WHAT SHOULD I BE WHEN I GROW UP? OCCUPATIONS AND UNEMPLOYMENT
OVER THE LIFE CYCLE

Martin Gervais
Nir Jaimovich
Henry E. Siu
Yaniv Yedid-Levi

What Should I Be When I Grow Up? Occupations and Unemployment over the Life Cycle
Martin Gervais, Nir Jaimovich, Henry E. Siu, and Yaniv Yedid-Levi
NBER Working Paper No. 20628
October 2014
JEL No. E0,J0

ABSTRACT

Why is unemployment higher for younger individuals? We address this question in a frictional model of the labor market that features learning about occupational fit. In order to learn the occupation in which they are most productive, workers sample occupations over their careers. Because young workers are more likely to be in matches that represent a poor occupational fit, they spend more time in transition between occupations. Through this mechanism, our model can replicate the observed age differences in unemployment which, as in the data, are due to differences in job separation rates.

Martin Gervais
University of Iowa
martin-gervais@uiowa.edu

Nir Jaimovich
Department of Economics
Duke University
213 Social Services Building
Durham, NC 27708
and NBER
njaimo@gmail.com

Henry E. Siu
Vancouver School of Economics
University of British Columbia
1873 East Mall #997
Vancouver, BC V6T 1Z1
Canada
and NBER
hankman@mail.ubc.ca

Yaniv Yedid-Levi
University of British Columbia
yaniv@mail.ubc.ca

1 Introduction

Labor market outcomes differ greatly for individuals of different ages. Unemployment rates are much higher for the young than for all others. For example, in the U.S., the unemployment rate for individuals aged 20–24 years old is approximately 2.5 times that of prime-aged individuals aged 45–54 years old; the unemployment rate of 25–34 year olds is about 50% greater than that of 45–54 year olds.[1]

In our view, explaining this fact is interesting in its own right, given that the age differences are so stark. Moreover, we believe that developing such an explanation is important for our understanding of aggregate labor market dynamics. For example, the results of Shimer (1998) indicate that compositional change in the labor force due to the baby boom generated a substantial fraction of the rise and fall in U.S. unemployment from the 1960s through to the end of the century.

In Section 2, we detail the declining age profile for unemployment observed in the postwar U.S. data. In particular, we discuss the role of age differences in the *job finding rate* (the rate at which individuals transition from unemployment to employment) and the *separation rate* (the transition rate from employment to unemployment) in shaping the age profile of unemployment. Job separation rates decline monotonically with age, generating a declining profile of unemployment. Job finding rates exhibit much less age variation, and in fact, decline with age. On its own, the declining pattern of job finding rates would counterfactually imply an increasing age profile of unemployment. Hence, unemployment rate differences are accounted for solely by age differences in separation rates.

Given this fact, Sections 3 and 4 present and characterize a model that focuses on differences in the separation rate. Our framework is a life-cycle version of the search-and-matching model of Diamond (1982), Mortensen (1982) and Pissarides (1985) (DMP, hereafter). The key life-cycle feature that we introduce is learning on the worker's part about her best occupational fit. Specifically, young workers enter the labor market not knowing the occupation that they are most productive in. We call this occupation the worker's "true calling".

In order to learn her true calling, a worker must sample occupational matches over her career. Upon entering the workforce, a worker searches for her first job in an occupation. Upon meeting a vacancy, a match is established. Over time, the worker and firm learn whether the current occupation is the worker's true calling. If it is not, the worker-firm pair

---

[1] See also Shimer (1998) and the references therein.

can either maintain the match or choose to separate. Upon separation, the worker seeks employment in a new occupation, having ruled out the previous occupation as being her true calling. As the worker samples more occupations and accumulates knowledge about her occupational fit, the probability of finding her true calling rises.

Match formation, learning, and separation are stochastic in our framework. As such, *ex ante* identical individuals experience different histories over time, as they move in and out of unemployment, and learn, more or less quickly, about their true calling. Despite this heterogeneity, workers can be summarized simply by their *type*: lower type workers have little information about their true calling, while higher types are closer to discovering it. Hence, the model generates an endogenous mapping between type and age.

This allows us to address the differences in unemployment between young and old workers. Since lower type matches are more likely to turn out to be bad matches, they are more likely to endogenously separate. Thus, lower type workers—who tend to be young—experience higher unemployment rates, as they are more likely to be in transition between occupations.

In Section 5, we calibrate our model and study its quantitative implications. Our model does a very good job of matching the observed age profile of unemployment. This is due to the fact that our model nearly replicates the declining age profile of separation rates observed in the data. In addition, our model does a good job of matching the age profile of occupational mobility. Finally, our model rationalizes a significant portion of the fall in aggregate unemployment observed since the mid-1970s, through aging of the labor force, as emphasized in the empirical results of Shimer (1998). Hence, we find that the "learning about occupational fit" mechanism in our model is important for understanding life-cycle and aggregate labor market dynamics.

Ours is not the first frictional labor market model to address age differences in unemployment. Jovanovic (1979) represents a seminal contribution to the literature on endogenous job separations. In that model, a worker's separation rate falls with tenure in a match. This arises because productivity in a worker-firm match is stochastic and observed imprecisely; longer tenured matches are those in which the worker and firm have learned that idiosyncratic match quality is high. Menzio et al. (2012) illustrate how this mechanism generates a declining age profile of separation in a life-cycle model: since young workers have drawn fewer matches, they are more likely to be in ones with lower idiosyncratic productivity, and hence, more likely to separate.

The key distinction in our framework is what is being learned in a match. In the

Jovanovic (1979) and Menzio et al. (2012) framework, agents learn about the quality of the match between worker and firm, or *firm-specific* human capital. Since this is independent across matches, separation from any one match imparts no information about a worker's productivity and separation rate in future matches.

In our model, agents learn about match quality between worker and occupation. This is a form of *general* human capital—knowledge of one's true calling—that worker's carry forward into future employment and unemployment spells. Separation from a match is due to the accumulation of this human capital, and hence generates a lower probability of separation and higher expected productivity in future matches. In contrast to the Jovanovic (1979) and Menzio et al. (2012) mechanism where young workers are more likely to be in matches with lower firm-specific productivity, our model's declining age profile of separation is because young workers are more likely to be in matches with lower *occupation-specific* productivity. Moreover, our model is consistent with the results of Kambourov and Manovskii (2009) who find that the importance of firm-specific effects in rationalizing life-cycle wage dynamics vanishes when occupation-specific and total experience effects are taken into account in Mincerian wage regressions.

Our framework is also related to the one-sided search problem studied by Neal (1999), in which employment relationships have two components: 'career' quality and match quality. One interpretation of our model is that individuals are searching for a career by sampling occupations. As such, our model is consistent with Neal (1999)'s empirical findings that career change falls with labor market experience. However, since that model abstracts from unemployment as a labor market state, it does not address the age profile of unemployment rates, our subject of principal interest.

More recently, Papageorgiou (2012) also studies the importance of learning about unobserved occupational ability in a search framework. However, because the phenomena of interest are different, a number of model features differ. Papageorgiou (2012) finds that a calibrated version of his model does well in accounting for life-cycle wage dynamics, residual wage inequality, and inter-occupational flows. By contrast, our paper studies the role of learning about occupational fit in accounting for the life-cycle profile of unemployment and separation rates, and aggregate unemployment dynamics. Both papers do well in accounting for the life-cycle pattern of occupational mobility.[2]

---

[2]Finally, see also Esteban-Pretel and Fujimoto (2011). In their model, older workers are assumed to be better able to observe idiosyncratic match productivity, and hence, reject poor matches before they are formed. In this way, their model generates a declining age profile of separation by construction. See also Gorry (2012) for a one-sided search model in which labor market experience provides a worker with better information about future match quality.

Table 1: Average Unemployment Rates by Age Group

	20–24	25–34	35–44	45–54	55–64
average (%)	10.45	6.37	4.81	4.22	4.01
normalized	2.48	1.51	1.14	1	0.95

Notes: data from the CPS, 1976:6–2012:11. The second row indicates average unemployment rates by age, relative to that of 45–54 year olds.

2 Data

In this section, we detail the empirical observations that motivate our work. We first document the large age differences in unemployment. We then present evidence on job finding and separation rates that informs our theoretical approach in Section 3.

We begin by analyzing the unemployment rate disaggregated by age, obtained from the Current Population Survey (CPS) for the period June 1976 – November 2012.[3] The first row of Table 1 displays the average unemployment rate during this period for different age groups. As is obvious, unemployment falls monotonically with age. The average unemployment rate for 20–24 year olds is 10.45% and falls to 4.01% for 55–64 year olds.

Moreover, the age differences are large. The second row presents the average unemployment rate for each age group, relative to that of 45–54 year olds. During this period, average unemployment for 20–24 year olds is 2.48 times that of 45–54 year olds. The average unemployment rate for 25–34 year olds is 51% greater than that of the prime-aged.

