Demographic Change and the Great Moderation in an Overlapping Generations Model with Matching Frictions*

Steven Lugauer University of Notre Dame Department of Economics 719 Flanner Hall Notre Dame, IN 46637 (574) 631-1432 slugauer@nd.edu

September 7, 2010

^{*}I thank Daniele Coen-Pirani for many helpful suggestions and comments. The paper has also benefited from conversations with Finn Kydland, Fallaw Sowell, and Stan Zin.

Demographic Change and the Great Moderation in an Overlapping Generations Model with Matching Frictions

Steven Lugauer University of Notre Dame Department of Economics 719 Flanner Hall Notre Dame, IN 46637 (574) 631-1432 slugauer@nd.edu Abstract:

The fraction of the labor force under 35, or youth share, has been correlated with cyclical GDP volatility over the past several decades. The youth share and business cycle fluctuations were high during the 1970's. Then, as the population aged, output volatility declined. I develop an overlapping generations model featuring search frictions and productivity shocks, in which the age distribution affects cyclical volatility through two channels. First, employment for younger workers fluctuates more, creating a composition effect. Second, inexperienced workers produce less, so firms decide how many jobs to create based on the age distribution. Young job searchers do not necessarily induce firms to post new vacancies. Both this endogenous response by firms and the composition effect increase aggregate volatility when the youth share is high. The model can replicate a large portion of the recent moderation, suggesting an important role for demographics in determining the magnitude of output volatility. The model also captures the pattern of steady state unemployment, steady state separation rates, and employment volatility by age group.

Keywords: Business Cycles, Demographics, Overlapping Generations, Job Search *JEL Classification:* E, J

1. Introduction

This paper develops a tractable overlapping generations (OLG) model in which variation of the age distribution can generate a substantial portion of the observed changes in aggregate cyclical volatility.¹ The model also does well replicating the observed differences in unemployment rates, job-separation rates, and employment volatility by age group.² Figure 1 plots a measure of U.S. gross domestic product (GDP) volatility against time. The graph also shows the fraction of the U.S. labor force under the age of 35, or youth share. The youth share was only about 48 percent in 1967, while GDP volatility was low. Then, the young baby-boom generation began to enter the labor market. By 1982 the youth share had risen to over 58 percent, and GDP volatility had dramatically increased. However, as the population aged, GDP volatility rapidly declined. This large reduction in cyclical volatility has been labeled the Great Moderation.

The model features a search friction. Workers and firms meet randomly and matching takes time. A worker-firm match can be good or bad. Good matches last longer on average. New young workers enter the labor force each period, and the oldest workers retire. Match output depends on the worker's age and a persistent aggregate productivity shock. The age distribution affects aggregate output volatility through two channels - a composition effect and the endogenous response by firms.

The composition effect occurs because employment for young workers fluctuates more than for older people over the cycle. Older workers are likely to be employed in good matches; they have had ample search time. Young workers frequently move in and out of employment because they tend to be in bad matches. Therefore, variation in the jobfinding rate generates more employment volatility for younger workers.³ When the youth share is large, all else constant, aggregate employment volatility is high. High employment volatility translates into high output volatility.

The search friction also contributes to the second channel connecting the age distribution to aggregate output volatility. In the model, firms decide how many jobs to create based on the job searchers' ages because young workers produce less output. To illustrate, consider a negative productivity shock. Expected revenues decrease, so companies post fewer vacancies and the job-finding rate goes down. Employment falls, especially among poorly-matched young workers. The number of people looking for jobs increases. If the labor force is relatively young, then the average productivity level among job searchers may decrease. Firms respond to a reduction in expected match output by posting even fewer vacancies, exacerbating the decline in employment. Thus, the endogenous response by firms propagates the original shock when the population is young. The model's second channel has not been considered in the literature before.

I examine the model's quantitative implications by choosing parameter values to target relevant worker flow statistics. The size of the youngest worker cohort in the model economy changes period-by-period to simulate the U.S. youth share over time. When the population is relatively old in the model economy, aggregate output volatility is low; when the youth share is high, output volatility is high. This relationship captures the main result; the model can replicate much of the observed cyclical volatility pattern. The labor-market based mechanism also can replicate three related empirical regularities: unemployment rates, separation rates, and employment volatility all tend to decrease with age.

The findings in Jaimovich and Siu (2009) help motivate my research question. Using panel-data methods, Jaimovich and Siu (2009) exploit variation in the timing and the magnitude of population changes across G7 countries to show that the age distribution has a (statistically and economically) significant effect on cyclical volatility. In other words, they provide evidence that the youth share is positively correlated with aggregate output volatility in several countries.⁴ I present a theory explaining the Jaimovich and Siu (2009) empirical results by explicitly modeling the labor market. The model includes search frictions and finite-lived workers, allowing for analysis of employment by age.

My approach differs from Jaimovich and Siu (2009) in several ways. First, in my model firms can react to the age distribution, creating general equilibrium effects; whereas, the reduced form approach of Jaimovich and Siu (2009) accounts for only the compositional effects. Second, differences in employment across age groups arise naturally in my framework as a consequence of the matching process and the life-cycle. Jaimovich and Siu (2009) do not model the source of employment volatility across age groups. Finally, Jaimovich and Siu (2009) estimate the age distribution effect averaged across many countries. My model can deliver a full time series given an evolving set of demographics for a specific country. Despite the methodological differences, my results broadly agree with those reported in Jaimovich and Siu (2009).

I borrow heavily from recent papers studying business cycles via Mortensen and Pissarides (1994) style search models, such as Shimer (2005) and Hall (2005). Standard matching models do not include the age distribution. Hence, I extend the search framework to an OLG setting to address the question at hand. Two earlier papers, Rios-Rull (1996) and Gomme, Rogerson, Rupert, and Wright (2004), have imbedded real business cycles in OLG models. Neither paper uses labor matching; although, Gomme, Rogerson, Rupert, and Wright (2004) suggest, but do not pursue, search frictions as a way to examine employment fluctuations.⁵

Finally, many papers address the recent large decline in aggregate volatility. Existing

theories fall into three categories: good luck, good policy, or a structural change in the economy (Stock and Watson 2002). Davis and Kahn (2008) detail why these three explanations fail to be convincing (also see Herrera and Pesavento (2009) for a nuanced analysis of the failure of monetary policy to account for the Great Moderation). Jaimovich and Siu (2009) add demographics as a fourth possibility. Davis and Kahn (2008) make no mention of the demographics hypothesis. My model supports the demographics hypothesis by showing how exogenous variation in the youth share could have caused a substantial portion of the reduction in cyclical volatility associated with the Great Moderation. Modeling how the age distribution affects aggregate fluctuations and demonstrating that this effect is quantitatively large are my two main contributions.

In Section 2, I present data on the youth share and aggregate cyclical volatility. Section 3 develops the model of the labor market. I explain my parameter choices in Section 4. In Section 5, I examine the results quantitatively. Section 6 contains additional discussion of the model's mechanism, and Section 7 concludes.

2. Youth Share and Cyclical Volatility Data

The employment data comes from the Current Population Survey (CPS), and the GDP data comes from the Bureau of Economic Analysis (BEA). I use seasonally adjusted quarterly observations from 1962 through the second quarter of 2006 restricted to individuals aged 16 to 54. The youth share equals the fraction of the labor force under the age of 35. I measure cyclical volatility at quarter t as the standard deviation of a 41-quarter window centered around quarter t of the de-trended, logged series. I remove the trend by applying the Hodrick-Prescott (HP) filter with smoothing parameter 1600 to the entire logged series. Then, I calculate the rolling standard deviation. This method is somewhat standard; see Jaimovich and Siu (2009) for example.

