

Why Did the Average Duration of Unemployment Become So Much Longer?*

Toshihiko Mukoyama[†]
University of Virginia
and
CIREQ

Ayşegül Şahin
Federal Reserve Bank of New York

November 2008

Abstract

There has been a substantial increase in the average duration of unemployment relative to the unemployment rate in the U.S. over the last 30 years. We evaluate the performance of a standard job-search model in explaining this phenomenon. In particular, we examine whether the increase in within-group wage inequality and the decline in the incidence of unemployment can account for the increase in unemployment duration. The results indicate that these two changes can explain a significant part of the increase over the last 30 years, although the model fails to match the behavior of unemployment duration during 1980s.

Keywords: unemployment duration, wage dispersion, job search model

JEL Classifications: E24, J64

*We are grateful to an anonymous referee, Mark Bilts, Yongsung Chang, Nikolay Gospodinov, Jeremy Greenwood, Nezih Güner, Mark Huggett, Susumu Imai, Sun-Bin Kim, Bob King, Per Krusell, Lance Lochner, Adi Mayer, Kristy Mayer, Jiyeon Oh, Rob Shimer, Roxanne Stanoprud, Sarah Tulman, Rob Valletta, Gianluca Violante, and seminar participants at the FRS Macro System Meeting, the Montreal Macro Study Group, the Rochester Wegmans Conference, and the SED Meetings for their comments and suggestions. We thank Chris Huckfeldt for excellent research assistance. We are grateful to Melissa Kearney for providing us with CPS May/ORG data. We also thank Emy Sok at the Bureau of Labor Statistics for the help with the BLS data. All errors are our own. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

[†]Corresponding author. Department of Economics, University of Virginia, P.O. Box 400182, Charlottesville, VA 22904-4182, U.S.A. E-mail: tm5hs@virginia.edu. Phone: +1-434-924-6751. Fax: +1-434-982-2904.

1 Introduction

Average unemployment duration and the unemployment rate generally move together; however, recent data indicate that this relationship is changing in the United States. Figure 1 shows the U.S. unemployment rate (left scale) and average unemployment duration (right scale) from 1948 to 2003.¹ Although they track each other closely from 1948 to 1990, we observe a clear break in this pattern during recent years. Specifically, the U.S. unemployment rate declined dramatically over the past 20 years, while average unemployment duration remained high through the 1990s. There are other studies that document the increase in the average duration of unemployment. Juhn, Murphy, and Topel (1991, 2002), Baumol and Wolff (1998), Valletta (1998, 2005), Abraham and Shimer (2001), and Machado, Portugal, and Guimaraes (2006) all point out that unemployment duration has become longer in recent years, despite low levels of unemployment.²

This pattern is more apparent if we look at the trends of the unemployment duration and unemployment rate. Figure 2 compares the Hodrick-Prescott (HP) trend of the two series.³ Casual observation of Figure 2 suggests that the difference between the trends has been particularly pronounced in recent years.⁴ The U.S. unemployment rate

¹The data in this paper are taken from Bureau of Labor Statistics (BLS) website, <http://www.bls.gov>, except where noted otherwise.

²Appendix A discusses the related literature in detail.

³For HP-filtering, we use the smoothing parameter $\lambda = 6.25$ for yearly data following Ravn and Uhlig's (2002) suggestion.

⁴In Appendix B, we show that the cyclical components of both series track each other very well,




Figure 1: Unemployment rate (left scale) and average unemployment duration (right scale).

Data source: Current Population Survey

declined dramatically over the past 20 years, while average unemployment duration remained high through the 1990s.

Figure 2 indicates that the increase in the average duration occurred around 1980.⁵ The average value from 1971 to 1980 is 11.5 weeks, while the average after 1981 is 15.5 weeks. Therefore, there has been an increase of about 4.0 weeks after the break.⁶

even in recent years.

⁵In Appendix B, we identify the break more formally and reach a similar conclusion.

⁶Unemployment duration increased from 13.5 weeks (from 1971 to 1980) to 16.9 weeks (from 1981 to 2002) for males and it increased from 10.7 weeks (from 1971 to 1980) to 13.5 weeks (from 1981 to 2002) for females.




Figure 2: Trends of unemployment rate (left scale) and average unemployment duration (right scale).

Data source: Current Population Survey

We investigate the economic reasons behind the increase in unemployment duration. In particular, we evaluate the performance of the standard job search model in explaining this phenomenon. We first document the two prominent changes in the U.S. labor market that took place recently: the increase in wage dispersion and the decline in the incidence of unemployment. We quantitatively examine how these changes affect unemployed workers' job-search behavior. We argue that these two changes can explain a significant part of the increase in unemployment duration in the last 30 years. However, the model fails to match the behavior of unemployment duration during the 1980s. We conduct some robustness analyses and discuss the generality of our results. Lastly, we examine the link between unemployment duration and the dispersion of wages at a more disaggregated level.

Juhn, Murphy, and Topel (1991) are among the first to point out that the duration of unemployment has been increasing during the recent years. They emphasize the labor-supply response to the change in the wage *level*. Their work can be considered complementary to ours, which emphasizes the change in wage *dispersion*. In her comments to Juhn, Murphy, and Topel (1991), Yellen (1991) argues that:

When there is wage dispersion, so that both good and bad jobs are available for workers with given skills, some workers will choose to remain unemployed, searching for good, rent-paying jobs, rather than work at the poor jobs that are readily available... The long-term unemployed are searching

for work for which they are qualified. In this interpretation, unemployment is a response to wage dispersion rather than to wage levels, contrary to the authors' labor supply function, in which labor supply depends only on wage levels. (pp. 129–130)

Our hypothesis parallels her argument—the recent change in wage distribution may have had a significant impact on unemployment duration.

An alternative explanation for the increase in unemployment duration is institutional change. In fact, the unemployment insurance system did change during the post-war period. Theoretically, if unemployment insurance became more generous, it may have lengthened the duration of unemployment. Baicker, Goldin, and Katz (1997) describe the changes in the unemployment insurance system since the 1930s. Although there were increases in coverage during the 1970s (mainly affecting public sector workers), they argue that the generosity of unemployment insurance has remained almost constant, and that the ratio of unemployment insurance claims to total unemployment has actually declined over the post-war era (see Figure 7.2 in their paper). Consistent with this observation, Burtless (1983) argues that the *insured* unemployment rate has been stable throughout the period from the 1950s to the 1970s.⁷ Therefore, changes in the unemployment insurance system are not likely to be an explanation for the

⁷Baumol and Wolff (1998) report that the insured coverage rate (the percent of unemployed workers receiving benefits) has actually dropped if the data is extended through the 1990s.

longer unemployment duration.⁸ Baumol and Wolff (1998) also support this view. They examine the effect of institutional changes on the duration of unemployment, and conclude that institutional factors like changes in the coverage and generosity of unemployment insurance, in the rate of unionization, and in the minimum wage cannot account for the observed increase in unemployment duration.

Another possible explanation is the change in the demographic composition of the U.S. labor force. In general, older workers tend to be unemployed longer, and women tend to be unemployed for a shorter period of time. In Appendix C, we examine whether demographic change can account for the increase in unemployment duration, by making age and gender compositions that are similar to Shimer (1998). We find that demographic change can explain only a minor part of the increase in unemployment duration.

Section 2 documents the changes in wage dispersion and the incidence of unemployment. Section 3 sets up a job-search model and quantitatively evaluates the effects of these changes on unemployment duration. Section 4 examines the robustness of the quantitative findings. Section 5 examines whether wage dispersion and unemployment duration are correlated at a disaggregated level. Section 6 concludes.

⁸Ehrenberg and Oaxaca (1976) estimate that moving from a welfare system without unemployment insurance to a system similar to the U.S. unemployment insurance system increases the duration of unemployment by one week for older (age 45-59) males, and by less than one week for other demographic groups. This result suggests that even if the increase in coverage increased the unemployment duration, its quantitative impact would be small.

2 Wage dispersion and the incidence of unemployment over time

In this section, we provide an overview of the evolution of wage dispersion and the incidence of unemployment over time. When we move on to calibrate our job-search model, we then use the statistics that are presented in this section.

Many labor economists have documented that there are substantial wage differentials among observationally equivalent workers. Mincer-style wage equations typically explain less than 30% of overall wage variation.⁹ The remaining variation, which is more than 70%, is often called residual (or within-group) wage inequality.¹⁰

Beginning with Katz and Murphy (1992) and Juhn, Murphy, and Pierce (1993), researchers have noticed that there has been a significant increase in residual wage inequality during recent decades.¹¹ In this subsection, we examine the evolution of the change in residual wage inequality, using the March Current Population Surveys (CPS) data covering 1971 to 2002. In Appendix D, we consider two more data sources—IPUMS Census Samples for 1960 to 2000 and the May CPS samples for 1973 to 1978 combined with the CPS Outgoing Rotation Group (ORG) files for 1979 to 2003.

We make use of the data set that Eckstein and Nagypál (2004) put together by

⁹See, for example, the review of the empirical evidence in Mortensen (2003).

¹⁰We use the terms “inequality” and “dispersion” interchangeably.

¹¹Some recent studies investigate the causes of this increase in within-group wage inequality. For example, Violante (2002) argues that the recent rapid investment-specific technological change is the major cause of the increase in within-group wage inequality.




Figure 3: 90% – 10% residual wage inequality, March CPS 1971-2002.

using annual data drawn from the March CPS covering the period between 1971 and 2002.¹² We follow Eckstein and Nagypál’s (2004) log wage regression to control for the observables.¹³ Figure 3 plots the path of the difference between the 90th and 10th percentiles of the residual wage distributions from two regressions, one for males and one for females. Residual wage inequality rose from 1971 to 2002 for both male and female workers. Male residual wage inequality has been increasing steadily (with the exception of a few years) during this 30 year period while female residual wage inequality did not change significantly during 1970s.

As described in the Appendix D, the other two data sources also show the increase in residual wage inequality over the recent years, although the magnitude of the change

¹²The data and their program are downloaded from <http://faculty.wcas.northwestern.edu/~een461/QRproject/>.

¹³See Eckstein and Nagypál (2004) for details. Our Census data regression in Appendix C controls for more detailed occupational groups.




Figure 4: Incidence of unemployment

varies substantially across data sources. When we check the robustness of our findings in Section 4, we examine the quantitative implications of these differences.