2.1 Job Finding and Separation Rates

What accounts for these large age differences? To address this question, we examine the age differences in job finding and separation rates. To understand the role of these transition rates, consider a simple labor market model where: (a) all unemployed workers transit to employment at the constant rate f, and remain unemployed otherwise, and (b) all employed workers transit to unemployment at the constant rate s, and remain employed otherwise. In this setting, the steady state unemployment rate, u, is given by:

$$u = \frac{s}{s + f}.$$

---

[3]We focus on this time period since this is the period for which micro-level data (used in the following subsection) is available. None of our substantive results regarding age differences are altered when data from 1948 to 1976 are included.

Table 2: Average Monthly Transition Rates by Age Group

	20–24	25–34	35–44	45–54	55–64
A. Shimer approach					
job finding (%)	39.54	34.24	31.17	28.42	27.34
separation (%)	5.58	2.70	1.79	1.39	1.26
B. direct approach					
job finding (%)	28.46	26.58	25.50	23.77	21.35
separation (%)	2.62	1.51	1.13	0.96	0.84

Notes: data from the CPS, 1976:6–2012:11. Panel A presents transition rates computed from data on stocks of employed and unemployed workers. Panel B presents transition rates computed from data on worker flows. See the main text for details.

Holding f constant across age groups, a decreasing age profile for unemployment would require a decreasing profile for s. Similarly, holding s constant across age groups, a decreasing age profile for unemployment would require an increasing profile for f.

We first calculate job finding and separation rates following the approach of Shimer (2005). This approach is well suited for our analysis, since it assumes that all workers transit solely between states of employment and unemployment, as in the statistical model discussed above, and the economic model presented in Section 3.[4] The approach uses monthly data on employment, unemployment, and short-term unemployment tabulated from the CPS. Disaggregated by age, this data is available beginning in June 1976. Panel A of Table 2 displays average job finding and separation rates by age during this period.

The first row of Table 2 indicates that job finding rates decrease monotonically with age. In the average month, 39.54% of unemployed 20–24 year olds transit to employment, a rate that is about 40% greater than that of 45–54 year olds. This represents a relatively small age difference when compared to separation rates, as we discuss below. Moreover, absent differences in separation rates, age differences in job finding rates would counterfactually imply an *increasing* age profile for unemployment.

Hence, age differences in unemployment are accounted for solely by differences in separation rates. Indeed, the second row of Table 2 indicates that the separation rate falls monotonically with age. Moreover, the age differences are large. For example, the separation rate for 20–24 year olds is 4.0 times that of 45–54 year olds. As such, we view the declining profile of separation rates as integral to any theory of age differences in unem-

---

[4]We refer the reader to Shimer (2005) for further methodological details.

ployment.[5]

The Shimer (2005) approach to computing transition rates assumes that workers do not transit in and out of the labor force, though such flows are observed in longitudinal data at the individual level. This may give rise to misleading conclusions regarding the roles of job finding and separation in rationalizing age differences in unemployment. This is particularly true if transition rates in and out of the labor force differ sufficiently by age.

As such, Panel B of Table 2 displays average transition rates measured directly from data on individual-level labor market flows. To do this, we use the monthly CPS files. Households in the CPS are surveyed for four consecutive months, then leave the sample during the next eight months, and are surveyed again for a final four months. Each household member is uniquely identified. As such, we obtain a longitudinal record for each person which allows us to track monthly transitions across states of employment, unemployment, and labor force non-participation.[6]

The first row of Panel B presents the job finding rate (from unemployment to employment). This falls monotonically with age, with the job finding rate of 20–24 year olds about 20% greater than that of 45–54 year olds. Again, absent differences in separation rates, these differences would counterfactually imply an increasing age profile for unemployment. The second row presents the separation rate (from employment to unemployment). As before, these fall monotonically with age, and the age differences are large. The separation rate for 20–24 year olds is 2.7 times that of 45–54 year olds.[7] Again, this evidence on transition rates indicates the critical role of separation rates in understanding age differences in unemployment.

2.2 Occupational Mobility

The mechanism underlying our theory is that individuals learn about their true calling in life by experiencing various occupations as they age. Following Kambourov and Manovskii (2008, 2009), we use data from the Panel Study of Income Dynamics (PSID) to document

---

[5]See also Shimer (1998) who emphasizes the importance of higher separation rates in understanding higher unemployment rates for the young.

[6]For a complete description of the construction of the longitudinal data, see Nekarda (2009).

[7]Finally, we note that the levels of transition rates naturally differ across the two measurement approaches, since the previous approach assumes away flows in and out of the labor force. See also Menzio et al. (2012) who study data from the SIPP for male high-school graduates, December 1995 to February 2000. They find that the job finding rate is essentially constant between the ages of 20 and 55, and only fall in a noticeable way between the ages of 55 and 64. Separation rates exhibit a marked fall throughout the age profile, most notably between the ages of 20 and 35, as in the CPS data.

Figure 1: Occupational Mobility by Age

Notes: Data from Kambourov and Manovskii (2008). Original data from the PSID, Retrospective Files, 1968–1980. The sample consists of male heads, non-self-employed, either employed or temporarily laid off.

the age-profile of occupational mobility.

As is well-known, measures of occupational mobility are subject to sizable measurement error. In order to mitigate this error, the PSID released Retrospective Occupation-Industry Supplemental Data Files (Retrospective Files hereafter) in 1999. The goal of this release was to make available retrospectively assigned 3-digit 1970 census codes to the reported occupations and industries of household heads and wives for the period 1968–80. Kambourov and Manovskii (2008, 2009) argue that the methodology used to construct the Retrospective Files minimizes the error in identifying true occupation switches.[8]

Figure 1 shows occupational mobility rates by age group for occupations at the 1-, 2-, and 3-digit level.[9] While the data and methodology underlying Figure 1 comes from Kambourov and Manovskii (2008), we adapt their exercise for the purpose of this study. First, we only use data from the period covered by the Retrospective files. Second, we use a broader sample, consisting of male heads of households who aren't self-employed and are

---

[8]For example, approximately half of occupation switches in the original data are identified as legitimate occupation switches in the Retrospective Files. A key aspect of the Retrospective files is that a single individual was responsible for coding the occupation of a particular individual over the entire sample period, thereby minimizing changes in interpretation of the occupation reported by a given individual over time.

[9]See Appendix B of Kambourov and Manovskii (2008) for details of the occupation classification system.

either employed or temporarily laid off.[10] The main difference with their sample is our inclusion of government workers and farmers, who tend to switch occupation less frequently than the rest of the sample.

Figure 1 reveals that occupational mobility declines sharply as people age.[11] For the young, occupational mobility rates are very high: more than 40% of individuals in the 20–24 year old age group switches occupation in a given year at the 3-digit level. Even at the 2-digit level, approximately one in three 20–24 year olds switch occupation annually. For prime-aged workers, about one in 10 or 12 individuals changes 2- or 3-digit occupation on a yearly basis.

These mobility rates can be used to calculate the number of occupations an individual experiences over her working life. For example, suppose that an individual enters the labor force and experiences her first occupation at age 20. Assuming a constant hazard rate between the ages of 20 and 24, the probability that she switches to a different 3-digit occupation during each of the first five years of her career is 41%; this would imply that she switches occupations 2.05 times during the first five years. Similarly, over the next 10 years the average switching probability is 24%, implying 2.40 occupation switches. Repeating this calculation for all the age groups yields that the average worker switches about 8.6 times, and hence works in 9.6 3-digit occupations over her career. A similar procedure indicates that the average person works in 7.3 and 6.6 occupations, at the 2- and 1-digit level, respectively.

3   Economic Environment

To analyze age differences in unemployment, we study a search-and-matching model of the labor market. The matching process between unemployed workers and vacancy posting firms is subject to a search friction. The ratio of vacancies to unemployed determines the economy's match probabilities, in a way that we make precise in the next subsection.

Workers differ in their knowledge of their best occupational fit. Specifically, there are M occupations in the economy that are identical, except in name. Each worker is best-suited

---

[10]The latter restriction implies that we do not include individuals who leave the labor force in the calculation of the occupational mobility rate. The mobility rate is the ratio of the number of individuals who switch occupations, divided by the total number of workers (i.e. the sum of 'switchers' and 'non-switchers'). Assume an individual who leaves the labor force (retires or goes to school). If we count this individual as a 'non-switcher' we generate a downward bias in the occupational mobility measure (because the denominator is larger). Excluding those who leave the labor force from the sample avoids this bias.

[11]See also Moscarini and Vella (2008) for evidence of declining occupational mobility by in CPS data.

for one occupation; that is, only one occupation, $m^* \in \{1, 2, ..., M\}$, is a best match, or the worker's "true calling." When a worker is employed in her true calling occupation, m^*, she produce f_G units of output. For simplicity, in all other $M - 1$ occupations, the worker-firm pair produces f_B units of output, with $f_B < f_G$.

