The Cyclicality of Hires, Separations, and Job-to-Job Transitions

Robert Shimer*
University of Chicago
shimer@uchicago.edu

September 30, 2004

Abstract

This paper uses readily accessible data on unemployment duration to measure the rate at which unemployed workers are hired and the rate at which they involuntarily separate from their job. In the United States, the hiring rate is strongly procyclical and the separation rate is nearly acyclical, particularly during the last two decades. Using the underlying microeconomic data, the paper shows that these results are not due to compositional changes in the pool of searching workers, nor are they due to movements of workers in and out of the labor force. It then uses the measure of the hiring and separation rate, together with a canonical model of job-to-job transitions, to predict the job-to-job transition rate. The results are quantitatively consistent with the best available data from the last decade.

*This paper was prepared for the Federal Reserve Bank of St. Louis's 29th annual economic policy conference, October 21–22, 2004. It is supported by grants from the National Science Foundation and the Sloan Foundation.

1 Introduction

The modern theory of unemployment recognizes the conceptual value of decomposing employment and unemployment fluctuations into changes in the rate at which workers are hired and changes in the rate at which workers separate from their old jobs. If most fluctuations are a consequence of changes in the separation rate, economists are naturally led to think that certain theories of the business cycle are more plausible than others. A well-known example is Lilien's (1982) 'sectoral shifts' hypothesis, which posits that recessions are periods when some industries suffer a sharp contraction, forcing workers to move to other sectors after experiencing a spell of unemployment. If this is correct, the separation rate should be countercyclical and the hiring rate approximately acyclical. Job-to-job transitions will also be countercyclical if some workers can switch sectors without an intervening jobless spell. Keynesian rigid wage models typically lead to the same conclusion. During recessions the (nominal) wage exceeds the (nominal) marginal product of labor and a (nominal) rigidity prevents firms from reducing their workers' wages. Instead, the firm lays off workers until the marginal product of labor rises to the prevailing wage rate, generating countercylical separations. The cyclicality of the hiring rate depends on whether the wage rigidity affects new employment relationships. If new employment relationships sidestep the rigidity, recessions are good times for firms to hire, and so the model predicts a countercyclical hiring rate.

An alternative view is that recessions are times when firms do not hire many workers. For example, in the presence of large hiring or firing costs, a firm may choose not to increase its separation rate above the natural attrition rate, but will instead reduce the rate at which it hires workers until the wage equals the marginal product of labor (Bentolila and Bertola 1990). Of course, the conventional view is that turnover costs are unlikely to be important in a flexible economy like that of the United States (Siebert 1997). Another possibility advocated by Hall (2004a) is that during a recession, the marginal product of labor declines but a social norm keeps the real wage unchanged. Because hiring costs are sunk, firms do not lay off their existing workers. Instead, they reduce their recruiting effort, lowering the rate at which unemployed workers are hired. Hall does not explain why the social norm leaves real wages unchanged, and indeed many other social norms exist in his model. More recent research has suggested that informational issues within the employment relationship may provide a satisfactory resolution of this puzzle (Kennan 2004, Shimer and Wright 2004).

The goal of this paper is to document the cyclicality of the hiring and separation rates in the United States from 1948 to 2004. As much as possible, I try to use aggregate data that are readily available so comparable measures can easily be constructed in other countries.[1] I show how data on unemployment duration constructed by the Bureau of Labor Statistics (BLS) from the Current Population Survey (CPS) can be used to measure both the hiring and separation results. I find that there are substantial fluctuations in the hiring rate—the monthly probability with which a typical unemployed worker finds a job—at business cycle frequencies (Figure 1), while the fluctuations in the separation rate—the monthly probability that a typical employed worker is forced to leave her job—is comparative acyclic (Figure 7). This is particularly true in the last two decades, during which period the separation rate has steadily declined despite two spikes in the unemployment rate. I also use the underlying household data from the CPS to show that these results are not a consequence of cyclical changes in the composition of the unemployed population, nor are they due to cyclical movements of workers in and out of the labor force. For example, Figure 4 shows that most of the increase in unemployment during 2001 and 2002 was due to growth in the number of 'other job losers'—job losers who are not on layoff—and that this group's hiring rate is extremely procyclical, rising by 50 percent from 1994 to 1999 and then falling by as much as 70 percent during the subsequent downturn.

I then put these measures of the hiring and separation rates in a canonical model of job-to-job movements to explore another prediction of procyclical hiring and acyclical separation rates. I assume workers switch employers between months t and $t+1$ either when they find a better job during month t or when they are forced to leave their old job during month t but manage to find a new one before the survey date in month $t+1$. The former corresponds to a voluntary quit while the latter is a subset of the layoffs. I feed the time series for the hiring and separation rates into this model and show that it predicts a strongly procyclical job-to-job transition rate (Figure 14), which is quantitatively consistent with available evidence from the last recession (Figure 15). In contrast, if separations were countercyclical and the hiring rate acyclic, as in the Keynesian or Sectoral Shifts view, the canonical model would predict a countercyclical job-to-job transition rate both because the number of 'involuntary job changers'—workers who are forced to leave their old job but find a new one within the month—increases during a recession and because the increase in the separation rate reduces the average quality of job matches, placing more workers in a position where they are eager to switch employers.

This paper is related to an older literature that decomposes unemployment fluctuations

[1] But so far I have only constructed these measures for the United States.

into changes in duration and changes in incidence (Sider 1982, Darby, Haltiwanger, and Plant 1986). Partially because of different methodologies and partially because of different sample periods, that literature left a much greater role for separations (or unemployment incidence) than I find here. Besides the use of newer data, this paper has a few contributions to offer. First, I try to use publicly available data whenever possible, focusing on unemployment duration data throughout. This makes it easy to verify and extend my results both to other time periods and potentially to other countries. Second, I stress the role of heterogeneity throughout my analysis, arguing that changes in the composition of the unemployed population do not drive these results. Third, I emphasize that because of time aggregation, the separation rate and transition rate from employment to unemployment are not the same. A household survey will observe a worker moving from employment to unemployment only if the worker both separates from her job and fails to find a new job before the next survey date. If the hiring rate is procyclical, then the probability that a worker who separates from her job fails to find another one before the survey date will be countercyclical, causing fluctuations in the transition rate from employment to unemployment even if the separation rate is constant. And finally, I connect my analysis of transitions between unemployment and employment to a canonical model of transitions between jobs. Perhaps my most surprising finding is that by feeding the time series that I construct for the hiring rate and the separation rate into this simple model, I recover a measure of job-to-job movements that is quantitatively consistent with the observed pattern in recent United States data. Whatever forces make it harder for an unemployed worker to find a job during a recession also seem to make it harder for an employed worker to find a better job.

Section 2 discusses measurement of the hiring rate, including compositional issues, in detail. Section 3 turns to the separation rate, demonstrating the importance of time-aggregation, particularly during the last two downturns. In both of these sections, I use a simple 'bathtub' model of the labor market, distinguishing between the flow of workers in and out of unemployment to identify the hiring and separation rates. Section 4 complicates the model by introducing a reason for job-to-job movements: some jobs are more desirable than others. This canonical model predicts the job-to-job transition rate as a function of current and past hiring and separation rates, but independent of model details like the distribution of jobs' desirability, and I show that this prediction is consistent with data from the last decade in the United States. Section 5 concludes.

2 The Hiring Rate

If all unemployed workers were equally likely to find a job in a given month, it would be straightforward to use unemployment duration data to compute the hiring rate. But in reality, some workers, e.g. the long-term unemployed or those on layoff, are less likely to find a job than others, making even the definition of *the* hiring rate murky. This section proposes a simple measure of the mean hiring rate h_t, defined as the fraction of workers who are unemployed in month t and employed in month $t+1$, using two easily accessible time-series: the number of unemployed workers and the number of workers unemployed for less than one month. I show that the mean hiring rate is procyclical and volatile.

I then compare my measure with two alternatives proposed by Shimer (2004) and Hall (2004a). If workers were homogeneous, the three measures would be identical, but since workers are heterogeneous, neither alternative measures the fraction of unemployed workers who find a job in month t. Nevertheless, it is reassuring that all three series have similar cyclical properties. I next argue that fluctuations in h_t are not driven primarily by changes in the composition of the unemployed population, as suggested by Darby, Haltiwanger, and Plant (1985) and (1986), but rather by changes in the hiring rate for each type of worker. Finally, I consider the bias in the hiring rate introduced by entry and exit from the labor force and explain why this too has quantitatively small effects on my conclusions.

2.1 A Simple Measure

To start, I maintain the fiction that all unemployed workers find a job with a common probability h_t in month t and that no unemployed worker exits the labor force. Let u_t denote the number of unemployed workers and u_t^s denote the number of short-term unemployed, those whose current spell is shorter than one month in duration. The number of unemployed workers next month is equal to the number of unemployed workers this month who fail to find a job plus the number of newly unemployed workers:

$$u_{t+1} = u_t(1 - h_t) + u_{t+1}^s. \tag{1}$$

This is easily solved for the hiring rate,

$$h_t = 1 - \frac{u_{t+1} - u_{t+1}^s}{u_t}. \tag{2}$$

This measure is attractive because it is easily computed using commonly available data. For the United States, the BLS publishes monthly time series based on the underlying microeconomic data from the CPS on the total number of unemployed workers and the number of workers unemployed for zero to four weeks, corresponding to u_t and u_t^s, respectively.[2]

A second attractive feature of this measure of the hiring rate is that if workers are heterogeneous, h_t measures the probability that the average unemployed worker in month t finds a job. Index the u_t unemployed workers at time t by $i \in \{1, \ldots, u_t\}$. Let h_t^i denote the probability that worker i finds a job in month t. Then assuming idiosyncratic uncertainty cancels out in the aggregate, one can generalize equation (1) to the case where h_t^i varies with i:

$$u_{t+1} = \sum_{i=1}^{u_t} (1 - h_t^i) + u_{t+1}^s.$$

Rearranging this equation gives

$$\frac{\sum_{i=1}^{u_t} h_t^i}{u_t} = 1 - \frac{u_{t+1} - u_{t+1}^s}{u_t}.$$

Since the right hand side of this expression is identical to the right hand side of equation (2), the left hand sides must be the same as well, i.e. h_t measures the mean hiring rate among unemployed workers in month t:

$$h_t = \frac{\sum_{i=1}^{u_t} h_t^i}{u_t}. \tag{3}$$

I can therefore measure the mean hiring rate using data on the the number of unemployed workers and the number of short-term unemployed workers even if every worker has a different, and potentially time-varying, hiring rate.

Figure 1 shows the time series for the hiring rate, constructed according to equation (2), using United States data from 1948 to 2004. For comparison, I also plot the unemployment rate. Three facts stand out. First, the hiring rate is high, averaging 0.44 per month during the past 56 years and never falling below 31 percent per month. Second, it is negatively correlated with unemployment at cyclical frequencies. For example, after detrending both data series using a low frequency filter,[3] the correlation between the two series is -0.94.

[2]The data can be downloaded from the BLS web site `http://www.bls.gov/` or from the St. Louis Federal Reserve Economic Database (FRED®II) `http://research.stlouisfed.org/fred2/`.

[3]Throughout this paper I time-aggregate underlying monthly data to get quarterly averages, removing substantial low frequency fluctuations that likely reflect measurement error in the CPS. I then detrend the quarterly data using an HP filter with smoothing parameter 10^5. This is a much lower frequency filter than is commonly used in business cycle analyses of quarterly data. A standard filter seems to remove much of

Third, the hiring rate is volatile, with a standard deviation of the detrended series equal to 0.12. For comparison, the standard deviation of the detrended unemployment rate is seventy percent larger, 0.20.

[Figure 1 about here.]

2.2 Alternative Measures

If unemployed workers were homogeneous, there would be other valid methods of constructing the hiring rate. In Shimer (2004), I proposed one that uses data on mean unemployment duration d_t. Mean unemployment duration in month $t+1$ can be expressed as a weighted average of the mean unemployment duration of previously-unemployed workers who failed to get a job in month t and the unemployment duration of newly-unemployed workers,

$$d_{t+1} = \frac{(d_t + 1)(1 - \tilde{h}_t)u_t + \big(u_{t+1} - (1 - \tilde{h}_t)u_t\big)}{u_{t+1}}, \tag{4}$$

where $\tilde{h}_t$ is the alternative measure of the hiring rate. There are $(1 - \tilde{h}_t)u_t$ unemployed workers, with mean unemployment duration d_t, who fail to get a job in month t. The mean unemployment duration for these workers increases by one month to d_t+1. In addition, there are $u_{t+1} - (1 - \tilde{h}_t)u_t$ newly unemployed workers in month $t+1$, each of whom has an unemployment duration of one month. This equation can be solved for the hiring rate as a function of the current and future mean unemployment duration and number of unemployed workers,

$$\tilde{h}_t = 1 - \frac{(d_{t+1} - 1)u_{t+1}}{d_t u_t}. \tag{5}$$

In steady state, $u_t = u_{t+1}$ and $d_t = d_{t+1}$, so equation (5) reduces to $\tilde{h} = 1/d$, a familiar relationship for a variable with a constant arrival rate. More generally, one can use mean unemployment duration and the number of unemployed workers, both constructed by the BLS from the CPS, to compute $\tilde{h}_t$ from 1948 to 2004.