Figure 1 plots the youth share and GDP volatility from 1967 to 2002. The two time series clearly move together. The work force was relatively old during the 1960's. The baby-boom generation entered the labor market during the 1970's, and the youth share increased to almost 60 percent by 1980. Then, as the population aged, the youth share decreased. GDP volatility displays a similar pattern. GDP volatility was relatively low during the 1960's. In the 1970's and early 1980's output fluctuations were high. However, as the youth share decreased, GDP volatility rapidly declined.

The standard deviation of the cyclical component of GDP from 1962–2006 is 1.49 percent; see Table 1. Table 1 also reports the standard deviation of the cyclical component of aggregate employment and employment by age group. Aggregate employment volatility has been lower than GDP volatility at 1.02 versus 1.49 percent. These numbers are based on employment's extensive margin. I have performed similar calculations based on annual total hours for 16–54 year-olds using CPS data from the March supplement. Since the observations are at an annual frequency, I set the HP filter to 10 and use a sliding 9year window. The resulting pattern of cyclical volatility of total hours is shown in Figure 2. Jaimovich and Siu (2009) also examine the volatility of total hours; their findings are similar to what I report here. Furthermore, Jaimovich and Siu (2009) document a large difference in volatility of total hours by age. Young workers experience more employment volatility over the cycle. I find the same relationship on the extensive margin. The standard deviation of the series of deviations from trend employment equals 1.35 percent for young workers (aged 16–34) and 0.72 percent for older workers (aged 35–54) in the CPS data.

The difference between young and old workers suggests a simple compositional explanation for the recent moderation in cyclical fluctuations. The youth share began to shrink around 1983. Consequently, older workers, who typically experience less employment volatility, made up a larger share of the labor force, and aggregate employment volatility declined. However, this simple compositional effect cannot entirely account for the changes in employment volatility. Figure 3a plots employment volatility over time with the data split into two age groups. Figure 3b contains aggregate employment for comparison.⁶ The within age group employment volatility for both young and old workers follows the same pattern as aggregate employment volatility and the youth share. The composition effect alone cannot account for changes in employment volatility within age groups. I argue that general equilibrium effects (e.g. the endogenous response by firms to the age distribution) drive the employment volatility changes within age groups.⁷

Overall, Table 1 and Figures 1–3 suggest that aggregate cyclical volatility is related to the age distribution. When the youth share was high, aggregate volatility was large. The remainder of this paper seeks to explain how the age distribution affects output volatility.

3. OLG Labor Search Model

The model imbeds overlapping generations into a discrete-time labor search environment. In the model, workers and firms maximize the present value of their own lifetime earnings by forming worker-to-firm matches. Workers receive wages when employed, and firms generate earnings by hiring a worker and producing. All workers age each period, with the oldest dying (or retiring) and a new cohort born into the labor force. The population's age distribution depends on the relative size of different generations. Workers can be heterogeneous in terms of age and employment status, but have identical preferences at birth. Workers flow between employment and unemployment because some existing matches end for exogenous reasons and new young workers enter the labor force each period. I use a value function formulation of the model.

The search model's defining feature is the existence of a friction in the labor market; finding a match takes time. Workers may remain unemployed for several periods, and firms must pay a per-period cost to post a vacancy and find a match. Free-entry into the labor market allows firms to drive the equilibrium value of posting to zero, with the matching rates playing the role of prices. I focus on the labor market and do not explicitly model other markets (e.g. goods and savings). All agents possess full knowledge of current aggregate state variables and form rational expectations about the future.

Events within a period unfold as follows. First, matched workers and firms produce together in one-to-one pairings. Output is a function of the worker's age and the current aggregate productivity shock. Second, some worker-firm pairs separate due to retirement, death, or match destruction. Third, firms post vacancies and randomly meet job searchers. Workers do not search while employed, but a separated worker can look for new job immediately.⁸ New matches produce in the next period and can be either good or bad in quality. A match is good with probability ϕ . Good matches last longer on average.

Agents do not observe match quality. Instead, workers and firms form beliefs over the probability their match will be destroyed contingent on how long they have been together. Agents update their beliefs using Bayes' Rule. The expected survival rate for a match of tenure T is:

$$heta^{T} = rac{\phi \left(heta^{g}
ight)^{T+1} + \left(1 - \phi
ight) \left(heta^{b}
ight)^{T+1}}{\phi \left(heta^{g}
ight)^{T} + \left(1 - \phi
ight) \left(heta^{b}
ight)^{T}} ,$$

where T indexes θ , θ^g is the survival rate for a good match, θ^b is the survival rate for a bad match, and $\theta^g > \theta^b$. Agents' beliefs are correct on average, but they never know the quality of their match for sure. A new match has tenure zero, denoted T_0 . The longer a pair stays together the more likely they have a good match. New jobs tend to have a shorter duration than older matches, as in Pries (2004). Neither ϕ nor θ^T change over the cycle.

3.1. Firms

Firms create vacancies at flow cost c and produce upon matching with a worker. Firms cannot age discriminate in terms of hiring or firing.⁹ In equilibrium, firms post vacancies until the expected profit from doing so equals zero. Equation (1) captures this free entry condition:

(1)
$$c = q \delta \lambda \sum_{a=1}^{\bar{a}-1} \frac{s_a}{\sum_{a=1}^{\bar{a}-1} s_a} \sum_{z'} \pi_{zz'} J(a+1, T_0, z').$$

In equation (1), q is the matching rate or probability a vacancy meets a worker. The matching rate decreases with the number of vacancies posted. The parameter δ denotes the discount factor. A worker lives to produce in the next period with probability λ .¹⁰ All workers retire at age $a = \bar{a}$, and a total of s_a workers with age a search for a job in the current period. Next period's values are primed. Given a current aggregate productivity shock of z, the shock in the following period equals z' with probability $\pi_{zz'}$. Firms place value $J(a + 1, T_0, z')$ on a new match with a worker of age a. Table 2 contains a list of the notation.

Equation (2) recursively defines the value of a matched firm:

(2)
$$J(a,T,z) = \beta z \xi_a + \theta^T \lambda \delta \sum_{z'} \pi_{zz'} J(a+1,T+1,z').$$

Each match produces $z\xi_a$ per period. Firms keep share β of the output; the rest goes to the worker. $J(\bar{a}, T, z) = \beta z \xi_{\bar{a}}$ due to the worker's impending retirement. A tenure T match is destroyed with probability $(1 - \theta^T)$. The labor input ξ_a depends on the worker's age, reflecting experience. For now, I assume productivity increases with age at a decreasing rate.¹¹ Workers with the same productivity receive equal wages.

Splitting period-by-period output ensures productive matches never voluntarily break apart. This stark wage rule has been used with search frictions before; see Acemoglu (1999) for example. Cooperative Nash bargaining over total match surplus is a common alternative for wage determination in search models. Bargaining over surplus requires agents to speculate on future job-finding rates. Dividing output into fixed shares does not require agents to form forward looking expectations. The difference is significant, as it greatly simplifies the model. I discuss wages further in Section 6.

The value (2) placed on a job, once filled, does not depend on the age distribution among job searchers. However, the number of jobs created does depend on the distribution. The key decision made by firms is how many vacancies, v, to post given the aggregate productivity shock, z, and the age distribution of searching workers, $\{s_a\}_{a=1}^{\bar{a}-1}$. In equilibrium, firms create jobs until the free entry condition (1) is satisfied.

