Reassessing the Ins and Outs of Unemployment*

Robert Shimer
University of Chicago
shimer@uchicago.edu

First Version: January 16, 2005
Current Version: June 21, 2005

Abstract

This paper uses readily accessible data to measure the probability that an employed worker becomes unemployed and the probability that an unemployed worker finds a job, the ins and outs of unemployment. The job finding probability is strongly procyclical and the separation probability 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. These results contradict the conventional wisdom that has guided the development of macroeconomic models of the labor market during the last fifteen years.

*My title borrows from Darby, Haltiwanger, and Plant (1986). I am grateful for comments from Fernando Alvarez, Gadi Barlevy, Francesco Belviso, Tito Boeri, Robert Hall, David Laibson, and Randall Wright and from seminar participants at the Bank of Italy, Bocconi University, the Chicago Fed, Harvard University, the St. Louis Fed, and the University of Texas-Austin on an earlier version of this paper. This paper is supported by grants from the National Science Foundation and the Sloan Foundation.

1 Introduction

This paper measures the probability that an employed worker becomes unemployed and the probability that an unemployed worker finds a job. Using United States data from 1948 to 2004, I find that there are substantial fluctuations in unemployed workers' job finding probability at business cycle frequencies, while employed workers' separation probability is comparative acyclic. This is particularly true in the last two decades, during which period the separation probability has steadily declined despite two spikes in the unemployment rate. In other words, virtually all of the increase in unemployment and decrease in employment during the 1991 and 2001 recessions was a consequence of a reduction in the job finding probability. If one wants to understand fluctuations in unemployment, one must understand fluctuations in the transition rate from unemployment to employment, the 'outs of unemployment'. This conclusion is in direct opposition to the conventional wisdom, built around research by Darby, Haltiwanger, and Plant (1985) and (1986), Blanchard and Diamond (1990), and Davis and Haltiwanger (1990) and (1992), that recessions are periods characterized primarily by high job loss rates.

I base my conclusion on novel but simple measures of the job finding and separation probabilities. These measures rely on two strong assumptions: workers neither enter nor exit the labor force but simply transit between employment and unemployment; and all workers are *ex ante* identical, and in particular in any period all unemployed workers have the same job finding probability and all employed workers have the same separation probability. Given these assumptions, I show that the probability that an unemployed worker finds a job during a period can be expressed as a simple function of the number of unemployed workers at the start of the period, the number of unemployed workers at the end of the period, and the number of unemployed workers at the end of the period who were employed at some point during the period ('short-term unemployment'). The probability that an employed worker separates from her job can be found using the same data and the number of employed workers at the start of the period. Simple calculations using these data give me my preferred measures of the job finding probability and separation probability, shown in Figure 1. My estimate of the job finding probability is strongly positively correlated with a measure of the vacancy-unemployment ratio (Figure 4) and hence consistent with the predictions of a simple matching function (Pissarides 1985).

It is not surprising that strong assumptions deliver strong results, so this paper also explores what happens if I relax these assumptions. Consider first the restriction that workers neither enter nor exit the labor force. Once I relax this assumption, I can no longer use publicly available aggregate data on employment, unemployment, and short-term employment to construct the job finding and separation probabilities. Instead, I follow a standard methodology (Abowd and Zellner 1985, Poterba and Summers 1986, Blanchard and Diamond 1990) and use microeconomic data on individuals' employment status in consecutive months from 1968 to 2004 to construct time series for the gross flow of workers between employment, unemployment, and inactivity (out of the labor force). I then compute the job finding probability for unemployed workers and the separation probability for employed workers from these data. The surprising finding is that although this changes the level of the job finding and separation probabilities, it scarcely affects their fluctuations (Figure 5).

I then relax the restriction that all workers are homogeneous. The first question that arises is what exactly *the* job finding probability measures if different workers have a different job finding probability. I show that my methodology measures the probability that the average worker who is unemployed at the start of period t finds a job during period t. Other alternatives would give an identical measure of the job finding probability if workers were homogeneous, but have a predictable bias if workers are heterogeneous. United States data are consistent with the predicted bias.

Another issue that arises when workers are heterogeneous is whether that heterogeneity can explain fluctuations in the job finding probability. Darby, Haltiwanger, and Plant (1985) and (1986) argue that the job finding probability declines during recessions because workers who are unemployed during recessions are different than workers who unemployed during expansions. According to this theory, recessions are periods when prime age workers suffer permanent job loss in particularly large numbers. Such workers have a low probability of finding a job, but they would have a low job finding probability regardless of when they become unemployed. Darby, Haltiwanger, and Plant argue that this compositional effect drives down the measured job finding probability during recessions, a possibility that Baker (1992) labelled the "heterogeneity hypothesis." I test this hypothesis by examining the compositional variation of the unemployment pool along several different dimensions and find scant evidence in support of it.

Many previous authors have measured the cyclicality of the job finding and separation probabilities, but this paper offers several contributions to the existing literature.[1] First, I use data from the long booms of the 1980s and 1990s, during which period the separation probability has become noticeably less cyclical. Second, I use publicly available data whenever possible, making it easy for others to verify my results, extend them as more data becomes available, and examine their consistency both within the United States and across countries.[2] Third, I emphasize the importance of time aggregation throughout the paper, working explicitly in a continuous time model in which data are available at discrete intervals. I argue that ignoring time aggregation will bias a researcher towards finding a countercyclical separation probability, because when the job finding probability falls, a worker who loses her job is more likely to experience a measured spell of unemployment. Fourth, I stress the potential role of heterogeneity throughout my analysis, arguing that changes in the composition of the unemployed population do not drive my results.

The rest of this paper proceeds as follows. Section 2 proposes new measures of the job finding and separation probabilities that use readily accessible data and avoid the time aggregation bias. I then discuss the behavior of the job finding and separation probabilities in the United States from 1948 to 2004 and show the close link between the job finding probability and the vacancy-unemployment ratio. Section 3 relaxes the assumption that workers never enter or exit the labor force. I use gross flow data to measure the probability that a worker who is in one employment state at the beginning of the month (employed, unemployed, or inactive) switches to another employment state by the end of the month. Since workers can go through multiple states within a month, I then adjust these measures for time aggregation to get the instantaneous transition rates between employment states. I find a strong correlation between this measure of the unemployment-employment transition probability and the job finding probability and between the employment-unemployment transition probability and the separation probability.

Section 4 examines the role of heterogeneity. First I show that the simple measure of the job finding probability measures the mean job finding probability for an unemployed worker. Other simple measures of the job finding probability would be identical if workers were homogeneous, but with heterogeneous workers

[1] A companion paper, Shimer (2005b), puts the facts established in this paper through a simple model of on-the-job search. The model predicts that the job-to-job transition rate should be procyclical, consistent with existing evidence. On the other hand, if the job finding rate were acyclic and the separation rate countercyclical, the model would counterfactually predict a countercyclical job-to-job transition rate.

[2] Most of the time series I construct in this paper and the programs I use to construct them are available online at `http://home.uchicago.edu/~shimer/data/flows/`.

these correspond to a weighted average of the job finding probability for unemployed workers, over-weighting certain groups of workers, e.g. the long-term unemployed. I then address Darby, Haltiwanger, and Plant's (1986) heterogeneity hypothesis. I confirm that the unemployment pool switches towards 'job losers not on layoff' during recessions, and that these workers always have an unusually low job finding probability. Nevertheless, this explains little of the overall fluctuations in the job finding probability. Other dimensions of heterogeneity—age, sex, race, marital status, education, and geographic region—contribute virtually nothing to explaining fluctuations in the job finding probability.

Finally, Section 5 discusses the conventional wisdom on the cyclicality of the job finding and separation probabilities, especially the evidence presented by Davis and Haltiwanger (1990) and (1992). I argue that this evidence has largely been misinterpreted and may shed little light on the question of interest in this paper. Finally, I argue that this misinterpretation has profoundly influenced the development of macroeconomic models of the labor market during the past 15 years, including such well-known papers as Mortensen and Pissarides (1994) and Caballero and Hammour (1994). Subsequent research has focused on the cause of job loss during recessions rather than the difficulty of finding a job. Section 6 concludes by mentioning some recent research that attempts to address this shortcoming.

2 Simple Measures

In this section, I develop simple measures of the job finding probability for unemployed workers F_t and the separation probability for employed workers S_t. I then use publicly available data from the Current Population Survey (CPS) to measure the two transition probabilities in the United States from 1948 to 2004. I find that the job finding probability is strongly procyclical while the separation probability is less cyclical and explains little of the overall fluctuations in employment and unemployment, particularly during the last two decades.

To obtain simple measures of the job finding and separation probabilities, it is necessary to make strong assumptions. Throughout this section, I ignore movements in and out of the labor force, so workers simply transition between employment and unemployment. I also assume that all unemployed workers find a job with probability F_t and all employed workers lose a job with probability

S_t during period t, ignoring any heterogeneity or duration dependence that makes some unemployed workers more likely to find and some employed workers less likely to lose a job within the period. Sections 3 and 4 argue that these assumptions do not significantly affect my conclusions.

2.1 Theory

I model a continuous time environment in which data are available only at discrete dates. For $t \in \{0, 1, 2, \ldots\}$, refer to the interval $[t, t+1)$ as 'period t.' The goal is to recover the job finding probability $F_t \in [0, 1]$ and separation probability $S_t \in [0, 1]$ during period t from commonly available data. I assume that during period t, all unemployed workers find a job according to a Poisson process with arrival rate $f_t \equiv -\log(1 - F_t) \geq 0$ and all employed workers lose their job according to a Poisson process with arrival rate $s_t \equiv -\log(1-S_t) \geq 0$. Throughout this paper, I refer to f_t and s_t as the job finding and separation *rates* and to F_t and S_t as the corresponding *probabilities*, i.e. F_t is the probability that a worker who begins period t unemployed finds at least one job during the period and similarly for S_t.

Fix $t \in \{0, 1, 2, \ldots\}$ and let $\tau \in [0, 1]$ be the time elapsed since the last measurement date. Let $e_{t+\tau}$ denote the number of employed workers at time $t + \tau$, $u_{t+\tau}$ denote the number of unemployed workers at time $t + \tau$, and $u_t^s(\tau)$ denote 'short term unemployment', workers who are unemployed at time $t + \tau$ but were employed at some time $t' \in [t, t + \tau]$. Note that $u_t^s(0) = 0$ for all t. It is convenient to define $u_{t+1}^s \equiv u_t^s(1)$ as the total amount of short term unemployment at the end of period t.

For $t \in \{0, 1, 2, \ldots\}$ and $\tau \in [0, 1)$, unemployment and short term unemployment evolve according to

$$\dot{u}_{t+\tau} = e_{t+\tau} s_t - u_{t+\tau} f_t \tag{1}$$

$$\dot{u}_t^s(\tau) = e_{t+\tau} s_t - u_t^s(\tau) f_t. \tag{2}$$

Unemployment increases when employed workers separate, at an instantaneous rate s_t, and decreases when unemployed workers find jobs, at an instantaneous rate f_t. Short term unemployment increases when employed workers separate and decreases when short term unemployed workers find jobs.