But if workers are heterogeneous, $\tilde{h}_t$ does not recover the mean hiring rate of unemployed workers. To see this, again index the u_t unemployed workers in month t by $i \in \{1, \ldots, u_t\}$. Suppose worker i has unemployment duration d_t^i and finds a job with probability h_t^i. By definition, the mean unemployment duration in month t is $d_t \equiv \frac{1}{u_t}\sum_{i=1}^{u_t} d_t^i$. Generalizing

the cyclical volatility in the variable of interest.

equation (4) to allow for heterogeneous workers, we find that mean unemployment duration in month $t+1$ will be

$$d_{t+1} = \frac{\sum_{i=1}^{u_t}(d_t^i+1)(1-h_t^i) + \left(u_{t+1} - \sum_{i=1}^{u_t}(1-h_t^i)\right)}{u_{t+1}}.$$

The first term in the numerator is the unemployment duration of a previously unemployed worker i if she remains unemployed in month $t+1$. The second term is the number of newly unemployed workers in month t. This is averaged across unemployed workers to get the mean unemployment duration next month. With some algebra, this equation may be rewritten as

$$\frac{\sum_{i=1}^{u_t} d_t^i h_t^i}{\sum_{i=1}^{u_t} d_t^i} = 1 - \frac{(d_{t+1}-1)u_{t+1}}{d_t u_t}.$$

Comparing this with equation (5) yields

$$\tilde{h}_t = \frac{\sum_{i=1}^{u_t} d_t^i h_t^i}{\sum_{i=1}^{u_t} d_t^i}.$$

This is a weighted average of the individual hiring rates h_t^i, where the weight accorded to individual i is her unemployment duration d_t^i. Compared to the true mean hiring rate, this measure overweights the long-term unemployed, which tends to reduce the measured hiring rate. In fact, $\tilde{h}_t$ averages approximately 35 percent from 1948 to 2004, significantly lower than the 44 percent for the mean hiring rate h_t.

Nevertheless, there is one important advantage to Shimer's (2004) measure. If unemployment were constant between months t and $t+1$, the measured hiring rate in equation (5) would reduce to

$$\tilde{\tilde{h}}_t = 1 - \frac{d_{t+1}-1}{d_t},$$

a function of unemployment duration alone. In fact, the pairwise correlations of the detrended series for h, $\tilde{h}$, and $\tilde{\tilde{h}}$ all exceed 0.87, indicating that fluctuations in the measured hiring rate, and hence the correlation between the hiring rate and the unemployment rate, are primarily driven by fluctuations in unemployment duration, not in the level of unemployment itself.

Hall (2004a) proposes a third measure of the hiring rate $\hat{h}_t$. Let u_t^m denote the number of medium-term unemployed workers. Because of data availability, this is defined as workers who have already experienced 1 or 2 months of unemployment. This is equal to the number of short-term unemployed in previous months who have failed to find a job:

$$u_{t+1}^m = \big(u_t^s + u_{t-1}^s(1-\hat{h}_{t-1})\big)(1-\hat{h}_t). \tag{6}$$

This is a first order difference equation for $\hat{h}$. With a reasonable initial guess, e.g. that $\hat{h}_t$, u_t^s, and u_t^m were constant before 1948, one can solve this equation forward and compute $\hat{h}_t$.[4] If all unemployed workers have the same hiring rate h_t at every point in time, Hall's (2004a) method will uncover that hiring rate. But if workers are heterogeneous, this measure captures only the hiring rate of the short-term unemployed and hence is likely to yield an estimate that exceeds the true mean hiring rate h_t. The data supports this hypothesis; the mean value from 1948 to 2004 of $\hat{h}_t$ is 53 percent, compared to 44 percent for h_t. Figure 2 shows all three proposed measures of the hiring rate. Reassuringly, their cyclical behavior is extremely similar.

[Figure 2 about here.]

2.3 Compositional Effects

Equation (3) indicates two distinct explanations for why the mean hiring rate h_t declines when unemployment is high. Either the hiring rate declines for each worker or the unemployment pool shifts disproportionately towards workers with a low hiring rate. Darby, Haltiwanger, and Plant (1985) and (1986) advance the second possibility in their exploration of the behavior of unemployment duration. They argue that there are two types of workers. The first type experiences frequent short spells of unemployment. The second type, including prime aged workers and those on layoff, experiences unemployment infrequently and takes a long time to find a new job. If recessions are periods when disproportionately many of the second type of worker lose their job, then the measured hiring rate may decrease even though h_t^i does not change for any particular worker. Following Baker (1992), I refer to this as the 'heterogeneity hypothesis'.[5]

[4] Hall (2004a) assumes instead that the economy is in steady state in each month, replacing equation (6) with

$$u_{t+1}^m = \big(u_t^s + u_{t-1}^s(1-\hat{\hat{h}}_t)\big)(1-\hat{\hat{h}}_t).$$

He then solves this equation explicitly for $\hat{h}_t$. Given the slow evolution of $\hat{h}_t$, the results are very similar.

[5] Dynarski and Sheffrin (1990) and Baker (1992) show that unemployment duration is strongly countercyclical, and so the hiring rate is strongly procyclical, for all workers conditional on a broad set of characteristics, including the reason for unemployment, census region, sex, race, education, and previous industry. This leads Baker (1992) to conclude that "the heterogeneity explanation of aggregate variation sheds little

To see whether this argument is quantitatively important, it is necessary to put some structure on it. One approach would be to assume that each individual i has a time-varying hiring rate h_t^i and use repeated spells of unemployment for particular individuals in order to check how her hiring rate depends on aggregate labor market conditions. Unfortunately, I am unaware of a reliable representative data set for the United States with repeated spells. Instead, I assume that workers can be divided into J different groups, indexed by $j \in \{1, \ldots, J\}$. For example, the groups may correspond to different reasons for unemployment: job losers, job leavers, re-entrants, or new entrants. I assume that all workers within a group are identical. More precisely, let $u_{t,j}$ be the number of unemployed workers with characteristic j in month t and $h_{t,j}$ be the hiring rate of those workers, computed using a type-dependent analog of equation (2). By definition, the aggregate hiring rate is

$$h_t = \frac{\sum_{j=1}^{J} u_{t,j} h_{t,j}}{\sum_{j=1}^{J} u_{t,j}}.$$

If Darby, Haltiwanger, and Plant's heterogeneity hypothesis is correct, fluctuations in the hiring rate are due primarily to changes in the shares $u_{t,j}$ rather than in the type-specific hiring rates $h_{t,j}$. To see whether this is the case, one can construct two hypothetical measures. Let $\bar{h}_t^{\text{comp}}$ denote the change in the hiring rate due to changes in the composition of the work force and $\bar{h}_t^{\text{real}}$ denote the 'real' changes due to changes in the hiring rate for each type of worker:

$$\bar{h}_t^{\text{comp}} \equiv \frac{\sum_{j=1}^{J} u_{t,j} \bar{h}_j}{\sum_{j=1}^{J} u_{t,j}} \qquad \text{and} \qquad \bar{h}_t^{\text{real}} \equiv \frac{\sum_{j=1}^{J} \bar{u}_j h_{t,j}}{\sum_{j=1}^{J} \bar{u}_j},$$

where $\bar{h}_j \equiv \sum_{t=1}^{T} h_{t,j}/T$ is the time-averaged hiring rate for type j workers and $\bar{u}_j \equiv \sum_{t=1}^{T} u_{t,j}/T$ is the average number of unemployed type j workers. If the heterogeneity hypothesis is correct $\bar{h}_t^{\text{comp}}$ should be strongly procyclical and $\bar{h}_t^{\text{real}}$ should be acyclical.[6] Note that in order to generate large fluctuations in $\bar{h}_t^{\text{real}}$, the data must have large differences in hiring rates for groups with substantially different cyclical unemployment risk. If individual hiring rates are too similar, composition effects will not generate substantial fluctuations in the aggregate hiring rate. If the composition of the unemployed population is not sufficiently

light on the nature of unemployment dynamics." (p. 320) Based on this type of evidence and on the fact that there is simply not enough measurable variation in the composition of the unemployed population to generate large movements in unemployment duration, van den Berg and van der Klaauw (2001) and Abbring, van den Berg, and van Ours (2002) reach a similar conclusion in their detailed analyses of French data.

[6]Another possibility is that the heterogeneity hypothesis is correct but I do not observe the relevant characteristics. This cannot be tested using cross-sectional data.

cyclical, the weights in equation (3) will not change.

I use the public use monthly CPS micro data from January 1994 to March 2004 to construct measures of the number of short-term unemployed workers and total unemployed workers in different groups.[7] I consider six different characteristics: sex, race (white or non-white), seven age groups (16–19, 20–24, 25–34, ..., 55–64, and 65 and over), six marital status categories, six reasons for unemployment (job loser on layoff, other job loser, temporary job ended, job leaver, re-entrant, and new-entrant), and nine census regions.[8] Although the data set is large, it is impractical to consider all $2 \times 2 \times 7 \times 6 \times 6 \times 9 = 9072$ possible groups simultaneously. Instead, I analyze each characteristic in isolation.

According to the aggregate data, the hiring rate rose steadily from 34 percent per month at the start of 1994 to 47 percent by the end of 1999 before falling to 30 percent by the first quarter of 2004. Figure 3 shows my measure of $\bar{h}_t^{\text{comp}}$ and $\bar{h}_t^{\text{real}}$ for the six different characteristics. Only changes in the composition of the reason for unemployment appreciable affect the hiring rate, raising it by 2 percentage points from 1994 to 2000 and then reducing it by a similar amount during the next three years.[9] Figure 4 delves into the source of this compositional change more deeply. The top panel indicates that an increase in the number of 'other job losers' (as opposed to 'job losers on layoff') explains most of the increase in unemployment during this period. The bottom panel shows that this group dependably has the lowest hiring rate, which explains the measured compositional change in the hiring rate. Nevertheless, the cyclical pattern of the monthly hiring rate for this group is, if anything, sharper than the cyclical pattern of the hiring rate for the rest of the population, rising from 25 percent in 1994 to almost 40 percent in 1999 and then declining by a factor of 2 by the start of 2000. Only job losers on layoff show little change in the hiring rate during this time period, but that probably reflects that layoffs typically last for two months. In short, there is scant evidence that changes in the composition of the labor force explains fluctuations in the hiring rate, since the hiring rate is strongly procyclical for all groups of workers.

[Figure 3 about here.]

[7] The data are available from `http://www.nber.org/data/cps_basic.html`. It should be possible to extend these results back to January 1976.

[8] I also examine five different education groups (high school dropouts, high school diploma, some college, bachelor's degree, some postgraduate education) for workers 25 and older. There is virtually no change in the composition of the unemployed population conditional on their education.

[9] Changes in the age distribution also explain about a one percent decline in the hiring rate in the later period. This appears to be because older workers are more likely to be 'other job losers', a fact that is already picked up in the panel on 'Reason for Unemployment.'

[Figure 4 about here.]

2.4 Entry and Exit from the Labor Force

I have so far assumed that unemployment spells end only when a worker finds a job. As the share of re-entrants in Figure 4 suggests, many spells end when a worker exits the labor force. To see how this affects the measured hiring rate, let x_t denote the rate at which an unemployed worker exits the labor force. Then one can extend equation (1) to allow for this possibility:

$$u_{t+1} = u_t(1 - h_t - x_t) + u^s_{t+1}.$$

Now a worker is unemployed both in month t and $t+1$ only if she is neither hired nor exits the labor force. This gives

$$h_t + x_t = 1 - \frac{u_{t+1} - u^s_{t+1}}{u_t}$$

In other words, the 'hiring rate' in equation (2) actually measures the total exit rate from unemployment, including the exit rate from the labor force.

To understand whether this affects the computed cyclicality of the hiring rate, it is necessary to know how x_t behaves. I use microdata from the January 1976 to March 2004 CPS to estimate x_t. More precisely, I take advantage of the fact that the CPS is a rotating panel, so each individual is in the survey for four consecutive months. I use standard techniques to match individual records across consecutive months using household identifiers, individual line numbers, the survey rotation group, race, age, and sex.[10] I then categorize each individual according to her current and lagged employment status: employed, unemployed, or inactive (out of the labor force). Finally, I weight and sum the data to compute the fraction of unemployed workers who exit the labor force in the subsequent month. The results are depicted in Figure 5. Note that the redesign of the survey instrument in January 1994 breaks the exit rate series, making the earlier and later data incomparable.

[Figure 5 about here.]

The figure shows that the fraction of unemployed workers that exit the labor force declines when the unemployment rate increases, but the decline is relatively small, on the order of four percent of unemployed workers per month. Figure 1 shows that the hiring rate fell

[10] The best known examples of this are Abowd and Zellner (1985) and Poterba and Summers (1986). They estimate measurement error in transition rates; unfortunately, I am unable to replicate their analysis with public use data.

by about four times as much during each of those episodes, so this calculation does not qualitatively affect the measured cyclicality of the hiring rate. Moreover, this result may be well reflect the same phenomenon as Figure 4, which shows that the cyclicality in the hiring rate is partially a consequence of the changing composition of the unemployment pool away from re-entrants and towards other job losers during cyclical downturns. Presumably re-entrants have a weaker labor force attachment than job losers, and so a decline in their share of the unemployed population will naturally reduce the transition rate from unemployment to inactivity. Accounting both for changes in the composition of the unemployed population and for changes in labor force attachment is likely to be duplicative.

3 The Separation Rate

The flip side of the hiring rate is the separation rate, the probability that an employed worker loses or leaves her job in a given month. I again start with the fiction that workers never enter or exit the labor force. This might suggest that the transition rate from employment to unemployment, or equivalently the short-term unemployment rate u_t^s, is a simple measure of the separation rate. But this measure ignores a potentially important issue arising from time aggregation: even if workers only search while unemployed, a worker may lose her job and find another one without ever being measured as unemployed by a CPS interviewer. To address this, I use information both on the transition rate from employment to unemployment and on the hiring rate to measure the separation rate. This time series for separations contributes little to the overall fluctuations in the unemployment rate, particularly during the lats two decades. I then incorporate movements in and out of the labor force into this analysis before discussing alternative data sources that have convinced many economists that fluctuations in the separation rate are an important source of volatility in the unemployment rate.

