Volume Title: NBER Macroeconomics Annual 1998, volume 13

Volume Author/Editor: Ben S. Bernanke and Julio Rotemberg, editors

Volume Publisher: MIT Press

Volume ISBN: 0-262-52271-3

Volume URL: http://www.nber.org/books/bern99-1

Publication Date: January 1999

Chapter Title: Why is the U.S. Unemployment Rate So Much Lower?

Chapter Author: Robert Shimer

Chapter URL: http://www.nber.org/chapters/c11245

Chapter pages in book: (p. 11 - 74)

Robert Shimer
PRINCETON UNIVERSITY

Why Is the U.S. Unemployment Rate So Much Lower?

In a well-known paper in one of the inaugural issues of the Brookings Papers, *Robert Hall posed the question, "Why is the Unemployment Rate So High at Full Employment?" (Hall, 1970). Hall, writing in the context of the 3.5% unemployment rate that prevailed in 1969, answered his question by explaining that the full-employment rate was so high because of the normal turnover that is inevitable in a dynamic economy. . . . Today [in 1986], four years into an economic recovery, the unemployment rate hovers around 7%. Over the past decade, it has averaged 7.6% and never fallen below 5.8%. . . . While some of the difference between recent and past levels of unemployment has resulted from cyclical developments, it is clear that a substantial increase in the normal or natural rate of unemployment has taken place. (Summers, 1986)*

1. *Introduction*

In November 1997, the U.S. unemployment rate fell to 4.6%, the lowest level since 1973. This monthly report was not a fluke; the aggregate unemployment has remained below 4.7% for six months, for the first

The title is intended to recall Hall (1970) ("Why is the Unemployment Rate So High at Full Employment?") and Summers (1986) ("Why is the Unemployment Rate So Very High Near Full Employment?"). The wording of this title reflects the recent decline in the unemployment rate in the United States. Explaining the cross-sectional behavior of unemployment, e.g. why the U.S. unemployment rate is so much lower than Europe's, is beyond the scope of this paper.

I have benefited from discussions with and comments and suggestions by Daron Acemoglu, Ben Bernanke, Alan Krueger, Greg Mankiw, Richard Rogerson, Julio Rotemberg, Martin Gonzalez Rozada, Robert Topel, Paul Willen, Mike Woodford, and seminar participants at Chicago, MIT, NYU, and the NBER Macroeconomics Annual 1998 Conference. Thanks also to Steve Haugen for help with some of the data, and to Mark Watson for providing me with Staiger, Stock, and Watson's (1997) series for the NAIRU. Financial support from the National Science Foundation Grant SBR-9709881 is gratefully acknowledged.

Figure 1 THE AGGREGATE UNEMPLOYMENT RATE IS AT THE LOWEST SUSTAINED LEVEL SINCE 1970, PART OF A DOWNWARD TREND SINCE THE EARLY 1980S

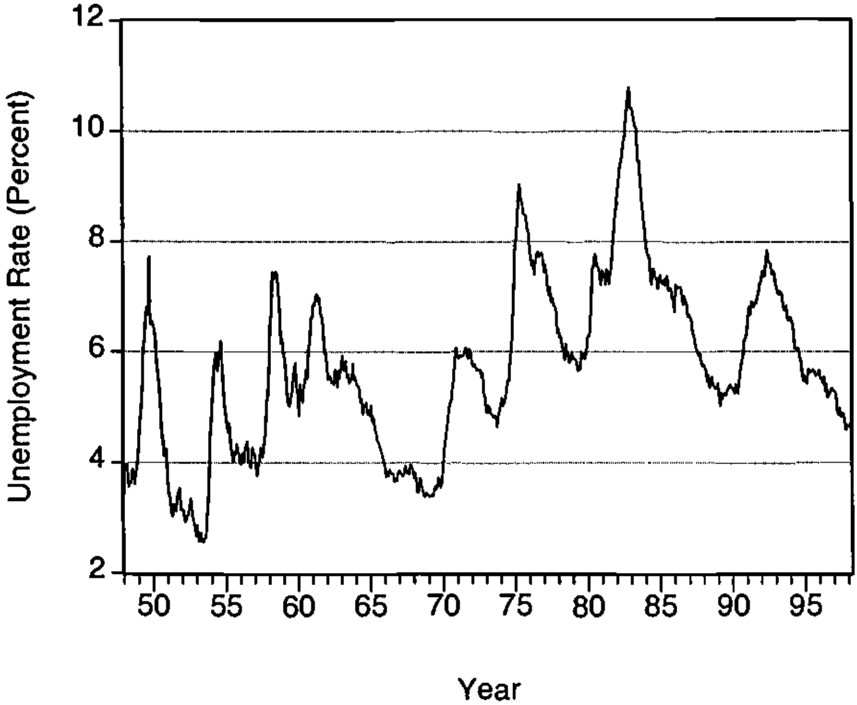

Computed from seasonally adjusted labor-force series based on the Current Population Survey.

time since 1970. Figure 1 documents the decrease in unemployment during the past 19 years, following decades of secular increases. The unemployment rate fell by 63 basis points (hundredths of a percentage point), from 5.66% just before the second oil shock in 1979 to 5.03% at the end of the long expansion in 1989. It has fallen by an additional 43 basis points since then. This prompted Federal Reserve Chairman Alan Greenspan to ask in his semiannual testimony to Congress whether there has been a structural change in the U.S. economy:

> *We do not know, nor do I suspect can anyone know, whether current developments are part of a once or twice in a century phenomenon that will carry productivity trends nationally and globally to a new higher track, or whether we are merely observing some unusual variations within the context of an otherwise generally conventional business cycle expansion. (Greenspan, 1997)*

This paper argues that "current developments are part of a once or twice in a century phenomenon," albeit one far more mundane than Greenspan suggests in his testimony.

The U.S. labor force has changed significantly during the past two decades, primarily due to the aging of the baby-boom generation. Because the teenage unemployment rate is several times higher than the aggregate unemployment rate, historical changes in the percentage of teenagers in the population have had a significant effect on the aggregate unemployment rate. I calculate that the entry of the baby boom into the labor market in the 1960s and 1970s raised the aggregate unemployment rate by about 2 percentage points. The subsequent aging of the baby boom has reduced it by about $1\frac{1}{2}$ percentage points. This is the bulk of the low-frequency fluctuation in unemployment since World War II.

This demographic story fits with other characteristics of the current expansion. For example, Greenspan (1997) notes that although consumers "indicate greater optimism about the economy," many do not perceive an unusually attractive labor market: "Persisting insecurity would help explain why measured personal savings rates have not declined as would have been expected . . . ." An explanation for these apparently contradictory opinions is that the probability of being unemployed conditional on demographic characteristics is higher now than during most other recent business-cycle peaks, particularly for men. A 42-year-old man sees that 3.7% of his peers are unemployed today, but remembers that when his father was about his age in 1973, only 2.0% of his father's peers were unemployed. He looks at his 18-year-old son, who is unemployed with 17.2% probability. He recalls that in 1973, only 13.9% of his peers were unemployed.

1.1 A PRELIMINARY CALCULATION

A simple way of establishing the magnitude of the covariance between the unemployment and the baby boom is through a linear regression of the aggregate unemployment rate at time t, U_t, on a constant and on the fraction of the working-age population (age 16–64) in its youth (age 16–24), YouthShare_t[1]:

$$U_t = \underset{(0.0049)}{0.0211} + \underset{(0.0255)}{0.1895}\ \text{YouthShare}_t + \epsilon_t, \qquad R^2 = 0.08.$$

(Standard errors in parentheses.) The youth share of the working-age population peaked at 23% in 1976, and has since declined to 16%. Thus

1. Thanks to Greg Mankiw for suggesting this calculation.

the declining youth population is correlated with a 130-basis-point reduction [0.1895 × (23% − 16%) ≈ 1.3%] in the unemployment rate. If we include a time trend in the regression, the coefficient on youth population falls to 0.1589 (standard error 0.0224), and so the point estimate of the impact of the declining youth population is 110 basis points.

There are two problems with these calculations. First, they demonstrate that there is a correlation between the youth population and the aggregate unemployment rate, but do not establish any causal mechanism. Second, the result is sensitive to the precise specification. For example, if we define the youth share to be the fraction of the working-age population between the ages of 16 and 34, the coefficient on youth share rises considerably. The same back-of-the-envelope calculation implies that the aging of the baby boom reduced the aggregate unemployment rate by about 270 basis points! Thus one should be hesitant in interpreting these regressions structurally.

1.2 PROJECT OUTLINE

The remainder of this paper gives a structural interpretation to the relationship between demographics and aggregate unemployment. Summers (1986) asserts, "There is no reason why the logic of adjusting for changes in labor force composition should be applied only to changes in" the age structure. The first goal of this paper (Section 2) is to document the disaggregated unemployment rate of different groups of workers. Using data based on the Current Population Survey, the source of official U.S. unemployment statistics, I calculate the unemployment rate of workers grouped by their observable characteristics—age, sex, race, and education.

Next I follow Perry (1970) and Gordon (1982) in calculating how much of the recent decline in unemployment is attributable to these demographic factors and how much of the decline would have happened if all demographic variables had remained constant. To perform this counterfactual exercise, I maintain the hypothesis that the unemployment rate of each group of workers is unaffected by demographics. Any change in unemployment for a group of workers would therefore have happened in the absence of demographic changes; it is a *genuine* change in unemployment. Any remaining changes in unemployment are *demographic*. I find that the changing age structure of the population reduced the unemployment rate by more than 75 basis points since the business-cycle peak in 1979. The increased participation of women has had virtually no effect on unemployment in the last two decades, while the increase in the nonwhite population raised unemployment moderately, by about 13 basis points. Finally, under the maintained hypothesis, the increase in education has been much more important: it has reduced the unemployment rate by another 99 basis points.

In summary, the aggregate unemployment rate fell by about 106 basis points since 1979. Under the maintained hypothesis, if labor-force demographics had remained unchanged, the aggregate unemployment rate would have *increased* by about 55 basis points during that 19-year period. I should not be asking why the aggregate unemployment rate is so low, but rather why the genuine unemployment rate is still so very high.

If this is a puzzle during the 1980s and 1990s, it is more so during the 1960s and 1970s, a time of rising aggregate unemployment. Summers (1986) calculates that increased education during the 1960s and 1970s should have reduced the unemployment rate by about one full percentage point, enough to outweigh all other demographic effects. He writes, "taking into account the changing composition of the labor force does not reduce and may even increase the size of the rise in unemployment [in the 1960s and 1970s] that must be explained." More generally, increases in education should have caused a sharp secular decline in unemployment in the United States and throughout the world over long time horizons. This has of course not happened, leading Summers to dismiss the relevance of demographic explanations of changes in unemployment.

A more rigorous method of dismissing demographic adjustments would be to invalidate the maintained hypothesis that the unemployment rate of different groups of workers is unaffected by demographics. The second goal of this paper is then to provide a framework for evaluating the hypothesis. I argue that there are good theoretical reasons to believe it is adequate with respect to changes in the age structure, but that it might be violated when there are changes in educational attainment.

Section 3 develops a simple model of youth unemployment with the key feature that the source is not that young workers have trouble finding jobs, but that new jobs are easily destroyed. Young workers are learning about their comparative advantage by experimenting, and so necessarily endure many brief unemployment spells. In contrast, many older workers are in extremely stable jobs. Now consider the effect of the baby boom. There are more young workers, and so more unemployment. If this gives rise to a proportional incentive to create jobs, it has no effect on the rate that young workers find jobs. The age-specific unemployment rate is unaffected by population dynamics, and it makes sense to demographically adjust the unemployment rate for age.

Section 4 points out that education may be quite different. First, employers may care about relative education more than the absolute level of education. Thus an increase in the fraction of college graduates may simply lead employers to increase the educational requirement of jobs.

This says that a shift in the education distribution may have no real effects. Second, educational choice is endogenous and correlated with (unobserved) ability. Abler workers are likely to have a lower unemployment rate for a given level of education, and an increase in education reduces the ability of the average worker with a given level of education. Therefore an increase in education will tend to raise the unemployment rate conditional on education, even if it has little or no effect on aggregate unemployment. A demographic adjustment for education would be unwarranted and potentially misleading.

Ultimately, the appropriateness of a demographic adjustment is an empirical issue, and so in Section 5 I return to the data. I first test whether changes in a group's size are correlated with changes in the group's relative unemployment rate. The reduction in the population of high-school dropouts is highly correlated with a relative increase in their unemployment rate. That is a prediction of the theory in Section 4, and implies that demographic adjustments for educational attainment are inappropriate.

Next I look at age. I find that when an age group gets larger, its unemployment rate increases relative to the aggregate unemployment rate. In particular, when the baby-boom generation was young, the youth unemployment rate increased. It correspondingly fell as the baby boom aged. In terms of the model of youth unemployment, job creation did not keep pace with the increase in young workers. The maintained hypothesis *understates* the effect of the baby boom on aggregate unemployment.

In light of this evidence, I construct a series for the effect of the baby boom on unemployment. My new hypothesis, supported by the model in Section 3 and the data in Section 5, is that the unemployment rate of prime-age workers was unaffected by the baby boom. Any movement in the aggregate unemployment rate that cannot be explained by movements in the prime-age unemployment rate is due to demographics. This includes both the direct effect under the original hypothesis of having more young workers, and the indirect effect that the baby boom apparently had on the youth unemployment rate. Figures 19 and 20 display the genuine and demographic fluctuations in unemployment. The baby boom explains a 190-basis-point increase in unemployment from 1954 to 1980 and a 150-basis-point decrease in unemployment from 1980 to 1993, which was moderated by about 30 basis points in the last five years. This is the bulk of the low-frequency fluctuations in unemployment since World War II.

1.3 RELATED LITERATURE

An older literature looks at whether changes in the age and sex composition of the labor force could explain the increase in unemployment in the

1960s and 1970s. Prominent papers include Perry (1970) and Gordon (1982). They found, as I confirm, that these variables have explanatory power. There are two significant differences between their papers and mine. First, they use a different method to calculate the demographic unemployment rate, as I explain in footnote 8. My demographic adjustment is suggested by the theory I develop in later sections of the paper, and requires less data than Perry's.

Second, these earlier papers focus on changes in the *nonaccelerating inflation rate of unemployment* (NAIRU), while I look at changes in the actual unemployment rate. I do this primarily for expository simplicity. If one has a model connecting equilibrium unemployment and inflation in mind, these two objectives are likely to be almost equivalent. If the actual unemployment rate requires a demographic adjustment, then so surely must the unemployment rate associated with no wage-push inflation. Conversely, if demographics have no effect on unemployment, then they should have no effect on the NAIRU, in the absence of some other channel connecting demographics and inflation. In support of this idea, I show at the end of this paper (Figure 21) that my demographic adjustment of the unemployment rate is remarkably similar to Staiger, Stock, and Watson's (1997) nonstructurally estimated series for the NAIRU.

Demographic adjustments to the unemployment rate have attracted less attention in recent years.[2] There are two apparent reasons. First, the aging of the baby boom should have led to a decline in the unemployment rate starting in 1980, but 1980–1986 marked a period of very high unemployment. In response, early proponents of demographic adjustments stopped making them. For example, Gordon (1997) writes, ". . . when I tested in the late 1980s to see whether the demographic changes of the 1980s . . . had reduced the NAIRU accordingly, I found that it had not. Without any justification other than its empirical performance, I arbitrarily set the textbook NAIRU equal to 6.0 percent for the entire period after 1978." This paper argues that Gordon was too quick to abandon his model. High unemployment in the early 1980s was a temporary phenomenon.

2. An exception is Council of Economic Advisors (1997). In a discussion of the NAIRU, the report notes, ". . . about 0.5 percentage point of the decline in the NAIRU since the early 1980s can be attributed to demographic changes. The single most important demographic change is the aging of the baby-boom generation: the United States now has a more mature labor force, with smaller representation of age groups that traditionally have higher unemployment rates." It is not clear what other demographic changes are considered, although education must surely not have been, since I show in Section 2.4 that college enrollment is more important than the aging baby boom. The reason the demographic adjustment is smaller in the Council's report than in this paper appears to be that they are focusing on a shorter time interval and ignoring the impact that the baby boom had on the youth unemployment rate.

Second, and more to the point of this paper, the validity of demographic adjustments came under attack by Summers (1986), who argued that if one is going to adjust the unemployment rate for changes in the age and sex composition of the labor force, one logically must adjust it for changes in education. He shows that the effect of a demographic adjustment for education in the 1970s is larger in magnitude and opposite in sign to the effects emphasized by Gordon (1982). He concludes that demographics increase the unexplained change in unemployment. An important goal of this paper is then to examine his premise, that education and age adjustments are equally sensible. Moreover, to the extent that the reader does not believe my conclusion that demographic adjustments for education are unwarranted, we must ask why the aggregate unemployment rate has not fallen steadily during the twentieth century.

Finally, many other papers seek to explain why the U.S. unemployment rate has fallen so much, or alternatively why it in fact is not low by historical standards. Making a complete list of proposed explanations is beyond the scope of this paper. I mention only one, chosen because it may be as large in magnitude and as fundamental as demographics. Juhn, Murphy, and Topel (1991) note that the labor-market participation of men has declined sharply over the last thirty years. For example, the participation of men aged 35–44 declined from 97% in 1968 to 93% today. Part of the reason for this has to do with how we measure unemployment. There has been an increase in the number of *discouraged* workers, who are without a job and available for work, but who view job search as hopeless. They are not counted as unemployed, although this omission may seem somewhat arbitrary. The implications of this can be added to the implications of the demographic adjustment, leading us to conclude that by historical standards, the U.S. *nonemployment* rate is quite high.

2. *Disaggregating Unemployment*

The U.S. Bureau of Labor Statistics (BLS) publishes its monthly unemployment report using data gathered in the Current Population Survey (CPS), a representative sample of about 50,000 households. These data can therefore also be used to calculate the unemployment rate of a group of workers sharing certain characteristics.[3] The BLS in fact publishes monthly official statistics on unemployment as a function of age, sex, and race.[4] Also, the BLS maintains annual observations of unemployment as a func-

3. Throughout this paper, I use statistics for the unemployment rate of the civilian labor force.
4. The data in this paper are available from the BLS Web site, `http://stats.bls.gov`, except where noted otherwise. The specific series used are available upon request.

tion of education. This section uses those data to characterize the unemployment experience of different groups of workers. I then reaggregate the data to construct series for the demographic and genuine components of unemployment.

2.1 AGE

After years of depression and war, the birth rate in the United States reached unprecedented levels from 1946 to 1964, the baby boom (Ventura, Martin, Mathews, and Clarke, 1996). In 1940, 8% of women aged 15–44 gave birth. At the peak of the baby boom, the birth rate rose to over 12%, but by 1975, it had fallen to less than 7%, where it has remained.