In each period, a mass of newly born workers enter the economy not knowing their true calling. A new worker has m^* randomly assigned into one of the M occupations with probability $1/M$. This assignment is distributed independently across all new workers. In order to learn whether any given occupation is their best fit, the worker must search, be matched, and work in that occupation. Learning about occupational fit in a match does not happen instantaneously. In each period of employment, the worker (and firm) learns whether it is the worker's true calling with probability $\lambda \in (0, 1)$.

Given stochastic match creation, destruction, and learning, workers are heterogeneous with respect to their labor market history. Worker heterogeneity can be summarized by a worker's *type*, $i \in \{1, \ldots, M\}$; here, i indicates the number of ill-suited occupations the worker has identified, plus one. A worker of type $i < M$ has previously sampled $i - 1$ ill-suited occupations, and therefore has $M - i + 1$ left to try; a worker of type M knows her true calling, m^* (either having sampled $M - 1$ ill-suited occupations, or having sampled fewer than that but gotten "lucky").

Workers understand that their true calling is uniformly distributed across all occupations, and use this to form their beliefs. Hence, an unemployed worker of type $i = 1$ has a flat prior over occupations, and randomly selects an occupation to search for employment. If that occupation turns out to be the wrong fit, the worker becomes type $i = 2$ and updates her prior according to Bayes' rule. She now understands that her true calling is uniformly distributed over the remaining $M - 1$ occupations, and randomly picks one of those occupations to search in the next time she is unemployed. Defining p_i as the probability, conditional on being matched in the i^{th} occupation, that the worker has found her true calling, we have that:

$$p_i = \frac{1}{M - i + 1}, \quad i = 1, \ldots, M.$$

To close the description of the environment, we must take a stand on "death/exit" from the economy. For simplicity, we assume that each worker faces a constant probability of remaining in the labor market from one period to the next, $\sigma \in (0, 1)$. We interpret β in what follows as the "true" discount factor times this survival probability.

3.1 Market Tightness

We assume that there are M labor markets in the economy, one for each occupation. All unemployed workers seeking employment in a particular occupation search in that occupation's market.

While a worker's type is known to the worker, it is unobservable by vacancy posting firms. Workers are unable to signal their type to potential employers. A firm wishing to hire a worker in a particular occupation posts a vacancy in that occupation's market, understanding that it may be matched with a worker of any type $i \in \{1, \ldots, M\}$. The probability the firm assigns to being matched with a type i worker is taken parametrically, and determined by the equilibrium distribution of unemployed worker types. In this sense, matching is *random* within each occupational market.[12] Upon matching, the worker's type is observable by both the worker and firm.

Given the symmetry assumptions of the model, all occupational markets are identical. We define market tightness in the representative market, θ, as the ratio of the number of vacancies maintained by firms to the number of workers looking for jobs in that occupation. While θ is an equilibrium object, it is taken parametrically by firms and workers.

We denote the probability that a worker will meet a vacancy—the job finding rate—by $\mu(\theta)$, where $\mu : \mathcal{R}_+ \to [0, 1]$ is a strictly increasing function with $\mu(0) = 0$. Similarly, we let $q(\theta)$ denote the probability that a firm with a vacancy meets a worker, where $q : \mathcal{R}_+ \to [0, 1]$ is a strictly decreasing function with $q(\theta) \to 1$ as $\theta \to 0$, and $q(\theta) = \mu(\theta)/\theta$.

Finally, we note that this structure of matching markets implies that all unemployed workers face the same job finding rate. This feature allows us to isolate the role of occupational learning in generating heterogeneity in separation rates. We find this to be an attractive feature of our approach, since age differences in separation fully account for age differences in unemployment, as discussed in Subsection 2.1.

3.2 Contractual Arrangement and Timing

We specify the worker's compensation in a match as being determined via Nash bargaining with fixed bargaining weights, as in Pissarides (1985). When an unemployed worker and a firm match, they begin producing output in the following period. In all periods that a worker and firm are matched, the compensation is bargained with complete knowledge of

---

[12]See also Pries (2008) who studies a model with random search and heterogeneous workers.

the worker's type.

3.3 Worker's Problem

Workers can either be unemployed or employed. Employed workers can either be: in a "good" match, working in their true calling occupation; in a "bad" match, in an occupation that is not the worker's true calling; or in a match of yet unknown occupational fit.

We define U_i as the value of being unemployed for a worker of type i:

$$U_i = z + \beta \left[\mu(\theta)W'_{L,i} + (1 - \mu(\theta))U'_i \right], \quad i = 1, \dots, M - 1. \tag{1}$$

Here, z is the flow value of unemployment, $W_{L,i}$ is the worker's value of being employed in her i^{th} occupation and learning whether it is her true calling, and primes $(')$ denote variables one period in the future. An unemployed worker transits to employment in the following period with probability $\mu(\theta)$.

Once employed, the type i worker's value while in the "learning phase" is given by:

$$W_{L,i} = \omega_{L,i} + \beta(1 - \lambda)\left[(1 - \delta)W'_{L,i} + \delta U'_i\right] +$$
$$\beta\lambda\left[p_i\left((1 - \delta)W'_M + \delta U'_M\right) + (1 - p_i)\left((1 - \delta)W'_{B,i} + \delta U'_{i+1}\right)\right]. \tag{2}$$

While employed in a match of unknown occupational fit, the worker earns per period compensation $\omega_{L,i}$, and learns whether this is her true calling with probability λ.[13] If the worker does not learn, then she remains as type i in the following period. If the match separates, with exogenous probability $\delta \in (0, 1)$, she returns to being unemployed with value U_i; otherwise, she continues with value $W_{L,i}$.

If the worker does learn about her fit in the current occupation, her continuation value is given by the square bracketed term in the second line of equation (2). With probability p_i, the worker has found her true calling, and she becomes type $i = M$. Otherwise, the current occupation is not her true calling. If the match does not exogenously separate, she continues as a worker in a bad occupational match with value $W_{B,i}$; if the match separates, she becomes unemployed of type $i + 1$, having eliminated an additional occupation as her true calling.

---

[13]Note that this constant hazard/learning rate can be viewed as the reduced form of a more elaborate signal extraction problem. Specifically, Pries (2004) demonstrates how a model where match output is observed with uniform measurement error gives rise to "all-or-nothing" learning, as in our model: with some probability (in our notation, $1 - \lambda$), observed output reveals nothing about the occupational fit, while with complementary probability, the signal is perfectly revealing.

The worker's value of being employed in her true calling is given by:

$$W_M = \omega_M + \beta \left[(1-\delta)W_M' + \delta U_M' \right].\tag{3}$$

Obviously, if the match exogenously separates, the worker retains her type M as she knows her true calling; in this case, the value of being unemployed is:

$$U_M = z + \beta \left[\mu(\theta)W_M' + (1-\mu(\theta))U_M' \right].\tag{4}$$

Finally, the value of being employed in a bad match is given by :

$$W_{B,i} = \max \left\{ \omega_{B,i} + \beta \left[(1-\delta)W_{B,i}' + \delta U_{i+1}' \right],\ U_{i+1} \right\}.\tag{5}$$

Note that the worker chooses either to remain in the match or, if it is higher, to be unemployed. In the event of separation – whether exogenous or endogenous – the worker becomes unemployed of type $i+1$ as she knows that her last occupation was not her true calling.

This formulation assumes that workers will never sample an occupation that they already know is not their true calling. We show later in Proposition 1 that this is indeed a feature of the equilibrium: workers who have sampled i occupations that are not their best fit never sample one of these again upon separation.

3.4   Firm's Problem

Firms—or more correctly, vacancies for a specific occupation—can be either filled or unfilled. When filled, a vacancy may be matched with a worker for whom the occupation is her true calling. In this case, the firm's value is given by:

$$J_M = f_G - \omega_M + \beta \left[(1-\delta)J_M' + \delta V' \right],\tag{6}$$

where f_G is the output in a true calling match, and ω_M is the compensation paid to the worker. In the case of separation, the vacancy becomes unfilled with value V.

When matched with a type i worker who is in a bad occupational fit, match output is given by f_B, and the firm's value is:

$$J_{B,i} = \max \left\{ f_B - \omega_{B,i} + \beta \left[(1-\delta)J_{B,i}' + \delta V' \right],\ V \right\}.\tag{7}$$

Firms in bad matches may choose to separate if it is in their best interest to do so.

