



What should I be when I grow up? Occupations and unemployment over the life cycle

Martin Gervais ^a, Nir Jaimovich ^{b,d,*}, Henry E. Siu ^{c,d}, Yaniv Yedid-Levi ^c

^a University of Iowa, United States

^b University of Southern California, United States

^c University of British Columbia, Canada

^d NBER, United States

ARTICLE INFO

Article history:

Received 13 February 2014

Received in revised form

17 August 2016

Accepted 18 August 2016

Available online 3 September 2016

Keywords:

Unemployment over the life cycle

Learning about occupational fit

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.

© 2016 Elsevier B.V. All rights reserved.

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 than 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 (see also [Shimer, 1998](#) and the references therein). As we discuss in [Section 2](#), these differences are accounted for solely by the declining age profile in the job separation rate, the rate at which individuals transition from employment to unemployment. Moreover, age differences in separation rates are also large. For example, the separation rate for 20–24 year olds is approximately four times than that of 45–54 year olds.

Though these age differences are well known, surprisingly little quantitative work has been done to account for them in the context of an equilibrium labor market model. In this paper, we present a model that focuses on differences in the separation rate. In our quantitative analysis, we find that the model does a good job of accounting for age differences in unemployment. The model also sheds light on the low-frequency evolution of aggregate unemployment experienced in the U.S., and rationalizes findings regarding wage dynamics discussed in the influential literature on job displacement (see, for instance, [Kletzer, 1998](#) and the references therein). Finally, the model replicates important empirical features of job mobility over the life-cycle that are at odds with models that rely on heterogeneity in match-specific productivity to generate age differences in unemployment.

* Corresponding author at: USC Marshall School of Business, Los Angeles, CA 90089, USA.

E-mail addresses: martin-gervais@uiowa.edu (M. Gervais), njaimo@gmail.com (N. Jaimovich), hankman@mail.ubc.ca (H.E. Siu), yaniv@mail.ubc.ca (Y. Yedid-Levi).

In Sections 3 and 4, we present and characterize our model, a life-cycle version of the search-and-matching framework 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 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 mapping 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. As documented by Kambourov and Manovskii (2008) and shown in the next section, this emphasis on the age profile of occupational mobility is highly supported by U.S. longitudinal data.

In Section 5, we study our framework's quantitative implications. The model does a very good job of replicating the observed age profile of unemployment. This is because 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.

In Sections 6 and 7, we consider two applications of our model. We first demonstrate that accounting for labor force aging enables the model to rationalize a significant portion of the fall in aggregate unemployment observed in the U.S. since the mid-1970s. We then illustrate how a simple extension allows the model to rationalize the heterogeneous responses of earnings and wages to match separations, as documented in the empirical literature on job displacement. Hence, taken together with the results in Section 5, we find that the "learning about occupational fit" mechanism is important for understanding a number of key features regarding life-cycle and aggregate labor market dynamics.

Our paper is not the first to emphasize the role of occupational fit in a labor market search framework. Our framework is 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's (1999) empirical findings that career change falls with labor market experience. However, since Neal's (1999) 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 (2014) also studies learning about unobserved occupational ability in a search framework. However, because the phenomena of interest are different, a number of model features differ. Papageorgiou (2014) 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, as well as aggregate unemployment dynamics. Both papers do well in accounting for the life-cycle pattern of occupational mobility.

Finally, learning about occupational fit is not the only mechanism that can address age differences in unemployment in a frictional labor market model. 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. Since productivity in a 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. (2016) illustrate how this mechanism generates a declining age profile of separation in a life-cycle model: since old workers have drawn more matches, they are more likely to be in ones with higher idiosyncratic productivity, and hence, less likely to separate.¹ By contrast, separation rates in our model fall as workers accumulate knowledge about their occupational fit, a form of human capital acquired through labor market experience. In Section 8, we discuss the key distinctions between these two mechanisms and, specifically, how our model rationalizes facts regarding job mobility and separation over the life-cycle that the match-specific productivity mechanism cannot.

¹ Finally, see also Esteban-Prete and Fujimoto (2014). 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.

Table 2
Average monthly transition rates by age group.

	20–24	25–34	35–44	45–54	55–64
<i>A. Shimer approach</i>					
Unemployment exit rate	39.54	34.24	31.17	28.42	27.34
Employment exit rate	5.58	2.70	1.79	1.39	1.26
<i>B. Direct approach</i>					
Unemployment to employment rate	28.46	26.58	25.50	23.77	21.35
Employment to unemployment rate	2.62	1.51	1.13	0.96	0.84

Notes: data from the CPS, 1976:6–2012:11. All transition rates expressed as percentages. 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.

2. Data

This section documents the empirical observations that motivate our work. We first document large age differences in unemployment. This is followed by evidence on job finding rates, separation rates, and occupational mobility that informs our theoretical approach in Section 3.

We begin with the unemployment rate disaggregated by age, obtained from the Current Population Survey (CPS) for the period June 1976–November 2012.² The first row of Table 1 displays the average unemployment rate for different age groups. 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. 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:

$$u = \frac{s}{s+f}. \quad (1)$$

Holding f constant across age groups, a decreasing age profile for unemployment would require a decreasing profile for s . 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. The approach uses monthly data on employment, unemployment, and short-term unemployment tabulated from the CPS.³ Disaggregated by age, this data is available beginning in June 1976. Panel A of Table 2 displays average transition 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 difference is small when compared to separation rates, as discussed 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. The second row of Table 2 indicates that the separation rate falls monotonically with age. Moreover, the age differences are large. For example,

² 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.

³ We refer the reader to Shimer (2005) for further methodological details.