One effect of the baby boom was a large increase in the fraction of young workers in the labor force in the late 1960s and 1970s, as Figure 2 documents. Because of the low birth rate during the Great Depression

Figure 2 DURING THE LATE 1960S AND 1970S, THE SHARE OF YOUNG WORKERS IN THE LABOR FORCE INCREASED DRAMATICALLY, AS THE PEAK OF THE BABY BOOM PASSED THROUGH ITS TEENS AND EARLY TWENTIES

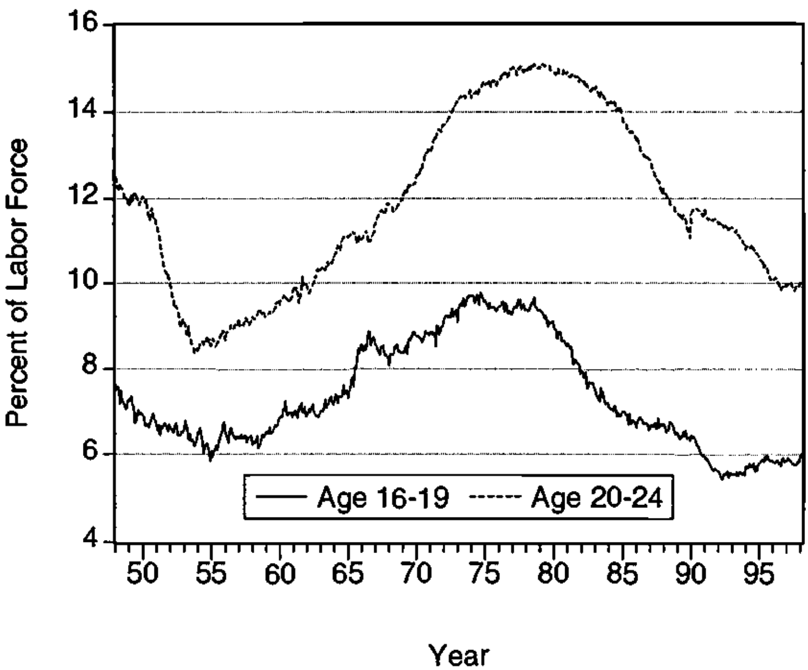

Computed from seasonally adjusted labor-force series based on the Current Population Survey.

Figure 3 THE UNEMPLOYMENT RATE OF TEENAGERS IS TYPICALLY AT LEAST 3 TIMES THE UNEMPLOYMENT RATE OF WORKERS OVER 35; THAT OF YOUNG ADULTS IS TYPICALLY ABOUT TWICE AS HIGH

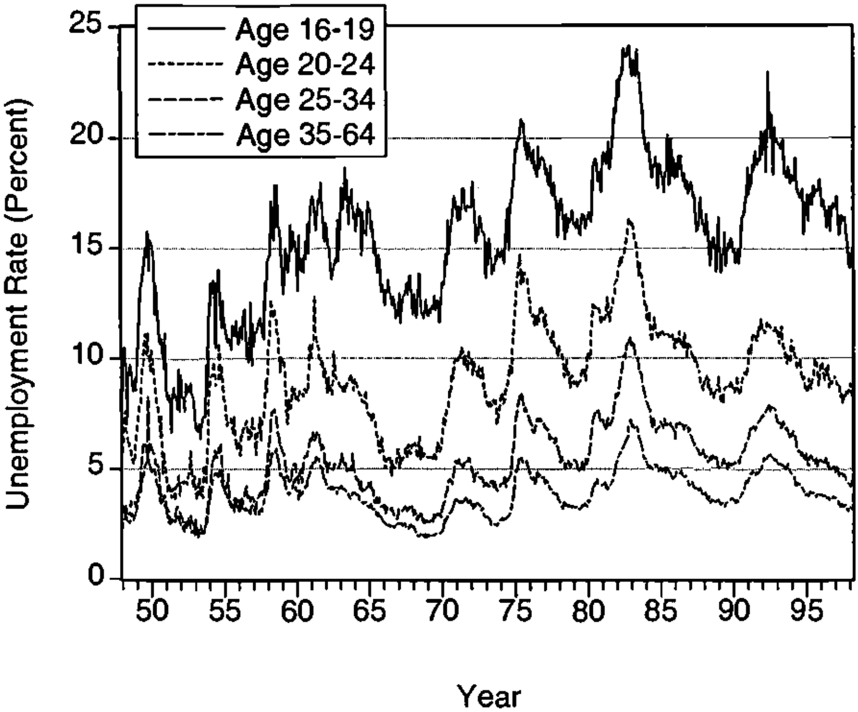

Computed from seasonally adjusted labor-force series based on the Current Population Survey.

and World War II, only 6.1% of the labor force were in their teens in 1958. This increased steadily during the next sixteen years, until by 1974, nearly 10% of the labor force was 16 to 19 years old. The teenage share then declined even more sharply, bottoming out at 5.4% in 1992, before climbing slightly in the last half decade. The share of young adults followed a similar pattern, with a natural small lag.

These demographic changes are important for aggregate unemployment, because the unemployment rate of young workers is much higher than the unemployment rate of adult workers. Figure 3 provides a graphical depiction of this fact. To be more rigorous, the unemployment rate of workers in age group i, $u_t(i)$, is described well by an ARMA(2,2) process:

Table 1 ARMA COEFFICIENTS FROM EQUATION (1), FROM OCTOBER 1957 TO NOVEMBER 1997

Age	$\bar{u}(i)$	$\alpha_1(i)$	$\alpha_2(i)$	$\theta_1(i)$	$\theta_2(i)$	$\sigma(i)$
16–19	0.169	1.620	−0.632	−1.061	0.341	0.008
	(0.009)	(0.070)	(0.068)	(0.062)	(0.044)	
20–24	0.095	1.828	−0.835	−1.048	0.264	0.005
	(0.007)	(0.064)	(0.063)	(0.061)	(0.035)	
25–34	0.056	1.824	−0.831	−0.888	0.212	0.002
	(0.005)	(0.078)	(0.076)	(0.083)	(0.049)	
35+	0.039	1.816	−0.824	−0.861	0.220	0.001
	(0.003)	(0.073)	(0.071)	(0.072)	(0.045)	

Newey–West standard errors in parentheses.

$$u_t(i) = \bar{u}(i) + \alpha_1(i)[u_{t-1}(i) - \bar{u}(i)] + \alpha_2(i)[u_{t-2}(i) - \bar{u}(i)] + \eta_t(i), \quad (1)$$
$$\eta_t(i) = \epsilon_t(i) + \theta_1(i)\epsilon_{t-1}(i) + \theta_2(i)\epsilon_{t-2}(i),$$

where $\bar{u}(i)$ is the unconditional expectation of the unemployment rate of group i, and $\epsilon_t(i)$ is white noise with variance $\sigma_t^2(i)$, allowing for heteroscedasticity. Table 1 shows the ARMA coefficients for four different groups of workers during the last forty years.[5] The coefficients have been stable and are remarkably consistent across groups. The unconditional unemployment rate of teenage workers is more than four times the unemployment rate of prime-age workers, while the unemployment rate of workers in their early twenties is nearly three times as high.

To quantify the importance of the changing age structure of the labor force, I divide the labor force into seven age groups: $I = \{16–19, 20–24, 25–34, 35–44, 45–54, 55–64, 65+\}$. Define $\omega_t(i)$ to be the fraction of workers who are in group i at time t, so $\sum_{i\in I}\omega_t(i) = 1$ for all t. Let $u_t(i)$ denote the unemployment rate of age group i at time t. Then the aggregate unemployment rate at time t is

$$U_t \equiv \sum_{i\in I} \omega_t(i)u_t(i). \quad (2)$$

There are two ways that aggregate unemployment might fall. First, the unemployment rate of different groups of workers, $u_t(i)$, might fall. Second, the population might shift towards groups with lower unemployment rates, so $\omega_t(i)$ increases for i with small $u_t(i)$ and decreases for i with large $u_t(i)$.

5. The years from 1948 to 1957 were characterized by higher-frequency cyclical fluctuations and the Korean War, so these results are sensitive to a choice of initial year before 1957. However, they are not sensitive to a choice of a later initial year.

I want to understand how much of that change would have happened if demographics had remained the same. I will refer to this as the *genuine* change in unemployment. A useful hypothesis is that if demographics had remained unchanged at some initial shares $\omega_{t_0}(\cdot)$, the disaggregate unemployment rates $u_t(\cdot)$ would have followed the same path that we observed from t_0 to t_1. This implies that the unemployment rate at time t_1 would have been

$$U^G_{t_1,t_0} \equiv \sum_{i \in I} \omega_{t_0}(i) u_{t_1}(i) \tag{3}$$

if demographics had remained the same from t_0 to t_1. The calculation of $U^G_{t_1,t_0}$ naturally depends on the choice of the base year t_0. An interesting candidate is August 1978 (78:08), as this is the demographically "worst" point in recent U.S. history. That is, $U^G_{t_1,78:08} \geq U_{t_1}$ for all t_1 since World War II, as shown in Figure 4.

$U^G_{t_1,78:08}$ rose by 208 basis points from the peak of the expansion in 1969 to the peak in 1979, a period during which the aggregate unemployment rate U_t rose by 229 basis points. Thus genuine unemployment changes

Figure 4 $U^G_{t_1,78:08}$ AND U_{t_1} FOR SEVEN AGE GROUPS

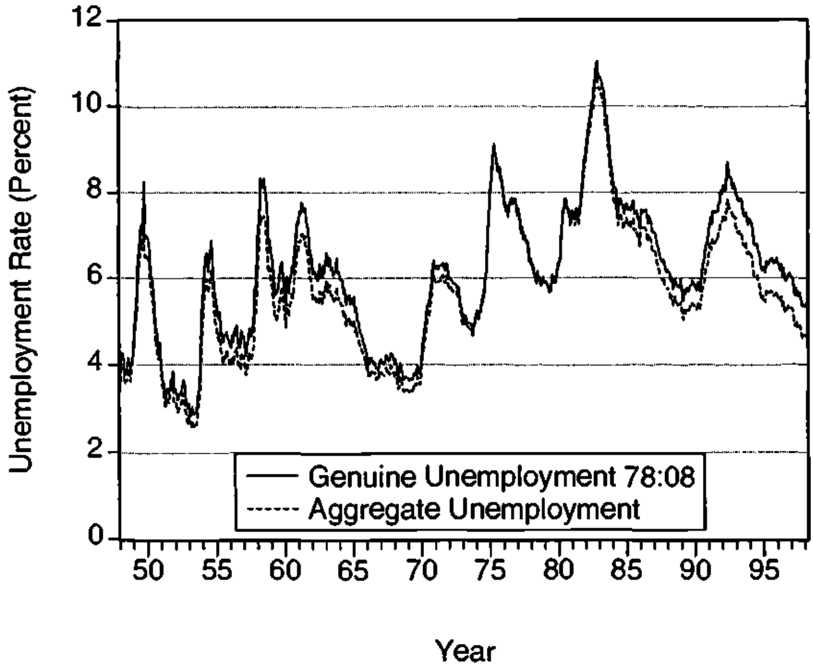

Less than 40% of the decline in aggregate unemployment during the last two decades is genuine.

Figure 5 DEMOGRAPHIC ADJUSTMENT FOR SEVEN AGE GROUPS

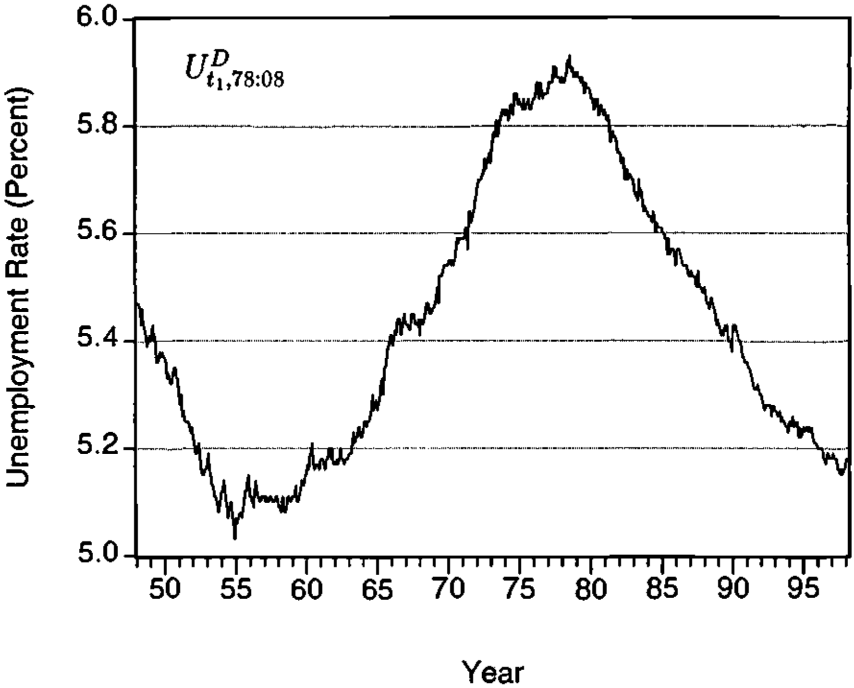

account for most of the action in the 1970s. In contrast, $U^C_{t_1,78:08}$ only declined by 25 basis points from 1979 to 1989 and by 12 basis points from 1989 to the present, about 40% of the total decline in aggregate unemployment in the first interval, and less than 30% in the second interval.[6]

The obvious alternative is to ask how much the unemployment rate changed because of demographics. Maintain our hypothesis that demographics do not affect disaggregate unemployment rates. Then if the only changes in the economy from t_0 to t_1 were demographic, the unemployment rate at t_1 would be

$$U^D_{t_1,t_0} \equiv \sum_{i \in I} \omega_{t_1}(i) u_{t_0}(i). \tag{4}$$

Changes in U^D are *demographic* unemployment changes. Figure 5 plots this series, again using a base of August 1978. $U^D_{t_1,78:08}$ rose by 90 basis

6. Using a different base year gives similar results. For example, $U^C_{t_1,54:12}$ (using the demographically "best" point as the base year) rose by 179 basis points in the 1970s and declined by 38 basis points in the last 19 years. $U^C_{t_1,97:11}$ rose by 185 basis points and then declined by 38.

points from its lowest level in 1954 to its peak in 1978, and has since declined by 78 basis points. About 65% of the decline in unemployment from the peak in 1979 to the peak in 1997 is picked up by $U^{D}_{t_1;78:08}$.[7]

One problem with these measures of demographic change is that they depend on a choice of base year. To avoid this issue, I introduce a *chain-weighted* measure of the change in unemployment attributable to demographics. Given an initial time period t_0, define for $t_1 > t_0$

$$\Delta_{t_1,t_0} = \sum_{t=t_0}^{t_1-1} \sum_{i \in I} [\omega_{t+1}(i) - \omega_t(i)] \frac{u_{t+1}(i) + u_t(i)}{2}. \tag{5}$$

$\Delta_{t+1,t_0} - \Delta_{t,t_0}$ reflects the change in demographics from t to $t + 1$. Thus Δ_{t_1,t_0} reflects the cumulative effect of changing demographics since period t_0. This series rose by 84 basis points from 1954 to its high point in August 1978, and has since fallen by 81 basis points (Figure 6). Since the 1979

Figure 6 DEMOGRAPHIC ADJUSTMENT FOR SEVEN AGE GROUPS

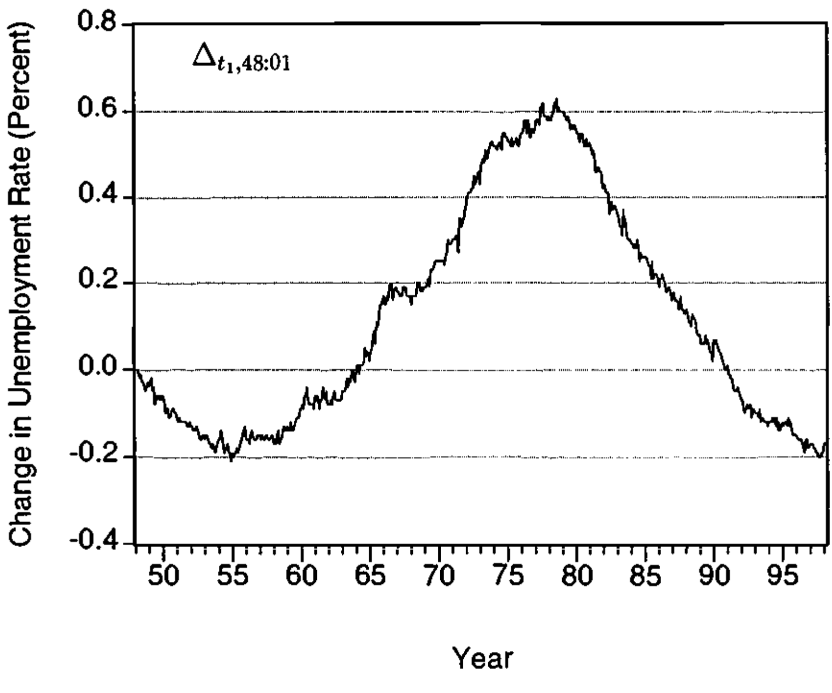

7. Different base years give similar results. We can also let $u_{t_0}(i)$ equal the unconditional expectation of group i's unemployment rate, as estimated by an ARMA(2,2) like equation (1). This implies a 91-basis-point increase in unemployment due to demographics from 1954 to 1978 and a subsequent 80-basis-point decline.

business cycle peak, $\Delta_{t_1,48:01}$ declined by 76 basis points, 68% of the total decline in unemployment.

Δ is not a perfect measure of demographics either,[8] since it may be affected by the cyclicality of labor market participation. For example, youth participation varies more with the business cycle, since prime-age workers do not move in and out of the labor market very easily.[9] During the 1960s and 1970s, when the youth population was increasing, this should have reduced the change in the weights $\omega(i)$ during recessions, since the secular increase in the share of youth was offset by their cyclical decrease in participation. It should have increased the change in the weights during expansions, since the two effects moved in the same direction. In calculating Δ, we multiply these weights by the current unemployment rate, which is higher during recessions. This would tend to moderate the change in Δ. This argument can be reversed to suggest that changes in Δ during the 1980s and 1990s are exaggerated. This may help explain why the simple demographically adjusted unemployment series U^D changed by more than the chain-weighted series Δ in the 1960s and 1970s, and by less in the 1980s and 1990s. Still, Figure 2 demonstrates that the secular shifts in labor-market shares swamp any cyclical variations, so this issue in unlikely to be quantitatively important.