To solve for the job finding probability, eliminate $e_{t+\tau} s_t$ between these equations, giving

$$\dot{u}_{t+\tau} = \dot{u}^s_t(\tau) - \big(u_{t+\tau} - u^s_t(\tau)\big) f_t$$

for $\tau \in [0,1)$. By construction, $u^s_t(0) = 0$, so given an initial condition for u_t, this differential equation can be solved for u_{t+1} and $u^s_{t+1} \equiv u^s_t(1)$:

$$u_{t+1} = (1 - F_t) u_t + u^s_{t+1} \tag{3}$$

The number of unemployed workers at date $t+1$ is equal to the number of unemployed workers at date t who do not find a job (fraction $1 - F_t = e^{-f_t}$) plus the u^s_{t+1} short term unemployed workers, those who are unemployed at date $t+1$ but held a job at some point during period t. Invert this,

$$F_t = 1 - \frac{u_{t+1} - u^s_{t+1}}{u_t}, \tag{4}$$

to express the job finding probability as a function of unemployment and short term unemployment.

One can also solve the differential equations (1) forward to obtain an implicit expression for the separation probability:

$$u_{t+1} = \frac{\big(1 - e^{-f_t - s_t}\big) s_t}{f_t + s_t} l_t + e^{-f_t - s_t} u_t, \tag{5}$$

where $l_t \equiv u_t + e_t$ is the size of the labor force during period t, which I assume is constant since I do not allow entry or exit from the labor force. Since $l_t > u_t$, the right hand side of this expression is increasing in s_t. Given the job finding probability from equation (4) and data on unemployment and employment, equation (5) uniquely defines the separation probability S_t.[3]

To understand equation (5), note first that if unemployment is constant during period t, the unemployment rate is determined by the ratio of the separation rate to the job finding rate, $\frac{u_t}{l_t} = \frac{s_t}{s_t + f_t}$, a standard formula. More generally, it helps to compare equation (5) with a discrete time model in which there is no possibility of both finding and losing a job within a period. In this case,

$$u_{t+1} = S_t e_t + (1 - F_t) u_t \tag{6}$$

[3]Shimer (2005a) measures the separation probability as $S_t = u^s_{t+1}/e_t\left(1 - \frac{1}{2}F_t\right)$, since a worker who loses her job has on average half a period to find a new one, and so does so with approximately probability $\frac{1}{2}F_t$. There is quantitatively very little difference between the two measures, but the one I use in this paper has more theoretical appeal.

A fraction S_t of employed workers lose their job and a fraction F_t of unemployed workers find a job during period t, determining the unemployment rate at the start of period $t+1$. When the time period is sufficiently short, or equivalently $s_t + f_t$ is sufficiently small, equation (5) converges to this simple expression. But with longer time periods, equation (5) allows workers to lose a job and find a new one, or vice versa, within the period.

The distinction between equations (5) and (6) is quantitatively important for measuring both the level of the separation probability and its cyclicality. When the job finding rate f_t is high, equation (5) captures the fact that a worker who loses her job is more likely to find a new one without experiencing a measured spell of unemployment. These separations are missed in equation (6), so the latter formula yields fewer separations and, more importantly for this paper, a negative bias in the measured correlation between the job finding and separation rates. Starting explicitly from a continuous time environment avoids this time aggregation bias.

2.2 Measurement

Since 1948, the Bureau of Labor Statistics (BLS) has published monthly data on employment, unemployment, and unemployment duration based on the CPS, downloadable from the BLS web site.[4] The measures of the number of employed and unemployed workers are standard, and I use these to quantify e_t and u_t. The survey also asks unemployed workers how long they have been unemployed and the BLS tabulates the number of unemployed workers with zero to four weeks duration. I use this as my measure of short term unemployment u_t^s from January 1948 to December 1993. Unfortunately, the redesign of the CPS instrument in 1994 introduced a significant discontinuity in the short term unemployment series (Abraham and Shimer 2001). Appendix A discusses a procedure to adjust the data after 1994.

Figure 1 shows the time series for the job finding probability F_t and separation probability S_t constructed according to equations (4) and (5) from 1948 to 2004. Several facts stand out. First, the job finding probability is high, averaging 46 percent over the post-war period. Second, it is variable, falling by about forty log points from peak to trough during recent decades. Third, the separation probability averaged 3.5 percentage points during the same period and was somewhat less volatile, particularly in recent years.

[4] `http://www.bls.gov/cps/`

To examine the cyclicality of the job finding and separation probabilities, recall that if unemployment were constant, $u_t = u_{t+1}$, equation (5) implies that the unemployment rate would be $\frac{u_t}{l_t} = \frac{s_t}{s_t+f_t}$. In fact, $\frac{s_t}{s_t+f_t}$ is a very good approximation to the end-of-month unemployment rate; in monthly data, the correlation between $\frac{u_{t+1}}{l_{t+1}}$ and $\frac{s_t}{s_t+f_t}$ is 0.99. I use this strong relationship to distinguish between the importance of fluctuations in the job finding and separation rates for fluctuations in unemployment. Let $\bar{f}$ and $\bar{s}$ denote the average values of f_t and s_t during the sample period and compute $\frac{\bar{s}}{\bar{s}+f_t}$ and $\frac{s_t}{s_t+\bar{f}}$ as measures of the contributions of fluctuations in the job finding and separation rates to overall fluctuations in the unemployment rate.

The top panel in Figure 2 shows that a decline in the job finding rate f_t contributed to every increase in the unemployment rate during the post-war period. The bottom panel shows that from 1948 to 1985, the separation rate tended to move with the unemployment, although it rarely explained more than half the fluctuation in unemployment. In the last two decades, however, the separation rate has varied little over the business cycle. One way to quantify this is to look at the comovement of the detrended data.[5] Over the entire post-war period, the correlation between the cyclical components of $\frac{u_{t+1}}{l_{t+1}}$ and $\frac{\bar{s}}{\bar{s}+f_t}$ is 0.97 while the correlation between $\frac{u_{t+1}}{l_{t+1}}$ and $\frac{s_t}{s_t+\bar{f}}$ is somewhat lower, 0.71. The latter correlation has fallen to 0.15 since 1986, while the former correlation is unchanged. Moreover, $\frac{\bar{s}}{\bar{s}+f_t}$ is relatively volatile, with a cyclical standard deviation equal to 0.78 times that of $\frac{u_{t+1}}{l_{t+1}}$. The relative standard deviation of $\frac{s_t}{s_t+\bar{f}}$ is just 0.35.

Although not the main topic of this paper, it seems worth commenting on the secular decline the separation probability since the early 1980s (Figure 1). This finding would appear to contradict a sizable literature that finds evidence for a constant or even increasing separation rate during the 1980s and early 1990s.[6] For example, Gottschalk and Moffitt (1999) write, "Almost all studies based on the various Current Population Surveys (CPS) supplements ... show little change in the overall separation rates through the early 1990s." Much of the difference appears to be due to differences in samples. For example,

[5] I time-aggregate the 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 the cyclical volatility in the variable of interest. I use this same filter throughout the paper.

[6] To my knowledge, no previous paper has studies the separation rate over a fifty year period. All of the studies cited in Gottschalk and Moffitt (1999) start in 1968 or later.

Gottschalk and Moffitt (1999) study married men age 20–62, while I examine the entire population. During the last two decades, the labor force has aged; since younger workers have the highest separation rates, this has reduced the separation rate. In addition, women have become increasingly attached to the labor force, further reducing turnover. Consistent with that view, Figure 3 indicates no trend in the separation probability for 25 to 54 year old men since 1976.[7] Nevertheless, fluctuations in the job finding and separation rates are about the same for this more homogeneous group as for the population at large.

2.3 The Matching Function

Pissarides (1985) proposes that the job finding rate f_t should be an increasing function of the ratio of job vacancies to unemployment, θ_t. More precisely, he argues that there is a matching function that gives the number of new employment relationships formed as a function of unemployment and vacancies. If the matching function is stable and has constant returns to scale, then the job finding rate will be increasing in the vacancy-unemployment ratio. I have already argued that the job finding rate is negatively correlated with unemployment. Moreover, the Beveridge Curve—the strong negative correlation between vacancies and unemployment at business cycle frequencies—has been noted in earlier research (Abraham and Katz 1986, Blanchard and Diamond 1989, Shimer 2005a). This suggests that one should indeed observe a positive relationship between the job finding rate and the vacancy-unemployment ratio.

To examine this, I use a crude but standard measure of vacancies, the Conference Board Help Wanted Advertising Index (Abraham 1987), downloadable from the Federal Reserve Bank of St. Louis's FRED®II database.[8] Since 1951, the Conference Board has constructed the index on a monthly basis by measuring the number of column inches of help wanted advertisements in the largest newspaper in 51 major metropolitan areas. Consolidation of the newspaper industry, changes in newsprint costs, legally mandated changes in advertising like equal employment opportunity laws, and the rise of the internet likely all affected the help wanted index. Fortunately, none of these should affect the cyclical behavior of the help wanted series.

Figure 4 shows that, even without addressing the measurement problems in the help wanted advertising index, there is a strong relationship between

[7] See Abraham and Shimer (2001) for a further discussion of the impact of demographic change on unemployment duration.

[8] `http://research.stlouisfed.org/fred2/series/helpwant/`.

the job finding rate f_t and the vacancy-unemployment ratio θ_t. In fact, the correlation between the unfiltered quarterly series for $\log f_t$ and $\log \theta_t$ is 0.92, and removing a low frequency trend raises the correlation to 0.96. Although this does not explain why the job finding rate is so cyclical, it rules out some possibilities. For example, suppose unemployed workers were, for some reason, less effective at finding jobs during recessions. Then the matching function would not be stable, nor would the relationship between the job finding rate and the vacancy-unemployment ratio. Instead, the obvious, if facile, explanation for the cyclicality of the job finding rate is that there are fewer jobs available during recessions.

3 Entry and Exit from the Labor Force

I have so far assumed that all workers are either unemployed or employed and ignored transitions in and out of the labor force. This section explores the importance of this restriction by examining the gross flow of workers between three labor market states, employment (E), unemployment (U), and inactivity (I).

3.1 Theory

As with the job finding and separation probabilities, I account for time aggregation bias by modelling a continuous time environment in which data are available only at discrete dates $t \in \{0, 1, 2, \ldots\}$. Let λ_t^{XY} denote the Poisson arrival rate of a shock that moves a worker from state $X \in \{E, U, I\}$ to state $Y \neq X$ during period t. $\Lambda_t^{XY} \equiv 1 - e^{-\lambda_t^{XY}}$ is the associated full-period transition probability.

I cannot measure the transition probabilities directly since workers may move through multiple states within a period. Instead, I have 'gross flow' data measuring the number of workers who were in state X at date t and are in state Y at date $t+1$. To see how this is useful, let $N_t^{XY}(\tau)$ denote the number of the workers who were in state $X \in \{E, U, I\}$ at date t and are in state $Y \in \{E, U, I\}$ at date $t+\tau$. Also define $n_t^{XY}(\tau) \equiv \frac{N_t^{XY}(\tau)}{\sum_Z N_t^{XZ}(\tau)}$, the associated share of workers who were in state X at t. Note that $N_t^{XY}(0) = n_t^{XY}(0) = 0$ for all $X \neq Y$. It is useful to think of a worker's state as including both her employment status at the last measurement date X and her current status Y, say XY. Then $n_t^{XY}(\tau)$ evolves according to a differential equation:

$$\dot{n}_t^{XY}(\tau) = \sum_Z n_t^{XZ}(\tau)\lambda_t^{ZY} - n_t^{XY}(\tau) \sum_Z \lambda_t^{YZ}. \tag{7}$$

The share of workers who are in state XY increases when a worker in some other state XZ transitions to XY and decreases when a worker in state XY transitions to XZ. All of these transition rates λ only depend on a worker's current employment status, i.e. Y or Z, and not on her start-of-period employment status X.