3.2. Workers

Workers receive share $(1 - \beta)$ of output and discount future wages by δ and λ . Matches survive into the next period with probability θ^T . If a match breaks apart, the worker can immediately search for a new job. New matches have tenure zero, T_0 . Given θ^T , workers' decisions are straightforward.¹² For any job-finding rate p, unemployed workers always search and accept any match, and an employed worker never voluntarily quits. These three claims are a direct consequence of the following proposition.¹³

Proposition: $F_a^{e,T} > F_a^u$ for all *a* and *T*, where $F_a^{e,T}$ and F_a^u stand for the beginning of period present value of all the future wages a worker with age *a* expects to receive when currently employed with tenure *T* and unemployed, respectively.

Proof by Induction: First, consider age $\bar{a} - n$, where n = 0. Unemployed workers receive $F_{\bar{a}}^{u} = 0$ for all z. Employed workers receive $F_{\bar{a}}^{e,T} = (1-\beta) z\xi_{\bar{a}}$, which is greater than zero for all z and all T including T_0 . Thus, a worker maximizing expected wages always searches, accepts any match, and never quits when n = 1 due to the possibility of obtaining positive payouts in the final period of life.

Second, assume $F_{\bar{a}-n}^{e,T} > F_{\bar{a}-n}^{u}$ for some *n*. Consider a worker with age $\bar{a} - n - 1$. The value placed on unemployment is:

$$\mathcal{F}^{u}_{\bar{a}-n-1} = \lambda \delta \left(p \mathcal{F}^{e,T_{0}}_{\bar{a}-n} + (1-p) \mathcal{F}^{u}_{\bar{a}-n} \right).$$

The value of employment is:

$$F_{\bar{a}-n-1}^{e,T-1} = (1-\beta) \, z\xi_{\bar{a}-n-1} + \lambda\delta\left(\theta^{T-1}F_{\bar{a}-n}^{e,T} + (1-\theta^{T-1})\left(pF_{\bar{a}-n}^{e,T_0} + (1-p)F_{\bar{a}-n}^u\right)\right),$$

which can be rewritten as:

$$\begin{split} \mathcal{F}_{\bar{a}-n-1}^{e,T-1} &= (1-\beta) \, z\xi_{\bar{a}-n-1} + \lambda \delta \theta^{T-1} \left(\mathcal{F}_{\bar{a}-n}^{e,T} - p\mathcal{F}_{\bar{a}-n}^{e,T_0} - (1-p) \, \mathcal{F}_{\bar{a}-n}^u \right) \\ &+ \lambda \delta \left(p\mathcal{F}_{\bar{a}-n}^{e,T_0} + (1-p) \, \mathcal{F}_{\bar{a}-n}^u \right) \\ &= (1-\beta) \, z\xi_{\bar{a}-n-1} + \lambda \delta \theta^{T-1} \left(\mathcal{F}_{\bar{a}-n}^{e,T} - p\mathcal{F}_{\bar{a}-n}^{e,T_0} - (1-p) \, \mathcal{F}_{\bar{a}-n}^u \right) + \mathcal{F}_{\bar{a}-n-1}^u \end{split}$$

Thus, $F_{\bar{a}-n-1}^{e,T-1} > F_{\bar{a}-n-1}^{u}$ whenever $F_{\bar{a}-n}^{e,T} > F_{\bar{a}-n}^{u}$, implying that workers with age $\bar{a} - n - 2$ always search, accept all matches, and never quit and concluding the proof by induction.

These worker choices do not depend on the aggregate state. Consequently, the worker side of the model does not enter into aggregate volatility considerations. I present the worker's value functions next to complete the model.

An age *a* worker places value $W(a, T, \{s_a\}_{a=1}^{\bar{a}-1}, z)$ on a match with tenure *T*:

(3)

$$W\left(a, T, \{s_{a}\}_{a=1}^{\bar{a}-1}, z\right) = (1-\beta) z\xi_{a}$$

$$\left(\begin{array}{c} \theta^{T}W\left(a+1, T+1, \{s_{a}'\}_{a=1}^{\bar{a}-1}, z'\right) \\ +p\left(1-\theta^{T}\right)W\left(a+1, T_{0}, \{s_{a}'\}_{a=1}^{\bar{a}-1}, z'\right) \\ +\left(1-p\right)\left(1-\theta^{T}\right)U\left(a+1, \{s_{a}'\}_{a=1}^{\bar{a}-1}, z'\right) \end{array} \right).$$

A worker whose match disintegrates finds a new employer with probability p and does a job-to-job transition. The job-finding rate increases with the number of vacancies posted. Thus, p depends on $\{s_a\}_{a=1}^{\bar{a}-1}$ through the free entry condition (1). With probability (1-p) a searching worker does not immediately meet a firm, so the worker stays unemployed. Equation (4) summarizes the value of being an unemployed worker:

(4)
$$U\left(a, \{s_a\}_{a=1}^{\bar{a}-1}, z\right) = \delta\lambda \sum_{z'} \pi_{zz'} \left(\begin{array}{c} pW\left(a+1, T_0, \{s'_a\}_{a=1}^{\bar{a}-1}, z'\right) \\ +(1-p)U\left(a+1, \{s'_a\}_{a=1}^{\bar{a}-1}, z'\right) \end{array} \right).$$

Unemployed workers find a job with tenure zero at rate p, but they receive zero income while searching.¹⁴ Workers retire at age $a = \bar{a}$, so $U\left(\bar{a}, \{s_a\}_{a=1}^{\bar{a}-1}, z\right) = 0$.

3.3. Stocks of Workers by Age

Let e_a^g stand for the stock of workers with age a in good matches. Similarly e_a^b denotes the number of workers aged a not searching for jobs and in bad matches. The following set of equations (5) update the worker stocks $\{s_a, e_a^g, e_a^b\}_{a=2}^{\bar{a}-1}$:

(5)
$$s'_{a} = (1-p)\lambda\left(s_{a-1} + (1-\theta^{g})e^{g}_{a-1} + (1-\theta^{b})e^{b}_{a-1}\right),$$
$$e^{g'}_{a} = \theta^{g}\lambda e^{g}_{a-1} + p\phi\lambda\left(s_{a-1} + (1-\theta^{g})e^{g}_{a-1} + (1-\theta^{b})e^{b}_{a-1}\right),$$
$$e^{b'}_{a} = \theta^{b}\lambda e^{b}_{a-1} + p(1-\phi)\lambda\left(s_{a-1} + (1-\theta^{g})e^{g}_{a-1} + (1-\theta^{b})e^{b}_{a-1}\right).$$

Some workers losing their jobs immediately form new matches. These job-to-job transitions show up as the terms $p(1 - \theta^g) e_{a-1}^g$ and $p(1 - \theta^b) e_{a-1}^b$ in the second and third lines of (5).

All workers start life as searchers. Thus, s_1 defines the size of a generation. To simplify notation, let:

$$S = \sum_{a=1}^{\bar{a}-1} s_a \, .$$

The stocks of workers change after separations take place and just prior to the matching process. For example, a worker making a job-to-job transition is counted in S for one period and not in e_a^g or e_a^b for that period. When calculating employment statistics, the worker is counted as employed.