Another important change that we focus on is the variation in the incidence of unemployment over time. Figure 4 plots the incidence of unemployment (defined as the number of the workers unemployed for less than five weeks divided by total employment) using BLS data for men and women. Before 1970, the incidence of unemployment was stable at around 0.025. It increased to 0.035 in the 1980s, then decreased to 0.020 in the 1990s. The relatively small value in recent years reflects the coexistence of a low unemployment rate and a long unemployment duration.¹⁴ Next, we formally examine whether these changes in the labor-market environment have had a significant impact

¹⁴In fact, Juhn, Murphy, and Topel (2002) argue that the decrease in the unemployment rate observed in the 1990s is driven almost entirely by the decreased incidence of unemployment. Even though the unemployment duration remained as high as its level in the 1980s, the lower rate of separation dragged the unemployment rate down to very low levels.

on the worker’s job-search behavior.

3 Model

In this section, we construct a McCall (1970)-style search model.¹⁵ Then, we feed the two important changes in the U.S. labor market—the increase in the wage dispersion¹⁶ and the decline in the incidence of unemployment—into the model and examine how the job-search behavior of the unemployed workers change.

3.1 Model setup

We consider a worker who is unemployed and searching for a job. Time is discrete. We assume that there is no borrowing or saving and that the period utility for an unemployed worker is U_s . For an employed worker receiving wage w , the momentary utility is $U_e(w) \equiv \ln(w)$.

An unemployed worker receives one wage offer each period and decides whether or not to accept it. If she accepts it, she works at that wage until she is separated from the job. If she rejects it, the search continues in the subsequent period. Separation

¹⁵There is a large body of empirical work based on this type of search model. See, for example, Wolpin (1987). Rogerson, Shimer, and Wright (2005) demonstrate that this model is a building block of many recent equilibrium search models, such as Mortensen and Pissarides (1994) and Burdett and Mortensen (1998). They also emphasize that this model can be interpreted as the equilibrium of a simple economy.

¹⁶There is a large body of theoretical literature that attempts to explain the existence of wage dispersion among workers of the same characteristics. The most popular approach is to utilize a model of search and matching in a frictional labor market (see, for example, Burdett and Mortensen 1998). We do not attempt to explain the wage dispersion in this paper—we take the wage dispersion as exogenously given.

occurs exogenously with probability $\alpha \in [0, 1)$ every period. If separation occurs, the worker is unemployed for at least one period.

We assume that the wage offer is independently and identically distributed and follows a lognormal distribution

$$\ln(w) \sim \mathcal{N}(m - \sigma^2/2, \sigma^2).$$

Therefore,

$$E[w] = e^m \text{ and } Var[w] = e^{2m}(e^{\sigma^2} - 1).$$

The worker's problem in each period is characterized by the following Bellman equation:

$$V(w) = \max \left\langle U_e(w) + \beta \left\{ (1 - \alpha)V(w) + \alpha \left[U_s + \beta \int V(w') dF(w') \right] \right\}, U_s + \beta \int V(w') dF(w') \right\rangle. \quad (1)$$

where $F(\cdot)$ is the distribution function of the wage offer and $V(w)$ is the value function of the worker with wage offer w . The implicit assumption here is that the worker views the economy as being in the steady state—she considers σ and α to be constant over time when she makes her decisions. (That is, changes in σ and α are unanticipated.)

While this assumption seems innocuous given that changes in σ and α are very slow relative to the length of each unemployment spell, it requires further evaluation. We check the robustness of our results to this assumption in Section 4 by solving a non-stationary version of the model.

The optimization problem in (1) has a simple reservation-wage property: the worker accepts the wage offer if w is above the reservation wage and rejects it if it is below the reservation wage. The reservation wage, \bar{w} , solves¹⁷

$$U_e(\bar{w}) - U_s = \frac{\beta}{1 - \beta(1 - \alpha)} \int_{\bar{w}}^{\infty} [U_e(w') - U_e(\bar{w})] dF(w'). \quad (2)$$

Let $\lambda \equiv F(\bar{w})$ be the probability that an unemployed worker is still unemployed next period.

Given that all unemployed workers are solving the optimization problem presented above, the dynamics of the aggregate unemployment rate, u_t , are governed by

$$u_{t+1} = \alpha(1 - u_t) + \lambda u_t,$$

where the first term on the right-hand side is the number of workers separated at time t , and the second term on the right-hand side is the number of workers who are unemployed at time t and who rejected the time- t job offer.¹⁸ In the steady-state, u_t is constant (call it \bar{u}), and

$$\bar{u} = \frac{\alpha}{1 + \alpha - \lambda}. \quad (3)$$

As is clear from (3), \bar{u} is increasing in both α and λ . We compute the average unemployment duration as $1/(1 - \lambda)$ and the unemployment rate as $\alpha/(1 + \alpha - \lambda)$ in our simulations.

¹⁷See the derivation in Appendix F.

¹⁸See Rogerson, Shimer, and Wright (2005, Section 3.3) for a discussion of these unemployment dynamics.

3.2 Calibration

Our calibration strategy is to set the parameters of the model so that it matches the observations in the initial period in our dataset. Here we use the March CPS dataset, which makes the model's starting year 1971.¹⁹ Then we vary α and σ over time and see how the average unemployment duration changes as these parameters change. As will become clear later, α can be pinned down directly from the data on the incidence of unemployment, while σ has to be calibrated indirectly so that the model solution matches the observed 90% – 10% residual wage inequality.

Equation (2) is scale-free in the sense that, if m is replaced by $m + \mu$, and U_s is replaced by $U_s + \mu$, the reservation wage \bar{w} becomes $\bar{w} \cdot e^\mu$, and λ remains the same. Therefore, we can normalize m by setting $m = 0$. In our benchmark calibration, we assume that mean wages do not change.²⁰

We set the length of one period as one month. Therefore, $\beta = 0.947^{1/12}$.²¹ The other parameters, U_s and σ , are set so that the following two conditions are satisfied.

1. The average unemployment duration, $1/(1 - \lambda)$, matches the unemployment duration in 1971: 12.3 weeks for males and 10.1 weeks for females.
2. The 90% – 10% log wage difference of employed workers in the model matches

¹⁹As will be explained in Section 4, we conduct the same experiment for the Census data and the CPS May/ORG data in Appendix G. For the Census data, the initial year is 1970 and for the CPS May/ORG it is 1973.

²⁰In the robustness analysis in Section 4, we also examine the implication of varying m over time.

²¹The value for the annual discount rate, β , is taken from Cooley and Prescott (1995).

the 90% – 10% log wage differences in 1971: 0.95 for both males and females. Note that when we compute the 90% – 10% log wage difference in the model, we look at the *accepted* wage offers, which requires us to calculate the solution of the model repeatedly until we match the observed 90% – 10% log wage difference in the data.

These two conditions imply $U_s = -5.58$ and $\sigma = 0.734$ for males and $U_s = -4.65$ and $\sigma = 0.680$ for females.

3.3 Qualitative predictions

Before we conduct the quantitative analysis, we summarize the theoretical predictions of the model.²²

Wage dispersion (σ): It is well known that in McCall-style search models with linear utility, a mean-preserving spread in the wage-offer distribution increases the reservation wage. This is because an unemployed worker will tend to wait longer to accept an offer since the option value of a job opportunity (i.e. the opportunity cost of accepting a job) increases with the variance of the wage offer. When the variance increases, the possibilities of receiving a very low wage and a very high wage both increase. An increase in the probability of a very low-wage offer does not affect the value of waiting, since those offers are always rejected anyway. The higher probability of a very high-wage offer, however, increases the value of waiting, since those are the

²²These results are standard, and therefore the proofs are omitted.

offers that workers accept. Therefore, the increase in variance increases the likelihood of a good job opportunity (raises the “option value” of waiting), causing unemployed workers to wait longer in hope of receiving one.²³

Incidence of unemployment (α): It is clear from (2) that \bar{w} is decreasing in α , so a higher α leads to a shorter duration. The intuition is that a worker becomes less selective about a job when the probability that the job will be terminated is high. The increase in the incidence of unemployment has two opposing effects on the unemployment rate: first the increase in α has the direct effect of increasing the unemployment rate since more workers become unemployed; second it has the indirect effect of decreasing unemployment since workers become less selective.

3.4 Results

Now we turn to the quantitative evaluation of the changes in wage dispersion and the incidence of unemployment. We compute the model for each year from 1971 to 2002 by changing the values of σ and α . We change σ so that the 90% – 10% log wage difference matches its value for each year in the data (shown in Figure 3). Similarly we set α to its value in the corresponding year in the data (shown in Figure 4).

Figure 5 presents the average unemployment duration in the data and model for males (right panel) and females (left panel) for different specifications of the model.

²³In our model, the utility function is concave and therefore there is an additional effect in the opposite direction due to consumption smoothing motive. In the calibrated version of the model, however, it turns out that the effect of the “option value” dominates. We thank Jiyeon Oh for pointing this out.




Figure 5: Unemployment duration for male (left) and female (right). “Data” refers to the March CPS data and “Trend” refers to the HP-filtered trend of the data. There are three model results: “Benchmark” (both α and σ are changed), α only, and σ only.

It shows that the benchmark model generates a significant increase in duration for both male and female workers. For males, average unemployment duration increases from 12.4 weeks (from 1971 to 1980) to 14.4 weeks (from 1981 to 2002). For females, it increases from 9.9 weeks (from 1971 to 1980) to 14.0 weeks (from 1981 to 2002).²⁴ Our benchmark model allows both wage dispersion (captured by σ) and incidence of unemployment (captured by α) to vary as they did in the data. In order to assess the importance of each parameter, we consider two alternative calibrations of the model. In particular, the first alternative calibration only varies α and keeps σ constant while the second alternative calibration varies σ and keeps α constant. In the first alternative calibration (only α changes), there is not much increase in unemployment duration

²⁴Recall that in the data, the duration increased by 3.4 weeks for males and 2.8 weeks for females (footnote 6).




Figure 6: Unemployment rate for male (left) and female (right), data (March CPS), model.

until the mid-1980s. Then unemployment duration starts increasing and stabilizes at a higher level in the mid-1990s. For males, the increase in unemployment duration implied by the change in α is not as high as in the data. For females, almost all of the increase in the data is captured by only changing α . As for the second alternative calibration (only σ changes), the model generates a monotonically increasing unemployment duration for both males and females. The effect of σ dominates the effect of α for males while the opposite is true for females.

Figure 6 presents the unemployment rate in the data and the implied unemployment rate \bar{u} in the benchmark model for males (right panel) and females (left panel). The model does not fully capture the cyclical fluctuations in the unemployment rate. In the model, the job finding rate does not vary sufficiently in the business cycle frequency.

As discussed by Hall (2005) and Shimer (2005), fluctuations in the job finding rate is the key to accounting for the cyclical fluctuations in the unemployment rate.