Given initial conditions and the restriction that $\sum_Z n_t^{XZ} = 1$, the differential equation system (7) can be solved for the six fractions $n_t^{XY}(1)$, $X \neq Y$, as functions of the six transition rates λ_t^{XY}, $X \neq Y$. To simplify this step, note that the system may be uncoupled into three two-dimensional linear differential equations, each of which depends on all six instantaneous transition rates. The resulting equations are messy and apparently cannot be solved analytically for the λ's.[9] Nevertheless, given data on the gross flow of workers from state X to state Y in period t, $N_t^{XY}(1)$, it is possible to compute the shares $n_t^{XY}(1)$ and then invert these equations numerically to recover the instantaneous transition rates λ_t^{XY} and hence the transition probabilities Λ_t^{XY}.

3.2 Measurement

To measure the gross flows $N_t^{XY}(1)$, I follow an approach adopted by many previous authors, perhaps most prominently by Blanchard and Diamond (1990).[10] The CPS is a rotating panel, with each household in the survey for four consecutive months. This makes it feasible to match as many as three-quarters of the survey records in the microdata files across months. Using these matched records, one can construct the gross flows.

Before 1976, I do not have access to the microdata and so I use Joe Ritter's tabulation of the gross flows from June 1967 to December 1975.[11] For the later

[9] In the two-state case with only employment and unemployment, the state equations can be written as

$$n_t^{XY}(1) = \lambda_t^{XY} \left(\frac{1 - e^{-\lambda_t^{XY} - \lambda_t^{YX}}}{\lambda_t^{XY} + \lambda_t^{YX}} \right)$$

for $X \neq Y$, so both instantaneous transition rates affect both gross flows. This can be inverted analytically to give

$$\lambda_t^{XY} = n_t^{XY}(1) \frac{-\log\left(1 - n_t^{XY}(1) - n_t^{YX}(1)\right)}{n_t^{XY}(1) + n_t^{YX}(1)}$$

for $X \neq Y$. In the three-state case, I cannot prove that the instantaneous transition rates are uniquely defined by gross flows, but for the values of n_t^{XY} in U.S. data, this does not appear to be an issue.

[10] See Abowd and Zellner (1985) and Poterba and Summers (1986) for discussions of measurement problems in gross flows data.

[11] I am grateful to Hoyt Bleakley for providing me with that data.

period, the monthly CPS public-use microdata are available from the NBER website.[12] I use these to construct my own time series for the gross flows. Starting with about 30 gigabytes of raw CPS data files, I match individual records from consecutive months using rotation groups, household identifiers, individual line numbers, race, sex, and age. I obtain more than 30 million matched records during the sample period, 92,361 in an average month. Using these, I compute the sample-weighted transition probabilities between employment states during the relevant month and seasonally adjust the time series using a ratio-to-moving average technique. This gives me series for the six gross flows $N_t^{XY}(1)$.[13] Finally, I adjust for time aggregation bias using the technique described in the previous subsection and recover time series for the instantaneous transition rates λ_t^{XY} and the transition probabilities Λ_t^{XY}.

The top panel in Figure 5 compares the job finding probability F_t, computed according to equation (4) from publicly available data on unemployment and short term unemployment, with the UE transition probability Λ_t^{UE}, computed using matched CPS files according to the procedure described in Section 3.1. Although the two series are constructed from entirely different data, their behavior is remarkably similar. They are equally volatile and their correlation is 0.94 in quarterly-averaged data. On the other hand, the job finding probability is consistently about 32 log points higher than the UE transition probability. This is probably because the former measure presumes that all workers exiting unemployment do so in order to take a job while the latter measure recognizes that some unemployment spells end when a worker exits the labor force. In any case, the level difference between the two probabilities is inconsequential for the cyclical behavior of the job finding probability. Gross worker flow data from the CPS confirm this paper's thesis that the job finding probability is strongly procyclical.

The bottom panel in Figure 5 shows the analogous comparison between the separation probability S_t and the EU transition probability Λ_t^{EU}. The correlation between the two series is 0.83 in quarterly-averaged data, with S_t average 59 log points higher than Λ_t^{EU}. Moreover, the amplitude of the fluctuations in both series at low frequencies is similar, although the EU transition probability

[12] <http://www.nber.org/data/cps_basic.html>. Unfortunately, there are a few gaps in the series due to changes in the household identifiers in the public-use files. It is impossible to match data for Dec. 1975/Jan. 1976, Dec. 1977/Jan. 1978, Jun. 1985/Jul. 1985, Sep. 1985/Oct. 1985, Dec. 1993/Jan. 1994, and May 1995/Jun. 1995 to Aug. 1995/Sep. 1995.

[13] Hoyt Bleakley also provided me with his independent estimates of gross flows from January 1976 to May 1993. During the overlapping period, the two series are virtually identical; the standard deviation of the log of the ratio of the two sets of series is less than 1 percent.

tends to fluctuate a bit more at business cycle frequencies. Notably, while the separation probability scarcely budged during the 1991 and 2001 recessions, the EU transition probability increased modestly.

To quantify the importance of changes in the six transition rates for fluctuations in the unemployment rate, it is again useful to do some steady state calculations. In steady state, the flows in and out of employment are equal, as are the flows in and out of unemployment:

$$(\lambda^{EU} + \lambda^{EI})e = \lambda^{UE}u + \lambda^{IE}i \quad \text{and} \quad (\lambda^{UE} + \lambda^{UI})u = \lambda^{EU}e + \lambda^{IU}i,$$

where e, u, and i are the number of employed, unemployed, and inactive individuals. Manipulate these equations to get

$$\begin{aligned} e &= k\big(\lambda^{UI}\lambda^{IE} + \lambda^{IU}\lambda^{UE} + \lambda^{IE}\lambda^{UE}\big) \\ u &= k\big(\lambda^{EI}\lambda^{IU} + \lambda^{IE}\lambda^{EU} + \lambda^{IU}\lambda^{EU}\big) \\ i &= k\big(\lambda^{EU}\lambda^{UI} + \lambda^{UE}\lambda^{EI} + \lambda^{UI}\lambda^{EI}\big), \end{aligned}$$

where k is a constant set so that e, u, and i sum to the relevant population.

In Section 2 I argued that $\frac{s_t}{f_t+s_t}$ is almost identical to the unemployment rate. Analogously, if the economy were in steady state at some date t, the unemployment rate in a three-state system would equal

$$\frac{\lambda_t^{EI}\lambda_t^{IU} + \lambda_t^{IE}\lambda_t^{EU} + \lambda_t^{IU}\lambda_t^{EU}}{\big(\lambda_t^{EI}\lambda_t^{IU} + \lambda_t^{IE}\lambda_t^{EU} + \lambda_t^{IU}\lambda_t^{EU}\big) + \big(\lambda_t^{UI}\lambda_t^{IE} + \lambda_t^{IU}\lambda_t^{UE} + \lambda_t^{IE}\lambda_t^{UE}\big)}.$$

This is also a good approximation. In quarterly-averaged data, the correlation between this steady state measure and next month's unemployment rate is 0.99.

This suggests a method of calculating the contribution of changes in each of the six transition rates to fluctuations in the unemployment rate. To be concrete, focus on the UI transition rate. Define

$$\begin{aligned} e_t^{UI} &= \lambda_t^{UI}\bar{\lambda}^{IE} + \bar{\lambda}^{IU}\bar{\lambda}^{UE} + \bar{\lambda}^{IE}\bar{\lambda}^{UE} \\ u_t^{UI} &= \bar{\lambda}^{EI}\bar{\lambda}^{IU} + \bar{\lambda}^{IE}\bar{\lambda}^{EU} + \bar{\lambda}^{IU}\bar{\lambda}^{EU} \\ i_t^{UI} &= \bar{\lambda}^{EU}\lambda_t^{UI} + \bar{\lambda}^{UE}\bar{\lambda}^{EI} + \lambda_t^{UI}\bar{\lambda}^{EI}, \end{aligned} \tag{8}$$

where $\bar{\lambda}^{XY}$ is the average XY transition rate from 1967 to 2004. That is, only λ_t^{UI} is permitted to vary over time, with the other five transition rates fixed at their average values. Then the contribution of fluctuations in the unemployment-inactivity transition rate to changes in the unemployment rate is $\frac{u_t^{UI}}{e_t^{UI}+u_t^{UI}}$. Calculate the contribution of the other five transition rates in a similar fashion.

Figure 6 shows the resulting time series, with the actual unemployment rate plotted for comparison. Fluctuations in the UE transition rate (middle left panel) are the most important factor in determining changes in the unemployment rate. In the 1970s and 1980s, the EU transition rate (top left) also rose during downturns, contributing to the increase in the unemployment rate; however, this factor has become much less important during the last two decades. On the other hand, a decrease in the UI transition rate (middle right) tends to raise the unemployment rate during downturns. This suggests that unemployed workers are more attached to the labor force during downturns than they are during expansions, a possibility I return to in Section 4 when I examine cyclical changes in the composition of the unemployed population. In recent decades, this last factor has been about as important as changes in the EU transition rate for explaining the fluctuations in the unemployment rate. The remaining three transition rates are irrelevant for fluctuations in the unemployment rate.

An advantage to looking at a system in which workers move in and out of the labor force is that I can distinguish between fluctuations in the unemployment rate $\frac{u_t}{e_t+u_t}$ and fluctuations in the employment-population ratio $\frac{e_t}{e_t+u_t+i_t}$. Following the same methodology, Figure 7 graphs the contribution of each of the six transition rates to fluctuations in the employment-population ratio. This picture is more muddled than Figure 6. For example, the low frequency trend in the employment-population ratio is driven primarily by a decline in the EI transition rate, which reflects an increase in women's labor force attachment (Abraham and Shimer 2001).

I quantify the effect of each transition rate on the employment-population ratio at business cycle frequencies by detrending the data and running a simple regression. Start with the UE transition rate. I construct $\frac{e_t^{UE}}{e_t^{UE}+u_t^{UE}+i_t^{UE}}$ as described above and detrend the quarterly average of the monthly time series using a low-frequency HP filter (smoothing parameter 10^5). I then regress this on the detrended employment-population ratio and examine the coefficient. I find that a 1 percentage point cyclical increase in the employment-population ratio is associated with a 1.04 percentage point increase in $\frac{e_t^{UE}}{e_t^{UE}+u_t^{UE}+i_t^{UE}}$, so UE fluctuations are critical for changes in the employment-population ratio. The second most important determinant is the IE transition rate (regression coefficient 0.64), which reflects the lower likelihood that an inactive worker finds a job during a downturn. Turning to measures of the separation rate, the EU transition rate tends to rise when the employment-population ratio falls (0.42), but this is mostly offset by a decline in the EI transition rate (-0.31). In net, the probability of leaving employment scarcely affects the employment-population ratio at business cycle frequencies, while fluctuations in the probability of finding a job drive both the unemployment rate and the employment-population ratio.

4 Heterogeneity

This section relaxes the assumption that all workers are homogeneous. I first show that if some workers are more likely to find a job than others, F_t measures the mean job finding probability among unemployed workers. Using other moments of the unemployment duration distribution, one can construct other weighted averages of the job finding probability for unemployed workers, all of which co-move with the job finding probability. I then ask why the job finding probability declines during recessions. Is it because all unemployed workers are less likely to find a job or because the type of workers who becomes unemployed during a recession is somehow different, less likely to find a job regardless of the stage of the business cycle, as Darby, Haltiwanger, and Plant (1985) and (1986) suggest? I find no evidence to support the latter 'heterogeneity hypothesis' (Baker 1992).