3.4. Matching Function and Equilibrium

I follow the literature and use a Cobb-Douglas matching function in vacancies and searchers with scale parameter A and elasticity σ , where m is the number of matches:

$$m = AS^{\sigma}v^{1-\sigma}$$

This function implies the following match probabilities:

(7)
$$p = A\left(\frac{v}{S}\right)^{1-\sigma},$$

(8)
$$q = A\left(\frac{S}{v}\right)^{\sigma}.$$

Given a vector of state variables $\{\{s_a\}_{a=1}^{\bar{a}-1}, z\}$, I define an equilibrium as a list $\{J(a, T, z)\}_{a=2}^{\bar{a}}$ and $\{W(a, T, \{s_a\}_{a=1}^{\bar{a}-1}, z)\}_{a=2}^{\bar{a}}$ for $T = 0...(\bar{a} - 2)$, and $\{U(a, \{s_a\}_{a=1}^{\bar{a}-1}, z)\}_{a=1}^{\bar{a}}$, $p(\{s_a\}_{a=1}^{\bar{a}-1}, z)$, and $q(\{s_a\}_{a=1}^{\bar{a}-1}, z)$ such that:

- 1. The free entry condition (1) holds
- 2. Firms' value functions satisfy equation (2)
- 3. Workers' value functions satisfy equations (3) and (4)
- 4. The match probabilities are given by (7) and (8).

3.5. Impact of a Productivity Shock

The impact of an aggregate shock depends on the age distribution among workers. The age distribution affects cyclical volatility in two connected ways. First, there is a composition effect. Second, there is an endogenous response by firms. The evolution of the worker stocks confounds an exact analytical representation. However, the next paragraphs characterize the model economy's reaction to a change in productivity, z.

Aggregate employment of young workers fluctuates more than for older workers. Consider how the stocks of workers evolve with age, from the set of equations (5). Changes in employment levels occur through p, the job-finding rate. If many searchers have age a, then variation in p has a large impact on next period's stocks of workers aged a + 1. Although note, the job-finding rate does not directly affect employed workers keeping their job, $\theta^b \lambda e_a^b$ and $\theta^g \lambda e_a^g$. The percent of workers employed increases with age because older workers have had longer to find a job and in particular a good job. In some sense e_a^g is an absorbing state, and employment volatility decreases with age. Thus, when there are many older workers in the economy, the impact of a shock is low, all else constant. This relationship generates the composition effect.

To simplify notation, define $J^{e}(a, z)$ as:

$$J^{e}(a,z) = \frac{1}{\beta} \sum_{z'} \pi_{zz'} J(a+1,T_0,z'),$$

and \widehat{c} as:

$$\widehat{c} = \frac{c}{A\beta}.$$

Then, the free entry condition (1) can be rewritten using the matching rate (8) to solve for the equilibrium number of vacancies:

(9)
$$v = \left[\frac{\delta\lambda}{\widehat{c}S^{1-\sigma}}\sum_{a=1}^{\overline{a}-1}s_a J^e(a,z)\right]^{\frac{1}{\sigma}}.$$

Consider the total impact of a sustained drop in aggregate productivity, z. The expected value of any match, $J^e(a, z)$, falls. Firms immediately cut back the number of vacancies posted according to equation (9). The job-finding rate, p, goes down according to equation (7). Existing matches continue to separate at the pre-shock rate. However, upon separating from an employer, workers are less likely to find a new job. Employment among young workers declines rapidly because they tend to be in bad, short-lived, matches. If there are many young workers in the economy, then the number of job-searchers increases quickly (the composition effect). The average new searcher has low productivity because ξ_a is small for young workers. Firms react by posting even fewer vacancies (the endogenous response by firms). The job-finding rate, p, decreases further. Employment spirals downward as the composition effect and the endogenous response fuel each other. Conversely, if there are many older workers in the labor force, then the new job searchers tend to be highly productive. Firms react by posting new vacancies, mitigating the original productivity shock.

Thus, the impact of a productivity shock on aggregate employment depends critically on the age distribution in the labor force. This feature of the model encapsulates the main result. A high youth share coincides with high aggregate volatility because of the composition effect and the endogenous response by firms. Next, I simulate the economy to quantitatively examine the theory by choosing parameter values using the general procedure outlined in Kydland and Prescott (1982).

4. Parameter Values

To select parameter values, I use a steady state of the model with the productivity parameter z set to one and a constant population. Table 3 summarizes the parameter choices. Each period represents one month. I base the survival rate on the average mortality rate reported in the U.S. Vital Statistics; $\lambda = 0.9998$.¹⁵ The parameter δ equals 0.9959, part way between the values used in Shimer (2005) and Hall (2005). This choice for δ gives an annual discount rate of 4.8 percent. I restrict agents to 39 years of working life; thus, $\bar{a} = 468$. The resulting youth share equals 49.89 percent, close to the U.S. mean from 1962 to 2006.

Parameter \hat{c} equals 9.455 to target a job finding rate of 0.42, about the percentage calculated in Nagypál (2004). In the model, if a match is destroyed, then the worker immediately searches for a new job. A worker losing his or her job in the current period finds a new employer at the same rate as other searchers because matching is random. Thus, nearly 42 percent of separations lead to job-to-job transitions, which is also close to the percentage reported in Nagypál (2004). In equation (9), \hat{c} depends on the flow cost to post a vacancy c, labor's share of output $(1 - \beta)$, and the matching function scale parameter A. Only the relative value of the three underlying parameters matters for the simulations below.¹⁶ If β equals the usual 0.36 and A the standard 0.45, then c = 1.53 is the present value of the expected revenue from posting a match in equilibrium.

I select the labor input by age, $\{\xi_a\}_{a=2}^{\bar{a}}$, based on individual-level data from the March CPS for the years 1962–2006. I use the fitted values from a regression of weekly wages on a constant, age, age squared, and indicators for gender, education, and race, and year fixed effects. More specifically, I obtain ordinary least squares estimates of d, f, g, and vector h

from:

$$w = d + f \times age + g \times age^2 + h \times X + \epsilon_s$$

where w equals logged annual real wage income divided by the number of weeks worked (mid-point of interval) reported in the CPS, and X contains variables on sex, race, education, and a full set of year fixed effects. I normalize $\xi_{a=2}$ to one. The estimated coefficients for age and age squared are statistically significant at the one percent level using heteroskedasticity robust standard errors. I calculate ξ_a from the estimates (denoted with a hat) as follows:

$$\xi_a = \xi_{a-1} \frac{\exp\left(\widehat{f} \times a\right) \exp\left(\widehat{g} \times a^2\right)}{\exp\left(\widehat{f} \times (a-1)\right) \exp\left(\widehat{g} \times (a-1)^2\right)}; \quad \text{for } a = 3...\overline{a}$$
$$\widehat{f} = 0.0695462; \quad \widehat{g} = -0.0007429; \quad \xi_{a=2} = 1.$$

This simple procedure delivers a set of parameter values consistent with the data. Figure 4 depicts $\{\xi_a\}_{a=2}^{\bar{a}}$. Considerable variation exists; prime age workers have twice the productivity of teens. The value decreases a little for the oldest workers. Gomme, Rogerson, Rupert, and Wright (2004) and Rios-Rull (1996) each calculate and use a similar set of values for the labor input by age.¹⁷

The matching function elasticity parameter σ equals 0.72 as in Shimer (2005). I assume good matches are not destroyed. I choose the probability of a match being good and the survival rate of bad matches to simultaneously target an unemployment rate of 6.10 percent and a monthly separation rate of 7.00 percent (Nagypál 2004). These targets require $\phi =$ 3.35 percent and $\theta^b = 71.01$ percent.

5. Quantitative Results

In quantitative simulations, the model economy generates the four results highlighted in the Introduction. First, the age distribution and output volatility in the model move together as in the data. The model also does well replicating the observed differences in unemployment rates, job-separation rates, and employment volatility by age group.