Overall, the model performs reasonably well in accounting for the long-run trends in the unemployment duration and the unemployment rate, while it fails to capture the cyclical properties of these variables. One shortcoming of the current model is that the offer arrival rate is given as constant. Allowing it to vary in cyclical frequency would improve the model performance along this dimension.

4 Robustness

In this section, we examine the robustness of our findings by considering two alternative data sources and two model extensions.

4.1 Different datasets

We consider two alternative datasets to calibrate the change in wage dispersion (σ) and examine the model's implications. We find out that the quantitative implications of the model change noticeably, depending on the data source used to calibrate our model.

The first alternative dataset is the IPUMS Census Samples on 1960, 1970, 1980, 1990, and 2000. The Census dataset has the advantage of having large samples and more detailed occupational controls. However, it is only available every ten years. The Census exhibits a large increase in residual wage inequality in recent years, similar to

the March CPS data.²⁵ Not surprisingly, when we use the 90% – 10% wage difference from the Census data as the measure of wage dispersion, we obtain a significant increase in the unemployment duration. The model implies an increase of 3.4 weeks for males and 8.3 weeks for females from 1970 to 2000.

The second alternative data source is the CPS May/ORG data. As is pointed out by Lemieux (2006), this dataset exhibits a much smaller increase in residual wage inequality. Naturally, when we use this dataset, our model generates a considerably smaller increase in unemployment duration. The model duration for male increases from 9.6 weeks (from 1973 to 1980) to 10.7 weeks (from 1981 to 2002).²⁶

The increase in unemployment duration in the Census data case is still significant, while the increase in the CPS May/ORG case is much smaller. The CPS May/ORG result is anticipated, given that the CPS May/ORG exhibits a much smaller increase in the residual wage dispersion than the other datasets (see Appendix D. This discrepancy is a reflection of the ongoing debate among labor economists about the nature and the size of the recent increase in wage inequality.²⁷ This debate is beyond the scope of this paper, and here we simply report the results from different datasets instead of taking a particular position.

²⁵See Appendix D.

²⁶In the data, the duration increases from 13.7 weeks (from 1973 to 1980) to 17.0 weeks (from 1981 to 2002).

²⁷See Autor, Katz, and Kearney (2005) and Lemieux (2006, 2007) for discussion.

4.2 Two extensions of the model

We consider two extensions of the basic model. Our benchmark model assumes that the economy is always in the steady state, and the worker expects that the economic environment will not change over time. In the first extension, we assume that the worker has perfect foresight about future changes in the economic environment. In this setup, we also examine the effect of changes in the average wage level over time. It turns out that these changes do not affect the main implications of the model. We then consider a model where the unemployed worker can affect the probability of receiving a job offer by exerting search effort. In this extension, the change in unemployment duration is somewhat smaller than in the benchmark model.

4.2.1 Perfect foresight and the effect of wage level

Instead of assuming that the worker treats α and σ as constant, we assume that the worker foresees the future changes in these parameters. In particular, we assume perfect foresight—the worker knows that α and σ evolve as they did in the data in Section 2.²⁸ When we re-calibrate the value of U_s to match the initial duration of unemployment, the behavior of the model duration of unemployment is almost identical to the benchmark model.

In addition, we examine the effect of the changes in the average wage level. We construct the average wage (compensation) level from the National Income and Product

²⁸Within a year, we linearly interpolate the values of α and σ that are used in Section 3.

Accounts (NIPA) following Sullivan (1997).²⁹ Then we change the value of m following this wage series (we assume a perfect foresight on m as well). We find that results are again very similar to the benchmark.

4.2.2 Endogenous search effort

Instead of assuming that the worker receives one job offer per period, we assume that the probability of receiving a job offer, p , is a function of the search effort level, a . Following Hopenhayn and Nicolini (1997), the period utility of the unemployed worker is assumed to be $U_s - a$. We assume that one period is 2 weeks and calibrate the $p(a)$ function so that $p(a) = 1/2$ at the optimal a in the initial period. As in Section 3, the worker expects σ and α to be constant over time. We calibrate this model by using the Census data and find out that the increase in the average unemployment duration is 2.3 weeks from 1970 to 2000 (recall that in the benchmark, the increase is 3.4 weeks). Thus, the effect of α and σ on unemployment duration is somewhat smaller. Since the decrease in α and the increase in σ both increase the value of receiving a wage offer, a increases over time. Consequently, $p(a)$ increases and average unemployment duration increases less relative to the benchmark case.

²⁹See Eckstein and Nagypál (2004) for the comparison between the NIPA compensation measure and the CPS wage measure.



Figure 7: Change in the unemployment duration and the dispersion of wages from 1970 to 2000 for different age and sex groups.

5 Examining disaggregated groups

In this section, we examine the link between unemployment duration and the dispersion of wages at a more disaggregated level. We consider two different types of groups—demographic groups and occupational groups. In Appendix H, we show the cross-sectional relationship between unemployment duration and residual wage inequality. They consistently show a positive relationship, as is predicted by the theory. A more interesting and relevant question for our macroeconomic observation is whether the groups which experienced higher increases in wage dispersion also experienced higher increases in unemployment duration.

First, Figure 7 shows relationship between the increase in unemployment duration from 1970 to 2000 and the increase in wage dispersion from 1970 to 2000, for each




Figure 8: Change in the unemployment duration and the dispersion of wages from the 1970s to the 2000s for different occupation groups.

demographic group. They show a positive relationship (the correlation coefficient is 0.58), which is consistent with our analysis.

Second, Figure 8 shows the relationship between the increase in unemployment duration from 1978 to 2002 and the increase in wage dispersion from 1970 to 2000 for different occupation groups.³⁰ This plot also exhibits a positive correlation, although the relationship is weaker than the case of the demographic groups (the correlation coefficient is 0.29).

6 Conclusion

In this paper, we examined the causes of the increase in the U.S. average unemployment duration in recent years. By quantitatively evaluating a search model, we showed that

³⁰The unemployment duration is computed from the CPS Monthly Basic Studies, as detailed in Appendix H.

both the decrease in the incidence of unemployment and the increase in wage dispersion can cause the average duration to be significantly longer.

Overall, the model performs well in replicating the long-run trends in the unemployment duration and the unemployment rate, while it fails to match the cyclical variations of these variables. In addition, our model does not account for the behavior of unemployment duration in the 1980s. Hall (2005) and Shimer (2005) argue that the key to understanding the cyclical movement of the unemployment rate is the change in the job finding rate. In the popular Diamond-Mortensen-Pissarides matching model, the arrival rate of job offers can vary due to the vacancy posting behavior of firms. We abstract from this aspect, assuming that the arrival rate of job offers are constant. An extension along this line is beyond the scope of this paper, but we believe that it has the potential to better match the data.

Another limitation of our paper is that we take the wage process and the incidence of unemployment as given. Clearly, wages and unemployment (both incidence and duration) are determined simultaneously in the labor market. Our analysis is a first step towards a better understanding of the interaction between them in the context of the U.S. economy in recent years. A more detailed and complete analysis of this interaction is an important future research agenda.

Appendix

A Discussion of related literature

Baumol and Wolff (1998) argue that technological progress can explain the increase in unemployment duration. When technological progress is more rapid, the relative cost of hiring a worker whose training cost is higher (e.g. unskilled and/or older workers) increases. As a consequence, the supply of jobs available to these workers decreases, thereby increasing their job search duration. Our theory does not rule out technological change as a main cause of the longer unemployment duration. In fact, Violante (2002) shows that the recent increase in embodied technological progress can explain a large part of the rise in wage dispersion that is not explained by worker characteristics. Our theory, combined with Violante's, provides another channel through which technological progress might affect unemployment duration.

Valletta (1998) examines how the incidence and duration of unemployment changed during 1967-1998 and shows that there is an upward trend in the duration of unemployment.³¹ He argues that the rising incidence and duration of permanent job loss can account for most of the increase in unemployment duration. His study does not explicitly analyze the underlying economic reasons of this phenomenon. He speculates that a rapid technological change, changing job search strategies, and measurement

³¹Valletta (2005) updates the analysis for 1977-2004 and discusses the measurement issues related to the duration of unemployment in more detail.

issues can be possible explanations.

Abraham and Shimer (2001) focus on demographic changes and changes in the measurement of unemployment duration as the sources of the observed increase in duration. They find that the redesign of the Current Population Survey (CPS) in 1994 explains a half-week of the increase in duration. They then examine two important demographic changes in the composition of the U.S. labor force that are likely to have increased the unemployment duration: the aging of the baby boom cohort and the increase in women's labor force attachment. They conclude that the first factor explains a half-week of the increase in duration. In total, therefore, they attribute one week of the increase to measurement issues and the aging of the baby boom cohort. They attribute the remainder of the increase to the increase in women's unemployment duration. There is no doubt that the increase in women's labor force attachment explains a part of the increase in the average unemployment duration. However, given that a significant increase in average duration can also be observed in samples that consist only of males, a substantial amount of the increase in duration remains unexplained. For example, Juhn, Murphy, and Topel (1991, 2002) point out that for the samples of prime-aged men, unemployment duration has become longer in recent years (see Figure 3 in their 1991 paper and Figure 11 in their 2002 paper).

A recent paper by Machado, Portugal, and Guimaraes (2006) examine the increase in unemployment duration, using data from the Displaced Worker Survey. They argue

that improved job-search technology can lead to an increase in unemployment duration. In particular, if screening technology is better, more-able individuals will face a higher probability of being hired, and less-able individuals will face a lower probability of getting job offers. This mechanism can increase the average duration of unemployment by causing the shorter durations to be shorter and longer durations to be longer.

B More on time series properties

In this section, we further characterize the time-series properties of average unemployment duration.

B.1 Cyclical movement of the unemployment rate and average unemployment duration

Figure 9 depicts the HP-detrended series of average unemployment duration. This shows that, at a high frequency, the two series move together, even in recent years (Baker, 1992).

