output per period, $\bar{x}$, is given by:

$$\tilde{\psi}(\bar{x}; p, t) = p\frac{1}{\frac{\sigma}{\sqrt{t}}\sqrt{2\pi}}e^{-\frac{1}{2}\left(\frac{\bar{x}-\mu_h}{\frac{\sigma}{\sqrt{t}}}\right)^2} + (1-p)\frac{1}{\frac{\sigma}{\sqrt{t}}\sqrt{2\pi}}e^{-\frac{1}{2}\left(\frac{\bar{x}-\mu_l}{\frac{\sigma}{\sqrt{t}}}\right)^2}$$

Using the same strategy, the posterior after observing the average output from t periods is computed as:

$$\tilde{f}(p, \bar{x}, t) = \frac{pe^{-\frac{1}{2}\left(\frac{\bar{x}-\mu_h}{\frac{\sigma}{\sqrt{t}}}\right)^2}}{pe^{-\frac{1}{2}\left(\frac{\bar{x}-\mu_h}{\frac{\sigma}{\sqrt{t}}}\right)^2} + (1-p)e^{-\frac{1}{2}\left(\frac{\bar{x}-\mu_l}{\frac{\sigma}{\sqrt{t}}}\right)^2}}$$

Again, inverting $\tilde{f}$ gives the value of $\bar{x}$ needed to generate posterior p': $\tilde{f}^{-1}(p', p, t) = \bar{x}$. Hence the p.d.f. of the H distribution, h, is given by:

$$h(p'|\tau) = \tilde{\psi}(\tilde{f}^{-1}(p', p0, \alpha\tau); p0, \alpha\tau)$$

where α and p_0 are parameters. $\alpha \in [0, 1]$ determines the fraction of experience that carries over from past jobs into information about new offers. p_0 is the prior percentage that any new job is good.

3.3 Parameters

To parameterize the model I assume that there are a large number of workers facing identical decision problems. Each of these workers faces a different history of idiosyncratic shocks. Averaging outcomes across workers, aggregate data can be constructed from the model. In

computations, I compare simulated data from the model over a 40 year career to actual worker outcomes. To solve the model, I restrict the maximum amount of experience to T periods so that the value functions can be solved backwards. With experience T the worker does not accumulate experience so her employed value function is given by:

$$V(p,T) = w(p) + \beta\delta U(T) + \beta(1-\delta)\int_0^1 \max\{U(\tau+1), V(p',T)\}G(dp'|p)$$

I also set the period length to be one month so that parameters can be chosen to match monthly data on job finding and separation rates in the United states.

With these restrictions, parameters are chosen to match employment statistics in the U.S. economy. To compute the model there are ten parameters that must be chosen: the maximum amount of experience T, the discount factor β, the job offer rate λ, the average output from a good match μ_h, the average output from a bad match μ_l, the probability that a match is good p_0, the variance of output noise σ, the proportion of experience used for new matches α, the exogenous separation rate δ, and the value of leisure b.

μ_h is normalized to one. Because the model period is one month, β is set to 0.9966 which corresponds to an annual interest rate of 4%. T is chosen to be large enough so that it will not impact individual decisions during the 40 years of work experience that I consider. I set $T = 480$.

The remaining parameters are chosen to match features of the decline in job finding and separation rates in the U.S. Figure 1 shows the decline in the job separation rate with age

in the U.S. for workers aged 18-57[3]. The separation has a sharp initial decline for 8-10 years followed by a gradual decline.

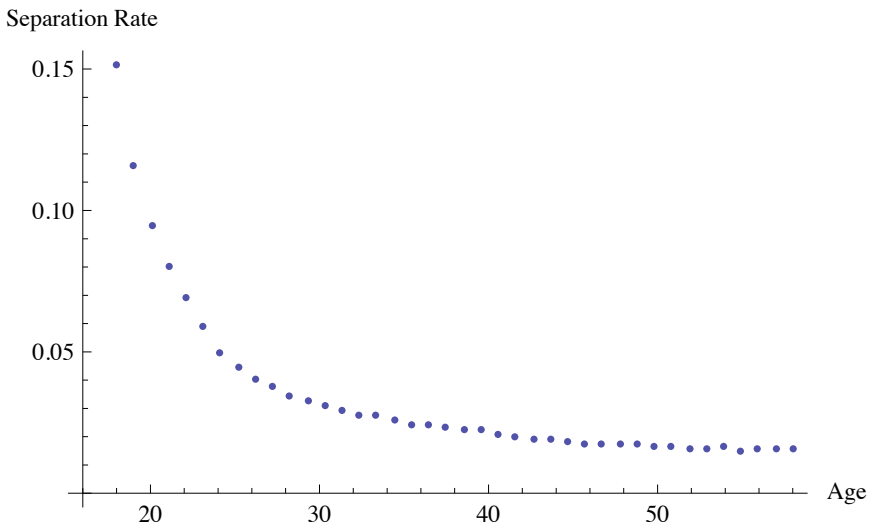

Figure 1: Average job separation rate by age for the U.S. economy.

Figure 2 shows the decline in the job finding rate. Similarly, the job finding rates fall fastest for the first 8-10 years, but the initial decline is less dramatic than the separation rate and finding rates continue to decline at a greater rate for the remainder of the workers' careers. Taken together, the steeper decline in the separation rate implies that the unemployment rate declines with age.

λ is chosen to to match the worker's rate of job offers. In the data, 16-year-old workers have the highest job finding rate of 0.612. According to the model, these workers with no experience should accept any job offered to them. Hence, to match this feature of the data,

[3]This data was constructed by Robert Shimer using CPS monthly microdata from 1976 to 2005. The procedure used follows Shimer (2007) to create a time series of job separation and finding rates for individuals of each age. The time series is then averaged to create average unemployment, job finding, and job separation rates for each age group. For additional details, please see Shimer (2007) and his webpage http://robert.shimer.googlepages.com/flows.