Data problems may bias both demographic adjustments to zero. Suppose we divide the population into only two age groups, 16–24 and 25+. We find that Δ rose by 74 basis points from 1954 to 1978, and then declined by 73 basis points from 1978 to 1997. Any other division into two groups gives smaller changes. For example, with age groups 16–19 and 20+, we only observe half of the change in Δ during both time intervals. It is not surprising that aggregation reduces the measured demographic changes, since it is precisely the differences in disaggre-

8. Both of my measures of the demographic unemployment rate differ from the measure suggested by Perry (1970) and used by him and by Gordon (1982). Perry and Gordon weight different groups by their members' total annual earnings, and construct an alternative measure of unemployment using these weights. The demographic adjustment is then the difference between the actual unemployment rate and this series. I show in Section 3 that a simple model justifies the use of U^G to represent genuine changes in unemployment, and U^D or Δ to represent demographic changes. My theory does not suggest Perry's demographic adjustment.

9. One way to quantify this is to look at the covariance between real GDP growth and labor-market participation growth for different age groups. If a group moves out of the labor market during recessions, this should be positive. From 1948 to 1979, the covariance for teenage workers was 8 times the covariance for workers aged 45–54 and 10 times the covariance for workers aged 55–64. For the remaining four groups of workers, the covariance was actually negative. Since 1980, the covariance between these two series has become positive for all groups. Still, the covariance is 3 times as large for teenagers as for workers 20–24, and at least 8 times larger than for the remaining five groups.

Figure 7 THE FRACTION OF THE LABOR FORCE THAT IS FEMALE HAS STEADILY INCREASED SINCE WORLD WAR II

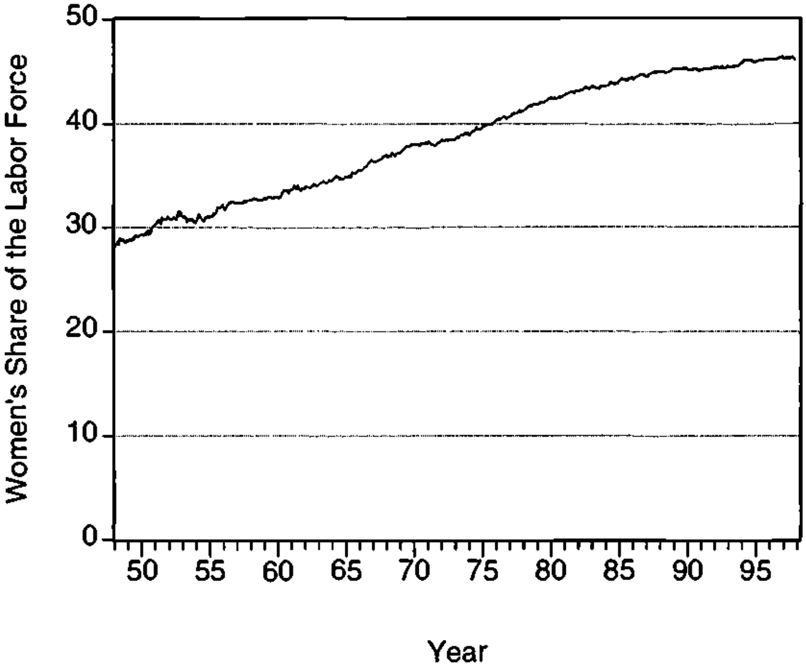

Computed from seasonally adjusted labor-force series based on the Current Population Survey.

gate unemployment rates and the changes in labor-market shares that result in demographic adjustments. If we could divide the population into more age groups, logic and evidence suggest that we would attribute more of the decline in unemployment to the changing age structure of the labor force.

In light of these caveats, under the maintained hypothesis, the aging of the baby boom explains at least 70% of the decline in unemployment since 1979, leaving about 30 or 40 basis points unexplained. The entry of the baby boom into the labor market had the opposite effect, but explains a smaller percentage of the larger increase in unemployment during the sixties and seventies.

2.2 SEX

Another dramatic trend has been the increasing participation of women in the U.S. labor market (Figure 7). This would have the potential to explain a change in the aggregate unemployment rate if women had a different unemployment rate than men. Historically this was the case, although the gap has disappeared during the last fifteen years (Figure 8). More precisely, both series are described by statistically identical ARMA(2,2) processes.

As this suggests, female participation cannot explain much of the change in unemployment. For example, divide workers up by sex and construct the chain-weighted series Δ as in equation (5). This explains a 19-basis-point rise in the unemployment rate from 1950 to 1979, and then virtually no change in the unemployment rate from 1979 to 1997 (Figure 9).

The behavior of U^D defined in equation (4) depends strongly on the choice of base year. If we choose a year when women's unemployment is lower (higher) than men's, then U^D monotonically declines (increases). This highlights the disadvantage of U^D when the relationship between disaggregate unemployment rates is unstable. For this reason I use Δ as my primary measure of demographic unemployment.

One might be tempted to add the demographic changes reported here

Figure 8 DURING THE 1960S AND 1970S, THE UNEMPLOYMENT RATE OF MEN WAS ABOUT 2 PERCENTAGE POINTS BELOW THAT OF WOMEN, BUT THAT GAP DISAPPEARED IN 1980

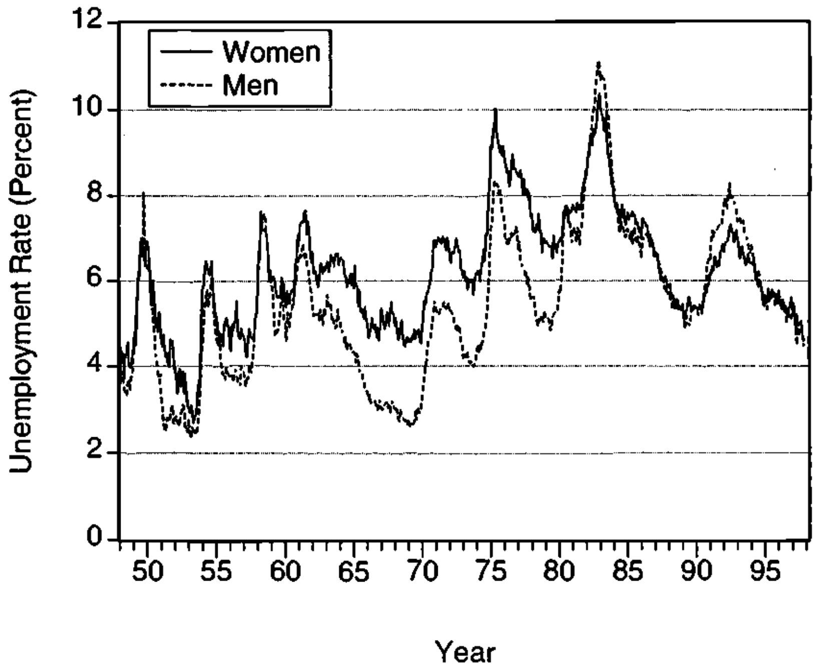

Computed from seasonally adjusted labor-force series based on the Current Population Survey.

Figure 9 DEMOGRAPHIC ADJUSTMENT FOR INCREASED FEMALE PARTICIPATION

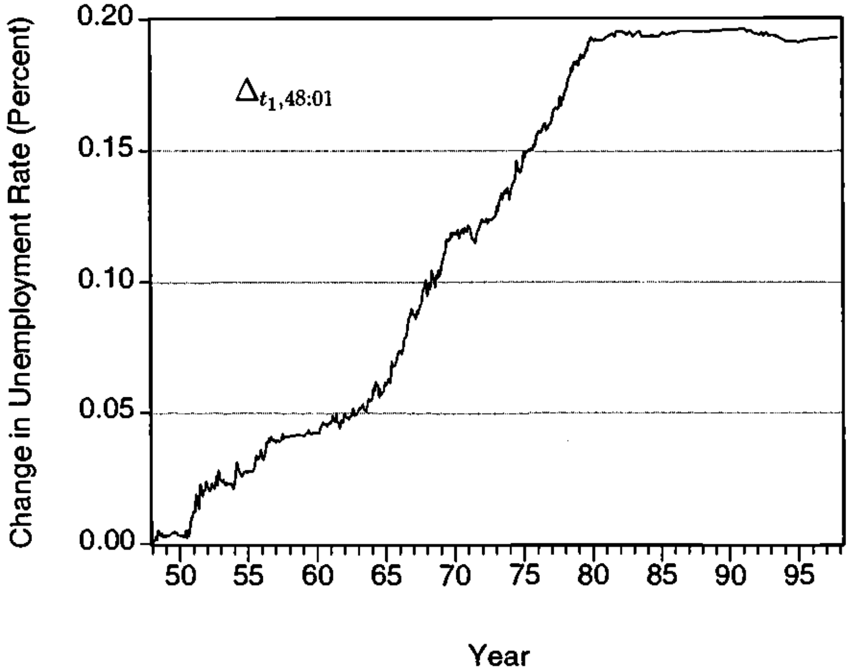

and in the first part of this section, in order to obtain the effects of the changes in the age and sex composition of the labor force—a 103-basis-point increase in unemployment from 1954 to 1978, followed by an 81-basis-point decline in unemployment from 1978 to 1997. That calculation is incorrect if there are any *mix effects* between age and sex. For example, the increase in participation has been most dramatic for prime-age women, who have a much lower unemployment rate than teenage women. Thus the increase in female participation may be double-counted as a relative decrease in teenage participation. Constructing Δ with fourteen age and sex groups, we find that age and sex jointly account for a 96-basis-point increase in the unemployment rate from 1954 to 1978, followed by an 80-basis-point decline in the unemployment rate from 1978 to 1997. Mix effects are quantitatively unimportant.

2.3 RACE

The fraction of the labor force that is white was nearly constant from 1948 to 1971. Since then, it has declined from 89% to 84%. Figure 10

Figure 10 THE FRACTION OF BLACKS HAS INCREASED BY LESS THAN 2%, WHILE THE FRACTION OF WHITES HAS DECLINED BY ABOUT 5%

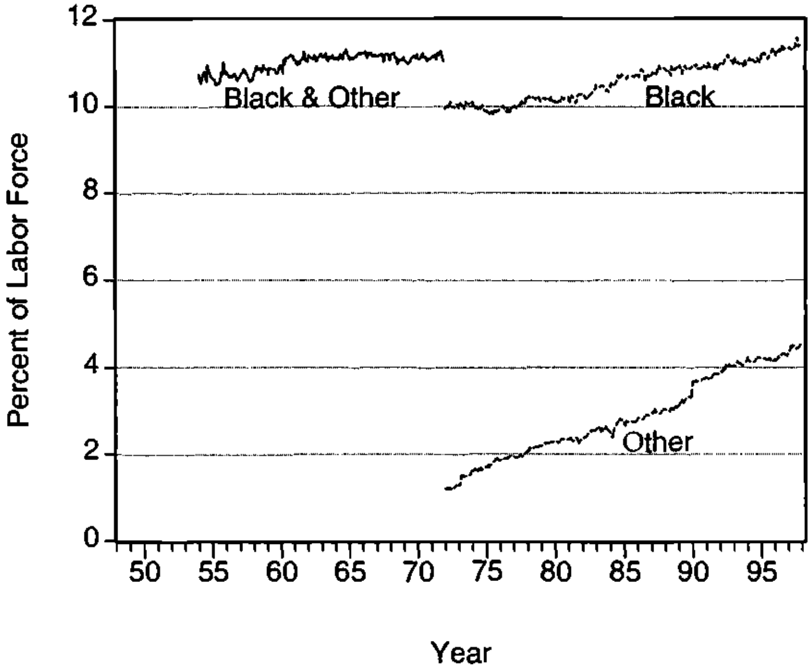

Computed from seasonally adjusted labor-force series based on the Current Population Survey. Data limitations only allow me to divide the labor force into two groups, white and other, before 1972.

shows that most of the increase in labor-force share has been for people who are neither white nor black. The white unemployment rate is only slightly lower than the unemployment rate for this group (Figure 11). The big gap between black and white unemployment rates should have had little effect on demographics, since the black labor-force share has been stable.

Dividing the labor force into three race categories, black, white, and other,[10] I find that Δ was constant from 1960 to 1977, and has since increased by 16 basis points (Figure 12). The changing racial composition has slightly mitigated the decline in unemployment in the last two decades.

10. Data limitations only allow me to divide the labor force into two groups, white and other, before 1972. Since whites and blacks account for 98.8% of the labor force in 1972, this is unlikely to be a problem. After 1972, separating out blacks and other nonwhites is quantitatively important.

Figure 11 THE BLACK UNEMPLOYMENT RATE IS MUCH HIGHER THAN THE WHITE UNEMPLOYMENT RATE

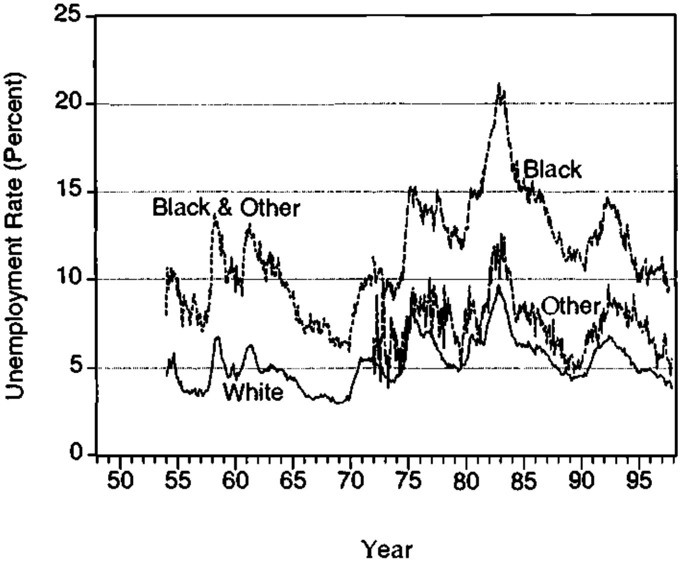

Computed from seasonally adjusted labor-force series based on the Current Population Survey. Data limitations only allow me to divide the labor force into two groups, white and other, before 1972.

Figure 12 DEMOGRAPHIC ADJUSTMENT FOR INCREASED NONWHITE PARTICIPATION

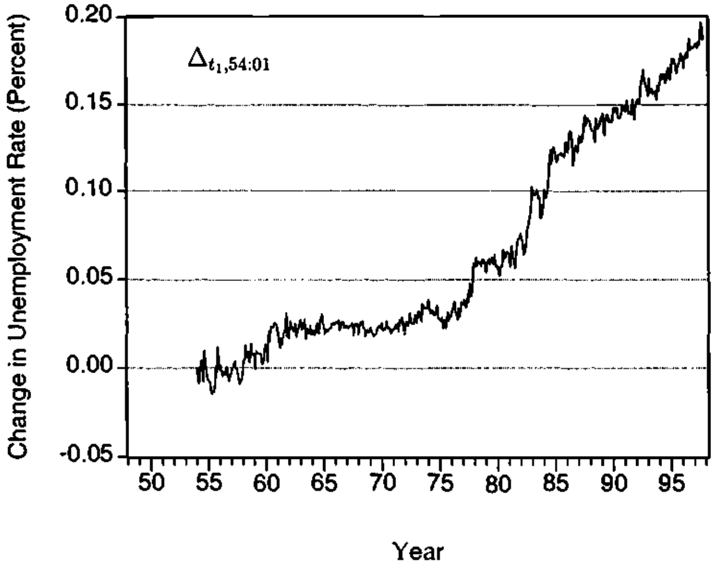

Computed from seasonally adjusted labor-force series based on the Current Population Survey.

I do not expect that it is race *per se* that leads nonwhites to have a high unemployment rate.[11] Instead, it is variables such as poor schooling and poverty. A demographic adjustment for the changing race composition would be misleading, if the relationship between race, quality of school, and wealth is changing over time. Since the magnitude of the race adjustment is not very large, this distinction is economically not very important. In practice, I do not adjust the unemployment rate for the racial composition of the labor force.

2.4 EDUCATION

The final trend that I examine is the increased education of the U.S. labor force (Figure 13). Since 1970, the percentage of workers[12] with at least some college education has increased from 26% to 56%. The percentage with only a high-school diploma increased from 38% in 1970 to 40% in the early 1980s, and has since fallen to 33%. The percentage with less than a high-school diploma fell steadily from 36% to 11% during this time period.

Also, less educated workers have a much higher unemployment rate (Figure 14). The unemployment rate of workers with less than a high-school diploma has been three to five times as high as the unemployment rate of workers with a college degree throughout this time period. The unemployment rate of those with a high-school diploma has been two to three times as high. These differences have increased over time.

I can again construct the variable Δ as in equation (5), using the four education categories (Figure 15). I find that Δ has fallen by 99 basis points since 1979.[13] This is larger than the decline that can be explained by the changing age structure, so changes in education appear to be a promising explanation for the recent decline in aggregate unemployment. In fact, the sum of an age and an education adjustment is a 175-basis-point decline in unemployment during the 1980s and 1990s, while the actual unemployment rate only fell by 106 basis points.

I will argue in the remainder of this paper that this education adjustment is misguided. As indirect evidence for this claim, one can perform

11. Employer preference for hiring whites rather than blacks would give one direct link between race and unemployment. One implication of the model in Section 4.1 is that if unemployment differences are due to discriminatory hiring, demographic adjustments to the unemployment rate are inappropriate.
12. Throughout this section, I look at workers between the ages of 25 and 64. A 16-year-old who works while in high school is probably quite different than an adult who dropped out of high school many years before. Since most workers have completed their education by age 25, I avoid complex aggregation issues by focusing on these workers.
13. The annual change of about six basis points remains the same if we omit 1993 and 1994, the years that were affected by the redesign of the CPS.

Figure 13 THE FRACTION OF WORKERS WITH AT LEAST SOME COLLEGE EDUCATION HAS INCREASED STEADILY SINCE 1970, AT THE EXPENSE OF WORKERS WITH LESS THAN A HIGH-SCHOOL DIPLOMA

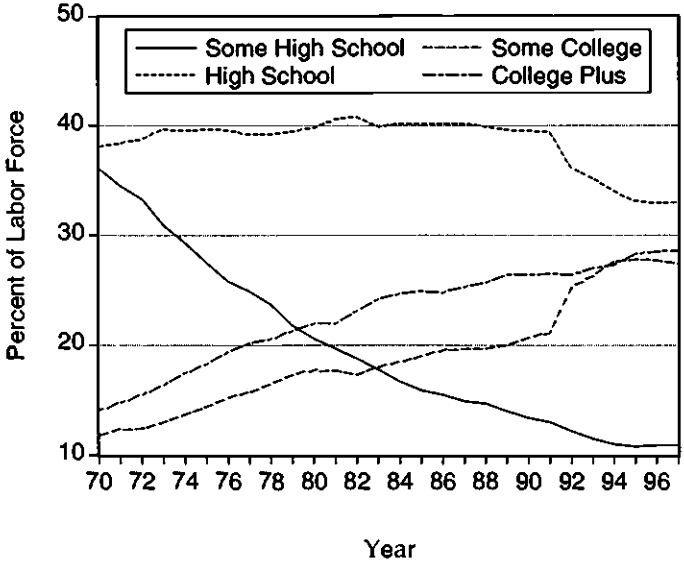

Data are for March of each year. Since 1992, data on educational attainment have been based on the highest diploma or degree received, rather than the number of years of school completed. Data from 1994 forward are not directly comparable with data for earlier years, due to the Current Population Survey redesign. I am grateful to Steve Haugen at the Bureau of Labor Statistics for providing me with these data.

the same exercise in previous decades. My data show a decline in Δ of almost 60 basis points during the 1970s. Similarly, Summers (1986) calculates that the increase in education in the 1960s reduced the aggregate unemployment rate at a similar rate, about 50 basis points during the decade.[14] As Summers pointed out, if education explains the unemployment decline during the last two decades, then it increases the mystery of why the unemployment rate was so very high in 1979.