B.2 Identifying the break more formally

Let D_t be the average duration of unemployment in year t . Across the entire sample period ($t = 1948$ to 2003), the average of D_t is 13.2 weeks. As we saw in Figure 1 in the main text, average D_t clearly increased during the sample period. To identify the timing of the regime shift, we apply the Sup F test by Andrews (1993). Let the mean of the whole sample be $m \equiv [\sum_{t=1948}^{2003} D_t]/T$, where $T = 56$ is the sample size. Let the

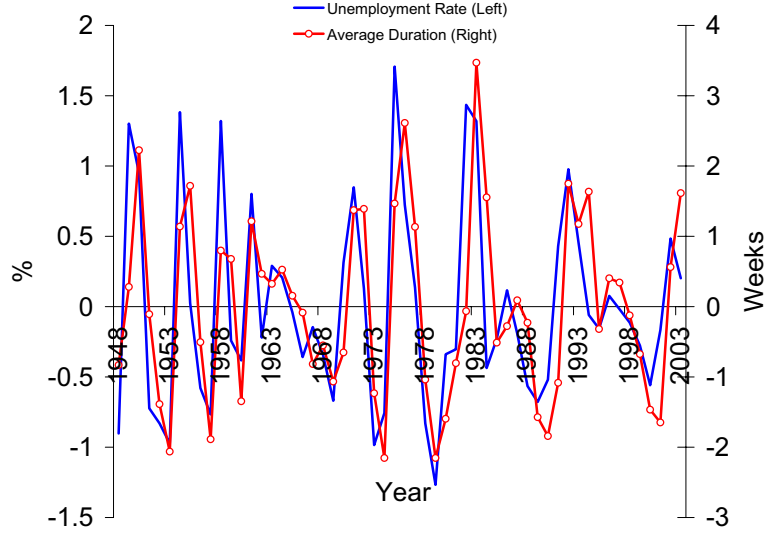


Figure 9: Detrended unemployment rate (left scale) and average unemployment duration (right scale).

Data source: Current Population Survey

sum of the square mean-deviation be $S \equiv \sum_{t=1948}^{2003} (D_t - m)^2$.

We would like to identify the timing of the increase in average D_t . To do this, we divide the sample period into two subperiods and denote the break point by τ . We then calculate the mean of the first subperiod, $m_1 \equiv [\sum_{t=1948}^{\tau} D_t]/T_1$. Here, the sample size is $T_1 = \tau - 1948 + 1$. Let S_1 be the sum of the square mean-deviation for the first subperiod: $S_1 \equiv \sum_{t=1948}^{\tau} (D_t - m_1)^2$. Similarly, for the second subperiod, $m_2 \equiv [\sum_{t=\tau+1}^{2003} D_t]/T_2$, where $T_2 = 2003 - \tau$ and $S_2 \equiv \sum_{t=\tau}^{2003} (D_t - m_2)^2$.




Figure 10: W statistic

We calculate a Wald-type statistic, $W \equiv T(S - S_1 - S_2)/(S_1 + S_2)$, for each possible τ . Following convention (Maddala and Kim 1998, p.395), we consider $\tau = 1956, \dots, 1995$ (disregarding the first and last 15%). A large value for W indicates a structural change.

Figure 10 plots the series of W . The value of W exceeds 25 between 1972 and 1982.³² Therefore, it seems likely that a structural change occurred during this period. W reaches its maximum in 1980.³³ The average of D_t for $t = 1948, \dots, 1980$ (m_1) is 11.5 weeks, while the average of D_t for $t = 1981, \dots, 2003$ (m_2) is 15.5 weeks. Therefore, the difference before and after the change ($m_2 - m_1$) when $\tau = 1980$ is 4.0 weeks. When we

³²The asymptotic critical value for 1% significance (against the null hypothesis of no structural change during the period) is 12.35 (Andrews 1993, Table 1).

³³When we specify the process of D_t as an AR(1) process, the W statistic calculated from the residuals reaches its peak in 1982.

select a τ from the period between 1972 and 1982 the difference in average duration, $m_2 - m_1$, takes a value of between 3.4 weeks and 4.0 weeks.

C Demographic adjustment

In this section, we examine whether the changes in demographic composition can account for the recent increase in average unemployment duration. In particular, we focus on the changes in the age and gender compositions.

First we divide the unemployed population into two gender groups, men and women. Then each group is divided into seven age groups: $A_m = \{ \text{men 16-19, men 20-24, men 25-34, men 35-44, men 45-54, men 55-64, men 65+} \}$ and $A_w = \{ \text{women 16-19, women 20-24, women 25-34, women 35-44, women 45-54, women 55-64, women 65+} \}$.³⁴ Therefore, we have fourteen demographic groups. Let $f_t(i)$ be the fraction of unemployed workers who are in group i at time t , and let $D_t(i)$ be the average unemployment duration for workers who are in group i at time t . Then, by definition, the average duration of unemployment at time t is

$$D_t = \sum_{i \in A_m \cup A_w} f_t(i) D_t(i). \quad (4)$$

Equation (4) implies that if the fraction of the group that experiences longer unemployment spells increases, then the average unemployment duration increases.

Table 1 shows the average unemployment duration for different age and sex groups

³⁴These age categories coincide with the age groups used by the BLS.

Men	16-19	20-24	25-34	35-44	45-54	55-64	65+
1970	6.7	7.9	9.6	10.5	12.0	14.1	17.4
2003	12.0	16.7	17.9	22.5	24.8	26.8	24.8
Women	16-19	20-24	25-34	35-44	45-54	55-64	65+
1970	6.1	7.1	8.2	8.8	10.1	10.2	11.5
2003	11.2	15.1	18.0	20.3	23.2	25.2	21.3

Table 1: Average duration of unemployment for each demographic group (weeks, year 1970 and 2003)

Data source: Current Population Survey

for 1970 and 2003. We first note that the increase in unemployment duration is common to all age and sex groups. In general, older workers experience a substantially longer unemployment duration than younger workers (except for the 65+ age group in 2003). Also, women’s durations tend to be shorter than men’s (except for the 25-34 age group in 2003). These observations suggest that changes in labor force composition might affect the average duration of unemployment if the fraction of demographic groups which experience longer durations increases. The composition of the U.S. labor force has indeed changed substantially from 1970 to 2003. Figure 11 shows the average age of unemployed workers from 1948-2003.³⁵ Since the 1970s, average age has increased, mainly due to the aging of the baby-boom cohort. Figure 12 shows the fraction of women in the unemployment pool. Here we also observe an increase over time. In the following, we explore the effect of these demographic changes on unemployment duration.

³⁵This figure is computed by assigning the middle value to each group and taking the weighted sum. We assigned an age of 70 to the 65+ group.




Figure 11: Average age of unemployed workers

We consider two thought experiments. The first is “What if people in each group behaved the same over time, but the demographic composition of the population changed as it actually did?” Second, “What if the demographic composition remained the same, but people in each group changed their behavior as they actually did?” The first experiment captures the effect of the change in the composition of the population, and the second experiment captures the effect of changing behavior.




Figure 12: The fraction (%) of women in unemployment pool

C.1 First experiment—pure composition effect

We first compute how much of the increase in unemployment duration was caused by demographic change. We examine the *demographic* change in unemployment duration³⁶ by computing

$$D_{t_1, t_0}^d = \sum_{i \in A_m \cup A_w} f_{t_1}(i) D_{t_0}(i)$$

for t_1 from 1948 to 2003. For robustness, we examine two base years (t_0): 1967 and 2003.

³⁶We follow Shimer's (1998) terminology.




Figure 13: D_{t_1, t_0}^d for $t_0 = 1967$ and $t_0 = 2003$

Figure 13 shows D_{t_1, t_0}^d for $t_0 = 1967$ and $t_0 = 2003$. Neither series has an apparent trend, and if we calculate the average of D_{t_1, t_0}^d for $t_1 < 1980$ and separately for $t_1 > 1980$ there is even a slight decline. Even when we take the difference between the smallest D_{t_1, t_0}^d ($t_1 = 1972$ or 1973) and the largest D_{t_1, t_0}^d in recent years ($t_1 = 2003$), it is less than two weeks, which is much smaller than the observed increase in Section 1 and Appendix B.

To see the effect of the adjustment more clearly, we subtract D_{t_1, t_0}^d (where $t_0 = 2003$) from the trend of average duration (shown in Figure 2 in the main text). Figure




Figure 14: The deviation of actual average duration (HP trend) from D_{t_1, t_0}^d ($t_0 = 2003$).

14 shows this series. The structural change is even more apparent than in previous figures—there is a clear shift in the mean between the late 1970s and the early 1980s.

C.2 Second experiment—changing behavior

Next, we compute how much of the increase in unemployment duration would have occurred if the demographics had remained the same. This part (the *genuine* change in unemployment duration) measures the change in average duration that demographic changes do not account for. Here, we hypothetically assume that demographics remained the same from period t_0 to t_1 , and that unemployment duration for each group

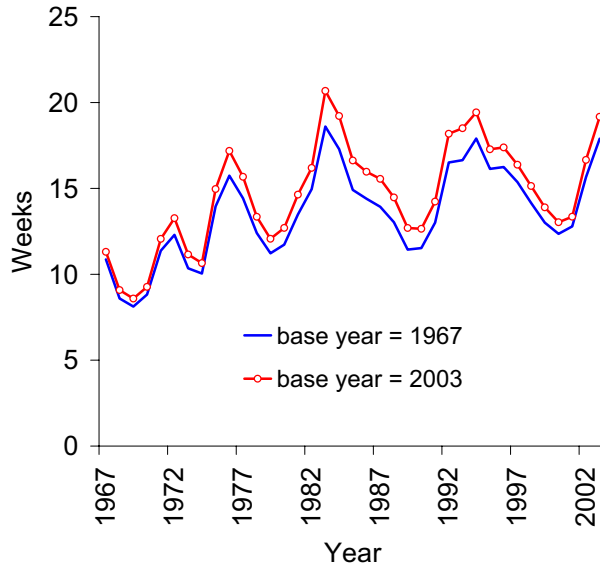


Figure 15: D_{t_1, t_0}^g for $t_0 = 1967$ and $t_0 = 2003$

$(D_t(i))$ followed the same pattern as in the data. Then, the unemployment duration would be

$$D_{t_1, t_0}^g = \sum_{i \in A_m \cup A_w} f_{t_0}(i) D_{t_1}(i).$$

Figure 15 shows D_{t_1, t_0}^g for $t_0 = 1967$ and $t_0 = 2003$.³⁷ They both show a trend that is very similar to that in Figure 1. In fact, the differences between the pre-1980 and post-1981 averages are 3.4 weeks ($t_0 = 1967$) and 3.9 weeks ($t_0 = 2003$).

³⁷The graph starts from 1967 due to data limitations.

D The change in residual wage inequality from Census data and CPS May/ORG

In this section, we repeat our empirical analysis of wage residual wage dispersion by using two alternative data sources. In particular we use data from the IPUMS Census Samples for 1960 to 2000 and the CPS May/ORG.