Fig. 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 separation rate for 20–24 year olds is 4.0 times that of 45–54 year olds. As such, the declining profile of separation rates is integral to any theory of age differences in unemployment.⁴

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. This may give rise to misleading conclusions, particularly if transition rates in and out of the labor force differ sufficiently by age. Panel B of Table 2 displays average transition rates measured from data on individual-level labor market flows. We use the longitudinal aspect of the monthly CPS files to track transitions across states of employment, unemployment, and labor force non-participation.⁵

The first row of Panel B presents the transition 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, these differences alone would imply an increasing age profile for unemployment. The second row presents the transition rate from employment to unemployment. As before, this falls monotonically with age, and the differences are large. The separation rate for 20–24 year olds is 2.7 times that of 45–54 year olds.⁶ Again, this 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 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 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, the PSID released Retrospective Occupation-Industry Supplemental Data Files (Retrospective Files hereafter) in 1999. The goal was to retrospectively assign 3-digit 1970 census codes to the reported occupations and industries of respondents for the period 1968–1980. Kambourov and Manovskii (2008, 2009) argue that the methodology minimizes the error in identifying true occupation switches.⁷

Fig. 1 shows the life-cycle profile of occupational mobility rates at the 1-, 2-, and 3-digit level.⁸ While the data and methodology underlying Fig. 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 household heads who are not self-employed and are either employed or temporarily laid-off.⁹ The main

⁴ See also Topel (1984) and Shimer (1998) who emphasize the importance of higher separation rates in understanding higher unemployment rates for the young.

⁵ 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. Since each household member is uniquely identified, we obtain a longitudinal record for each person to construct monthly transitions. For a complete description of the longitudinal data, see Nekarda (2009).

⁶ Note that the levels of transition rates naturally differ across the two measurement approaches. See also Menzio et al. (2016) 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.

⁷ For example, approximately half of occupation switches in the original data are identified as true occupation switches in the Retrospective Files. A key aspect of the methodology 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 an individual over time.

⁸ See online appendix of Kambourov and Manovskii (2008) for details of the occupation classification system.

⁹ This implies that we do not include individuals who leave the labor force in the calculation. 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”). If we count an individual who leaves the labor force as a non-switcher we generate a downward bias in the occupational mobility measure. Excluding those who leave the labor force from the sample avoids this bias.

Table 3
Rates of occupational change by age group.

	Aggregate	20–24	25–34	35–44	45–54	55–64
<i>A. From unemployment</i>						
1-digit (%)	51.24	60.67	50.01	44.42	42.06	39.01
2-digit (%)	55.30	64.81	54.17	48.25	45.88	42.83
3-digit (%)	68.88	77.67	68.68	62.89	60.61	57.35
<i>B. From employment</i>						
1-digit (%)	3.32	5.39	3.52	3.01	2.79	2.53

Notes: data from the CPS, 1994:01–2013:11. Panel A (Panel B) reports the fraction of workers who report a different occupation across consecutive months, conditional on being unemployed (employed) in the first month, and employed in the second.

difference from their sample is our inclusion of government workers and farmers, who tend to switch occupation less frequently than others.

Fig. 1 reveals that occupational mobility declines sharply as people age (see also Moscarini and Vella, 2008 for evidence from CPS data). For the young, mobility rates are very high: more than 40% of individuals in the 20–24 year old age group switch 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 change 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. 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. The same procedure indicates that the average person works in 7.3 and 6.6 occupations, at the 2- and 1-digit level, respectively.

This supports the notion that individuals learn about their true calling over the life cycle. In the model of Section 3, workers who switch occupations do so through intervening spells of unemployment. As such, we investigate data on occupational mobility conditional on transiting from unemployment to employment. We use the individual-level, matched CPS data during the 1994–2013 period and calculate, among those who transit from unemployment to employment, the fraction whose current occupation differs from the occupation during their previous employment spell.¹⁰ Panel A of Table 3 displays these statistics at the 1-, 2-, and 3-digit level. A number of results are worth noting.

First, the probability that a worker switches occupations across employment spells when she transits through unemployment is very high. On average, approximately 50% of such transitions involve a 1-digit level occupation switch, and nearly 70% involve a 3-digit occupation switch. We also find that on average, the occupation switching probability from unemployment-to-employment is approximately 15 times that of the occupation switching probability for employment-to-employment transitions at the 1-digit level (displayed in Panel B), and at the 3-digit level as well (not displayed for brevity).¹¹ Most importantly, we find that the rate of occupation switching is decreasing with age: for instance, at the 1-digit level, it falls from about 61% for 20–24 year olds to 39% for 55–64 year olds. Combined with the fact that the young experience much higher separation rates, this evidence is directly supportive of our mechanism: since the young are less likely to have found their true calling (relative to the old), they are also much more likely to experience occupation switching via a spell of unemployment.

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

¹⁰ We use data beginning in 1994 when the redesign of the CPS significantly improved data quality (due to the introduction of dependent coding) regarding occupation classification at the monthly frequency (see Kambourov and Manovskii, 2013 and Moscarini and Thomsson, 2007).

¹¹ Since the pool of employed individuals is much larger than the pool of unemployed individuals, the number of occupational switches from job-to-job is nevertheless non-negligible.

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. To learn whether a 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 $\lambda \in (0, 1)$.¹²

Given stochastic match creation, destruction, and learning, workers are heterogeneous with respect to their labor market history. This heterogeneity can be summarized by a worker's *type*, $i \in \{1, \dots, M\}$; 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”). As such, a worker's type summarizes her knowledge about her occupational fit, a form of human capital accumulated (stochastically) with labor market experience.

Workers understand that their true calling is uniformly distributed across occupations, and use this to form beliefs. An unemployed worker of type $i = 1$ has a flat prior over all 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. Her true calling is now uniformly distributed over the remaining $M - 1$ occupations, and randomly picks one of those to search in, the next time she is unemployed. Defining π_i as the probability, conditional on being matched in the i th occupation, that the worker has found her true calling, we have that:

$$\pi_i = \frac{1}{M - i + 1}, \quad i = 1, \dots, M. \quad (2)$$

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

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 $i \in \{1, \dots, M\}$. 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. 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 denote market tightness in the representative market, θ , as the ratio of the number of vacancies maintained by firms to the number of workers searching in that occupation. While θ is an equilibrium object, it is taken parametrically by agents. The probability that a worker will meet a vacancy—the job finding rate—is denoted $p(\theta)$, where $p: \mathcal{R}_+ \rightarrow [0, 1]$ is a strictly increasing function with $p(0) = 0$. Similarly, $q(\theta)$ denotes the probability that a vacancy meets a worker, where $q: \mathcal{R}_+ \rightarrow [0, 1]$ is a strictly decreasing function with $q(\theta) \rightarrow 1$ as $\theta \rightarrow 0$, and $q(\theta) = p(\theta)/\theta$.