5.1. Steady State

Table 4 reports unemployment rates by age group for the CPS data and for the steady state of the model. Teenagers and young adults have higher unemployment rates than older workers. The model captures the basic facts. For example, over 17 percent of teenagers are unemployed in the model, but only about 2 percent of the oldest group are out of work.

Table 5 contains total monthly separations by age group. The U.S. data reported in Table 5 originates from Nagypál (2004). Separations by age in the steady state of the model economy display the same pattern as in the data. Young workers are more likely to separate from their employer. Only 2.6 percent of the 45–54 year old age group separates from their employer per period in the model, while 16.6 percent of teenagers separate from their job every month.¹⁸

The differences in separation rates and unemployment rates across age groups arise in the model economy because older workers have had more time to find good quality matches, as captured by the equations (5) governing the stocks of workers. In contrast, young people begin life in unemployment and frequently move in and out of employment. The key parameters in (5) are θ^b and ϕ relative to θ^g , but the steady state results are not sensitive to the exact values used as long as the parameters are selected to match the overall average unemployment and separation rates and $\theta^b < \theta^g$. For example, if θ^g is lowered, then ϕ can be raised to match the separation rate and θ^b has to be lowered to match the unemployment rate; leaving the pattern of unemployment and separation rates by age virtually unchanged.

5.2. Business Cycles with Variable Youth Share

The parameter z takes two values $z^h = 1.027$ and $z^l = 0.973$ and evolves according to the following transition matrix:

$$\Pi = \begin{bmatrix} \pi_{hh} = 0.981 & \pi_{lh} = 0.019 \\ \pi_{hl} = 0.019 & \pi_{ll} = 0.981 \end{bmatrix}$$

This Markov process matches the standard deviation (about 1.6 percent) of the cyclical component of the model output to that of U.S. GDP from 1962–2001.¹⁹ I run the model with a constant population for several hundred periods to expunge the influence of the initial conditions (the steady state). Then, I simulate the economy by altering the size of the youngest generation. Each month, a new shock is drawn, and I change s_1 to approximate the pattern of the U.S. youth share.²⁰ I simulate 160 quarters of data, roughly corresponding to the years 1962–2001. I repeat the entire process 100 times and report on the average across the simulations.

I calculate cyclical output volatility for the model generated time series with the same

procedure I used for the U.S. GDP data. Output volatility at quarter t is the standard deviation of a 41-quarter window centered around quarter t of the de-trended, logged series of total output. I remove the trend using the HP filter with smoothing parameter 1600.

Figure 5 plots the youth share and aggregate output volatility for the simulation. Just as in the U.S. data (see Figure 1), output volatility rises with the youth share, then falls rapidly as the youth share declines. Without the exogenous variation in the youth share the magnitude of the cyclical output volatility in the model economy would be constant. The large swings in GDP volatility, therefore, suggest that the age distribution plays an important role in determining the size of cyclical fluctuations. Figure 5 represents the main quantitative result. The model can replicate the general pattern of output volatility observed over the past several decades.

In Figure 5, output volatility peaks before the youth share.²¹ Output volatility leads the youth share because output volatility depends on the whole age distribution, and the initial increase in the youth share comes from the youngest, most 'volatile', workers. To illustrate, compare the mid-1970s to the late-1980s in Figure 5. Each has similar values for the youth share even though the 1970s had a much higher proportion of teenagers in the labor force. Consequently, the model economy has much higher output volatility in the mid-1970s. Figure 1 shows the same relationship for the US data, although GDP volatility's lead on the youth share is less pronounced. I take the similarity of this lead / lag relationship as a piece of corroborating evidence for the way the age distribution affects output volatility.

To get a sense of scale, I compare the demographic-induced reduction in cyclical volatility in the model to the Great Moderation. GDP volatility decreased by about 50 percent after 1984 (see Figure 1). In the simulation, output volatility falls by about 9 percent over the same time period. Thus, by this calculation, changes in the age distribution can account for 18 percent of the decline in output volatility associated with the Great Moderation.

My estimate of the importance of demographic change to the Great Moderation is in line with the results reported in Jaimovich and Siu (2009). Both studies find a large role for changes in the age distribution in the recent moderation. Jaimovich and Siu (2009) conclude that demographics can explain 10 - 21 percent of the fall in GDP volatility.²² The results from my model simulation are on the high side of the Jaimovich and Siu (2009) estimates because Jaimovich and Siu (2009) only consider compositional effects and have no mechanism for firms to react to changes in the age distribution over the cycle. My model captures the additional impact on aggregate cyclical output volatility due to the endogenous response by firms. As mentioned, this second channel represents a novel contribution of the paper.

To compare the relative importance of each channel (the composition effect and the endogenous response by firms) consider the following. Young people experience greater employment fluctuations over the cycle than older workers in both the US data and in the model (see Table 6).²³ Thus, aggregate volatility is higher when there are more young people in the labor force - the composition effect. Multiplying the reduction in the youth share (10 percentage points) over the time period of the Great Moderation by the difference in employment volatility across the two age groups (0.49) provides an estimate of the mechanical contribution of the composition effect. By this rough calculation, the composition effect accounts for about one-half of the decline in aggregate cyclical volatility in the model. I attribute the remaining drop in employment volatility to the endogenous response by firms.

Jaimovich and Siu (2009) and Jaimovich, Pruitt, and Siu (2010) both suggest a mechanism based on experience-capital complementarity to connect the age distribution to cyclical fluctuations. The idea assumes exogenous 'age group' differences without a true lifecycle process. I build a richer model of the labor market that includes search frictions and explicitly models the aging process, allowing for analysis of employment by age. Differences in employment across age groups arise naturally in my framework as a consequence of the life-cycle interacting with the labor market. Young workers are less likely to be in a good match precisely because of where they are in the life-cycle. That older people have had more time on average to search for a job is undoubtedly true in the real-world, so I model that process. Labor-capital complementarity may also evolve with age, but the approach suggested above does not model how the evolution takes place, leaving the connection between aging and output volatility unclear.

To summarize, my model can reproduce the observed pattern of output volatility. This finding represents the main result. The swings in cyclical volatility caused by the demographic changes appear to be quite large when measured against the recent decline in aggregate fluctuations. The model also generates differences in unemployment rates, jobseparation rates, and employment volatility by age.

6. Discussion

First, this section elaborates on the importance of separation rates by age. Then, I discuss wage bargaining, age discrimination, and on-the-job search. Lastly, I document how the economy reacts to a one time permanent change in the aggregate productivity parameter.

6.1. Separations Over the Cycle

The cyclicality of the separation rate is the subject of an on-going debate.²⁴ Shimer (2005) and Hall (2005) argue that the separation rate is relatively acyclical compared to the job-finding rate, while more recent papers by Fujita and Ramey (2009) and Elsby, Michaels, and Solon (2009) use CPS data to show that cyclicality in the separation rate accounts for as much as 50% of the movement in unemployment over the cycle. I use fixed and exogenous match destruction rates $(1 - \theta^b)$ and $(1 - \theta^g)$; however, the separation rate does change over the cycle in my model, particularly for the youngest workers. The separation rate changes because the mix of good and bad matches depends on the history of aggregate shocks.

Young workers in a recession (i.e. entering the labor market when the job-finding rate is low) have had few opportunities to find a good job in the model economy and therefore have high separation rates. During an expansion, the separation rate for young workers is much lower because more young workers have good jobs. Older workers have only small changes in their separation rates at business cycle frequencies because they tend to have good jobs already based on a lifetime of job searching.²⁵ The difference in the cyclicality of the separation rates between young and old helps explain why the model's propagation of the aggregate shock is stronger when the population is younger. When a large young cohort is born, the firm's job-posting incentive decreases and the cyclicality of the separation rate increases for young workers. Thus, the endogenous response by firms is more sensitive to the aggregate shock when the population is young because of the proportionally larger young worker flows in and out of employment. I further illustrate this point by comparing two economies with different age distributions in Section 6.5.