D.1 Census data: 1960-2000

First we exploit wage data from the IPUMS Census Samples for 1960 to 2000. We use hourly wage as a measure of earnings and calculate the hourly wage by dividing total salary by the product of hours worked per week and total weeks worked.³⁸ Since our goal is to have an accurate measure of residual wage inequality, we try to eliminate all wage variation due to ex-ante differences across individuals in observable characteristics. To this end, we control for the most commonly used human capital and demographic indicators in Mincer-type regressions. Specifically we add controls for education, experience, race and gender for individuals. In addition, unemployed people looking for work are affected by the wage offer distribution that they are facing in the labor market in which they are searching for a job. Most unemployed individuals limit their search to a specific region and occupation. To be consistent with this interpretation, we also control for occupation and geographical location.³⁹ We use Census

³⁸See Appendix E (the next section) for a detailed description of our treatment of the Census data.

³⁹Hornstein, Krusell, and Violante (2006) consider narrower geographical and occupational definitions than we do. This is because they only use the 1990 Census, so they are not constrained by small

regions as a measure of geographic location. All Censuses report occupation categories for the respondents and we use these variables to control for individuals' occupations.

To obtain a measure of residual wage inequality, we run an OLS regression on the logarithm of hourly wage on gender, four race indicators, experience, experience squared, four education indicators, Census region dummies, and occupation categories for each Census year. We then compute the difference between the 90th and 10th percentiles of the residual. As can be seen in Figure 16, residual wage inequality increased from 1960 to 1970, did not change much from 1970 to 1980, increased from 1980 to 1990, and did not change much from 1990 to 2000. We also run a similar regression separately for men and women and compute the 90% – 10% residual wage inequality for males and females separately.⁴⁰ Figure 16 shows that residual wage inequality increased for both male and female workers from 1960 to 2000. The difference between the 90th and 10th percentiles of the residual wage distribution increased from 0.85 to 1.07 for males while for females it increased from 0.74 to 1.03. As Figure 16 shows, male residual wage inequality increased from 1970 to 1980 while female wage inequality declined in the same time period.

sample problems. On the other hand, we analyze how wage dispersion changed from 1960 to 2000 by using five Censuses. The earlier Censuses constrain us to use broader occupation and geographical groups.

⁴⁰The R-squares of the regressions we considered are generally between 0.3 and 0.4, typical for Mincer-type regressions.




Figure 16: 90% – 10% residual wage inequality, the U.S. Census 1960-2000.

D.2 The CPS May/ORG: 1973-2003

There is an ongoing debate in labor economics related to the magnitude, timing, and reasons of the recent increase in wage inequality.⁴¹ In particular, Lemieux (2006, 2007) argues that some widely accepted facts about wage inequality growth are not robust to the choice of data. For example, one disputed issue related to wage inequality is whether or not within-group wage inequality had been growing in the 1970s, as argued by Juhn, Murphy and Pierce (1993). This finding was based on data from the March supplement of the CPS, which was used in our main text. Lemieux (2007) summarizes the debate related to this observation and argues that by using data on hourly wages from the May and Outgoing Rotation Group (ORG) Supplements of the CPS (CPS May/ORG) instead, one can find that within-group wage inequality

⁴¹See for example Autor, Katz, and Kearney, 2005 and Lemieux, 2006, 2007 for a detailed discussion.

did not increase for men and declined for women during the 1970s (see also DiNardo, Fortin, and Lemieux (1996) and Card and Lemieux (1996)). The debate about the increase in wage inequality is not only about the 1970s; there are various other aspects of wage inequality growth that are under disputed: its nature (*episodic* or *gradual*), its reasons (*change in labor force composition* versus *skill-biased technological change*) etc. We believe that a thorough analysis of this debate is beyond the scope of our paper. However, for the purpose of robustness we also present the evidence from the CPS May/ORG. To this end, we are using the data set put together by Autor, Katz, and Kearney (2005). They use May CPS samples for 1973 to 1978 combined with the CPS ORG files for 1979 to 2003. We replicate their Figure 2, Panel B, of Male 90% – 10% residual inequality in our Figure 17.⁴² As Figure 17 shows residual male inequality increase is mostly concentrated to 1980s according to data from the CPS May/ORG.

E Census data description

We use 1960 1% sample, 1970 Form 2 Metro sample, 1980 5% sample, 1990 5% sample, and 2000 5% sample. Sample sizes are 342,296 for 1960, 405,709 for 1970, 2,726,286 for 1980, 3,342,556 for 1990, and 3,590,464 for 2000.

⁴²As explained in Autor, Katz, and Kearney (2005) and Lemieux (2005), both March and May/ORG CPS surveys have limitations in terms of availability and consistency of data on measures of wage. The March CPS data lack a point-in-time wage measure and thereby hourly wages must be computed by dividing annual earnings by the product of weeks worked last year and usual weekly hours last year. The May/ORG samples provide more accurate measures of the hourly wage distribution but cover a shorter time period than the March CPS.




Figure 17: 90% – 10% Residual wage inequality, CPS May/ORG: 1973-2003

We create four education groups that indicate whether an individual has less than a high school degree, has a high school degree, has completed less than four years of college, or has completed four or more years of college. Observations representing individuals who are currently in school, younger than 20 or older than 60, self-employed, or currently in armed forces are not included.

Income is top-coded at 25,000 dollars in 1960, 50,000 dollars in 1970, 75,000 dollars in 1980, 140,000 dollars in 1990 and 175,000 dollars in 2000. We exclude all observations where individuals' incomes are top-coded. Hourly wage is calculated by dividing total salary by the product of hours worked per week and total weeks worked. We use *wkwork1* and *wkwork2* as measure of total weeks worked. For hours worked we use *uhrwork* (usual hours worked per week) or *hrwork2* (hours worked last week) depending on the Census year. Once hourly wages are calculated, we drop observations

representing individuals who were earning an hourly wage below the effective federal minimum wage in the corresponding year. The minimum wage was \$1.00/hr in 1960, \$1.3/hr in 1970, \$2.90/hr in 1980, \$3.35/hr in 1990 and \$5.15/hr in 2000. Education measures are drawn from the variable *educrec*, which records educational attainment in increments. We use *educrec* in the creation of two different variables. We create dummy variables indicating whether an individual has less than a high school degree, has a high school degree, has completed less than four years of college, or has completed four or more years of college. We create a measure of an individual's years of labor market experience with a variable equal to age minus years of education minus six. We use occupation categories *occ99* that are consistent across different Census years.

F Derivation of the equation (2)

First, we establish that the solution to (1) has a reservation-wage property. The second term of the max does not depend on w . Suppose that $V(\cdot)$ is a continuous and nondecreasing function. Since $\lim_{x \rightarrow 0} U_e(x) = -\infty$ and $\lim_{x \rightarrow \infty} U_e(x) = \infty$ (and because of the continuity), there exists at least one value of \bar{w} which satisfies

$$U_e(\bar{w}) + \beta \left\{ (1 - \alpha)V(\bar{w}) + \alpha \left[U_s + \beta \int V(w')dF(w') \right] \right\} = U_s + \beta \int V(w')dF(w').$$

With such a \bar{w} , the optimal choice for the unemployed worker is to accept when $w \geq \bar{w}$ and to reject when $w < \bar{w}$.

In (1), with the supposition that $V(\cdot)$ is continuous and nondecreasing, the right

hand side is also continuous and nondecreasing. It is also straightforward to see that (1), seen as a mapping, is a contraction mapping. Therefore, by the standard argument it can be established that $V(\cdot)$ is in fact continuous and nondecreasing.

When $w \geq \bar{w}$, solving (1) for $V(w)$ yields

$$V(w) = \frac{U_e(w) + \beta\alpha [U_s + \beta \int V(w')dF(w')]}{1 - \beta(1 - \alpha)}. \quad (5)$$

Since this is equal to $U_s + \beta \int V(w')dF(w')$ when $w = \bar{w}$,

$$\frac{U_e(\bar{w}) + \beta\alpha [U_s + \beta \int V(w')dF(w')]}{1 - \beta(1 - \alpha)} = U_s + \beta \int V(w')dF(w'),$$

and therefore

$$\frac{U_e(\bar{w})}{1 - \beta} = U_s + \beta \int V(w')dF(w'). \quad (6)$$

This can be rewritten as

$$\begin{aligned} \frac{U_e(\bar{w})}{1 - \beta} &= U_s + \beta \int_0^{\bar{w}} V(w')dF(w') + \beta \int_{\bar{w}}^{\infty} V(w')dF(w') \\ &= U_s + \beta \int_0^{\bar{w}} \frac{U_e(\bar{w})}{1 - \beta} dF(w') + \beta \int_{\bar{w}}^{\infty} \frac{U_e(w') + \beta\alpha [U_e(\bar{w})/(1 - \beta)]}{1 - \beta(1 - \alpha)} dF(w'), \end{aligned} \quad (7)$$

where the second equality uses the fact that $V(w) = U_s + \beta \int V(w')dF(w')$ when $w \leq \bar{w}$, (5), and (6). The left-hand side can be decomposed as

$$\frac{U_e(\bar{w})}{1 - \beta} = \int_0^{\bar{w}} \frac{U_e(\bar{w})}{1 - \beta} dF(w') + \int_{\bar{w}}^{\infty} \frac{U_e(\bar{w})}{1 - \beta} dF(w').$$

Applying this to the left-hand side of (7), subtracting $\beta \int_{\bar{w}}^{\infty} U_e(\bar{w})dF(w')/(1 - \beta)$ from both sides, and rearranging yields (2). Since the left-hand side of (2) is strictly increasing in \bar{w} and the right-hand side of (2) is nonincreasing in \bar{w} , the reservation wage \bar{w} is unique.

G Robustness

In this section, we provide a detailed description of the robustness experiments described in Section 4.

G.1 Different datasets

G.1.1 Census data

Here we use the residual wage dispersion obtained from the Census data. First, we calibrate U_s and σ so that

1. The average unemployment duration, $1/(1 - \lambda)$, matches the unemployment duration in 1970: 12.1 weeks for males and 9.7 weeks for females.
2. The 90% – 10% log wage difference of employed workers matches the 90% – 10% log wage differences in 1970: 0.90 for males and 0.84 for females.

These two conditions imply $U_s = -5.35$ and $\sigma = 0.685$ for males and $u_s = -4.20$ and $\sigma = 0.596$ for females. Then we set α to its value in the corresponding Census year by following Figure 4. Similarly we change σ such that the 90% – 10% log wage difference matches its value in each Census year that we presented in Figure 16.

Table 2 shows the unemployment duration and the unemployment rate from Census the data and the model. The model predicts lower levels of duration of unemployment and higher rates of unemployment for both males and females in 1980 than in 1970. That is because 1980 is characterized by high separation rates and low wage dispersion.