3.1 A Simple Measure

If a worker becomes unemployed in month $t+1$ whenever she loses her job in month t, then the separation rate would be the ratio of newly unemployed workers in month $t+1$ to employed workers in month t. I measure newly unemployed workers using short-term unemployment u_{t+1}^s and use the BLS's measure of employment based on the CPS for employment e_t, giving a separation rate $\tilde{s}_t = u_{t+1}^s/e_t$.[11] Figure 6 shows that this time series is countercyclical,

[11] Hall (2004a) discusses this measure extensively.

rising in anticipation of each peak in the unemployment rate. In fact, equation (1) indicates that by definition this measure explains all of the fluctuations in unemployment not directly attributable to changes in the hiring rate h_t.

[Figure 6 about here.]

If the CPS observed each unemployment incident, $\tilde{s}_t$ would provide a satisfactory measure of the separation rate. But it is confounded by a time-aggregation problem: there is a non-trivial probability that a worker who loses her job after the survey date in one month will find another job before the next survey date. Moreover, that probability is increasing in the hiring rate. This means that a decline in the hiring rate will naturally raise short-term unemployment and hence the simple separation rate even if the 'true' separation rate s_t is constant. In light of the evidence presented in Section 2 that the hiring rate is procyclical, time aggregation is likely to induce countercyclical fluctuations in the simple separation rate.

3.2 Time Aggregation

I instead estimate the probability that a separation results in a measured employment to unemployment transition. If a worker separates from a job with a fraction x months remaining before the survey date, the probability that she fails to find another job is $(1-h_t)^x$, where h_t is her full-month hiring rate. If furthermore workers are equally likely to separate at any time during the month, this implies that a fraction

$$\int_0^1 (1-h_t)^x \, dx = -\frac{h_t}{\log(1-h_t)} \tag{7}$$

of workers who experience a separation in month t are measured as unemployed in month $t+1$. This is a decreasing function of the hiring rate h_t. When h_t is small, this is nearly equal to $1-\frac{h_t}{2}$, although the approximation is increasingly poor for values of the hiring close to 1. Equivalently, we can deduce the separation rate from

$$u^s_{t+1} = -\frac{e_t s_t h_t}{\log(1-h_t)}, \tag{8}$$

A worker is employed in month t and unemployed in month $t+1$ if she separates from her job (fraction s_t of the workers) and fails to find another job within the month (fraction $-h_t/\log(1-h_t)$ of the workers). This can be inverted to solve for the separation rate s_t.

Now suppose workers are heterogeneous. The same logic implies

$$u_{t+1}^s = -\sum_{i=1}^{e_t} \frac{s_t^i h_t^i}{\log(1-h_t^i)},$$

where employed individuals $i \in \{1, \ldots, e_t\}$ are identified using superscripts. Unfortunately, unless the joint distribution of s_t^i and h_t^i are known, it is impossible to invert this equation to solve for the mean separation rate, $s_t = \sum_{i=1}^{e_t} s_t^i / e_t$.[12] I therefore am forced to assume that all workers who experience a separation have the same full-month hiring rate. Rather than using the mean hiring rate h_t, however, it seems more reasonable to use Hall's (2004a) measure of the rate for short-term unemployed workers, $\hat{h}_t$ defined in equation (6).[13] Then the previous logic still implies

$$s_t = \frac{\sum_{i=1}^{e_t} s_t^i}{e_t} = -\frac{u_{t+1}^s \log(1-\hat{h}_t)}{e_t \hat{h}_t}, \tag{9}$$

a trivial generalization of the model with homogeneous workers.

I construct s_t according to equation (9) using measures of the number of employed and short-term unemployed workers and Hall's measure of the hiring rate, which is itself a function of the short-term and medium-term unemployment rate. Figure 7 shows the results. Relative to the simple measure, the true separation rate reflects workers who lose a job but quickly find another one, and so is naturally somewhat higher. It is also somewhat less volatile, with a standard deviation about trend of 8.4 percent rather than 11.4 percent. The difference is most noticeable during the last two decades when the separation rate fell almost constantly with scarcely any interruption for cyclical downturns.[14] Virtually all of the increase in the unemployment rate during the recessions in 1991 and 2001 was

[12] If s_t^i and h_t^i are independent, then

$$s_t = \frac{u_{t+1}^s}{-\sum_{i=1}^{e_t} h_t^i / \log(1-h_t^i)}.$$

$-h/\log(1-h)$ is a concave function, so Jensen's inequality implies $-h_t/\log(1-h_t) > -\sum_{i=1}^{e_t} h_t^i/\log(1-h_t^i)$. Then the above formula probably understates the separation rate. This affects the interpretation of the results only if the cross-sectional variance in the h_t^i, and so the importance of Jensen's inequality, varies cyclically.

[13] The results are insensitive to this choice. For example, one can use h_t or even the hiring rate for 'other job losers' shown in Figure 4. All of these series fluctuate substantially at business cycle frequencies.

[14] The drop between 1993 and 1994 is notable. Abraham and Shimer (2001) explain that a redesign of the CPS instrument in January 1994 reduced the measured short-term unemployment rate. From equation (9), this reduced the measured separation rate, which appears to explain the sharp drop in the measured separation rate at that time. This implies that the data before and after the redesign are incomparable.

therefore a consequence of the reducing hiring rate. This manifested itself directly in longer unemployment duration and indirectly in higher unemployment incidence, with the latter occurring because a reduction in the hiring rate raised the probability that a worker becomes unemployed following a separation from her old employer.

[Figure 7 about here.]

A transformation of variables makes this point even clearer. Suppose the hiring rates h_t and $\hat{h}_t$ were constant between 1948 and 2004 at their average value $\bar{h}$ and $\bar{\hat{h}}$, but the separation rate followed the actual pattern depicted in Figure 7. Then combining equations (1) and (8), the unemployment rate $\tilde{u}_t$ would have evolved according to

$$\tilde{u}_{t+1} = \tilde{u}_t(1 - \bar{h}) - \frac{(1 - \tilde{u}_t)s_t\bar{\hat{h}}}{\log(1 - \bar{\hat{h}})}. \tag{10}$$

I fix the initial unemployment rate at its January 1948 level and then let it evolve according to this equation in the subsequent years to see how much of the fluctuations in unemployment are driven by changes in the separation rate. Figure 8 shows the results. In the immediate post-war period, fluctuations in the separation rate explain most of the fluctuations in the unemployment rate, but the pattern has changed. Even during the 1982 recession, the separation rate raised the unemployment rate by perhaps one percentage point, and in later recessions it had virtually no impact.

[Figure 8 about here.]

This calculation may in fact understate the indirect effect of a reduction in the hiring rate on the measured separation rate. If workers receive advanced notification of layoffs,[15] then on average they have more than half a month to find a new job before becoming unemployed, which increases the dependence of the simple separation rate u^s_{t+1}/e_t on the hiring rate. For example, if half the workers get one month's advanced notice and start searching for a new job immediately, the probability that a separation results in unemployment rises to $\frac{-(2-h_t)h_t}{2\log(1-h_t)}$, which is about twice as responsive to changes in the hiring rate as the earlier

[15] Ruhm (1992) examines the frequency of advanced notification using the 1988 Displaced Workers Survey. He shows that 53 percent of workers who experience a permanent job loss due to a plant closure or layoff received advanced notification. Moreover, workers who receive advanced notification have shorter average unemployment spells.

expression $-h_t/\log(1-h_t)$. This reinforces my conclusion that increases in the separation rate have not contributed to recent cyclical increases in the unemployment rate.

The reader may be concerned that if decreases in the hiring rate raise the simply-measured separation rate, then increases in the separation rate should likewise reduce the measured hiring rate. In other words, isn't part of the measured cyclicality of h_t a consequence of cyclicality in s_t because of time-aggregation? There are two reasons why this is probably not relevant. First, the measured hiring rate uses only information on the number of workers who are unemployed in month t and $t+1$ and the number of short-duration unemployed workers in month $t+1$ (equation 1). If a worker who has a brief job between the survey dates in months t and $t+1$ correctly reports that her unemployment duration is less than four weeks, she will be counted among the short-duration unemployed in month $t+1$ and I will correctly infer that she was hired and then experienced a separation.[16] Second, the bias introduced by short jobs is likely to be small. The logic behind equation (7) implies that a fraction $-s_t/\log(1-s_t)$ of the workers who find a job lose it before the next survey. If the separation rate increases from 4 to 5 percent per month, a magnitude comparable to the worst recessions in the 1970s and 1980s, the fraction of workers who experience such a brief employment spell rises from 2.0 to 2.5 percent. Even if the CPS records none of these workers' employment spells, the increase in separations causes a negligible 0.5 percent decrease in the measured hiring rate.

3.3 Entry and Exit from the Labor Force

Entry and exit from the labor force may confound my measure of the separation rate. A reduction in the transition rate from inactivity to unemployment reduces the short-duration unemployment rate u_t^s, which reduces the measured separation rate (equation 9). Conversely, I would fail to measure an increase in the separation rate that induces workers to transit from employment to inactivity since equation (9) only uses information on unemployment. To quantify the importance of these effects, I again use matched CPS files to examine the cyclical behavior of these transition rates.

Figure 9 indicates little cyclicality in the rate at which employed workers exit the labor force, so the latter bias is unlikely to be important. On the other hand, Figure 10 shows

[16] Before the redesign of the CPS instrument in 1994, workers who were unemployed in months t and $t+1$ were asked their unemployment duration in both months, and so the worker could claim to have had a short intervening job. The redesigned instrument uses dependent interviewing to infer the unemployment duration of a worker who is unemployed both in months t and $t+1$, implicitly assuming that there was no intervening employment spell. Thus after 1994, the argument in the text is inapplicable.

that the transition rate from inactivity to unemployment is countercyclical, so some of the measured short-term unemployed are entering or reentering the labor force and have not recently experienced a separation. Presumably this is related to fluctuations in the hiring rate, which affect the ease with which inactive workers move directly into employment. But in any case, it implies that using the short-term unemployment rate to measure the separation rate probably leads me to overstate the cyclicality of the variable of interest. There is no evidence that separations are countercyclical, particularly during the last two decades.

[Figure 9 about here.]

[Figure 10 about here.]

3.4 Reconciliation with Alternative Evidence

Given the evidence presented in this section, why do so many economists believe that the separation rate is strongly countercyclical?[17] Part of the reason is a failure to distinguish between an increase in unemployment inflows due to an increase in separations and that due to a decrease in hiring. For example, at least since Abowd and Zellner (1985) and Poterba and Summers (1986), many authors have used matched CPS data to examine the transition rate between employment, unemployment, and inactivity. But to my knowledge none has accounted for the fact that a decrease in the hiring rate indirectly raises the measured transition rate from employment to unemployment.

Another common piece of evidence comes from unemployment insurance claims. These are useful in part because they are available at a high frequency (weekly) and in a timely manner. Figure 11 shows that initial claims for unemployment insurance rise sharply during recessions.[18] This is partially attributable to a reduction in the fraction of workers who experience a separation but manage to get a new job without experiencing an unemployment spell. It is probably exacerbated by an increase in the unemployment insurance take-up rate

[17] A standard reference is Hall (1995), who writes "...brief, sharp episodes of primary job loss are followed by long periods of slowly rebuilding employment relationships over the business cycle. Although the case is far from complete, I believe that these events in the labor market play an important part in the persistence of high unemployment and low output long after the initial shock that triggers a recession.". He has since been convinced that the evidence is indeed far from complete, writing more recently, "...in the modern U.S. economy, recessions are not times of unusual job loss. New data on separations show them to be remarkably constant from peak to trough. Bursts of job loss had some role in earlier recessions, but are still mostly a side issue for the reason just mentioned—a burst is quickly reabsorbed because of high job-finding rates." (Hall 2004b)

[18] The data are available from `http://workforcesecurity.doleta.gov/unemploy/claims.asp`.

during downturns, as would occur if there are fixed costs of starting to collect unemployment insurance and workers rationally anticipate longer unemployment spells during downturns. Figure 11 also shows that continuing claims increase during downturns. Again, this is consistent with constant separations since a reduction in hiring raises the fraction of initial claims that continue several months later. Similarly, mass layoff and extended mass layoff statistics are based on unemployment insurance claims and therefore probably have the same bias.

[Figure 11 about here.]

Perhaps the best known evidence on the cyclicality of separations comes from Davis, Haltiwanger, and Schuh (1996), who compute measures of job creation and destruction in manufacturing. Job creation is defined as the increase in employment at expanding business establishments and job destruction is the decrease in employment at contracting business establishments. An important conclusion that comes out of this research is that "Job destruction rises dramatically during recessions, whereas job creation initially declines by a relatively modest amount." (Davis, Haltiwanger, and Schuh 1996, p. 31) There are a few reasons to be cautious with the interpretation of this conclusion, however. First, Davis, Haltiwanger, and Schuh (1996) focus exclusively on manufacturing establishments, although more recent work has extended their methodology to cover the entire labor market since 1990. Faberman (2004) uses that data to show that job destruction was more volatile than job creation in the 1991 recession, but the pattern reversed in 2001. Second, firms can destroy jobs either by firing workers or by not hiring to replace workers who leave. One way to distinguish these alternatives is to look at establishments that shutdown, which is clearly evidence of firms firing workers. Davis, Haltiwanger, and Schuh (1996) conclude that "shutdowns do not account for an unusually large fraction of job destruction during recessions." (p. 34) This means that spikes in job destruction are consistent with the view advanced in this paper that there were only small increases in the separation rate of employed workers during those downturns.