Finally, a firm may be matched with a type i worker who is learning whether the current occupation is her true calling. In this case, the firm's value is given by:

$$J_{L,i} = p_i f_G + (1 - p_i) f_B - \omega_{L,i} +$$
$$\beta \Big[\lambda(1 - \delta) \big[p_i J'_M + (1 - p_i) J'_{B,i} \big] + (1 - \lambda)(1 - \delta) J'_{L,i} + \delta V' \Big]. \qquad (8)$$

This is composed of the contemporaneous profit (expected match output minus the worker's compensation) plus the continuation value.[14] With probability λ, the worker and firm learn about the occupational fit. With probability p_i, this is the worker's true calling, and (conditional on surviving) the match continues with value J_M; otherwise, this is an ill-suited occupational match and the continuation value is $J_{B,i}$.

We assume that there is a large number of firms who can potentially post vacancies in any of the M occupational markets. Doing so requires the payment of a vacancy posting cost, $\kappa > 0$. The value of posting a vacancy in the representative market is defined as:

$$V = -\kappa + \beta \left[q(\theta) \left[\sum_{i=1}^{M-1} \pi_i J'_{L,i} + \pi_M J'_M \right] + (1 - q(\theta)) V' \right]. \qquad (9)$$

An unfilled vacancy is matched with a worker with probability $q(\theta)$. Upon filling the vacancy, the firm learns the type of worker that it has matched with. As such, the continuation value, conditional upon matching, is probability weighted across the M worker types. Here, π_i denotes the firm's probability of matching with a worker of type i, with $\sum_{i=1}^{M} \pi_i = 1$.

4 Characterizing Equilibrium

4.1 Definition of Equilibrium

A steady state equilibrium with Nash bargaining is a collection of value functions $\Big(\{J_{B,i}, J_{L,i}\}_{i=1}^{M-1}$, $J_M, V; \{U_i, W_{B,i}, W_{L,i}\}_{i=1}^{M-1}, U_M, W_M \Big)$, compensations $\Big(\{\omega_{L,i}, \omega_{B,i}\}_{i=1}^{M-1}, \omega_M \Big)$, a probability distribution over unemployed workers $\{\pi_i\}_{i=1}^{M}$, and a tightness ratio θ such that:

1. workers are optimizing. That is, workers that are matched prefer to remain matched rather than be unemployed, $W_{L,i} > U_i, W_{B,i} \geq U_{i+1}, i = 1, \ldots, M-1$, and $W_M > U_M$;

---

[14]We have specified output in a learning match to equal its expected value. This ensures consistency with the expectation used in the firm's determination of vacancy creation, equation (9). As discussed in the previous footnote, our model is isomorphic to one in which output in the learning phase of a match is observed with measurement error.

2. firms are optimizing. That is, firms that are matched prefer to remain matched as opposed to maintaining a vacancy, $J_{L,i}, J_M > V, J_{B,i} \geq V, i = 1, \ldots, M$;

3. compensations solve the generalized Nash bargaining problems displayed below;

4. the probability distribution over unemployed workers is consistent with individual behavior and the implied laws-of-motion across labor market states and worker types;[15] and

5. the free entry condition is satisfied. That is, $V = 0$.

4.2   Compensation

As is standard in the literature, we assume that compensation is determined via generalized Nash bargaining.[16]   Let τ denote the bargaining weight of workers. The compensation earned by workers in good matches solves the following problem:

$$\omega_M = \arg\max \ (W_M - U_M)^\tau (J_M)^{1-\tau}.$$

The solution takes the form:

$$\omega_M = \tau f_G + (1 - \tau)\left[z + \mu(\theta)\beta(W_M - U_M)\right].$$

The interpretation of this is standard. In a good match, the worker is paid a fraction, τ, of the match output, plus a fraction, $1 - \tau$, of the value of unemployment. The value of unemployment consists of both the flow value, z, and the *option value* of being unemployed.

The compensation for workers who are in their i^{th} bad match solves:

$$\omega_{B,i} = \arg\max \ (W_{B,i} - U_{i+1})^\tau (J_{B,i})^{1-\tau}.$$

For these workers, the outside option of remaining in the match is being unemployed of type $i + 1$. As such:

$$\omega_{B,i} = \tau f_B + (1 - \tau)\left[z + \mu(\theta)\beta(W_{L,i+1} - U_{i+1})\right].$$

Finally, the compensation for type i workers in a learning-phase match solves:

$$\omega_{L,i} = \arg\max \ (W_{L,i} - U_i)^\tau (J_{L,i})^{1-\tau}.$$

---

[15]See the Appendix for details.

[16]See the Appendix for details on all derivations presented in this subsection. Since we focus on steady states, we omit the $'$ on next period variables both here and in the appendix.

For these workers, the threat point is U_i: if agreement is not reached in bargaining, they return to unemployment having not learned anything about their type. The solution for $\omega_{L,i}$ is given by:

$$\omega_{L,i} = \tau \left[p_i f_G + (1 - p_i) f_B \right] + (1 - \tau) \left[z + \mu(\theta)\beta \left(W_{L,i} - U_i \right) \right]$$
$$- (1 - \tau)\beta\lambda \left[p_i U_M + (1 - p_i)U_{i+1} - U_i \right].$$

Relative to ω_M and $\omega_{B,i}$, $\omega_{L,i}$ contains a new term, $\beta\lambda \left[p_i U'_M + (1 - p_i)U'_{i+1} - U'_i \right]$. We refer to this as the *information value*. By working in the match, the worker learns about her fit in the current occupation with probability λ. Conditional on learning, she learns that it is her true calling with probability p_i; with complementary probability she learns that her type (when next unemployed) is $i + 1$. Hence, the information value is the discounted, expected gain to working and augmenting her threat point in future bargaining.

Moreover, the information value is *positive*. This follows as an immediate corollary of Proposition 1, which we establish below. Hence, the information value term represents a "pay cut" to the worker relative to the standard Nash bargaining solution. Because the worker has this additional incentive to work—to advance toward her true calling—she is willing to accept this pay cut in her negotiated compensation.

4.3 Characterizing Steady State

In this section we establish that in any steady state equilibrium, the value of unemployment increases with workers' type, that is, $U_i < U_{i+1}$, $i = 1, \ldots, M - 1$. This result is useful not only as it leads to a sharp characterization of steady state equilibria, but also because unemployed workers in such equilibria do not have an incentive to misrepresent their type.

Proposition 1 *Assume that all bad matches endogenously separate.*[17] *In any steady state equilibrium, $U_i < U_{i+1}$ for $i = 1, 2, \ldots, M - 1$.*

The proof is contained in the Appendix. A number of results follow immediately from this Proposition, collected in the following Corollary:

Corollary 1 *Assume that all bad matches endogenously separate. Then*

1. *$J_{L,i+1} > J_{L,i}$;*

---

[17]This assumption holds in all equilibria we compute in Section 5.

2. $W_{L,i+1} - U_{i+1} > W_{L,i} - U_i$;

3. $W_{L,i+1} - W_{L,i} > U_{i+1} - U_i > 0$, so $W_{L,i+1} > W_{L,i}$;

4. $p_i U_M + (1 - p_i)U_{i+1} - U_i > 0$.

5 Numerical Analysis

In Subsection 5.1, we discuss the calibration of our model to the U.S. data. Subsection 5.2 presents results regarding our model's ability to match the life-cycle labor market facts, and Subsection 5.3 discusses robustness. Subsection 5.4 considers a counterfactual exercise, to determine the role of age composition change in the labor market in the secular fall in unemployment observed since the mid-1970s.

5.1 Calibration

Many of our model features are standard to the DMP literature, so our stategy is to maintain comparability wherever possible. The model is calibrated to a monthly frequency. As such, the discount factor is set to $\beta = 0.996$ to accord with an annual risk free rate of 5%.

We assume that the matching function in each occupational market is Cobb-Douglas:

$$\mu(\theta) = \theta q(\theta) = \theta^\alpha.$$

Summarizing a large literature that directly estimates the matching function using aggregate data, Petrongolo and Pissarides (2001) establish a plausible range for α of $0.3 - 0.5$. Refining the inference approaches of Shimer (2005) and Mortensen and Nagypál (2007), Brügemann (2008) obtains a range of $0.37-0.46$. In our calibration, we specify $\alpha = 0.4$ to be near the mid point of these ranges (see also Pissarides (2009)). For comparability with previous work, we specify the parameter in the Nash bargaining problem as $\tau = 1 - \alpha$.

The aggregate unemployment rate in the model is determined by the aggregate job finding rate and the endogenous distribution of separation rates. We set the exogenous separation rate parameter to $\delta = 0.012$, as is standard in the literature. From the free entry condition (equation (9) with $V = 0$), the vacancy cost, κ, pins down $q(\theta)$, and hence, the aggregate job finding rate, $\mu(\theta)$. Given this, we set κ so that the steady state unemployment rate among the model's 20 to 64 year old workers is 5.75%.[18] This target corresponds to

---

[18]As discussed in Shimer (2005), the exact value of κ is irrelevant. By introducing a multiplicative constant,

the average aggregate unemployment rate observed among 20–64 year olds in the the CPS data, 1976:6 to 2012:11.