	Male						Female					
	Parameters		Duration		URate		Parameters		Duration		URate	
Year	α	σ	Data	Model	Data	Model	α	σ	Data	Model	Data	Model
1970	0.023	0.685	12.1	12.1	5.3%	6.5%	0.038	0.596	9.7	9.7	6.9%	8.5%
1980	0.030	0.685	18.0	10.4	8.1%	7.3%	0.040	0.570	13.1	9.2	8.2%	8.4%
1990	0.025	0.840	17.6	13.1	6.6%	7.6%	0.028	0.820	14.2	14.3	6.2%	9.1%
2000	0.020	0.880	17.1	15.5	5.3%	7.2%	0.020	0.880	16.0	18.0	5.1%	8.3%

Table 2: Unemployment duration (in weeks) and unemployment rate (URate) from Census the data and the benchmark model.

High incidence of unemployment lowers unemployment duration and increases the unemployment rate. The model predicts higher durations of unemployment for 1990 and 2000 for both males and females than in 1970. The effect of declining separations and rising wage inequality both work in the same direction. However, the effect on the unemployment rate is opposite. Declining separations act to lower the unemployment rate, while increasing dispersion acts to increase it. Compared to the March CPS case in the main text, the increase is somewhat lower for males.

We also perform two additional experiments to assess the relative quantitative importance of the change in wage dispersion (σ) and change in the incidence of unemployment (α). Table 3 presents the summary statistics from the model when we only change α and keep σ constant. When the incidence of unemployment was high in the 1980s, duration of unemployment was lower. Conversely, when α was low in 2000, duration of unemployment was higher. The model implies that for a given job offer arrival rate, the duration of unemployment is lower if the probability that the job

Year	Male				Female			
	Duration		URate		Duration		URate	
	Data	Model	Data	Model	Data	Model	Data	Model
1970	12.1	12.1	5.3%	6.5%	9.7	9.7	6.9%	8.5%
1980	18.0	10.5	8.1%	7.3%	13.1	9.5	8.2%	8.6%
1990	17.6	11.6	6.6%	6.8%	14.2	11.6	6.2%	7.5%
2000	17.1	13.1	5.3%	6.1%	16.0	13.9	5.1%	6.5%

Table 3: Duration of unemployment (in weeks) and unemployment rate (Urate) from the Census data and model 2 (change α only).

Year	Male				Female			
	Duration		URate		Duration		URate	
	Data	Model	Data	Model	Data	Model	Data	Model
1970	12.1	12.1	5.3%	6.5%	9.7	9.7	6.9%	8.5%
1980	18.0	12.5	8.1%	6.7%	13.1	9.5	8.2%	8.3%
1990	17.6	14.1	6.6%	7.5%	14.2	11.4	6.2%	9.8%
2000	17.1	14.1	5.3%	7.5%	16.0	11.6	5.1%	9.9%

Table 4: Unemployment duration (in weeks) and unemployment rate (Urate) from the Census data and model 3 (change σ only).

will be terminated is higher. As for the unemployment rate, the model generates the decline in the unemployment rate in 2000 even though the magnitude of the decline is smaller than in the data.

Table 4 presents the summary statistics from the model when we only change σ and keep α constant. Unemployment duration increases as σ increases. Comparison of Tables 3 and 4 show that the effect of changing σ is higher than the effect of changing α for males while the opposite is true for females.

G.1.2 The CPS May/ORG calibration

In this subsection we repeat the experiment using the residual wage inequality measure that we obtained from the CPS May/ORG data. We examine only male workers. First we calibrate U_s and σ so that

1. The average unemployment duration, $1/(1 - \lambda)$, matches the unemployment duration in 1973 which is 11.2 weeks.
2. The 90% – 10% log wage difference of employed workers matches the 90% – 10% log wage differences in 1973 which is 0.99.

These two conditions imply $U_s = -7.5$ and $\sigma = 0.764$. Then we set α to its value in the corresponding year by following Figure 4. Similarly we change σ such that the 90% – 10% log wage difference matches its value for each year that we presented in Figure 17. We compute the model for each year from 1973 to 2003.

Figure 18 shows the unemployment duration and the unemployment rate implied by the model. In the CPS May/ORG calibration, the implied increase in the unemployment duration is more modest than the Census and March CPS calibrations.

G.2 Alternative model specifications

In Section 4, we considered a stationary version of the McCall model. Specifically we calibrated the deterministic steady state of the McCall search model to each year in our calibration exercises. This approach assumes that the changes in wage dispersion

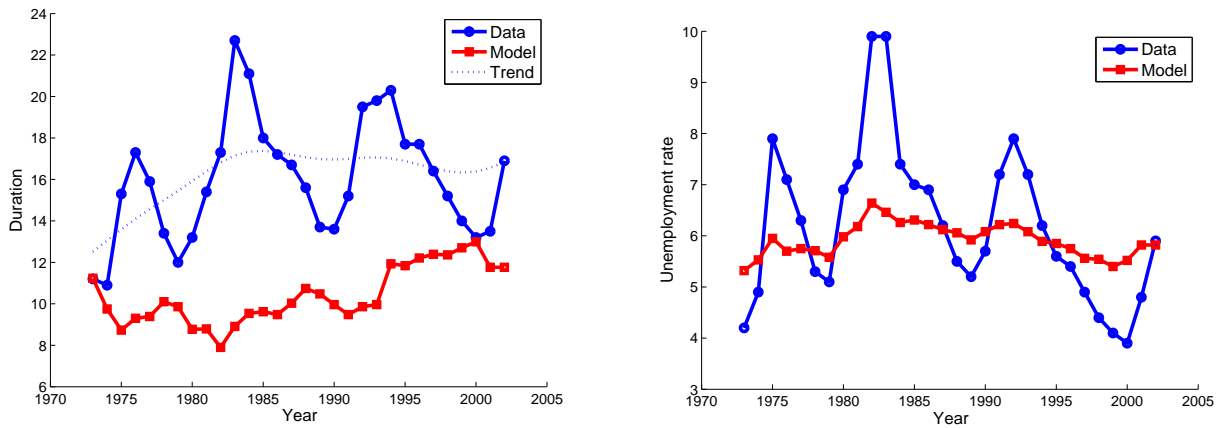


Figure 18: Unemployment duration (left) and unemployment rate (right) for males, data (CPS May/ORG), model.

and incidence of unemployment are unanticipated. In this section, we examine the robustness of this assumption by solving a non-stationary version of the model. After we formulate the non-stationary version of the model, we feed in the series of α and σ used in Section 2 and let the agent have perfect foresight about the future values of α and σ . In this non-stationary version of the model, we also allow for the mean wages to change as they did in the data in order to examine the effect of the change in the level of wages on search behavior of unemployed agents.

As another extension, we consider a version of the model with endogenous search intensity. In the benchmark model, the frequency of receiving a job offer is exogenous and is not affected by changes in the economic conditions. However, search effort can be affected by changes in separation rates or wage dispersion since these parameters

affect the return from receiving a job offer. To examine the quantitative importance of the endogenous search effort margin, we set up a version of the McCall model with endogenous search effort.

G.2.1 Perfect foresight and the change in average wage

Here we set up a non-stationary version of the model in Section 3. The setup is similar except that the value of α and the wage distribution $F(w)$ is allowed to change over time in the dynamic programming problem. We solve the model backwards. Since α changes over time, we denote it as α_t with the time subscript t . The distribution of the offered wages, in particular the mean and the variance of them, changes over time, so we denote the distribution function as $F^t(w)$. We assume that the worker has perfect foresight.

We assume that the parameter stays the same after the final period of our sample T . Thus, at time T , the agent solves the following problem:

$$V_T(w) = \max \left\langle \begin{array}{l} U_e(w) + \beta \{ (1 - \alpha_T) V_T(w) + \alpha_T [U_s + \beta \int V_T(w') dF^T(w')] \} , \\ U_s + \beta \int V_T(w') dF^T(w') \end{array} \right\rangle .$$

This is the stationary dynamic programming problem and can be solved similarly to Section 3. (In practice, we solve this optimization problem by value function iteration.)

We first find the solution to this problem.

Then, given $V_T(w)$, we solve the problem backwards. At time $T - 1$, the problem

becomes

$$V_{T-1}(w) = \max \left\langle \begin{array}{l} U_e(w) + \beta \{ (1 - \alpha_T)V_T(w) + \alpha_T [U_s + \beta \int V_T(w')dF^T(w')] \} , \\ U_s + \beta \int V_T(w')dF^T(w') \end{array} \right\rangle .$$

In periods $t = T - 2, \dots, 1$, the problem can be written as

$$V_t(w) = \max \left\langle \begin{array}{l} U_e(w) + \beta \{ (1 - \alpha_{t+1})V_{t+1}(w) + \alpha_{t+1} [U_s + \beta \int V_{t+2}(w')dF^{t+2}(w')] \} , \\ U_s + \beta \int V_{t+1}(w')dF^{t+1}(w') \end{array} \right\rangle .$$

Let λ_t be the probability that the job offer is rejected by the agent at time t .

Then the implied (steady-state) average unemployment duration can be calculated as $1/(1 - \lambda_t)$, and the implied (steady-state) unemployment rate is calculated by $u_t = \alpha_t/(1 - \alpha_t + \lambda_t)$.

We compute the real wages by following the methodology suggested by Sullivan (1997). The nominal wage measure underlying real hourly compensation is derived from total business sector compensation figures from the National Income and Product Accounts (NIPA) and the BLS series on hours worked by all persons in the business sector.⁴³ The resulting real hourly compensation is shown in the left panel of Figure 19.

We solve for two different versions of the non-stationary model. The first version assumes that the average wage is constant and only α and σ change as in the March CPS calibration. The second version assumes that in addition to α and σ , m also changes over time. In particular, we calibrate the path of m so that the average wage

⁴³The deflator for the real hourly compensation series is the CPIU-X1, which eliminates upward bias in the CPIU.