Finally, the Job Openings and Labor Turnover Survey (JOLTS) provides some new evidence that the total separation rate actually fell during the last recession. Since December 2000, this survey has asked business establishments how many workers they added to their payrolls during the previous month, how many workers left their payrolls, and whether those workers were laid off or quit. Figure 12 suggests that both new hires and separations fell as the United States labor market remained weak in 2002 and 2003. More tellingly, Figure 13 shows only a brief small spike in layoffs just after the terrorist attack in September 2001, while quits fell steadily during this period. It is notable that both new hires and quits fell by approximately 0.6 percentage points between 2001 and 2003. A plausible explanation for this pattern is a decline in job-to-job movements, which would naturally cause fewer quits and fewer new hires. Of course, using unemployment duration data I cannot hope to measure job-to-job quits, but the next section explores job-to-job movements in depth.

[Figure 12 about here.]

[Figure 13 about here.]

4 Job-to-Job Transitions

This section extends the basic model of transitions between employment and unemployment to allow employed workers to search for better jobs. As before, unemployed workers find a job with probability h_t in month t. Employed workers lose their job with probability s_t but find another one before the next survey with probability $-\hat{h}_t / \log(1-\hat{h}_t)$. I also introduce a reason for voluntary job-to-job transitions: I assume that jobs are of different 'quality' $\tilde{z}$, an index summarizing all of the job's pecuniary and non-pecuniary aspects, and allow employed workers to search for better jobs. Suppose an employed worker finds a job at rate h_t^e and accepts it if the new job's quality $\tilde{z}'$ exceeds the old job's quality $\tilde{z}$. The critical assumption is that when a worker finds a job, the quality is drawn from a time-invariant continuous distribution $F(\tilde{z})$ with support $[0, \bar{z}]$. This gives a canonical model of job-to-job transitions where workers switch jobs whenever they have an opportunity to improve their job quality.

To compute the job-to-job transition rate, it is necessary to keep track of the distribution of employed workers across job qualities, say $\tilde{G}_t(\tilde{z})$ in month t. This evolves according to a simple difference equation:

$$\tilde{G}_{t+1}(\tilde{z})e_{t+1} = \tilde{G}_t(\tilde{z})e_t(1-s_t)\big(1-h_t^e(1-\tilde{F}(\tilde{z}))\big) + \left(u_t h_t + e_t s_t \left(1 + \frac{\hat{h}_t}{\log(1-\hat{h}_t)}\right)\right)\tilde{F}(\tilde{z}). \tag{11}$$

The left hand side is the number of employed workers with job quality less than $\tilde{z}$ in month $t+1$. This is equal to the number of workers in that situation in month t, $\tilde{G}_t(\tilde{z})e_t$, who do not experience a separation (probability $1-s_t$) and do not find a job with quality in excess of $\tilde{z}$ (probability $1-h_t^e(1-\tilde{F}(\tilde{z}))$), plus the number of unemployed workers who find a job with quality less than $\tilde{z}$, $u_t h_t \tilde{F}(\tilde{z})$, plus the number of employed workers who separate but get a new job with quality below $\tilde{z}$ within the month, $e_t s_t\big(1 + \hat{h}_t / \log(1-\hat{h}_t)\big)\tilde{F}(\tilde{z})$.[19]

Note that $\tilde{G}_t(\bar{z}) = 1$ for all t, since all employed workers produce less than $\bar{z}$. Then evaluating equation (11) at $\tilde{z} = \bar{z}$ gives

$$e_{t+1} = e_t + u_t h_t + \frac{e_t s_t \hat{h}_t}{\log(1-\hat{h}_t)},$$

which states that the number of employed workers next month is equal to the number of employed workers this month plus the unemployed workers who are hired minus the employed workers who lose their job and fail to find a new one. I use this to eliminate e_{t+1} from equation (11):

$$\tilde{G}_{t+1}(\tilde{z}) = \frac{\tilde{G}_t(\tilde{z}) e_t (1-s_t)\big(1 - h_t^e(1-\tilde{F}(\tilde{z}))\big) + \left(u_t h_t + e_t s_t \left(1 + \frac{\hat{h}_t}{\log(1-\hat{h}_t)}\right)\right)\tilde{F}(\tilde{z})}{e_t + u_t h_t + \frac{e_t s_t \hat{h}_t}{\log(1-\hat{h}_t)}}.$$

Given an initial guess of the distribution $\tilde{G}$, a time-invariant distribution function $\tilde{F}$, and time series of employment, unemployment, and separation and hiring rates, one can compute all future distributions $\tilde{G}$. Then given this, the fraction of employed workers who switch employers between months t and $t+1$ is

$$j_t = h_t^e \int_0^{\bar{z}} (1-\tilde{F}(\tilde{z}))\tilde{G}_t'(\tilde{z}) d\tilde{z} + s_t \left(1 + \frac{\hat{h}_t}{\log(1-\hat{h}_t)}\right).$$

This is the sum of the fraction of 'voluntary' job switchers who quit their job to take a better one and the fraction of 'involuntary' switchers who succeed in finding another job despite suffering an involuntary separation.

The main difficulty with measuring j_t is that the quality distribution $\tilde{F}$ is unobservable. There is an easy way around this: rather than indexing a job opportunity by its quality $\tilde{z}$ drawn from the latent distribution $\tilde{F}$, I can represent it by its percentile in the quality distribution, z, which by definition is distributed uniformly on $[0,1]$. Then the distribution

[19] Implicit in this expression is an assumption that a worker cannot find two jobs within a month. To allow for that possibility, it is easiest to express the evolution of the distribution $\tilde{G}$ as a differential equation in continuous time, with constant hiring and separation rates between survey dates. Although the analysis is somewhat more cumbersome, the results are quantitatively unchanged.

of workers' normalized quality satisfies

$$G_{t+1}(z) = \frac{G_t(z)e_t(1-s_t)\big(1-h_t^e(1-z)\big) + \left(u_t h_t + e_t s_t \left(1+\frac{\hat{h}_t}{\log(1-\hat{h}_t)}\right)\right) z}{e_t + u_t h_t + \frac{e_t s_t \hat{h}_t}{\log(1-\hat{h}_t)}}, \tag{12}$$

with $G_t(F(\tilde{z})) = \tilde{G}_t(\tilde{z})$ for all $\tilde{z}$. The job-to-job transition rate is

$$\begin{aligned} j_t &= h_t^e \int_0^1 (1-z)G_t'(z)\,dz + s_t\left(1+\frac{\hat{h}_t}{\log(1-\hat{h}_t)}\right) \\ &= h_t^e \int_0^1 G_t(z)dz + s_t\left(1+\frac{\hat{h}_t}{\log(1-\hat{h}_t)}\right), \end{aligned} \tag{13}$$

where the second equality uses integration-by-parts. I assume that the rate at which an employed worker finds a new job is proportional to the hiring rate, $h_t^e = \alpha h_t$, where the constant α can be varied to get match the average empirical frequency of job-to-job movements. Then given an initial choice of the distribution G and data on current and past employment, unemployment, hiring rates, and separation rates, it is straightforward to compute the time series for job-to-job movements.[20] From equation (13), the frequency of voluntary job-to-job transitions is high when the hiring rate h_t is high or when employed workers are in bad jobs, so $\int_0^1 G_t(z)dz$ is low. The latter occurs if hiring rates have been low or separation rates have been high in the recent past.

In steady state, one can solve equation (12) for $G_{t+1}(z)$ and then integrate to obtain the job-to-job transition rate. The resulting expression depends only on the job finding rate of employed workers h_t^e and newly laid-off workers $\hat{h}$ and on the separation rate s:

$$j = \frac{s}{1-s}\left(\left(1+\frac{s}{h^e(1-s)}\right)\log\left(\frac{h^e(1-s)}{s}+1\right)-1\right) + s\left(1+\frac{\hat{h}}{\log(1-\hat{h})}\right).$$

Not surprisingly, the steady state job-to-job transition rate is increasing in the hiring rates h^e and $\hat{h}$; if it is easier to get a job, more workers will change jobs either voluntarily or involuntarily following a separation. The transition rate is also increasing in the separation rate. An increase in the separation rate reduces the duration of employment spells, leaving most workers at lower rungs in the job ladder. Such workers are willing to accept more outside

[20] It is easy to check the sensitivity of results to the initial choice of G. In practice, the effects disappear after a few years.

job opportunities, and so the job-to-job transition rate is higher. Although this calculation is incorrect out of steady state, it suggests that if the separation rate were strongly countercyclical and the hiring rate acyclical, this simple model would predict a countercyclical job-to-job transition rate.

Out of steady state, I use BLS time series for employment and unemployment, the hiring rate, the hiring rate constructed as in equation (2) and (6), the separation rate constructed using equation (9), and a guess at the relative efficiency of search on- and off-the-job α to compute the distribution $G(z)$ from equation (12) in each month from 1948 to 2004. Using that, equation (13) immediately yields the job-to-job transition rate. Regardless of the choice of α, in an average month approximately 1.1 percent of employed workers switch jobs involuntarily, because they experience a separation but manage to locate a new employer within the month. But the frequency of voluntary job-to-job transitions depends on α. Figure 14 plots the total job-to-job transition rate for $\alpha = 0.2$, with the unemployment rate graphed for comparison. On average, the job-to-job transition rate is 3.8 percent per month, with about 71 percent accounted for by voluntary transitions. Those levels are sensitive to the choice of α. For example, lowering it to $\alpha = 0.1$ reduces the voluntary job-to-job transition rate to 1.6 percent per month, 60 percent of total job-to-job transitions, but has little effect on the cyclicality of the job-to-job transition rate.[21] The figure also shows that total job-to-job transitions are fairly strongly negatively correlated with the unemployment rate, which means that the lower hiring rate (which reduces both voluntary and involuntary job-to-job transitions) outweighs the higher separation rate (which raises involuntary transitions) during a typical downturn.

[Figure 14 about here.]

How does this compare to actual United States data on job-to-job transition rates? The JOLTS data in Figures 12 and 13 show a 0.6 percentage point decline in both new hires and quits between 2001 and 2003, an effect that I previously suggested might reflect a decline in job-to-job movements. Figure 14 in fact shows a slightly larger decline in the predicted job-to-job transition rate over that time period. A more direct measure of the job-to-job transition rate comes from the microeconomic data underlying the CPS. In 1994, the CPS instrument began using dependent interviewing, asking respondents who had been

[21] If $\alpha = 0.2$, the model predicts that between the first quarters of 2000 and 2003, the job-to-job transition rate should have fallen from 3.5 percent to 2.5 percent. With $\alpha = 0.1$, the model predicts a similar relative decline, from 2.5 to 1.7 percent, over the same time period.

surveyed in the previous month, "Last month, it was reported that you worked for x. Do you still work for x (at your main job)?" Their answers are recorded in the public use files. Following Fallick and Fleischman (2004), I use the fraction of employed workers who answer this question affirmatively, weighted by the CPS final weights, to compute the empirical job-to-job transition rate.[22] Figure 15 shows the empirical behavior of job-to-job transitions and the prediction of the theoretical model. The fact that the levels are approximately correct is due to a judicious choice of the relative intensity of on-the-job search α. But the underlying data on hiring rates and separation rates drive the fluctuations in the theoretical series. Although the theory over predicts job-to-job transitions during the boom years from 1998 to 2000, it is otherwise very successful at explaining the timing and magnitude of the decline in job-to-job transitions from 2001 to 2002 and levelling off in 2003. I await the release of additional public use CPS files to see whether the data support the theoretical prediction that job-to-job transitions should have increased during 2004. But the JOLTS data in Figures 12 and 13 provide some evidence that this was the case. According to this employer survey, new hires and quits both increased during the second half of 2003 and the first half of 2004.

[Figure 15 about here.]

In summary, if downturns were periods with high separation rates and normal job finding rates, the canonical model would predict an increase in job-to-job transitions in 2001–2003. This would have occurred both because of an increase in the number of workers who suffered a separation but managed to find a new job and because the increase in separations would have reduced the age of matches and hence their quality, causing more voluntary job-to-job transitions. The fact that job-to-job transitions fell is qualitatively and quantitatively consistent with the evidence that employed workers, like unemployed workers, found it harder to obtain a job during the recession and did not experience a large increase in the separation rate.

5 Conclusions

This paper argues that business cycle fluctuations in unemployment are primarily a consequence of changes in the probability that an unemployed worker finds a job within a month,

[22] There are five possible answers: same employer as last month, new employer, refused, don't know, and blank. I drop the last three categories, accounting for 10.3 percent of employed workers.

the hiring rate. Changes in the separation rate do not explain any of the observed unemployment fluctuations during the last two decades. A canonical model of on-the-job search suggests that in such an environment, job-to-job transitions should be procyclical, consistent with recent evidence from JOLTS and the CPS.

The question remains, why is the hiring rate so low during recessions? A facile answer is that firms create fewer jobs, as measured by vacancy rates either from the Conference Board Help Wanted Advertising Index or the Job Openings and Labor Turnover Survey, relative to the number of unemployed workers. For example, from 1951 to 2004, the correlation between the detrended ratio of help wanted advertising to unemployment and the detrended job finding rate was 0.87.[23] Although there is as yet no generally accepted explanation for why there are periods of time when firms create so few vacancies, this paper provides some guidance by ruling out some possibilities that a priori seemed reasonable. For example, recessions are not periods when the unemployed population is particularly unattractive to firms, as suggested by this paper's findings that compositional changes in the unemployed population are not very large and that employed workers also have trouble getting new jobs during recessions. Likewise, recessions are probably not periods when some unemployed workers get discouraged and reduce their search effort: both the transition rate from unemployment to inactivity and the transition rate from inactivity to unemployment are countercyclical; and the hiring rate declines uniformly for all workers, regardless of demographic characteristics. Explanations based on the idea that 'wages are too high' (Hall 2004a, Kennan 2004, Shimer and Wright 2004) seem more promising.

