

Changes in Unemployment Duration and Labor Force Attachment*

Katharine G. Abraham
Bureau of Labor Statistics

Robert Shimer
Princeton University and NBER

August 29, 2001

Abstract

This paper accounts for the observed increase in unemployment duration relative to the unemployment rate in the U.S. over the past thirty years, typified by the record low level of short-term unemployment. We show that part of the increase is due to changes in how duration is measured, a consequence of the 1994 Current Population Survey redesign. Another part is due to the passage of the baby boomers into their prime working years. After accounting for these shifts, most of the remaining increase in unemployment duration relative to the unemployment rate is concentrated among women, whose unemployment rate has fallen sharply in the last two decades while their unemployment duration has increased. Using labor market transition data, we show that this is a consequence of the increase in women's labor force attachment.

*We are grateful to Giuseppe Bertola, Robert Solow, David Weiman, and participants in the Sustainable Employment Initiative conference for their comments and to Fran Horvath, Randy Ilg, Rowena Johnson, Bob McIntire, and Anne Polivka at the Bureau of Labor Statistics (BLS) for their help compiling the data. Ron Tucker of the Census Bureau and Clyde Tucker of the BLS provided useful information concerning the Current Population Survey redesign. Joydeep Roy provided valuable research assistance. Shimer acknowledges financial support from National Science Foundation grant SES-0079345 and the hospitality of the University of Chicago while part of this paper was written. Please address correspondence to shimer@princeton.edu.

1 Introduction

Between the late 1960s and the late 1970s, the U.S. unemployment rate trended consistently upwards. The unemployment rate attained at successive cyclic peaks increased from 3.4 percent in September 1968 to 4.6 percent in October 1973, 5.6 percent in May 1979, and 7.2 percent in April 1981 (see Figure 1). In November 1982, the worst point of the 1981–1982 recession, the unemployment rate reached a postwar high of 10.8 percent. Since that time, however, unemployment has fallen back towards the levels of the late 1960s. By March 1989, the unemployment rate had dropped to 5.0 percent, and, from October 1999 through the end of 2000, it hovered near 4.0 percent.

During most business cycles, there is a strong correlation between the unemployment rate and the average length of an unemployment spell, so, not surprisingly, unemployment durations increased steadily during the 1970s. But, subsequently, unemployment durations have not fallen nearly as much as the decline in the unemployment rate might have led one to expect. For example, the top left panel of Figure 2 shows that, historically, the mean duration of in-progress unemployment spells,¹ the solid line, closely tracked the unemployment rate, the dashed line (depicted on a different scale). As the unemployment rate increased from business-cycle peak to business-cycle peak during the 1970s, mean unemployment duration followed suit, rising from 8 to 11 weeks during the decade. But, while the unemployment rate has subsequently declined, mean unemployment duration has not fallen commensurately. As a result, the mean duration of an unemployment spell was fifty percent longer at the end of the 1990s than at the end of the 1960s. The remaining panels of Figure 2 explore how the unemployment-duration distribution has shifted over time.

¹Throughout this paper we use data on the duration of in-progress unemployment spells rather than on the duration of completed spells.

The short-term unemployment rate, defined as the fraction of the labor force that has been unemployed for 0 to 4 weeks, is at its lowest sustained rate since the early 1950s, forty percent lower than the levels that prevailed from 1975 to 1985 and significantly lower than the levels that one would expect from the historical relation between the short-term and the aggregate unemployment rates. The long-term (15- to 26-week) unemployment rate has also declined steadily during the last two expansions, although it remains slightly high by historical standards. But most of the action comes from the increase in the very long-term (more than 26-week) unemployment rate. In 1969, persons unemployed half a year or more accounted for just 0.16 percent of the labor force; at the peak of each of the last three expansions, three times as large a share of the labor force had similarly long spells of unemployment.²

This paper seeks to explain the breakdown in the historical relation between the aggregate unemployment rate and measures of unemployment duration. There are several reasons why understanding this change is important. First, if workers are risk averse and labor-income risk cannot be insured, welfare is lower when unemployment duration is longer, holding the unemployment rate constant. Essentially, longer unemployment durations load more uninsurable risk onto individuals. Lucky workers never lose their jobs, so these workers have a lower marginal utility of income than do their less fortunate peers, who suffer long spells of joblessness. This suggests that ignoring unemployment durations and focusing only on the low unemployment rate at the end of the 1990s may overstate the strength of the U.S. economy. Similarly, the long-term unemployed may lose skills or contact with the labor market, generating hysteresis (Blanchard and Summers 1987), and blocking the possibility of further reductions in the unemployment rate. This will also limit the downward pressure

²Summers (1986) and Katz and Krueger (1999) have also commented on the increase in long-term unemployment as a share of total unemployment.

that the long-term unemployed place on wages, with the result that, in an economy with many long-term unemployed, the Phillips curve may lie to the right of that in an economy with shorter unemployment durations. On the other hand, the existence of long unemployment durations suggests a positive role for government intervention in the labor market. For example, retraining programs might help in reintroducing the long-term unemployed to regular labor-market activity, thereby reducing unemployment duration without increasing the incidence of unemployment. Section 2 discusses these issues in more detail.

We then turn to our main task, trying to understand the source of the recent increase in unemployment duration. Each subsequent section focuses on one important factor — Section 3 on measurement issues, Section 4 on the aging of the baby boom, and Section 5 on labor-force attachment.

Measurement Issues. In 1994, the Bureau of Labor Statistics (BLS) redesigned the Current Population Survey (CPS), the instrument used to measure the unemployment rate and unemployment duration. There are reasons to suspect that this redesign reduced the measured incidence of short-term unemployment. We conclude that the redesign is likely to explain a decline of about half a percentage point in the short-term (0- to 4-week) unemployment rate and an increase of about a half week in mean unemployment duration, although it does not explain the increase in long-term unemployment.

Aging of the Baby Boom. The mass of the U.S. population has shifted into age groups that report longer unemployment spells. If one assumes that this demographic shift has affected only the share of workers in different age groups, not unemployment duration conditional on age, then the aging of the baby boom can explain another

half percentage point decline in the short-term unemployment rate between 1980 and 2000 and another half-week increase in mean unemployment duration over the same time period. It does little, however, to explain the rise in long-term unemployment.

Labor-Force Attachment. Most of the secular increase in unemployment duration is accounted for by women, whose unemployment duration has risen to approach the male level. At the same time, women's unemployment rate has declined towards the men's rate since 1980. This pattern suggests the possibility that increases in women's attachment to the labor force may be responsible for the aggregate trends. Labor-force attachment affects unemployment incidence and unemployment duration in at least two ways. Workers who have a stronger attachment to the labor force tend to stay unemployed when they lose a job, rather than dropping out of the labor force. This raises both the unemployment rate and unemployment duration. On the other hand, because these workers are unlikely to quit their jobs and exit the labor force, they can build up stable employment relations with a minimal incidence of unemployment. This reduces the unemployment rate and may also raise unemployment duration by reducing the pool of workers who chronically transition out of the labor force from unemployment.

We examine evidence on changes in the transition rates of workers across labor-market states (employment, unemployment, and out of the labor force) in order to explore the role that women's increasing labor-force attachment has played in the rise in unemployment duration relative to the unemployment rate. We find that the declining exit rate of employed women from the labor force is quantitatively important for explaining both the decline in women's overall unemployment rate and the increase in their unemployment duration.

To summarize our findings, changes in measurement and changes in demographics

each explain part of the increase in unemployment duration relative to the unemployment rate during recent years, especially for men. For women, an important part of the explanation is the increase in their attachment to the labor market, which has reduced their unemployment rate while raising their unemployment duration.

2 Theory

2.1 Unemployment Duration and Risk

In representative-agent models of the business cycle, labor-market activity is summarized by the number of hours worked. Lucas (1987) has forcefully argued that, in such an economy, there is little value to reducing the variance of output.³ But representative-agent models may significantly understate the cost of recessions if the burden falls on a small subset of the population. According to this logic, the unemployment rate is an important measure of economic activity in part because it highlights the distributional consequences of recessions. For example, Gruber (1997) and Browning and Crossley (2001) have shown that workers with low asset holdings substantially reduce their consumption following job loss, although the effect is mitigated by unemployment insurance.

Unemployment duration, not just the unemployment rate, is an important determinant of the distributional consequences of recessions. At one extreme, if unemployment spells are very brief, workers can easily use a small stock of savings to smooth

³Lucas asks how much consumption an individual would be willing to give up in order to eliminate any variation of output around its deterministic trend. This presumes, not that recessions are periods when output is below trend and that expansions are periods when output is on trend, but that business cycle fluctuations are symmetrical deviations around a trend. Lucas's argument is based on assumptions about individual preferences, in particular, about the extent of risk-aversion. Alvarez and Jermann (2000) reach a similar conclusion using evidence from asset prices without having to impose strong restrictions on preferences. They conclude that rational agents would forgo about half a percent of their consumption in order to eliminate business-cycle-frequency fluctuations.

consumption across these spells, and hours worked will be a good measure of economic activity. At the other extreme, if unemployment spells never end, no stock of savings will be large enough to allow for consumption smoothing, and average measures will be inadequate for describing individual activity.

To get at this idea more formally, we consider an economy inhabited by rational workers who live for two periods.⁴ Workers inelastically supply labor and seek to maximize their expected utility from consumption $U_1(c_1) + U_2(c_2)$, where $c_1 \geq 0$ and $c_2 \geq 0$ are consumption in the first and second period of life. We assume that the utility function is concave, so workers are risk-averse. In each period, a fraction u of them are unemployed and earn an unemployment benefit b , and the remaining $1 - u$ are employed at a fixed wage $w > b$. If a worker is unemployed in the first period of her life, there is a probability λ_{ue} that she will be employed in the second period; and, symmetrically, a worker who is employed in the first period may be unemployed in the second period with probability λ_{eu} . λ_{ue} is inversely related to the mean duration of unemployment (which can vary between 1 and 2 periods in this simple model) since it indicates how likely an unemployed worker is to find a job. In order to ensure that the unemployment rate in the second period is u as well, the number of workers who find a job after the first period must equal the number of workers who lose a job, $u\lambda_{ue} \equiv (1 - u)\lambda_{eu}$.

A worker faces a lifetime budget constraint. She is able to borrow and lend at a fixed real interest rate r , so, if she earns y_1 in the first period and y_2 in the second period of her life, consumption must satisfy $(1 + r)c_1 + c_2 = (1 + r)y_1 + y_2$. We assume that c_1 is chosen after y_1 is realized and that c_2 is chosen after y_2 is realized. An optimal consumption plan will depend on the worker's fortune in the labor market

⁴The ideas contained here do not depend on the assumption of short-lived workers, although, in a model with a longer time horizon, recurrence of unemployment, not just unemployment duration, would be important for welfare.

as well as on her expectation about future labor market prospects. For now, think of the unemployment rate (u) and the probability that an unemployed person finds a job in the second period (λ_{ue}) as fixed. A worker who is employed in the first period will choose to consume c_1^e , while an unemployed worker will consume c_1^u . In the second period, consumption will depend on the entire employment history, giving four possible values — c_2^{uu} , c_2^{ue} , c_2^{eu} , and c_2^{ee} — for a worker who is always unemployed, unemployed then employed, and so on. The highest level of consumption in the second period will be that enjoyed by a worker who was employed in both periods, while the lowest level of consumption will be that experienced by a worker who was always unemployed. It is not generally possible to rank c_2^{ue} and c_2^{eu} except to say that they both lie in the interval between c_2^{uu} and c_2^{ee} .

It is important to realize that a worker can set a complete contingent path for consumption, that is, choose the six consumption levels c_1^u , c_1^e , c_2^{uu} , and so on, before realizing any labor-market outcome. Then she simply faces a lottery over the possible consumption levels. When she is young, the chance that she consumes c_1^u is u and the chance that she consumes c_1^e is $1 - u$, independent of unemployment duration. But, when she is old, the chance that she consumes c_2^{uu} is $u(1 - \lambda_{ue})$, the product of the probability that she is unemployed when young and the conditional probability that she does not move to employment when old. Given the inverse relation between λ_{ue} and mean unemployment duration, this is increasing in mean duration for a fixed unemployment rate. Likewise, the chance that she consumes c_2^{ee} is $(1 - u)(1 - \lambda_{eu}) = 1 - u(1 + \lambda_{ue})$, where we use the fact that $u\lambda_{ue} \equiv (1 - u)\lambda_{eu}$ to simplify the expression. Again, this is increasing in mean unemployment duration for a fixed unemployment rate. The chance of each of the intermediate events, consuming c_2^{ue} or c_2^{eu} , is $u\lambda_{ue}$, decreasing in mean unemployment duration.

But now it is easy to show that shorter average unemployment duration for a

fixed unemployment rate must raise the worker's expected utility.⁵ She can clearly afford the same consumption in the same states as before, so suppose for now that her consumption plan conditional on employment status does not change. Then a higher level of λ_{ue} is the opposite of a mean-preserving spread. It has no effect on first-period consumption, which remains a lottery placing probability u on c_1^u and otherwise yielding c_1^e . But, in the second period, it shifts probability mass away from the extreme consumption levels c_2^{uu} and c_2^{ee} and towards the intermediate levels. That is, the distribution of consumption in the second period has the same mean, independent of λ_{ue} ,⁶ but the distribution with a higher value of λ_{ue} second-order stochastically dominates the distribution with a lower value. Then it follows immediately from [Rothschild and Stiglitz \(1970\)](#) that, holding the consumption plan fixed, any risk-averse worker is better off in an economy with shorter mean unemployment duration. Intuitively, the utility gain from consuming c_2^{ue} rather than c_2^{uu} exceeds the utility loss from consuming c_2^{eu} rather than c_2^{ee} . Since the unemployment rate is the same in the first and second periods, the probability of consuming c_2^{ue} must equal that of consuming c_2^{eu} . But having higher transition rates λ_{ue} and λ_{eu} for a fixed unemployment rate raises the probability of each event and, therefore, raises expected utility. All this ignores the fact that, in an economy with shorter mean unemployment duration, the worker can reoptimize her consumption plan. Doing so will trivially yield still higher utility. This establishes that, for a given unemployment rate, workers prefer a shorter mean unemployment duration.

⁵This argument ignores any general-equilibrium effect on wages or interest rates. We thank Fernando Alvarez for suggesting this method of proof.

⁶The expected value of consumption in the second period is given by expected lifetime income minus expected first-period consumption, or $(2+r)(ub + (1-u)w) - (1+r)(uc_1^u + (1-u)c_1^e)$, which is independent of λ_{ue} .

2.2 Unemployment Duration and Hysteresis

Unemployment duration is an excellent predictor of whether a worker will find a job. Figure 3 shows that the fraction of unemployed workers with a given unemployment duration who find a job in the following month is a decreasing function of duration.⁷ From 1976 to 2000, this probability exceeds fifty percent for workers at the shortest end of the duration distribution and falls to barely ten percent for workers who have been unemployed for ninety weeks or more.

Broadly speaking, there are two possible explanations for this pattern. First, workers may be heterogeneous. Unemployed workers with low exit probabilities are dynamically sorted into long-term unemployment. According to this theory, changes in mean unemployment duration simply reflect changes in the composition of job losers. But there is some empirical evidence that workers who initially have a high probability of finding a job are less likely to find one if they have been unemployed for a long time (see Abbring, van den Berg, and van Ours 2000). That is, unemployment exhibits hysteresis (Blanchard and Summers 1987). This will be the case if skills atrophy during unemployment (Pissarides 1992), if long-term unemployment stigmatizes workers (Blanchard and Diamond 1994), or if the long-term unemployed lose contact with social networks (Granovetter 1974, Montgomery 1991) and so do not know where to look for good employment opportunities.

To the extent that these models help explain the decreasing hazard rate of finding a job depicted in Figure 3, then there may also be a role for government intervention

⁷We constructed this figure by matching individual observations in the CPS across months. Appendix B discusses the matching process and its limitations in detail. Ideally, one would construct this hazard rate directly from the observed cross-sectional distribution of unemployed workers' unemployment duration, but measurement errors, discussed in more detail in Section 3, make this impossible. Crudely put, it would appear that workers who are unemployed for 51 weeks have a negative probability of finding a job in the next week since so many more workers report 52 weeks of unemployment.

in retraining the long-term unemployed. For example, [Becker \(1964\)](#) predicts that workers will bear the costs of general human capital. But, in the presence of credit constraints, this may not be a realistic possibility for the long-term unemployed, creating a role for subsidized training programs. Likewise, a firm with imperfect information about worker quality may use unemployment duration as a screening device. Although it knows that the long-term unemployed would be willing to accept a low wage in return for a job, the firm also realizes that, if it manages to hire a good worker from the long-term unemployed population, the market will quickly perceive the worker's high marginal product and drive up her wage. That is, other firms will free-ride, reducing the firm's information acquisition below its efficient level. The government can alleviate this problem by running training programs designed in part to establish participants' ability to partake in environments similar to the workplace. In a similar vein, [Caplin and Leahy \(2000\)](#) argue that equilibrium unemployment duration may be excessively long since, in an environment characterized by social networks, unemployed workers do not internalize all the informational benefits of maintaining contact with the labor market. If this is correct, search subsidies for the long-term unemployed may improve welfare by reconnecting groups of discouraged workers with the labor market.