There are two possible solutions to this issue. First, the problem could be ameliorated by the mix effects between age and education. Unemployment among college graduates may have been rising in the 1950s and 1960s because most college graduates were young. This solution does not appear promising, because in looking at education, I have restricted

14. Summers effectively calculates the change in the gap between the actual unemployment rate and the genuine unemployment rate $U^G_{t_1,t_0}$, where the base year t_0 is 1965.

Figure 14 THE UNEMPLOYMENT RATE OF LESS-EDUCATED WORKERS IS MUCH HIGHER THAN THE UNEMPLOYMENT RATE OF MORE-EDUCATED WORKERS; THE GAP HAS BEEN GROWING RECENTLY

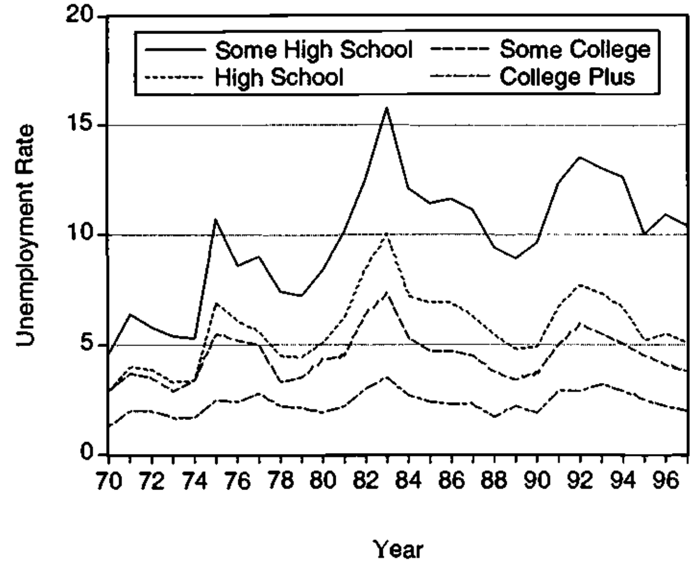

Data are for March of each year. Since 1992, data on educational attainment have been based on the highest diploma or degree received, rather than the number of years of school completed. Data from 1994 forward are not directly comparable with data for earlier years due to the Current Population Survey redesign. I am grateful to Steve Haugen at the Bureau of Labor Statistics for providing me with these data.

Figure 15 DEMOGRAPHIC ADJUSTMENT FOR INCREASED EDUCATION

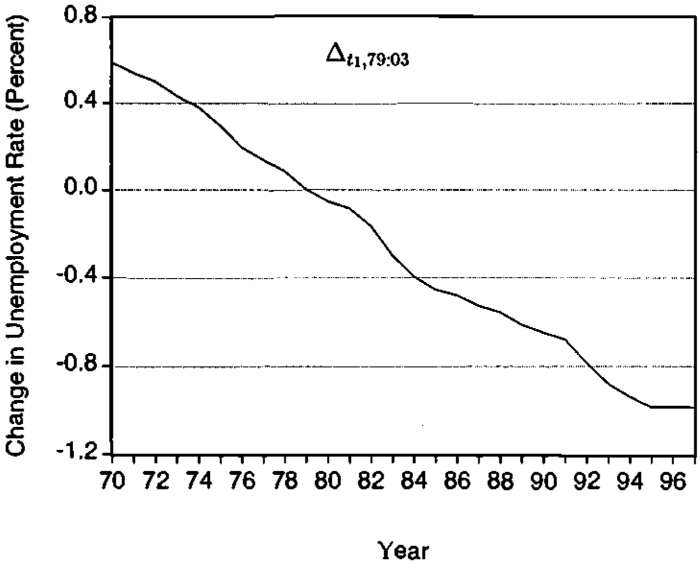

attention to workers aged 25–64. Empirically, the relationship between age and unemployment is weak for these workers, and so there is unlikely to be much overlap between the two demographic adjustments. Second, the maintained hypothesis, that changes in education do not affect unemployment conditional on education, may be false. The next two sections address the theoretical foundations of the maintained hypothesis, and the final section evaluates it empirically.

3. *Youth Unemployment*

Empirically, changes in the age and education composition of the labor force account for large changes in aggregate unemployment under the maintained hypothesis. Other demographic changes do not have as much potential explanatory power. Thus the remainder of this paper focuses exclusively on age and education.

This section develops a benchmark model of youth unemployment, in order to explore whether changes in the age structure of the population affect unemployment rates conditional on age. The model illustrates conditions under which the maintained hypothesis is satisfied, while simultaneously helping us to understand the bias that would be introduced by plausible violations of the hypothesis.

The premise of the model is that young workers do not have a particularly hard time finding jobs, but instead they have trouble keeping jobs. This is motivated by the fact that the mean and median unemployment durations in the United States are increasing functions of age. For example, the median unemployment duration of unemployed teenage workers during 1997 averaged 5.6 weeks, and the mean was 10.3 weeks. For workers between 55 and 64 years old, the median unemployment duration was 10.3 weeks and the mean was 21.9. This monotonic order surprisingly holds for the six working age groups (16–19, 20–24, 25–34, 35–44, 45–54, and 55–64) and for every year since 1976, when the Bureau of Labor Statistics began reporting unemployment duration data.[15] Since more young workers are unemployed, but those who lose their jobs stay unemployed for shorter times, it follows that young workers are much more likely to lose their jobs.[16]

15. The fact that older workers stay unemployed for longer may be due to their superior access to unemployment insurance. This possibility goes beyond the model here.
16. There is some controversy about this conclusion. Clark and Summers (1982) argue that youth unemployment duration is reduced by unemployed workers leaving and reentering the labor force, and that the source of youth unemployment is a core group of young workers who cannot find jobs.

I assume that young people lose their jobs more frequently not because they are young, but because they are inexperienced. More precisely, the hazard rate of job loss is decreasing in the length of time since the worker was last unemployed or out of the labor force, her *job tenure.* Thus, when an older worker loses her job, she is thrust back into the same situation as a younger worker. There are a number of theoretical justifications for this assumption. A new employee may be unsuitable for her job, but this can only be learned by trial and error. After surviving an apprenticeship, her job security increases. Moreover, as she stays on the job for a long period of time, she acquires specific human capital, making firing costly. Even if she quits her job to take a new one, she will tend to do so only if she expects it will enhance her job security. There is also an empirical justification for the assumption: it matches U.S. data.

3.1 MODELING JOB LOSS

I do not take a stand on why the hazard rate of job loss is decreasing, but instead set up a simple model with that property. A representative worker is "born" unemployed. She looks for a job, and is hired with flow probability M. That is, during any interval of length t, an unemployed worker is hired with probability $1 - e^{-Mt}$. When she is hired, her marginal product is initially equal to some constant s. Thereafter, productivity follows a random walk in continuous time: the productivity after t periods is $x(t) = s + \sigma Z(t)$, where $Z(t)$ is a standard Brownian motion. Thus after $\tau < t$ periods, $x(t)$ is a normally distributed random variable with mean $x(\tau)$ and variance $\sigma^2(t - \tau)$. I assume that the job is destroyed when productivity falls below some threshold $\underline{x} < s$.[17] Following a separation, she looks for a new job, finding one with flow probability M. The productivity in the new job is again initially equal to s, and again follows a random walk. This repeats forever. The unemployment rate of workers at age t is equal to the probability that this representative worker is unemployed at age t. Later in this section, I endogenize the separation threshold $\underline{x}$ and the hiring rate M. However, it is simpler at this stage to treat these two variables as exogenous.

Under these assumptions, a match ends within t periods with probability

$$F(t) = 2\Phi\left(-\frac{s - \underline{x}}{\sigma\sqrt{t}}\right), \tag{6}$$

17. If $\underline{x} \geq s$, all matches are destroyed immediately upon creation, an uninteresting case.

where Φ is the cumulative distribution of a standard normal random variable. Figure 16 depicts the match survival probability for one particular value of $(s - \underline{x})/\sigma$.

The intuition for this result is the *reflection principle* (Karatzas and Shreve, 1991). Momentarily ignore the fact that matches are terminated when productivity reaches the threshold $\underline{x}$. Consider a sample path that first reaches $\underline{x}$ after t' periods. Given this, productivity at some time $t > t'$ is distributed normally with mean $\underline{x}$. Thus among sample paths that first reach the separation threshold after t' periods, only half of them have productivity less than $\underline{x}$ after $t > t'$ periods. Since t' was chosen arbitrarily, this is true for any $t' \in [0,t]$. Thus for sample paths that reached the separation threshold within t periods, exactly half of them have productivity less than the separation threshold after t periods. The fraction of sample paths that are below the separation threshold after t periods is

$$\Phi\left(-\frac{s - \underline{x}}{\sigma\sqrt{t}}\right).$$

Figure 16 THE PROBABILITY OF A MATCH SURVIVING FOR t PERIODS

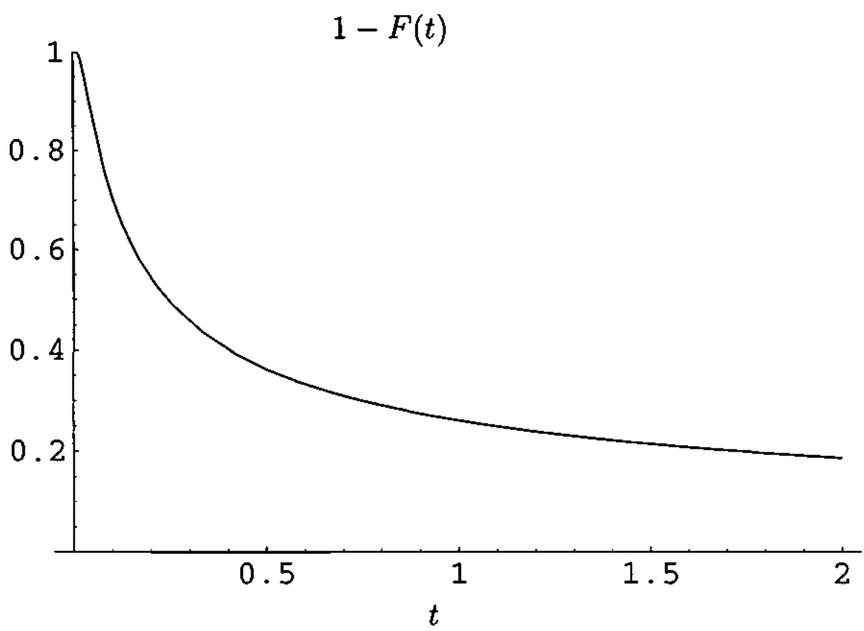

Drawn for $(s - \underline{x})/\sigma = \frac{1}{3}$.

Figure 17 THE HAZARD RATE FOR MATCH SEPARATIONS AFTER t PERIODS

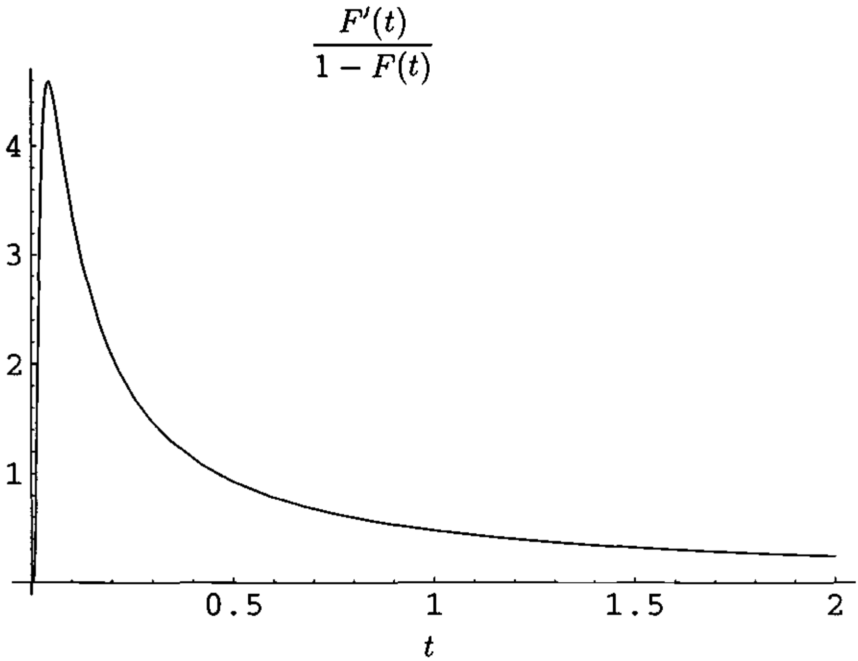

Drawn for $(s - \underline{x})/\sigma = \frac{1}{3}$.

This must be half of the fraction of sample paths that reach the separation threshold within t periods, and hence half the fraction of matches that are terminated within t periods, $F(t)/2$.

The hazard rate of separation in a t-period-old match is $F'(t)/[1 - F(t)]$, depicted in Figure 17. This is increasing for small values of t, since a sample path is unlikely to fall from s to $\underline{x}$ in a very short period of time. It is decreasing for large values of t, since, conditional on a match having survived for a long period of time, the productivity is likely to be much larger than the separation threshold.

3.2 THE RELATIONSHIP BETWEEN UNEMPLOYMENT AND AGE

Let $u(t)$ denote the probability that the representative worker is unemployed at age t. This must satisfy a backward-looking differential equation

$$\dot{u}(t) = -Mu(t) + M\int_0^t u(\tau)F'(t - \tau)\, d\tau \tag{7}$$

with initial condition $u(0) = 1$. Unemployment decreases due to match creation and increases due to match separations. The flow probability that a worker is hired at age t is the flow probability that an unemployed worker is hired M times the probability that she is actually unemployed at t. The flow of separations depends on earlier hiring probabilities. She was hired at age $\tau \in [0,t]$ with flow probability $Mu(\tau)$. The flow probability that this job ends at t is $F'(t - \tau)$, where F is the probability that a match does not survive, given in equation (6).

One can show that $u(t)$ converges to zero over time, because everyone winds up in very good matches eventually. Also, simulations show that $u(t)$ is monotonically decreasing (Figure 18), although, since the separation hazard rate is nonmonotonic, I cannot prove this analytically.

A simpler model in which the flow probability of a separation is a constant $\delta > 0$ (e.g. Mortensen and Pissarides, 1994) would deliver the same qualitative predictions as this model. However, quantitatively the models are extremely different. In the Mortensen–Pissarides model, the unemployment rate moves exponentially towards its steady-state value; half the gap between the actual and steady-state unemployment

Figure 18 THE UNEMPLOYMENT RATE OF A WORKER AS A FUNCTION OF HER LABOR-MARKET EXPERIENCE

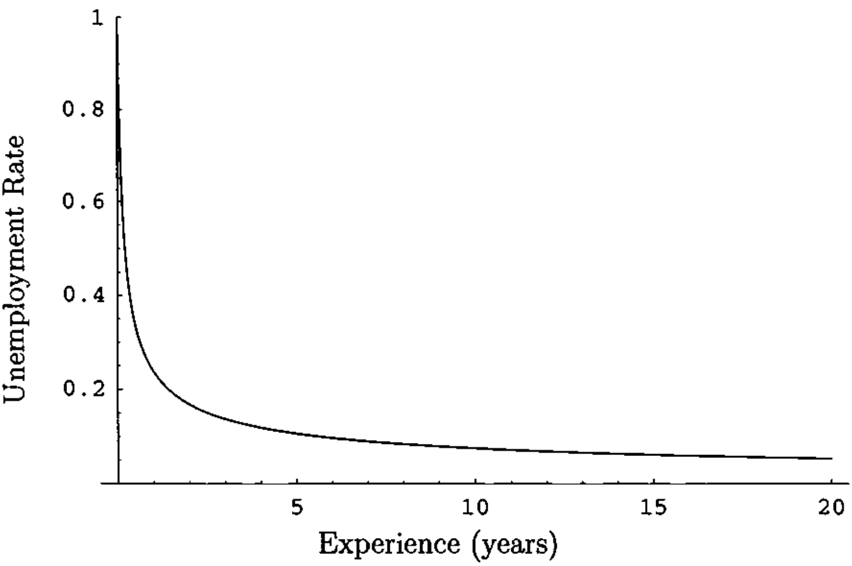

Drawn for $(s - \underline{x})/\sigma = \frac{1}{3}$ and $M = 5$.

rate is closed in (log 2)/(δ + M) periods. For plausible parameters, this implies the unemployment rate of workers with one year of labor-force experience will be indistinguishable from the unemployment rate of prime-age workers. This would be an inadequate explanation for youth unemployment.

The (mostly) declining hazard rate of job loss in this model implies a much slower decline in unemployment as a function of age. In the example depicted in Figure 18, the unemployment rate of workers with 1 year's labor-market experience is 24%. It declines to 12% after 4 years and 6% after 16 years. The eventual steady-state unemployment rate is zero. Thus there is a quantitatively strong and persistent relationship between age and unemployment in this model which is absent from the simpler model. This is why a declining hazard rate of job loss is a necessary ingredient in a model that explains youth unemployment through differential rates of job destruction. The interesting point is that it appears to be a sufficient ingredient as well.

3.3 DEMOGRAPHIC ADJUSTMENT

I am now in a position to ask whether the demographic adjustments U^D and Δ and the genuine unemployment rate U^G appropriately reflect the impact of changes in the age structure of the population. Let $N(t)$ denote the population measure of infinite-lived workers at time t. Each worker's life proceeds as described above. To avoid introducing aggregate uncertainty, assume that all the stochastic processes in the economy are independent. That is, the hiring probabilities of different workers are independent, and the stochastic process for the productivity of worker i in job j, $Z_{ij}(t)$, is independently distributed.

The sole source of aggregate fluctuations is variations in the population growth rate, $\dot{N}(t)/N(t)$. Applying equation (2) and the law of large numbers, the aggregate unemployment rate must equal

$$U_t = \int_0^\infty u_t(\tau) \frac{\dot{N}(t-\tau)}{N(t)} \, d\tau,$$

where $u_t(\tau)$ is the unemployment rate of τ-year-old workers in period t, as given by the solution to differential equation (7), $\dot{N}(t-\tau)$ workers were born τ periods ago, and the current population is $N(t)$. Thus the labor-market share of τ-period-old workers is $\omega_t(\tau) = \dot{N}(t-\tau)/N(t)$.