References

Abbring, Jaap, Gerard van den Berg, and Jan van Ours (2002): "The Anatomy of Unemployment Dynamics," *European Economic Review*, 46(10), 1785–1824.

Abowd, John, and Arnold Zellner (1985): "Estimating Gross Labor-Force Flows," *Journal of Business and Economic Statistics*, 3(3), 254–283.

[23]Help-wanted advertising data can be downloaded from the BLS web site `http://www.bls.gov/` or from the St. Louis Federal Reserve Economic Database (FRED®II) `http://research.stlouisfed.org/fred2/`. I detrend the quarterly data using an HP filter with smoothing parameter 10^5.

Abraham, Katharine, and Robert Shimer (2001): "Changes in Unemployment Duration and Labor-Force Attachment," in *The Roaring Nineties*, ed. by Alan Krueger, and Robert Solow, pp. 367–420. Russel Sage Foundation and Century Foundation.

Baker, Michael (1992): "Unemployment Duration: Compositional Effects and Cyclical Variability," *American Economic Review*, 82(1), 313–321.

Bentolila, Samuel, and Giuseppe Bertola (1990): "Firing Costs and Labour Demand: How Bad is Eurosclerosis," *Review of Economic Studies*, 57(3), 381–402.

Darby, Michael, John Haltiwanger, and Mark Plant (1985): "Unemployment Rate Dynamics and Persistent Unemployment under Rational Expectations," *American Economic Review*, 75(4), 614–637.

——— (1986): "The Ins and Outs of Unemployment: The Ins Win," NBER Working Paper 1997.

Davis, Steven, John Haltiwanger, and Scott Schuh (1996): *Job Creation and Destruction*. MIT Press, Cambridge, MA.

Dynarski, Mark, and Steven Sheffrin (1990): "The Behavior of Unemployment Durations over the Cycle," *Reveiw of Economics and Statistics*, 72(2), 350–356.

Faberman, Jason (2004): "Gross Job Flows over the Past Two Business Cycles: Not all "Recoveries" are Created Equal"," BLS Working Paper 372.

Fallick, Bruce, and Charles Fleischman (2004): "Employer-to-Employer Flows in the U.S. Labor Market: The Complete Picture of Gross Worker Flows," Federal Reserve Board, Finance and Economics Discussion Series Working Paper 2004-34.

Hall, Robert (1995): "Lost Jobs," *Brookings Papers on Economic Activity*, 1, 221–256.

——— (2004a): "Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market," Stanford Mimeo, April 5, 2004.

——— (2004b): "The Labor Marekt is Key to Understanding the Business Cycle," Stanford Mimeo, September 23, 2004.

Kennan, John (2004): "Private Information, Wage Bargaining and Employment Fluctuations," University of Wisconsin-Madison Mimeo, April 2004.

Lilien, David (1982): "Sectoral Shifts and Cyclical Unemployment," *Journal of Political Economy*, 90(4), 777–793.

Poterba, James, and Lawrence Summers (1986): "Reporting Errors and Labor Market Dynamics," *Econometrica*, 54(6), 1319–1338.

Ruhm, Christopher (1992): "Advance Notice and Postdisplacement Joblessness," *Journal of Labor Economics*, 10(1), 1–32.

Shimer, Robert (2004): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, forthcoming.

Shimer, Robert, and Randall Wright (2004): "Competitive Search Equilibrium with Asymmetric Information," University of Chicago Mimeo, August 26, 2004.

Sider, Hal (1982): "Unemployment Duration and Incidence: 1968–1982," *American Economic Review*, 75(3), 461–472.

Siebert, Horst (1997): "Labor Market Rigidities: At the Root of Unemployment in Europe," *Journal of Economic Perspectives*, 11(3), 37–54.

van den Berg, Gerard, and Bas van der Klaauw (2001): "Combining Micro and Macro Unemployment Duration Data," *Journal of Econometrics*, 102(2), 271–309.

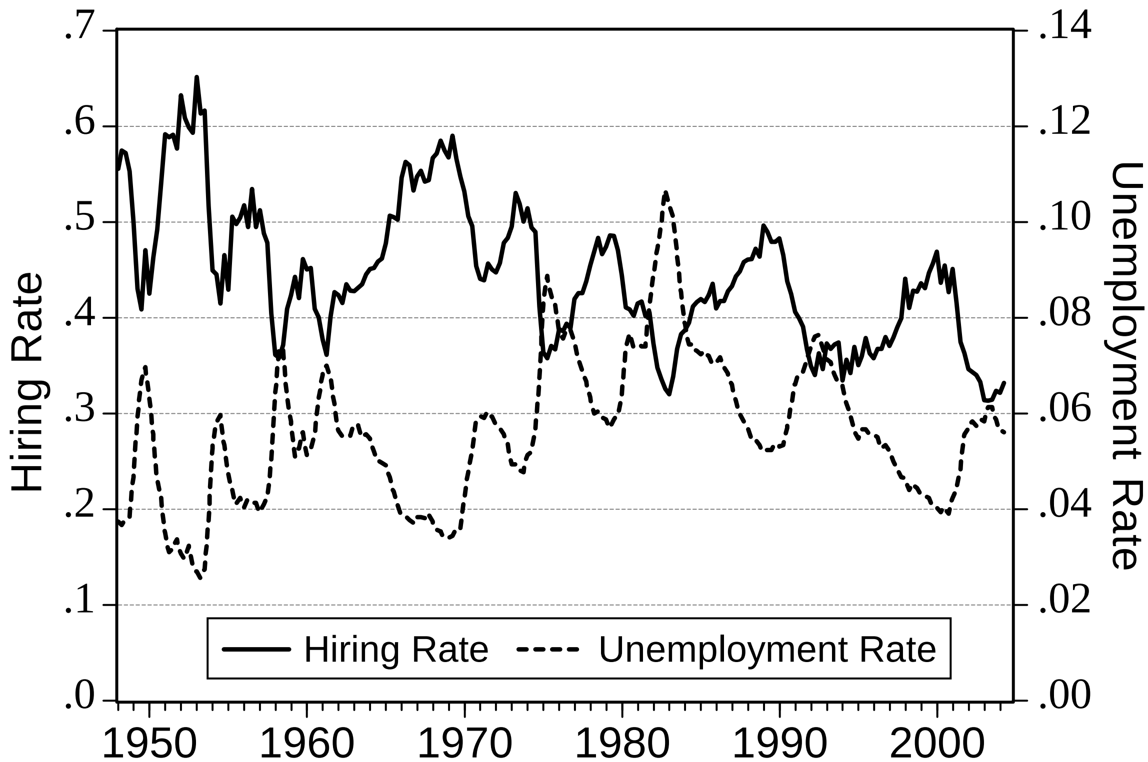

Figure 1: Hiring Rate and Unemployment Rate, United States, 1948Q1–2004Q2, quarterly average of monthly data. The hiring rate is constructed from unemployment and short-term unemployment according to equation (2). Unemployment and unemployment duration are computed by the BLS and seasonally adjusted.

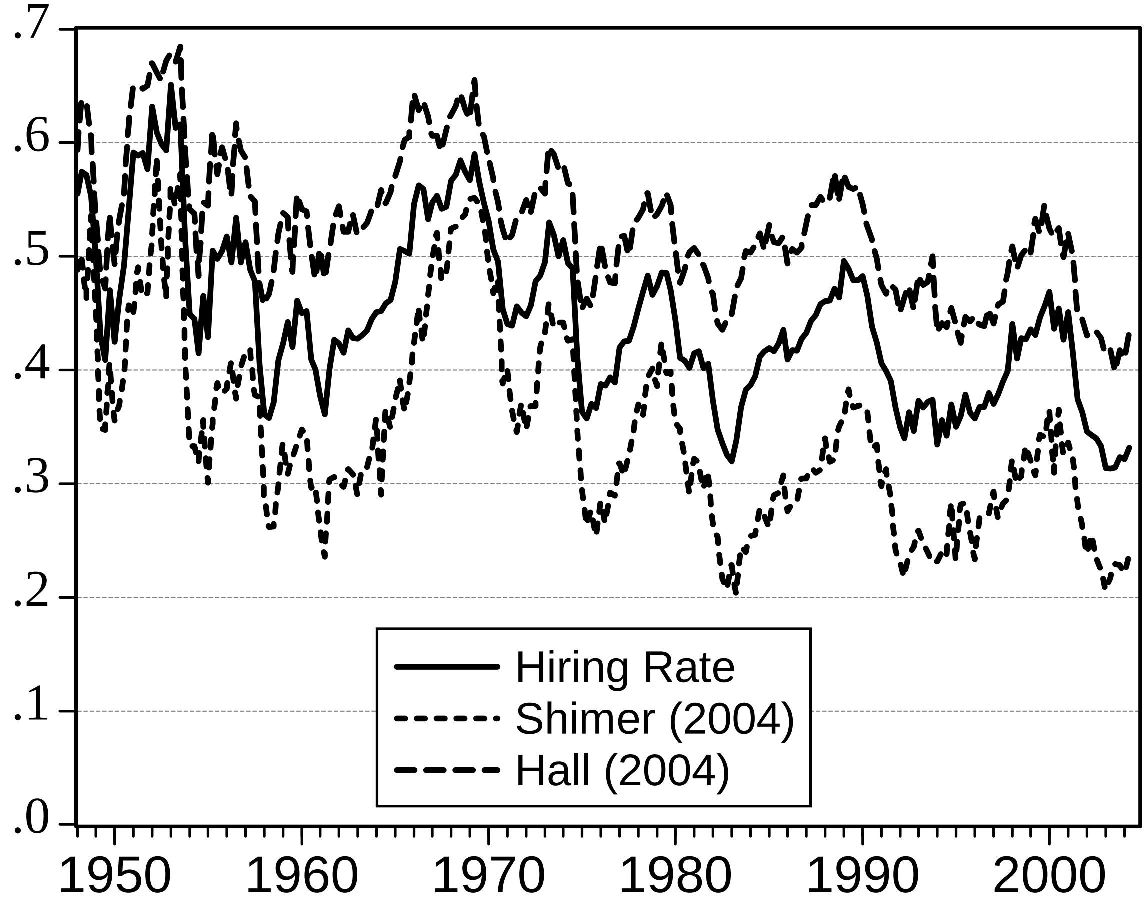

Figure 2: Three Measures of the Hiring Rate, United States, 1948Q1–2004Q2, quarterly average of monthly data. The basic hiring rate is constructed from unemployment and short-term unemployment according to equation (2). Shimer's (2004) measure is constructed from unemployment and unemployment duration according to equation (5). Hall's (2004) measure is constructed from short-term and medium-term unemployment according to equation (6). All data are computed by the BLS and seasonally adjusted.