2.3 Unemployment Duration and Wage Pressure

The decreasing relation between unemployment duration and job-finding hazard rates depicted in Figure 3 also suggests the possibility that the long-term unemployed may put less downward pressure on wages than do the short-term unemployed. If this is the case, an increase in unemployment duration will tend to shift the Phillips curve

to the right and raise the NAIRU (nonaccelerating inflation rate of unemployment).⁸ To assess the quantitative significance of this effect, define an index of unemployment pressure on the labor market equal to the unemployment rate at time t multiplied by the job-finding probability of the average unemployed worker at time t . We calculate the latter term using the average job-finding rate conditional on unemployment duration from 1976 to 2000, weighted by the unemployment-duration distribution prevailing at t . Figure 4 depicts the results. The index of unemployment pressure closely tracks the aggregate unemployment rate, although the cyclic variation is slightly muted. It thus seems unlikely that shifts in the unemployment-duration distribution have caused substantive shifts in the Phillips curve.

3 Measurement Issues

An important issue with respect to analyzing movements in unemployment and unemployment duration over time is whether these data have been affected by the major redesign of the CPS that was introduced in January 1994. This redesign included important improvements in the survey questionnaire as well as the conversion of the survey from a pencil-and-paper instrument to a computerized instrument. Available evidence suggests that the redesign had no significant effect on the overall unemployment rate or on the unemployment rates for age/sex groups, with the limited exception that it appears to have raised measured unemployment rates for persons aged 55 and older (see Polivka and Miller 1998). The proper conclusions to draw with respect to the measurement of unemployment duration, however, are less clear.

Two important changes in the design of the CPS could have affected the measurement of unemployment duration. The first change is the use of dependent inter-

⁸A decrease in the number of out-of-the-labor-force workers moving directly into employment could have a similar effect, but we find no evidence that such a shift has occurred.

viewing. Households included in the CPS sample are interviewed for four consecutive months, are out of the sample for the next eight months, and then are interviewed for an additional four consecutive months. Prior to the redesign, an individual who was reported as unemployed was always asked the duration of her unemployment spell. Since the redesign, the unemployment duration of a worker who is unemployed in consecutive months is calculated automatically on the basis of the spell length recorded for the earlier month and the number of weeks between the two months' survey reference periods. Among other possible effects, this change can be expected to have reduced the number of people recorded as having spells of 0 to 4 weeks duration. The second change is to allow individuals to report unemployment duration in months or years rather than only in weeks, although interviewers are instructed to ask for duration in weeks for anyone reporting four or fewer months in their current unemployment spell. It is not a priori obvious what effect this change should have had.

In Appendix A, we discuss one approach to analyzing the effect of the CPS redesign, the use of a parallel survey that was constructed explicitly for this purpose (Polivka and Miller 1998). For reasons discussed in the Appendix, however, we are hesitant to accept the results this approach yields for unemployment duration at face value. Instead, we attempt to assess the likely effect of the CPS redesign on unemployment duration by directly measuring the quantitative significance of the design changes. Consider the effect of dependent interviewing. Before the redesign, there was no difference in the way in which unemployment duration was measured for different rotation groups. But, after the redesign, unemployment duration was measured differently for the “incoming rotation groups” (the first and the fifth months in the sample) than it was for the rest of the sample. Substantial differences between reported unemployment durations have subsequently emerged. According to time

series that we constructed from the CPS, since the redesign about 6.5 percent more of the unemployed workers in the incoming rotation groups report that their current unemployment spell has lasted for less than four weeks, compared to the full sample (Figure 5). There were no meaningful differences across rotation groups before the CPS redesign, so this is almost certainly a redesign effect. More to the point, from 1994 to 2000, the mean unemployment duration averaged 15.20 weeks for the full CPS sample but only 14.83 weeks for the incoming rotation groups. Most of this difference is accounted for by short-duration unemployment (0 to 4 weeks), which averaged 1.97 percent of the labor force for the full sample after the CPS redesign but 2.40 percent for the incoming rotation groups. On the other hand, the rates of long-term unemployment (15 to 26 weeks) and very long-term unemployment (27 weeks and longer) were essentially unaffected by the switch to dependent interviewing. They averaged 0.68 and 0.79 percent, respectively, in the full sample and 0.67 and 0.78 percent, respectively, for the incoming rotation groups. Although it is likely that the full sample estimates are more accurate (and certainly are based on a much larger sample), the estimates for the incoming rotation groups are more comparable with the pre-1994 data.⁹ Throughout the remainder of our analysis, we use only data for the incoming rotation groups from 1994 forward.

Unfortunately, the absence of an obvious control group means that we cannot perform a similar assessment of the other relevant redesign change, allowing individuals to report their unemployment duration in months or years instead of only in weeks. The redesign resulted in a marked increase in the frequency with which certain un-

⁹In some cases, an individual observed as unemployed in two successive CPS interviews might in fact briefly have held a job in the intervening period, in which case she would have been properly recorded as having had a very short unemployment duration in the second month. Research conducted as part of the CPS redesign process suggested this to be a relatively rare event (Polivka and Miller 1998) and, in any case, there is a real question as to whether holding a job for only a very short time should be considered as having broken an otherwise continuing unemployment spell.

employment durations were reported.¹⁰ For example, from 1994 to 2000, 98.5 percent of unemployed workers in the incoming rotation groups who reported that their unemployment spell had lasted between 49 and 55 weeks said that their spell had been in progress for exactly 52 weeks (one year). During a roughly comparable period a decade earlier, only 80.7 percent of these unemployment spells were reported as having lasted exactly 52 weeks. Similarly, the share of 23- to 29-week spells that were reported as lasting exactly 26 weeks (6 months) increased from 34.3 to 91.3 percent, and the share of 36- to 42-week spells that were reported as lasting exactly 39 weeks (9 months) increased from 11.0 to 88.6 percent. Other changes were less dramatic.

Obviously, when offered the option to report their unemployment duration in months or years rather than in weeks, many respondents choose a round number. This need not bias the mean unemployment duration if as many workers report 52 rather than 51 weeks as report 52 rather than 53 weeks, but, in practice, that may not have occurred. The fact that there are more workers with shorter unemployment durations implies that symmetrical rounding errors like the one described above will result in an increase in measured mean unemployment duration. On the other hand, the rounding errors need not be symmetrical (Baker 1992). For example, it is plausible that a worker is as likely to report a 45-week unemployment spell as lasting for nine months as she is to report it lasting for one year. This may further bias the measured unemployment duration, although the direction of the bias is not *ex ante* obvious.

Fortunately, any effect of the CPS redesign on rounding errors is less likely to alter measured short-term and long-term unemployment rates. Interviewers are instructed to ask for unemployment duration in weeks whenever a respondent reports a duration of less than four months. To the extent that this is in fact done, all individuals should

¹⁰Baker (1992) discusses adjustments to reported unemployment durations to account for ‘digit preference’, but his analysis — conducted before the CPS redesign — does not help with our problem, the effect of the redesign on reported unemployment duration.

be correctly assigned to two key duration categories, 0 to 4 weeks and 5 to 14 weeks (short- and medium-term unemployment). The number of spells that have lasted more than 27 weeks, however, may be underestimated if some respondents report a six-month (26-week) duration rather than a slightly higher number.¹¹ We sidestep this by extending the definition of *long-term unemployment* to include durations from 15 to 32 weeks, and thus define *very long-term unemployment* to include spells that exceed 33 weeks.

To increase our confidence that we are not missing some other important effects of the CPS redesign, we explore whether the relation between the incoming rotation groups' unemployment duration and other cyclic labor-market variables changed in 1994. For this to be a meaningful exercise, the other cyclic indicators must be ones that were not themselves affected by the CPS redesign. We focus on a measure of the employment-population ratio constructed using employment from the BLS establishment survey (the Current Employment Statistics, or CES, survey) and the Census Bureau's projection of the civilian noninstitutional population aged 16 and older that serves as the CPS control total. The employment-population ratio trends upward over time, so we apply a Hodrick-Prescott (HP) filter to the time series.¹² We similarly detrend the seasonally adjusted mean unemployment duration as well as the short-term unemployment rate (0 to 4 week), the long-term unemployment rate (15 to 32 week), and the very long-term unemployment rate (33 or more weeks) unemployment rates.¹³

¹¹Note that this bias would tend to reduce the measured very long-term unemployment rate and so in any case cannot explain the surprisingly high rates at the end of the 1990s noted in the introduction.

¹²We set the smoothing parameter to 1,440,000, one hundred times larger than the standard value with monthly data. This does a good job of capturing the low-frequency movements in the various time series without erasing the cyclic variation. Lower values of the smoothing parameter yield similar results. On the other hand, owing to its substantial upward trend during the sample period, the raw time series for the employment-population ratio does a poor job of predicting the unemployment rate.

¹³From 1968 to 1993, we calculate the numbers using a monthly time series containing the number of unemployed workers with each week's duration, constructed by the BLS. After 1994, we

There is a strong negative correlation between the employment-population ratio and measures of longer-duration unemployment: -0.85 with the mean unemployment duration; -0.89 with the long-term unemployment rate; and -0.90 with the very long-term unemployment rate. The correlation with the short-term unemployment rate is weaker but still negative (-0.55).

To test for a break in a CPS-based time series in January 1994, we regress the series on the employment-population ratio using data from 1968 through 1993. We then forecast the series from 1994 through 2000 and plot the residual forecast errors. If there were a break in the time series, we would expect the residuals to look different after the break and, in particular, to jump in January 1994. The results are shown in Figure 6. The regressions perform remarkably well out of the sample period. In only one case, the very long-term unemployment rate, is there any indication of a discontinuity in January 1994, and even then the change is not unusual by historical standards. A reasonable measure of the possible effect of the redesign is the mean value of the residuals during the first half of 1994.¹⁴ The unemployment rates show no systematic pattern, with the short-term rate 0.02 percent below, the long-term rate 0.13 percent below, and the very long-term rate 0.13 percent above the expected levels, all small values given the pattern of variation in the residuals. The mean unemployment duration is further from trend, 1.42 weeks above what one would expect in the first half of 1994, but that is not much different than the deviation in the last half of 1993, when it was 0.85 weeks above trend. Because the point estimates are imprecise, we cannot reject the null hypothesis that the CPS redesign had no effect on the measured mean unemployment duration in the incoming rotation

constructed our own time series from the CPS using only the incoming rotation groups. The data are seasonally adjusted using a ratio-to-moving-average procedure. The required data are unavailable before 1968.

¹⁴Filtering the data diminishes our ability to recognize redesign effects many years after the redesign. This should not be a problem during the first year.

groups.

In sum, although it had no significant effect on the unemployment rate, the CPS redesign indisputably affected measured unemployment duration. Only the effects associated with the introduction of dependent interviewing can be identified with confidence, and we control for those in what follows. The results just discussed suggest that the full effect of the redesign on mean unemployment duration could have been somewhat larger, but we have no strong empirical basis for making additional adjustments and so choose not to do so. Figure 7 summarizes these findings, depicting our preferred time series for mean unemployment duration and the short-, long-, and very long-term unemployment rates, with the standard time series depicted as a dashed line.¹⁵ While measurement issues explain much of the recent decline in the short-term unemployment rate, they do not explain the persistently high level of very long-term unemployment or mean unemployment duration.

4 Aging of the Baby Boom

The youth unemployment rate is much higher than the prime-age unemployment rate. Viewing the aggregate unemployment rate as a weighted average of the unemployment rate of workers in different age cohorts, part of the recent decline in unemployment may be attributed to a simple compositional effect, a consequence of the aging of the baby boom (Shimer 1998). What is true for the aggregate unemployment rate is even

¹⁵The standard time series differ from our preferred time series before the CPS redesign as well as after. The large differences in the long- and very long-term unemployment rates are due to differences in definitions: the solid lines indicate the 15- to 32-week and 33 or more week unemployment rates, while the dashed lines indicate the 15- to 26-week and 27 or more week unemployment rates. The small differences in the mean unemployment duration and the short-term unemployment rate are due to changes in population weights associated with the 1980 and 1990 census. In order to construct the desired time series for the long- and very long-term unemployment rates, we had to use the original weights based on the 1980 census, while the official time series use adjusted weights based on the 1990 census.

more true for the short-term unemployment rate. Young workers rarely suffer long spells of unemployment,¹⁶ so, from 1968 to 2000, the fraction of teenagers unemployed for less than four weeks in a typical month was 6.5 times as high as the comparable fraction of 35- to 54-year-old workers. The short-term unemployment rate for 20- to 24-year olds was 3.2 times as high, and for 25- to 34-year olds was 1.6 times as high. Since 1980, the share of the labor force accounted for by 35- to 54-year-olds has risen by 13.2 percent, mostly at the expense of younger workers. The aging of the baby-boom generation during the last two decades has therefore shifted the population into age groups that suffer fewer short unemployment spells.

Under the assumption that the aging of the baby boom has not affected the unemployment probability or unemployment duration of workers conditional on their age, we can measure the effect of this demographic transition on the overall unemployment rate, the short-, long-, and very long-term unemployment rates, and the mean duration of unemployment. To illustrate, let $u_{i,t}$ denote the unemployment rate of age group i in year t and $\omega_{i,t}$ denote the labor force share of that age group in the same year. Then the unemployment rate is just a weighted average of the unemployment rates of workers of different ages, $\sum_{i=1}^N \omega_{i,t} u_{i,t}$, where N is the number of age groups.¹⁷ Under the assumption described previously, the unemployment rate in year t would have been $u_t^s \equiv \sum_{i=1}^N \omega_{i,s} u_{i,t}$ if the labor-force shares had remained at their year- s levels. We refer to this as the *age-adjusted* unemployment rate. Similar calculations can be carried out for the other variables of interest. To apply this procedure, we must select a base year s . To the extent that the differences across age groups in the

¹⁶This is partly because of how unemployment duration is measured. Young workers who do not have a job often engage in periodic job search, thus moving in and out of the labor force (Clark and Summers 1982).

¹⁷We use the $N = 7$ “standard” age groups — 16-19, 20-24, 25-34, 35-44, 45-54, 55-64, and over 65 — for these calculations. In practice, further refinements have little effect on demographic adjustments.

unemployment rates $u_{i,t}$ are stable from one year to the next, as is generally the case, the choice of base year is of no real consequence. In what follows, we consistently use 1980 as the base year for our age adjustments, but we have obtained similar results with other base years.

Figure 8 shows the effects of age adjustments on the overall unemployment rate. As documented in [Shimer \(1998\)](#), the age-adjusted unemployment rate has been higher relative to rates observed in the past than has the unadjusted rates more commonly the focus of attention. Between 1980 and 1992, 0.67 percentage points of the decline in the aggregate unemployment rate could be attributed to the aging of the baby boom. Some of this trend has been reversed in the last few years, but, comparing across business-cycle peaks, the aging of the baby boom contributed to a 0.21 percentage point increase in the aggregate unemployment rate from 1969 to 1979, a 0.44 percentage point decline over the next decade, and a further 0.14 percentage point decline from 1989 to 2000.¹⁸

The age-related decline in unemployment has been concentrated at the short end of the unemployment-duration spectrum. Figure 9 shows that the age-adjusted short-term unemployment rate was 2.1 percent in 1969 and 2.5 percent in 1973 but that it has not subsequently fallen back to these historic lows. In contrast, age-adjusted mean unemployment duration has increased modestly from expansion to expansion during the last three decades. The age-adjusted long-term unemployment rate has declined somewhat during the 1980s and 1990s, but the age-adjusted very long-term unemployment rate has not declined at all. Taking the movements in overall unem-

¹⁸Whether this demographic adjustment is quantitatively significant is a matter for debate. Within this volume, [Bertola, Blau, and Kahn \(2001\)](#) and [Blank and Shapiro \(2001\)](#) emphasize demographic trends, but [Staiger, Stock, and Watson \(2001\)](#) argue that demographics are unimportant for understanding recent shifts in the Phillips curve, and [Cohen, Dickens, and Posen \(2001\)](#) claim that demographics cannot explain sudden sharp jumps in the Beveridge curve. Our view is that demographic trends are important for understanding shifts beginning in 1980 but not very important for changes during the last decade.

ployment shown in Figure 8 and the movements in the duration measures shown in Figure 9 together, neither our measurement adjustments nor our age-structure adjustments alter the qualitative conclusion that long-duration unemployment has been high in recent years compared to the aggregate unemployment rate.