Since the only shocks in this model are demographic, all fluctuations in unemployment should be attributed to demographics. It is easy to verify conditions under which our constructions from Section 2 behave correctly. Use (3), (4), and the continuous time limit of (5):

$$U^G_{t_1,t_0} = \int_0^\infty u_{t_1}(\tau) \frac{\dot N(t_0 - \tau)}{N(t_0)} \, d\tau,$$

$$U^D_{t_1,t_0} = \int_0^\infty u_{t_0}(\tau) \frac{\dot N(t_1 - \tau)}{N(t_1)} \, d\tau,$$

$$\Delta_{t_1,t_0} = \int_{t_0}^{t_1} \int_0^\infty u_t(\tau) \frac{d}{dt} \left(\frac{\dot N(t - \tau)}{N(t)} \right) d\tau \, dt.$$

Now suppose the unemployment rate of workers conditional on their age does not vary over time: $u_t(\tau) = u_{t_0}(\tau)$ for all t, t_0, and τ. Simplifying the equations above yields three conclusions:

1. The genuine unemployment rate is constant over time t_1 for fixed initial time t_0: $U^G_{t_1,t_0} \equiv U_{t_0}$, so none of the changes in unemployment are genuine.
2. The demographic unemployment rate tracks the actual unemployment rate, $U^D_{t_1,t_0} \equiv U_{t_1}$, so all of the changes in unemployment are demographic.
3. The chain-weighted unemployment rate tracks the change in the unemployment rate, $\Delta_{t_1,t_0} \equiv U_{t_1} - U_{t_0}$, again stating that all of the changes in unemployment are demographic.

All three constructions recognize that the change in aggregate unemployment is demographic. On the other hand, if the age-specific unemployment rates depend on the age distribution, and so vary over time, demographic adjustments are misleading.

In summary, the question of whether these constructions capture demographic changes well is equivalent to the question of whether the disaggregate unemployment rates $u(\tau)$ are independent of the age distribution.[18] By equations (6) and (7), this requires that the age distribution not affect the hiring rate M and the threshold $\underline{x}$.[19]

18. This is certainly not the only model that ensures that disaggregate unemployment rates satisfy this property. For example, if heterogeneous workers choose to segregate themselves into separate labor markets, as in Acemoglu and Shimer (1997), a change in the distribution of workers changes the size of each labor market, but has no effect on unemployment rates. I do not use this segregation model to study youth unemployment, since it focuses attention on job creation. As discussed before, the primary cause of youth unemployment is the high rate of job destruction.

19. I assume that the technological parameters s and σ are unaffected by demographics.

3.4 CLOSING THE MODEL

To see whether this is a reasonable assumption, I endogenize these two variables, following the general methodology of Pissarides (1985). I have already discussed workers' lifetime in some detail, and here I briefly outline the environment faced by firms. I assume firms use a constant-returns-to-scale technology, and so without loss of generality I impose that a firm can employ at most one worker. It can either have a vacancy (and be looking for a worker) or have a filled job (and be producing). When a firm has a filled job, its productivity follows a random walk, as described before. The first issue is the determination of the separation threshold $\underline{x}$. My solution assumes that the economy is in steady state, but these results can all be generalized to nonstationary environments.

The Separation Threshold. An employment relationship is terminated when it is in the mutual interest of the worker and her employer.[20] Let $\mathcal{W}(x)$ denote the expected present value of the firm's profits plus the worker's wages, discounted at the interest rate $r > 0$, as a function of current productivity x. Let $\mathcal{U}$ denote the worker's value following a separation, the expected present value of her future wages while she is unemployed; and $\mathcal{V}$ denote the firm's value, the expected present value of its future profits while it has a vacancy. These values do not depend on x, as this represents an idiosyncratic or match-specific shock. Then the match ends if $\mathcal{W}(x) \leqslant \mathcal{U} + \mathcal{V}$. Otherwise, the productivity process implies $\mathcal{W}(x)$ satisfies a second-order differential equation:

$$r\mathcal{W}(x) = x + \tfrac{1}{2}\sigma^2\mathcal{W}''(x).$$

The flow value of a match comes from its current productivity x plus the capital gain from changes in x. Although there is no drift in the stochastic process for productivity, variability of x raises the value if $\mathcal{W}$ is a convex function. The expected value of $\mathcal{W}$ next period is larger than $\mathcal{W}$ evaluated at the expected productivity, an application of Jensen's inequality.

The general solution to this differential equation is

$$r\mathcal{W}(x) = x + k_1e^{-(\sqrt{2r}/\sigma)x} + k_2e^{(\sqrt{2r}/\sigma)x}$$

20. A generous interpretation of this model would recognize that it allows for job-to-job movement. Suppose an employed worker takes a new job when it is in the mutual interest of her current and future employers and herself. This may be ensured by the structure of her contract, which can include breach-of-contract penalties, such as vested pension plans (Diamond and Maskin, 1979, 1981). Then $x(t)$ should be interpreted as the productivity after t periods, assuming the worker follows an optimal sequence of job-to-job moves.

for constants k_1 and k_2. We require two terminal conditions. First, if x is very large, the threat of a separation is remote. An additional unit of productivity must raise the value of a match by $1/r$. We have $\lim_{x\to\infty} \mathcal{W}'(x) = 1/r$, so $k_2 = 0$.

Next observe that $\mathcal{W}$ is weakly increasing.[21] This verifies the optimality of a threshold rule. Now if productivity x is close to the threshold $\underline{x}$, it will fall below $\underline{x}$ in a short interval of time with probability close to 1. This property of Brownian motions implies that a small change in productivity near the threshold $\underline{x}$ will have almost no effect on the present value, $\mathcal{W}'(\underline{x}) = 0$. Thus

$$k_1 = \frac{\sigma}{\sqrt{2r}} e^{(\sqrt{2r}/\sigma)\underline{x}},$$

so

$$r\mathcal{W}(x) = x + \frac{\sigma}{\sqrt{2r}} e^{-(\sqrt{2r}/\sigma)(x-\underline{x})}. \tag{8}$$

Finally, the value of the match at the threshold must equal the sum of the values after a separation, $\mathcal{W}(\underline{x}) \equiv \mathcal{U} + \mathcal{V}$:

$$\underline{x} = r(\mathcal{U} + \mathcal{V}) - \frac{\sigma}{\sqrt{2r}}. \tag{9}$$

Matches survive even after productivity passes below the myopic threshold $x = r(\mathcal{U} + \mathcal{V})$. The reason is that a separation destroys the option to take advantage of future productivity increases. The option is worth more when agents are patient (r is small) and when the variability of output is large (σ is large), since productivity is more likely to increase substantially in the near future.

Job Search and the Hiring Rate. I next turn to the job search process. This will allow me to determine the continuation values $\mathcal{U}$ and $\mathcal{V}$ and the hiring rate M.

An important question is how wages are determined in this environment. There is generally surplus when a new match is created, since the worker and firm are jointly better off matching instead of waiting for

21. To prove this formally, use the fact that the optimal stopping rule for one value of x can also be used for a higher value x'. This must yield at least as high a value. Using an optimal termination rule must yield a still-higher value.

new partners. There are a number of ways that the surplus can be divided.[22] At this point choosing how is not very important. However, later in the paper it is convenient to use a multilateral bargaining rule (Shimer, 1997), and so I introduce it now.

Multilateral bargaining focuses on the mechanics of the matching process. Every unemployed worker targets her job search towards one vacancy, selected at random. With exogenous flow probability π, the vacancy closes. However, that does not guarantee the worker a job. Other unemployed workers may also be applying for the job opening, giving the firm an opportunity to hire one of the competing applicants. The multilateral bargaining rule is that the firm sells the job to the highest bidder. This can equivalently be thought of as Bertrand competition or a second-price auction between job applicants.[23] As a result, a worker's wage depends on the circumstances under which she was hired.

Since all workers are equally likely to apply for all job openings, the number of applicants for a given vacancy is a Poisson random variable with expectation q, equal to the ratio of the measure of unemployed workers to the measure of firms with vacancies. That is, the probability that $n \in \{0, 1, 2, \ldots\}$ other applicants compete for a particular job is $q^n e^{-q}/n!$. In particular, when a worker's desired job vacancy closes, there are no other applicants with probability e^{-q}. In this event, she is hired and is able to extract all the surplus from the match. Her value jumps by $\mathcal{W}(s) - \mathcal{V} - \mathcal{U}$ upon being hired, where $\mathcal{W}(s)$ is the value of a new match. She may still be hired if there are other job applicants; however, if this happens, the firm is able to extract all the surplus. The firm's value jumps by $\mathcal{W}(s) - \mathcal{V} - \mathcal{U}$, while all the applicants' values are constant whether or not they are hired.

Putting this together, the value of an unemployed worker satisfies

$$r\mathcal{U} = y + \pi e^{-q}[\mathcal{W}(s) - \mathcal{V} - \mathcal{U}]. \tag{10}$$

22. For example, Pissarides (1985) sparked a literature that assumes wages are set by *Nash bargaining*. A worker and firm divide the difference between the value of a match and the sum of their outside options. Let $\beta \in [0,1]$ denote the worker's bargaining power. Her value increases by $\beta[\mathcal{W}(s) - \mathcal{V} - \mathcal{U}]$ when she gets a job. The value of a firm jumps by $(1 - \beta)[\mathcal{W}(s) - \mathcal{V} - \mathcal{U}]$. The conclusions in this section would be unchanged by this bargaining rule.

23. This wage-setting procedure is equivalent to other job auctions (Shimer, 1997). For example, workers may not know the number of competing applicants. They bid for jobs by committing to a wage if they are hired, and the firm accepts the worker who demands the lowest wage. Although the analysis is more complex, since one must solve for the equilibrium (mixed) bidding strategy, the conclusions are unchanged. In particular, $\mathcal{U}$ and $\mathcal{V}$ are the same as with this wage-setting rule. This is essentially an application of the revenue equivalence theorem (Riley and Samuelson, 1981).

Here y represents a worker's exogenous unemployment benefit or value of leisure. With flow probability e^{-q} the worker is the only applicant for a vacancy when it closes, and so enjoys a capital gain. Similarly, the value of a vacant firm satisfies

$$r\mathcal{V} = -c + \pi[1 - e^{-q}(1 + q)][\mathcal{W}(s) - \mathcal{V} - \mathcal{U}] \tag{11}$$

where c is the cost of maintaining an open vacancy. When the vacancy closes, there is a probability e^{-q} of there being no applicants. The firm must reopen the vacancy. With probability qe^{-q}, the firm receives one application, and so hires the applicant; but she keeps the whole surplus. Otherwise the firm enjoys a capital gain. Equations (8), (9), (10), and (11) can be solved for the endogenous variables $\underline{x}$, $\mathcal{U}$, $\mathcal{V}$, and $\mathcal{W}$ as functions of the various parameters and the unemployment–vacancy ratio q.

Finally, since conditional on a vacancy closing with n other applicants, a worker is hired with probability $1/(n + 1)$, the flow rate at which she is hired is

$$M = \pi \sum_{n=0}^{\infty} \frac{q^n e^{-q}}{(n + 1)!} = \pi \frac{1 - e^{-q}}{q}. \tag{12}$$

Thus M is a simple function of the unemployment–vacancy ratio q as well. If q does not depend on the age distribution of the population, the disaggregate unemployment rates do not depend on the age distribution, and a demographic adjustment for the baby boom is justified.

3.5 CONCLUSION

What does it mean for q to be independent of the age distribution? The most sensible interpretation is that job creation and labor demand are perfectly elastic. If free entry drives the value of a vacant job, $\mathcal{V}$, to zero, one can solve for the equilibrium values of q and $\mathcal{U}$ with no reference to the age distribution. Essentially, changes in labor supply are accommodated by changes in labor demand, leaving the hiring rate M and separation threshold $\underline{x}$ unchanged.

The evidence on the elasticity of labor demand is mixed. More populous countries do not have higher unemployment rates, so the cross-sectional evidence is that labor demand is elastic. But it is less clear whether this is true within countries over medium time horizons. Katz and Murphy (1992) argue that from 1965 to 1980, wages tended to decrease for a group of workers when the size of that group increased, evidence for relatively stable labor demand curves. During the 1980s, wages tended to increase, suggesting that labor demand might have *overreacted* to labor-supply shifts. Also, Krueger and Pischke (1997) argue that the U.S. "employment miracle" is due to the ability of the labor market to absorb large changes in supply. They cite the wage *decreases* for young workers at the time that the baby bust entered the labor force as supporting evidence. Another famous example is the large migration of Cubans to Miami, which had little effect on wages in Miami (Card, 1990). Thus the benchmark model offers a simple, plausible condition under which the baby boom should have had no effect on disaggregate unemployment rates, and the demographic adjustment performed in Section 2 is justified.

4. *Low Skilled Unemployment*

I now turn to the source of low skilled unemployment. Skilled workers have a lower unemployment rate because their skills are most useful while they are employed. The gap between their productivity while employed and while unemployed is large, and so the cost of searching for a job is small compared to the potential reward. Similarly, firms will spend much more effort recruiting skilled workers. There is even an industry ("headhunters") that attempts to match skilled workers with jobs. Reinforcing this effect, the large gap between skilled workers' productivity s and reservation wage $\underline{x}$ implies that they are much less likely to be fired during downturns. In short, skilled workers are rarely fired, and have a relatively easy time finding a job when they are.

Since more-educated workers tend to be more skilled, increases in education raise the fraction of the population in low unemployment categories. This suggests a demographic adjustment is merited. There are at least two objections to this reasoning. First, the absolute level of education may be less important than relative education attainment. An increase in the percentage of job seekers with a college degree not only increases the competition for jobs among college graduates, but puts workers with only a high-school diploma at a further disadvantage. Thus the unemployment rate of all education levels may increase when all workers get more education, even if the aggregate unemployment rate is unchanged. A demographic adjustment would incorrectly interpret this as a demographic reduction in unemployment and an offsetting genuine increase, and so is misleading.

Second, the fact that more-educated workers tend to be more skilled does not imply that increases in education raise the skill level of the labor force. This point is crystallized by Spence's (1973) job-market signaling model. If education is more costly for low-ability workers, it may be used to indicate ability to future employers, and hence may be undertaken even if it conveys no direct benefit. Changes in the cost of education or the return to ability may then alter workers' education decision, but this has no effect on the skill distribution. More generally, if education is positively correlated with ability, the average college graduate today is less able than the average college graduate was twenty years ago. Again, a demographic adjustment for education overstates the true effect of a change in education.

4.1 A MODEL OF RELATIVE EDUCATIONAL ATTAINMENT

I begin by showing that if employers only care about workers' relative educational attainment, demographic adjustments are extremely misleading.[24] I extend the model in Section 3 by assuming that when workers are born, they are randomly and independently assigned an educational attainment or schooling level $s \in S \subset \mathbb{R}_+$. Let $g : S \mapsto \mathbb{R}_+$ represent the atomless density of educational attainment.

I make two simplifications to the basic model. First, the population growth rate is a constant $\dot{N}(t)/N(t) \equiv \gamma > 0$. Second, when a type-$s$ worker gets a job, her productivity is equal to s and is constant thereafter. In the language of Section 3, $\sigma = 0$. As a result, jobs last forever, and I can ignore job destruction. With this assumption, it is obviously not true that more-educated workers lose their jobs less frequently. However, it simplifies the analysis considerably: equation (8) implies that the value of a type-s job is $\mathcal{W}^*(s) \equiv s/r$. This helps elucidate the main point, which is about relative rates of job creation. The analysis here could be extended to allow $\sigma > 0$ without qualitatively changing the results.

Balanced-Growth Path. I look for a balanced growth path. On this path, the type-dependent flow matching probabilities $M(s)$ are constant over time, so the unemployment rate of type-s workers at age t is $e^{-M(s)t}$. Using the fact that age-t workers constitute a density $\gamma e^{-\gamma t}$ of the labor force at any point in time, the unemployment rate of the average type-s worker, aggregating across age cohorts, is

$$\bar{u}(s) = \frac{\gamma}{\gamma + M(s)}. \tag{13}$$

24. This is related to the point made by Blanchard and Diamond (1994) that if firms prefer to hire workers who have been unemployed for less time, a worker's relative, not absolute, unemployment duration affects her probability of finding a job. The model here is an extension of Blanchard and Diamond (1995) and Shimer (1997).

This is decreasing in the flow matching probability for obvious reasons. It is increasing in the population growth rate, since fast population growth increases the relative proportion of young workers.

It is also convenient to define the fraction of unemployed workers with schooling less than s:

$$\theta(s) = \frac{\int_0^s \tilde{u}(s')g(s')\, ds'}{U}, \tag{14}$$

where $U \equiv \int_s \tilde{u}(s)g(s)\, ds$ is the aggregate unemployment rate.

Matching Probabilities. The job search process is a generalization of the process described in Section 3.3. In this stylized model, there is only one type of job, so college graduates and high-school dropouts compete directly against each other for jobs. This represents a more realistic world of *interlinked competition,* in which college graduates compete against workers with some college education, and workers with some college education apply for other jobs, in which they compete against high-school graduates, who apply for still other jobs in which they compete against dropouts.

Job vacancies close with flow probability π. The number of other applicants is a Poisson random variable with expectation q. I assume that the most educated applicant always gets the job, and verify in the next paragraph that this ranking rule is an equilibrium.[25] Thus s is hired with flow probability

$$M(s) = \pi e^{-q[1-\theta(s)]}. \tag{15}$$

Since $\theta(s)$ is monotonically increasing, less-educated workers are hired less frequently, and so have a higher unemployment rate.

Now return to the assumption that the most educated applicant is always hired. To verify this, I must generalize the multilateral bargaining game to this environment. If there is only one applicant s for a job, she retains the entire match surplus $\mathcal{W}^*(s) - \mathcal{U}(s) - \mathcal{V}$. If there are multiple applicants, the firm cannot extract the most educated applicant's entire

25. One can also prove that the equilibrium is unique. See Shimer (1997), which also argues that for a wide range of wage determination procedures, more-productive workers will always be hired in preference to less-productive ones. This contrasts with Blanchard and Diamond (1995), which shows that with the Nash bargaining rule described in footnote 24, lower-quality workers may be ranked ahead of higher-quality ones. The important distinction is whether the presence of workers who are not hired affects the wage of workers who are hired.

value, but it instead extracts the value that it would get if it hired the second most educated applicant and retained the entire surplus. If the second highest applicant's type is s', then s enjoys a gain $[\mathcal{W}^*(s) - \mathcal{U}(s)] - [\mathcal{W}^*(s') - \mathcal{U}(s')]$. For s to be willing to outbid s', this must be positive.