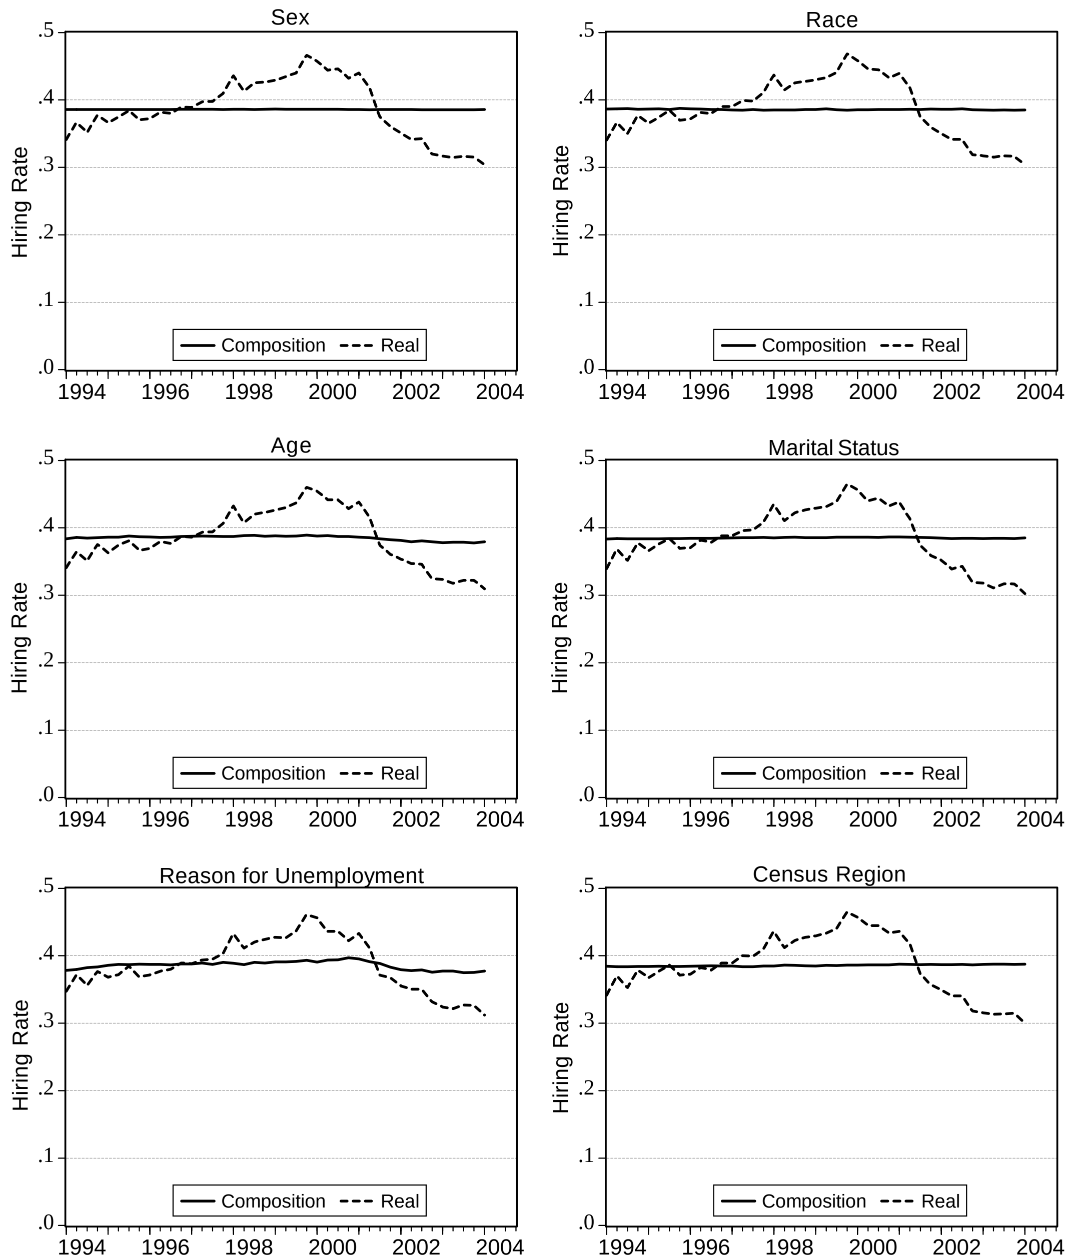

Figure 3: Six measures of the 'compositional' and 'real' component of changes in the hiring rate, $\bar{h}_t^{\text{comp}}$ and $\bar{h}_t^{\text{real}}$, respectively, United States, 1994Q1–2004Q2, quarterly average of monthly data. Each figure uses different characteristics: sex, race (white or nonwhite), seven age groups, six marital status categories, six reasons for unemployment (job loser on layoff, other job loser, temporary job ended, job leaver, re-entrant, and new-entrant), and nine census regions. The underlying data are constructed from the monthly CPS, seasonally adjusted, and averaged within quarters.

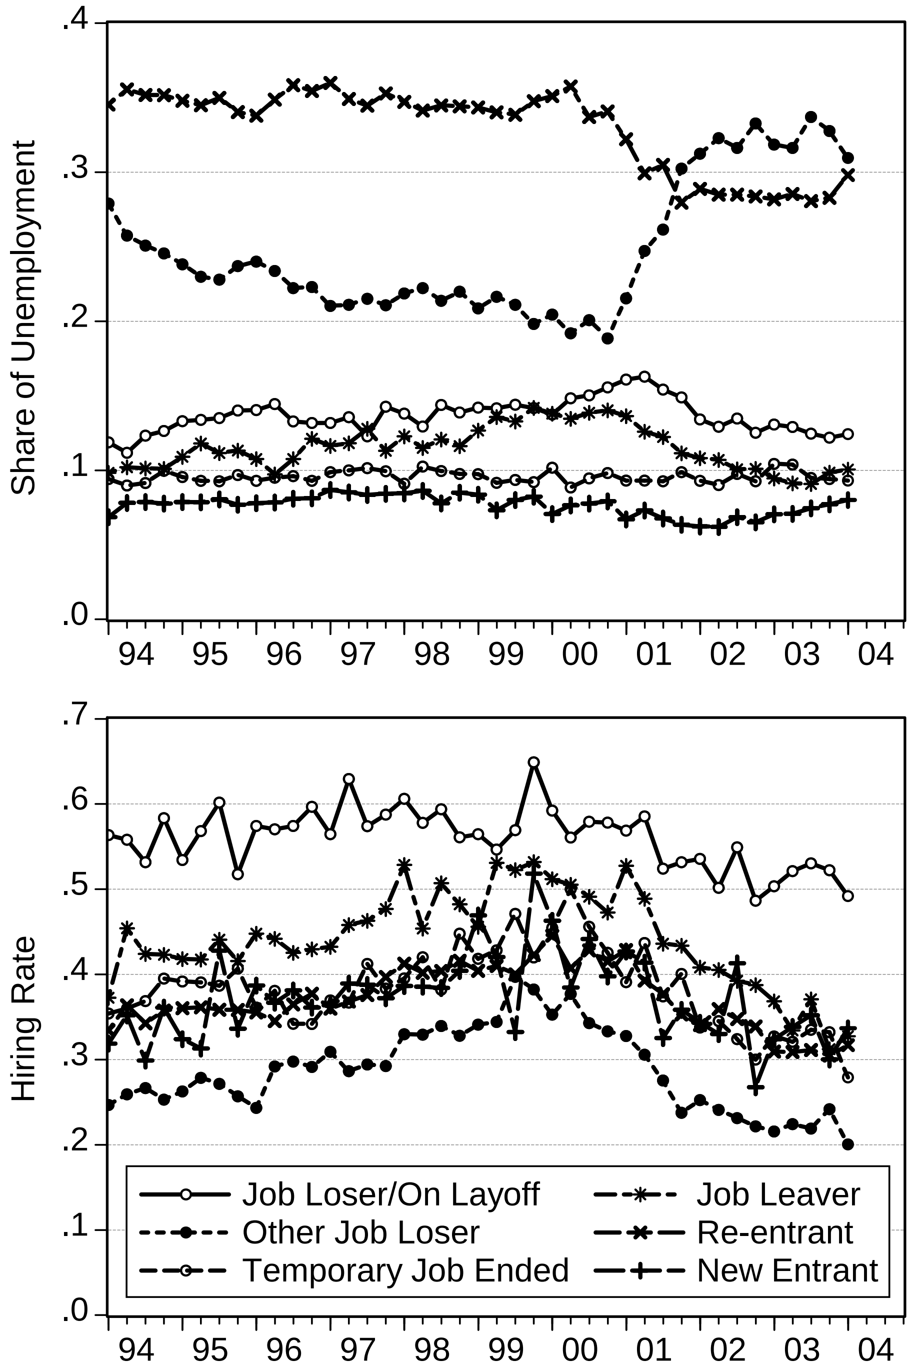

Figure 4: The Share of Unemployment and the Hiring Rate for Six Different Reasons for Unemployment, United States, 1994Q1–2004Q2, quarterly average of monthly data. The underlying data are constructed from the monthly CPS, seasonally adjusted, and averaged within quarters.

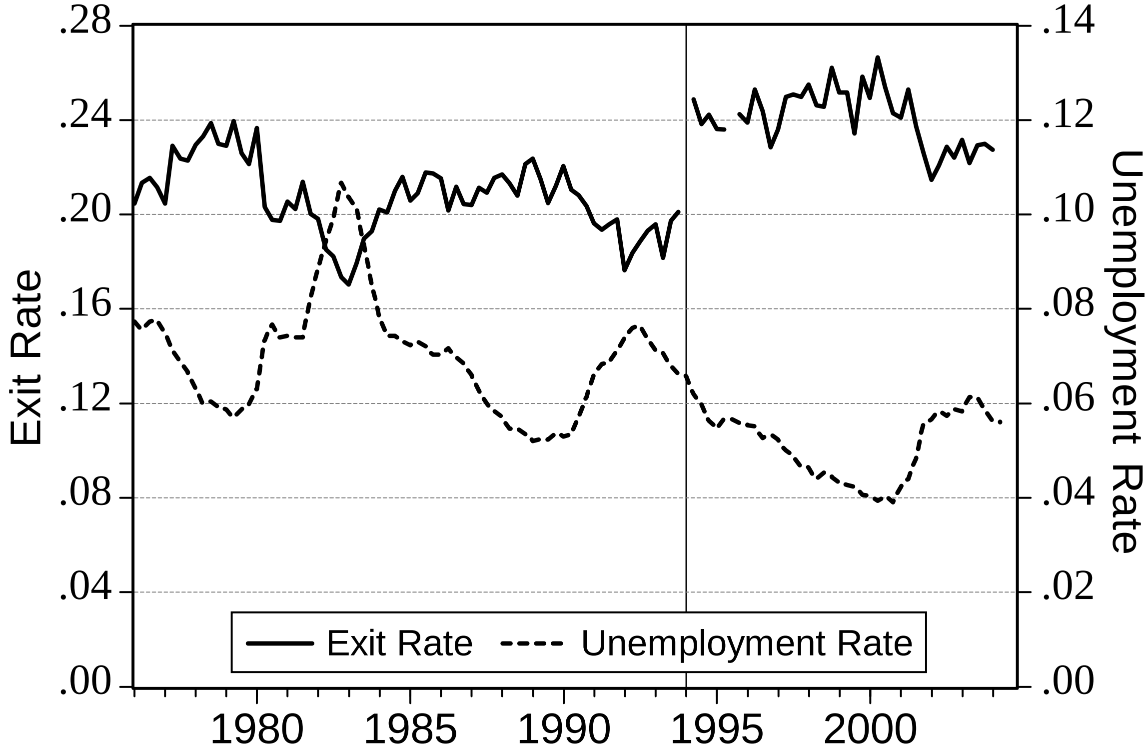

Figure 5: The Unemployment-Inactive Transition Rate, United States, 1976Q1–2004Q1, quarterly average of seasonally adjusted monthly data. The transition rate is constructed using matched records from the CPS. The matching algorithm uses household identifiers, individual line numbers, rotation group, race, age, and sex. The survey redesign makes the data before and after 1994Q1 incomparable.

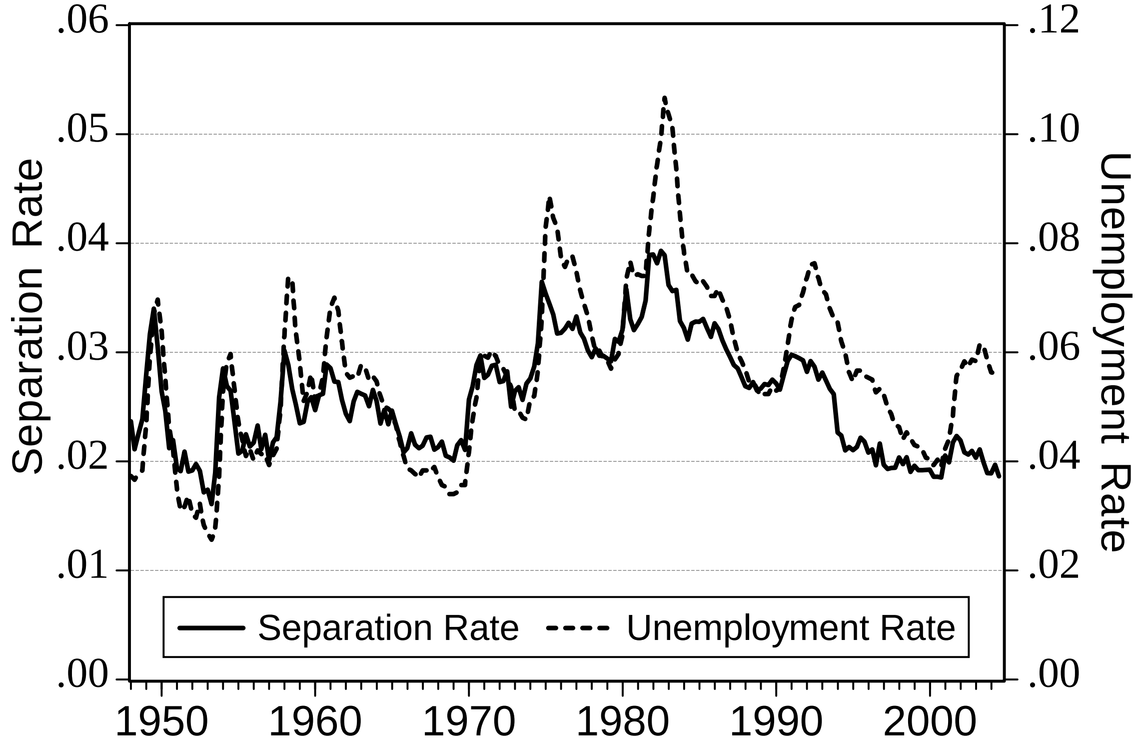

Figure 6: A Simple Measure of the Separation Rate and Unemployment Rate, United States, 1948Q1–2004Q2, quarterly average of monthly data. The separation rate is constructed from employment and short-term unemployment. Employment and short-term unemployment are computed by the BLS and seasonally adjusted.