5 Labor Force Attachment

5.1 Women’s and Men’s Labor Market Outcomes

Our analysis so far has ignored an important factor in determining labor-market outcomes, the worker’s sex. From 1965 through 1979, women’s age-adjusted unemployment rate was on average nearly two percentage points higher than men’s rate (Figure 10). Since then, the two rates have converged, with an average difference from 1980 to 2000 of less than 0.1 percentage point, although women’s unemployment rate has been somewhat less variable than men’s at business-cycle frequencies. Given the close link between the unemployment rate and unemployment duration, one might expect that men’s and women’s age-adjusted unemployment durations would also have tracked each other quite closely since 1980. Figures 11 and 12, which graph men’s and women’s age-adjusted mean unemployment duration and short-, long-, and very long-term unemployment rates, show that this has not been the case. Men have consistently experienced much less short-term unemployment than have women and hence have had a much longer mean unemployment duration. This observation enables us to refine our understanding of the unemployment-duration puzzle.

Men’s unemployment rate and unemployment duration have had a relatively stable relation during the past three decades (Figure 11). After accounting for age, men’s unemployment rate reached 3.0 percent at the end of the 1960s before rising to 5.1

percent in 1979 and 5.6 percent in 1989. By 2000, it had fallen to 4.4 percent, which is by no means an unprecedented level. Likewise, men's mean unemployment duration rose from 8.3 weeks in 1969 to over 12 weeks in 1979, where it has remained during subsequent expansions. Looking at means masks a significant secular increase in men's short-term unemployment rate, which has been offset by an increase in the prevalence of very long-term unemployment. Given these changes, male unemployment is higher at both extremes of the unemployment-duration distribution than would have been expected given the level of the aggregate unemployment rate, all after adjustment for measurement and age-structure changes. This "bimodality" is consistent with the growing dispersion of labor-market outcomes documented in other contexts, for example, the growth in the dispersion of labor-market earnings (Katz and Murphy 1992). It would be interesting to find out whether the increasing concentration of male unemployment at the short and long ends of the unemployment-duration spectrum can be explained in a similar manner.

In contrast, women's age-adjusted unemployment rate is at the lowest level since the 1950s, having fallen from 6.8 percent during the expansion at the end of the 1970s, to 5.8 percent during the 1980s expansion, to 4.7 percent by the end of the 1990s. Figure 12 shows that there has not been any drop in women's age-adjusted unemployment duration during this time period. Instead, women's mean unemployment duration increased from 9.5 weeks in 1979 and 9.4 weeks in 1989 to 11.0 weeks in 1990, primarily a consequence of a sharp drop in women's short-term unemployment rate and a sustained high level in women's very long-term unemployment rate. The historic link between women's unemployment rate and women's unemployment duration broke down during the 1980s and 1990s, and measurement and demographics do not explain why.

One hint is that, in all cases, men's and women's unemployment durations have

converged toward a common level. Our first thought was that this pattern might be linked to changes in the industry and occupation distribution of women’s employment, but further investigation showed that unemployment durations do not differ enough across these categories to explain the increase in women’s unemployment duration. Instead, we explore another big shift that has occurred during the last two decades, the increase in women’s attachment to the labor force.

5.2 Labor-Market Transitions

To understand labor-market attachment, we must analyze not only transitions between employment (E) and unemployment (U) but also movements in and out of the labor force (N). We construct a time series for women’s and men’s age-adjusted labor-market-transition rates following the methodology discussed in Appendix B and display them in Figure 13. These data are available for the period from 1976 through 2000. The first thing that stands out in this figure is that, for both men and women, the transition probability from E to U (hereafter the EU transition rate λ_{eu} , with similar notation for the other transition rates) and the NU transition rate λ_{nu} are at their lowest levels in twenty-five years, while the UE transition rate λ_{ue} is at its highest level. In an accounting sense, this is a major part of the explanation for the current low level of unemployment. In comparing women’s and men’s transition rates, there is no evidence of convergence in λ_{eu} or λ_{ue} , but the remaining four transition rates, which all involve the labor-market-participation decision, have steadily converged during the past twenty-five years. Women are now much less likely to exit the labor market directly from employment, a sign of increased attachment, while men have become somewhat more likely to exit the labor market from unemployment and less likely to reenter the labor force once they leave it, two signs of decreased

labor force attachment.

Next, we show how to use the transition data to construct measures of the aggregate unemployment rate and the short-term unemployment rate. Recall that any Markov transition matrix implies a unique steady-state distribution of workers across the three labor-market states E, U, and N, given by the appropriately normalized eigenvector associated with the eigenvalue of one:

$$e = k(\lambda_{nu}\lambda_{ue} + \lambda_{un}\lambda_{ne} + \lambda_{ue}\lambda_{ne})$$

$$u = k(\lambda_{ne}\lambda_{eu} + \lambda_{en}\lambda_{nu} + \lambda_{eu}\lambda_{nu})$$

$$n = k(\lambda_{ue}\lambda_{en} + \lambda_{eu}\lambda_{un} + \lambda_{en}\lambda_{un})$$

where e denotes the fraction of the population that is employed, u the fraction that is unemployed, and n the fraction that is not in the labor force, and k is a proportionality constant that ensures $e + u + n = 1$. From this distribution, we can calculate the “implied steady-state unemployment rate” $u/(e + u)$. Of course, there is no reason to believe that the economy must always be in steady state or, equivalently, that the actual unemployment rate equals the implied unemployment rate. But the top row of Figure 14 documents that, in practice, the implied unemployment rate does a remarkable job of tracking the actual age-adjusted unemployment rate both for women and for men. The mean absolute difference between the two time series is 0.14 percentage points for women and 0.17 percentage points for men.

We can also use the labor-market-flow data to construct a measure of the short-term unemployment rate. A fraction $\lambda_{ue} + \lambda_{un}$ of unemployed workers exit unemployment in an average month and, in steady-state, this must equal the fraction of unemployed workers who are in their first month of unemployment. The product of this and the unemployment rate implied by the labor-market-transition rates should,

therefore, give us an independent measure of the short-term unemployment rate. The bottom row of Figure 14 shows that the implied measure of short-term unemployment is on average about half a percentage point below the standard time series. The most plausible interpretation of this result harks back to our analysis of the effect of the CPS redesign on the measured unemployment rate. Recall that the incoming rotation groups report almost half a percentage point more short-term unemployment than did the full CPS in post-1994 data. This reflects the switch to dependent interviewing, with the result that a worker who is unemployed in consecutive months cannot report less than a five-week unemployment duration in the second month. Our construction of the short-term unemployment rate from labor market flow data is analogous to the switch to dependent interviewing since a worker who is unemployed in consecutive months cannot be counted as short-term unemployed. This suggests that our time series for short-term unemployment is internally consistent but that it does not measure the same quantity as the short-term unemployment rate reported elsewhere in the paper. Rather, it measures something more like what the redesigned CPS measures for workers who are not in the incoming rotation groups. This view is also quantitatively reasonable. The gap between the short-term unemployment rate constructed using transition data and the actual rate averaged 0.69 percentage point between 1994 and 2000, while the short term unemployment rate for the incoming rotation group was on average 0.63 percentage point lower than the short term unemployment rate in the other rotation groups over the same time period.¹⁹

In any case, the basic puzzle of women's unemployment duration remains in the time series implied by labor-market-transition rates: women's implied unemployment

¹⁹We do not have a good explanation for the widening gap between men's implied and actual short-term unemployment rates after 1994, but we do not believe that this is a redesign effect. The gap has grown slowly over time, while we would expect a redesign effect to have emerged suddenly in 1994.

rate fell from 7.0 percent in 1979 to 4.6 percent in 2000, a decline of 33 percent. Under ordinary circumstances, we would expect this to be reflected almost exclusively in declines at longer unemployment durations. But, instead, women’s implied short-term unemployment rate fell from 2.8 to 2.0 percent during the same time period, a decline of 29 percent. In contrast, men’s implied unemployment rate fell by half as much, 16 percent, over the same time period, while their implied short-term unemployment rate fell by just 7 percent.²⁰

To see whether this is related to the increase in women’s labor-force attachment (or to the decrease in men’s labor-force attachment), we conduct a counterfactual exercise whereby we “shut down” changes in the transitions between various labor-market states in turn. For example, we take women’s actual λ_{ue} and λ_{eu} time series but fix the other four transition rates at their 1979 level.²¹ Using the constructed sequence of Markov matrices, we calculate the implied aggregate and short-term unemployment rates. We perform similar counterfactual experiments by allowing for time-variation only in λ_{un} and λ_{nu} and similarly for time-variation only in λ_{en} and λ_{ne} . By plotting the resulting time series, we can see which of the changes in transition rates are important for understanding the secular variation in aggregate and short-term unemployment rates.

Consider first women’s aggregate unemployment, depicted in the top left panel of Figure 15. The solid line shows the persistent decline in women’s implied age-adjusted unemployment rate. The remaining lines indicate the contribution of the

²⁰Our choice of base year affects these numbers — particularly for men, who suffered a sharp increase in unemployment in 1980 — but not the qualitative result that the decline in women’s unemployment has been registered disproportionately at short unemployment durations. From 1980 to 2000, men’s unemployment rate fell by 41 percent, and men’s short-term unemployment rate fell by 26 percent. The decline in women’s unemployment rate was smaller (39 percent), but the decline in their short-term unemployment rate was larger (32 percent) over that time period.

²¹Our exercise is similar in spirit to [Pissarides \(1986\)](#), who examines the contribution of increases in λ_{eu} and decreases in λ_{ue} to the increase in British unemployment during the 1970s and early 1980s.

three other pairs of transition rates. Changes in λ_{un} and λ_{nu} are not particularly important for understanding the decline in women's unemployment rate, although they do help explain some of the business-cycle fluctuations in the rate. Not surprisingly, the increase in λ_{ue} and the decrease in λ_{eu} reduced the incidence of female unemployment, and their fluctuations are also important contributors to short-run variation in the unemployment rate. But the most surprising finding is the equally important contribution of changes in flows between E and N, with both the decrease in λ_{en} and the increase in λ_{ne} contributing to the secular reduction in women's unemployment.

This finding is even more apparent when we look at the counterfactual time series for women's short-term unemployment rate in the bottom left panel of Figure 15. Over half the trend decline is explained by changes in λ_{en} and λ_{ne} , much more than changes in the other transition rates can explain. The sharp decline in λ_{en} for women presumably reflects an increase in their attachment to the labor market. It was much less common for a woman to quit her job and exit the labor force in the 1990s than it was two decades earlier. To the extent that this implies that women today are better able to build long-term employment relations, the decline in women's unemployment rate is readily understandable. Worker-flow data give a unique perspective on this fact.

Performing a similar experiment with men's unemployment rates highlights the importance of labor force attachment (right panels of Figure 15). The *increase* in λ_{en} for men has *raised* men's unemployment rate by over half a percentage point and men's short-term unemployment rate by over 0.2 percentage point. Variations in the other transition rates are responsible for the observed decline in unemployment. Again, it is not surprising that the decrease in λ_{eu} and the increase in λ_{ue} reduced unemployment, although it is curious that these variables are responsible for virtually all the cyclic fluctuations in men's aggregate and short-term unemployment rate. In

contrast, for women, λ_{nu} and λ_{un} also contribute significantly to the cyclic movements in employment. More interesting is that the decline in λ_{nu} is equally important for understanding male unemployment. This reflects the well-known decrease in men's labor-market-participation rate over this time period (Juhn, Murphy, and Topel 1991) and provides some justification for the belief that the current low levels of unemployment in the United States are at least partially a consequence of the decline in men's participation rate.

To summarize the importance of labor-force attachment, it is easiest to focus on what would have happened had there been no change in the flows in and out of the labor market, only in λ_{eu} and λ_{ue} . Between 1979 and 2000, women's unemployment rate would have fallen by 16 percent (from 7.0 to 5.8 percent of the labor force), rather than by 33 percent (to 4.6 percent of the labor force), while their short-term unemployment rate would have fallen by only 7 percent, rather than by 29 percent. In contrast, over the same time period, men's unemployment rate would have fallen by 12 percent, rather than by 16 percent with all flow-rate changes taken into account, while their short-term unemployment rate would have fallen by 8 percent, rather than by 7 percent. This suggests that, in the absence of changes in labor-force attachment, men's unemployment duration would have increased relative to women's over this period. Conversely, it suggests that changes in labor-force attachment are responsible for the observed increase in women's unemployment duration relative to their unemployment rate during the last two decades.

6 Conclusion

This paper started with the observation that there has been a breakdown in the historic relation between the unemployment rate and unemployment duration. We

then showed that, for men, much of the relative increase in unemployment duration can be explained by measurement issues related to the 1994 CPS redesign and by the aging of the baby boom. For women, however, these two factors do not have much explanatory power. Instead, we used labor-market-transition rates to show that much of the decline in women's unemployment rate since the late 1970s is due to the increase in their attachment to the labor force and that this likewise explains why the decline has been concentrated at the short end of the unemployment-duration distribution. There are reasons to believe that this process is coming to an end. Women's and men's labor-market experiences have become much more similar and, in particular, their labor market transition rates are converging towards a common level.

The increase in women's labor-force attachment has manifested itself in a variety of other ways. Periodic job-tenure supplements to the CPS indicate that, after accounting for the aging of the labor force, the fraction of employed women who have worked in one job for at least ten years increased from 25 to 27 percent between 1983 and 2000. Over the same time period, the age-adjusted fraction of employed men in the same job actually fell, from 39 to 32 percent (Figure 16). Many authors have linked the changing occupation distribution of female employment to the increased likelihood that women will be continuously employed, arguing that women who expect to be in the labor market for an extended period are more likely to find it attractive to enter occupations requiring substantial on-the-job training. Similarly, to the extent that employers bear the costs of job training, they may be more likely to consider women for such jobs when they believe that their tenure will be longer (Blau and Kahn 2000). Moreover, as shown by Blau and Kahn (1997) and O'Neill and Polachek (1993), increases in working women's job experience as compared to men's have contributed substantially to the narrowing of the gap between men's and women's earnings. The labor-market-transition rate data that we have examined

indicate that the growing stability of women's employment relations is also central to understanding why women's aggregate and short-term unemployment rates have declined while women's unemployment duration has increased. Conversely, men's weakening labor-force attachment appears to be important for understanding the relatively high levels of aggregate and short-term unemployment for men. Although the trend towards declining male labor-force participation has abated in recent years, its resumption would exacerbate these effects.

Appendix

A. Parallel-Survey Analysis of the CPS Redesign

Polivka and Miller (1998) provide a comprehensive and careful analysis of a parallel survey that was designed to assess the effects of the CPS redesign. From July 1992 through December 1993, the parallel survey sample was administered the redesigned CPS instrument. From January 1994 through May 1994, the parallel survey sample was administered the old, pencil-and-paper CPS instrument. With appropriate identifying assumptions, estimates based on the parallel survey can be used in conjunction with those from the official CPS to assess the effect of the CPS redesign on various variables of interest.

Polivka and Miller select the following as their preferred specification for assessing the effects of the CPS redesign:

$$Y_{it} = \mu_t + \delta_p p_{it} + \delta_m m_{it} + \varepsilon_{it} \quad (1)$$

where Y_{it} is the variable to be estimated, p_{it} is a dummy variable that equals 1 for parallel survey estimates and 0 for official CPS estimates, m_{it} is a dummy variable that equals 1 for estimates based on data collected using the new CPS methodology and 0 for estimates based on data collected using the old CPS methodology, ε_{it} is the equation error, the μ_t are time period effects, and δ_p and δ_m are parameters to be estimated. We are particularly interested in δ_m , which is intended to capture the magnitude of the effect of the redesign on the measured value of the dependent variable Y_{it} . Polivka and Miller fit the model using estimates for the period October 1992 through May 1994 derived from both the parallel survey and the official CPS.

Note that this specification allows estimates based on parallel survey data to differ from estimates based on data collected using the same methods as part of the official CPS, but identifies the parallel survey effect by constraining it to be the same before and after January 1994.²²

Using this methodology, Polivka and Miller conclude that the redesign did not significantly affect the unemployment rate, although they estimate a significant parallel survey effect δ_p . Figure 17 plots seasonally-adjusted versions of the unadjusted series Polivka and Miller use to fit their unemployment rate model.²³ As is true of the unadjusted data, the parallel survey unemployment estimates exceed the CPS estimates in all months. Had the higher parallel survey rates observed prior to January 1994 represented a redesign effect, the parallel survey and the CPS lines should have crossed between December 1993 and January 1994, as the parallel survey estimate fell and the CPS estimate rose.

We use similar models to investigate the effect of the redesign on the short-term (0- to 4-week), long-term (15- to 26-week) and very long-term (27 or more weeks) unemployment rates. The estimated values of associated with these three dependent variables imply that the redesign reduced the measured short-term unemployment rate by 0.46 percentage point; raised the long-term unemployment rate by 0.17 percentage point; and raised the very long-term unemployment rate by 0.25 percentage point. Application of the model to assess the effect of the redesign on the mean

²²The specification shown in equation (1) treats the CPS redesign effect as additive. Polivka and Miller also estimate models in which the CPS redesign effect is treated as multiplicative. The same general comments would apply to both models and the qualitative conclusions to be drawn from them are similar.