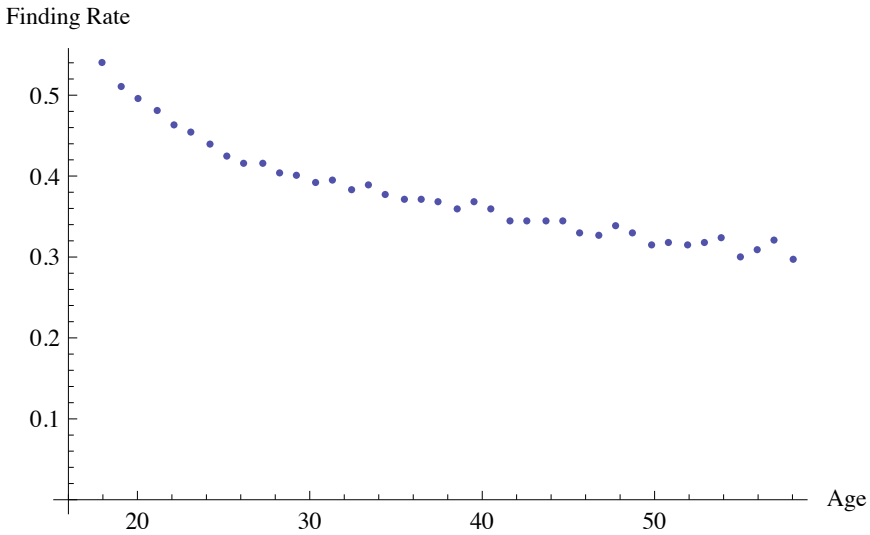

Figure 2: Average job finding rate by age for the U.S. economy.

I set $\lambda = 0.61$.

Workers with an infinite amount of experience workers are able to perfectly distinguish between good and bad jobs. In this case, they would accept only good jobs so their job finding rate is given by $p_0\lambda$. I use this hypothetical limit to set the value of p_0. The job finding rates data implies that the lowest job finding rate is 0.283 for 58 year old workers. Assuming that they are perfectly distinguishing between good and bad job offers, I set $p_0 = 0.46$. In this sense, p_0 determines the magnitude of total decline in the job finding rate over an individual's lifetime.

I next choose the output for bad matches, μ_l. μ_l is chosen to determine the amount of wage growth generated by the model. The evolution of p is fully determined by the signal to noise ratio: $\frac{\mu_h - \mu_l}{\sigma}$. Hence given the normalization of μ_h, for any choice of the μ_l there is a value of σ that generates identical information for the worker. σ and b are then calibrated

to determine worker search behavior while I set μ_l to determine the dispersion of wages in the model. μ_l provides a lower bound on possible wage realizations while μ_h is the upper bound. $\mu_l = 0$ is chosen so that wages increase by 66% over the life cycle. This implies that most of wage life cycle wage growth in the data is accounted for by sorting and learning in the model, but there is still room for other factors like human capital accumulation to play a role.

Next, σ is the amount of output noise. Higher values of σ imply that workers learn slowly about the quality of their matches. In the limit, $\sigma = 0$ implies that workers perfectly observe the quality of the match with one observation while as $\sigma \to \infty$ workers have no learning. $\sigma = 2$ is chosen so that the peak of job separations matches that found in the data. Higher values of σ imply that workers learn more slowly. Slower learning implies that it takes longer to distinguish bad matches, and therefore workers are willing to stay in initial matches longer before quitting.

α determines the amount of experience that carries over in learning about new job opportunities. It is natural to restrict α to be in in $[0, 1]$. $\alpha = 0$ is analogous to the standard Jovanovic (1979) model where individuals learn nothing about future jobs and the employment is a pure renewal process. $\alpha = 1$ is the limit where all learning carries over to future jobs. Higher values of α imply that workers learn faster about future jobs and therefore have a steeper decline in both job finding and separation rates. With the model period set to be a month, $\alpha = \frac{1}{30}$. This corresponds to getting on average one month worth of information about a new job for every two and a half years of labor market experience. This parameter

is chosen so that the model matches the curvature in the decline in separation rates. Higher values of α predict a steeper initial decline followed by less learning later. This parameter is sensitive to the choice of σ. The chosen value of σ implies that individuals learn quickly by observing output. To create enough curvature on the separation rates I choose a low value of α.

δ is the rate of exogenous job separations. It determines the level of job separations in the model. Using data on average monthly job finding probabilities by age in the population, the lowest observed number in the data is 0.014 for 59-year-olds. This should be an upper bound on the value of δ. I choose $\delta = 0.009$ so that the level of separations for experienced workers matches the data.

The final parameter that must be set is b. This parameter determines the relative desirability of being employed in a bad job compared to searching for a new job. Higher values of b make unemployment more attractive. $b = 0.35$ is set to match the average level of job finding rates over the career. The model is not highly responsive to changes in the value of b.

Table 1 summarizes the chosen parameters and their values.

4 Simulation Results

After computing the value functions and reservation probabilities for workers at each experience level I simulate employment outcomes for individual workers. When simulating the model, 80% of workers start off employed while 20% start off unemployed. This matches the

Name	Parameter	Value
Max Experience	T	480
Discount Factor	β	0.9966
Job Offer Rate	λ	0.61
Good Output	μ_h	1
Bad Output	μ_l	0
Probability of Good Job	p_0	0.46
Output Noise	σ	2
Experience Rate	α	$\frac{1}{30}$
Exogenous Separation Rate	δ	0.009
Value of Leisure	b	0.35

Table 1: Calibrated values of the model parameters.

steady state level of unemployment that is generated from a job finding rate of 0.15 and a job separation rate of 0.61. Starting with some workers employed avoids having higher than normal initial levels of unemployment. When generating outcomes I keep track of employment, job finding, job separation, wages, tenure, and total experience. I simulate the model for 10,000 workers and calculate average outcomes from the date that workers enter the labor force. To compare average outcomes with labor force data, I construct average outcomes by age by entering workers into the labor market at the age they get their first full time employment. Topel and Ward (1992) compute the percentage of workers who enter the labor force at a given age by assuming that workers enter when they attain their first employment that lasts at least 2 quarters. This measure leaves out workers who take summer jobs and then return to school. Table 2 replicates their table showing the percentage of workers who enter the labor force at each age. When constructing the data from the model I assume that all workers in the ≤ 18 category enter at age 18 and that all workers in the ≥ 25 category enter at age 25.

≤ 18	19	20	21	22	23	24	≥ 25
29.6	24.9	18.8	11.4	8.1	4.8	1.7	0.7

Table 2: Percent of populations first employment spell by age. Taken from Topel and Ward (1992).

For the remainder of the section, I compare results from the calibrated model where $\alpha = \frac{1}{30}$ to the data and to the model where $\alpha = 0$. $\alpha = 0$ is an interesting benchmark as it corresponds to the standard Jovanovic (1979) model where there is no learning about future matches. This highlights the novel effects of learning. I find that this model is better able to match employment data than the standard one.