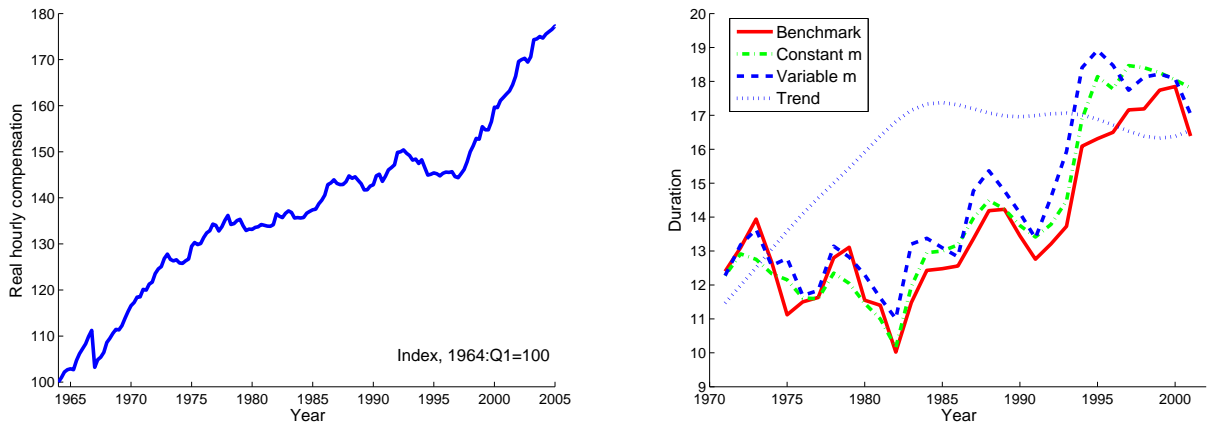


Figure 19: Real hourly compensation, 1964-2005 (left panel) and comparison of the benchmark model with the non-stationary model (right panel).

(e^m) matches Figure 19. Within one year, we linearly interpolate the parameter values.

As the right panel of Figure 19 shows, the two versions of the non-stationary model behave very similarly to the benchmark model. The divergence between the benchmark model and the non-stationary model is higher when the wage level increases steeply in the second half of the 1990s.

G.2.2 Endogenous search effort

We extend the benchmark model to allow for an endogenous job offer probability which depends on agents' job-search effort. We follow Hopenhayn and Nicolini (1997) and assume that the probability of receiving an offer, p , is a function of the search effort, a which we denote as $p(a)$. Following Hopenhayn and Nicolini, we assume that a amount of search effort generates the same amount of utility loss. Thus the Bellman equation

of an employed worker is:

$$V(w) = \max \langle U_e(w) + \beta [(1 - \alpha)V(w) + \alpha V_s], V_s \rangle, \quad (8)$$

where V_s is the value of an unemployed worker (“searcher”) and satisfies

$$V_s = \max_a \left\{ U_s - a + \beta \left[p(a) \int V(w') dF(w') + (1 - p(a)) V_s \right] \right\}. \quad (9)$$

Since V_s is a number, in (8), it is straightforward to show that $V(w)$ is increasing and continuous in w . Therefore, the decision in (8) has a threshold property: there is a \bar{w} where

$$U_e(w) + \beta [(1 - \alpha)V(w) + \alpha V_s] \geq V_s$$

if and only if $w \geq \bar{w}$. The formula for determining the reservation wage \bar{w} (for the optimally chosen search effort level a^*) is:

$$U_e(\bar{w}) - (U_s - a^*) = \frac{\beta p(a^*)}{1 - \beta(1 - \alpha)} \int_{\bar{w}}^{\infty} [U_e(w') - U_e(\bar{w})] dF(w'). \quad (10)$$

To derive this, first we solve for V_s in (9):

$$V_s = \tilde{U}_s + \tilde{\beta} \int V(w') dF(w'),$$

where $\tilde{U}_s \equiv (U_s - a^*) / (1 - \beta(1 - p(a^*)))$ and $\tilde{\beta} = \beta p(a^*) / (1 - \beta(1 - p(a^*)))$. Then we can apply the same steps as before.

The optimality condition for a^* is

$$p'(a^*) \left(\int V(w') dF(w') - V_s \right) = 1. \quad (11)$$

From (8), when $w < \bar{w}$, $v(w) = v_s$. Thus, (11) can be rewritten as

$$p'(a^*) \int_{\bar{w}}^{\infty} (V(w') - V_s) dF(w') = 1. \quad (12)$$

From (8), when $w \geq \bar{w}$,

$$V(w) = \frac{U_e(w) + \beta\alpha V_s}{1 - \beta(1 - \alpha)}.$$

Since $V(w) = V_s$ with $w = \bar{w}$,

$$\frac{U_e(\bar{w}) + \beta\alpha V_s}{1 - \beta(1 - \alpha)} = V_s$$

holds, therefore

$$V_s = \frac{U_e(\bar{w})}{1 - \beta}.$$

Using these, (13) can be rewritten as:

$$p'(a^*) \int_{\bar{w}}^{\infty} \frac{U_e(w') - U_e(w)}{1 - \beta(1 - \alpha)} dF(w') = 1. \quad (13)$$

Using (10),

$$\frac{p'(a^*)}{p(a^*)} (U_e(\bar{w}) - (U_s - a^*)) = 1. \quad (14)$$

Therefore, equations (10) and (14) are two equations to solve for two numbers, \bar{w} , and a^* . We follow Hopenhayn and Nicolini (1997) and specify

$$p(a) = 1 - e^{-ra}.$$

Thus, we have one additional parameter, r , in this model. To be consistent with the benchmark model, we set one period as 2 weeks and set r so that $p(a^*) = 1/2$ in the benchmark. Then, on average, the worker receives one job offer every 4 weeks.

	Duration		
Year	Data	Benchmark	Endogenous a
1970	12.1	12.1	12.1
1980	18.0	10.4	10.5
1990	17.6	13.1	12.7
2000	17.1	15.5	14.4

Table 5: Comparison of the benchmark model with the endogenous effort model.

We calibrate the endogenous job-effort model by following the Census calibration for male workers. As Table 5 shows, the duration of unemployment increases relatively less in the model with endogenous effort, even though the increase is still substantial. To understand the reaction of the job search effort to changes in parameters, we perform two counterfactual experiments. In the first, we let α change as it did in the data, keeping σ constant at its 1970 value. In the second, we let σ change as it did in the data, keeping α constant at its 1970 value. The results are presented in Table 6. Unemployed workers' job search effort increases as α decreases. The return from a job is higher when jobs are less likely to be terminated, therefore workers exert a higher effort to obtain a job offer. Similarly, the return to receiving a job offer is higher when the wage dispersion is larger, and the workers' job-search efforts increase as a response. Since both the decrease in α and the increase in σ increase the effort of the unemployed, the job-finding rate is higher under the endogenous effort model. Consequently, the increase in the unemployment duration is milder in the endogenous effort model.

	Parameters		Effort (a)		
Year	α	σ	Benchmark	α only	σ only
1970	0.023	0.826	1.784	1.784	1.784
1980	0.030	0.835	1.701	1.709	1.789
1990	0.025	0.990	1.822	1.772	1.850
2000	0.020	1.020	1.903	1.840	1.859

Table 6: Effect of the endogenous effort, comparison of the models.

H Disaggregated groups

H.1 Age and sex groups

In Appendix C, we examined unemployment duration for different demographic groups. We saw that, in general, older workers experience a substantially longer unemployment duration than younger workers, and women’s durations tend to be shorter than men’s. Our model predicts that there is a close association between unemployment duration and wage dispersion. Thus, it is natural to expect that the groups with higher wage dispersion experience longer unemployment durations.

To examine this prediction, we first compute the within-group wage dispersion for different age and sex groups. To obtain a good measure of wage dispersion for observationally equivalent workers within each age/sex group, we regress log hourly wage on seven race indicators, four education indicators, Census region dummies, and occupation group dummies, running one regression for 1970 and another for 2000. We then compute the statistics regarding the residuals by age group and gender. We

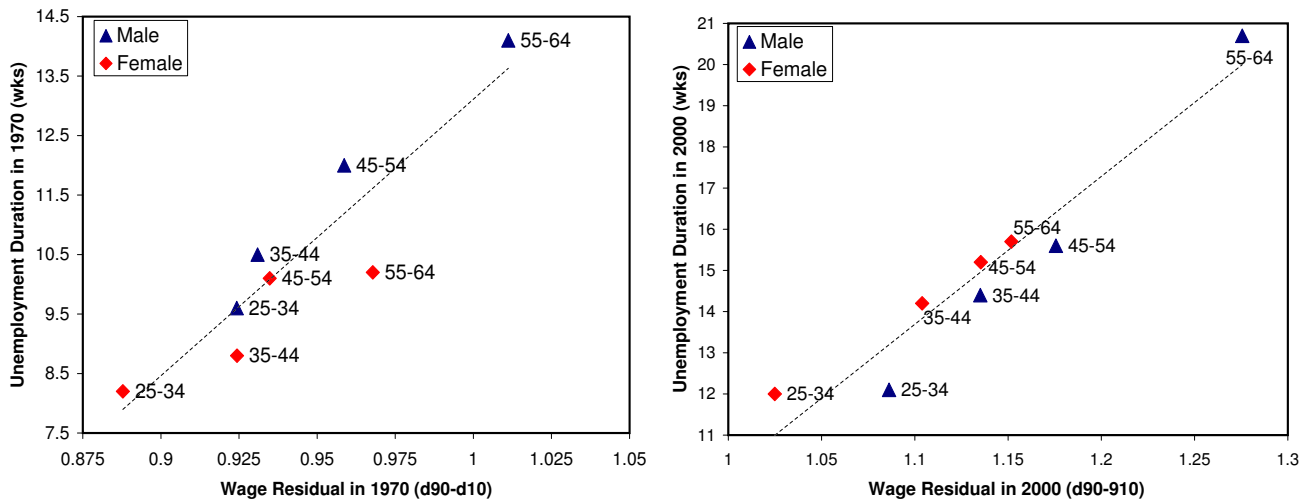


Figure 20: Left: Duration of unemployment in 1970 and dispersion of wages from the 1970 Census for different age and sex groups. Right: Duration of unemployment in 2000 and dispersion of wages from the 2000 Census for different age and sex groups.

define the age groups as follows: 25-34, 35-44, 45-54, and 55-64.⁴⁴ We then plot the difference between the 90th and 10th percentiles of the residual wage distribution for each age/sex group and the unemployment duration for these groups. Figure 20 shows the relationship between unemployment duration and wage dispersion for 1970 and 2000.⁴⁵ Demographic groups that face higher wage dispersion also experience higher durations of unemployment (the correlation coefficient is 0.92 for 1970 and 0.96 for 2000). In particular, female workers and younger workers face lower wage dispersion and lower unemployment durations.

⁴⁴We focus on prime-age workers, since the labor supply behavior of younger and older workers are more influenced by other factors.

⁴⁵The wage dispersion calculations are made by using the 1970 and 2000 Censuses. The unemployment duration data for 1970 and 2000 are taken from the BLS.

H.2 Occupation groups

Here, we regress log hourly wage on seven race indicators, two sex dummies, experience, experience squared, four education indicators, and Census region dummies. We then create 32 occupation groups that are consistent with both the 1970 and 2000 Censuses and categorize each individual occupation into one of the 32 groups.⁴⁶ We group the residuals by the 32 occupation groups.