²³Because only 20 months of parallel survey data are available, it was not possible to construct seasonal factors directly from the parallel survey time series. Instead, we apply multiplicative seasonal factors based on official CPS data to adjust the parallel survey data, and then use these seasonally-adjusted estimates throughout in what follows. Although the CPS seasonal factors capture the seasonal movements in the parallel survey data imperfectly, this imperfect adjustment helps to highlight the nonseasonal movements in which we are interested. None of our conclusions are affected by the use of adjusted rather than unadjusted data.

duration of unemployment (in weeks) yields an estimated value for δ_m of 2.37.

An examination of the data underlying these estimates leaves us comfortable with the conclusion regarding the impact of the redesign on the incidence of short-duration unemployment, but raises questions about the implied impact on the other unemployment duration measures. The four panels of Figure 18 plot the parallel survey values and the corresponding official CPS values used to estimate each of the four unemployment duration models.²⁴ The top right panel, which shows the short-term unemployment rate, looks exactly as one might have expected. The parallel survey short-term unemployment rate jumps upward at the point of transition from the new to the old survey protocol, consistent with the expectation that the new survey protocol should have reduced the number of people reporting unemployment spells of 0 to 4 weeks duration. Presumably because the effects of dependent interviewing in the new protocol are not felt until the second month a person is interviewed, the corresponding drop in the official CPS estimates of the incidence of short-duration unemployment follows with a one month lag. The decrease in the CPS short-term unemployment rate is accompanied by a roughly offsetting increase in the 5- to 14-week unemployment rate series (not shown in the figure), although there is no offsetting decrease in the corresponding series in the parallel survey.

The bottom left panel displays the long-term (15- to 26-week) unemployment rate series calculated using data from the two surveys. Consistent with the existence of a redesign effect on these data, the parallel survey and official CPS estimates move in opposite directions between December 1993 and January 1994. On the other hand, the large swings in the gap between the two series during both the pre-transition period and the post-transition period seem inconsistent with the simple model hypothesized by Polivka and Miller, raising questions about its applicability.

²⁴We are grateful to Anne Polivka for making the data that underlie these figures available to us.

We are most troubled by the plot of mean unemployment duration in the top left panel and the plot of the very long-term unemployment rate in the bottom right panel. A redesign effect should produce opposing movements in the parallel survey and the official CPS data. Instead, at the point of transition between survey protocols, there is a sharp drop in the two parallel survey series, while the two official CPS series remain essentially unchanged.

In principle, declines in the true level of mean unemployment duration and the true incidence of very long-term unemployment between December 1993 and January 1994 could be responsible for this pattern. Such declines would have magnified the negative measurement effect of moving to the old methodology in the parallel survey and masked the positive measurement effect of moving to the new methodology in the official CPS. In practice, however, movements of the necessary magnitude are implausibly large. The December 1993-January 1994 drop in (seasonally adjusted) mean unemployment duration implied by our estimates of the Polivka-Miller model is 2.24 weeks.²⁵ To put this in context, the standard deviation of the month-to-month change in the official seasonally adjusted CPS mean unemployment duration series is just 0.52 weeks and the largest absolute change ever observed over the 1948 to 2000 period is only 1.9 weeks. The corresponding implied drop in the very long-term unemployment rate is 0.33 percentage point. The month-to-month change in the official CPS very long-term unemployment rate series has a standard deviation of just 0.06 percentage point and the largest post-1948 absolute change in the same series is only 0.26 percentage point. In contrast, the implied December-1993-to-January-1994 movements in both the short-term and the long-term unemployment rates are well within those series' one-standard-deviation bounds.

²⁵This is calculated as the difference between the January 1994 and the December 1993 μ_t values in the specification shown in equation (1).

One possible explanation is that the post-January 1994 parallel survey did not yield unemployment duration estimates fully comparable to those that continuation of the old pencil-and-paper CPS would have produced. There are some noteworthy differences in how the two surveys were administered. Fewer than half of the interviewers for the parallel survey from January 1994 forward had any prior experience with the pencil-and-paper instrument. In addition, a substantial share of the post-January 1994 parallel survey interviewers had recent experience with the computerized CPS instrument, which could have contaminated the responses they collected to the questions on the old CPS in unknown ways. This may have been less of an issue with the core CPS questions on which interviewer training presumably concentrated, but could have been a more important issue with the questions on unemployment duration. Although it is difficult to say exactly how these or other differences might have affected the survey estimates, it is well known that seemingly minor differences in survey administration can have a significant effect.²⁶ A further consideration is that, as compared to many of the other labor market variables for which Polivka and Miller report redesign effects, the sample sizes underlying all of the unemployment duration estimates are very small, and the estimates of mean unemployment duration, in particular, are unusually susceptible to outliers. Because we are not fully comfortable with using the parallel survey data to identify how the redesign affected the unemployment duration estimates, especially the estimates for mean unemployment duration and the incidence of very long-term unemployment, we pursue a different approach to measuring redesign effects in the text.

²⁶BLS-Census (1993) discusses a number of differences between the pre-January-1994 CPS and the post-January-1994 parallel survey that could have produced differences in the resulting estimates. This document also contains information on the projected mix of parallel survey interviewer experience that has been cited by other authors (for example, Polivka and Miller 1998) but turns out not to match the actual mix. The information on interviewer mix we cite was provided by Ron Tucker of the Census Bureau.

B. Construction of Worker Transition Data

Measures of worker transitions take advantage of the rotating-panel aspect of the CPS by matching workers' labor-market states across months. This requires access to the CPS public-use microdata, which we have only since 1976. Furthermore, there are a few months when we are not able to match workers across months because of survey redesigns or confidentiality suppressions in the available public-use data files.²⁷ We use the matched worker data to construct the probability that a worker who is in labor-market state X in month $t - 1$ is in state Y in month t , the XY transition rate in month t or, equivalently, λ_{xy} . Although this is the standard measure of worker transitions in the United States, it is important to recognize that it has a number of shortcomings. First, it is relatively noisy. In any month, three-quarters of the CPS households were in the CPS in the previous month and, theoretically, should be matchable. Because of sample attrition, including individuals who move between sample months and therefore are no longer eligible to be interviewed, and mistakes in recording data elements such as age or sex that are essential to the matching algorithm, matched files in practice contain about 70 percent of the observations in the unmatched files.²⁸ But, since some labor-market transitions (for example, men's UN transitions) are relatively rare, they cannot be estimated very precisely, particularly for small demographic groups. We boost our sample size by working with

²⁷We used a household identifier, individual line number, race, sex, and age in our matching algorithm. The values for each of these variables were required to be the same in month t and month $t - 1$, except that we also accepted a value for age in month t equal to one more than the value in month $t - 1$. The months for which matches could not be constructed due to redesigns or suppressions are January 1978, July 1985, October 1985, January 1994, and June - October 1995. In addition, matched data are unavailable for January 1976 since we do not have the required "last-month" file, December 1975.

²⁸This implies that about 7 percent of the eligible observations (5/75) cannot be matched across months, significantly less than the 15 percent reported by [Abowd and Zellner \(1985\)](#). The source of the difference is unclear, but it seems unlikely that we are matching observations that should not be matched. Our matching algorithm yields transition rates that are almost identical to those computed by the BLS, while one would expect over-matching to yield additional erroneous transitions.

annual average data.

Perhaps more important, some fraction of measured transitions are spurious. A misreported labor-market status in one month generates two erroneous transitions. In a careful analysis of data from 1977 to 1982, [Abowd and Zellner \(1985\)](#) estimate that as many as 40 percent of labor-force transitions are spurious. Unfortunately, [Abowd and Zellner's](#) analysis has not been replicated using more recent data, and to do so would go well beyond the scope of our analysis. As a consequence, we do not know whether the spurious transition rate has increased or decreased over time. With no contradictory evidence, we assume that the bias in estimated transitions is unchanged over time, so we adjust our estimated labor-force-transition rates using the percentage changes in table 5 of [Abowd and Zellner \(1985\)](#), in particular, the column labeled *Classification & Margin Error to Unadjusted*.

The 1994 CPS redesign may have affected the measured exit rate of unemployed workers from the labor force. There is no obvious reason why this should be so, but an examination of the time series of labor-market transitions indicates an unusually large and persistent increase in λ_{un} between December 1993 and February 1994. To get at this more precisely, we regress each detrended labor-market-transition series on the detrended CES employment-population ratio and look at the residuals for evidence of a break in the normal relation between worker transitions and the state of the business cycle, exactly as in our analysis of the effect of the redesign on measures of unemployment duration. [Figure 19](#) graphically displays the results for the six transition rates. λ_{un} takes a significant jump, with the residual increasing from -0.6 percent in December 1993 to 0.9 percent in February 1994 and then further increasing during the next six months. Comparing the last six months of 1993 to the first six available months of 1994 (February to July), the average value of the residual increases by a full 3.15 percentage points, twice as large as any comparable change in

the residual.²⁹ Although it is conceivable that this represents a structural shift in the relation between λ_{un} and the employment-population ratio, we feel fairly conservative in reducing λ_{un} by 3.15 percentage points after 1994. Since there is no evidence of a change in λ_{ue} , we offset this with an increase in the probability that an unemployed worker remains unemployed.

Figure 19 suggests that λ_{nu} may also have increased with the redesign, although the absolute size of the jump is smaller. The six-month change in the residual peaks at about 0.34 percent in January 1994, fifty percent larger than any other comparable change, and slightly more than one-tenth of the estimated size of the λ_{un} residual. Although the evidence here is somewhat weaker than it is for λ_{un} , the finding is consistent with Polivka and Miller’s (1998) evidence that the CPS redesign did not affect the measured stock of workers in each employment state. To understand why, observe that the worker-transition data imply a unique steady-state distribution of workers across employment states, the solution to a system of three equations:

$$\begin{bmatrix} \lambda_{ee} & \lambda_{ue} & \lambda_{ne} \\ \lambda_{eu} & \lambda_{uu} & \lambda_{nu} \\ \lambda_{en} & \lambda_{un} & \lambda_{nn} \end{bmatrix} \cdot \begin{bmatrix} e \\ u \\ n \end{bmatrix} = \begin{bmatrix} e \\ u \\ n \end{bmatrix} ,$$

where e denotes the fraction of the population that is employed, u the fraction that is unemployed, and n the fraction that is not in the labor force. The transition matrix is Markov, with columns summing to 1, so has one eigenvalue equal to 1. The solution to these equations is therefore given by the eigenvector associated with that eigenvalue. Moreover, for any change in λ_{un} to $\hat{\lambda}_{un}$ accompanied by a change in λ_{uu} that keeps the matrix Markov, it is possible to find a change in λ_{nu} together

²⁹The residual gradually reverts to zero over the next three years. This is because we are examining the relationship between detrended time series. Eventually, the trend catches up with and erases any effects of the CPS redesign.

with an appropriate change in λ_{nn} that generates the same ergodic distribution of workers across employment states, that is, a change in λ_{nu} to $\hat{\lambda}_{nu}$ that solves $\lambda_{eu} \times e + \hat{\lambda}_{uu} \times u + \hat{\lambda}_{nu} \times n = u$, together with a change in λ_{nn} that keeps the transition matrix Markov. Since accounting for the CPS redesign raises our measure of λ_{uu} by 3.15 percentage points but should not have affected the overall unemployment rate, it must reduce our measure of λ_{nu} by $3.15 \times u/n$ percentage points. With the number of unemployed persons averaging approximately one-ninth of the not-in-the-labor-force population during the past several years, this is almost perfectly consistent with our independent finding that the CPS redesign has increased the measured value of λ_{nu} by 0.34 percentage points from 1994 to 2000. A similar analysis finds no evidence of changes in any of the other transitions, so we account for the CPS redesign by reducing λ_{un} by 3.15 percentage points and raising λ_{nu} by 0.34 percentage points throughout our analysis.

We use this methodology to construct separate transition rates for women and men in the usual seven age groups. We then aggregate these into time series for women’s and men’s age-adjusted transition rates, weighting λ_{xy} for age group i by the share of age group i in the initial labor-market state X in the base year 1980. Finally, we adjust the data according to Table 5 in [Abowd and Zellner \(1985\)](#).³⁰ We then test and do not reject the hypothesis that the CPS redesign had the same effect on age-adjusted values of λ_{un} and λ_{nu} for men and women as it did on the aggregate transition rates, so we subtract 3.15 and 0.34 percentage points, respectively, from these transition rates after 1994.³¹

³⁰[Abowd and Zellner](#) report three different adjustments — for “all persons”, “male,” and “female”. The numbers that we report use the “all persons” adjustment because the available evidence suggests a convergence between men’s and women’s behavior since 1981, the end of the period that [Abowd and Zellner](#) study. In practice, this choice has little effect on our conclusions.

³¹These numbers are consistent with [Frazis \(1996\)](#), who, using an entirely different methodology, concluded that the CPS redesign raised λ_{un} by at least 2.8 percentage points and λ_{nu} by about 0.4 percentage points.

References

- Abbring, Jaap, Gerard van den Berg, and Jan van Ours. 2000. "Business Cycles and Compositional Variation in U.S. Unemployment." Unpublished paper. Free University, Amsterdam. Forthcoming, *Journal of Business and Economic Statistics*.
- Abowd, John, and Arnold Zellner. 1985. "Estimating Gross Labor-Force Flows." *Journal of Business and Economic Statistics*, 3(3): 254–283.
- Alvarez, Fernando, and Urban Jermann. 2000. "Using Asset Prices to Measure the Cost of Business Cycles." Unpublished paper. University of Chicago.
- Baker, Michael. 1992. "Digit Preference in CPS Unemployment Data." *Economics Letters*, 39(1): 117–121.
- Becker, Gary. 1964. *Human Capital*. University of Chicago Press, Chicago, IL.
- Bertola, Giuseppe, Francine Blau, and Lawrence Kahn. 2001. "Comparative Analysis of Labor Market Outcomes: Lessons for the US from International Long-Run Evidence." Unpublished paper. Cornell University. Prepared for the *Sustainable Employment Initiative*.
- Blanchard, Olivier, and Peter Diamond. 1994. "Ranking, Unemployment Duration, And Wages." *Review of Economic Studies*, 61(3): 417–434.
- Blanchard, Olivier, and Lawrence Summers. 1987. "Hysteresis in Unemployment." *European Economic Review*, 31(1–2): 288–295.
- Blank, Rebecca, and Matthew Shapiro. 2001. "Labor and the Sustainability of Output and Productivity Growth." Unpublished paper. University of Michigan. Prepared for the *Sustainable Employment Initiative*.
- Blau, Francine, and Lawrence Kahn. 1997. "Swimming Upstream: Trends in the Gender Wage Differential in the 1980s." *Journal of Labor Economics*, 15(1): 1–42.
- . . 2000. "Gender Differences in Pay." Unpublished paper. Cornell University.
- BLS-Census, CPS-Overlap Analysis Steering Committee. 1993. "Implications of Survey Design Differences between the Current CPS and the Proposed New Parallel Survey." Memorandum to Katharine Abraham and others dated December 22. Washington, D.C.: U.S. Bureau of Labor Statistics.
- Browning, Martin, and Thomas Crossley. 2001. "Unemployment Insurance Benefit Levels and Consumption Changes." *Journal Of Public Economics*, 80(1): 1–23.
- Caplin, Andrew, and John Leahy. 2000. "Mass Layoffs and Unemployment." *Journal of Monetary Economics*, 46(1): 121–142.

- Clark, Kim, and Lawrence Summers. 1982. "The Dynamics of Youth Unemployment." in *The Youth Labor Market Problem: Its Nature, Causes, and Consequences*, edited by Richard Freeman, and David Wise, pp. 199–235. Chicago University Press, Chicago.
- Cohen, Jessica, William Dickens, and Adam Posen. 2001. "Have the New Human Resource Management Practices Lowered the Unemployment Rate?." Unpublished paper. Brookings Institution. Prepared for the *Sustainable Employment Initiative*.
- Frazis, Harley. 1996. "Note on CPS Redesign Effects on Gross Flows." Unpublished paper. U.S. Bureau of Labor Statistics.
- Granovetter, Mark. 1974. *Getting A Job: A Study of Contacts and Careers*. Harvard University Press, Cambridge, Massachusetts.
- Gruber, John. 1997. "The Consumption Smoothing Benefit of Unemployment Insurance." *American Economic Review*, 87(1): 192–205.
- Juhn, Chinhui, Kevin Murphy, and Robert Topel. 1991. "Why Has the Natural Rate of Unemployment Increased over Time?." *Brookings Papers on Economic Activity*, (2): 75–142.
- Katz, Lawrence, and Alan Krueger. 1999. "The High-Pressure U.S. Labor Market of the 1990s." *Brookings Papers on Economic Activity*, (1): 1–87.
- Katz, Lawrence, and Kevin Murphy. 1992. "Changes in Relative Wages, 1963-1987: Supply and Demand Factors." *Quarterly Journal of Economics*, 107(1): 35–78.
- Lucas, Robert. 1987. *Models of Business Cycles*. Basil Blackwell, New York.
- Montgomery, James. 1991. "Social Networks and Labor-Market Outcomes: Towards an Economic Analysis." *American Economic Review*, 18(5): 1408–1418.
- O'Neill, June, and Solomon Polachek. 1993. "Why the Gender Gap in Wages Narrowed During the 1980s." *Journal of Labor Economics*, 11(1): 205–229, Part 1.
- Pissarides, Christopher. 1986. "Employment and Vacancies in Britain." *Economic Policy*, (3): 500–559.
- . 1992. "Loss of Skill During Unemployment and the Persistence of Employment Shocks." *Quarterly Journal of Economics*, 107(4): 1371–1391.
- Polivka, Anne, and Stephen Miller. 1998. "The CPS After the Redesign." in *Labor Statistics Measurement Issues*, edited by John Haltiwanger, Marilyn Manser, and Robert Topel, pp. 249–289. University of Chicago Press, Chicago.