4.1 Unemployment

It is well known that young workers face higher unemployment rates than prime aged workers. The model is able to capture this decline in unemployment with age.

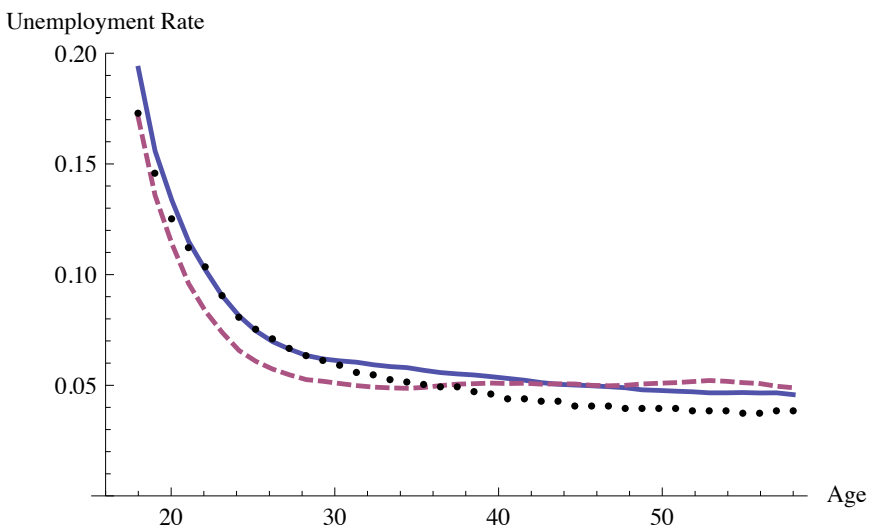

Figure 3: Average unemployment rate by age. Data: dots; Calibrated Model: Line; $\alpha = 0$ Model: Dashed.

Figure 3 shows the average annual unemployment rate by age. The dots depict the decline in unemployment found in the data where the solid and dashed lines depict the results from the model with $\alpha = \frac{1}{30}$ and $\alpha = 0$ respectively. The data show a steady decline in unemployment with age. Unemployment declines from about 17% for 18-year-old workers to between 3.5 and 4% for prime aged workers. The calibrated model captures a similar decline over the life cycle, with 18-year-old workers experiencing unemployment of 19% and declining to 4.4%. There are two mechanisms in the model that generate this decline. First, as workers are in the labor market for longer they are more likely to sort themselves into good jobs. These jobs have a lower rate of separation. Hence, fewer workers are unemployed. A second mechanism whereby workers are more likely to be in good jobs with more experience is the learning mechanism. Experienced workers are better able to distinguish between good and bad jobs so are more likely to start off in good jobs than inexperienced workers.

The model with $\alpha = 0$ captures much of the decline in unemployment. Unemployment drops rapidly during the first 10 years of work experience with little to no decline later in life. In this case, only the mechanism is present. The figure shows that workers are able to sort themselves into good jobs within the first 10 years of work. To generate a continued decline in unemployment rates the learning mechanism is necessary.

To compare the fits of the two models with the data I construct a measure of the goodness of fit:

$$Fit = 1 - \frac{\sum_{a=18}^{57}(\varepsilon_a - \bar{\varepsilon})^2}{\sum_{a=18}^{57}(y_a - \bar{y})^2}$$

This is similar to an R^2 measure, where ε_a is the difference between the model and the data

for age a, $\bar{\varepsilon}$ is the average difference, y_a is the level of the data for age a, and $\bar{y}$ is the average level of the data. So the numerator give the sum of squared errors between the data and the model and the denominator gives the sum of squared deviations in the data.

Both models are able to account for most of the decline in the unemployment rate with age. The calibrated model fits slightly better with a measure of 0.98. The model with $\alpha = 0$ has a measure of 0.88.

4.2 Labor Market Flows

While both the calibrated model and the $\alpha = 0$ model capture the decline in unemployment, they have different implications for labor market flows. It is well known that turnover declines with age (see Clark and Summers (1982)). Job separations exhibit a sharp initial drop and continue to decline with age. It is also the case that job finding rates decline. The effectiveness of the two models in capturing the decline in separation rates is shown in Figure 4

Figure 4 shows the decline in separations for both models compared with the data. Both models have an initial decline in the separation rate that is steeper than the data. The calibrated model accounts for nearly the entire magnitude of the decline. In contrast, the model with $\alpha = 0$ is unable to capture the entire variation in the separation rate with age as the separation rate becomes flat after the first 10 years. The mechanism for this is the same as with unemployment. Both models are characterized by the selection of individuals into good matches as they spend more time in the labor market. However, the calibrated model

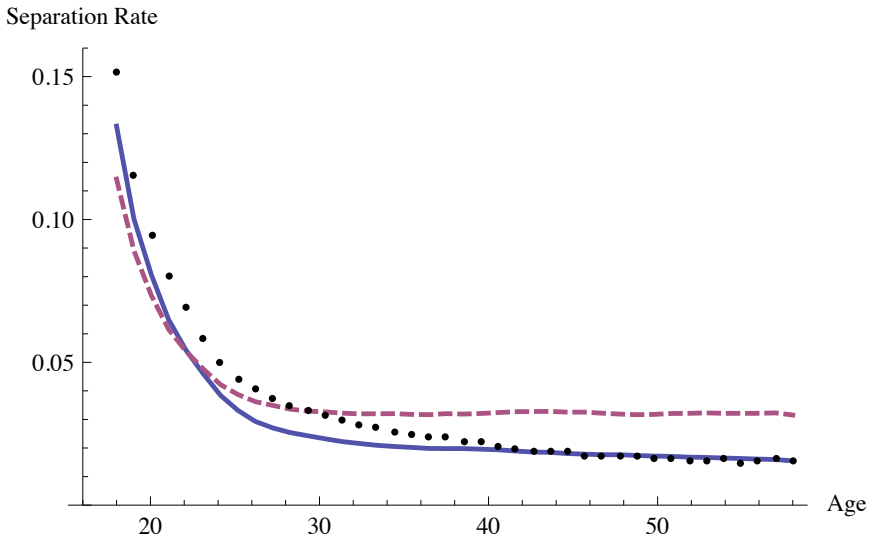

Figure 4: Average job separation rate by age. Data: dots; Calibrated Model: Line; $\alpha = 0$ Model: Dashed.

is able to capture the continued decline in the separation rate as workers continue to learn about the quality of future matches once the effects from selection have ended.