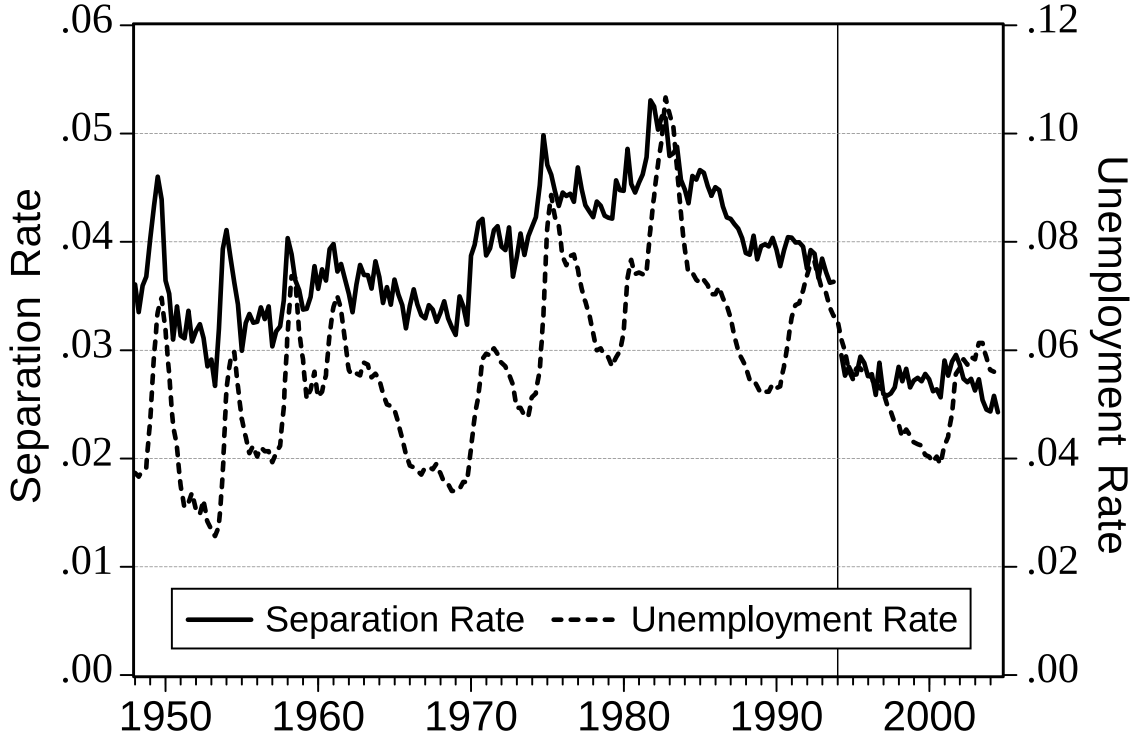

Figure 7: The Separation Rate and Unemployment Rate, United States, 1948Q1–2004Q2, quarterly average of monthly data. The separation rate is constructed from employment, short-term unemployment, and Hall's hiring rate according to equation (9). Employment and short-term unemployment are computed by the BLS and seasonally adjusted. Hall's hiring rate is constructed using short-term and medium-term unemployment data from the CPS as the solution to equation (6). The survey redesign makes the data before and after 1994Q1 incomparable.

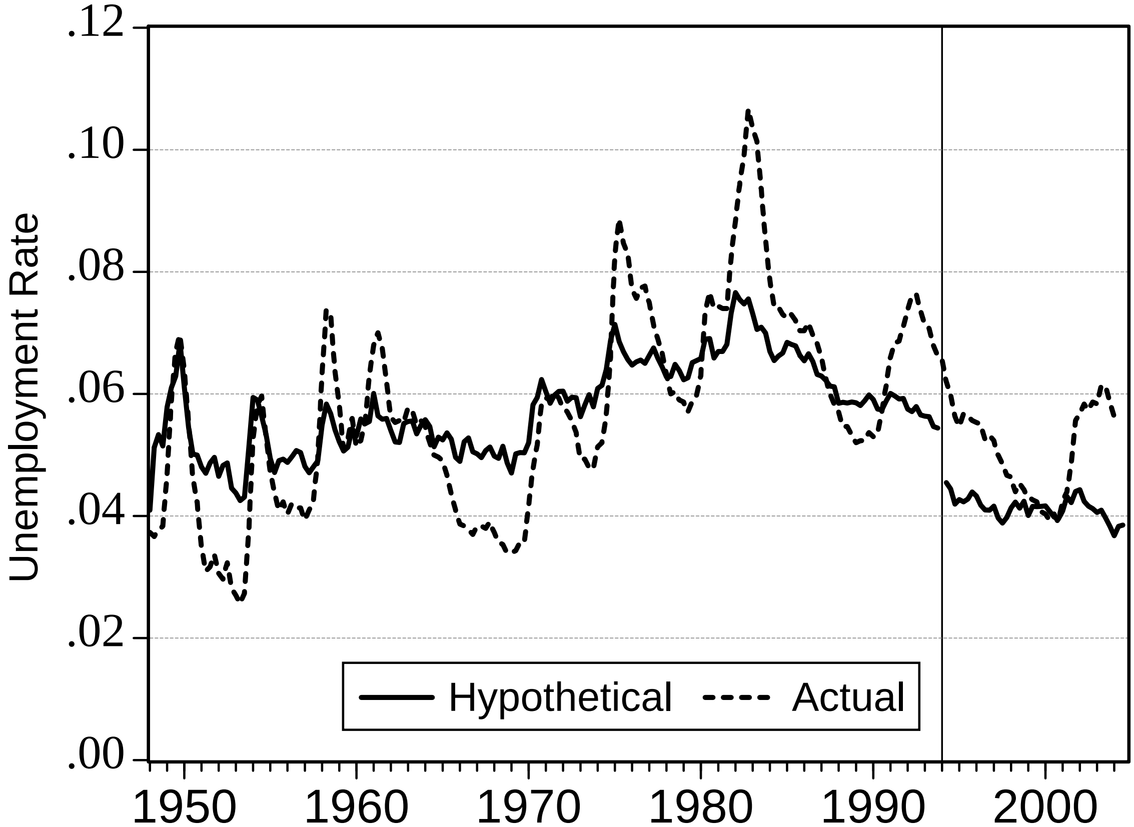

Figure 8: Hypothetical and Actual Unemployment Rates, United States, 1948Q1–2004Q2, quarterly average of monthly data. The hypothetical unemployment rate is constructed using equation (10). The separation rate is constructed from employment, short-term unemployment, and Hall's hiring rate according to equation (9). Employment and short-term unemployment are computed by the BLS and seasonally adjusted. Hall's hiring rate is constructed using short-term and medium-term unemployment data from the CPS as the solution to equation (6). The survey redesign makes the data before and after 1994Q1 incomparable.

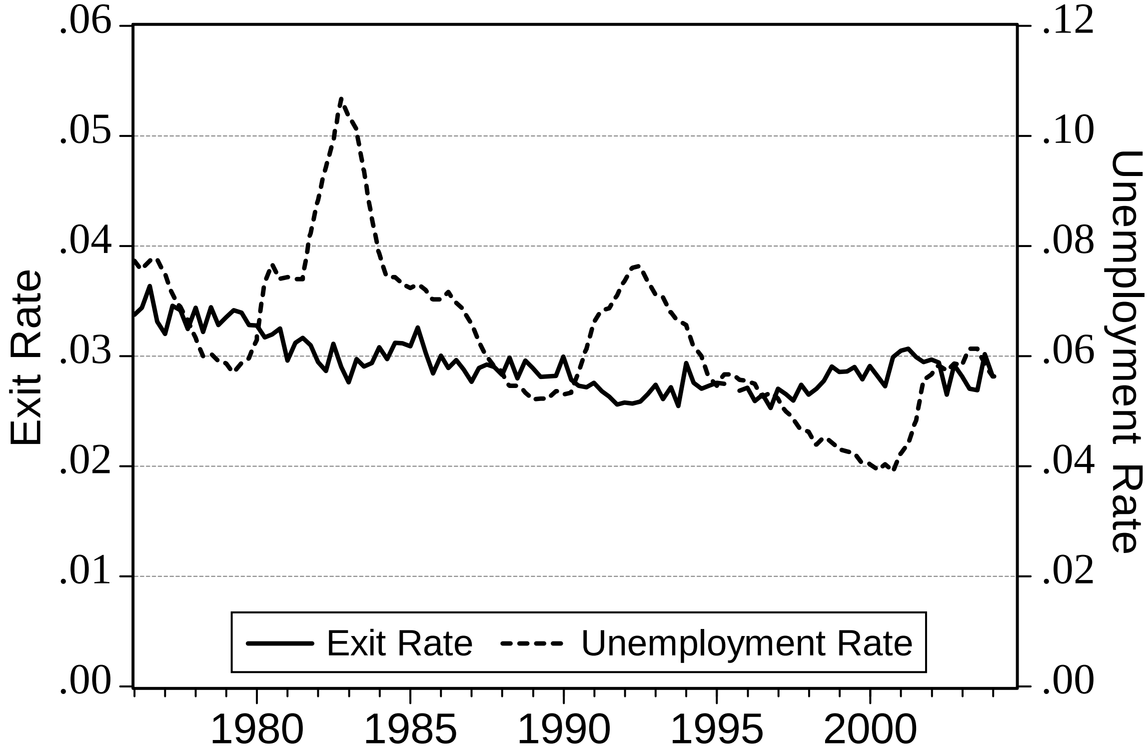

Figure 9: The Employment-Inactive Transition Rate, United States, 1976Q1–2004Q1, quarterly average of seasonally adjusted monthly data. The transition rate is constructed using matched records from the CPS. The matching algorithm uses household identifiers, individual line numbers, rotation group, race, age, and sex.

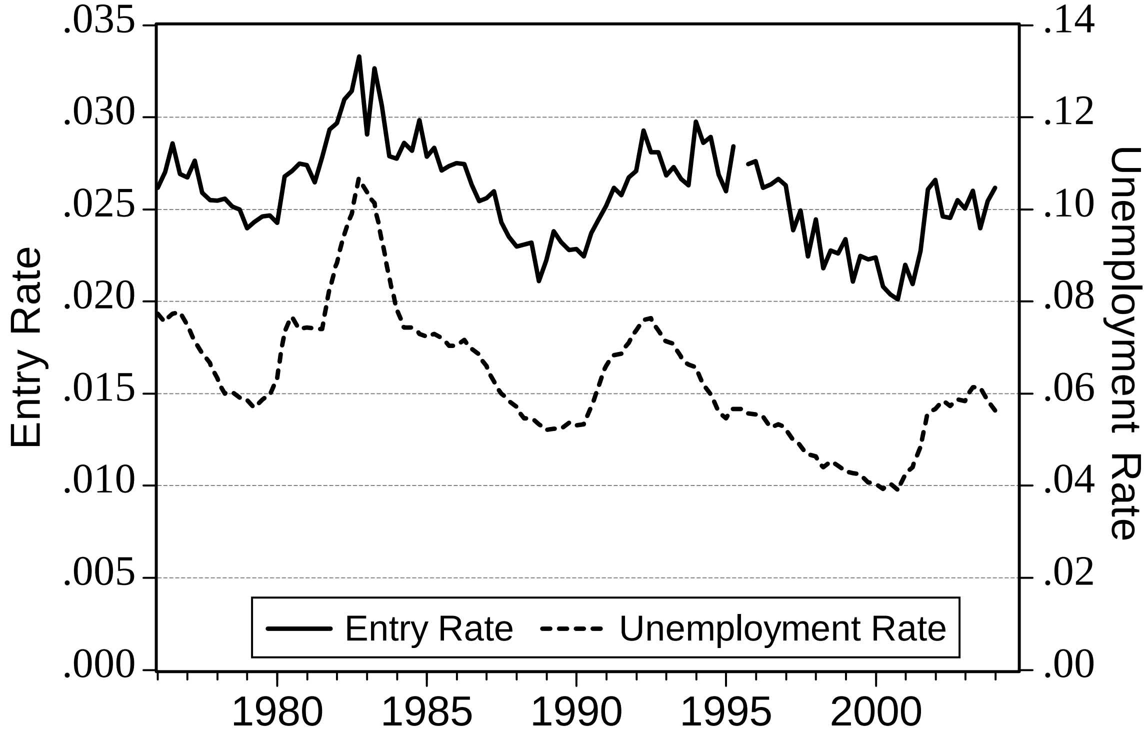

Figure 10: The Inactive-Unemployment Transition Rate, United States, 1976Q1–2004Q1, quarterly average of seasonally adjusted monthly data. The transition rate is constructed using matched records from the CPS. The matching algorithm uses household identifiers, individual line numbers, rotation group, race, age, and sex.

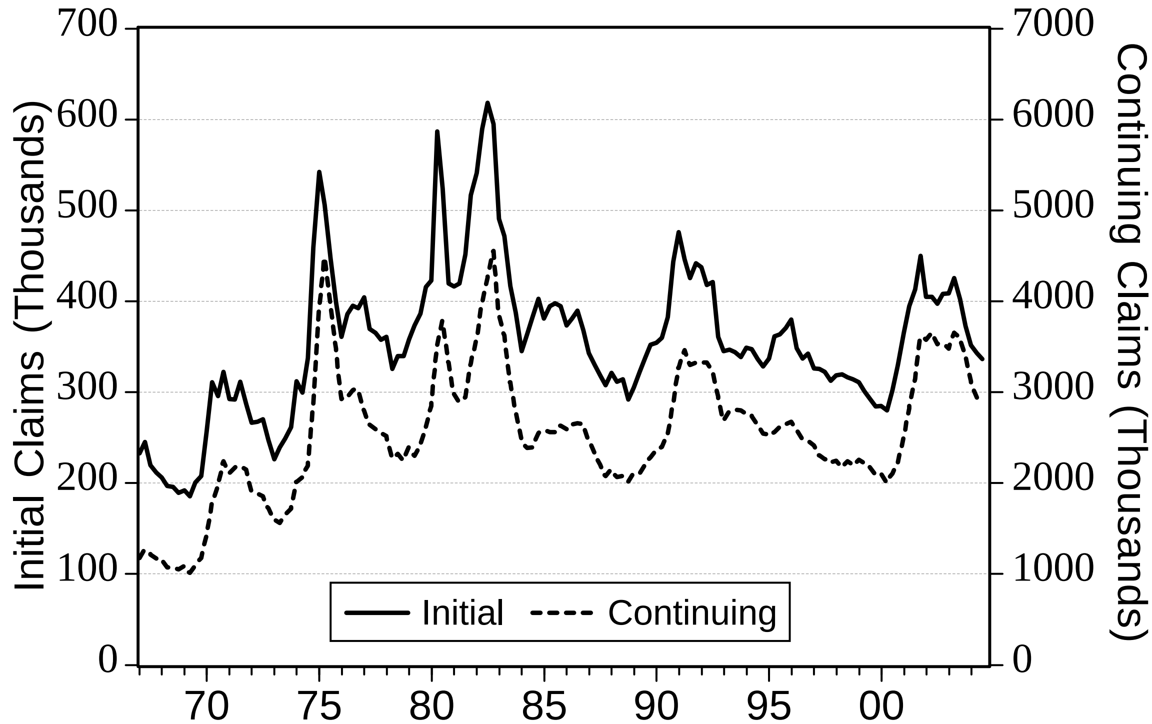

Figure 11: Initial and Continuing Unemployment Insurance Claims, United States, 1967Q1–2004Q3, quarterly average of seasonally adjusted weekly data. The data are available from `http://workforcesecurity.doleta.gov/unemploy/claims.asp`.