- Rothschild, Michael, and Joseph Stiglitz. 1970. "Increasing Risk I: A Definition." *Journal of Economic Theory*, 2(3): 225–43.
- Shimer, Robert. 1998. "Why is the U.S. Unemployment Rate So Much Lower?." in *NBER Macroeconomics Annual*, edited by Ben Bernanke, and Julio Rotemberg, vol. 13, pp. 11–61. MIT Press, Cambridge, MA.
- Staiger, Douglas, James Stock, and Mark Watson. 2001. "Prices, Wages, and the U.S. NAIRU in the 1990s." Unpublished paper. Dartmouth University. Prepared for the *Sustainable Employment Initiative*.
- Summers, Lawrence. 1986. "Why is the Unemployment Rate So Very High Near Full Employment?." *Brookings Papers on Economic Activity*, (2): 339–383.

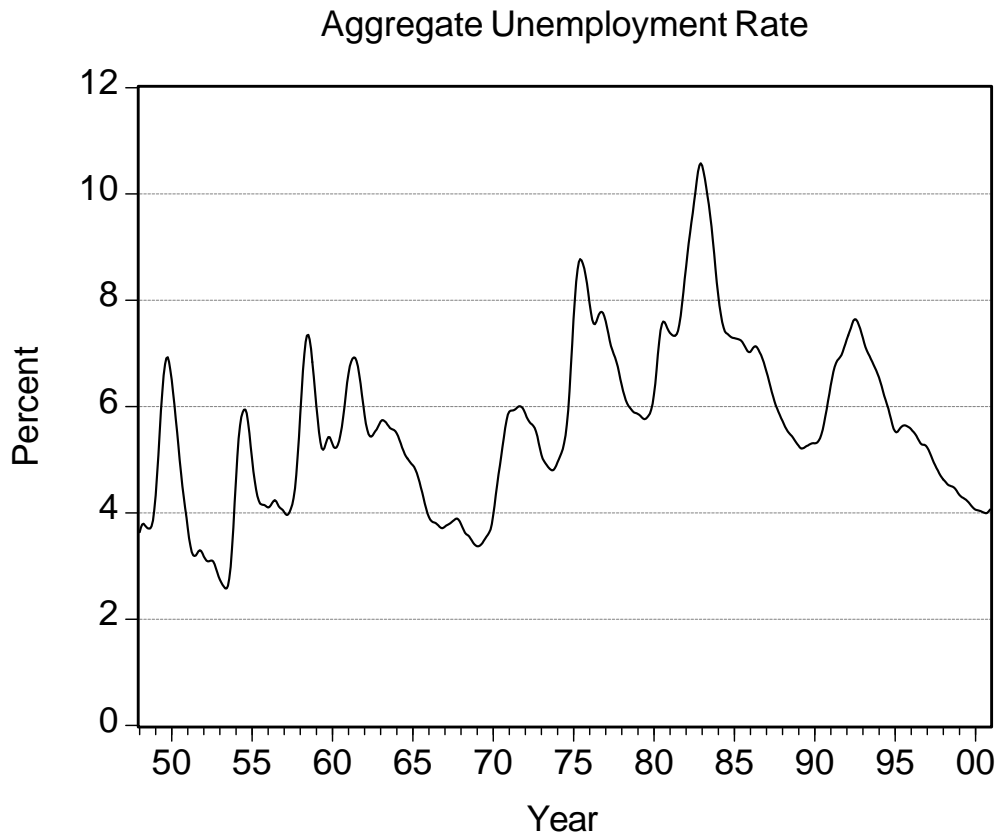


Figure 1: Monthly fluctuations were smoothed using a Hodrick-Prescott filter with parameter 10. Data are from official Bureau of Labor Statistics time series.

Aggregate Unemployment Duration and Rates

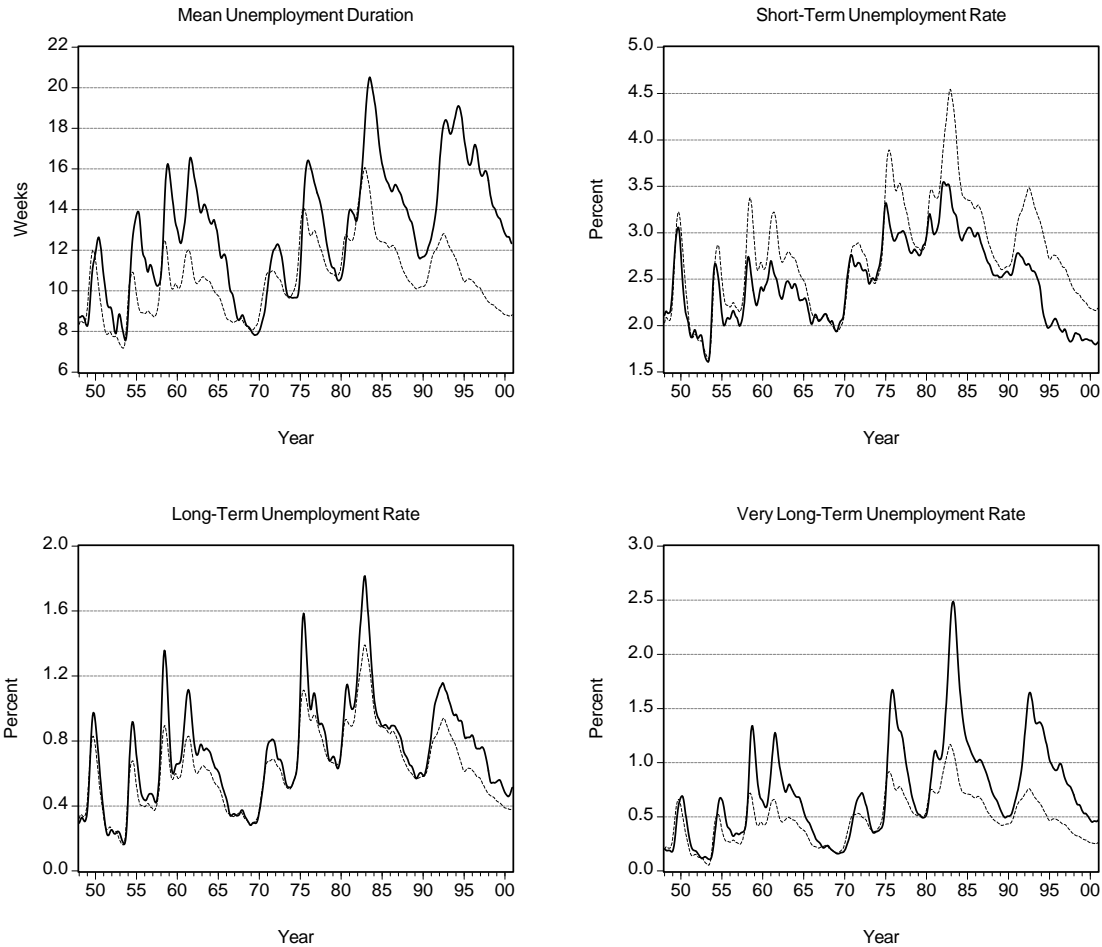


Figure 2: The top left panel shows the mean duration of unemployment spells currently in progress. The remaining panels show the share of the labor force with 0 to 4, 15 to 26, and 27 or more weeks unemployment duration (short-term, long-term, and very long-term unemployment rates, respectively). The dashed line in each panel shows the aggregate unemployment rate on a different scale. Monthly fluctuations were smoothed using an Hodrick-Prescott filter with parameter 10. Data are from official Bureau of Labor Statistics time series.

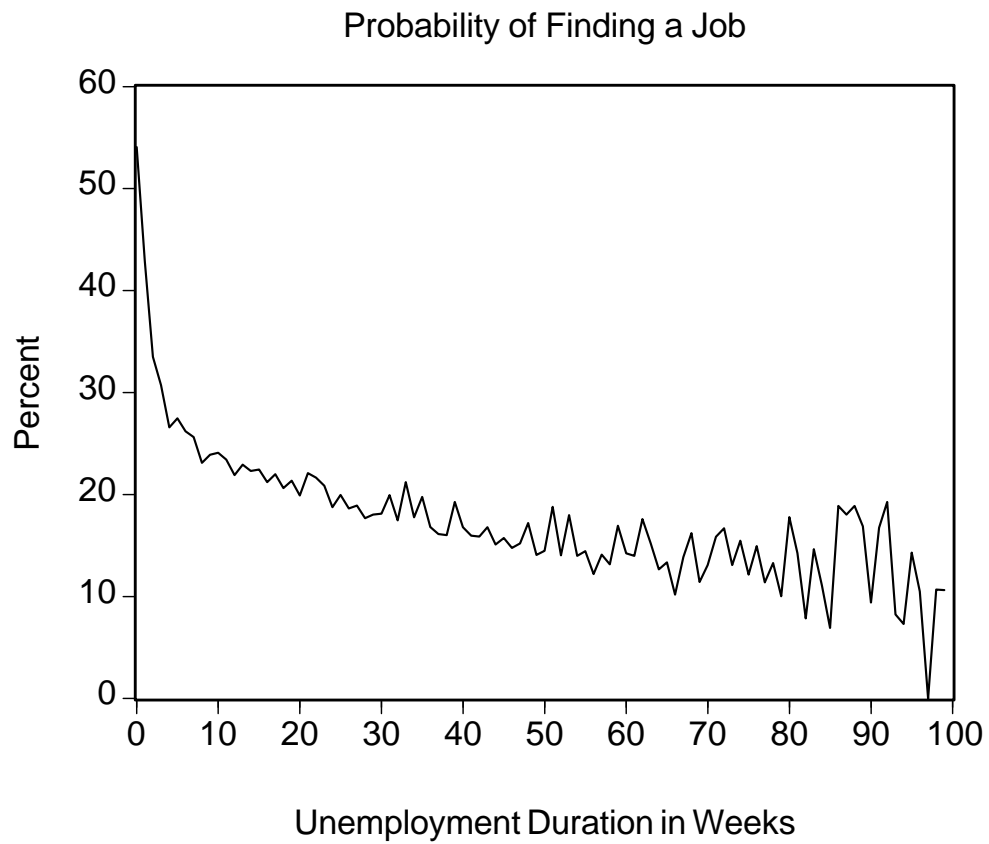


Figure 3: This figure shows the average probability of an unemployed worker becoming employed in the following month as a function of her unemployment duration, 1976 to 2000. Data are our calculations from the CPS.

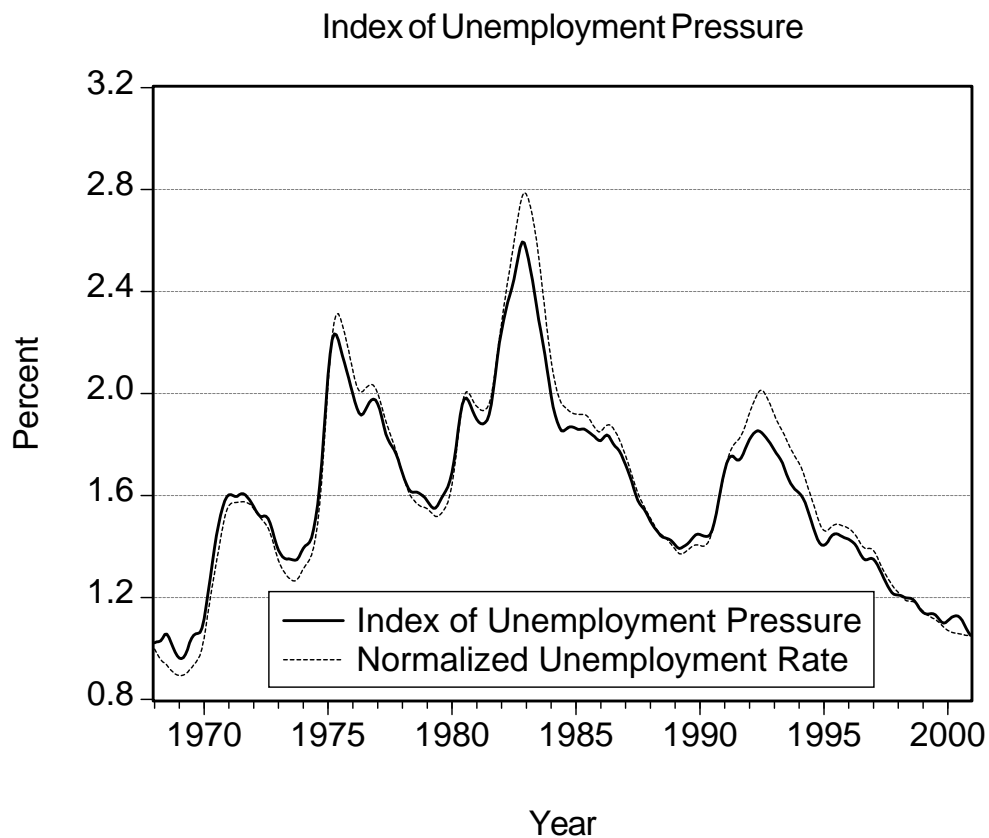


Figure 4: This figure shows the index of unemployment pressure (the solid line) and the normalized unemployment rate (the dashed line). The definition of the index of unemployment pressure is given in the text. The unemployment rate is normalized so it has the same average value as the index of unemployment pressure. Monthly fluctuations were smoothed using an Hodrick-Prescott filter with parameter 10. Data are from official Bureau of Labor Statistics time series and the authors' calculations from the CPS.

Unemployment Duration CDFs: Incoming Rotation Group minus Full Sample

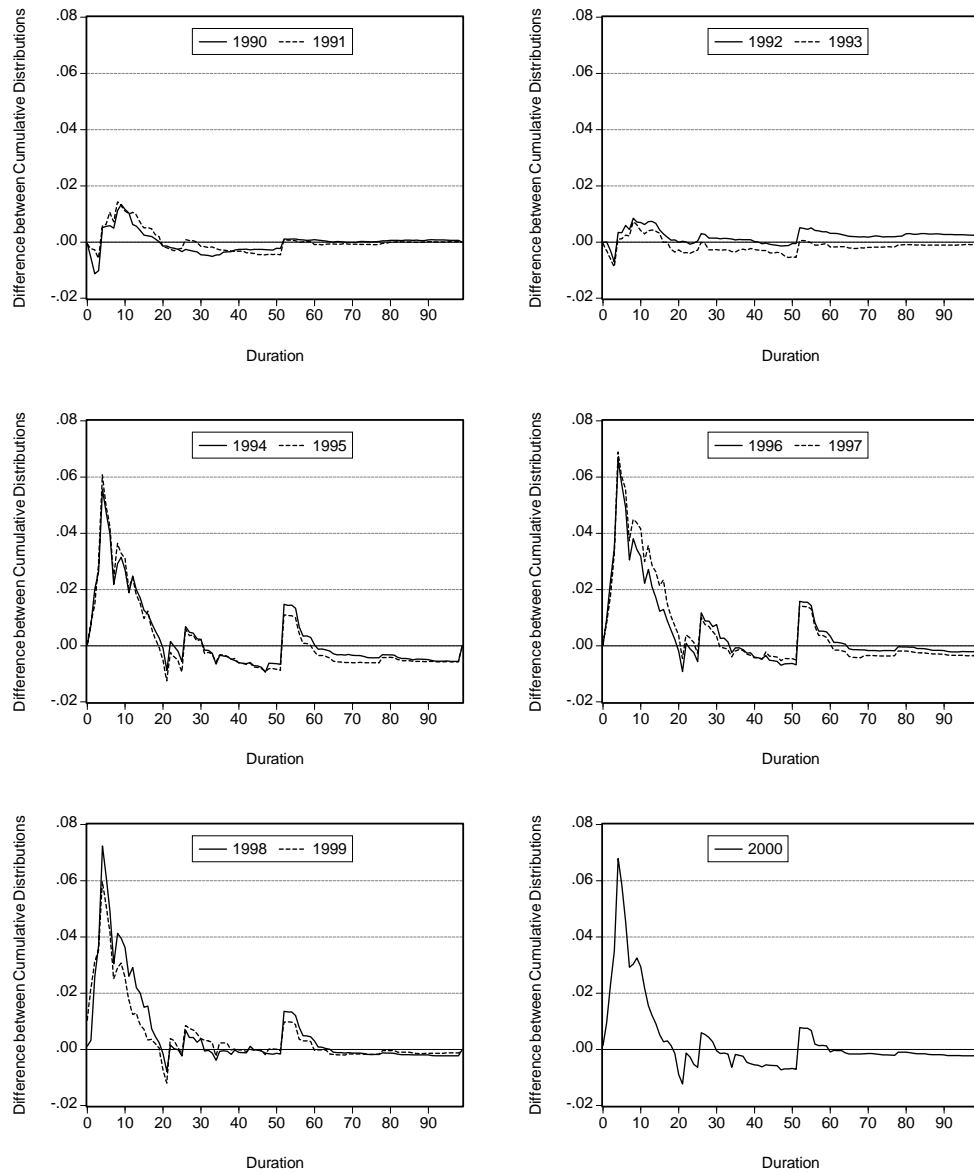


Figure 5: The difference between the cumulative distribution of unemployment durations for individuals in the first and fifth months in the sample compared to the entire survey sample in the four years before and the seven years since the CPS redesign. Figures are annual averages. The cumulative distribution is the fraction of unemployed workers with unemployment duration less than or equal to x weeks. Data are our calculations from the CPS.