The model's aggregate unemployment rate depends on the age distribution of endogenous separation rates. This, in turn, depends on the model's age distribution of workers. As such, we set the labor market survival probability, σ, to best match the age distribution of the labor force observed in the data. Specifically, we minimize the sum of squared residuals between the model implied labor force shares in the 20–24, 25–34, 35–44, 45–54, and 55–64 age bins, with the average values of those shares in the CPS data, 1976:6 to 2012:11. This results in a value of $\sigma = 0.9984$.[19] As an additional check, we find that this survival probability implies an average worker age of 39.02 years in the model, which matches exactly (to the second decimal place) the average age found in the data.

Our calibration of z, the flow value of unemployment, follows the strategy of Hall and Milgrom (2008), Mortensen and Nagypál (2007), and Pissarides (2009). Specifically, we interpret z as being composed of two components: a value of leisure or home production, and a value associated with unemployment benefits. As in their work, the return to leisure/home production is equated to 43% of the average return to market work. Given this target, the model's Nash bargained compensation, and the steady state distribution of worker types, we set $z = 0.56$. This implies an unemployment benefit replacement rate of 35% for the lowest type ($i = 1$) workers, and 20% for the highest type ($i = M$) workers; this accords with the range of replacement rates reported by Hall and Milgrom (2008).

We calibrate M based on the data on occupational mobility of Kambourov and Manovskii (2008). Taking the age profile of occupational switching rates as hazard rates implies that the average individual works in 9.6 three-digit occupations over her career. Matching this statistic in our model would require $M = 19$.[20] However, this interpretation of the occupational mobility data might overstate our model's ability to match age difference in separation; this occurs if some of observed occupational switches happen without an intervening spell of unemployment, such as internal promotions for example. Given this, we choose to be conservative and set $M = 10$ in our baseline calibration, so that the average worker experiences 5 occupations over her career.

Relative to the standard DMP model, our model emphasizes the role of learning about one's true calling that occurs through the sampling of occupations. Knowledge about oc-

---

ξ, to the matching function, κ can be scaled by a factor of x and ξ by a factor of x^{α}, leaving the job finding rate, and hence the aggregate unemployment rate, unchanged.

[19] The exact age distribution generated by the model is discussed in detail in subsection 5.4.

[20] Defining occupations at the 1-digit level would imply setting $M = 13$.

cupational fit represents a form of human capital that is acquired through labor market experience. Given this, it is natural to calibrate the remaining novel parameters of our model—f_G, f_B, and λ—to match the empirical life-cycle earnings profile, as estimated by Murphy and Welch (1990) and others. That is, in our baseline calibration, we assume that the returns to occupational learning match exactly the returns to labor market experience.

We choose f_G/f_B and λ to match two key properties of the return to experience. First, the maximal lifetime wage gain for a typical worker represents an approximate doubling of earnings. Second, this doubling occurs after the typical worker has accumulated 25 years of experience. Normalizing $f_G = 1$ and matching these statistics in our model requires setting $f_B = 0.57$ and $\lambda = 1/18$.[21] This implies that a worker's match productivity is 75% higher in her true calling than in any other occupation, and that it takes, on average, 18 months in a match in order to learn the occupational fit.

In our results, we assume that workers are "born" into the workforce at the age of 19.5 years old. We allow 6 months to elapse before tracking labor market outcomes, so that the youngest worker in the age-specific statistics we report is 20 years old. Results for the baseline calibration are contained in the next subsection, and robustness of our results to variation in M, f_G/f_B, and λ is investigated following that.

5.2 Results

The first row of Panel A, Table 3 reproduces the unemployment rate by age displayed in Table 1. The second row displays the age profile of unemployment generated by the baseline calibration of our model. Occupational learning over the life cycle implies that unemployment falls as workers age and find their true calling. Through this mechanism, our model does a very good job at matching the observed age profile of unemployment.

The assumption of random search within an occupational market implies that the job finding rate of all workers is identical. As such, all age differences in unemployment are driven by differences in separation rates. In our model, separation rates differ across workers because they face different endogenous separation probabilities. In our baseline calibration, all workers choose to separate if they learn that their current match is of poor occupational fit. That is, $W_{B,i} = U_{i+1}$ for all $i = 1, \ldots, M-1$ in equation (5). Only type $i = M$ matches separate at the exogenous rate δ; all other matches – namely, learning-phase matches –

---

[21]Technically, given the non-deterministic nature of our model, the expected life-cycle profile of earnings only asymptotes to the theoretical maximal value as time approaches ∞. Hence, our calibration procedure requires that earnings come within 0.25% of fully doubling at the 25 year horizon.

Table 3: Labor Market Statistics by Age Group: Data and Model

	20–24	25–34	35–44	45–54	55–64
A. unemployment rate					
U.S. data (%)	10.45	6.37	4.81	4.22	4.01
model (%)	12.21	7.33	3.58	3.06	3.04
B. separation rate					
U.S. data (%)	5.58	2.70	1.79	1.39	1.26
model (%)	5.33	3.03	1.42	1.21	1.20

Notes: U.S. data from the CPS. Model statistics computed from the theoretical steady state distribution of the baseline calibration.

separate at rate $\delta + (1 - \delta)\lambda(1 - p_i) > \delta$. Hence, our model obtains a declining age profile of separation because older workers are more likely to have found their true calling.

This can be seen in Panel B of Table 3 where we display separation rates by age.[22] In the baseline calibration of our model, young workers aged 20–24 years old tend not to have found their true calling; as such, they face a separation rate of 5.33%. On the other hand, old workers aged 55–64 have essentially all found their true calling; their separation rate of 1.20% is essentially identical to the calibrated exogenous separation rate, δ. Moreover, our model does a good job of accounting for the age profile of separation rates. For example, the separation rate of 20–24 year olds is 4.4 times that of either the 45–54 or 55–64 year olds. In the U.S. data, this ratio is 4.0 and 4.4, respectively. Hence, the reason our model overstates age differences in unemployment rates is because of our simplifying assumption of identical job finding rates.

Finally, we note that our model does a good job of replicating the age profile of occupational mobility. Figure 2 displays the occupational switching rate at the annual frequency for different age groups in our model; this is analogous to Figure 1 derived from the PSID data studied by Kambourov and Manovskii (2008). For young workers aged 20–24 years old, the model generates an occupational switching rate of 32% which is very close to the rate found at the 2-digit occupational level in the data. However, occupational mobility falls faster in the model relative to the data: by the time workers reach prime age, essentially everyone has found their true calling, and the occupational switching rate is near zero.

---

[22]For the U.S. data, we report separation rates as calculated in Shimer (2005), as this method accords with our model's assumption that workers do not transit in and out of labor force participation during their working-age life.

Figure 2: Occupational Mobility by Age: Model

Notes: Probability of a worker switching occupations, annual frequency.

5.3 Robustness

Here we explore the robustness of our results to variations from the baseline calibration. We first consider the effect of changing M, the number of potential occupations. To see that M affects the ability to generate age differences in unemployment, consider the case of $M = 1$: the model collapses to the standard representative worker DMP model (augmented with a constant survival probability).

Panel C of Table 4 addresses this exercise. Here, we reduce the number of occupations to $M = 7$; this accords with the implied number of 1-digit occupations worked by the average individual using the PSID data. Doing so reduces the amount of time a worker spends over her life cycle searching for her true calling. In these results, the rest of the parameter values are reset to maintain the remaining calibration targets discussed in Subsection 5.1.

Despite the 30% reduction in the number of potential occupations, the model still delivers sizeable age differences in unemployment and separation. Young workers in the model experience an unemployment rate that is 3.6 times that of 45–54 year olds. In terms of separation rates, the model delivers a ratio of 20–24 year olds to 45–54 year olds of 4.0,

Table 4: Labor Market Statistics by Age Group: Robustness

	20–24	25–34	35–44	45–54	55–64
A. U.S. data					
unemployment rate	10.45	6.37	4.81	4.22	4.01
separation rate	5.58	2.70	1.79	1.39	1.26
B. benchmark model					
unemployment rate	12.21	7.33	3.58	3.06	3.04
separation rate	5.33	3.03	1.42	1.21	1.20
C. $M = 7$					
unemployment rate	12.83	6.21	3.64	3.53	3.53
separation rate	4.83	2.17	1.24	1.20	1.20
D. $\lambda = 1/12$					
unemployment rate	14.85	6.33	3.07	3.01	3.01
separation rate	6.75	2.62	1.22	1.20	1.20
E. $f_B = 0.95$					
unemployment rate	12.21	7.33	3.58	3.06	3.04
separation rate	5.33	3.03	1.42	1.21	1.20

Notes: U.S. data from the CPS. Model statistics computed from the theoretical steady state distribution for various parameter specifications.

matching the ratio of 4.0 observed in the data.