4.1 Accounting for Heterogeneity

Suppose unemployed workers are heterogeneous. For example, long term unemployment may diminish a worker's prospect of finding a job. Alternatively, some time-invariant characteristic may affect the job finding probability, so a dynamic selection process makes it appear that the long term unemployed are less likely to find a job. In its most general form, one can model heterogeneity in the job finding probability by indexing the u_t unemployed workers at time t by $i \in \{1, \ldots, u_t\}$ and letting F_t^i denote the probability that worker i finds a job during month t. Then one can generalize equation (3) to the case where F_t^i varies with i:

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

where I assume for simplicity that the randomness in the outcome of the job finding process cancels out in the aggregate so u_{t+1} is not a random variable.

End-of-month unemployment is equal to the number of unemployed workers who fail to obtain a job within the month, $\sum_{i=1}^{u_t}(1-F_t^i)$, plus the number of workers who are unemployed at the end of the month but held a job at some time during the month, u_{t+1}^s. Rearrange to get

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

Comparing this with equation (4) gives

$$F_t = \frac{\sum_{i=1}^{u_t} F_t^i}{u_t},$$

so F_t is the mean job finding probability among workers who are unemployed at date t.

If unemployed workers were homogeneous, there would be other valid methods of constructing the job finding probability. Mean unemployment duration in month $t+1$, d_{t+1}, would be 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,

$$d_{t+1} = \frac{(d_t+1)(1-D_t)u_t + \big(u_{t+1} - (1-D_t)u_t\big)}{u_{t+1}}, \tag{9}$$

where D_t is the job finding probability for a worker who is unemployed in month t and 'D' indicates that this measure of the job finding probability is constructed using mean unemployment *duration* data. There are $(1-D_t)u_t$ 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 $u_{t+1}-(1-D_t)u_t$ newly unemployed workers in month $t+1$, each of whom has an unemployment duration of one month. This equation can be solved for the job finding probability as a function of the current and future mean unemployment duration and the number of unemployed workers,

$$D_t = 1 - \frac{(d_{t+1}-1)u_{t+1}}{d_t u_t}. \tag{10}$$

In steady state, $u_t = u_{t+1}$ and $d_t = d_{t+1}$, so equation (10) reduces to $D = 1/d$, a familiar relationship for a variable with a constant arrival rate.

Heterogeneity throws this calculation off. 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 F_t^i. By definition, mean unemployment duration in month t is $d_t \equiv \frac{1}{u_t}\sum_{i=1}^{u_t} d_t^i$. Generalizing equation (9) to allow for heterogeneous workers, we find that mean unemployment duration in month $t+1$ will be

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

or equivalently,

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

Comparing this with equation (10) yields $D_t = \frac{\sum_{i=1}^{u_t} d_t^i F_t^i}{\sum_{i=1}^{u_t} d_t^i}$, a weighted average of the individual job finding probabilities F_t^i, where the weight accorded to individual i is her unemployment duration d_t^i. Compared to the mean job finding probability F_t, this measure over-weights the long term unemployed. Since in practice the job finding probability falls with unemployment duration, one would expect that D_t to be smaller than F_t.

Hall (2005a) proposes a third measure of the job finding probability. Let u_t^m denote the number of medium term unemployed workers, defined because of data limitations as workers who have experienced 5 to 14 weeks (1 to 2 months) of unemployment. This is equal to the number of short term unemployed in previous months who have failed to find a job:

$$u_{t+1}^m = \big(u_t^s + u_{t-1}^s(1-M_{t-1})\big)(1-M_t). \tag{11}$$

This is a first order difference equation for M, where 'M' is a mnemonic for *medium* term unemployment. With a reasonable initial guess, e.g. that M_t, u_t^s, and u_t^m were constant before 1948, one can solve this equation forward for M. If all unemployed workers have the same job finding probability at every point in time, this will uncover that probability. But if workers are heterogeneous, this measure captures only the job finding probability of the short term unemployed and hence is likely to yield an estimate that exceeds the mean job finding probability F_t.

Figure 8 examines these predictions empirically using publicly available BLS time series constructed from the CPS. I use standard time series for the number of employed and unemployed workers; multiply mean unemployment duration, published in terms of weeks, by $\frac{12}{52}$ to convert it to monthly terms; and adjust short- and medium term unemployment for the effects of the CPS redesign, as discussed in Appendix A. Even though each series is constructed from different moments of the unemployment duration distribution, their cyclical behavior is similar and their levels line up as predicted. The mean value from 1948 to 2004 of F_t is 46 percent, in between the corresponding means for M_t (55 percent) and D_t (35 percent). I conclude that while heterogeneity complicates the definition of the 'the' job finding rate, it does not alter the conclusion that the job finding rate is procyclical.

4.2 The Heterogeneity Hypothesis

There are two distinct explanations for why the job finding probability F_t is procyclical: either the job finding probability declines for each worker or the unemployment pool shifts disproportionately towards workers with a low job finding probability. Darby, Haltiwanger, and Plant (1985) and (1986) advance the second possibility in their exploration of the cyclical 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 job finding probability will fall even if F_t^i does not change for any particular worker. Following Baker (1992), I refer to this as the 'heterogeneity hypothesis'.[14]

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 job finding probability F_t^i and use repeated spells of unemployment for particular individuals in order to check how her job finding probability depends on aggregate labor market conditions. Unfortunately, I am unaware of a large reliable representative data set for the United States that contains information on repeated unemployment spells.

Instead, I assume that workers can be divided into J different groups, in-

[14]Dynarski and Sheffrin (1990) and Baker (1992) show that unemployment duration is strongly countercyclical, and so the job finding probability 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, p. 320) to conclude that "the heterogeneity explanation of aggregate variation sheds little light on the nature of unemployment dynamics." 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.

dexed 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 $F_{t,j}$ be the job finding probability of those workers, computed using a type-dependent analog of equation (4). If Darby, Haltiwanger, and Plant's heterogeneity hypothesis is correct, fluctuations in the job finding probability, $F_t = \frac{\sum_j u_{t,j} F_{t,j}}{\sum_j u_{t,j}}$, are due primarily to changes in the shares $u_{t,j}$ rather than in the type-specific job finding probability $F_{t,j}$.

To quantify this, one can construct two hypothetical measures. Let F_t^{comp} denote the change in the job finding probability due to changes in the composition of the work force and F_t^{real} denote the "real" changes due to changes in the job finding probability for each type of worker:

$$F_t^{\text{comp}} \equiv \frac{\sum_j u_{t,j} \bar{F}_j}{\sum_j u_{t,j}} \qquad \text{and} \qquad F_t^{\text{real}} \equiv \frac{\sum_j \bar{u}_j F_{t,j}}{\sum_j \bar{u}_j},$$

where $\bar{F}_j \equiv \frac{1}{T} \sum_{t=1}^{T} F_{t,j}$ is the time-averaged job finding probability for type j workers and $\bar{u}_j \equiv \frac{1}{T} \sum_{t=1}^{T} u_{t,j}$ is the time-averaged number of unemployed type j workers. If the heterogeneity hypothesis is correct, F_t^{comp} should be strongly procyclical and F_t^{real} should be acyclical. Note that in order to generate large fluctuations in F_t^{comp}, there must be large differences in the average job finding probability of groups with substantially different cyclical fluctuations in their unemployment rates. If average job finding probabilities are too similar, composition effects will not generate substantial fluctuations in the aggregate job finding probability. If the composition of the unemployed population is not sufficiently cyclical, the weights will not change.

I construct measures of the number of short term unemployed workers and total unemployed workers in different demographic groups from the public-use monthly CPS microdata from January 1976 to January 2005.[15] I use these to measure the type-specific job finding probabilities $F_{t,j}$. I consider seven different dimensions of heterogeneity: seven age groups (16–19, 20–24, 25–34, 35–44, 45–54, 55–64, and 65 and over), sex, race (white or nonwhite), four marital status categories (spouse present, spouse absent or separated, widowed or divorced, never married), five reasons for unemployment (job loser on layoff, other job loser, job leaver, re-entrant, and new-entrant), nine census regions, and five

[15]Following Appendix A, I use only the incoming rotation groups after 1994.

education categories (high school dropouts, high school diploma, some college, bachelor's degree, some postgraduate education, only for workers age 25 and over). I analyze each dimension of heterogeneity in isolation.

The best case for the heterogeneity hypothesis is made by looking at changes in the fraction of workers reporting different reasons for unemployment, the focus of Figure 9. The top panel shows that in an average month between 1976 and 2004, a job loser not on layoff found a job with 0.31 probability, much lower than the probability for all other unemployed workers, which averaged 0.48. The bottom panel shows the share of job losers not on layoff in the unemployed population. The correlation between this share and the job finding rate for this group is -0.72. This pattern has the potential to generate fluctuations in the composition component of the job finding probability. In fact, this measure of F_t^{comp} averaged 41.3 percent in 1992 but rose to 43.6 percent in 2000 before falling back to 41.4 percent in 2003. But although these changes are noticeable and systematic, they explain little of the overall change in the job finding probability. By comparison, F_t^{real} rose from 36.5 percent in 1992 to 50.0 percent in 2000 and fell to 37.1 percent in 2003.

Figure 10 shows my measure of the "real" changes in the job finding probability F_t^{real} (solid lines) and compositional changes F_t^{comp} (dashed lines) for the seven different dimensions. Each figure shows that virtually all of the change in the job finding probability is "real." I conclude that changes in the composition of the unemployed population explain little of the overall fluctuations in the job finding probability.[16]

5 The Conventional Wisdom

This section serves two purposes: first, to describe the conventional wisdom on the cyclicality of the job finding and the separation probabilities; and second, to explain the consequences of the conventional wisdom for the development of macroeconomic models of the labor market.

[16]Changes in the age distribution also lead to some variation in the job finding probability, particularly at low frequencies. 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.'

5.1 Review of the Existing Evidence

The facts that I describe in this paper are in direct opposition to the conventional wisdom.[17] From their analysis of gross worker and job flows, Blanchard and Diamond (1990, p. 87) conclude that "The amplitude of fluctuations in the flow out of employment is larger than that of the flow into employment. This, in turn, implies a much larger amplitude of the underlying fluctuations in job destruction than of job creation." In their 1996 book, Davis, Haltiwanger, and Schuh, building on research by Davis and Haltiwanger (1990) and (1992), conclude that evidence from the United States manufacturing sector indicates that "job destruction rises dramatically during recessions, whereas job creation initially declines by a relatively modest amount." (Davis, Haltiwanger, and Schuh 1996, p. 31) The conventional wisdom based on this type of evidence is eloquently summarized by the title of Darby, Haltiwanger, and Plant (1986): "The Ins and Outs of Unemployment: The Ins Win."

Figure 11 shows Davis and Haltiwanger's quarterly data from 1972 to 1993, with job creation defined as the net increase in employment at expanding business establishments and job destruction as the net decrease in employment at contracting business establishments. Clearly job destruction is more volatile than job creation in this data set, rising during each of the major recessions in the 1970s and 1980s.[18] But there are at least three reasons why this does not say much about the cyclicality of the job finding and separation probabilities.

First, firms can destroy jobs either by firing workers or by not hiring to replace workers who leave. The former represents an increase in separations while the latter leads to a decrease in the job finding probability. One way to distinguish these alternatives is to look at establishments that shut down, which is clearly evidence of firms firing workers. Davis, Haltiwanger, and Schuh (1996, p. 34) conclude that "shutdowns do not account for an unusually large fraction of job destruction during recessions." This means that spikes in job destruction are consistent with the view advanced in this paper that there are only small

[17]Sider (1982) studies the cyclicality of unemployment incidence and duration. If workers are homogeneous and the economy is in steady state, unemployment incidence is equivalent to the separation rate and unemployment duration is the inverse of the job finding probability. He concludes that "changes in duration play a very important role in explaining ... fluctuations and trends in total unemployment." (Sider 1982, p. 461) This paper therefore argues for a return to this older wisdom.