6.2. Wages

An equilibrium in the model economy essentially consists of firms posting vacancies until they satisfy the free entry condition (1). The simplicity of the solution is due in part to the wage setting rule. Wages equal a fixed share of output as in Acemoglu (1999), Shimer (2001), and Nagypál (2006). Cooperative Nash bargaining over total surplus is the main alternative method used to determine wages in matching models. However, bargaining over surplus could create a counter factual wage distribution in an OLG environment. Young workers may require higher wages than older workers because young workers live longer, creating a large outside option. Thus, the least productive workers might receive the most compensation. Wages would be a function of age rather than just productivity, and young workers would be paid more than older workers net of productivity differences.

The ability of a standard matching model with Nash bargaining to capture the observed business-cycle-frequency fluctuations in unemployment and vacancies is a matter of debate; see Shimer (2005) for example. Recent work downplays the value of unemployment (Hall and Milgrom 2008) and bargaining power (Cahuc, Postel-Vinay, and Robin 2006) for wage determination. My wage rule avoids some of the problems associated with Nash bargaining, but the wages in my model are more volatile than in the data. As already mentioned, the wage mechanism and the information structure over match quality simplify the model. Agents do not have to form expectations over future matchfinding rates. Even though the economy-wide employment stocks are endogenously determined, there is no need to calculate a fixed point rational expectations equilibrium. Future values of the endogenously determined state variables do not enter into agents' decisions, and I am able to generate interesting insights into the labor market using simple surplus splitting.

A more complicated wage mechanism is unlikely to change my results. Consider wages based on the worker's outside option like in Nash bargaining. The output produced by an older worker is high in the present, so a change in current productivity has a relatively large effect on older workers and their outside option. Firms must adjust wages accordingly. A young worker's value comes from future output. The current state has a smaller impact on the worker's outside option. Wages for young workers would change less than the wages of older workers over the cycle. This relationship makes firms more sensitive to aggregate productivity shocks when there are many young workers (similarly to the argument put forward in Hall (2005) regarding wages). Therefore, employment volatility might be even more closely tied to the age distribution if wages were based on the worker's outside option.

6.3. Age Discrimination and Targeted Search

I offer three justifications for assuming firms cannot age discriminate. First, most forms of age discrimination are illegal under the Federal Age Discrimination in Employment Act of 1967. Second, after paying the posting cost, a match with a worker of any age has a positive value in the model; whereas, not finding a match has an equilibrium value of zero. Thus, once matched, firms would rather stay with their match than search for an

older, more productive, worker. Third, using a single labor market for all workers keeps the model tractable and provides the reason why firms care about the entire age distribution.

In reality, firms rarely draw from the whole available labor force. Instead, firms create vacancies for workers with certain skills, experiences, etc. Targeting particular traits is one reason why finding a worker takes time. The matching function (6), as in most search models, captures this friction without explicitly considering the many different labor markets operating in the real world. However, as just mentioned, random matching in a single market is the reason a firm's posting decision depends on the age distribution in my model.

Consider the opposite set-up where each worker type (age) searches in a different labor market, which has its own matching function, unemployment rate, and so on.²⁶ Within any particular market there would be no endogenous response by the firms; however, firms would have to decide what market to enter. The age distribution matters for this decision. The model in this paper can be thought of as approximating both the decision over what type of workers to target and the endogenous response by firms to the age distribution in a simple way that avoids dealing with multiple labor markets. Although, extending the model to include multiple labor markets could yield further insights into how the age distribution affects the business cycle.

6.4. On-the-Job Search

The model does not account for an employed worker's decision to search for a new job while remaining with his or her current employer.²⁷ Given the wage and information structure, workers never benefit, in expectation, from leaving their job. Equation (3) (and the preceding proposition) shows why. The expected value of $W(a, T, \{s_a\}_{a=1}^{\bar{a}-1}, z)$ is greater

than the expected value of $W(a, T_0, \{s_a\}_{a=1}^{\bar{a}-1}, z)$ for all values of T > 0 because the value of a match increases with tenure. No worker would voluntarily leave a job to take a new position, and if there is any cost associated with searching, then no worker would search while employed.

6.5. Labor Market Mechanism

In the full dynamic model, the shocks are transitory; however, there exists a high level of persistence. The following experiment approximates the impact of a change in productivity, at least in the first few periods after a shock. The intention is to provide further insight into the labor market based mechanism by mimicking an impulse response function.

Beginning from the steady state, I increase z by one percent. Figure 6 shows how employment responds after the permanent change. Panel (a) plots the percent difference from the steady state employment level in the months following the shock. The shock occurs in period three. Agents do not know the productivity shock will occur beforehand, but once it happens they know the change is permanent. Employment immediately increases because firms post more vacancies according to equation (9). Then, employment continues to increase as the stocks of workers adjust and firms respond to the new pool of available workers. Panel (b) examines the response by age group. Employment among young workers increases about twice as much as for older workers, in units of percent change. This difference by age group agrees with the data; see footnote 3, Figure 3, and Table 6.

Panels (c) and (d) contain the same information as (a) and (b). Panels (c) and (d) also depict the response of an economy with a survival rate of $\lambda = 0.9978$ (versus 0.9998 in (a) and (b)). This economy has a youth share of 61.37 percent (versus 49.89 percent).

The other parameters are unchanged. Employment jumps up considerably more for the economy with the higher youth share; the change in employment is about 30 percent greater. The within age group responses are also bigger in the economy with the larger youth share (just as suggested in Section 6.1). Thus, this simple experiment indicates that younger populations have higher employment volatility because of both the composition effect and the endogenous response by firms.

7. Final Remarks

Aggregate GDP volatility has been positively correlated with the youth share over the past fifty years. In this paper, I developed a tractable framework that demonstrates how exogenous variation in the age distribution relates to the changes in business cycle volatility. The OLG model features search frictions, idiosyncratic match quality, and aggregate productivity shocks. There are two ways the age distribution affects output volatility in the model economy. First, employment for the young fluctuates more than for older workers, creating a simple composition effect. Second, firms decide how many jobs to create based on the age and productivity profile of the available labor force. Young inexperienced job searchers do not induce firms to post new vacancies. This endogenous response by firms also increases cyclical volatility when the youth share is high. The model can reproduce the general shape of the aggregate volatility pattern observed over the past few decades, generating a substantial portion of the decline in output volatility associated with the Great Moderation. The findings should be of interest to policy makers seeking to understand the amplification and propagation of productivity shocks. Plus, the model environment could be expanded to study other life-cycle issues such as social security reform.

I focused on the age distribution of workers because it evolves naturally and exogenously to the business cycle. Firms also have an interest in the productivity distribution of available workers. In the model presented here, the age and productivity distributions are the same; however, in the data the availability of high-skill workers at any age has increased over time. Also, the labor force participation of married women has increased dramatically. Future research could seek to understand how these and other demographic changes affect the business cycle.

Notes

¹The term volatility refers to the magnitude of the variations from trend at business cycle frequencies. I measure GDP volatility at quarter t as the standard deviation of a 41-quarter window centered around quarter t of the detrended, logged series of total output. See Section 2 for details.