In our second robustness exercise, we increase the learning rate to $\lambda = 1/12$. This represents a 50% increase in the learning rate relative to the baseline calibration, and implies that workers learn their occupational fit in a match in 12 months, on average. Increasing λ has the effect of reducing the amount of time a worker spends over her life cycle searching for her true calling. The results are displayed in Panel D of Table 4.

Not surprisingly, increasing λ causes workers to spend more of the early part of their career unemployed, as they transition between occupations. Relative to the baseline calibration, the unemployment and separation rates of 20–24 year olds are 25% higher. Beyond this age group, workers are more likely to have found their true calling, and age differences in unemployment and separation rates among 25 to 64 year old workers are compressed compared to the baseline calibration.

Finally, we increase the value of f_B from 0.57 to 0.95, leaving $f_G = 1$. This implies that the productivity gain from working at one's true calling relative to all other occupations is only 5%. As displayed in Panel E of Table 4, the model's predictions for unemployment and

separation rates by age are unchanged relative to the baseline calibration. For even a small compensation gain of discovering one's true calling, workers are willing to leave matches that turn out to be poor occupational fits in favor of unemployment. Hence, our model's results are extremely robust to the calibration target for f_G/f_B.

5.4 Counterfactual Exercise

The empirical results of Shimer (1998) indicate that the age composition of the labor force has a causal impact on the level of aggregate unemployment. Specifically, the entrance of the postwar baby boom generation into the workforce in the late-1970s and their subsequent aging accounts for a substantial fraction of the rise and fall in U.S. unemployment observed in the past 50 years. In this subsection, we address this issue within the context of our structural model of life-cycle unemployment. We determine what fraction of the observed change in aggregate unemployment can be accounted for by the "learning about occupational fit" mechanism embodied in our analysis.

Specifically, we keep all parameters of the model identical to the benchmark calibration, except for the value of the labor market survival probability, σ. As discussed in subsection 5.1, σ determines the model's age distribution of the labor force. By varying this parameter, we are able to "tilt" the age distribution to resemble that of the late-1970s, when the baby boom first entered the labor force, and compare the resulting unemployment rate predictions to those when the age distribution resembles the recent period.

Panel A of Table 5 replicates the U.S. data averaged over the 1976:6–2012:11 period, as well as results for the model calibrated to this time period, as presented in Table 3. It also provides additional information pertaining to the age distribution of the labor force. Given that the model has only one parameter, σ, to match the age distribution, the benchmark calibration does a surprisingly good job at replicating the U.S. data.

The first two rows of Panel B present U.S. data averaged over 1976–1983, a period featuring the entrance of the baby boom generation to the workforce. This is evidenced by the larger share of the labor force aged 35 years and younger, relative to 1976–2012 as a whole. This early period is also marked by higher aggregate unemployment.

Since this is a short, eight year time period, a simple time-series average of the unemployment rate is unlikely to eliminate the effects of business cycle fluctuations. This is particularly true since 1976–1983 featured many years of recession and relatively few "boom" years. As a result, we use the HP filter to eliminate the cyclical component of the

Table 5: Labor Market Statistics by Age Group: Counterfactual

	20–24	25–34	35–44	45–54	55–64	20–64
A. 1976–2012						
U.S. Data						
labor force share	12.90	27.90	25.88	21.01	12.31	100
unemployment rate	10.45	6.37	4.81	4.22	4.01	5.75
Model						
labor force share	15.68	27.29	22.65	18.79	15.59	100
unemployment rate	12.21	7.33	3.58	3.06	3.04	5.75
B. 1976–1983						
U.S. Data						
labor force share	16.77	30.50	21.69	18.27	12.78	100
unemployment rate	11.87	7.22	5.01	4.37	4.04	6.59
Model						
labor force share	17.77	29.23	22.47	17.27	13.27	100
unemployment rate	12.57	7.60	3.71	3.15	3.13	6.25
C. 2002–2009						
U.S. Data						
labor force share	10.95	23.75	26.05	25.16	14.10	100
unemployment rate	9.75	5.80	4.53	4.01	3.84	5.17
Model						
labor force share	13.29	24.80	22.59	20.58	18.74	100
unemployment rate	11.96	7.13	3.50	2.99	2.98	5.32

Notes: U.S. data from the CPS. Model statistics computed from the theoretical steady state distribution for various parameter specifications.

unemployment data.[23] Given our interest in "steady state" analysis and the impact of low-frequency change in demographics, we present the average unemployment rate of the HP filtered trend. As Panel B of Table 5 indicates, aggregate unemployment was high in this period: 6.59%, as compared to 5.75% during 1976–2012. This is true of each age-specific unemployment rate as well.

The bottom two rows of Panel B present results from our model calibrated to this time period. In particular, we lower the survival probability to $\sigma = 0.9978$ in order to minimize the sum of squared residuals between the model's labor force shares in each of the five age bins with the average values observed in the data. As a check, we find that this value of σ

---

[23]Since the HP filter performs poorly near the endpoints of time series, we extend the unemployment rate data to 1950 before applying the filter.

implies an average worker age of 37.85 years in the model as compared to an average age of 37.90 in the U.S. data.

Panel B indicates that the model does well in replicating the higher aggregate unemployment rate observed during this period. Lowering σ to match the younger labor force composition raises the 20–64 year old unemployment rate to 6.25%. This is close to the value of 6.59% found in the data.

Panel C of Table 5 presents results for the period 2002–2009. Obviously, the labor force is noticeably older given the aging of the baby boom. The aggregate unemployment rate in the data is also noticeably lower.[24] The bottom rows of Panel C present results from the model. To best match the labor force age distribution during this period, we set $\sigma = 0.9992$. This implies an average worker age of 40.50 in the model as compared to 40.54 in the U.S. data. Again, the model does a good job of accounting for observed unemployment. Aging of the workforce causes the aggregate unemployment rate to fall to 5.32%, as compared to 5.17% in the data.

Hence, aging of the labor force—in a manner that best matches that observed from the mid-1970s to present—causes aggregate unemployment to fall by 0.92% in our model.[25] This compares with a fall of 1.42% in the U.S. data. As such, the model rationalizes 65% of the observed change in aggregate unemployment.

Moreover, the change in aggregate unemployment generated by the model is not due simply to the mechanical effect of changing labor force shares. This can be seen by comparing the age-specific unemployment rates across Panels B and C. As the labor force ages, the unemployment rate of 20–24 year olds falls from 12.57% to 11.96%; the unemployment rate of 45–54 year olds falls from 3.15% to 2.99%. Indeed, this "behavioral" implication of our model is consistent with the U.S. data: labor force aging has been accompanied by lower unemployment rates for workers of all age groups.

In our model, this equilibrium effect on unemployment operates through the free entry condition (equation (9) with $V = 0$). As the labor force ages, the endogenous distribution of unemployed worker types shifts, so that there are fewer low-type workers, and more workers who know their true calling. Hence, for any given vacancy cost, κ, firms are willing

---

[24]Again, given the poor performance of the HP filter at endpoints of time series, we have chosen to end our analysis period in 2009, and not consider the filtered data from the last three years of the sample. Including the data from the Great Recession would increase the observed unemployment rates. Given our calibrated results, discussed shortly, this would improve the model's ability to explain the observed unemployment rate changes via demographic change.

[25]As a point of reference, our results are similar to the empirical results of Shimer (1998), who finds that aging generates a 1.20% fall in aggregate unemployment from the mid-1970s to the late-1990s.

to create more vacancies per unemployed worker, since the expected surplus from a match is greater.[26] This results in a higher job finding rate, and lower unemployment rate, among all workers.

Indeed, it is possible to quantify the importance of this effect on the unemployment rate change generated by the model. Taking the age-specific unemployment rates in Panel B, and weighting them by the 2002–2009 labor force shares of Panel C would imply an aggregate unemployment rate of 5.63%. Hence, of the 0.92% fall in unemployment across the two periods, the mechanical effect of age composition change accounts for $(6.25 - 5.63) \div 0.92 = 67\%$. The other 33% is due to the equilibrium effect of increased job creation. As such, the model embodies a non-trivial internal propagation mechanism to labor force composition change.

6 Conclusion

In this paper, we study one of the key life cycle labor market facts: that unemployment declines as a function of age. We propose a simple model of occupational learning that accounts for this fact. Young workers, who are less likely to have found their "true calling," are more likely to separate from employment matches. Hence, our model correctly predicts that age differences in unemployment rates are due to age differences in separation rates. Our calibrated model does a very good job at quantitatively replicating the age differences in unemployment and separation rates observed in the U.S. data. Moreover, labor force aging in our model accounts for a significant fraction of the fall in aggregate unemployment observed in the past 35 years.

---

[26]See also Pries (2008) who emphasizes the importance of compositional change in the pool of unemployed workers in explaining business cycle fluctuations in unemployment.