Finally, it is worth noting 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 [Section 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 of a match, compensation is bargained with complete knowledge of 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 their true calling; or in a match of yet unknown occupational fit.

¹² This constant hazard/learning rate can be viewed as the reduced form of a 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, observed output reveals nothing about the occupational fit, while with complementary probability, the signal is perfectly revealing.

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

$$U_i = z + \beta [p(\theta)W'_{L,i} + (1-p(\theta))U'_i], \quad i = 1, \dots, M-1. \quad (3)$$

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 with probability $p(\theta)$, the job finding rate.

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

$$W_{L,i} = \omega_{L,i} + \beta(1-\lambda) \left[(1-\delta)W'_{L,i} + \delta U'_i \right] + \beta\lambda \left[\pi_i \left((1-\delta)W'_M + \delta U'_M \right) + (1-\pi_i) \left((1-\delta)W'_{B,i} + \delta U'_{i+1} \right) \right]. \quad (4)$$

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 . If the worker does not learn, then she remains as type i in the following period.

If the worker learns about her fit in the current occupation, her continuation value is given by the square bracketed term in Eq. (4). With probability π_i , the worker has found her true calling, and 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, with exogenous probability $\delta \in (0, 1)$, she becomes unemployed of type $i+1$, having eliminated an additional occupation as her true calling.

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

$$W_M = \omega_M + \beta \left[(1-\delta)W'_M + \delta U'_M \right]. \quad (5)$$

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:

$$U_M = z + \beta [p(\theta)W'_M + (1-p(\theta))U'_M]. \quad (6)$$

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

$$W_{B,i} = \max \left\{ \omega_{B,i} + \beta \left[(1-\delta)W'_{B,i} + \delta U'_{i+1} \right], U_{i+1} \right\}. \quad (7)$$

Note that the worker chooses either to remain in the match or, if preferred, 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:

$$J_M = f_G - \omega_M + \beta \left[(1-\delta)J'_M + \delta V' \right], \quad (8)$$

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:

$$J_{B,i} = \max \left\{ f_B - \omega_{B,i} + \beta \left[(1-\delta)J'_{B,i} + \delta V' \right], V \right\}. \quad (9)$$

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. We specify the 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, Eq. (11). In this case, the firm's value is given by:

$$J_{L,i} = \pi_i f_G + (1-\pi_i) f_B - \omega_{L,i} + \beta \left[\lambda(1-\delta) \left[\pi_i J'_M + (1-\pi_i) J'_{B,i} \right] + (1-\lambda)(1-\delta) \left[J'_{L,i} + \delta V' \right] \right]. \quad (10)$$

This is composed of the contemporaneous profit (expected match output minus the worker's compensation) plus the continuation value.¹³ With probability λ , the worker and firm learn about the occupational fit. With probability π_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

¹³ As discussed in footnote 12, our model is isomorphic to one in which output in the learning phase of a match is observed with measurement error.

market is defined as:

$$V = -\kappa + \beta \left[q(\theta) \left[\sum_{i=1}^{M-1} \phi_i J'_{L,i} + \phi_M J'_M \right] + (1 - q(\theta)) V' \right]. \tag{11}$$

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

4. Equilibrium and calibration

In this section we define an equilibrium for our economy and offer a short though sharp characterization of stationary equilibria. A detailed description of the parameterization of the economy closes this section.

4.1. Definition of equilibrium

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

1. workers are optimizing: $W_{L,i} > U_i, W_{B,i} \geq U_{i+1}, i = 1, \dots, M-1$, and $W_M > U_M$;
2. firms are optimizing: $J_{L,i}, J_M > V, J_{B,i} \geq V, i = 1, \dots, M$;
3. compensations solve generalized Nash bargaining problems, discussed 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 (see online appendix for details); and
5. the free entry condition is satisfied. That is, $V=0$.

We assume that compensation in a match is determined via generalized Nash bargaining. As this is standard in the literature, details are provided in the online appendix. It is worth noting, however, that the compensation in the learning phase of a match, $\omega_{L,i}$, contains a new term relative to the standard DMP model. We refer to this term, $\beta\lambda[\pi_i U_M' + (1 - \pi_i)U_{i+1}' - U_i']$, as the *information value* of the match. This captures the discounted, expected gain to working and learning about the current occupational fit, and thus augmenting the worker's threat point in future bargaining. Moreover, the information value is *positive*; this follows as an immediate corollary of [Proposition 1](#), which is established below.

4.2. Characterizing equilibrium

Here we establish that in any stationary equilibrium, the value of unemployment increases with workers' type: $U_i < U_{i+1}, i = 1, \dots, M-1$. This result is useful not only as it leads to a sharp characterization of stationary equilibria, but also because unemployed workers in such equilibria have an incentive to become informed about their type or true calling.

Proposition 1. Assume that all bad matches endogenously separate, i.e., $W_{B,i} = U_{i+1}$ for all $i = 1, \dots, M-1$ in Eq. (7). In any stationary equilibrium, $U_i < U_{i+1}$ for $i = 1, 2, \dots, M-1$.

The proof is contained in the online 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: (a) $J_{L,i+1} > J_{L,i}$; (b) $W_{L,i+1} - U_{i+1} > W_{L,i} - U_i$; (c) $W_{L,i+1} - W_{L,i} > U_{i+1} - U_i > 0$, so $W_{L,i+1} > W_{L,i}$; and (d) $\pi_i U_M + (1 - \pi_i)U_{i+1} - U_i > 0$.

This corollary not only states that both firm and worker surplus monotonically increase with workers' type, but also that the information value of any match in a learning phase is strictly positive.