Residual Unemployment Duration and Rates

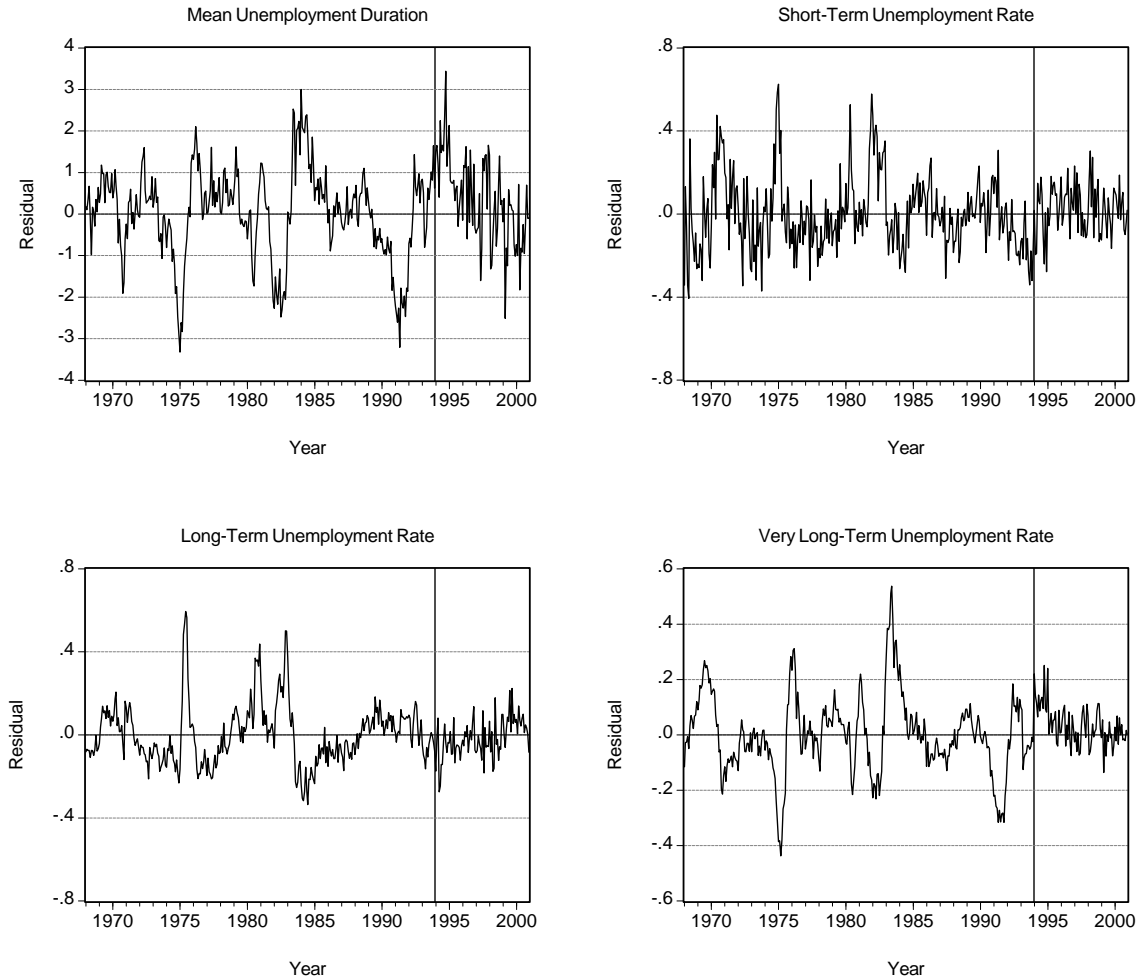


Figure 6: Residuals from regressions of the detrended mean unemployment duration, short-term unemployment rate, long-term unemployment rate, and very long-term unemployment rate on the detrended employment-population ratio. The employment level in the latter ratio comes from the Current Employment Statistics (CES) survey. The population is the over-16 population used by the CPS. We computed the dependent variables from the CPS. They are seasonally adjusted using a ratio to moving average procedure. For 1994 through 2000, only the incoming rotation groups are included. All series are detrended using an HP filter with smoothing parameter 1,440,000, which eliminates only very low frequency fluctuations. Vertical lines indicate the timing of the CPS redesign.

Unemployment Duration and Rates Adjusted for Measurement Changes

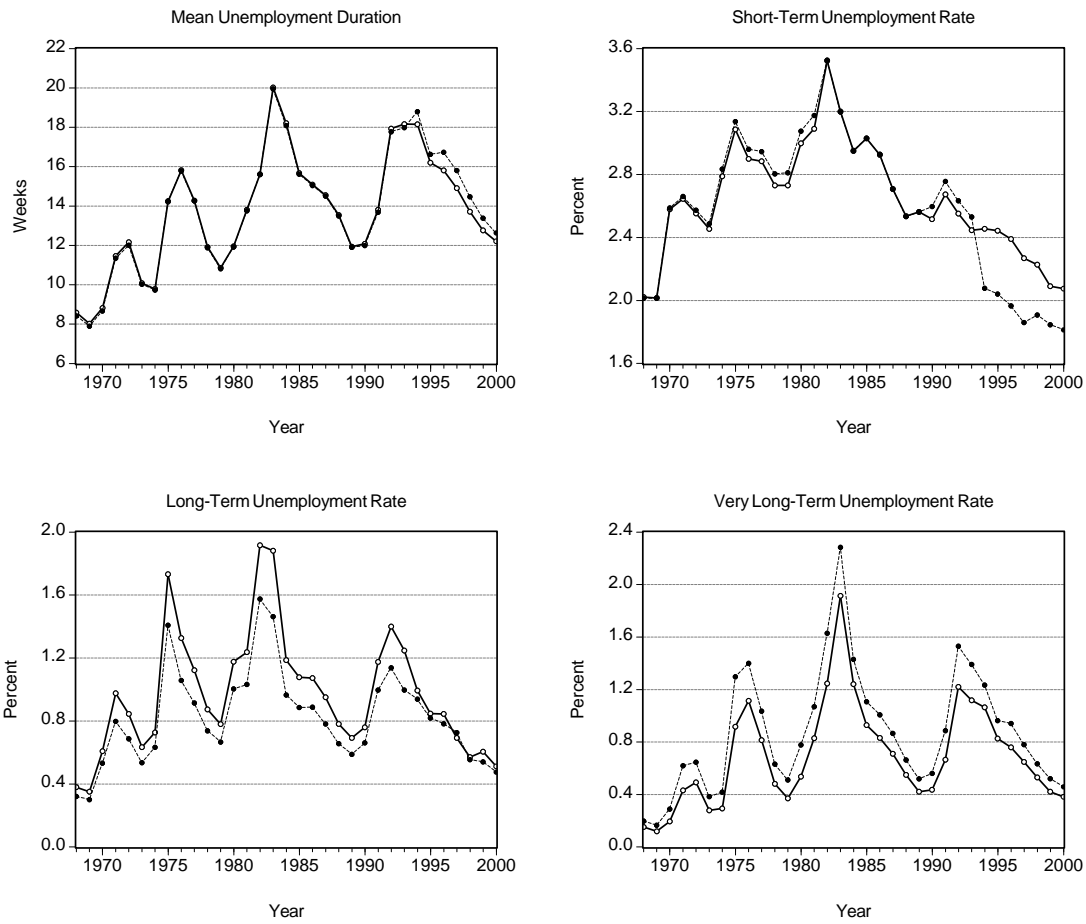


Figure 7: The top left panel shows the mean duration of unemployment spells currently in progress. The remaining panels show the share of the labor force with 0 to 4, 15 to 32, and 33 or more weeks unemployment duration (short-term, long-term, and very long-term unemployment rates, respectively). We constructed the solid line in each panel from the CPS, except for values before 1975, which are constructed from BLS weekly unemployment-duration data. From 1994 to 2000, the time series use only the incoming rotation groups. The dashed line indicates the corresponding official time series, with long-term unemployment defined as 15 to 26 weeks and very long-term unemployment defined as 27 or more weeks (see Figure 2).



Figure 8: This figure shows the unemployment rate adjusted for measurement changes and changes in the age composition of the labor force (the solid line) and unemployment rate adjusted only for measurement changes (the dashed line). Data are from official Bureau of Labor Statistics time series and our computations.

Age-Adjusted Unemployment Duration and Rates

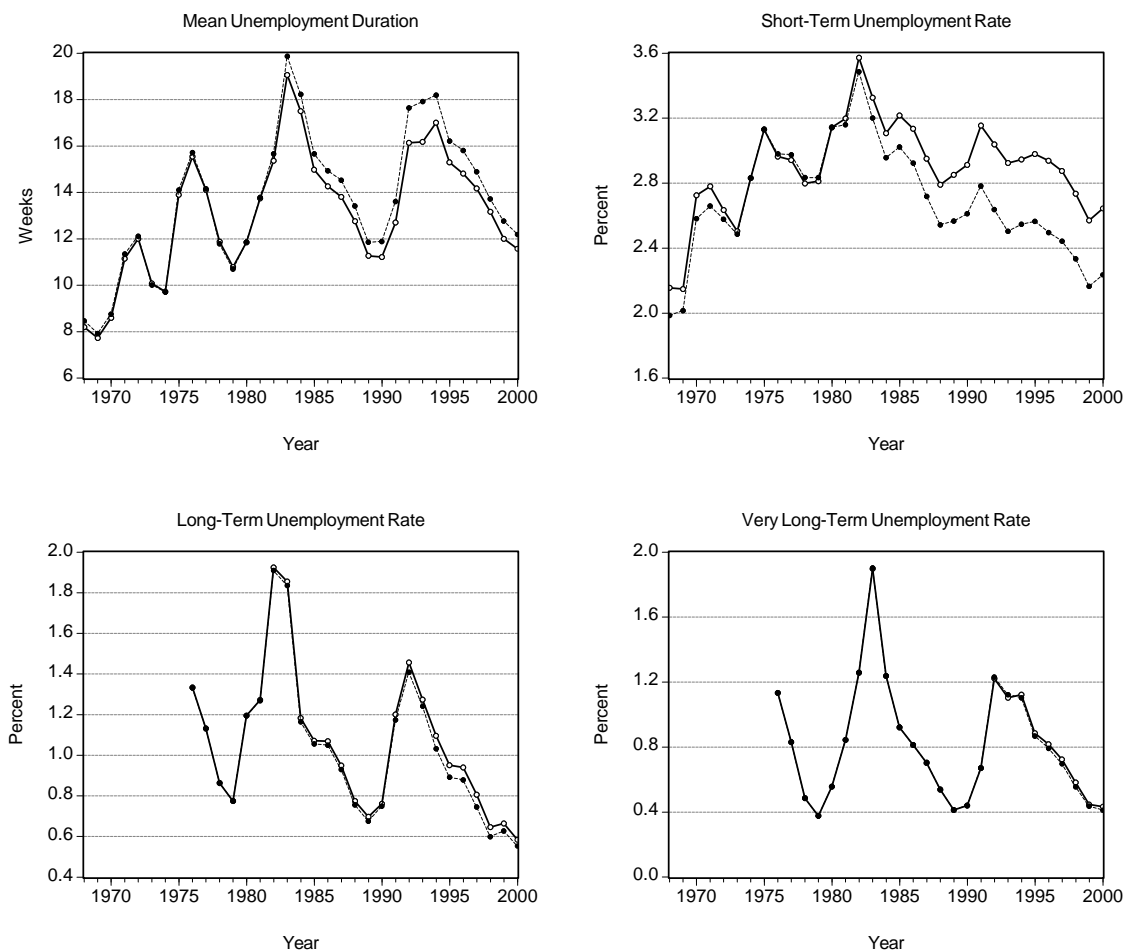


Figure 9: The top left panel shows the mean duration of unemployment spells currently in progress. The remaining panels show the share of the labor force with 0 to 4, 15 to 32, and 33 or more weeks unemployment duration (short-term, long-term, and very long-term unemployment rates, respectively). The solid line in each panel shows measurement- and age-adjusted time series, while the dashed line indicates the series adjusted only for measurement changes (see Figure 7). The time series from 1968 through 1975 were constructed using unpublished BLS data; information on the number of people unemployed 15 to 32 and 33 plus weeks was not available by age. From 1976 forward, the time series were constructed directly from the CPS. From 1994 forward, the time series use only the incoming rotation groups.

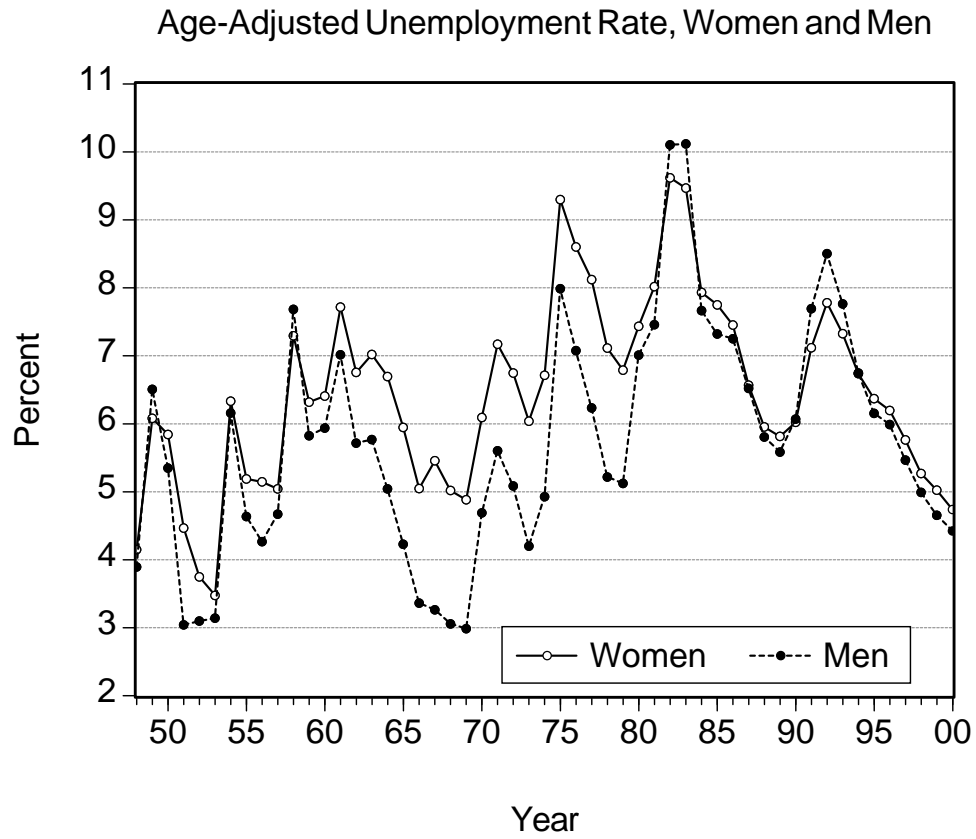


Figure 10: Unemployment data are from the official Bureau of Labor Statistics time series and our calculations.