²The term separation refers to the breakup of a worker-firm pair. In my model, separations include retirements, deaths, and exogenous match destruction, and match destruction can result in the worker making a job-to-job transition or becoming an unemployed searcher.

³Empirically, employment volatility among teenagers and young adults is more than twice that of prime age workers. Clark and Summers (1981) was the first paper to report employment volatility by age group. See also, Rios-Rull (1996), Gomme, Rogerson, Rupert, and Wright (2004), and Jaimovich and Siu (2009). Jaimovich and Siu (2009) point out that employment fluctuations for the oldest workers (55+) do not occur at business cycle frequencies. Since I focus on the cycle and old workers constitute a small portion of the labor force, I consider only workers aged 16–54.

⁴The U.S. and Japan make for a compelling comparision. The youth share and aggregate volatility in Japan both decreased in the 1960's; meanwhile in the U.S., the youth share and volatility were increasing.

⁵See Andolfatto (1996) for a search model with business cycles.

⁶Note, Figures 3a and 3b go back to 1953. I use all available data rather than truncating the time series, which contain interesting information on another cycle.

⁷I have carried out similar analysis with the data separated by gender and by race. The same pattern emerges. Employment volatility for both males and females, white and non-white, was high in the 1970's, when the youth share was at its zenith.

⁸Section 6 contains further discussion of job-to-job transitions and on the job search.

⁹Section 6 contains a discussion of age discrimination and targeted search.

¹⁰Time discounting and mortality are standard in many OLG models, so I include them to conform with the existing literature. Note, λ affects the age distribution while δ does not, but neither one is quantitatively important for the results reported in Section 5.

¹¹I find empirical support for this assumption in Section 4, where I select $\{\xi_a\}_{a=2}^{\bar{a}}$ so the model delivers wages by age group consistent with CPS wage income data. For higher

values of a, ξ_a does start to decrease.

¹²Match quality does not directly impact any of a firm's decisions because the value of a firm's outside option always equals zero in equilibrium. However, the information structure over θ^T simplifies the worker side of the model. If a worker knew for certain that he or she had a bad match, then the worker might be tempted to quit in order to search for a good match. In the full knowledge scenario, young workers would be more likely to leave a bad match than older workers. Older workers care less about a job's potential duration because they are closer to retirement. Thus, young workers would move in and out of employment at an even greater frequency relative to older workers, potentially strengthening my mechanism. However, solving the model would be difficult, so I assume agents update their beliefs over time. Several papers use similar assumptions about match quality; see Tasci (2006) and Pries and Rogerson (2005).

¹³To simplify algebra, I assume quitting requires the worker to remain unemployed for one period. The proposition also applies when workers can immediately search because the job-finding rate is less than one and θ^T increases with T.

¹⁴Setting unemployment flow income to zero is an innocuous normalization as long as employment pays more than unemployment in all states of the world.

¹⁵The Centers for Disease Control and Prevention web site contains the U.S. Vital Statistics information on mortality by age. There are differences in death rates across age groups. People aged 20–24 survive to the next month with probability 0.9999 on average; whereas, 50–54 year-olds face a survival rate of 0.9996. I do not account for this difference across age groups, which seems small compared to productivity differences.

¹⁶See the calibration in Shimer (2005) for more on this point.

¹⁷There have been several attempts to estimate the returns to experience. For example, Altonji and Williams (1998) estimate that the return to 10 years of experience on log wages ranges from 0.06 to 0.14. The increase in log wages for 10 years of experience using my calibration averages about 0.12.

¹⁸While the model generated data reported in Table 4 and Table 5 captures the descending pattern in unemployment and separations by age, the fit is not perfect. Unemployment and separation rates are too low for older workers. The model does not address several life cycle issues such as declining participation due to retirement or illness, which could affect the results along these dimensions.

¹⁹Tasci (2006) uses a similar productivity process to calibrate a matching model with a monthly frequency.

²⁰I choose s_1 to minimize the sum of the squared differences between the actual youth

share and the youth share in the model economy using quarterly observations.

²¹I thank the Associate Editor for noticing the youth share's lead on output volatility.

²²Lugauer (2010) makes a comparable estimate based on demographic and aggregate output volatility differences across US States.

²³Employment volatility decreases with age in the model economy in a similar way as unemployment and separations. I aggregate the data into two age groups to correspond with the calculation below and with Table 1 and Figure 3. The results for finer age groups are available upon request.

²⁴I thank Richard Rogerson and an anonymous Referee for helping me better understand the importance of separations by age group over the cycle and for helping frame Section 6.1.

²⁵In the model, the observed probability of job loss decreases as the population ages. Davis and Kahn (2008) provide some empirical evidence that the risk of job loss has indeed decreased over the same time period as the Great Moderation.

²⁶If firms pay a single cost to post to all age groups or if workers can match with any vacancy, then the model returns to the set-up used throughout the paper.

²⁷However, workers do make job-to-job transitions in the model economy. These transitions could be interpreted as capturing the worker flows associated with on-the-job search. In the simulation, the model delivers a large number of job-to-job transistions per month, in line with the data. Thus, the model is not incompatible with on-the-job search, even though it does not explicitly consider the worker's decision to search while employed.

References

- ACEMOGLU, D. (1999): "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence," *American Economic Review*, 89, 1259–1278.
- ALTONJI, J. G., AND N. WILLIAMS (1998): "The Effects of Labor Market Experience, Job Seniority, and Job Mobility on Wage Growth," *Research in Labor Economics*, 17, 233–276.
- ANDOLFATTO, D. (1996): "Business Cycles and Labor-Market Search," *American Economic Review*, 86(1), 112–132.
- CAHUC, P., F. POSTEL-VINAY, AND J.-M. ROBIN (2006): "Wage Bargaining with Onthe-Job Search: Theory and Evidence," *Econometrica*, 74(2), 323–364.
- CLARK, K. B., AND L. H. SUMMERS (1981): "Demographic Differences in Cyclical Employment Variation," *Journal of Human Resources*, 16, 61–79.
- DAVIS, S. J., AND J. A. KAHN (2008): "Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels," *Journal of Economic Perspectives*, 22(4), 155–180.
- ELSBY, M. W. L., R. MICHAELS, AND G. SOLON (2009): "The Ins and Outs of Cyclical Unemployment," *American Economic Journal: Macroeconomics*, 1(1), 84–110.
- FUJITA, S., AND G. RAMEY (2009): "The Cyclicality of Separation and Job Finding Rates," *International Economic Review*, 50(2), 415–430.

- GOMME, P., R. ROGERSON, P. RUPERT, AND R. WRIGHT (2004): "The Business Cycle and the Life Cycle," *NBER Macro Annual*, 19, 415–461.
- HALL, R. E. (2005): "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95(1), 50–65.
- HALL, R. E., AND P. R. MILGROM (2008): "The Limited Influence of Unemployment on the Wage Bargain," *American Economic Review*, 98(4), 1653–1674.
- HERRERA, A. M., AND E. PESAVENTO (2009): "Oil Price Shocks, Systematic Monetary Policy, and the "Great Moderation"," *Macroeconomic Dynamics*, 13, 107–137.
- JAIMOVICH, N., S. PRUITT, AND H. SIU (2010): "The Demand for Youth: Implications for the Hours Volatility Puzzle," *Working Paper*.
- JAIMOVICH, N., AND H. E. SIU (2009): "The Young, the Old, and the Restless: Demographics and Business Cycle Volatility," *American Economic Review*, 99(3), 804–826.
- KYDLAND, F. E., AND E. C. PRESCOTT (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, 50, 1345–1370.
- LUGAUER, S. (2010): "Estimating the Effect of the Age Distribution on Cyclical Output Volatility Across the United States," *Working Paper*.
- MORTENSEN, D. T., AND C. A. PISSARIDES (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61(3), 397–415.
- NAGYPÁL, É. (2004): "Worker Reallocation over the Business Cycle: The Importance of Job-to-Job Transitions," *Working Paper*.