Finally, we note that when all bad matches separate, only type $i=M$ matches separate at the exogenous rate δ . All other matches—namely, learning-phase matches—separate at rate $\delta + (1 - \delta)\lambda(1 - \pi_i) > \delta$. Moreover, $\pi_{i+1} > \pi_i$. Hence, our model obtains a declining age profile of separation because older workers are more likely to be of higher type.

4.3. Calibration

Many of our model features are standard to the DMP literature, so our strategy is to maintain comparability wherever possible. The model is calibrated to a monthly frequency. 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:

$$p(\theta) = \theta q(\theta) = \theta^\alpha. \quad (12)$$

Summarizing the 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 \(2005\)](#), [Mortensen and Nagypál \(2007\)](#) and [Brügemann \(2008\)](#) obtains a range of 0.37–0.46. We specify $\alpha = 0.4$ to be near the mid point of these ranges (see also [Pissarides, 2009](#)). For comparability with previous work, the parameter in the Nash bargaining problem is specified as $\tau = 1 - \alpha$.

To calibrate δ and κ , two aggregate moments are targeted. First, the aggregate unemployment rate among the model's 20–64 year old workers is set to equal 5.75%. This corresponds to the average unemployment rate observed among 20–64 year olds in the CPS data, between 1976:6 to 2012:11. Second, the aggregate job finding rate is set to equal 38.5% in order to match the job finding rate observed during the same time period in the U.S. data. These two moments pin down the parameters to be $\delta = 0.012$ and $\kappa = 0.235$.¹⁴

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. The survival probability, σ , is set 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 their average values in the CPS data, 1976:6 to 2012:11. This results in a value of $\sigma = 0.9984$. As a check, this implies an average worker age of 39.02 years in the model, which matches the average age found in the data to the second decimal place.

Our calibration of z , the flow value of unemployment, follows the strategy of [Hall and Milgrom \(2008\)](#), [Mortensen and Nagypál \(2007\)](#), and [Pissarides \(2009\)](#). 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, the model's Nash bargained compensation, and the stationary 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 \(2008\)](#).

In the model exposition of [Section 3](#), we defined M as the total number of occupations in the economy. Taking this literally would imply an unreasonably large value; for instance, there are currently about 500 census occupation codes at the 3-digit level. A more nuanced reading of the model reveals that M is the number of occupations an individual views as being in her set of potential “best fits” with her skills/tastes when she first enters the labor market.¹⁵ Since M is clearly not observable, we calibrate it 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 $M = 19$; defining occupations at the 1-digit level would imply setting $M = 13$. 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.¹⁶

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 occupational 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. 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

¹⁴ More precisely, we normalize θ to one, and use a match efficiency parameter η in the matching technology, so that $p(\theta) = \eta\theta^\alpha$, with $\eta = 0.385$. For any value of δ , then, the free entry condition implies a value for κ . We then iterate on δ until the model replicates an unemployment rate of 5.75% among the 20–64 year olds.

¹⁵ Under this interpretation, individuals are homogenous in M but can be heterogeneous in the identities of those M occupations. In this way, the model can be rationalized with the many ($> M$) occupations observed in the data; all that is required is that the distribution of occupation identities across the continuum of individuals' choice sets is consistent with the distribution of agent types within an occupation, so that the firm's value of vacancy creation satisfies Eq. (11).

¹⁶ We note that [Topel and Ward \(1992\)](#) find that the typical male high school graduate holds approximately 10 or 11 jobs throughout his career. Given the different datasets, time periods, and sample selection criteria employed, their statistics on job mobility need not correspond with ours on occupational mobility. Nonetheless, it is worth noting the similarity in mobility, and that the number of jobs exceeds the number of occupations (since not all job changes involve an occupation change).

Table 4

Labor market statistics by age group: data and model.

	20–24	25–34	35–44	45–54	55–64
<i>A. Unemployment rate</i>					
U.S. data (%)	10.45	6.37	4.81	4.22	4.01
Model (%)	12.16	7.30	3.57	3.04	3.03
<i>B. Separation rate</i>					
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 stationary distribution of the baseline calibration.

accumulated 25 years of experience. Normalizing $f_G=1$ and matching these statistics in our model requires setting $f_B=0.57$ and $\lambda=1/18$.¹⁷ 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.

5. Accounting for age differences in unemployment

Section 5.1 presents results for the baseline calibration of our model, specifically, its ability to match the life-cycle labor market facts. Section 5.2 discusses robustness of our results to variation in M , f_G/f_B , and λ .

5.1. Results

The first row of Panel A, Table 4 reproduces the unemployment rate by age displayed in Table 1. The second row displays the age profile of unemployment generated by the 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. 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. Hence, as discussed in Section 4.2, 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 4 where we display separation rates by age.¹⁸ 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. Fig. 2 displays the occupational switching rate at the annual frequency for different age groups in our model; this is analogous to Fig. 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. We address this shortcoming in an extension of the model presented in Section 7.

5.2. 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

¹⁷ 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.

¹⁸ 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.




Fig. 2. Occupational mobility by age: model. *Notes:* Probability of a worker switching occupations, annual frequency. This is constructed by simulating the model economy with a large number of individuals who are “interviewed” on a yearly basis, as in the PSID.

Table 5

Labor market statistics by age group: robustness.

	20–24	25–34	35–44	45–54	55–64
<i>A. U.S. data</i>					
Unemployment rate	10.45	6.37	4.81	4.22	4.01
Separation rate	5.58	2.70	1.79	1.39	1.26
<i>B. Benchmark model</i>					
Unemployment rate	12.16	7.30	3.57	3.04	3.03
Separation rate	5.33	3.03	1.42	1.21	1.20
<i>C. $M=7$</i>					
Unemployment rate	11.71	6.10	3.99	3.91	3.91
Separation rate	5.11	2.50	1.60	1.57	1.57
<i>D. $M=13$</i>					
Unemployment rate	12.12	8.14	3.77	2.31	2.18
Separation rate	5.31	3.41	1.51	0.91	0.86
<i>E. $\lambda = 1/12$</i>					
Unemployment rate	14.87	6.30	3.02	2.96	2.96
Separation rate	6.72	2.59	1.20	1.17	1.17

Notes: U.S. data from the CPS. Model statistics computed from the theoretical stationary distribution for various parameter specifications. M denotes the total number of occupations and λ represents the probability of learning whether your current occupation is your true calling.

unemployment, consider the case of $M=1$: the model collapses to the standard representative worker DMP model (augmented with a constant survival probability). Panels C and D of Table 5 address this exercise. First, 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.¹⁹

Despite the 30% reduction in the number of potential occupations, the model still delivers sizable age differences in unemployment and separation. Young workers in the model experience an unemployment rate that is 3.0 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 3.25, close to the ratio of 4.0 observed in the data. For similar reasons, Panel D of Table 5 shows that increasing the number of potential occupations to $M=13$ amplifies the age differences in unemployment and separation.