Using this bargaining rule, one can prove that

$$r\mathcal{U}'(s) = M(s)[\mathcal{W}^{*\prime}(s) - \mathcal{U}'(s)]. \tag{16}$$

Intuitively, the gain to a small increase in schooling ds is that with probability $M(s + ds)$, a type $s + ds$ worker is hired, and in that event she keeps an extra bit of surplus $[\mathcal{W}^{*\prime}(s) - \mathcal{U}'(s)]\, ds$. The fact that she is hired more often, $M(s + ds) - M(s) \approx M'(s)\, ds > 0$, is a second-order effect. To the extent that $s + ds$ is hired more often than s, this happens only if she is hired in preference to some s' between s and $s + ds$. The value of the job to s' is nearly the same as the value to $s + ds$, so multilateral bargaining holds $s + ds$ nearly to her reservation value.

The return to education, $\mathcal{U}'(s)$, is strictly positive but less than the productivity increase $\mathcal{W}^{*\prime}(s) \equiv 1/r$. This implies that the match surplus $\mathcal{W}^*(s) - \mathcal{U}(s) - \mathcal{V}$ is increasing in s, so more-educated workers will always outbid less-educated ones. The ranking rule described here is in fact an equilibrium, and so the probability that a type-s worker is hired is given by (15).

Effect of Demographics. Now I can address the effect of an increase in education on unemployment. I begin by showing that the aggregate unemployment rate does not depend on demographics g. Rewriting the education-specific unemployment rate (13) as $\tilde{u}(s) = 1 - M(s)\tilde{u}(s)/\gamma(s)$, the aggregate unemployment rate satisfies

$$U = \int_s \left(1 - \frac{M(s)\tilde{u}(s)}{\gamma}\right) g(s)\, ds.$$

Expanding $M(s)$ and then using the fact that $\theta'(s) = \tilde{u}(s)g(s)/U$ yields

$$U = 1 - \frac{\pi}{\gamma}\int_s e^{-q[1-\theta(s)]}\tilde{u}(s)g(s)\, ds = 1 - \frac{\pi}{\gamma}\frac{e^{-q}}{q} U.$$

This is easily solved for U with no reference to the schooling distribution g. Since changes in schooling have no effect on unemployment, any demographic adjustment is completely spurious.

Despite this, under the maintained hypothesis, demographic changes would lead us to make significant demographic adjustments. This is because demographic changes affect disaggregate unemployment rates, a violation of the maintained hypothesis. Consider the effect of a first-order stochastic dominating "improvement" in the schooling density from g_1 to g_2.[26] That is, for all $s \in \tilde{S} \subseteq S$,

$$\int_0^s g_1(s')\, ds' > \int_0^s g_2(s')\, ds' \tag{17}$$

with a weak inequality for $s \in S \setminus \tilde{S}$. I will prove that this raises all the type-contingent unemployment rates:

$$\tilde{u}_1(s) < \tilde{u}_2(s) \qquad \text{for all} \quad s \in \tilde{S}. \tag{18}$$

I omit the similar proof that the unemployment rates are equal for other $s \in S \setminus \tilde{S}$.

According to (13), my claim is equivalent to $M_1(s) > M_2(s)$ for all s; by (15), that is equivalent to $\theta_1(s) > \theta_2(s)$ for all s. So in order to find a contradiction, assume that $\theta_1(s) \leqslant \theta_2(s)$ for some s. By (14) and the fact that aggregate unemployment U is unchanged,

$$\int_0^s \tilde{u}_1(s')g_1(s')\, ds' \leqslant \int_0^s \tilde{u}_2(s')g_2(s')\, ds'.$$

Subtract this from (17), multiply both sides by γ, and then again apply the relationship $\gamma[1 - \tilde{u}_i(s')] = M_i(s')\tilde{u}_i(s')$:

$$\int_0^s M_1(s')\tilde{u}_1(s')g_1(s')\, ds' > \int_0^s M_2(s')\tilde{u}_2(s')g_2(s')\, ds'.$$

Now replace $M_i(s')$ using (15), and evaluate the resulting integrals to find $e^{-q[1-\theta_1(s)]} > e^{-q[1-\theta_2(s)]}$ or $\theta_1(s) > \theta_2(s)$, contradicting our hypothesis and establishing (18).

Now suppose one naïvely calculated the change in the "genuine unemployment rate" as in Section 2:

$$U_{2,1}^G - U_{1,1}^G = \int_S g_1(s)[\tilde{u}_2(s) - \tilde{u}_1(s)]\, ds = \int_{\tilde{S}} g_1(s)[\tilde{u}_2(s) - \tilde{u}_1(s)]\, ds > 0,$$

where the inequality is an application of (18). One would mistakenly interpret the increase in type-specific unemployment rates as a genuine increase in unemployment. Conversely, one would find that the "demographic" decline in unemployment exactly offsets the "genuine" in-

26. I perform a comparative statics exercise here, though recent decades are probably better characterized by a transition to a new steady state. These results also hold along the transition path.

crease: $U^D_{2,1} - U^D_{1,1} = -(U^G_{2,1} - U^G_{1,1}) < 0$. But we just proved that the aggregate unemployment rate is unaffected by demographics! Thus this model illustrates a case in which demographic adjustments are theoretical nonsense and extremely misleading.

The Difference between Age and Education. An important question is why this model describes the effect of changes in the schooling distribution, but cannot be altered to describe the effect of changes in the age distribution. That is, another potential explanation for youth unemployment is that firms hire older workers in preference to younger ones. The problem with this explanation is that it has the counterfactual prediction that older workers have a shorter unemployment duration. As I discussed at the start of Section 3, youth unemployment is a consequence of the short duration of jobs, not the length of unemployment spells. Thus an argument that "relative age not absolute age" matters appears incorrect. Section 5 also provides empirical evidence that changes in schooling have very different effects than changes in the age distribution.

4.2 ENDOGENOUS EDUCATION

My analysis has glossed over one important issue: an increase in skills raises the value of a vacancy in the ranking model. In treating q as exogenous, I have implicitly assumed that labor demand is perfectly inelastic in the long run. The analytical proof of this is an algebraic mess, but the intuition is clear. The value of a firm is increasing in the skill of the second best applicant, since surplus is increasing in skill by equation (16). The expected skill level of this applicant is higher when the labor force is more skilled.

Making q endogenous and $\mathcal{V}$ exogenous through an elastic job creation condition clouds the analysis considerably. With elastic labor demand, the increase in the value of a vacancy will cause entry, which will reduce aggregate unemployment. Some of the disaggregate unemployment rates will decrease, although many will continue to increase. A demographic adjustment is still misleading, since one would conclude that there had been a sharp demographic decline in unemployment, partially offset by a genuine increase. However, it would be less misleading than the analysis in the text suggests.

Ability Bias. Fortunately, this issue is probably not very relevant, because a change in education is unlikely to have as large an effect on unemployment as Figure 14 suggests. The key reason for this, and another important difference between age and education, is that educational choice is endogenous. A number of theories predict that more-able workers will opt for more education, so some of the unemployment rate differential attributed to education is in fact due to ability differences. Since we cannot observe workers' ability in the Current Population Survey (CPS), this problem is not easily corrected.

The basic issue is illustrated with another small extension to the model. When workers are born, they are endowed exogenously with an ability a drawn from a known density $\tilde{g}$. They are then given an option to invest in a unit of education (go to college) at unit cost. If they go to college, they are left with skill $s = a + b$. If they do not make the investment (drop out of high school), $s = a$. The decision is irreversible. The rest of the model is unchanged. In particular, a worker's productivity is equal to her skill.

A college degree is worth more to a more able worker. The reason is that she will be employed more frequently, and therefore she makes better use of her educational investment. More formally, (16) implies that $\mathcal{U}'$ is increasing, since M is increasing. Convexity of $\mathcal{U}$ implies that the slope of the secant, $[\mathcal{U}(a + b) - \mathcal{U}(a)]/b$, is increasing in ability a. There is a threshold $\bar{a}$ such that workers with higher ability go to college, while those with lower ability drop out.

Since workers who go to college are more able and better educated, they are more skilled. Following the logic of the first part of this section, they are always hired in preference to dropouts, and so have a lower unemployment rate. However, this is not because they went to college. In this simple model, any college graduate is more able than all the dropouts, and so would have been hired in preference to them even if she had not gone to college. Likewise, if a high-school dropout went to college, she would still be ranked below college graduates, and so would be unemployed more frequently. More to the point, despite the fact that the returns to education are positive and that college graduates have a lower unemployment rate than dropouts, the unemployment rate of college graduates would be the same if none of them went to college; and the unemployment rate of high-school dropouts would be the same if all of them went to college.

To see why this matters, suppose there is a decrease in the cost of education or an increase in the returns to education. The threshold for attending college will fall. Since the marginal college graduate is of lower ability than the rest of the college graduates, this reduces the quality of the average college graduate. Since she is of higher ability than all the dropouts, it also reduces the quality of the average dropout. If we could observe workers' ability, we would realize that neither relative rankings nor unemployment rates conditional on ability have changed. We would conclude that there had been no change in unemployment, demographic or genuine. But since we observe education, not ability, we would see that unemployment has increased for both groups. A demographic adjustment would mistake this for a genuine increase in unemployment offset by a demographic decline. The fact that ability is unobservable leads to a further bias in the demographically adjusted unemployment rate.

Education as a Signal. Spence's (1973) signaling model affords an extreme case where education has no effect on skills, and so a change in the education distribution has no effect on the value of a vacancy.[27] People pay the high cost of a college degree and forgo years of labor income, because if they did not use this costly signal, they would be unable to obtain a high wage.

Suppose firms cannot observe applicants' skill but can observe how much they have invested in education, now a continuous choice variable $b \geqslant 0$. Applicants make wage demands, and the firm hires the applicant who it expects will yield the most profit, retaining the expected profit from the second most attractive applicant. To do this, the firm must form beliefs about each applicant's skill based on its knowledge of her education and the relationship between skill and education. The most interesting case is a separating equilibrium, in which more-skilled workers choose more education: $b = B(s)$ for some strictly increasing function B. This implies that upon observing education b, the firm can simply invert the education schedule B to calculate the implied skill. The firm believes that it has complete information, and so we can apply the analysis of the case with exogenous skill (Section 4.1). Firms always hire the most educated applicant, so the hiring rate as a function of education must satisfy an analogue of equation (15):

$$\tilde{M}(b) = \pi e^{-q[1-\theta(B^{-1}(b))]},$$

where $B^{-1}(b)$ is the skill of a worker who chooses b units of education in equilibrium.

What keeps an unskilled worker from getting a lot of education and claiming that she is skilled? If she did so, she would be hired more rapidly, and more importantly, would be able to demand higher wages. The marginal value of an additional unit of education for an unemployed worker is the analogue of equation (16):

27. Recent papers by Farber and Gibbons (1996) and especially Altonji and Pierret (1997) provide empirical support for the notion that firms discriminate statistically on the basis of education, so education indeed serves as a costly signal. However, they conclude that education raises productivity as well.

$$\tilde{\mathcal{U}}'(b) = \frac{\tilde{M}(b)}{r + \tilde{M}(b)} \mathcal{W}'(B^{-1}(b)).$$

This does not depend on a worker's skill.[28] Thus to sustain a separating equilibrium, we require that the cost of education depend on a worker's skill $C(b,s)$. More precisely, we need a single-crossing property that the marginal cost of education is lower for abler workers. Given this restriction, a simple revealed-preference argument implies that high-skilled workers will opt for more education,[29] and strictly so if C is a continuously differentiable function (Edlin and Shannon, 1996).

Consider the effect of an increase in the returns to skill (skill-biased technical change), modeled as an increase in $\mathcal{W}'(s)$. To restore a separating equilibrium, all workers must endure more of the costly signal. Otherwise, it would pay unskilled workers to imitate the education choice of skilled workers. It is clear that this change in education can have no real effect on the economy, since it is simply a change in the expenditure on the costly signal. In particular, it gives firms neither more nor less of an incentive to create jobs. Despite this, because we cannot observe the decline in skill conditional on education, we would measure it as an increase in unemployment conditional on education. This is then misinterpreted as a genuine increase in unemployment, offset by the increase in expenditures on education, the costly signal.

5. *How Much do Demographics Explain?*

I have shown that there are theoretical reasons why demographic adjustments to the unemployment rate may or may not be appropriate, but ultimately this is an empirical question. If changes in a group's labor-market share $\omega(i)$ do not affect any disaggregate unemployment rate $u(j)$, then $U^{C}_{t_1,t_0}$ is an accurate measure of what the unemployment rate would be at time t_1 if the demographics looked as they did in period t_0. The difference $U_{t_1} - U^{C}_{t_1,t_0}$ measures how much the unemployment rate increased due to demographics. Similarly, $U^{D}_{t_1,t_0}$ is an accurate measure of what the unemployment rate would be if the only changes had been demographic, so $U^{D}_{t_1,t_0} - U_{t_0}$ is another measure of how much the unem-

28. I am implicitly assuming that firms cannot punish workers who "lie" about their skill by choosing the "wrong" amount of education, $b \neq B(s)$. Firms must learn a worker's skill when they observe her productivity, and so I am imposing that contracts are incomplete.

29. Proof: Take $s > s'$ who choose education b and b' respectively. Revealed preference implies $\tilde{U}(b) - C(b, s) \geqslant \tilde{U}(b') - C(b', s)$ and $\tilde{U}(b') - C(b', s') \geqslant \tilde{U}(b) - C(b, s')$. Add these inequalities, and use the single-crossing property to prove $b \geqslant b'$.

ployment rate increased due to demographics. To the extent that $U_{t_1} - U^G_{t_1,t_0} \neq U^D_{t_1,t_0} - U_{t_0}$, the quantities U^G and U^D are poor measures of genuine and demographic unemployment.

With a little algebraic manipulation, one finds that

$$(U_{t_1} - U^G_{t_1,t_0}) - (U^D_{t_1,t_0} - U_{t_0}) = \sum_{i\in I} [\omega_{t_1}(i) - \omega_{t_0}(i)][u_{t_1}(i) - u_{t_0}(i)]$$
$$= \sum_{i\in I} [\omega_{t_1}(i) - \omega_{t_0}(i)]\{[u_{t_1}(i) - U_{t_1}] - [u_{t_0}(i) - U_{t_0}]\}.$$

If this number is positive, groups that increase their labor-market share tend to have *relative* increases in unemployment.[30] The problem with using this as a measure of the quality of demographic adjustments is that if demographic changes or relative unemployment-rate changes are small, this inner product (covariance) will be small. Therefore I construct a measure that normalizes by the size of these changes, analogous to a correlation:

$$\rho = \frac{(\vec{\omega}_{t_1} - \vec{\omega}_{t_0}) \cdot (\vec{u}_{t_1} - \vec{u}_{t_0})}{|\vec{\omega}_{t_1} - \vec{\omega}_{t_0}|\,|\vec{u}_{t_1} - \vec{u}_{t_0}|} \in [+1,1],$$

where $\vec{\omega}$ and $\vec{u}$ are the vectors of labor-market shares and disaggregate unemployment rates, and the vertical bars indicate the Euclidean length of the indicated vectors. If ρ is positive, then there is a relatively large increase in unemployment for groups that grow relatively larger. If ρ is negative, then groups that grow larger had a relative decline in unemployment. Only if $\rho = 0$ do U^G and U^D have the desired interpretations.

Table 2 shows the value of ρ obtained by dividing the population according to age and education. These estimates are fairly robust to changes in the time period. For example, changing the initial or terminal time by one year does not change the sign of any of the entries.

Education. I discuss the negative results for education, before moving on to the positive results for age. The second column of Table 2 shows that when an education category gets smaller, its members' unemployment rate increases relatively more. This reflects the disproportionate increase in unemployment for high-school dropouts (Figure 13). This is inconsistent with a model in which the reduced supply of high-school dropouts has made them a scarce resource.

However, it is consistent with the models suggested in Section 4. First,

30. This is similar to the method used by Katz and Murphy (1992) to test for relative demand shifts.

Table 2 MEASURES OF THE CORRELATION BETWEEN A GROUP'S GROWTH IN LABOR-MARKET SHARE AND ITS RELATIVE GROWTH IN UNEMPLOYMENT

	Correlation	
Year	*Age*	*Education*
1957–69	0.35	—
1969–79	0.29	−0.48
1979–89	0.44	−0.77
1989–97	−0.09	−0.24

Comparisons are March to March for the indicated years. Data availability precludes calculating this for education from 1957 to 1970; the first entry is for 1970–1979.

the ranking model tells us that the increase in a skill group's unemployment rate for a decrease in the relative ranking is larger in magnitude for groups that are unemployed more frequently (low-skilled workers). To see this, differentiate equations (13) and (15):

$$\frac{\partial \tilde{u}(s)}{\partial \theta(s)} = -q\tilde{u}(s)[1 - \tilde{u}(s)].$$

Since $\tilde{u}(s)$ is decreasing in skill, the magnitude of this partial derivative is decreasing in s as long as the skill group's unemployment rate is less than 50%. Intuitively, the most-skilled workers find jobs almost instantly, and so if they fall slightly down the ladder, they will suffer only a slight increase in unemployment (but a larger percentage increase).

A second explanation for the negative correlation is that when high-school dropouts constituted almost 40% of the labor force in 1970, lack of a high-school degree did not signal anything special. In contrast, now all but the very least-skilled, lowest-ability workers stay in high school, so that there is a precipitous drop in the quality of high-school graduates. In summary, the maintained hypothesis that changes in education have no effect on unemployment rates conditional on education is false, making a demographic adjustment for changes in education severely overestimate the actual impact of changes in education. I conclude that the logic of adjusting for changes in labor-force composition does not apply to changes in education, contrary to Summers (1986).

Age. The first column of Table 2 shows that from 1957 to 1989, age groups that increased in size had a correspondingly larger increase in unemployment. Looking back at Figure 2, this is saying that the youth unemployment rate increased disproportionately from 1957 to 1978, and then decreased disproportionately from 1978 to 1989. Since 1989, there has been a negative correlation between changes in labor-market share and changes in unemployment. However, the magnitude is much smaller, and is exactly zero for the interval 1988–1997.

The maintained hypothesis, that changes in age structure do not affect disaggregate unemployment rates, appears to be false. The baby boom caused an increase in youth unemployment, and so estimates like Δ and U^D *understate* the size of the demographic unemployment increase. The question is, by how much?