In computing the fit of both models with the data it is clear that the calibrated model outperforms the standard model. In the calibrated case the fit is 0.96 where when $\alpha = 0$ the fit is 0.80. So while the $\alpha = 0$ model fits the data quite well, the full calibration is able to account for almost all of the variation in the data.

Job finding rates also decline with age. They, however, decline more slowly as the more rapid decrease in separations leads to declining unemployment over the life cycle. New to the model is the ability to account for this decline in job finding rates. Since experience allows individuals to learn about the quality of new matches, experienced workers can be selective about which jobs they choose to accept.

Figure 5 shows the decline in job finding rates. The solid line shows that the decline

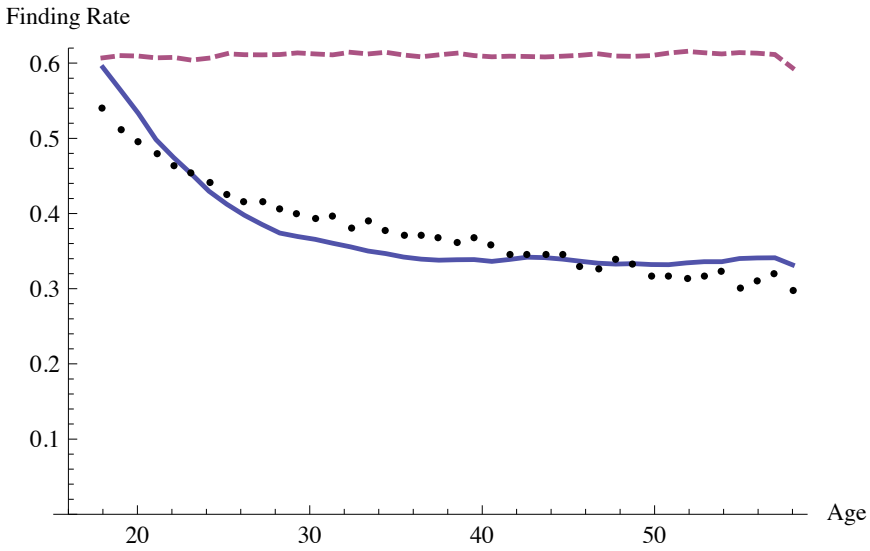

Figure 5: Average job finding rate by age. Data: dots; Calibrated Model: Line; $\alpha = 0$ Model: Dashed.

from the calibrated model is initially steeper with less decline later in the career than found in the data. Despite the difference in shape of the decline, the model with learning better matches the decline in finding rates. The dashed line shows the model when $\alpha = 0$. In this case, the job finding rate is simply $\lambda = .61$ as the worker is willing to accept any job offered since she is unable to distinguish between them. The standard model is unable to deliver any decline in the separation rates.

Comparing the fit between the two models reveals that the calibrated model has a fit of 0.82 compared to a fit of $0.03 \approx 0$ in the standard model. The negative outcome from the standard model is a result of noise from the simulations. It should deliver an exact zero as job finding rates are λ for all ages.

4.3 Separations by Tenure

One of the primary motivations for Jovanovic's (1979) paper was to explain the declining hazard rate of unemployment with tenure. This model continues to capture the negative relationship between separations and tenure. The primary explanation for this feature in the model is that within any job, workers learn their productivity as tenure increases.

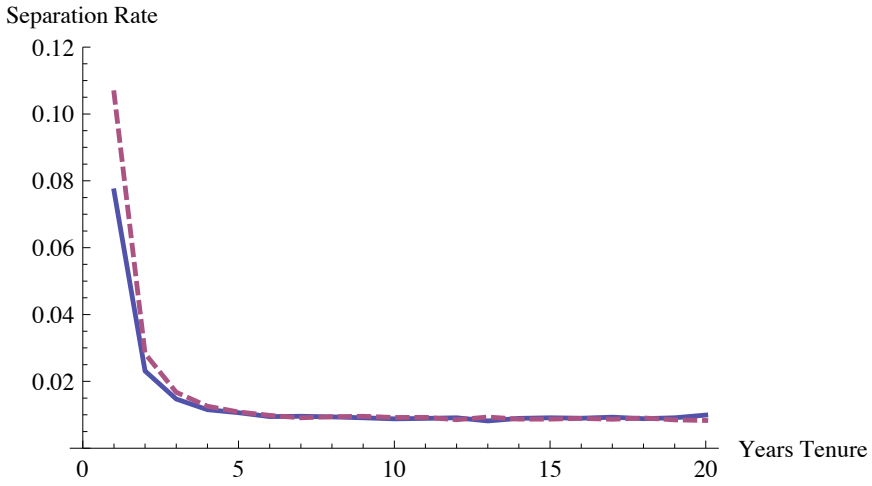

Figure 6: Average separation finding rate by years of tenure from the model. Calibrated Model: Line; $\alpha = 0$ Model: Dashed.

Figure 6 depicts the average monthly separation rate by year of tenure for agents in the two models. It shows that the probability of separation is more than twice as high during the first year on a job than during any other year. As tenure accumulates, the job separation rate continues to decline slowly after the first five years. The decline in both models is very similar to that found in Moscarini (2003), which compares well with the data.

4.4 Wage Growth

Topel and Ward (1992) document a number of features of wage profiles during worker's first 10 years of experience. They document that the first 10 years of the career account for two-thirds of lifetime wage growth. These job changes can explain about one-third of wage growth. Moreover, wages on the job approximate a random walk. The model qualitatively replicates the behavior of wages over the life cycle.

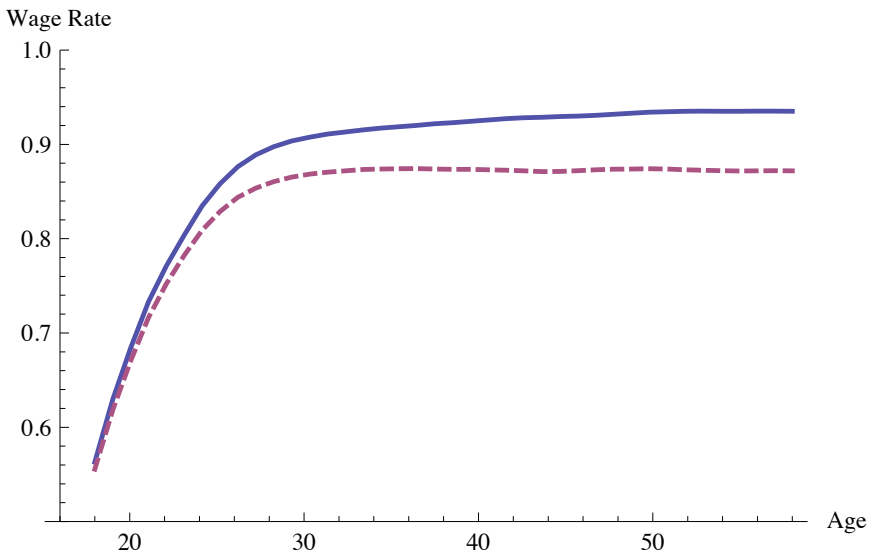

Figure 7: Average wages by age. Calibrated Model: Line; $\alpha = 0$ Model: Dashed.