In our second robustness exercise, we increase the learning rate to $\lambda = 1/12$. This represents a 50% increase relative to the baseline calibration, so that workers learn their occupational fit in a match in 12 months, on average. Increasing λ has the

¹⁹ In all robustness exercises, the rest of the parameter values are reset to maintain the remaining calibration targets discussed in Section 4.3. Specifically, we keep targeting the job finding rate of 38.5% by recalibrating κ , and set δ to match a 5.75% working-age unemployment rate.

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 E of Table 5.

Not surprisingly, increasing λ causes workers to spend more of the early part of their career unemployed, as they transition between occupations. Relative to Panel B, 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 among 25 to 64 year old workers are compressed compared to the baseline calibration.

Finally, we note that our main results are essentially unaffected by the value of f_B relative to f_C (recall that we normalize $f_C = 1$). This follows from the fact that the allocation is invariant to this ratio as long as matches endogenously separate when an occupation is revealed to be a poor fit. Under our benchmark calibration, increasing f_B relative to f_C has no impact on the unemployment nor the separation rates as long as $f_B/f_C < 0.92$.²⁰ This invariance result can be derived analytically, and demonstrated simply when $M = 2$. In this case, bad matches endogenously separate as long as

$$f_B < z + \beta p(\theta)\tau \left[\frac{f_C - z}{1 - \beta(1 - \delta - p(\theta)\tau)} \right],$$

highlighting that equilibrium allocations are invariant as long as f_B satisfies a cutoff rule.

5.3. Discussion

Our model has been kept intentionally simple. As a result, there are obvious extensions to be fruitfully considered, both theoretically and quantitatively. For instance, upon learning about occupational fit within a match, agents receive a perfectly informative signal about whether the worker has found her true calling. If instead signals were noisy, separation decisions would be a richer function of prior and posterior beliefs, and might lead to workers switching back to previously sampled occupations, as observed in the data (see, for instance, Papanageorgiou, 2014 for analysis along this dimension). Similarly, the only form of human capital being accumulated in our model is knowledge of one's true calling. On the other hand, if workers accumulated occupation-specific or firm-specific human capital on the job, workers may choose not to switch occupations, even if they learn that they are not in their true calling (see, for instance, Gervais et al., 2015 and Wee, 2016). Finally, we note that all occupation switching in our model occurs through unemployment. This would obviously differ in a model with on-the-job search (see, for instance, Menzio et al., 2016).

6. Labor force aging and aggregate unemployment

Age differences in unemployment imply that changes in the age composition of the labor force can influence the aggregate unemployment rate. Specifically, the entrance of the baby boom generation into the U.S. workforce and its subsequent aging potentially accounts for part of the low-frequency movements in aggregate unemployment. Indeed, based on the time series evidence, Shimer (1998) shows that compositional change due to the baby boom generated a substantial fraction of the rise and fall in unemployment from the 1960s through to the end of the century. More recently, Shimer (2001) and Foote (2007) exploit cross-state variation in the share of young workers and find mixed results.²¹

Given this, we ask 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 , which 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 6 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 4. 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. This is when the baby boom entered the workforce, as evidenced by the larger share of the labor force aged 34 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 unemployment data.²² Given our interest in “steady state” analysis and the impact of low-frequency change in

²⁰ While the allocation is invariant to changes in f_B/f_C in the relevant range, such is obviously not the case for wages.

²¹ Specifically, Shimer (2001) finds that a state's unemployment rate falls when its youth share rises. Foote (2007) shows that this result disappears once the presence of spatial correlation in state-level data is taken into account; indeed, Foote (2007) finds point estimates for the youth share effect on unemployment that are positive and aligned with the findings of Shimer (1998), though with large standard errors.

²² Since the HP filter performs poorly near the endpoints of time series, we extend the unemployment rate data to 1950 before applying the filter.

Table 6

Labor market statistics by age group: counterfactual.

	20–24	25–34	35–44	45–54	55–64	20–64
<i>A. 1976–2012</i>						
<i>U.S. Data</i>						
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
<i>Model</i>						
Labor force share	15.68	27.29	22.65	18.79	15.59	100
Unemployment rate	12.16	7.30	3.57	3.04	3.03	5.75
<i>B. 1976–1983</i>						
<i>U.S. Data</i>						
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
<i>Model</i>						
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
<i>C. 2002–2009</i>						
<i>U.S. Data</i>						
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
<i>Model</i>						
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 stationary distribution for various parameter specifications.

demographics, we present the average unemployment rate of the HP filtered trend. The rightmost column of Panel B indicates that 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. 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 age bin and the average values observed in the data. As a check, we find that this value of σ implies an average worker age of 37.85 years in the model as compared to 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 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 6 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 is also noticeably lower.²³ The bottom rows 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 late-1970's to present—causes aggregate unemployment to fall by 0.93 pp in our model.²⁴ This compares with a fall of 1.42 pp in the U.S. data. As such, the model rationalizes 65% of the observed change in aggregate unemployment.