The theory in Section 3 provides some guidance. Suppose labor demand is not perfectly elastic, even in the long run. When the baby boom entered the labor market, there was a sharp increase in the number of unemployed workers, without a correspondingly large increase in the number of vacancies. From equation (12), the increase in the equilibrium unemployment–vacancy ratio q reduced the hiring rate M. However, it had an ambiguous effect on the separation threshold $\underline{x}$. To see why, recall that by (9), the separation threshold is increasing in $\mathcal{U} + \mathcal{V}$. Adding (10) and (11) implies[31]

$$r(\mathcal{U} + \mathcal{V}) = y - c + \pi(1 - qe^{-q})[\mathcal{W}(s) - \mathcal{U} - \mathcal{V}].$$

Thus if qe^{-q} is increasing in q, as is the case for $q < 1$, an increase in the relative number of unemployed workers reduces $\mathcal{U} + \mathcal{V}$ and $\underline{x}$. If $q > 1$, then qe^{-q} is decreasing in q, and so the separation threshold is increasing in q.

As the baby boom was gradually absorbed into the employed population, the unemployment–vacancy ratio slowly returned to its previous levels, restoring the old hiring rate and separation threshold. In short, the entry of the baby boom led to a temporary reduction in hiring rates, and perhaps a small change in separations. For prime-age workers entrenched in stable jobs, this had little effect. However, it led to two decades of high youth unemployment.

A reasonable alternative hypothesis is that the unemployment rate of prime-age workers, say the unemployment rate of workers aged 35–64,

31. Interpreting this expression is complicated by the fact that $\mathcal{W}(s)$ depends on $\mathcal{U} + \mathcal{V}$ through (8) and (9). Still, the joint gain from creating a new job, $\mathcal{W}(s) - \mathcal{U} - \mathcal{V}$, is decreasing in $\mathcal{U} + \mathcal{V}$.

Figure 19 GENUINE UNEMPLOYMENT, DEFINED AS THE PART OF UNEMPLOYMENT THAT CAN BE PREDICTED FROM PRIME-AGE UNEMPLOYMENT

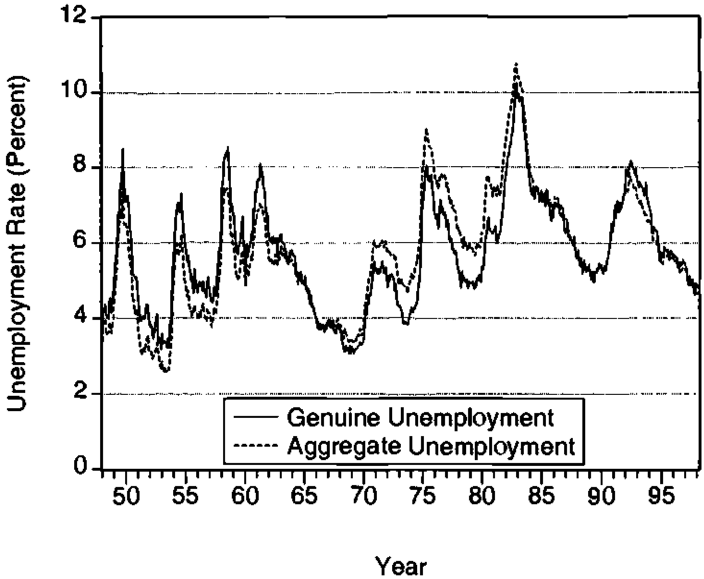

Prime-age workers are aged 35–65.

was unaffected by the baby boom.[32] The part of aggregate unemployment that can be predicted by this series is genuine. The orthogonal residual is demographic. That is, I regress aggregate unemployment on a constant and the prime-age unemployment rate, yielding

$$U_t = \underset{(0.0010)}{0.0051} + \underset{(0.0246)}{1.3531}\ u_t(\text{prime}) + \epsilon_t, \qquad R^2 = 0.834$$

(standard errors in parentheses). Genuine unemployment is 0.0051 + $1.3531u_t$(prime) (Figure 19). Notably, this is approximately the same in 1957, 1979, 1989, and 1998, as the prime-age unemployment rate has not changed over long time horizons. The demographic unemployment rate

32. If q is much smaller than 1, the prime-age unemployment rate may have dropped, because the decrease in job destruction from the reduced threshold outweighed the decrease in job creation from the reduced hiring rate. Under these conditions, the alternative hypothesis overstates the effect of the baby boom on unemployment. However, for $q \geqslant 1$, the unemployment rate of prime-age workers unambiguously increased because of the baby boom. Job destruction increased and job creation fell. Then the analysis here understates the effect of the baby boom on unemployment.

at time t is the residual $\hat{\epsilon}_t$ (Figure 20). The shape of the series looks very much like the age-adjusted series in Figures 5 and 6, but the magnitude of the adjustment is much larger. A Hodrick–Prescott (HP) filter of the residuals removes the high-frequency fluctuations and shows that the baby boom explains about a 180-basis-point increase in demographic unemployment from 1959 to 1980, and more than a 145-basis-point decline in demographic unemployment from 1980 to 1993, offset by a slight increase in the last five years.

Despite the crudeness of this measure of demographic unemployment, the result is robust. I illustrate this in three ways. First, the definition of prime age is not very important. Using the unemployment rate of workers aged 35–54 or 35+ as our measure of prime-age unemployment has no discernable effect on the demographic unemployment rate. Including younger workers aged 25–34 in our measure reduces the magnitude of the residuals by about half, as would be expected if these workers still have some characteristics of youth.

Second, omitting extreme cyclical fluctuations in unemployment does not affect the result. I regressed an HP-filtered series for aggregate unem-

Figure 20 CHANGES IN DEMOGRAPHIC UNEMPLOYMENT, DEFINED AS THE UNPREDICTABLE RESIDUAL, TOGETHER WITH A HODRICK–PRESCOTT FILTER OF THE RESIDUALS

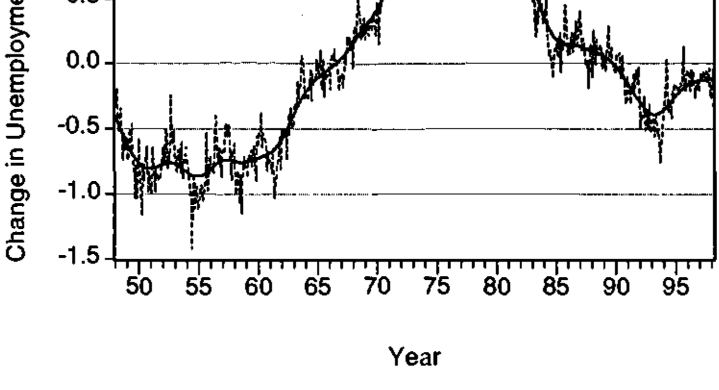

Prime-age workers are aged 35–65.

Figure 21 A SMOOTHED SERIES OF DEMOGRAPHIC UNEMPLOYMENT FITS EASILY WITHIN STAIGER, STOCK, AND WATSON'S 95% CONFIDENCE INTERVAL FOR THE NAIRU

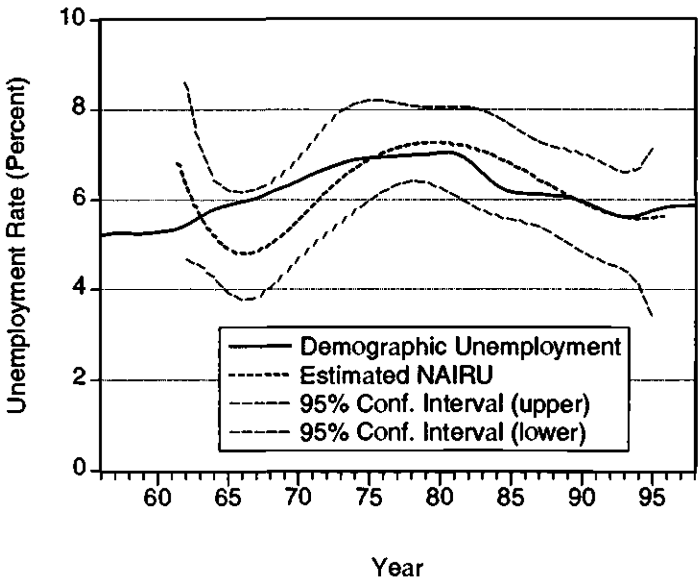

Thanks to Mark Watson for providing me with the calculation of the NAIRU.

ployment on an HP-filtered series for prime-age unemployment.[33] The residuals are nearly identical to the HP-filtered residuals in the basic regression.

Third, since female participation has increased enormously during these decades, I regressed aggregate unemployment on the unemployment rate of men aged 35–64. The peak of the HP-filtered residual increased by about 50 basis points compared to the basic regression, without any other substantial effect. According to this calculation, the demographic portion of unemployment rose by 250 basis points from 1953 to 1976, and has subsequently fallen by 200 basis points.

Any variant of these numbers is large. For example, compare my series for the demographic component of unemployment with Staiger, Stock, and Watson's (1997) nonstructurally estimated series for the NAIRU, which presumably represents some sort of equilibrium rate of unemployment. Figure 21 reproduces Staiger, Stock, and Watson's (1997) point estimate for the NAIRU and their 95% confidence interval. I then juxta-

33. Thanks to Richard Rogerson for suggesting this calculation.

pose the HP-filtered residuals from Figure 20, renormalizing the mean of the residuals to 6%, which represents long-run average unemployment. My series for the demographic component of unemployment fits easily within the confidence interval, and actually reproduces their point estimate of the NAIRU quite well, except during the Vietnam War.

I conclude that changes in the age composition of the labor force explain the bulk of the rise in unemployment during the fifties and sixties, and the subsequent decline in the eighties and nineties. However, the simplest demographic story, that a group's unemployment rate is unaffected by the size of that group, is inadequate. The entry of the baby boom into the labor force led to an increase in youth unemployment, multiplying the size of the demographic shock from a direct effect of about 80 basis points (the estimate of the change in Δ from Section 2) to nearly 200 basis points. The subsequent effect of the aging of the baby boom has been almost as important, again multiplying the size of the demographic shock from about 80 basis points to at least 120. The answer to why the U.S. unemployment rate is so much lower is that the population is so much older.

REFERENCES

Acemoglu, D., and R. Shimer. (1997). Efficient unemployment insurance. Mimeo.

Altonji, J., and C. Pierret. (1997). Employer learning and statistical discrimination. Northwestern University. Mimeo.

Blanchard, O., and P. Diamond. (1994). Ranking, unemployment duration, and wages. *Review of Economic Studies* 61(3):417–434.

———, and ———. (1995). Ranking, unemployment duration, and training costs. M.I.T. Mimeo.

Card, D. (1990). The impact of the Mariel boatlift on the Miami labor market. *Industrial and Labor Relations Review* 43(2):245–247.

Clark, K., and L. Summers. (1982). The dynamics of youth unemployment. In *The Youth Labor Market Problem: Its Nature, Causes, and Consequences*, R. Freeman and D. Wise (eds.). Chicago: University of Chicago Press, pp. 199–235.

Council of Economic Advisors. (1997). *Economic Report of the President*. Washington: U.S. Government Printing Office.

Diamond, P., and E. Maskin. (1979). An equilibrium analysis of search and breach of contract, I: Steady states. *Bell Journal of Economics* 10:282–316.

———, and ———. (1981). An equilibrium analysis of search and breach of contract, II: A non-steady state example. *Journal of Economic Theory* 25:165–195.

Edlin, A., and C. Shannon. (1996). Strict monotonicity in comparative statics. University of California, Berkeley. Mimeo.

Farber, H., and R. Gibbons. (1996). Learning and wage dynamics. *Quart. Journal of Economics*, 111(4):1007–1047.

Gordon, R. (1982). Inflation, flexible exchange rates, and the natural rate of unemployment. In *Workers, Jobs and Inflation*, M. Baily (ed.). Washington: Brookings Institute, pp. 89–152.

———. (1997). The time-varying NAIRU and its implications for economic policy. *Journal of Economic Perspectives,* 11(1):11–32.

Greenspan, A. (1997). Monetary policy testimony and report to the congress. Humphrey–Hawkins Testimony, July 22.

Hall, R. (1970). Why is the unemployment rate so high at full employment. *Brookings Papers on Economic Activity,* 33:369–402.

Juhn, C., K. M. Murphy, and R. Topel. (1991). Why has the natural rate of unemployment increased over time. *Brookings Papers on Economic Activity,* 22:75–142.

Karatzas, I., and S. Shreve. (1991). *Brownian Motion and Stochastic Calculus,* 2nd ed. Graduate Texts in Mathematics. New York: Springer-Verlag.

Katz, L., and K. Murphy. (1992). Changes in relative wages, 1963–1987: Supply and demand factors. *Quarterly Journal of Economics,* 107(1):35–78.

Krueger, A., and J.-S. Pischke. (1997). Observations and conjectures on the U.S. employment miracle. Princeton University. Mimeo.

Mortensen, D., and C. Pissarides. (1994). Job creation and job destruction in the theory of unemployment. *Review of Economic Studies* 61:397–415.

Perry, G. (1970). Changing labor markets and inflation. *Brookings Papers on Economic Activity* 33:411–441.

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

Riley, J., and W. Samuelson. (1981). Optimal auctions. *American Economic Review,* 71(3):381–392.

Shimer, R. (1997). Do good guys come in first? How wage determination affects the ranking of job applicants. Mimeo.

Spence, M. (1973). Job market signaling. *Quarterly Journal of Economics* 87(3):355–374.

Staiger, D., J. Stock, and M. Watson. (1997). The NAIRU, unemployment, and monetary policy. *Journal of Economic Perspectives,* 11(1):33–49.

Summers, L. (1986). Why is the unemployment rate so very high near full employment? *Brookings Papers on Economic Activity* 22:339–383.

Ventura, S., J. Martin, T. Mathews, and S. Clarke. (1996). Advance report of final natality statistics, 1994. *Monthly Vital Statistics Report* 44(11 Supplement):1–87.

Comment

RICHARD ROGERSON
University of Pennsylvania

The recent decrease in the U.S. unemployment rate to levels not seen in almost three decades has attracted much attention from policymakers and media analysts. Following nearly two decades of secular increases in the unemployment rate, many are inclined to view this recent development as evidence of some fundamental improvement in the state of the U.S. economy, due either to improvements in productivity, good policymaking, or both. Shimer deflates this view by arguing that much of the secular

Table 1 AGGREGATE AND DISAGGREGATED U.S. UNEMPLOYMENT

	Unemployment Rate (%)			
Year	*Total*	*Male 16–19*	*Male 35–44*	*Male 55–64*
1952	3.0	8.9	1.9	2.4
1957	4.3	12.4	2.8	3.5
1960	5.5	15.3	3.8	4.6
1969	3.5	11.4	1.5	1.8
1974	5.6	15.6	2.6	2.6
1979	5.8	15.9	2.9	2.7
1989	5.3	15.9	3.7	3.5
1997	4.9	16.9	3.6	3.1

change in unemployment over the last three decades can be attributed to the simple process of the aging of the baby-boom generation.

I am largely persuaded by the author's main point. I will begin my comments by offering a slightly different presentation of the facts. This serves two purposes: first, it provides direct support for the important role that compositional effects have had on the aggregate unemployment rate, and second, it highlights another important secular trend in the labor market that recent discussions seem to gloss over.

Consider Table 1. The years for this table were chosen to correspond to similar points of the business cycle, with the exception of 1997, which is simply the most recent (complete) year in the current expansion. The first column, which displays the aggregate unemployment rate, presents the following picture. After World War II there was a secular increase in unemployment, followed by a substantial drop during the 1960s. This was followed by another period of secular increase, which continued into the 1980s. In the most recent period the unemployment rate has declined significantly. It is natural to ask what lies behind these low-frequency waves in the aggregate unemployment rate. This paper is specifically concerned with what has given rise to the most recent wave.

Consider first, however, the data in the next three columns of Table 1, which show unemployment rates for males of three different age groups. It is quite remarkable that while there is some evidence of a decrease in unemployment for some of these groups in the 1990s, each of them shows a substantial increase relative to 1979—despite the fact that the aggregate unemployment rate has increased by almost a full percentage point. In particular, the first wave seems quite different from the second wave in that the 1960s witnessed not only a decrease in aggregate unemployment but also a decrease in the unemployment rate for each of the age groups displayed.

Table 2 pursues this issue a bit further by showing some additional information for males aged 35–44, which is traditionally the group with the strongest attachment to the labor force and hence of particular interest.

It is remarkable that if one looks at the employment-to-population ratio for this group, there is no secular decline in the post-1969 period. The basic message one gets from this table is that there is no puzzle about declining "unemployment"—rather, the puzzle suggested by these data is how to account for the uninterrupted secular decrease in *employment* among prime-aged males since 1969. Though there has been a slight decrease in the unemployment rate for this group since 1989, it has been accompanied by a substantial drop in both the participation rate and the employment-to-population ratio. Again, the second wave appears very different from the first wave, as the 1960s brought a substantial increase in the employment-to-population ratio of prime-aged males. An important question that I do not have space to go into here concerns the interpretation of the decreased participation rate. I think an argument can be made that the nonparticipating individuals are largely those with poor labor-market opportunities who have stopped searching, and hence are best viewed as discouraged workers.

Finally, Table 3 provides some information on international comparisons using data from the OECD. It is interesting that, in contrast to the picture painted by the aggregate unemployment rate numbers, the U.S. and Europe are basically experiencing a similar phenomenon in terms of employment ratios, though it is clearly more intense in Europe.

Table 2 U.S. LABOR-MARKET DATA FOR MALES AGED 35–44

Year	*Unemployment Rate (%)*	*Participation Rate (%)*	*Employment/ Population (%)*
1952	1.9	97.8	96.0
1957	2.8	97.9	95.1
1960	3.8	97.7	94.0
1969	1.5	96.9	95.5
1974	2.6	96.0	93.5
1979	2.9	95.8	93.0
1989	3.7	95.3	92.2
1997	3.6	92.6	88.9

The data surveyed above lead one to ask three questions:

1. Why has there been a secular decline in employment-to-population ratios for prime-age males in the post-1969 period?
2. Why has this decrease been so much larger in some countries than others?
3. Why does the aggregate unemployment rate behave so differently after 1969 than do the age-specific unemployment rates?

I think it is fair to say that although much work has addressed the first two questions, there is presently no clear consensus on their answers. The present paper is effectively focused on the third question, and having seen the various facts displayed above, it is not surprising that the entry of the baby boom into the labor force and its subsequent aging figure prominently in the answer. The main contribution of this paper, then, is to quantify the effect of the baby boom on the evolution of the unemployment rate over the last three decades.

Young people have much higher unemployment rates than do old people. In trying to assess the dynamic impact of a large cohort on aggregate unemployment, there are two channels to consider. The first is a direct effect—even if the presence of a large cohort has no effect on the labor-market history of any individual, there will be a direct effect on aggregate unemployment due to the change in the age composition of the labor force. The second channel involves potential general equilibrium effects, whereby the presence of a large cohort may affect the labor-market history of each individual, in particular their unemployment experiences. For example, a large group of young workers entering the labor force may plausibly be expected to lead to higher unemployment rates among younger workers as they compete with more people for

Table 3 INTERNATIONAL COMPARISONS FOR MALES 25–54

Year	United States	Germany	France
	Unemployment Rate (%)		
1979	3.4	2.0	3.2
1995	4.4	6.3	8.8
	Employment/Population(%)		
1979	91.2	93.0	93.3
1995	87.6	87.3	86.6

entry-level positions. For other age groups the effects may be more difficult to assess, even at a qualitative level. Intuitively, it would seem to depend on the extent to which older and younger workers are substitutes or complements.