Figure 7 shows the average annual wages by age from both models. I calibrate the $\alpha = \frac{1}{30}$ model so that wages grow by 66%; the pattern of wage growth from the model is endogenous. The model generates rapid wage growth during the first 10 years of experience and then levels off. Wage growth in during the first 10 years in the calibrated model accounts for about 89% of total wage growth instead of the 66% found in the data. The standard model does not generate as much wage growth as the calibrated model. Wages grow rapidly for the first 10

years then stop growing completely. The standard model delivers wages that grow by only 55% over the life cycle. The first ten years accounts for 93% of total wage growth. Again, the secondary learning effect is crucial to deliver continued wage growth in the model.

5 Minimum Wages

Finally, I explore the effects of wage restrictions in the model and find that having endogenously determined job finding rates is essential to understanding the implications of minimum wages for young workers' employment outcomes. To explore minimum wages, it is necessary to introduce wage setting into the model. I follow the convention from Jovanovic (1979) that there are competitive firms who compete for the services of a worker. Specifically, workers receive job offers from industries composed of many firms that compete over the services of the worker. Therefore, the worker is paid her expected output in each period. The worker is paid their expected marginal product each period on the job[4].

While the worker flows in the model are calibrated to match the employment data for the U.S., the nature of these flows varies across countries. Cohen et al. (1997) compare French and U.S. labor markets. They find that young workers in France have much higher unemployment rates than those in the U.S. When breaking down the factors that contribute

[4]Alternately, I could introduce firm worker matches and bargaining as in Moscarini (2005). This model introduces a few more modeling choices as workers with different levels of experience have different values of unemployment. The main choice is whether to use one matching function for all workers or a separate matching function for each worker type. If one matching function is used, the overall distribution of worker types must be tracked while if multiple are used this is not an issue. This method endogenizes the job offer rate λ and the choices of the matching function have different implications. One matching function implies that distortions on any worker can have effects for all workers' search behavior, while having multiple matching functions localizes these effects. There is no clear reason to choose one over the other, though it is interesting to think about how policies alter matching rates.

to this difference, they show that the U.S. is characterized by more rapid job finding and separation rates than France. Table 3 and Table 4 show the distribution of unemployment durations in the U.S. and France respectively using O.E.C.D. Labour Statistics. The differences are striking. In both countries, young workers have shorter durations that prime aged workers, but unemployment durations are much longer in France. Over 40% of unemployed workers aged 15-24 in the U.S. are able to find a job within a month while only 7% find a job in France. Additionally, only about 8% of unemployment spells for 15-24 year-old workers in the U.S. last over a year where in france it is over 23%. Pries and Rogerson (2005) show that labor market policies can have large impacts on labor market flows.

Duration	Total	Age 15-24	Age 25-54
$<$ 1 month	4.4	7.0	3.7
1-3 months	18.3	29.1	15.7
3-6 months	16.7	21.8	15.9
6 months - 1 year	19.7	18.5	20.7
$>$ 1 year	40.9	23.6	44.0

Table 3: Incidence of unemployment by duration and age in France from OECD Labor Statistics.

Duration	Total	Age 15-24	Age 25-54
$<$ 1 month	33.1	41.0	29.7
1-3 months	29.2	31.9	28.4
3-6 months	15.9	12.9	17.3
6 months - 1 year	9.2	6.3	10.3
$>$ 1 year	12.7	7.9	14.2

Table 4: Incidence of unemployment by duration and age in the United States from OECD Labor Statistics.

To understand the importance of endogenous job finding rates in predicting the effect

of labor market policies, the implications of minimum wages are examined. The levels of minimum wages vary dramatically across countries. In the U.S., minimum wages are low and have been declining in real terms for much of the last 25 years. In 2002-2003 the ratio of the minimum wage to the median in the U.S. economy was 0.335. In contrast, several European countries have minimums that are higher than forty percent of the median wage. In particular, the French minimum wage is 0.62 of the median. While differences in wage dispersion across countries may make these figures difficult to compare directly, minimum wages in Europe are much more restrictive than in the U.S. In the U.S. only about 1.5% of workers are paid the minimum wage compared with 14% in France.

I compare the effects of a high minimum wage in the calibrated model and in the model where $\alpha = 0$. Workers enter the labor force employed with one unit of experience. Giving workers a unit of experience to start means that their initial job offers are not drawn from a degenerate distribution. Therefore, I can consider minimum wages that are higher than the wage from a job with believed to be good with probability p_0 without shutting down all employment in the model. To restrict wages in the model, I set a lower bound on the probability of a good match of jobs that workers are allowed to take. A minimum wage of $\underline{p} = 0.45$ is considered for both the calibrated model and the model with $\alpha = 0$. This places a large restriction on the jobs that workers are able to accept. I find that the model without learning is unable to generate realistic employment effects from the minimum wage because there is no decline in the job finding rate. In the model with experience, the initial decline in job finding rates for young workers generates the pattern of unemployment found

in European countries with high minimum wages.