The changes in aggregate unemployment rates shown in Table 6, both for the actual and the model-generated data, can be decomposed into two components: one that is *mechanical*, due simply to changes in labor force shares, and an equilibrium or *behavioral* response. Comparing the age-specific unemployment rates across Panels B and C highlights that the change in aggregate unemployment generated by the model is not simply mechanical. As the labor force ages, the unemployment rate of 20–24 year olds falls from 12.57% to 11.96%; that of 45–54 year olds falls from 3.15% to 2.99%. 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 operates through the free entry condition (Eq. (11)). As the labor force ages, the endogenous distribution of unemployed worker types shifts: there are fewer low-type workers, and more workers who know their true calling. Hence, for any given vacancy cost, firms are willing to create more vacancies per unemployed

²³ Again, given the properties of the HP filter at endpoints of time series, we end our analysis in 2009, and omit the filtered data from the last three years of the sample. Including the Great Recession would increase the observed unemployment rates. Given our calibrated results, this would improve the model's ability to explain the observed unemployment rate changes via demographic change.

²⁴ As a point of reference, our results are similar to the empirical results of Shimer (1998), who finds that aging generates a 1.20 pp fall in aggregate unemployment from the mid-1970s to the late-1990s.

worker, since the expected surplus from a match is greater.²⁵ This results in a higher job finding rate, and lower unemployment rate, among all workers.

Indeed, it is possible to quantify the relative importance of the two effects on the unemployment rate change generated by the model. The mechanical component takes the age-specific unemployment rates in Panel B, and weights them by the 2002–2009 labor force shares of Panel C, implying an aggregate unemployment rate of 5.63%. Hence, of the 93 basis point fall in unemployment across the two periods, the mechanical effect of age composition change accounts for $(6.25 - 5.63)/0.93 = 67\%$. The other 33% is due to the equilibrium effect of increased job creation.

Repeating this calculation with the U.S. data reveals that the effect of labor force aging implies a 60 basis point fall between 1976–1983 and 2002–2009.²⁶ Given the observed fall in aggregate unemployment, the mechanical component accounts for 43% of the change in the U.S. data. As such, while the model clearly embodies a non-trivial internal propagation mechanism to labor force composition change, it understates the actual behavioral response.

7. Job displacement and occupational mobility

In this section, we consider a simple extension of our model to study the individual-level earnings response to job loss. We illustrate how the mechanism linking job separation, unemployment, and occupational mobility over the life-cycle is capable of rationalizing the heterogeneity of responses documented in the literature on *job displacement*.

Following the seminal work of Jacobson et al. (1993), a number of studies have documented the large and persistent effects of job displacement on earnings and wages, especially for prime-aged workers and those with significant job tenure. We very briefly summarize those findings here.²⁷ Though, the precise concept of a displacement varies from paper to paper, the common defining feature is that it involves an involuntary job separation. In the period of the separation, the average earnings loss is on the order of 25%; much of this effect is due to the direct effect of lost employment. However, even 5 or 10 years later, the average effect of the displacement on earnings is still negative, indicating that workers tend to earn lower wages when subsequently re-employed.

Relative to this average effect, the literature also documents substantial heterogeneity. For instance, from the Displaced Worker Survey, Kletzer (1998) finds that while a third of re-employed workers earn substantially less than on their pre-displacement job, an equal or larger fraction earn the same or more on their new job. Moreover, Stevens (1997) shows that the displacement effects covary with individual characteristics. Specifically, using data from the PSID, Stevens (1997) finds that large and persistent losses are concentrated among workers who switch occupations across spells of unemployment: wage losses of workers who return to the same pre-displacement occupation are less pronounced and recover quickly.²⁸ Furthermore, those who experience persistent wage losses are also more likely to be those who experience subsequent job losses following the initial displacement.

Our framework emphasizing the relationship between job separation and occupational mobility is well suited to addressing these facts. To begin, note that the benchmark model of Section 3 already rationalizes the substantial fraction of workers who experience re-employment wage gains or negligible wage losses following job separation. This is because a worker's type is non-decreasing, and wages are increasing in type ($\omega_M > \omega_{L,i}$ for all $i < M$, and $\omega_{L,j} > \omega_{L,i}$ for $j > i$). To capture occasional large, discrete wage losses, we extend the model by allowing for *occupational displacement* shocks. For simplicity, we assume that only individuals who have found their true-calling (type M) are subject to such shocks.

Specifically, in each period type M individuals face a probability, ψ , that they lose their occupational fit. One interpretation is that from time-to-time, workers lose the skill/taste to perform their true calling occupation. Alternatively, this shock could represent a plant closure or mass layoff that causes a type M worker's occupation to disappear or become obsolete in her local labor market. When hit by the displacement shock, the worker's type is reset: she gets “knocked down” the type ladder to a type $x < M$, and must search again for her new true-calling occupation.²⁹ Given this, the worker switches occupation upon re-employment and experiences a discrete wage loss relative to her previous job. Moreover, given that she is re-employed in an occupation that is likely to be a poor occupational fit, she faces a higher probability of subsequent job separations as she learns about her (new) true calling. All of these features are consistent with the results of Stevens (1997).

²⁵ See also Pries (2008) who emphasizes the importance of compositional change in the pool of unemployed workers in explaining business cycle fluctuations in unemployment. Note also that the model of Shimer (2001) embodies the opposite prediction to ours: incentives for vacancy creation are higher in younger labor markets. This is due to his assumption of match-specific productivity, as opposed to occupation-specific productivity in our analysis.

²⁶ Using a similar methodology for the 16+ population, Elsby and Hobijn (2010) find that the mechanical component accounts for an approximate 100 basis point fall from the 1976–1983 period to 2009.

²⁷ See, for example, Kletzer (1998) for a comprehensive survey.

²⁸ Stevens (1997) finds that the same result applies when one conditions on whether the individual switches industry after job displacement; see also the results of Jacobson et al. (1993). It should be noted that, likely due to sample size limitations of the PSID, Stevens (1997) does not present results when occupation and industry switching are considered simultaneously. Hence, these results may be partially driven by confounding variation.

²⁹ For example, $x=1$ would correspond to a situation where an individual's labor market experience is irrelevant to identifying her new occupational fit. By contrast, $x > 1$ would be a situation where an individual's past experience imparts some knowledge relevant to determining her new true calling.

Table 7
Labor market statistics by age group with displacement.