How large is the direct effect? The basic idea is that there is an age-specific unemployment rate for each age, and that these can be used in combination with the age distribution of the labor force to determine the magnitude of the direct effect as the baby-boom cohort enters the labor force and ages. One issue however, is how to determine these age-specific unemployment rates. A quick look at the time-series data reveals that the gap between unemployment of young and old workers is not constant and in fact is subject to a substantial amount of variation. To deal with this issue Shimer carries out a couple of different exercises, one using the unemployment rates from some arbitrary benchmark year, and another using a chain-weighted series. The point estimate that he produces is roughly 0.8%, i.e., the direct effect of the baby boom was to cause an increase of the unemployment rate by 0.8% prior to 1979, and then an 0.8% decrease subsequently.

Another approach is to assess the range of magnitudes for the direct effect, rather than trying to produce a single point estimate. For example, over the period 1950–1980, which includes the period in which the baby-boom cohort enters the labor force, the minimum difference between the unemployment rates of those 16–19 and those 20 and older was 5 percentage points, which obtained in 1953, and the maximum difference was 12.5 percentage points, which obtained in 1976. The increase in the share of 16–19-year-olds in the labor force was approximately 8%. It follows that the minimum direct effect is about 0.4%, while the maximum was about 1%. Shimer's preferred point estimate of 0.8% is not that far from the midpoint of this interval. Redoing this calculation with three age groups (16–19, 20–24, and 25 and older) would also increase the effect somewhat. Using the years 1953 and 1976 as above, and allowing for an increase of 4% in the labor-force share of 20–24-year-olds, the range of estimates for the direct effect is between 0.5% and 1.4%.

At this point I think Shimer has effectively made his main point—basically that the changing age distribution between 1960 and the present may have led to a secular increase and subsequent decrease on the order of 1 percentage point. From Table 1, it is clear that this is a quantitatively significant development.

Shimer goes on to ask how significant the indirect effect may be, in an attempt to give a more definitive estimate for the effect of the baby boom on aggregate unemployment. Here, however, I do not find the analysis very persuasive. The basic objective is to ascertain to what extent the entry of the baby boom into the labor force may have had an additional effect on total unemployment by increasing the unemployment rate of young workers. The main difficulty here is the typical one for this type of exercise: the baby boom entering the labor force is only one of many changes that have potentially affected the unemployment experiences of various labor-market groups, and it is very difficult to sort out the contributions of various shocks by running a single regression.

Shimer's methodology is to assume that the prime-age unemployment rate was unaffected by the baby boom and that the effect of any factors other than the baby boom on aggregate unemployment can be captured by their effect on prime-age male unemployment. It follows that any part of unemployment that is not "explained" in a regression sense by prime-age unemployment can be attributed to the baby boom.

Unfortunately, this methodology seems inadequate. First, as indicated by Table 2, since the unemployment rate for prime-aged males is not necessarily a good indicator of secular patterns in the labor-market experience of these individuals, it is not clear exactly what this unemployment rate is proxying for in terms of secular developments. Second, without naming the other factors that are being considered, how can we judge the reasonableness of the assumption that the baby boom can be assumed to account for the entire residual in the regression equation? Changes in length of schooling, fertility, and marriage rates would all presumably affect the relationship between prime-age unemployment and total unemployment, but these effects should not be attributed to the size of the baby-boom cohort. Additionally, the literature on wage inequality has documented that there were substantial changes in the return to experience since 1970, which may also affect the relative employment prospects of young and prime-aged individuals.

The paper also raises some additional issues of interest, though I will comment on only one of them. It concerns whether the distribution of educational attainment should be viewed analogously to the age distribution from the perspective of accounting for changes in unemployment. Shimer argues not, for the reason that changes in the distribution of education have direct effects that may be largely offsetting to the indirect effects induced by compositional changes. While I don't disagree with the conclusion that he draws, I feel that the analysis ignores a key difference between changes in the age distribution and changes in the educational-attainment distribution: Whereas it seems appropriate to view the baby boom as an exogenous event, this does not seem to be appropriate for thinking about changes in educational attainment. Rather, changes in educational attainment are themselves a response to changes in the economy which affect the incentive for individuals to acquire education. It seems less interesting to ask what would have happened to aggregate unemployment if educational attainment had changed but there had been no change in the factors which affect the choice of educational attainment.

In summary, I think that Shimer has argued persuasively that the changing age composition of the U.S. labor force associated with the baby-boom cohort has had a significant impact on the evolution of the U.S. aggregate unemployment rate. In particular, much of the apparent decline since 1979 may be attributed to this factor. However, while much attention is currently focused on the recent decreases in aggregate unemployment, it is at the expense of attention that should be focused on accounting for what appears to be an underlying secular increase in "true" unemployment since 1969.

Comment

ROBERT TOPEL
University of Chicago and NBER

Robert Shimer has written an intriguing paper on an important topic. While I don't agree with all that he does, I think the bottom line of his paper is largely correct. I also think that the problem he studies has broader importance than he attributes to it.

Unemployment rates are interesting to economists and the general public because they are a sort of summary statistic for the current state of the labor market, and even for the success or failure of economic policies. Given that they are used in these ways, it would be nice if they were comparable over time. Then policymakers who congratulate themselves for being in office when unemployment reached "its lowest level since 19-whatever" would at least be saying something. Part of Shimer's point is that unemployment rates are not comparable over time. The basis of his conclusion is that demographics change, so that unemployment rises and falls as high-unemployment groups become more or less prominent in the labor force. In the 1990s, he points out, high-unemployment youths are a smaller portion of the labor force than in the 1970s. This means that aggregate unemployment will be lower, so that aggregate unemployment does not measure the same thing today as it did 20 years ago. Indeed, Shimer provides strong evidence that changes in labor-force shares of young workers—driven by the baby boom and subsequent baby bust—explain most of the low-frequency changes in measured unemployment from the 1960s to today. On this point, which is his main conclusion, I think he is right.

As an empirical matter Shimer also argues—again I think correctly—that other changes in demographic characteristics of the labor force do not have the same impact as age. Specifically, he points out that secular increases in the average educational attainment of the workforce do not reduce unemployment, as one might expect from the negative cross-sectional relationship between education and unemployment.

Should we be surprised by these findings? Suppose that in 1979 we had asked an economist the following two questions. (Question 1 is easy.)

1. With the aging of the baby-boom generation, the proportion of youths in the labor force will fall over the next two decades. Youths are known to have higher unemployment. Other things being equal, will this demographic change cause aggregate unemployment to be lower in 1998 than it is today?
2. Due to increased educational attainment in successive cohorts, the workforce in 1998 will be more educated than it is today. More-educated workers are known to have lower unemployment rates. Will this demographic change cause aggregate unemployment to be lower in 1998 than it is today?

I think that the right answers, in 1979 and today, are "yes" to question 1 and "probably not" to question 2. Those are the answers provided by Shimer as well, though we differ over the reasons why.

Why "yes" to the first question? In addition to being less skilled, young workers are known to have weaker labor-force attachments and high turnover rates, the latter because of the matching process that characterizes early careers. These features of young workers imply higher unemployment rates for them, as the sampled jobs of early careers are less likely to survive. Most importantly, these facts about youth won't change much simply because there are more or fewer youths at a particular point in time. So the cross-sectional relationship between age and unemployment is likely to hold in time series too. More youths as a proportion of the labor force means higher aggregate unemployment.

Why "probably not" to the second question? An initial reaction might be that educated workers have lower unemployment, so an increase in average education should cause aggregate unemployment to fall. There is some merit to that response, at least in the short run. But reflection on the past would give an empirically minded economist pause. Surely the labor force in 1979 is much more educated than it was in, say, 1879. Yet there is no evidence of a century-long decline in unemployment. Looking across countries, average unemployment rates bear no obvious relationship to schooling levels. There is clearly something different about education, as a demographic indicator of unemployment, as compared to age.

In Shimer's analysis, this difference follows from a search technology in which employers care about relative schooling levels of workers. When multiple workers apply for the same job, the one with more schooling gets it, after engaging in a multilateral bargaining game that determines the division of surplus. Giving everyone more schooling won't change relative rankings, so equilibrium unemployment isn't affected. To me, the real reason seems much simpler than that. Indeed, I think the broad facts about unemployment documented here would be true in an economy of independent farmers who bargain with no one. These facts have vastly more to do with labor-supply behavior than with search and bargaining games.

Why is education different than age? I don't believe it's because employers care about relative educational attainments, nor do I think that signaling models of education have much credibility in this context, especially over long periods of time. Indeed, I believe it's because rising education—as a measure of human capital investment—*is* an indicator of long-term productivity growth. In most models that fit broad macroeconomic facts, productivity growth will have no effect on the natural rate of unemployment, because wages and reservation wages rise together.

A touch of formalism helps. Let the utility of consumption and leisure be from the class in which income and substitution effects are offsetting, for example Cobb–Douglas. Then long-run desired labor supply is independent of the wage, and the reservation wage is

$$\ln W_R = a + \ln C - \ln L, \tag{1}$$

where C is consumption and L is nonworking time. With L constant, the reservation wage is proportional to consumption. Over long periods of rising productivity, the elasticity of consumption with respect to the wage is unity, so the wage and value of nonworking time move in unison. All aspects of labor supply, including unemployment, are constant over time. If education is human capital, and human capital is productive, then a rightward shift in the distribution of education should have no impact on the natural rate.

This is only part of the story, however. One of the most prominent facts about unemployment is that *at any point in time* unemployment is higher among the least productive. Another way of putting this is that while the long-run relationship between productivity and the natural rate is zero, the cross-sectional relationship is always negative. How do we reconcile these seemingly contradictory facts?

Figure 1 OLF, UNEMPLOYMENT, AND NONEMPLOYMENT RATES: 1967–1994

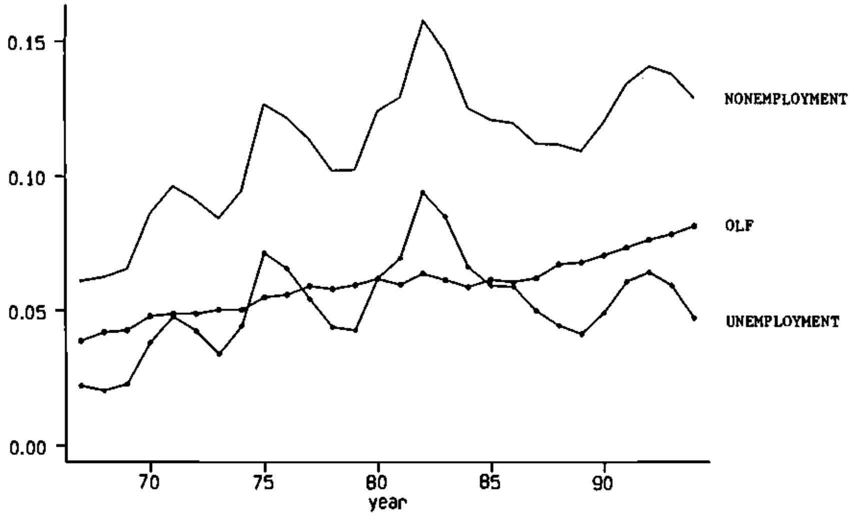

In the context of labor supply and desired consumption in equation (1), it must be the case that the cross-sectional elasticity of consumption with respect to the wage is smaller than 1.0, though the aggregate time-series elasticity is 1.0. This will be satisfied in the cross section so long as assets and other forms of nonwage income (including family and other sources of support) do not vary proportionally with wages. For example, suppose that high-school graduates earn 50% less than college graduates. If the gap in nonwage incomes is smaller than 50%, the condition is satisfied. Then in any cross section, college graduates will have lower unemployment (higher labor supply), but growth that raises wage and nonwage incomes by proportional amounts will leave the labor supply of each group unchanged.

I strongly suspect that this condition is satisfied. Juhn, Murphy, and Topel (1991) provide some evidence related to this point. They find that during the 1970s and 1980s, while real wages and employment of low-skill workers were falling, the household incomes of those workers were nearly constant. In other words, low-skill workers had sources of nonwage income that did not fall in proportion to wages, so their desired labor supply declined. This underlies our contention that declining labor-market opportunities for less-skilled workers—reflected in declining wages—explained rising unemployment and nonparticipation after 1970.

I agree with Shimer's conclusion that today's low measured unemployment rate is not directly comparable to unemployment rates of the past, and that labor-market conditions are not as strong as such comparisons might suggest. Shimer rests much of his conclusion on simple demographics—the aging of the baby-boom generation reduced unemployment—but I think the conclusion is broader and stronger. In my view, an oddity of Shimer's analysis is that it focuses so narrowly on unemployment. Labor-force participation and changes in labor demand conditions are conspicuous by their absence. For prime-aged men, the fact is that while unemployment rates have been falling, their labor-force participation rates have continued to decline, *especially among the least skilled.*

Murphy and Topel (1997) argue that long-term changes in labor demands have changed the returns to work, most notably among the least skilled. We argue that this decline in labor-market opportunities drove declining employment rates among the least skilled, which raised both unemployment and the rate of labor-force withdrawal. Since unemployment data alone exclude workers who have withdrawn from the labor force for market-driven reasons, they miss a key part of the story.

How big are these effects? Figures 1 and 2 show the data. Among workers with 1–30 years of potential experience, the nonparticipation

Figure 2 OLF RATE BY INTERVALS OF THE WAGE DISTRIBUTION

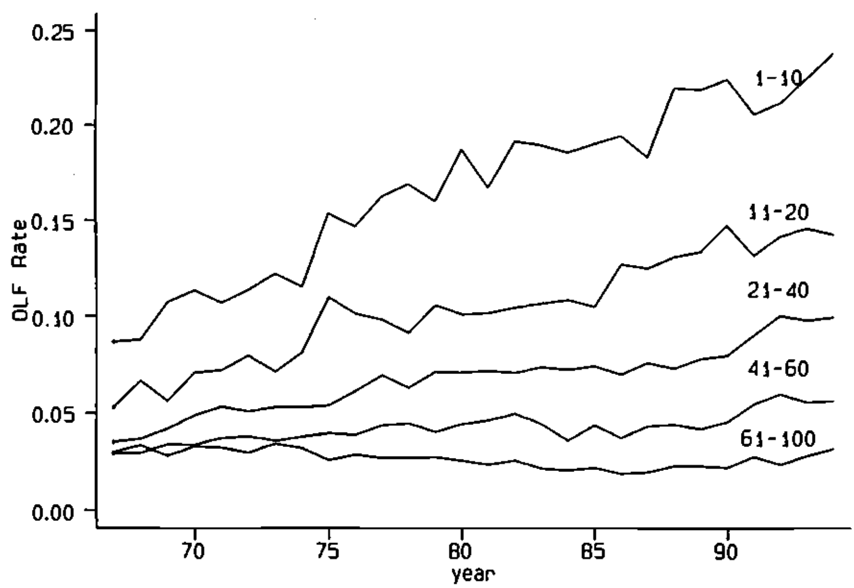

rate rose by over 2 percentage points between 1979 (Shimer's benchmark year) and 1994 (the last year for which I did the calculations). This increase was fairly steady over time. Among workers in the bottom decile of the wage distribution the increase was much more dramatic, from about 16% in 1979 to 24% in 1994 (see Figure 2).

Why is this important? As I noted above, many economists and almost all of the general public think of the unemployment rate as a summary measure of labor-market conditions, and of even the success or failure of economic policy. If these labor-force withdrawals are driven by long-term changes in labor demand, as the evidence strongly suggests, then measured *unemployment* misses a large part of the market-driven reason for declining *employment*. As Robert Shimer argues from his analysis, this means that unemployment rates of the 1990s are different animals than the rates of earlier decades. But the reasons go well beyond simple demographics. I think the evidence is strong that unemployment rates have become poorer measures of the joblessness that is caused by labor demand conditions. How much poorer? To compare current conditions with 1979, I think a useful rule of thumb might be to add two percentage points to the current unemployment rate, at least for men.

This leads to a final point. If labor-force withdrawals mean that "true" unemployment is higher than measured unemployment, concerns about accelerating inflation at today's "low" unemployment rates may be misplaced. While I'm no fan of the Phillips curve, and I've never come close to estimating a NAIRU, the distinction I've made here may be useful to those who are, and do.

REFERENCES

Juhn, C., K. M. Murphy, and R. Topel. (1991). Why has the natural rate of unemployment increased over time? *Brookings Papers on Economic Activity*, 2:75–142.

Murphy, K. M., and R. Topel. (1997). Unemployment and nonemployment. *American Economic Review* 87:295–300.

Discussion

There was considerable interest in whether Shimer's results for the United States might generalize to other countries: Pierre-Olivier Gourinchas wondered how European unemployment might look if corrected for baby-boom effects in those countries. Ben Bernanke suggested that a systematic comparison across industrialized countries of age effects on aggregate unemployment would be a useful extension of this research; he noted that Japan, with its dramatic demographic changes in the postwar period, is a particularly interesting case. Shimer expressed skepticism about the potential of demographic effects to explain much of the European experience in particular, noting the importance there of labor-market regulation and other structural factors.

John Shea thought that it would be interesting to analyze some conventional empirical relationships, such as the Beveridge curve and the Phillips curve, using demographically adjusted unemployment rates. For example, how do vacancy rates covary with the adjusted unemployment rates? In a Phillips-curve context, are changes in unemployment arising from demographic factors related to inflation in a way different from other changes in unemployment? On a related point, Herschel Grossman noted that the cyclical peaks considered in the paper may not correspond to peaks of the inflation cycle; it would be interesting to compare values of group-specific unemployment rates at inflation peaks.

Gregory Mankiw suggested that Shimer's use of the prime-age unemployment rate in his attempt to assess the importance of indirect demographic effects (such as the effects of generational crowding) may not be appropriate. He suggested that, instead, age-group-specific unemployment rates be regressed on a variable such as the share of young people in the population, with the fitted values being used to estimate the trends arising from demographic factors. Mankiw also expressed surprise that such a straightforward explanation of the decline in U.S. unemployment had not appeared previously in the literature. Shimer credited the *Economic Report of the President* with having noted earlier the possible link between demographic changes and the unemployment decline.

There was general agreement with Robert Topel's point, made in his formal comment, that it is as important to explain changes in employment and participation rates as changes in unemployment rates. Shimer responded that simultaneous modeling of employment, participation, and unemployment was beyond the scope of the current paper, but indicated that demographic changes, including possibly changes in educational attainment, might well be useful in explaining the trends both in employment and in participation as well as in unemployment.