[18]In a recent working paper, Davis, Faberman, and Haltiwanger (2005) construct a measure of job creation and job destruction back to 1947 (see their Figure 5). Although job destruction is more volatile than job creation in the 1960s, curiously they find that job creation and destruction were equally volatile in the 1950s.

increases in the separation probability of employed workers during downturns. Most contractions in employment are achieved by firms choosing to hire fewer workers, reducing workers' job finding probability.

Second, Davis and Haltiwanger focus exclusively on manufacturing establishments, a shrinking portion of aggregate employment. Foote (1998) uses Michigan data to show that job destruction is more volatile than job creation only in the manufacturing sector and argues that Davis and Haltiwanger's measures are biased by underlying trend employment growth. A new BLS survey, Business Employment Dynamics (BED), extends the Davis-Haltiwanger methodology to cover the entire labor market and provides some confirmation for Foote's theory. Figure 12 indicates that there was a brief spike in job destruction during the 2001 recession, but this was quickly reversed. Job creation fell immediately and has subsequently remained somewhat lower than normal.[19]

Third, Boeri (1996) compares the cyclicality of job creation and job destruction in several countries, showing that job destruction is more volatile than job creation only in the United States. He argues that this is an artefact of how the data are measured. Davis and Haltiwanger only measure job creation and destruction in establishments with more than five employees, but Boeri argues that small firms account for a large share of the cyclical volatility in job creation.

There are also shortcomings in the existing literature on gross worker flows, starting with its failure to address time aggregation. To my knowledge, none of the previous research using matched CPS data to measure gross worker flows between employment, unemployment, and inactivity has accounted for the fact that a decrease in the job finding probability indirectly raises the measured transition rate from employment to unemployment.

Another distinction between this paper and much of the gross flows literature is that while I measure the *probability* that an unemployed worker finds a job or an employed worker separates, Abowd and Zellner (1985), Poterba and Summers (1986), Blanchard and Diamond (1990), and much subsequent research has measured the *number* of workers who switch employment status in a given month. In fact, even after accounting for time aggregation, the decline in the job finding probability almost exactly offsets the increase in the number of unemployed workers at business cycle frequencies, so the number of unemployed workers who find a job in a month shows little cyclicality.

I focus here on the job finding probability because the notion of how difficult

[19]On the other hand, Faberman (2004) extends the BED survey back to 1990 and argues that job destruction was more volatile than job creation in the 1991 recession.

it is for an unemployed worker to find a job is a key input into models of job search such as those described in Pissarides (2000). For example, models of job search based on Pissarides's (1985) matching function predict that the job finding probability should depend directly on the vacancy-unemployment ratio via the matching function. The vacancy-unemployment ratio, in turn, depends only on exogenous variables. I am unaware of any coherent theory which predicts that the number of workers finding a job should depend only on exogenous variables.

5.2 Implications for Theoretical Models

The belief that separations drive unemployment fluctuations has dominated the recent development of macroeconomic models of the labor market. Mortensen and Pissarides (1994) extend Pissarides's (1985) model of an endogenous job finding probability to allow for idiosyncratic productivity shocks. Under reasonable conditions, an adverse aggregate shock raises the idiosyncratic threshold for maintaining an employment relationship, leading to the termination of many job matches. As a result, the model predicts that the time series of separations should be significantly more volatile than that of the number of workers finding jobs. Nevertheless, Mortensen and Pissarides (1994, pp. 412–413) are cautious, noting that "although empirical evidence on the cyclical issue is inconclusive, these results are consistent with Davis and Haltiwanger's (1990, 199[2]) findings." Over time, this caution has been lost. For example, Cole and Rogerson (1999) accept Davis and Haltiwanger's job creation and job destruction facts at face value in their reduced-form analysis of the implications of the Mortensen and Pissarides (1994) model.

Caballero and Hammour's (1994) model of creative destruction shows that if firms face a linear adjustment cost in hiring, fluctuations in the job finding probability will account for all of employment fluctuations. But because this contradicts the Davis-Haltiwanger and Blanchard-Diamond evidence, Caballero and Hammour (1994, p. 1352) argue that there must be strong convexities in hiring costs, and so conclude that recessions are "times of 'cleansing,' when outdated or relatively unprofitable techniques and products are pruned out of the productive system. . . ."[20] Koenders and Rogerson (2004) reason similarly in their analysis of 'jobless recoveries' that employment reductions during reces-

[20]More recently, Caballero and Hammour (2005) have argued that job destruction falls after a recession so that "*cumulatively*, recessions result in *reduced* rather than *increased* restructuring."

sions are due to firms postponing organizational restructuring until the end of an expansion. The longer the expansion, the more jobs that must be destroyed during the subsequent reorganization, resulting in a jobless recovery after prolonged expansions. In particularly, their model counterfactually predicts a surge of separations during 1991 and 2001 recessions.

Hall (1995) builds on the Davis-Haltiwanger and Blanchard-Diamond evidence to argue that that spikes in separations can generate persistent employment fluctuations: "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." (Hall 1995, p. 221)[21] Following this logic, Pries (2004) develops a model in which workers go through numerous short-term jobs before returning to a long-term employment relationship. This results in a persistent rise in the separation probability and gradual decline in the unemployment rate after a recession. Ramey and Watson (1997) propose a model of the business cycle with two-sided asymmetric information in which a transitory adverse shock induces a persistent rise in separations. den Haan, Ramey, and Watson (2000) examine how fluctuations in the separation probability can propagate and amplify shocks in a real business cycle model augmented with search frictions in the style of Mortensen and Pissarides (1994).

6 Conclusion

This paper measures the job finding and separation probabilities in the United States from 1948 to 2004. Throughout the time period, the job finding probability is strongly procyclical, while the separation probability was weakly countercyclical until the mid 1980s and more recently has been acyclic through two downturns in the labor market. These findings sharply contradict the conventional wisdom that fluctuations in the separation probability (or in job destruction) are the key to understanding the business cycle.

Recent research has attempted to come to terms with these facts. Shimer

[21] But Hall has since recanted, 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).

(2005a) argues that the Pissarides (1985) framework is ideally suited for thinking about these issues, but concludes that this standard matching model is incapable of generating large fluctuations in the job finding probability in response to shocks of a plausible magnitude. Hall (2005b) suggests that if wages in new employment relationships do not respond to labor market conditions, perhaps because of a social norm, the model can match the business cycle facts. More recent papers have proposed particularly mechanisms that generate this sort of wage rigidity. Kennan (2004) considers asymmetric information, showing how high job finding probabilities may be driven by firms' desire to hire workers during periods when their information rents are particularly large. Menzio (2004) suggests that firms attempt to hide business cycle frequency fluctuations in productivity in order to avoid giving wage increases to their existing workforce. This keeps wages relatively constant in the face of large fluctuations in the job finding probability.

Other researchers have focused on relaxing different aspects of the matching model. Nagypál (2004) extends the model to allow for search both employed and by unemployed workers. She argues that an unemployed worker might be willing to take a job even if he knows that better ones are readily available, while an employed worker only takes a job if it is better than his existing opportunity. If turnover is costly, this means that firms will prefer to hire employed workers. Since such workers are relatively plentiful during books, firms create more job openings during booms, raising (employed and unemployed) workers' job finding probability. Hall (2004a) similarly argues that workers' self-screening may affect firms' recruiting costs. When jobs are plentiful, workers only apply for jobs that are a good match with their skills, so most job applicants are worth hiring. But when the job finding probability is low, workers apply for any job they learn about, even if they are a poor match. This raises firms' recruiting costs, which reinforces the original shock that reduced the job finding probability. Whether any of these models ultimately explains the observed fluctuations in the job finding probability remains an open question. But it is clear that explaining these fluctuations is now at the center of the research agenda at the frontiers of macro and labor economics.

Appendix

A Measurement of Short-Term Unemployment

To measure short term unemployment, I rely on workers' self-reported duration of an in-progress unemployment spell. Unfortunately, the CPS instrument was redesigned in January 1994, changing how the unemployment duration question was asked (Abraham and Shimer 2001).[22] Recall that the CPS is a rotating panel. Each household is in the CPS for four consecutive months (rotation groups 1 to 4), out for eight months, and then in again for four more months (rotation groups 5 to 8). This means that in any month, approximately three-quarters of the households in the survey were also interviewed in the previous month.

Until 1994, unemployed workers in all eight rotation groups were asked how long they had been unemployed. But since then, the CPS has not asked a worker who is unemployed in consecutive months the duration of her unemployment spell in the second month. Instead, the BLS calculates unemployment duration in the second month as the sum of unemployment duration in the first month plus the intervening number of weeks. Thus prior to 1994, the CPS measure of short term unemployment should capture the total number of unemployed workers who were employed at any point during the preceding month, while after the redesign, short term unemployment only captures workers who transition from employment at one survey date to unemployment at the next survey date.[23]

There is no theoretical reason to prefer one measure to the other; however, the method I use to measure the job finding and separation probability in Section 2 relies on the pre-1994 measure of short term unemployment. In any case, the goal of this paper is to obtain a consistent time series for the job finding probability. To obtain one, note that one would expect that the redesign of the CPS instrument would not affect measured unemployment duration in rotation groups 1 and 5, the 'incoming rotation groups', since these workers are always asked their unemployment duration, but would reduce the measured short term unemployment rate in the remaining six rotation groups.

To see this empirically, I measure short term unemployment using CPS

[22] See Polivka and Miller (1998) for a thorough analysis of the redesign of the CPS instrument.

[23] The post-1994 methodology also prevents respondents from erroneously reporting short unemployment duration month after month.

microdata from January 1976 to January 2004.[24] In an average month from January 1976 to December 1993, short term unemployment accounted for 41.6 percent of total unemployment in the full CPS and 41.7 percent in the incoming rotation groups, an insignificant difference. From January 1994 to January 2005, however, short term unemployment accounted for 37.9 percent of unemployment in the full sample but 44.2 percent in the incoming rotation groups, an economically and statistically significant difference. Put differently, the short term unemployment rate in the full CPS fell discontinuously in January 1994, while it remained roughly constant in the incoming rotation groups.

In this paper I use short term unemployment from the full sample from 1948 to 1993 and then use only the incoming rotation groups in the later period.[25] More precisely, I first use the CPS microdata to compute the fraction of short term unemployed workers in the incoming rotation groups in each month from 1976 to 2004. I seasonally adjust this series using the Census's X-12-ARIMA algorithm with an additive seasonal factor. I then replace the standard measure of short-term unemployment with the product of the number of unemployed workers in the full CPS sample and the short-term unemployment share from 1994 to 2004.[26] This eliminates the discontinuity associated with the redesign of the CPS.[27]

I use a similar method to construct medium term unemployment and mean unemployment duration in Section 4. The only drawback to these procedures is that the reduced sample makes these measures slightly noisier than those using the full sample, an issue that is discernible in many of the figures in this paper.

---

[24]<http://www.nber.org/data/cps_basic.html>

[25]In January 1994, all unemployed workers were asked their unemployment duration, the last month in which this occurred. I start my adjustment a month earlier than necessary, using only the incoming rotation groups on and after January 1994, to coincide with the date of the CPS redesign.

[26]I multiply the number of unemployed workers from the full sample by the unemployment share from the incoming rotation groups to avoid another issue with the CPS. From 1976 to 2004, the unemployment rate in the first rotation group averaged 0.4 percentage points more than in the full sample. See Solon (1986) for a detailed discussion of rotation group biases in the CPS.