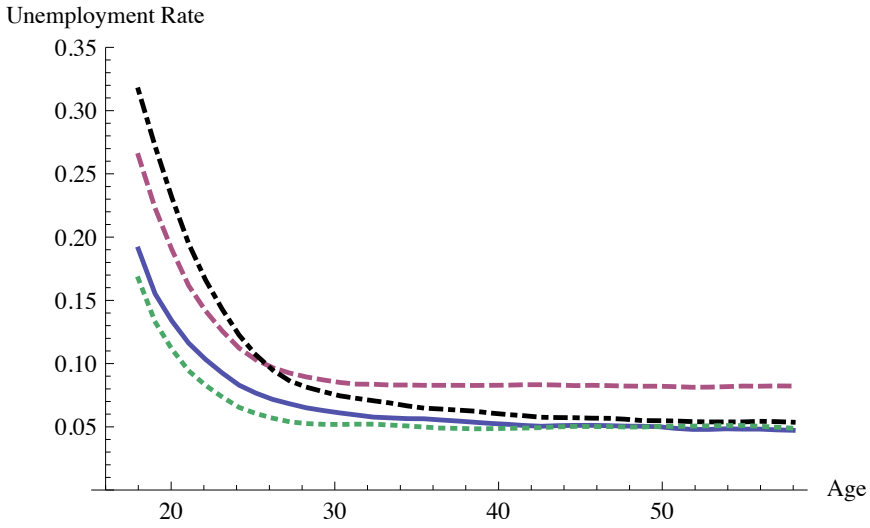

Figure 8: Unemployment rate by age for calibrated model with no minimum wage (solid), calibrated model with minimum (dot-dashed), $\alpha = 0$ model with no minimum (dotted), and $\alpha = 0$ model with minimum (dashed).

I first evaluate the effects of minimum wages on unemployment. Figure 8 compares the unemployment rate by age for both models with unemployment from the calibrated model with no minimum wage. While unemployment increases in both models, the pattern of the change is dramatically different. The dot-dashed line shows that the employment effects from the calibrated model are large for young workers then converge to the model with no minimum wage. In contrast, the dashed line shows that in the model with $\alpha = 0$ the minimum wage causes unemployment to be uniformly higher throughout the worker's life. The calibrated model much more closely resembles patterns in unemployment in European countries with high minimum wages. In these countries, young workers have higher rates of unemployment than in the U.S. but prime aged workers have similar unemployment rates.

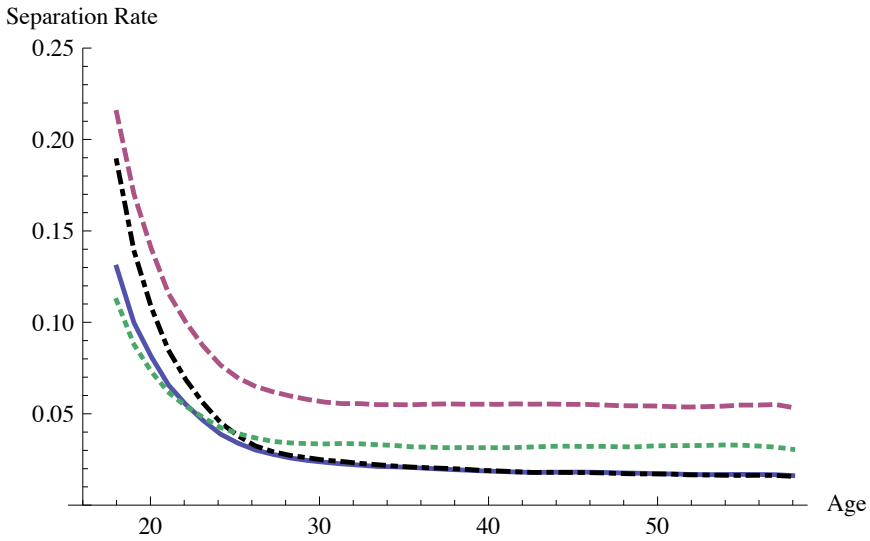

Figure 9: Job Separation rate by age for calibrated model with no minimum wage (solid), calibrated model with minimum (dot-dashed), $\alpha = 0$ model with no minimum (dotted), and $\alpha = 0$ model with minimum (dashed).

To understand the origin of the differences in unemployment between the two models, I examine their predictions on job separation and finding rates. Figure 9 shows the effect of minimum wages on job separation rates. Here the pattern is again similar to the unemployment rates. For the calibrated model, the minimum wage causes higher separations. In contrast to the unemployment figure, the differences in separations from the model with no minimum wages lasts less than 10 years. The dashed line shows that minimum wages increase the separation rate in the $\alpha = 0$ model for all ages. Without a minimum wage workers are free to gain experience on a job with low productivity. Minimum wage restrictions prevent this early accumulation of experience. These high initial separation rates are important in causing the increase in unemployment for young workers in both models.

Finally, the effects of minimum wages on job finding rates are explored. Figure 10 shows

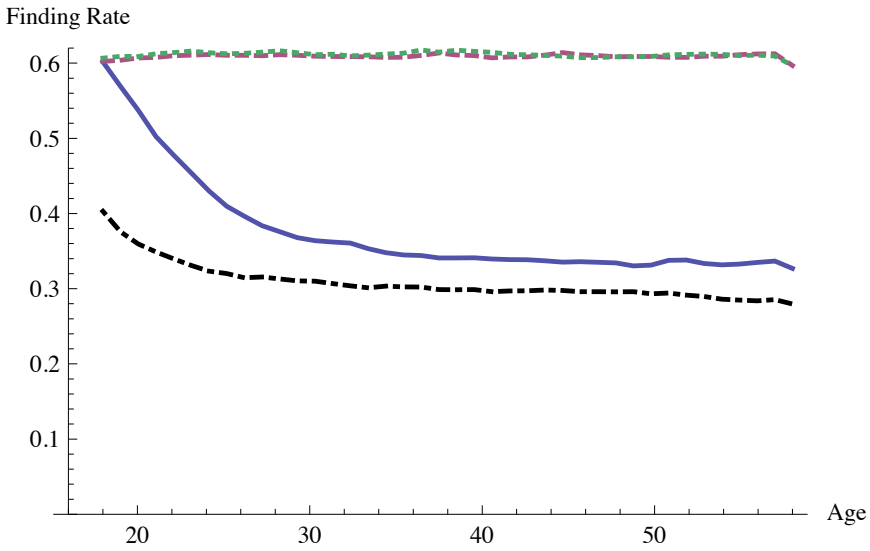

Figure 10: Job Finding rate by age for calibrated model with no minimum wage (solid), calibrated model with minimum (dot-dashed), $\alpha = 0$ model with no minimum (dotted), and $\alpha = 0$ model with minimum (dashed).

job finding rates for both models. Here the predictions of the two models are drastically different. The calibrated model predicts a decline in job finding rates, but job findings remain fixed in the model where $\alpha = 0$. The failure of job finding rates to move makes the model without experience unable to generate accurate predictions of the effects of minimum wages across countries.