A   Distribution of Unemployed Workers

Here we provide details on the distribution of unemployed workers across types. We begin with the laws-of-motion for employment and unemployment in steady state.[27] Letting u_1 denote the number of unemployed workers of type 1:

$$u_1 = (1 - \sigma) + \sigma\left[(1 - \mu(\theta))u_1 + \delta(1 - \lambda)n_{L,1}\right],$$

where $n_{L,1}$ denotes the number of employed workers of type 1 learning about the current occupational fit. In general:

$$n_{L,i} = \sigma\left[\mu(\theta)u_i + (1 - \lambda)(1 - \delta)n_{L,i}\right],$$

for $i = 1, \ldots, M - 1$.

The number of type i workers employed in matches of poor occupational fit is given by:

$$n_{B,i} = \sigma\left[(1 - \delta_{B,i})n_{B,i} + \lambda(1 - p_i)(1 - \delta_{B,i})n_{L,i}\right].$$

Here, $\delta_{B,i} = 1$ if type i bad matches endogenously separate, and $\delta_{B,i} = \delta$ otherwise. The number of unemployed workers of type $i = 2, \ldots, M - 1$ is given by:

$$u_i = \sigma\left[(1 - \mu(\theta))u_i + \delta_{B,i-1}n_{B,i-1} + \lambda(1 - p_{i-1})\delta_{B,i-1}n_{L,i-1} + (1 - \lambda)\delta N_{L,i}\right].$$

Finally, for type $i = M$ workers:

$$n_M = \sigma\left[(1 - \delta)n_M + \mu(\theta)u_M + \sum_{i=1}^{M-1}\lambda p_i(1 - \delta)n_{L,i}\right],$$

$$u_M = \sigma\left[(1 - \mu(\theta))u_M + \delta n_M + \delta_{B,M-1}n_{B,M-1} + \lambda(1 - p_{M-1})\delta_{B,M-1}n_{L,M-1} + \delta\sum_{i=1}^{M-1}\lambda p_i N_{L,i}\right].$$

Given these definitions:

$$\pi_i = \frac{u_i}{\sum_{j=1}^{M} u_j}, \quad i = 1, \ldots, M.$$

---

[27]Out of steady state, variables on the left hand side of all equations would be denoted with a prime ($'$).

B Compensations

There are three functions for compensation in the model. ω_M is the compensation paid to workers who found their "true calling". $\omega_{B,i}$ is the compensation for workers who are employed in their i^{th} occupation, and learned that this is not their true calling. $\omega_{L,i}$ is the compensation for workers who are employed in the i^{th} occupation, and have yet to learn whether this is their true calling.

We solve for these compensation functions by maximizing the values of the Nash products specified in Section 4.2. Given that τ represents the worker's bargaining share, and imposing the equilibrium condition $V = 0$, the solution for any compensation satisfies

$$J = (1 - \tau)S$$

where J is the surplus to the firm from a match, and S is the total surplus from a match, i.e. the sum of firm and worker surplus. We use the Bellman values defined in our model to express S and derive the steady state compensation for each type of match.

Compensation for Type-M Workers in "Good" Match

First note that $S_M = J_M + W_M - U_M$. Therefore we have:

$$
\begin{aligned}
J_M &= (1 - \tau)S_M \\
f_G - \omega_M + \beta(1 - \delta)J_M &= (1 - \tau)\left[f_G - z + \beta(1 - \delta - \tau\mu(\theta))S_M\right] \\
\omega_M &= \tau f_G + (1 - \tau)z + \beta(1 - \delta)J_M - (1 - \tau)\beta(1 - \delta)S_M - (1 - \tau)\beta(-\tau\mu(\theta))S_M \\
&= \tau f_G + (1 - \tau)z + \beta(1 - \tau)\tau\mu(\theta)S_M \\
&= \tau f_G + (1 - \tau)z + \beta(1 - \tau)\mu(\theta)[W_M - U_M],
\end{aligned}
$$

where in the last step we use the result that $W_M - U_M = \tau S_M$.

Compensation for Type-i Workers in "Bad" Match

The total surplus from a bad match is

$$
\begin{aligned}
S_{B,i} &= J_{B,i} + W_{B,i} - U_{i+1} \\
&= f_B - z + \beta(1 - \delta)S_{B,i} - \beta\mu(\theta)\tau S_{L,i+1}.
\end{aligned}
$$

When positive, this surplus can be used to solve for $\omega_{B,i}$ as follows:

$$
\begin{aligned}
J_{B,i} &= (1 - \tau)S_{B,i} \\
f_B - \omega_{B,i} + \beta(1 - \delta)J_{B,i} &= (1 - \tau)\left[f_B - z + \left(\beta(1 - \delta)S_{B,i} - \beta\mu(\theta)\tau S_{L,i+1}\right)\right] \\
f_B - \omega_{B,i} &= (1 - \tau)\left[f_B - z - \beta\mu(\theta)\tau S_{L,i+1}\right] \\
\omega_{B,i} &= \tau f_B + (1 - \tau)z + (1 - \tau)\beta\mu(\theta)\tau S_{L,i+1} \\
\omega_{B,i} &= \tau f_B + (1 - \tau)z + (1 - \tau)\beta\mu(\theta)\left(W_{L,i+1} - U_{i+1}\right).
\end{aligned}
$$

Compensation for Type-i Workers in a "Learning-Phase" Match

Let $\bar{f}_i = p_i f_G + (1 - p_i) f_B$. The total surplus generated by a match in the learning phase is

$$
\begin{aligned}
S_{L,i} \;=\; & J_{L,i} + W_{L,i} - U_i \\[4pt]
=\; & \bar{f}_i - \omega_{L,i} + \beta\lambda p_i(1-\delta)J_M + \beta\lambda(1-p_i)(1-\delta)J_{B,i} + \beta(1-\lambda)(1-\delta)J_{L,i} \\
& + \omega_{L,i} + \beta\lambda p_i(1-\delta)W_M + \beta\lambda p_i \delta U_M \\
& + \beta\lambda(1-p_i)(1-\delta)W_{B,i} + \beta\lambda(1-p_i)\delta U_{i+1} \\
& + \beta(1-\lambda)(1-\delta)W_{L,i} + \beta(1-\lambda)\delta U_i \\
& - z - \beta\mu(\theta)W_{L,i} - \beta(1-\mu(\theta))U_i \\[4pt]
=\; & \bar{f}_i - z + \beta\lambda p_i(1-\delta)\left[J_M + W_M - U_M\right] + \beta\lambda p_i U_M \\
& + \beta\lambda(1-p_i)(1-\delta)\left[J_{B,i} + W_{B,i} - U_{i+1}\right] + \beta\lambda(1-p_i)U_{i+1} \\
& + \beta(1-\lambda)(1-\delta)\left[J_{L,i} + W_{L,i} - U_i\right] + \beta(1-\lambda)U_i \\
& - \beta\mu(\theta)\left[W_{L,i} - U_i\right] - \beta U_i \\[4pt]
=\; & \bar{f}_i - z + \beta\lambda p_i(1-\delta)S_M + \beta\lambda(1-p_i)(1-\delta)S_{B,i} \\
& + \beta(1-\lambda)(1-\delta)S_{L,i} - \beta\mu(\theta)\tau S_{L,i} \\
& + \beta\lambda p_i U_M + \beta\lambda(1-p_i)U_{i+1} + \beta U_i - \beta\lambda U_i - \beta U_i \\[4pt]
=\; & \bar{f}_i - z + \beta\lambda p_i(1-\delta)S_M + \beta\lambda(1-p_i)(1-\delta)S_{B,i} \\
& + \beta(1-\lambda)(1-\delta)S_{L,i} - \beta\mu(\theta)\tau S_{L,i} \\
& + \beta\lambda\left[p_i U_M + (1-p_i)U_{i+1} - U_i\right].
\end{aligned}
$$

With this, we can calculate $\omega_{L,i}$:

$$
J_{L,i} = (1-\tau)S_{L,i}
$$

$$
\begin{aligned}
\bar{f}_i - \omega_{L,i} &+ \beta\lambda p_i(1-\delta)J_M + \beta\lambda(1-p_i)(1-\delta)J_{B,i} + \beta(1-\lambda)(1-\delta)J_{L,i} \\
&= (1-\tau)\Big\{\bar{f}_i - z + \beta\lambda p_i(1-\delta)S_M + \beta\lambda(1-p_i)(1-\delta)S_{B,i} + \beta(1-\lambda)(1-\delta)S_{L,i} \\
&\qquad - \beta\mu(\theta)\tau S_{L,i} + \beta\lambda\left[p_i U_M + (1-p_i)U_{i+1} - U_i\right]\Big\}
\end{aligned}
$$

$$
\omega_{L,i} = \tau\bar{f}_i + (1-\tau)z + (1-\tau)\beta\mu(\theta)\left[W_{L,i} - U_i\right] - (1-\tau)\beta\lambda\left[p_i U_M + (1-p_i)U_{i+1} - U_i\right].
$$