[27]I have also tried multiplying the standard series for short-term unemployment by a constant, 1.1, after 1994. This delivers very similar results.

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.

Abraham, Katharine (1987): "Help-Wanted Advertising, Job Vacancies, and Unemployment," *Brookings Papers on Economic Activity*, 1, 207–243.

Abraham, Katharine, and Lawrence Katz (1986): "Cyclical Unemployment: Sectoral Shifts or Aggregate Disturbances?," *Journal of Political Economy*, 94(3), 507–522.

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. Russell Sage Foundation and Century Foundation.

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

Blanchard, Olivier, and Peter Diamond (1989): "The Beveridge Curve," *Brookings Papers on Economic Activity*, 1, 1–60.

——— (1990): "The Cyclical Behavior of the Gross Flows of U.S. Workers," *Brookings Papers on Economic Activity*, 2, 85–143.

Boeri, Tito (1996): "Is Job Turnover Countercyclical?," *Journal of Labor Economics*, 14(4), 603–625.

Caballero, Ricardo, and Mohamad Hammour (1994): "The Cleansing Effect of Recessions," *American Economic Review*, 84(5), 1350–1368.

——— (2005): "The Cost of Recessions Revisted: A Reverse-Liquidationist View," *Review of Economic Studies*, 79(2), 313–341.

Cole, Harold, and Richard Rogerson (1999): "Can the Mortensen-Pissarides Matching Model Match the Business-Cycle Facts?," *International Economic Review*, 40(4), 933–959.

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, and John Haltiwanger (1990): "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," in *NBER Macroeconomics Annual*, ed. by Olivier Blanchard, and Stanley Fischer, pp. 123–168. MIT Press.

——— (1992): "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, 107(3), 818–863.

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

Davis, Steven J., R. Jason Faberman, and John Haltiwanger (2005): "The Flow Approach to Labor Markets: New Data Sources Micro-Macro Links, and the Recent Downturns," Mimeo, May 8.

den Haan, Wouter, Garey Ramey, and Joel Watson (2000): "Job Destruction and Propagation of Shocks," *American Economic Review*, 90(3), 482–498.

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.

Foote, Christopher (1998): "Trend Employment Growth and the Bunching of Job Creation and Destruction," *Quarterly Economic Review*, 113(3), 809–834.

Gottschalk, Peter, and Robert Moffitt (1999): "Changes in Job Instability and Insecurity Using Monthly Survey Data," *Journal of Labor Economics*, 17(4 part 2), S91–S126.

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

——— (2004a): "The Amplification of Unemployment Fluctuations through Self-Selection," Stanford Mimeo, December 21, 2004.

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

——— (2005a): "Employment Efficiency and Sticky Wages: Evidence from Flows in the Labor Market," *Review of Economics and Statistics*, forthcoming.

——— (2005b): "Employment Fluctuations with Equilibrium Wage Stickiness," *American Economic Review*, 95(1), 50–65.

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

Koenders, Kathryn, and Richard Rogerson (2004): "Organizational Dynamics Over the Business Cycle: A View on Jobless Recoveries," Arizona State mimeo, prepared for the Federal Reserve Bank of St. Louis's 29th annual Economic Policy Conference.

Menzio, Guido (2004): "High Frequency Wage Rigidity," Northwestern Mimeo, December 22, 2004.

Mortensen, Dale, and Christopher Pissarides (1994): "Job Creation and Job Destruction in the Theory of Unemployment," *Review of Economic Studies*, 61(3), 397–415.

Nagypál, Éva (2004): "Amplification of Productivity Shocks: Why Vacancies Don't Like to Hire the Unemployed?," Northwestern University Mimeo.

Pissarides, Christopher (1985): "Short-Run Equlibrium Dynamics of Unemployment, Vacancies, and Real Wages," *American Economic Review*, 75(4), 676–690.

——— (2000): *Equilibrium Unemployment Theory*. MIT Press, Cambridge, MA, second edn.

Polivka, Anne, and Stephen Miller (1998): "The CPS After the Redesign," in *Labor Statistics Measurement Issues*, ed. by John Haltiwanger, Marilyn Manser, and Robert Topel, pp. 249–289. University of Chicago Press, Chicago.

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

Pries, Michael (2004): "Persistence of Employment Fluctuations: A Model of Recurring Job Loss," *Review of Economic Studies*, 1(71), 193–215.

Ramey, Garey, and Joel Watson (1997): "Contractual Fragility, Job Destruction, and Business Cycles," *Quarterly Economic Review*, 112(3), 873–911.

Shimer, Robert (2005a): "The Cyclical Behavior of Equilibrium Unemployment and Vacancies," *American Economic Review*, 95(1), 25–49.

——— (2005b): "The Cyclicality of Hires, Separations, and Job-to-Job Transitions," University of Chicago Mimeo, prepared for the Federal Reserve Bank of St. Louis's 29th annual Economic Policy Conference.

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

Solon, Gary (1986): "Effects of Rotation Group Bias on Estimation of Unemployment," *Journal of Business and Economic Statistics*, 4(1), 105–109.

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

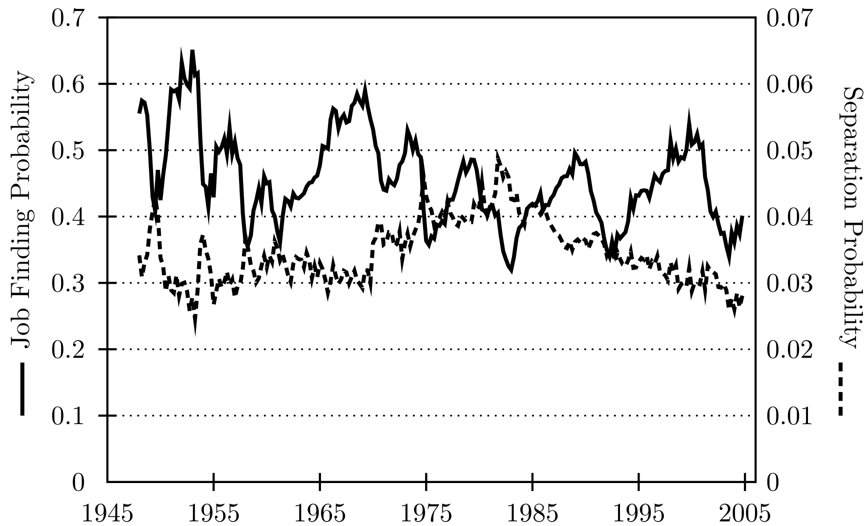

Figure 1: Job Finding and Separation Probabilities, 1948Q1–2004Q4, quarterly average of monthly data. The job finding probability is constructed from unemployment and short term unemployment according to equation (4). The separation probability is constructed from employment, unemployment, and the job finding probability according to equation (5). Employment, unemployment, and short term unemployment data are constructed by the BLS from the CPS and seasonally adjusted. Short term unemployment data are adjusted for the 1994 CPS redesign as described in Appendix A.

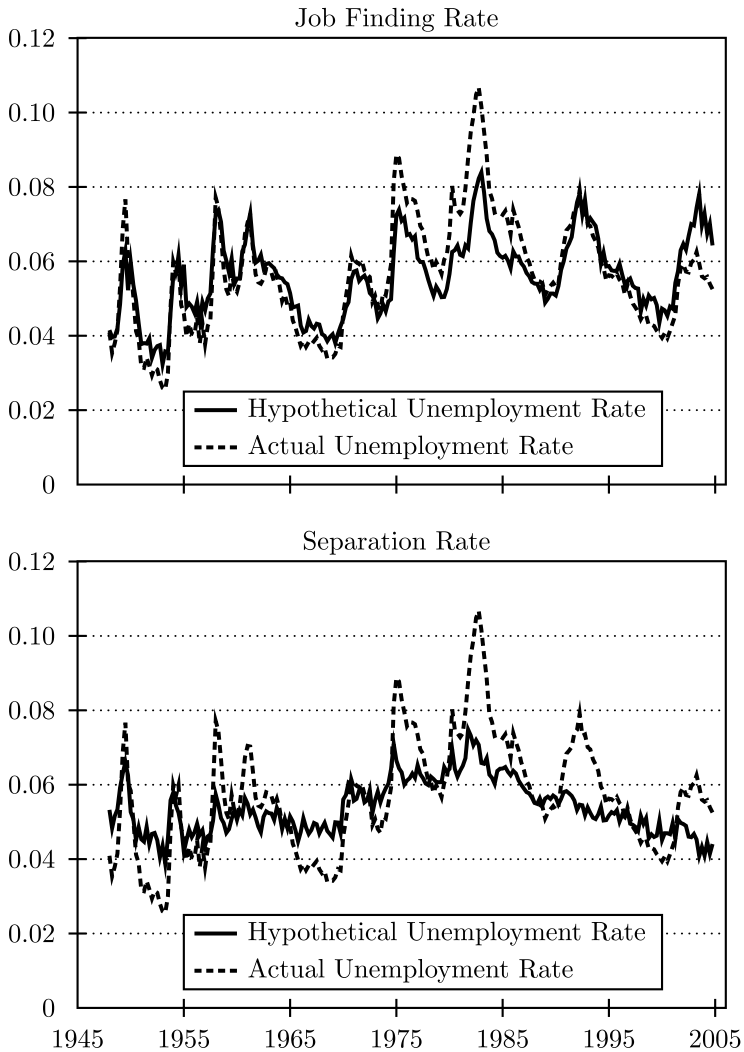

Figure 2: Contribution of Fluctuations in the Job Finding and Separation Rates to Fluctuations in the Unemployment Rate, 1948Q1–2004Q4, quarterly average of monthly data. The job finding rate f_t is constructed from unemployment and short term unemployment according to equation (4). The separation rate s_t is constructed from employment, unemployment, and the job finding rate according to equation (5). The top panel shows the hypothetical unemployment rate if there were only fluctuations in the job finding rate, $\bar{s}/(\bar{s}+f_t)$, and the bottom panel shows the corresponding unemployment rate with only fluctuations in the separation rate, $s_t/(s_t+\bar{f})$. Both panels show the actual unemployment rate for comparison. Employment, unemployment, and short term unemployment data are constructed by the BLS from the CPS and seasonally adjusted. Short term unemployment data are adjusted for the 1994 CPS redesign as described in Appendix A.