Age	No Minimum	Standard MW	$\alpha = 0$ no Minimum	$\alpha = 0$ MW
15-24	11.9	19.4	9.9	16.5
25-34	6.2	7.6	5.2	8.7
35-44	5.2	6.0	4.9	8.3
45-54	4.9	5.5	5.0	8.2

Table 5: Unemployment rate by age band from the calibrated model with no minimum wage, the calibrated model with a minimum wage, and the $\alpha = 0$ model with and without a minimum wage.

The endogenous decline in finding and separation rates is crucial to understanding the

pattern of unemployment changes with minimum wages. Table 5 computes average unemployment rates from the model for different age groups. In computing the table it is assumed that there are equal numbers of people at each age within each age band. Demographic changes are not accounted for. Higher levels of minimum wages are shown to have large effects on unemployment for workers aged 15-24 and modest effects for 25-34 year olds in the model with experience. In contrast, the unemployment differences for young workers are not as large in the model where $\alpha = 0$, and the effects do not completely disappear with age.

Age	U.S.	France
15-24	13.0	20.9
25-34	5.8	9.7
35-44	4.6	6.9
45-54	4.1	6.1

Table 6: Unemployment rate by age band for U.S and France. Average value for male workers from 2003-2004 from OECD.

The patterns of unemployment as a result of minimum wages are consistent with cross-country data. Table 6 shows the average unemployment rate by age band for U.S. and French male workers in 2003-2004. In the U.S. where workers face a modest minimum wage, unemployment is 13% for the youngest age group and converges quickly to under 6%. In contrast, in France where there is a high minimum wage young workers face unemployment rates above 20% that declines slowly with unemployment rates of almost 10% for workers aged 25-34 before to between six and seven percent for older workers. This pattern is consistent across European countries that typically have high minimum wages. These countries have

higher rates of unemployment than the U.S. and these differences are concentrated in young workers. Understanding these differences in employment outcomes requires a model that is able to capture key changes in job finding and separation rates over the life cycle.

Gorry (2008) provides a more detailed quantitative analysis of the effects of minimum wages on youth employment outcomes. Taking the decline of job finding and separation rates as a feature of the economy, he shows that minimum wages can account for a significant portion of the differences in employment rates between the U.S. and European countries.

6 Lasting Effects

In this section, I evaluate the model's predictions on the effects of a worker who graduates in a bad economy. Thus far I have assumed that the job offer rate λ is constant throughout the individual's working life. However, recent research shows that job finding rates vary substantially over the business cycle. Shimer (2007) decomposes fluctuations in unemployment since 1948 and finds that three-quarters of the fluctuations are accounted for by changes in job finding rate. Similarly, Hall (2005) argues that jobs are difficult to find during recessions because of low job finding rates rather that high job separations.

Given the fluctuations in job finding rate, I consider the effects of a worker who enters the labor force in a bad economy by evaluating the effect of facing low job finding rates for the first two years of work. In particular, I assume that for the first two years after entering the labor force the exogenous job offer rate is half the calibrated value followed by an unexpected permanent change back to the level in the original calibration. In this section, workers enter

the labor market unemployed with no past experience. Career earnings and employment outcomes are compared for workers who face this initial low job offer rate with workers who enjoy high job finding rates for their entire career.

First, I find that workers who enter the labor force in bad times have only a modest decrease in accumulated experience over their career. On average, graduating in a weak economy implies a loss of 0.86 months of experience during the first year and 1.8 months during the first two years. After that, the difference in experience grows to just 2.70 months over 40 years in the labor market. Workers who face poor job prospects in their first two years of labor market experience lose a modest amount of experience during these years and almost none after.

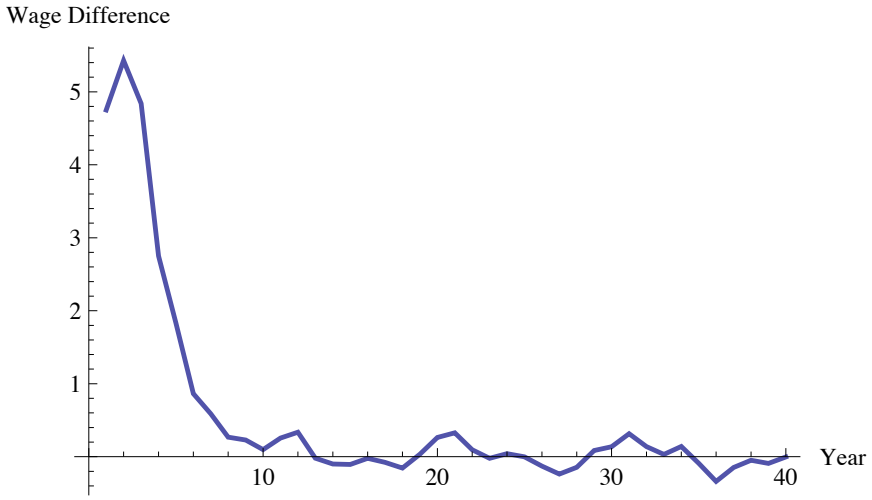

Figure 11: Log difference in wages of workers with normal finding rates compared to those who graduate in a weak economy by year of labor market experience.

Figure 11 shows the log difference in wages between workers who face the calibrated job offer rate for their entire life and workers who graduate in a bad economy. It shows that

wage losses grow to over five percent during the years when workers face diminished job finding prospects then decline to nearly zero in the next eight to ten years. Workers' wage losses grow during the first two years when they face poor job finding prospects. These initial job losses remain as workers have lower experience from these years, but the wage outcomes converge back to the levels for workers who never face poor job finding prospects. The fluctuations after the first 10 years in the graph are noise from the simulations.

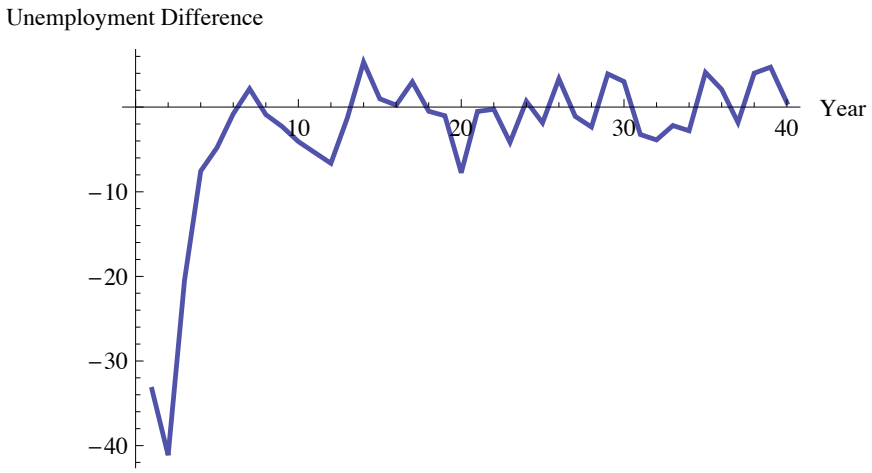

Figure 12: Log difference in unemployment rate of workers with normal finding rates compared to those who graduate in a weak economy by year of labor market experience.