	20–24	25–34	35–44	45–54	55–64
<i>A. U.S. data</i>					
Unemployment rate	10.45	6.37	4.81	4.22	4.01
Separation rate	5.58	2.70	1.79	1.39	1.26
<i>B. Benchmark model</i>					
Unemployment rate	12.16	7.30	3.57	3.04	3.03
Separation rate	5.33	3.03	1.42	1.21	1.20
<i>C. With occupational displacement</i>					
Unemployment rate	11.08	6.86	3.95	3.63	3.63
Separation rate	4.80	2.84	1.58	1.45	1.45

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

For brevity, we discuss only the type M value functions affected by introducing the occupational displacement shock, and make the remaining details of the model available upon request. The values of being an unemployed and employed type M worker become:

$$U_M = z + \beta[(1 - \psi)(p(\theta)W'_M + (1 - p(\theta))U'_M) + \psi U'_x],$$

$$W_M = w_M + \beta[(1 - \psi)(\delta U'_M + (1 - \delta)W'_M) + \psi U'_x],$$

where U_x denotes the value of being unemployed as type x , and x is the type to which a worker who has been hit by the shock resets to. Similarly, the value of a firm matched with a type M worker becomes:

$$J_M = f_G - w_M + \beta[(1 - \psi)((1 - \delta)J'_M + \delta V') + \psi V'].$$

To analyze the quantitative implications of this model, we calibrate the two new parameters, ψ and x , to key findings from the literature. Kletzer (1998) indicates that approximately one-third of workers experience large wage losses (losses greater than 25%) when they regain employment following a job displacement (i.e., involuntary separation). Data from the Job Openings and Labor Turnover Survey (JOLTS) from 2000:12–2014:11 indicate that around 40% of job separations are involuntary. Since the distinction between various types of separations does not exist in our model, we set ψ such that 13.3% (one-third of 40%) of workers who become unemployed in any given period experience occupational displacement. We adjust the exogenous separation probability, δ , to maintain an aggregate unemployment rate of 5.75% among the working-age population. Finally, we set $x = M - 3$. This implies that workers who are displaced experience a 35% wage cut upon re-employment. Moreover, wage losses are persistent: 5 years after separation, wages are on average 10% lower than pre-displacement levels, consistent with the evidence provided by Stevens (1997).

Table 7 displays labor market statistics by age for the extended model, along with the statistics for the U.S. data and benchmark model (originally displayed in Table 4). As is clear, the life-cycle profiles for unemployment and separation rates are very similar to that obtained by the benchmark model; as such, the model extension continues to do a very good job of capturing observed age differences in unemployment and separation.

Fig. 3 displays occupational mobility by age in the extended model. While all age-groups are affected by the introduction of occupational displacement, this phenomenon primarily affects prime-age and older workers, since they are most likely to have found their true calling. As a result, more than 6% of workers aged 45 and older switch occupation on an annual basis. This closely resembles the switching rate observed in the PSID data, as shown in Fig. 1, and represents a marked increase over the benchmark model, as displayed in Fig. 2. Hence, this simple extension of the model is not only consistent with the findings from the job displacement literature, but also improves upon the benchmark model in rationalizing the life-cycle profile of occupational mobility, while simultaneously capturing the life-cycle profile of unemployment.

8. Comparison with match quality models

As noted in the Introduction, our model is not the only one capable of rationalizing the life-cycle profile of unemployment. Menzio et al. (2016) discuss how a model featuring idiosyncratic worker–firm match productivity, as in Jovanovic (1979), can deliver declining unemployment and job separation by age. Since older workers have had more opportunities to draw from the match quality distribution, they are more likely to be in a match with higher idiosyncratic productivity. Here, we discuss the key distinctions between this match quality mechanism and “learning about occupational fit” as featured in our model.

First, as discussed in Section 7, our model is able to capture the heterogeneity in re-employment wage effects of job displacement shocks. In particular, our model naturally rationalizes the covariation of large wage losses with occupational switching across spells of unemployment, as documented in Stevens (1997). By contrast, models focused on heterogeneity in worker–firm match quality such as Menzio et al. (2016), simply cannot speak to this feature of the data as there is no well-defined notion of an “occupation” that is conceptually distinct from a match or “job.”




Fig. 3. Occupational mobility by age with displaced workers. *Notes:* Probability of a worker switching occupations, annual frequency. This is constructed by simulating the model economy with a large number of individuals who are “interviewed” on a yearly basis, as in the PSID.

More fundamentally, the mechanisms differ in terms of what is being learned in a match. In the [Jovanovic \(1979\)](#) framework, agents learn about the match quality between worker and firm, or *firm-specific* human capital. In our model, agents learn about the quality of the match between worker and occupation. This knowledge of one’s true calling is a form of *general* human capital that workers carry forward into the future. In contrast to the [Jovanovic \(1979\)](#) and [Menzio et al. \(2016\)](#) 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. As such, 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. Note that while [Pavan \(2011\)](#) also finds that the returns to tenure in a career (industry-occupation pair) are important in accounting for life cycle mobility and wage growth patterns, he finds a bigger role for match specific returns than [Kambourov and Manovskii \(2009\)](#).

This difference in the nature of human capital also has important implications for the life-cycle pattern of job mobility. In our model, separation from a match is due to the accumulation of general human capital (knowledge about occupational fit). This generates a lower probability of separation in future matches. By contrast, when human capital is firm-specific, a match separation implies the loss of this human capital, and imparts no information about a worker’s separation rate in future matches. This “resetting” property of unemployment is squarely at odds with the data. In a highly influential paper, [Topel and Ward \(1992\)](#) document that, conditional on starting a new job, the separation rate from that job is a decreasing function of prior labor market experience. For instance, for workers with at least 8 years of labor market experience, the probability that a new job lasts at least three years is 2.2 times greater than for workers with at most 1 year of experience. In models of idiosyncratic worker-firm match productivity, all unemployed workers are identical. Hence, this mechanism predicts no difference in the expected duration of new jobs as a function of prior experience. By contrast, our model is consistent with [Topel and Ward’s \(1992\)](#) finding: workers with greater labor market experience know more about their true calling occupation, and hence, face a lower separation probability in any new match.