C   Proof of Proposition 1

We start by showing that $U_M > U_{M-1}$. Assume, by contradiction, that $U_M \leq U_{M-1}$. From the value of unemployment, we have

$$W_{L,i} - U_i = \frac{(1-\beta)U_i}{\beta\mu(\theta)}, \quad i = 1 \ldots M,$$

where $W_{L,M} \equiv W_M$. So $U_M \leq U_{M-1}$ implies that $W_M - U_M \leq W_{L,M-1} - U_{M-1}$. Using equilibrium wages,

$$
\begin{aligned}
\omega_M - \omega_{L,M-1} &= \tau(f_G - \bar{f}_{M-1}) + (1-\tau)\mu(\theta)\beta[(W_M - U_M) - (W_{L,M-1} - U_{M-1})] + \lambda(1-\tau)\beta(U_M - U_{M-1}) \\
&= \tau(f_G - \bar{f}_{M-1}) + (1-\tau)(1-\beta)(U_M - U_{M-1}) + \lambda(1-\tau)\beta(U_M - U_{M-1}) \\
&= \tau(f_G - \bar{f}_{M-1}) + (1-\tau)(1-\beta(1-\lambda))(U_M - U_{M-1}).
\end{aligned}
$$

Now we can write the value of type M firm as

$$
\begin{aligned}
J_M &= f_G - \omega_G + \beta(1-\delta)J_M \pm \beta(1-\lambda)(1-\delta)J_M \\
&= \frac{f_G - \omega_G + \beta(1-\delta)\lambda J_M}{1 - \beta(1-\lambda)(1-\delta)},
\end{aligned}
$$

and that for type $M-1$ as

$$J_{L,M-1} = \frac{\bar{f}_{M-1} - \omega_{L,M-1} + \beta\lambda(1-\delta)p_{M-1}J_M}{1 - \beta(1-\lambda)(1-\delta)},$$

where $\bar{f}_i = p_i f_G + (1 - p_i)f_B$. Then we have

$$
\begin{aligned}
(1 - \beta(1-\lambda)(1-\delta))&(J_M - J_{L,M-1}) \\
&= (f_G - \bar{f}_{M-1}) - (\omega_G - \omega_{L,M-1}) + (\beta(1-\delta) - \beta\lambda(1-\delta)p_{M-1})J_M \\
&= (f_G - \bar{f}_{M-1}) - (\omega_G - \omega_{L,M-1}) + \beta(1-\delta)\lambda(1 - p_{M-1})J_M \\
&= (1-\tau)(f_G - \bar{f}_{M-1}) - (1-\tau)(1-\beta(1-\lambda))(U_M - U_{M-1}) + \beta(1-\delta)\lambda(1 - p_{M-1})J_M,
\end{aligned}
$$

that is, $J_M > J_{L,M-1}$. This is a contradiction of Nash bargaining: since $\tau J_M = (1-\tau)(W_M - U_M)$ and $\tau J_{L,M-1} = (1-\tau)(W_{L,M-1} - U_{M-1})$, we have

$$0 < \tau[J_M - J_{L,M-1}] = (1-\tau)[(W_M - U_M) - (W_{L,M-1} - U_{M-1})] \leq 0.$$

The rest of the proof is by induction. Suppose that $U_{i+1} < U_{i+2} \leq U_M$ (the weak inequality is necessary to encompass the case where $i = M - 2$). By contradiction, assume that $U_i \geq U_{i+1}$. From the value of unemployment, we know that $U_i \geq U_{i+1}$ implies that $W_{L,i} - U_i \geq W_{L,i+1} - U_{i+1}$. Using the same strategy as above, we have

$$
\begin{aligned}
\omega_{L,i+1} - \omega_{L,i} &= \tau(\bar{f}_{i+1} - \bar{f}_i) + (1-\tau)(1-\beta)(U_{i+1} - U_i) \\
&\quad - (1-\tau)\lambda\beta[p_{i+1}(U_M - U_{i+2}) + (U_{i+2} - U_{i+1}) - p_i(U_M - U_{i+1}) - (U_{i+1} - U_i)] \\
&= \tau(\bar{f}_{i+1} - \bar{f}_i) + (1-\tau)(1-\beta(1-\lambda))(U_{i+1} - U_i) \\
&\quad - (1-\tau)\lambda\beta[(p_{i+1} - p_i)U_M + (1 - p_{i+1})U_{i+2} - (1 - p_i)U_{i+1}].
\end{aligned}
$$

Finally, the firms' value functions implies that

$$
\begin{aligned}
(1 - \beta(1 - \lambda)(1 - \delta))&(J_{L,i+1} - J_{L,i}) \\
&= (\bar{f}_{i+1} - \bar{f}_i) - (\omega_{L,i+1} - \omega_{L,i}) + (p_{i+1} - p_i)\beta\lambda(1 - \delta)J_M \\
&= (1 - \tau)(\bar{f}_{i+1} - \bar{f}_i) - (1 - \tau)(1 - \beta(1 - \lambda))(U_{i+1} - U_i) \\
&\quad + (1 - \tau)\lambda\beta[(p_{i+1} - p_i)U_M + (1 - p_{i+1})U_{i+2}) - (1 - p_i)U_{i+1}] \\
&\geq (1 - \tau)(\bar{f}_{i+1} - \bar{f}_i) - (1 - \tau)(1 - \beta(1 - \lambda))(U_{i+1} - U_i) \\
&\quad + (1 - \tau)\lambda\beta[(p_{i+1} - p_i)U_{i+2} + (1 - p_{i+1})U_{i+2}) - (1 - p_i)U_{i+1}] \\
&= (1 - \tau)(\bar{f}_{i+1} - \bar{f}_i) - (1 - \tau)(1 - \beta(1 - \lambda))(U_{i+1} - U_i) \\
&\quad + (1 - \tau)\lambda\beta(1 - p_i)(U_{i+2} - U_{i+1}) \\
&\geq 0,
\end{aligned}
$$

again violating Nash Bargaining. It follows that $U_i < U_{i+1}$, $i = 1, \ldots, M - 1$.

References

Brügemann, B. (2008). What elasticity of the matching function is consistent with U.S. aggregate labor market data? Unpublished manuscript.

Diamond, P. A. (1982). Aggregate demand management in search equilibrium. *Journal of Political Economy 90*(5), 881–894.

Esteban-Pretel, J. and J. Fujimoto (2011). Life-cycle labor search with stochastic match quality. Unpublished manuscript.

Gorry, A. (2012). Experience and worker flows. Unpublished manuscript.

Hall, R. E. and P. R. Milgrom (2008). The limited influence of unemployment on the wage bargain. *American Economic Review 98*(4), 1653–1674.

Jovanovic, B. (1979). Job matching and the theory of turnover. *Journal of Political Economy 87*(5, Part 1), 972–990.

Kambourov, G. and I. Manovskii (2008). Rising occupational and industry mobility in the United States: 1968–97. *International Economic Review 49*(1), 41–79.

Kambourov, G. and I. Manovskii (2009). Occupational specificity of human capital. *International Economic Review 50*(1), 63–115.

Menzio, G., I. A. Telyukova, and L. Visschers (2012). Directed search over the life cycle. NBER Working Paper No. 17746.

Mortensen, D. T. (1982). Property rights and efficiency in mating, racing, and related games. *American Economic Review 72*(5), 968–979.

Mortensen, D. T. and É. Nagypál (2007). More on unemployment and vacancy fluctuations. *Review of Economic Dynamics 10*, 327–347.

Moscarini, G. and F. Vella (2008). Occupational mobility and the business cycle. Unpublished manuscript.

Murphy, K. M. and F. Welch (1990). Empirical age-earnings profiles. *Journal of Labor Economics 8*(2), 202–229.

Neal, D. (1999). The complexity of job mobility among young men. *Journal of Labor Economics 17*(2), 237–261.

Nekarda, C. J. (2009, May). A longitudinal analysis of the Current Population Survey: Assessing the cyclical bias of geographic mobility.

Papageorgiou, T. (2012). Learning your comparative advantages. Unpublished manuscript.

Petrongolo, B. and C. A. Pissarides (2001). Looking into the black box: A survey of the matching function. *Journal of Economic Literature 39*(2), 390–431.

Pissarides, C. A. (1985). Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *American Economic Review 75*(4), 676–690.

Pissarides, C. A. (2009). The unemployment volatility puzzle: Is wage stickiness the answer? *Econometrica 77*(5), 1339–1369.

Pries, M. (2004). Persistence of employment fluctuations: A model of recurring job loss. *Review of Economic Studies 71*, 193–215.

Pries, M. (2008). Worker heterogeneity and labor market volatility in matching models. *Review of Economic Dynamics 11*(3), 664–678.

Shimer, R. (1998). Why is the U.S. unemployment rate so much lower? *NBER Macroeconomics Annual 13*, 11–74.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review 95*(1), 25–49.