Next, I examine the employment effects of graduating in a recession. Figure 12 shows the log difference in unemployment between the two models. Graduating in a bad economy implies that unemployment rates are about 40% higher during the downturn. However, once job offer rates return to normal, unemployment rates quickly revert to normal levels. This occurs within two to three years. Workers who were unable to find a job initially have lower experience and hence have more motivation to quickly find a job and gain experience when prospects improve. Hence, there are no lasting effects on unemployment.

The results from the model are consistent with a growing empirical literature that studies the effects of unemployment on young workers. Kahn (2006) studies the effects of graduating from college in a bad economy using NLSY data between 1979 and 1988. She finds large negative wage effects of graduating in a bad economy but no lasting effects on labor supply. Similarly, Oreopoulos et al. (2005) study the effects of graduating college in a recession using Canadian university-employer-employee matched data from 1982 to 1999. They also find significant wage effects that fade after 8-10 years with little impact on time worked by those who faced high unemployment early in their career.

7 Conclusion

This paper presents a model of learning that has novel implications for workers' job finding rates. Workers' learning about the quality of their match is not only important for observed outcomes like wages and employment durations while employed; it is also important for their behavior while unemployed. This insight motivates the model where experience gives workers both knowledge about the quality of their current job and the ability to distinguish between good and bad jobs when unemployed.

A model with learning about both the quality of the current match and future matches has rich implications for labor market outcomes. It is consistent with the age profiles of unemployment, job finding rates, job separation rates, hazard rates of separation with tenure, wage dispersion, and wage growth. Having a model that has consistent prediction about a broad range of labor outcomes makes it ideal to analyze the effects of policy on these

outcomes. This paper explores the effects of graduating in a bad economy and the implications for wage restriction on employment outcomes. The model generates results that are consistent with empirical findings in both cases: early losses of experience have lasting consequences for wages but not for employment and high minimum wages imply high rates of unemployment for younger workers.

The model will be fruitful for further studies on the effects of labor market policies. In particular, the model has implications for optimal unemployment insurance. Current policy in the U.S. has different payments depending on past experience, but all workers receive the same duration of unemployment benefits. This model has strong implications for the duration of unemployment based on past experience. This model provides a mechanism to study the effects of making the duration of unemployment benefits dependent on past experience.

References

Altonji, Joseph G. and Nicolas Williams, "Do wages rise with job seniority? A reassessment," *Industrial and Labor Relations Review*, April 2005, *58* (3), 370–397.

Altonji, Joseph G and Robert A Shakotko, "Do Wages Rise with Job Seniority?," *Review of Economic Studies*, July 1987, *54* (3), 437–59.

Bartel, Ann P. and George J. Borjas, "Wage Growth and Job Turnover: An Empirical Analysis," NBER Working Papers 0285, National Bureau of Economic Research, Inc September 1982.

Clark, Kim B and Lawrence H Summers, "Labour Force Participation: Timing and Persistence," *Review of Economic Studies*, 1982, *49* (5), 825–44.

Cohen, Daniel, Arnaud Lefranc, and Gilles Saint-Paul, "French unemployment: a transatlantic perspective," *Economic Policy*, October 1997, *12* (25), 265–292.

Gibbons, Robert, Lawrence F. Katz, Thomas Lemieux, and Daniel Parent, "Comparative Advantage, Learning, and Sectoral Wage Determination," *Journal of Labor Economics*, October 2005, *23* (4), 681–724.

Gorry, Aspen, "Labor Market Connections, Minimum Wages, and Youth Employment," May 2008. Unpublished, University of Chicago.

Hall, Robert E., "Job Loss, Job Finding, and the Unemployment in the U.S. Economy over the Past Fifty Years," in Mark Gertler and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 20*, Cambridge, MA: MIT Press, 2005, pp. 101–37.

Jovanovic, Boyan, "Job Matching and the Theory of Turnover," *Journal of Political Economy*, October 1979, *87* (5), 972–90.

Kahn, Lisa B., "The Long-Term Labor Market Consequences of Graduating from College in a Bad Economy," September 2006. Unpublished, Harvard University.

Mincer, Jacob and Boyan Jovanovic, "Labor Mobility and Wages," NBER Working Papers 0357, National Bureau of Economic Research, Inc August 1982.

Moscarini, Giuseppe, "Skill and Luck in the Theory of Turnover," February 2003. Unpublished, Yale University.

_ , "Job Matching and the Wage Distribution," *Econometrica*, 2005, *73* (2), 481–516.

Neal, Derek, "The Complexity of Job Mobility among Young Men," *Journal of Labor Economics*, April 1999, *17* (2), 237–61.

Oreopoulos, Phil, Till von Wachter, and Andrew Heisz, "The Short- and Long-Term Career Effects of Graduating in a Recession: Hysteresis and Heterogeneity in the Market for College Graduates," 2005. Unpublished, Columbia University.

Papageorgiou, Theodore, "Learning Your Comparative Advantage," December 2007. Unpublished, Yale University.

Pavan, Ronni, "Career Choice and Wage Growth," July 2006. Unpublished, University of Rochester.

_ , "The Role of Career Choice in Understanding Job Mobility," July 2007. Unpublished, University of Rochester.

Pries, Michael and Richard Rogerson, "Hiring Policies, Labor Market Institutions, and Labor Market Flows," *Journal of Political Economy*, August 2005, *113* (4), 811–839.

Shimer, Robert, "Reassessing the Ins and Outs of Unemployment," October 2007. Unpublished, University of Chicago.

Topel, Robert H, "Specific Capital, Mobility, and Wages: Wages Rise with Job Seniority," *Journal of Political Economy*, February 1991, *99* (1), 145–76.

_ **and Michael P Ward**, "Job Mobility and the Careers of Young Men," *The Quarterly Journal of Economics*, 1992, *107* (2), 439–79.