To summarize, our model of learning about occupational fit is consistent with the empirical evidence on job displacement, Mincerian wage regressions, and the relationship between job mobility and labor market experience. This allows us to draw a clear distinction between our framework and those emphasizing learning about idiosyncratic match productivity.

9. 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 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. Finally, once augmented with the possibility occupational displacement, the model is consistent with several salient facts, namely that displaced workers who switch occupation experience large, persistent wage losses upon

reemployment and that these individuals tend to experience several unemployment spells following a displacement. In addition, the extended model produces a life-cycle profile of occupational mobility that closely matches the U.S. data.

Acknowledgements

We thank the Editor, Referees, Gadi Barlevy, Paul Beaudry, Bjorn Brügemann, Mick Devereux, Aspen Gorry, Michael Pries, Richard Rogerson, Shouyong Shi, as well as seminar participants at various institutions and conferences for helpful discussion, and Gueorgui Kambourov and Iouri Manovskii for providing us with their data on occupational mobility. Siu and Yedid-Levi thank the Social Sciences and Humanities Research Council of Canada for support.

Appendix A. Supplementary data

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.jmoneco.2016.08.003>.

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. *J. Polit. Econ.* 90 (5), 881–894.
- Elsby, M.W.L., Hobijn, B., Şahin, A., 2010. The labor market in the great recession. *Brookings Pap. Econ. Act.* 2, 1–48.
- Esteban-Pretel, J., Fujimoto, J., 2014. Life-cycle labor search with stochastic match quality. *Int. Econ. Rev.* 55 (2), 575–599.
- Foote, C., 2007. Space and time in macroeconomic panel data: Young workers and state-level unemployment revisited. Federal Reserve of Boston, manuscript.
- Gervais, M., Jaimovich, N., Siu, H.E., Yedid-Levi, Y., 2015. Technological learning and labor market dynamics. *Int. Econ. Rev.* 56 (1), 27–53.
- Gorry, A., 2012. Experience and worker flows. Unpublished manuscript.
- Hall, R.E., Milgrom, P.R., 2008. The limited influence of unemployment on the wage bargain. *Am. Econ. Rev.* 98 (4), 1653–1674.
- Jacobson, L.S., LaLonde, R.J., Sullivan, D.G., 1993. Earnings losses of displaced workers. *Am. Econ. Rev.* 83 (4), 685–709.
- Jovanovic, B., 1979. Job matching and the theory of turnover. *J. Polit. Econ.* 87 (5, Part 1), 972–990.
- Kambourov, G., Manovskii, I., 2008. Rising occupational and industry mobility in the United States: 1968–97. *Int. Econ. Rev.* 49 (1), 41–79.
- Kambourov, G., Manovskii, I., 2009. Occupational specificity of human capital. *Int. Econ. Rev.* 50 (1), 63–115.
- Kambourov, G., Manovskii, I., 2013. A cautionary note on using (March) Current Population Survey and Panel Study of Income Dynamics data to study worker mobility. *Macroecon. Dyn.* 17 (1), 172–194.
- Kletzer, L.G., 1998. Job displacement. *J. Econ. Perspect.* 12 (1), 115–136.
- Menzio, G., Telyukova, I.A., Visschers, L., 2016. Directed search over the life cycle. *Rev. Econ. Dyn.* 19 (1), 38–62.
- Mortensen, D.T., 1982. Property rights and efficiency in mating, racing, and related games. *Am. Econ. Rev.* 72 (5), 968–979.
- Mortensen, D.T., Nagypál, É., 2007. More on unemployment and vacancy fluctuations. *Rev. Econ. Dyn.* 10 (3), 327–347.
- Moscarini, G., Thomsson, K., 2007. Occupational and job mobility in the US. *Scand. J. Econ.* 109 (4), 807–836.
- Moscarini, G., Vella, F., 2008. Occupational mobility and the business cycle. Unpublished manuscript.
- Murphy, K.M., Welch, F., 1990. Empirical age-earnings profiles. *J. Labor Econ.* 8 (2), 202–229.
- Neal, D., 1999. The complexity of job mobility among young men. *J. Labor Econ.* 17 (2), 237–261.
- Nekarda, C.J., 2009. A longitudinal analysis of the Current Population Survey: Assessing the cyclical bias of geographic mobility. Unpublished manuscript.
- Papageorgiou, T., 2014. Learning your comparative advantages. *Rev. Econ. Stud.* 81 (3), 1263–1295.
- Pavan, R., 2011. Career choice and wage growth. *J. Labor Econ.* 29 (3), 549–587.
- Petrongolo, B., Pissarides, C.A., 2001. Looking into the black box: a survey of the matching function. *J. Econ. Lit.* 39 (2), 390–431.
- Pissarides, C.A., 1985. Short-run equilibrium dynamics of unemployment, vacancies, and real wages. *Am. Econ. Rev.* 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. *Rev. Econ. Stud.* 71 (1), 193–215.
- Pries, M., 2008. Worker heterogeneity and labor market volatility in matching models. *Rev. Econ. Dyn.* 11 (3), 664–678.
- Shimer, R., 1998. Why is the U.S. unemployment rate so much lower? *NBER Macroecon. Annu.* 13, 11–74.
- Shimer, R., 2001. The impact of young workers on the aggregate labor market. *Q. J. Econ.* 116 (3), 969–1007.
- Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. *Am. Econ. Rev.* 95 (1), 25–49.
- Stevens, A.H., 1997. Persistent effects of job displacement: the importance of multiple job losses. *J. Labor Econ.* 15 (1), 165–188.
- Topel, R., 1984. Equilibrium earnings, turnover, and unemployment: new evidence. *J. Labor Econ.* 2 (4), 500–522.
- Topel, R., Ward, M., 1992. Job mobility and the careers of young men. *Q. J. Econ.* 107 (2), 439–479.
- Wee, S.L., 2016. Delayed learning and human capital accumulation: The cost of entering the job market during a recession. Unpublished manuscript.