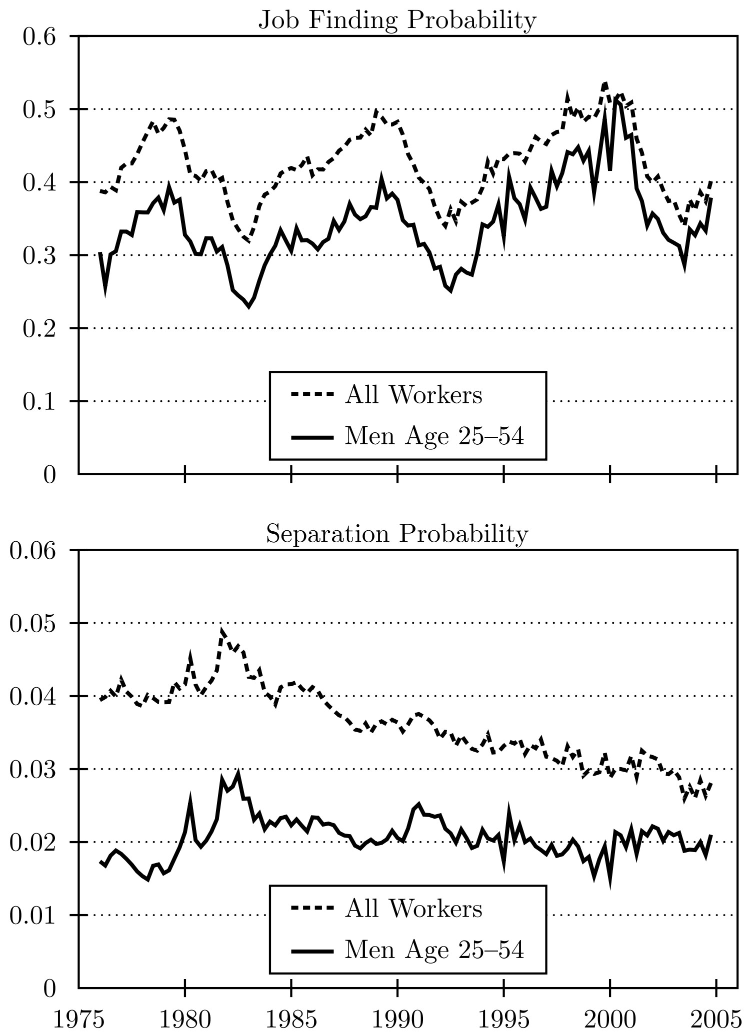

Figure 3: Job Finding and Separation Probabilities, Prime Age Men, 1976Q1–2004Q4, quarterly average of monthly data. The job finding probability is constructed from unemployment and short term unemployment according to equation (4). The separation probability is constructed from employment, unemployment, and the job finding probability according to equation (5). Employment, unemployment, and short term unemployment data are constructed by the BLS from the CPS and seasonally adjusted. Short term unemployment data are adjusted for the 1994 CPS redesign as described in Appendix A.

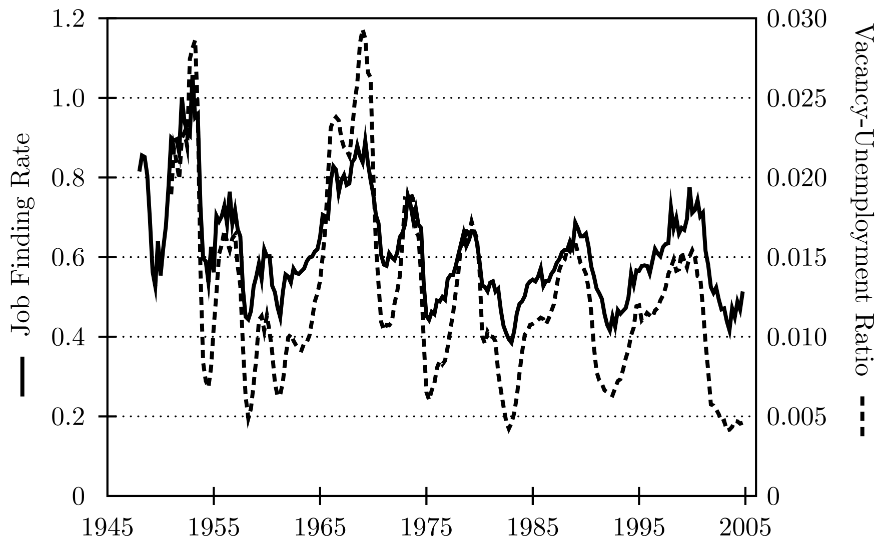

Figure 4: Job Finding Rate and Vacancy-Unemployment Ratio, 1948Q1–2004Q4, quarterly average of monthly data. The job finding rate f_t is constructed from unemployment and short term unemployment according to equation (4). The vacancy-unemployment ratio is the ratio of the Help Wanted Advertising Index to unemployment, measured in index units per thousand workers. Unemployment, and short term unemployment data are constructed by the BLS from the CPS and seasonally adjusted. Short term unemployment data are adjusted for the 1994 CPS redesign as described in Appendix A. The Help Wanted Advertising Index is constructed by the Conference Board and seasonally adjusted.

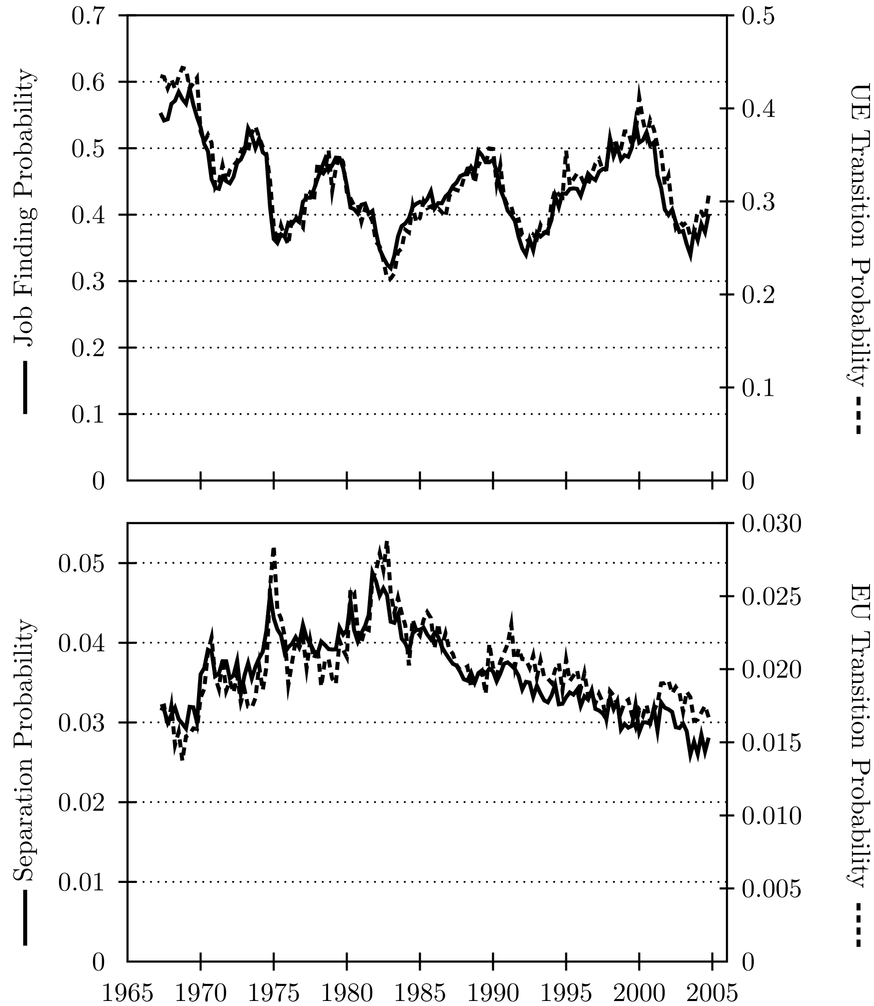

Figure 5: Alternative Measures of the Job Finding and Separation Probabilities, 1967Q2–2004Q4, quarterly average of monthly data. The job finding probability is constructed from unemployment and short term unemployment according to equation (4). The separation probability is constructed from employment, unemployment, and the job finding probability according to equation (5). Employment, unemployment, and short term unemployment data are constructed by the BLS from the CPS and seasonally adjusted. The gross flows are computed from matched CPS microdata files by Joe Ritter (1967Q2–1975Q4) and by the author (1976Q1–2004Q4), seasonally adjusted using a ratio to moving average, and then used to infer the transition probabilities following the procedure described in Section 3.1. Short term unemployment data are adjusted for the 1994 CPS redesign as described in Appendix A.

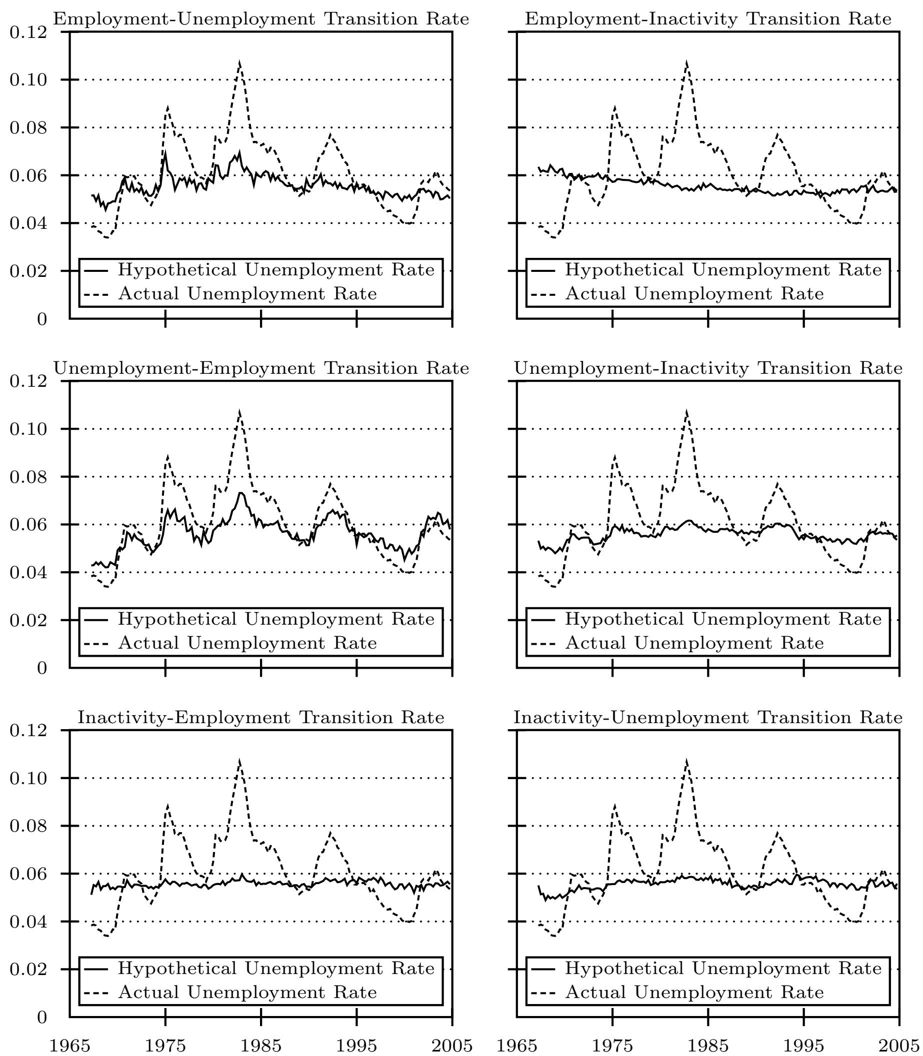

Figure 6: Contributions of Fluctuations in the Instantaneous Transition Rates to Fluctuations in the Unemployment Rate, 1967Q2–2004Q4, quarterly average of monthly data. The gross flows are computed from matched CPS microdata files by Joe Ritter (1967Q3–1975Q4) and by the author (1976Q1–2004Q4), seasonally adjusted using a ratio to moving average, and then used to infer the transition rates following the procedure described in Section 3.1. The contributions to the unemployment rate are inferred as in equation (8). Each panel shows the actual unemployment rate for comparison.