We then examine unemployment duration for the same occupation groups. Since the BLS does not report unemployment duration by occupation categories, we compute unemployment duration from the CPS Monthly Basic studies for the 12 months of 1978 and of 2002.⁴⁷ Due to data limitations, we cannot compute unemployment duration prior to 1976. We choose 1978 and 2002 because the average unemployment duration in 1978 is the same as pre-1982 and the average unemployment duration for 2002 is the same as post-1982. We use the CPS to examine average within-spell unemployment duration pooled over the entire year's sample organized by occupation groups.

Figure 21 shows the relationship between unemployment duration and wage dispersion for the 1970s and 2000s. Occupation groups that face higher wage dispersion also experience longer durations of unemployment both in the 1970s and 2000s (the correlation coefficient is 0.61 for 1970 and 0.35 for 2000).

⁴⁶We create these broader occupation groups so that individuals who search in similar labor markets are in the same occupation group.

⁴⁷See Appendix I (the next section) for the detailed description of this data analysis.

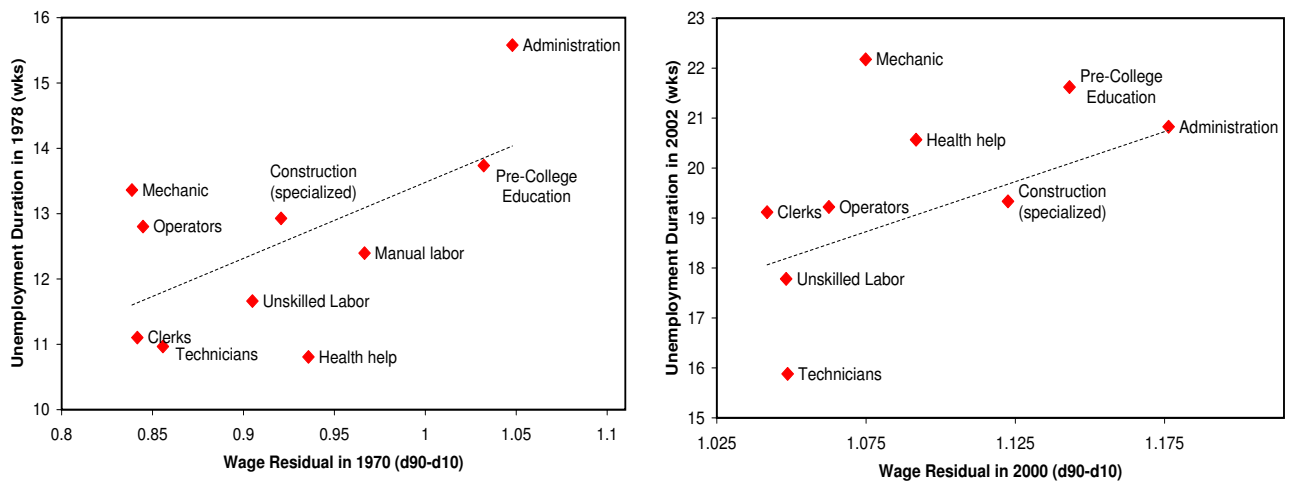


Figure 21: Left: Duration of unemployment in 1970 and dispersion of wages from the 1970 Census for different occupation groups. Right: Duration of unemployment in 2002 and dispersion of wages from the 2000 Census for different occupation groups.

I CPS Data Description

We compute the duration of unemployment from the CPS Monthly Basic studies, for the 12 months of 1978 and of 2002. The measure of unemployment duration is *wksun*. Since we pool observations across consecutive months, single individuals might contribute multiple unemployment values to the mean due to the nature of the CPS rotation scheme. We assume that the observation of an individual for a given month corresponds to a random draw of a representative agent from some stationary distribution, and that our pooling technique is valid. There are minor differences between how the duration of unemployment is recorded in 1978 and in 2002. In 1978, an individual is assigned a positive value of *wksun* if he or she is unemployed and

looking for work ($mlr = 3$). The universe for *wksun* is slightly more narrowly defined in 2002; an individual must be not only unemployed and looking for work, but also satisfy various other criteria, including whether they were able to work in the week previous to the survey (*b4wk*) and whether or not they have held a job previously (measured by *b4wk* and *whenlj*, and satisfying the condition that $1 \leq whenlj \leq 3$). Thus, individuals in the 1978 sample who have never held a job previous to the start of their job search and who would not have been able to hold a job in the week prior to the survey have no counterparts in the 2002 survey, so this might introduce some small distortion into the comparison of the mean computations.

References

- [1] Abraham, K.G., Shimer, R., 2001. Changes in unemployment duration and labor force attachment. *The Roaring Nineties: Can Full Employment Be Sustained?* In: Krueger, A. B., Solow, R.M. (Eds.). New York: Russell Sage Foundation, 367-420.
- [2] Andrews, D.W.K., 1993. Tests for parameter instability and structural change with unknown change point. *Econometrica* 61, 821-856.
- [3] Autor, D.H., Katz, L.F., Kearney, M.S., 2005. Trends in U.S. wage inequality: re-assessing the revisionists. NBER Working Paper No. 11627.
- [4] Baicker, K., Goldin, C. Katz, L.F., 1998. A distinctive system: origins and impact of U.S. unemployment compensation. In: Bordo, M.D., Goldin, C., White, E.N. (Eds.). *The Defining Moment: the Great Depression and the American Economy in the Twentieth Century*. Chicago: University of Chicago Press, 227-263.
- [5] Baker, M., 1992. Unemployment duration: compositional effects and cyclical variability. *American Economic Review* 82, 313-321.
- [6] Baumol, W.J., Wolff, E.N., 1998. Speed of technical progress and length of the average interjob period. Jarome Levy Economics Institute Working Paper No. 237, May.

- [7] Burdett, K., Mortensen, D.T., 1998. Wage differentials, employer size, and unemployment. *International Economic Review* 39, 257-273.
- [8] Burtless, G., 1983. Why is insured unemployment so low? *Brookings Papers on Economic Activity* 1, 225-249.
- [9] Card, D., Lemieux, T., 1996. Wage dispersion, returns to skill, and black-white wage differentials. *Journal of Econometrics* 74, 319-361.
- [10] Cooley, T.F., Prescott, E.C., 1995. Economic growth and business cycles. In: Cooley, T.F. (Ed.). *Frontiers of Business Cycle Research*, Princeton: Princeton University Press, 1-38.
- [11] DiNardo, J., Fortin, N.M., Lemieux, T., 1996. Labor market institutions and the distribution of wages, 1973-1992: a semiparametric approach. *Econometrica* 64, 1001-1046.
- [12] Eckstein, Z., Nagypál, É., 2004. The evolution of U.S. earnings inequality: 1961-2002. *Federal Reserve Bank of Minneapolis Quarterly Review* 28, 10-29.
- [13] Ehrenberg, R., Oaxaca, R.L., 1976. Unemployment insurance, duration of unemployment and subsequent wage gain. *American Economic Review* 66, 754-766.
- [14] Hall, R.E., 2005. Job loss, job-finding, and unemployment in the U.S. economy over the past fifty years. *NBER Macroeconomics Annual*, 101-137.

- [15] Hopenhayn, H.A., Nicolini, J., 1997. Optimal unemployment insurance. *Journal of Political Economy* 105, 412-438.
- [16] Hornstein, A., Krusell, P., Violante, G.L., 2006. Frictional wage dispersion in search models: A quantitative assessment. Federal Reserve Bank of Richmond, Princeton University, and New York University, mimeo.
- [17] Juhn, C., Murphy, K.M., Pierce, B., 1993. Wage inequality and the rise in returns to skill. *Journal of Political Economy* 101, 410-442.
- [18] Juhn, C., Murphy, K.M., Topel, R.H., 1991. Why has the natural rate of unemployment increased over time? *Brookings Papers on Economic Activity* 2, 75-142.
- [19] Juhn, C., Murphy, K.M., Topel, R.H., 2002. Current unemployment, historically contemplated. *Brookings Papers on Economic Activity* 1, 79-116.
- [20] Katz, L.F., Murphy, K.M., 1992. Changes in relative wages, 1963-1987: supply and demand factors. *Quarterly Journal of Economics* 107, 35-78.
- [21] Lemieux, T., 2006. Increasing residual wage inequality: composition effects, noisy data, or rising demand for skill? *American Economic Review* 96, 461-498.
- [22] Lemieux, T., 2007. The changing nature of wage inequality. NBER Working Paper No. 13523.

- [23] Machado, J.A.F., Portugal, P., Guimaraes, J., 2006. U.S. unemployment duration: has long become longer or short become shorter? IZA Discussion Paper No. 2174.
- [24] Maddala, G.S., Kim, I., 1998. Unit Roots, Cointegration, and Structural Change. Cambridge: Cambridge University Press.
- [25] McCall, J.J., 1970. Economics of information and job search. Quarterly Journal of Economics 84, 113-126.
- [26] Mortensen, D.T., 2003. Wage Dispersion: Why Are Similar Workers Paid Differently? Cambridge: MIT Press.
- [27] Mortensen, D.T., Pissarides, C.A., 1994. Job creation and job destruction in the theory of unemployment. Review of Economic Studies 61, 397-415.
- [28] Ravn, M., Uhlig, H., 2002. On adjusting the HP-Filter for the frequency of observations. Review of Economics and Statistics 84, 371-376.
- [29] Rogerson, R., Shimer, R., Wright, R., 2005. Search-theoretic models of the labor market: a survey. Journal of Economic Literature 43, 959-988.
- [30] Shimer, R., 1998. Why is the U.S. unemployment rate so much lower? In: Bernanke, B.S., Rotemberg, J.J. (Eds.). NBER Macroeconomics Annual, 1998. Cambridge: MIT Press, 11-61.

- [31] Shimer, R., 2005. The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review* 95, 25-49.
- [32] Sullivan, D., 1997. Trends in real wage growth. *Chicago Fed Letter* 1997, Number 115.
- [33] Valletta, R.G., 1998. Changes in the structure and duration of U.S. unemployment, 1967-1998. *Federal Reserve Bank of San Francisco Economic Review* 1998, Number 3.
- [34] Valletta, R.G., 2005. Rising unemployment duration in the United States: causes and consequences. *Federal Reserve Bank of San Francisco*, mimeo.
- [35] Violante, G.L., 2002. Technological acceleration, skill transferability, and the rise in residual inequality. *Quarterly Journal of Economics* 117, 297-338.
- [36] Wolpin, K.I., 1987. Estimating a structural search model: the transition from school to work. *Econometrica* 55, 801-817.
- [37] Yellen, J.L., 1991. Comments and discussion. *Brookings Papers on Economic Activity* 2, 127-133.