Age-Adjusted Unemployment Duration and Rates: Men

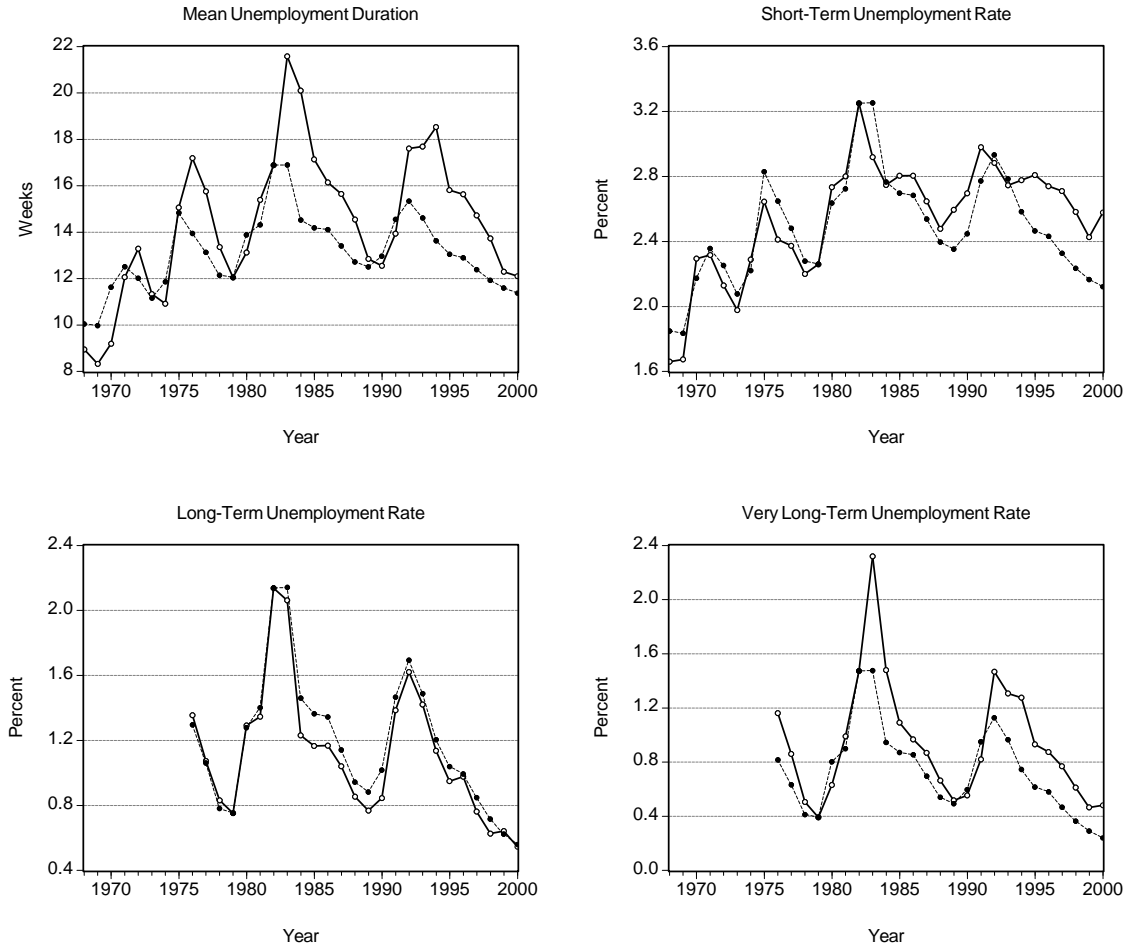


Figure 11: The top left panel shows the measurement- and age-adjusted mean duration of unemployment spells currently in progress for men. The remaining panels show the similarly adjusted shares of men in the labor force with 0 to 4, 15 to 32, and 33 or more weeks unemployment duration (short-term, long-term, and very long-term unemployment rates, respectively). The dashed line in each panel shows men's adjusted unemployment rate on a different scale. Unemployment data are from the official Bureau of Labor Statistics time series and our calculations. Duration data from 1968 through 1975 were constructed using unpublished BLS data; information on the number of people unemployed 15 to 32 and 33 or more weeks was not available by age. From 1976 forward, the time series were constructed directly from the CPS. From 1994 forward, the time series use only the incoming rotation groups.

Age-Adjusted Unemployment Duration and Rates: Women

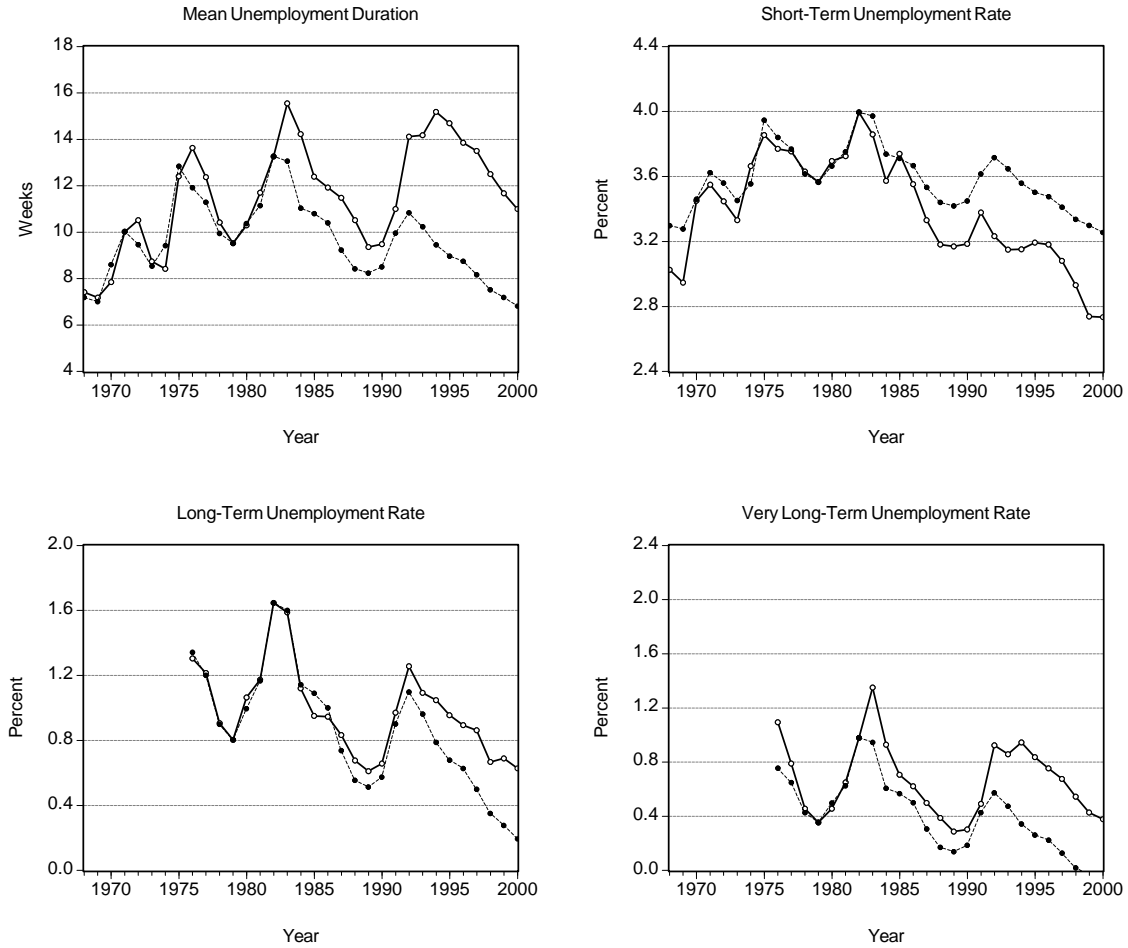


Figure 12: The top left panel shows the measurement- and age-adjusted mean duration of unemployment spells currently in progress for women. The remaining panels show the similarly adjusted shares of women in the labor force with 0 to 4, 15 to 32, and 33 or more weeks unemployment duration (short-term, long-term, and very long-term unemployment rates, respectively). The dashed line in each panel shows women's adjusted unemployment rate on a different scale. Unemployment data are from the official Bureau of Labor Statistics time series and our calculations. Duration data from 1968 through 1975 were constructed using unpublished BLS data; information on the number of people unemployed 15 to 32 and 33 or more weeks was not available by age. From 1976 forward, the time series were constructed directly from the CPS. From 1994 forward, the time series use only the incoming rotation groups.

Age-Adjusted Monthly Labor Force Transition Rates

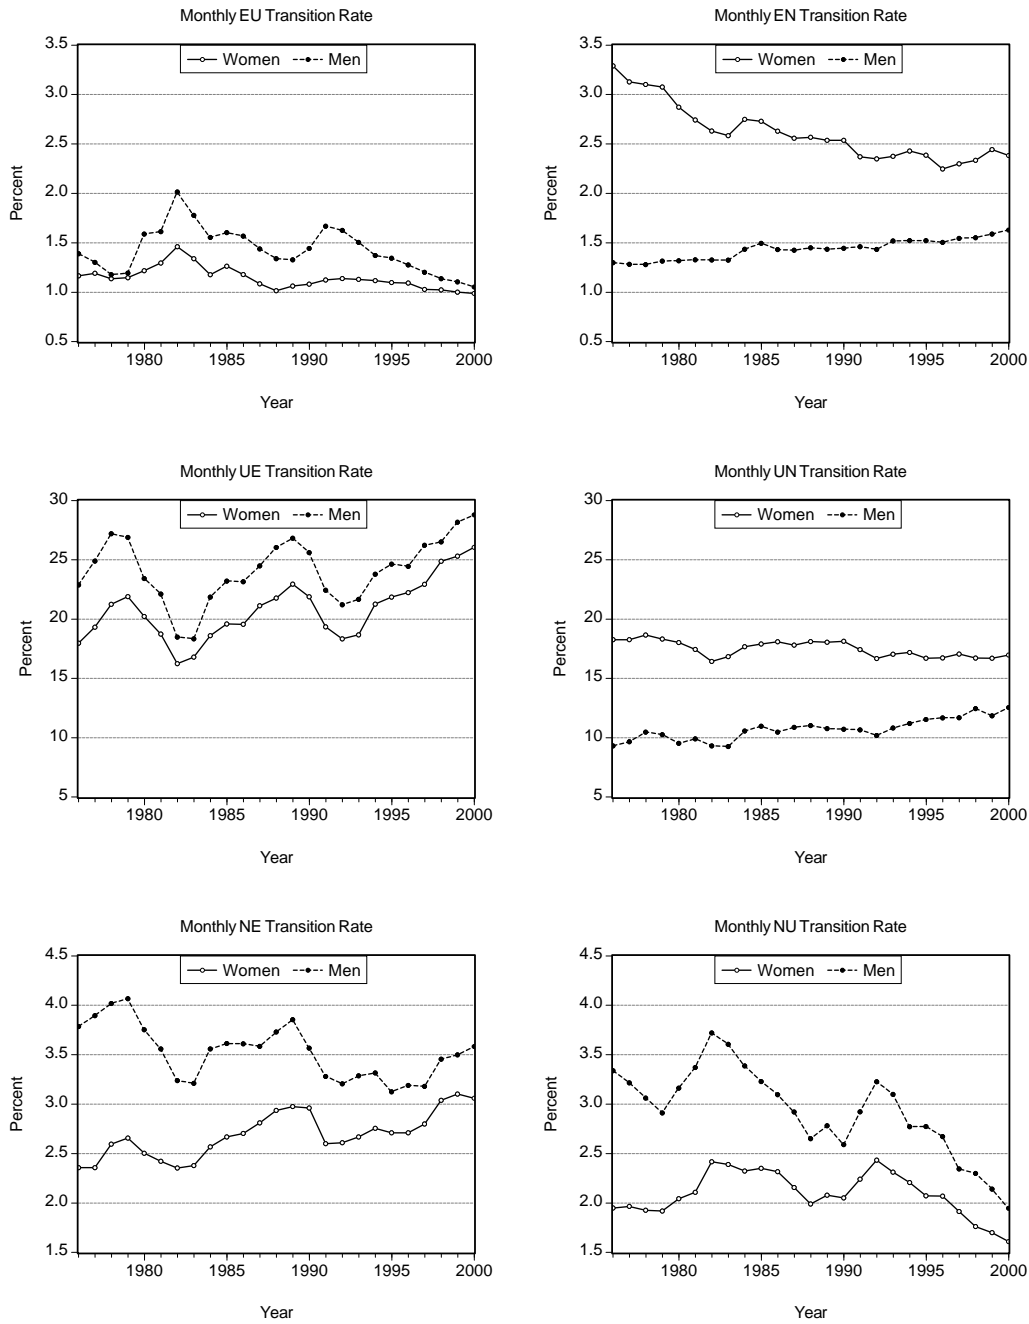


Figure 13: The figure shows women’s and men’s labor-market transitions from 1976 to 2000. All transition rates are adjusted using the Abowd-Zellner correction. λ_{un} and λ_{nu} are adjusted for the effects of the 1994 CPS redesign (see Appendix B for details). Data are our calculations from the CPS.

Unemployment Rates Implied by Labor Market Transitions

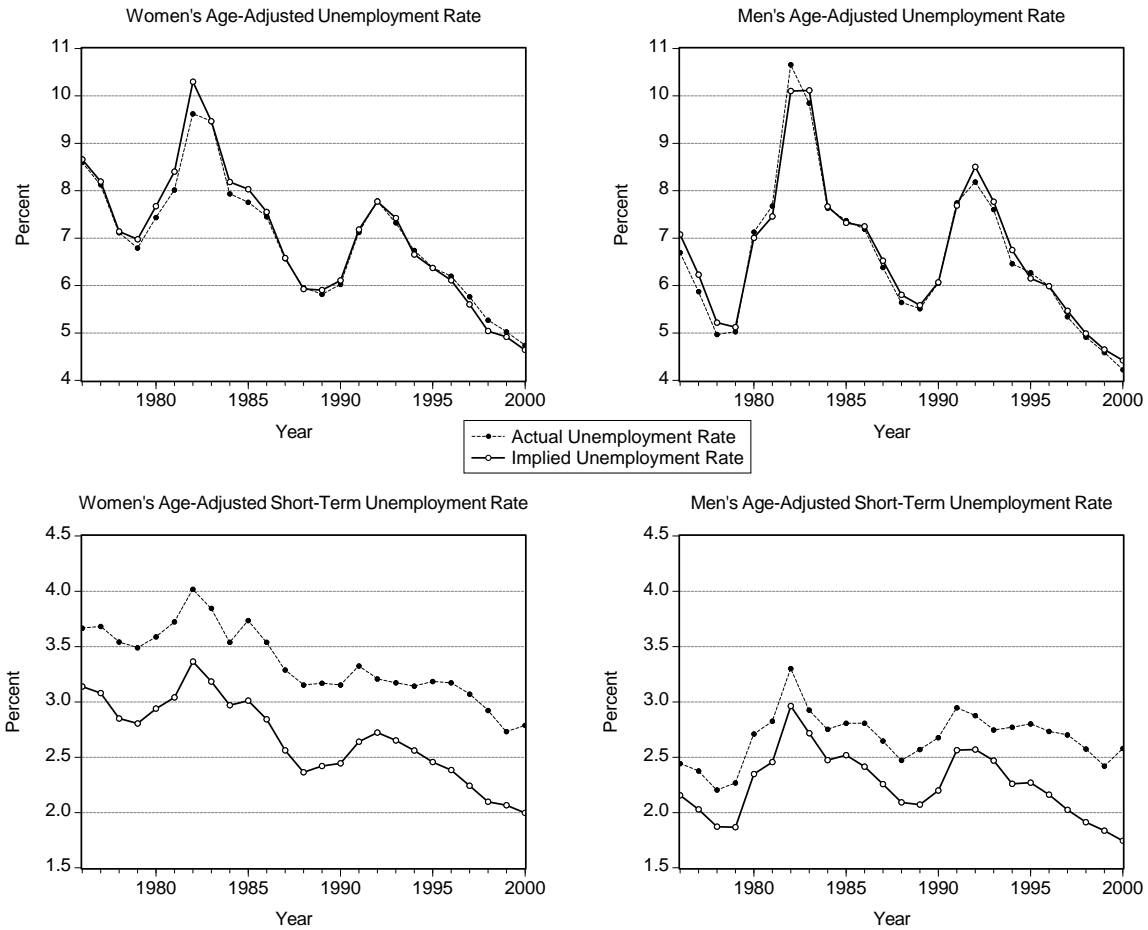


Figure 14: The top row shows the steady state unemployment rate implied by age-adjusted labor-market transition rates and the actual age-adjusted unemployment rate for women and men, 1976 to 2000. The bottom row shows implied and actual short-term unemployment rates. All flows are adjusted using the Abowd-Zellner correction. λ_{un} and λ_{nu} are adjusted for the effects of the 1994 CPS redesign (see Appendix B for details). Data are our calculations from the CPS.