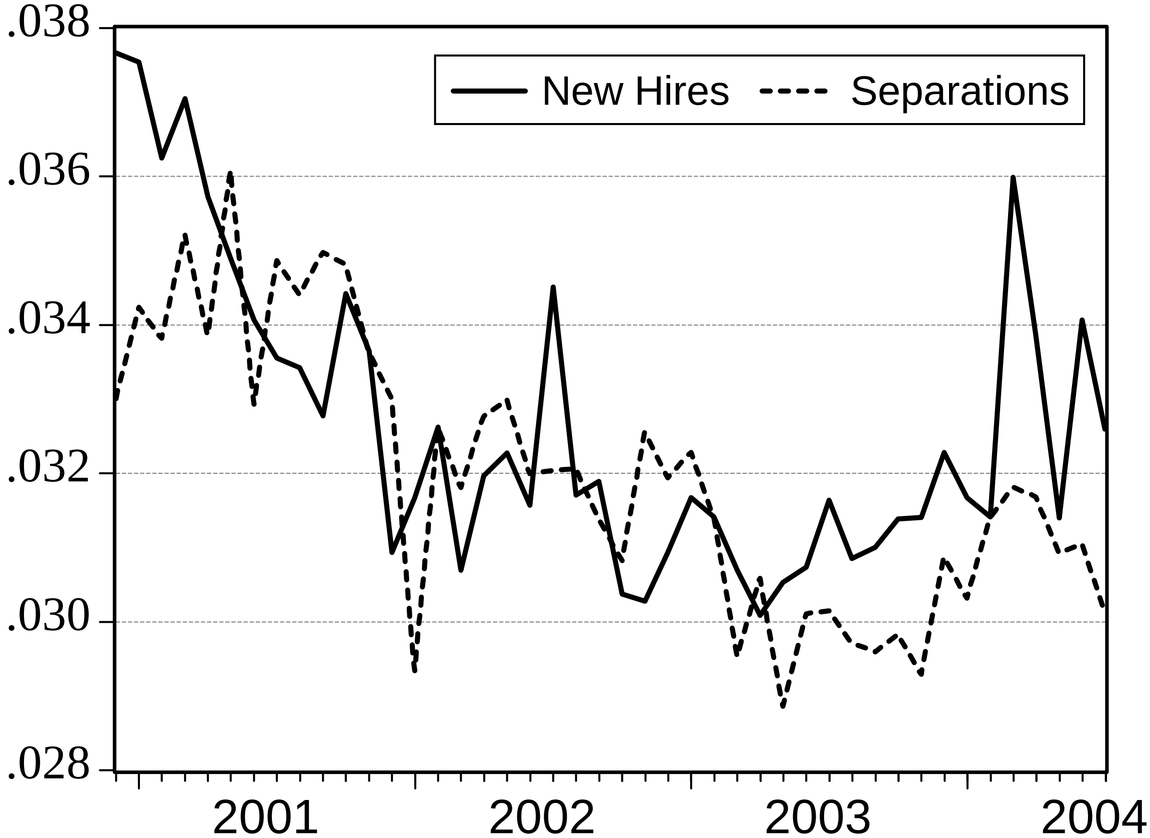

Figure 12: New Hires and Total Separations as a Percent of Employment, United States, December 2000–July 2004, seasonally adjusted. The data are constructed by the BLS as part of the JOLTS program.

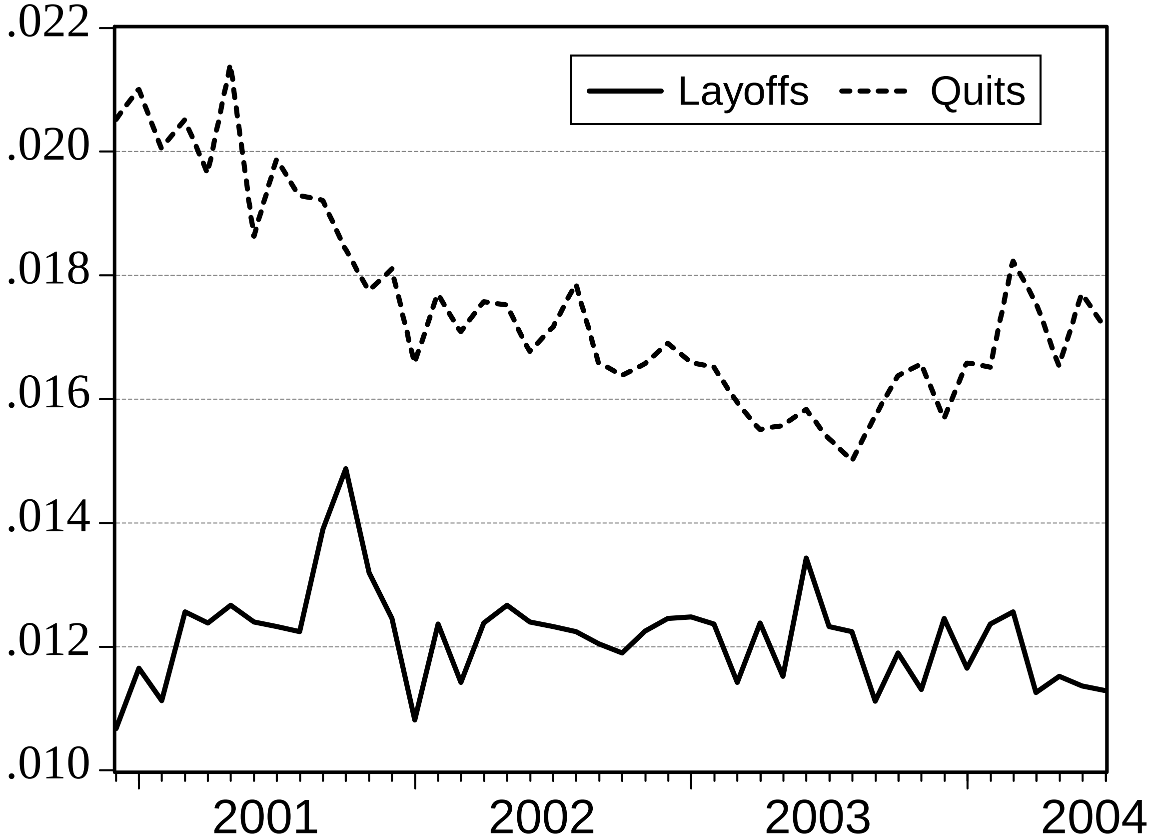

Figure 13: Layoffs and Quits as a Percent of Employment, United States, December 2000–July 2004, seasonally adjusted. The data are constructed by the BLS as part of the JOLTS program.

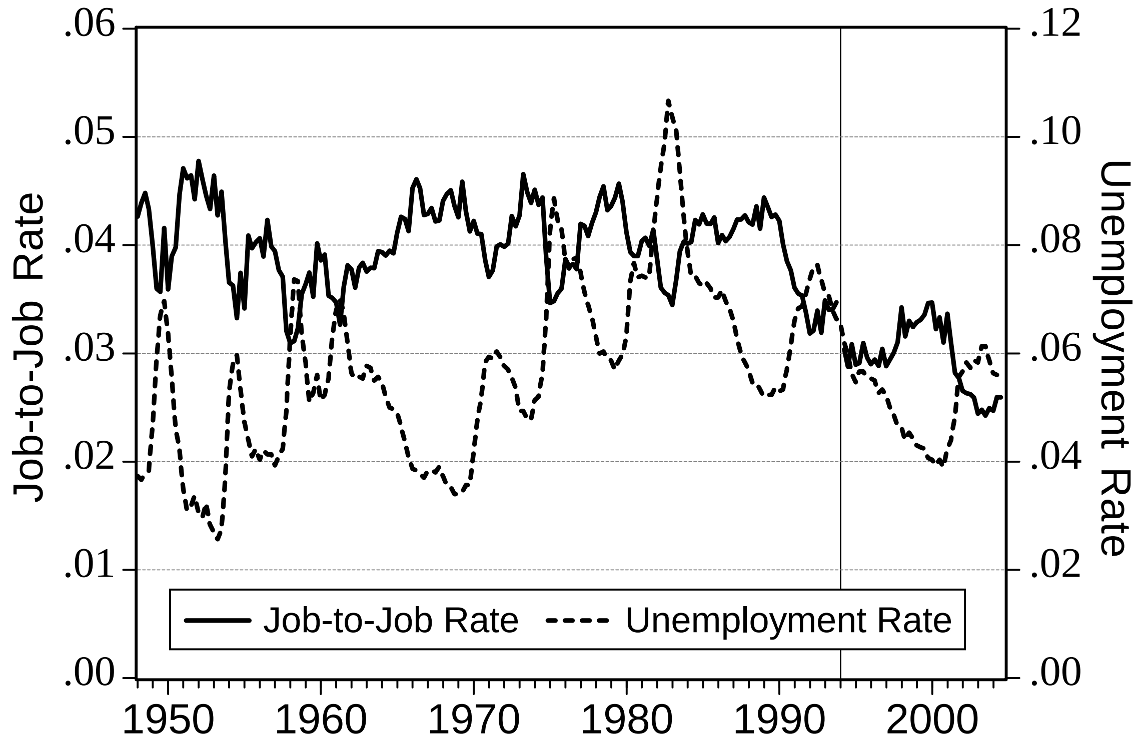

Figure 14: The Job-to-Job Transition Rate and Unemployment Rate, United States, 1948Q1–2004Q2, quarterly average of monthly data. The job-to-job transition rate is computed using equations (12) and (13) from underlying seasonally-adjusted BLS data from the CPS on employment and unemployment and different distributions. The survey redesign makes the data before and after 1994Q1 incomparable.

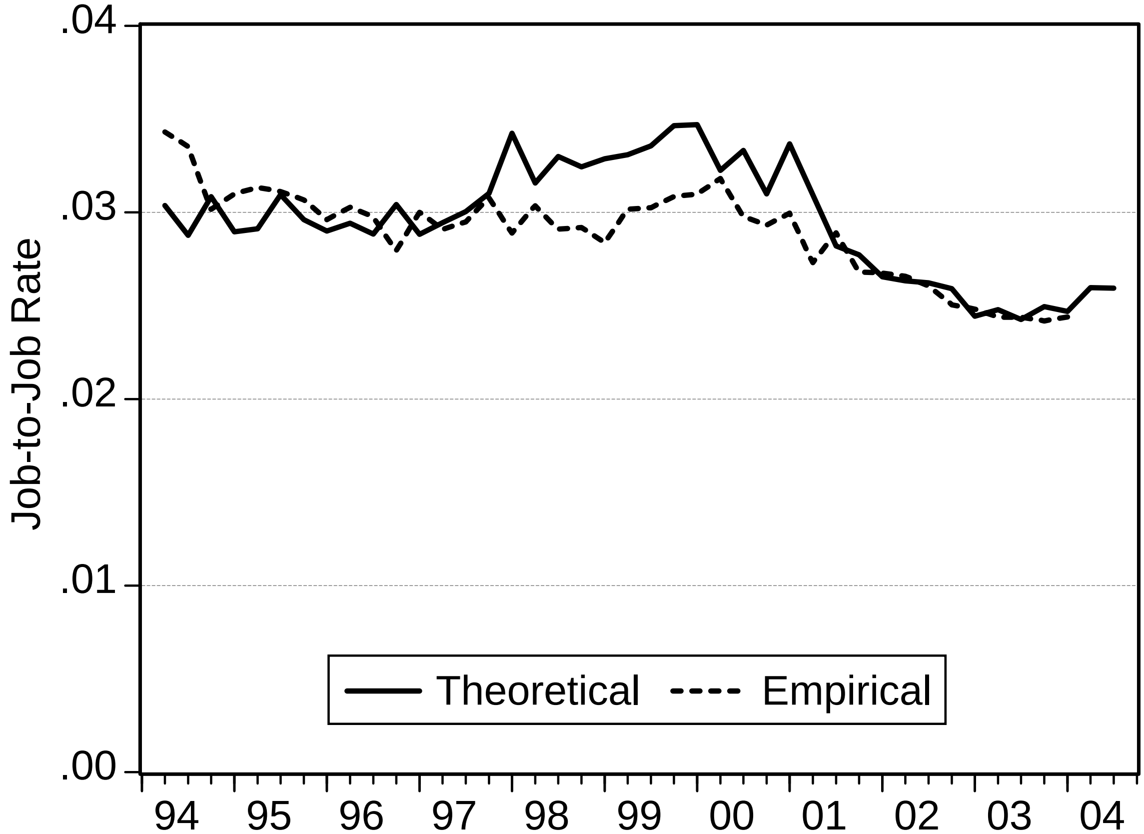

Figure 15: Theoretical and Empirical Job-to-Job Transition Rate, United States, 1994Q1–2004Q2, quarterly average of monthly data. The theoretical job-to-job transition rate is computed using equations (12) and (13) from underlying seasonally-adjusted BLS data from the CPS on employment and unemployment and different distributions. The empirical job-to-job transition rate is computed from public use CPS micro data as the fraction of employed workers who have a 'new employer' rather than the 'same employer as last month'.