- (2006): "Amplification of Productivity Shocks: Why Don't Vacancies Like to Hire the Unemployed?," *Structural Models of Wage and Employment Dynamics*, 275 of Contributions to Economic Analysis, 481–506.
- PRIES, M., AND R. ROGERSON (2005): "Hiring Policies, Labor Market Institutions, and Labor Market Flows," *Journal of Political Economy*, 113(4), 811–839.
- PRIES, M. J. (2004): "Persistence of Employment Fluctuations: A Model of Recurring Job Loss," *Review of Economic Studies*, 71, 193–215.
- RIOS-RULL, J.-V. (1996): "Life-Cycle Economies and Aggregate Fluctuations," *Review* of *Economic Studies*, 63, 465–489.
- SHIMER, R. (2001): "The Impact of Young Workers on the Aggregate Labor Market," *Quarterly Journal of Economics*, 116(3), 969–1007.
- (2005): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95(1), 25–49.
- STOCK, J. H., AND M. W. WATSON (2002): "Has the Business Cycle Changed and Why?," *NBER Macroeconomics Annual*, pp. 159–218.
- TASCI, M. (2006): "On-the-Job Search and Labor Market Reallocation," Working Paper.

8. Tables and Figures

Table 1: Cyclical Volatility

Description	1962-2006	<u>1962–1984</u>	<u>1985–2006</u>
U.S. GDP	1.49%	1.88%	0.91%
Employment volatili	ty by age group	:	
16–54	1.02%	1.24%	0.72%
16–34	1.35%	1.66%	0.93%
35–54	0.72%	0.83%	0.59%

I constructed Table 1 using quarterly CPS and BEA data from 1962–2006. Cyclical volatility equals the standard deviation of the entire HP filtered, logged, quarterly series expressed in levels. I removed the trend from each series using the HP filter with smoothing parameter 1600. See Section 2 for more details.

Table 2: Notation

Exogenous

Endogenous

Symbol	Description	Symbol	Description
a	Age of worker	e_a	Employed aged a
ā	Retirement age	J^e	Expected value matched firm
A	Matching function scale	J	Value matched firm
С	Cost to post vacancy	m	Number of matches
\widehat{c}	$\frac{c}{A\beta}$, normalized posting cost	p	Job-finding rate
T	Tenure	q	Matching rate
z	Aggregate productivity shock	s_a	Searchers aged a
β	Firm's share of output	S	Total number of searchers
δ	Time discount parameter	U	Value unemployed worker
θ	Match survival rate	v	Number of vacancies
λ	Survival rate	W	Value employed worker
ξ_a	Productivity by age		
π	Markov transition probability		

- σ Matching function parameter
- ϕ Good match probability

Table 3: Parameter Values

Symbol	Description	Value	Target / Reason
ā	Retirement age	468	Work for 39 years
\hat{c}	Normalized posting cost	9.4550	42% job-finding rate
z	Aggregate productivity	1.0000	Normalization
δ	Discount factor	0.9959	4.8% annual discount rate
$ heta^b$	Match survival rate	0.7101	7% separation rate
λ	Survival rate	0.9998	Mortality rate
ξ_a	Productivity by age	*	Fit to CPS data, 1962–2006
σ	Matching function elasticity	0.7200	Shimer (2005)
ϕ	Probability good match	0.0335	6.1% unemployment rate
	Youth share	0.4989	

Table 3 lists the parameter values used in the steady state analysis and dynamic simulations.

Age Group	U.S. Data	Steady State Model
16–19	16.1%	17.8%
20–24	9.4%	10.2%
25–34	5.6%	6.0%
35–44	4.2%	3.2%
45–54	3.7%	2.1%
16–54	6.1%	6.1%

Table 4: Unemployment Rates by Age Group

I calculated the unemployment rate for each age group using CPS data from 1948 through the second quarter of 2006. The parameter value choices for the steady state model can be found in Table 3.

Age Group	U.S. Data	Steady State Model
15–19	20.2%	16.6%
20–24	11.6%	12.1%
25–29	6.9%	8.3%
30–34	5.7%	5.9%
35–44	4.8%	3.9%
45–54	4.3%	2.6%
15–54	7.0%	7.0%

Table 5: Separation Rates by Age Group

Table 5 reports the average monthly separations as a fraction of employment by age group. The U.S. data originates from Table 1 in Nagypál (2004), which was created from CPS data. Note the first age group for the model is aged 16–19. In the model, separations include retirements, deaths, and match destructions. The parameter value choices for the steady state model can be found in Table 3.

 Table 6: Employment Volatility by Age Group

Age Group	U.S. Data	Model
16–34	1.35%	1.10%
35–54	0.72%	0.61%

I constructed Table 6 using CPS data and the data generated from the model as detailed in Section 5. Employment volatility is the standard deviation of the de-trended, logged, quarterly employment series expressed in levels. I remove the trend from each series using the HP filter with smoothing parameter 1600. Figure 1: Youth Share and GDP Volatility (1967-2001)



I constructed Figure 1 using CPS and BEA data from 1962—2006. The youth share equals the fraction of workers aged 16—54 under the age of 35. GDP volatility at quarter *t* is the standard deviation of a 41-quarter window centered around quarter *t* of the de-trended, logged, quarterly series of U.S. GDP. I removed the trend using the HP filter with the smoothing parameter set to 1600.

Figure 2: Youth Share and Hours Volatility (1966-2002)



I constructed Figure 2 using CPS data from 1962—2006. The youth share equals the fraction of workers aged 16—54 under the age of 35. Hours volatility at year *t* is the standard deviation of a 9-year window centered around year *t* of the de-trended, logged, annual series of aggregate hours. I removed the trend using the HP filter with the smoothing parameter set to 10.

Figure 3: Youth Share and Employment Volatility by Demographic Group

3a: By Age Group 3b: Aggregate Volatility 60% 2.15% 16-34 Youth Share Youth Share Youth Share 54% 1.55% Employment 48% 0.95% 35-54 0.35% 42% 1953 1965 1977 1989 2001 1953 1965 1977 1989 2001

Figures 3a and 3b were constructed using CPS data from 1948-2007. The youth share (solid line) equals the fraction of workers aged 16-54 under the age of 35. Employment volatility at quarter *t* is the standard deviation of a 41-quarter window centered around quarter *t* of the HP filtered, logged, quarterly total employment series.

Figure 4: *Productivity by Age*



Figure 4 depicts the labor inputs by age used in the simulations.



I constructed Figure 5 using the same method as Figure 1 and 160 quarters of model generated data. The size of the youngest cohort in the model was chosen to match the observed youth share (% of labor force under 35) pattern. Output volatility at quarter *t* is the standard deviation of a 41-quarter window centered around quarter *t* of the HP filtered, logged, quarterly series of total aggregate output.

Figure 6: *Response to Productivity Increase*



Figures 6a-b track the percent change in employment (from a steady state with youth share 49.98%) after a permanent 1% increase in productivity. The change occurs in month 3. The youth are aged 16—34; the old (dashed line) are 35—54. Figures 6c-d also contain the response of an economy with a higher youth share. The red (lighter / thicker) lines track the employment response when the youth share is 61.37%. The black (darker / thinner) lines track the response in an economy with share 49.98%.