Effect of Counterfactual Experiments on Unemployment Rates

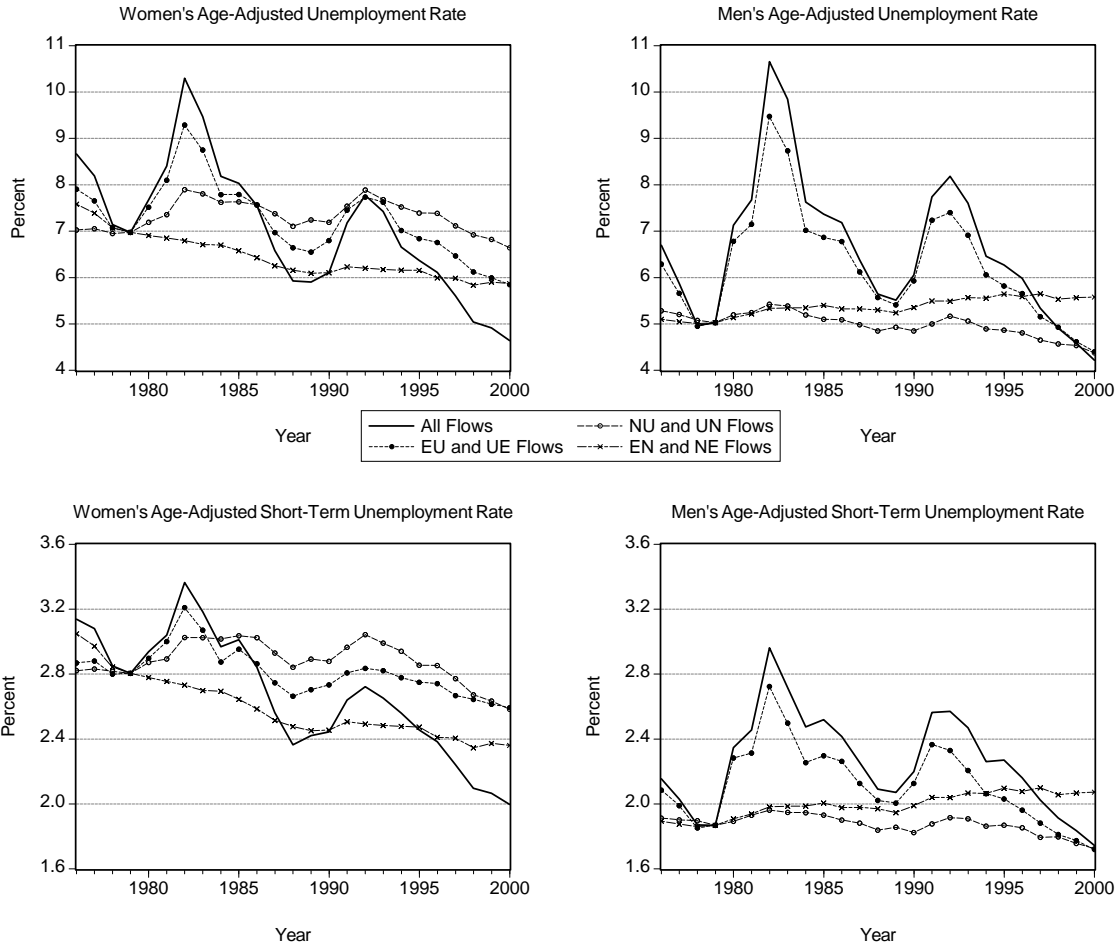


Figure 15: The figure shows the steady state unemployment rate and short-term unemployment rate implied by age-adjusted labor market flows, women and men, 1976 to 2000, allowing for time-variation in all transitions, or only in λ_{eu} and λ_{ue} , in λ_{un} and λ_{nu} , or in λ_{en} and λ_{ne} . All flows are adjusted using the Abowd-Zellner correction. λ_{un} and λ_{nu} are adjusted for the effects of the 1994 CPS redesign (for details, see Appendix B). Data are our calculations from the CPS.

Age-Adjusted Share of Employed Workers
with at least Ten Years Tenure

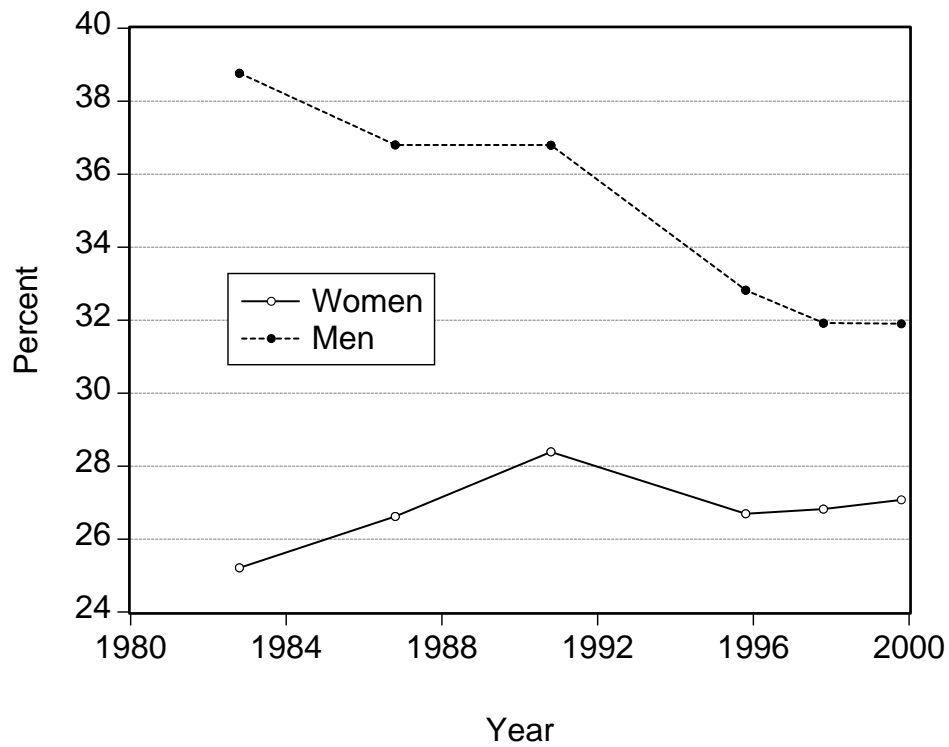


Figure 16: Age-adjustment weights use the 1983 employment shares of nine age groups: 25-29, 30-34, . . . , 60-64, and 65 plus. Data are our calculations from the January 1983, January 1987, January 1991, February 1996, February 1998, and February 2000 Job Tenure Supplements to the CPS.

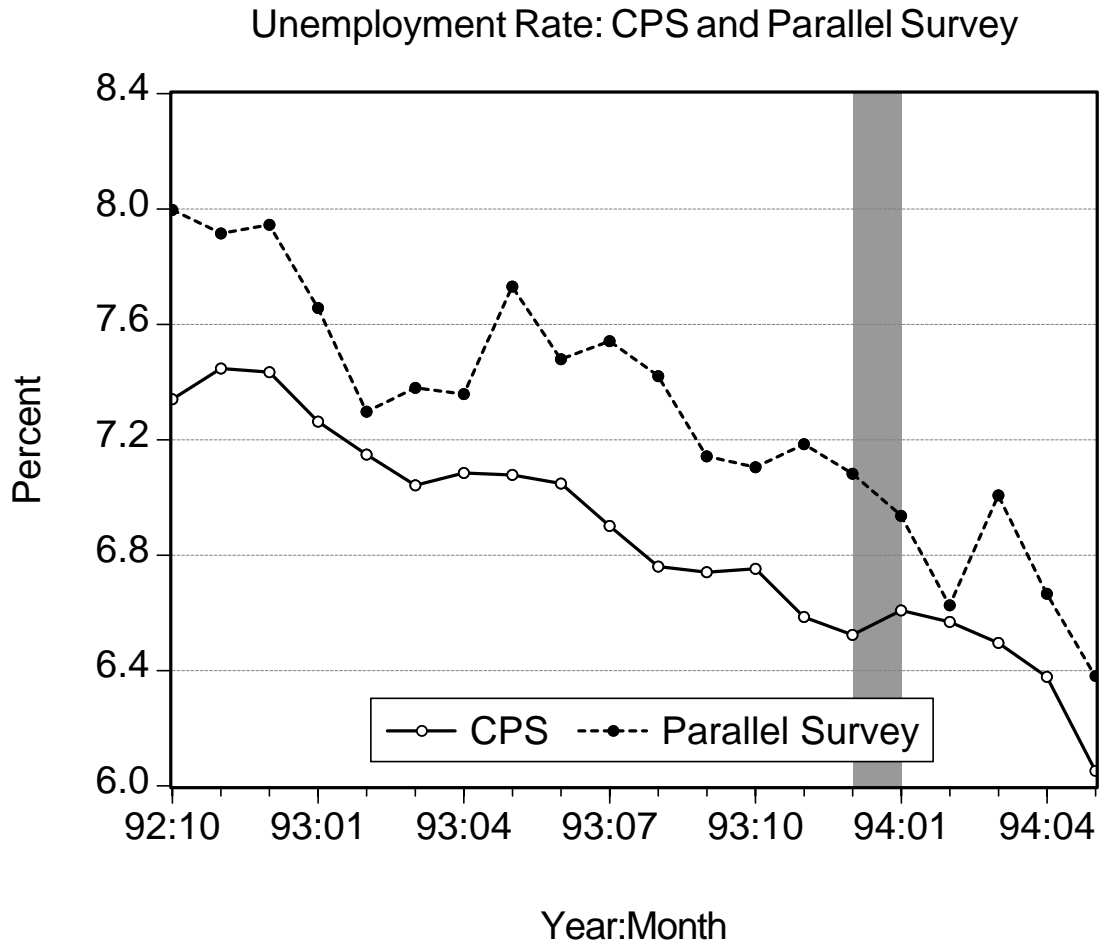


Figure 17: This figure shows the seasonally-adjusted unemployment rate in the CPS and in the parallel survey designed to assess the effect of the CPS redesign implemented in January 1994. The vertical bar indicates the timing of the CPS redesign. Anne Polivka extracted the parallel survey data and we constructed the time series.

Unemployment Duration and Rates: CPS and Parallel Survey

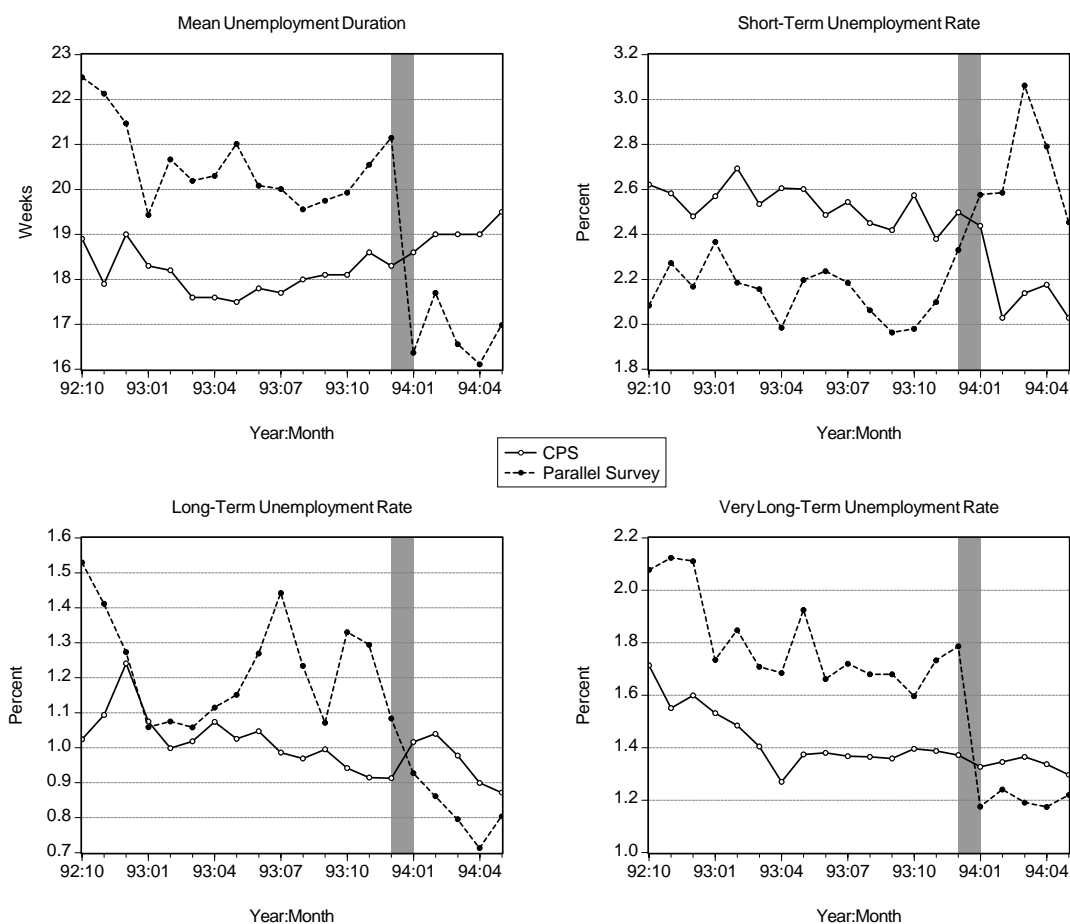


Figure 18: The top left panel shows the mean duration of unemployment spells currently in progress. The remaining panels show the share of the labor force with 0 to 4, 15 to 26, and 27 or more weeks unemployment duration (short-term, long-term, and very long-term unemployment rates, respectively). The solid line in each panel shows the seasonally-adjusted time series in the CPS, while the dashed line indicates the seasonally-adjusted time series in the parallel survey designed to assess the effect of the CPS redesign implemented in January 1994. The vertical bar indicates the timing of the CPS redesign. Anne Polivka extracted the parallel survey data and we constructed the time series.

Residual Worker Flows

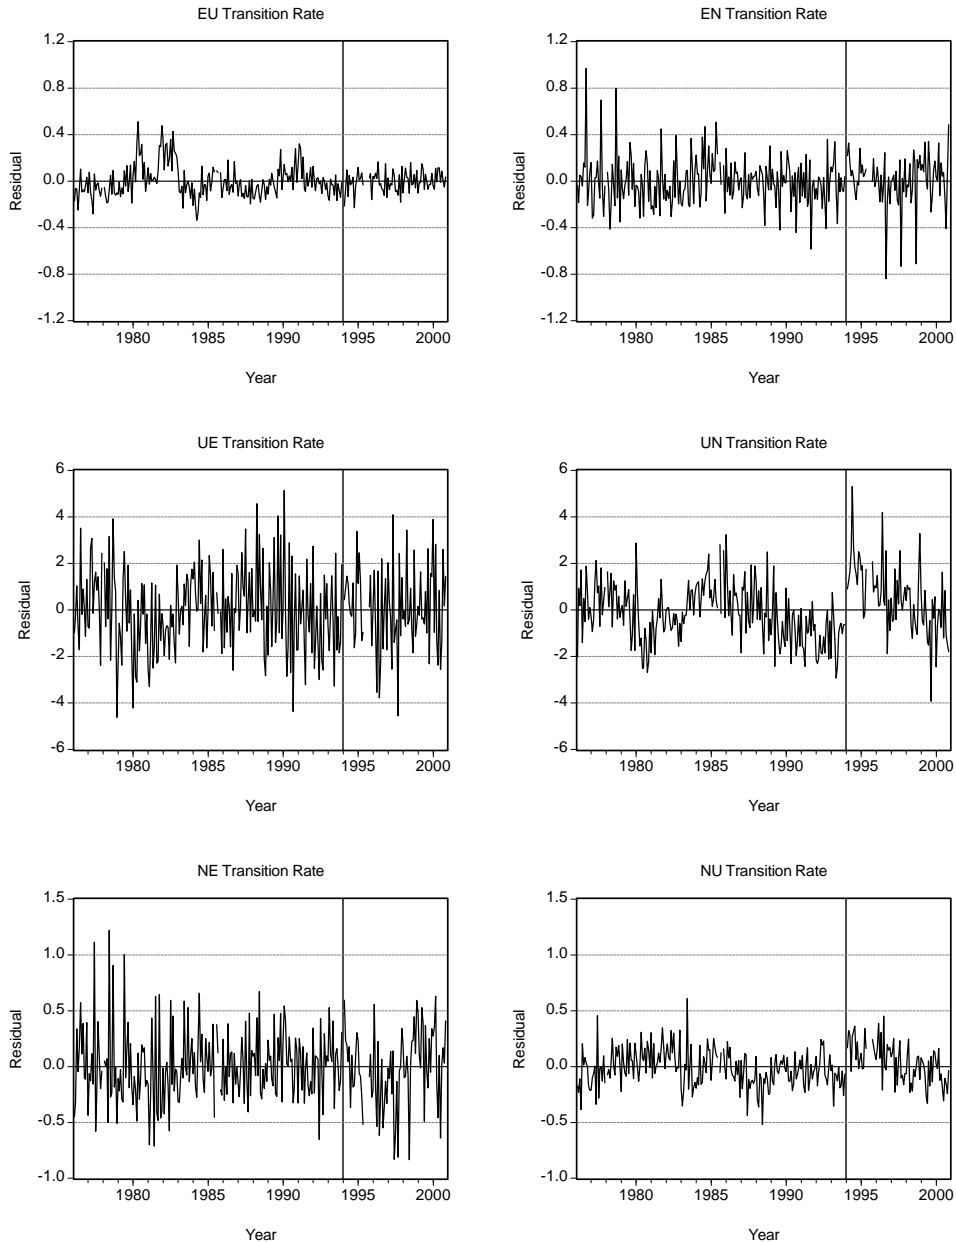


Figure 19: This figure shows residuals from regressions of detrended worker-flow rates on the detrended employment-population ratio. Worker-flow rates are our calculations from the CPS. The employment level in the latter ratio comes from the CES survey. The population is the over-16 population used for the CPS weights. All series are seasonally adjusted and detrended using a HP filter with smoothing parameter 1,440,000, which eliminates only very low-frequency fluctuations. Vertical lines indicate the timing of the CPS redesign. Breaks in the time series indicate missing observations.