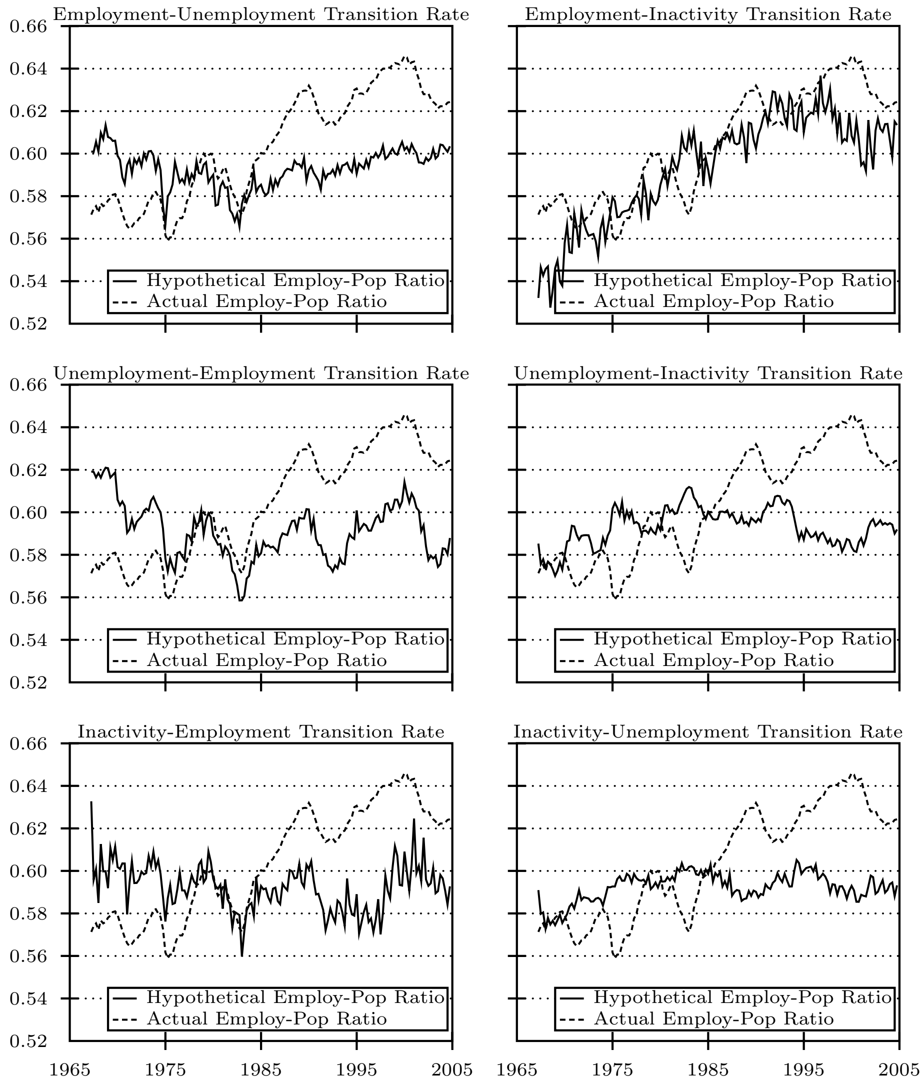

Figure 7: Contributions of Fluctuations in the Instantaneous Transition Rates to Fluctuations in the employment population Ratio, 1967Q2–2004Q4, quarterly average of monthly data. The gross flows are computed from matched CPS microdata files by Joe Ritter (1967Q3–1975Q4) and by the author (1976Q1–2004Q4), seasonally adjusted using a ratio to moving average, and then used to infer the transition rates following the procedure described in Section 3.1. The contributions to the employment-population ratio are inferred as in equation (8). Each panel shows the actual employment-population for comparison.

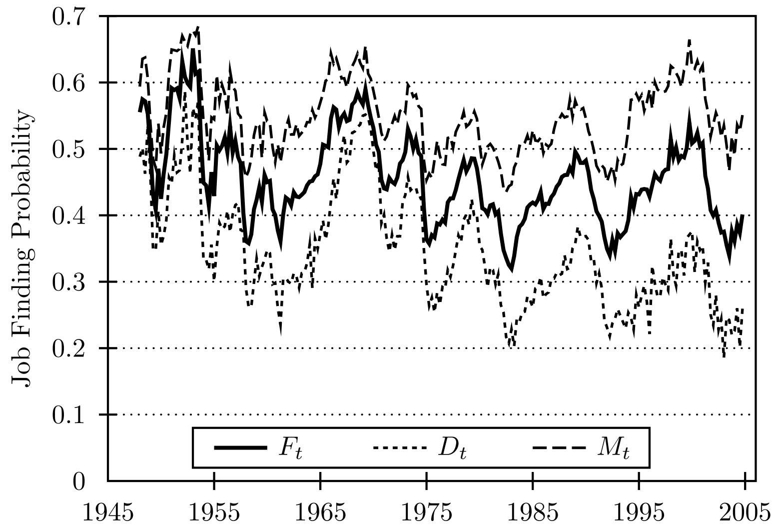

Figure 8: Three Measures of the Job Finding Probability, United States, 1948Q1–2004Q1, quarterly average of monthly data. The job finding probability F_t is constructed from unemployment and short term unemployment according to equation (4). The alternative measures D_t and M_t are constructed from mean unemployment duration data (equation 10) and short and medium term unemployment data (equation 11), respectively. All data are constructed by the BLS and seasonally adjusted. Mean unemployment duration and short and medium term unemployment data are adjusted for the 1994 CPS redesign as described in Appendix A.

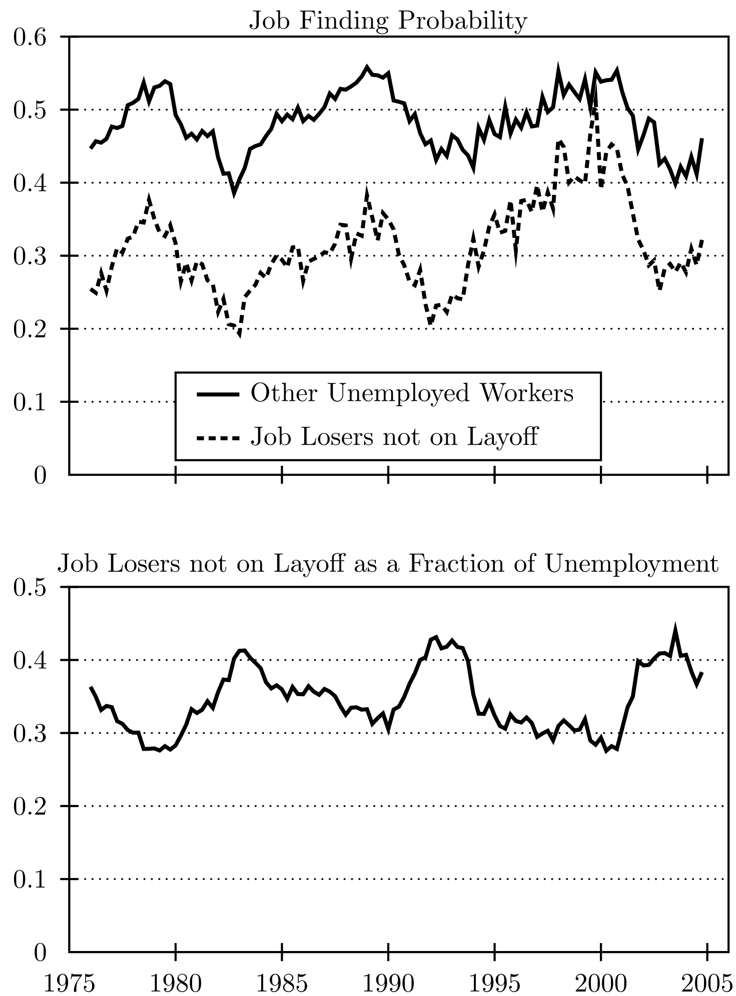

Figure 9: Fluctuations in the Job Finding Probability and Unemployment Share of Job Losers not on Layoff, United States, 1976Q1–2004Q4, quarterly average of monthly data. The underlying data are constructed from the monthly CPS, seasonally adjusted and adjusted for the 1994 CPS redesign as described in Appendix A, and averaged within quarters.

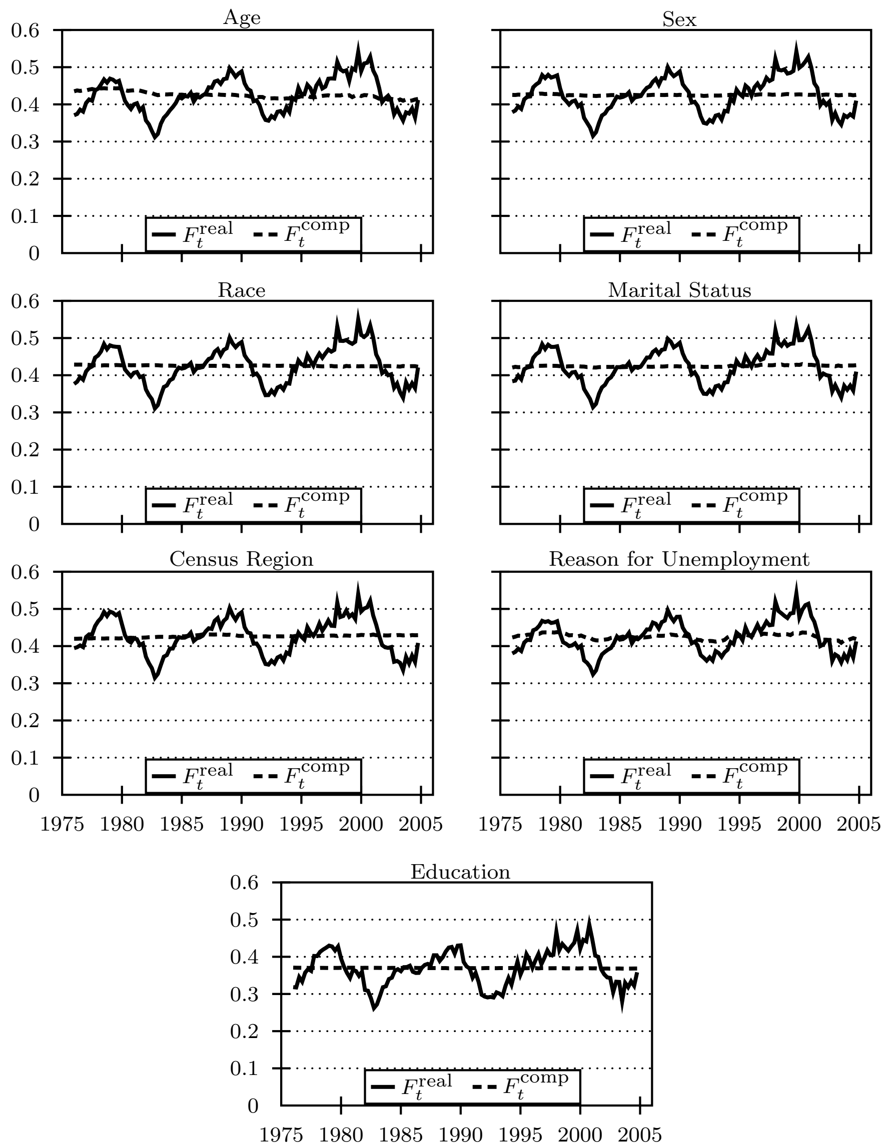

Figure 10: Seven measures of the 'compositional' and 'real' component of changes in the job finding probability, F_t^{comp} and F_t^{real}, respectively, United States, 1976Q1–2004Q4, quarterly average of monthly data. Each figure uses different characteristics: age (7 groups), sex, race (white or nonwhite), marital status (married spouse present, spouse absent or separated, divorced or widowed, never married), census region (9 regions), reason for unemployment (job loser on layoff, other job loser, job leaver, re-entrant, or new entrant), and education (5 groups, age 25 and over). The underlying data are constructed from the monthly CPS, seasonally adjusted and adjusted for the 1994 CPS redesign as described in Appendix A, and averaged within quarters.

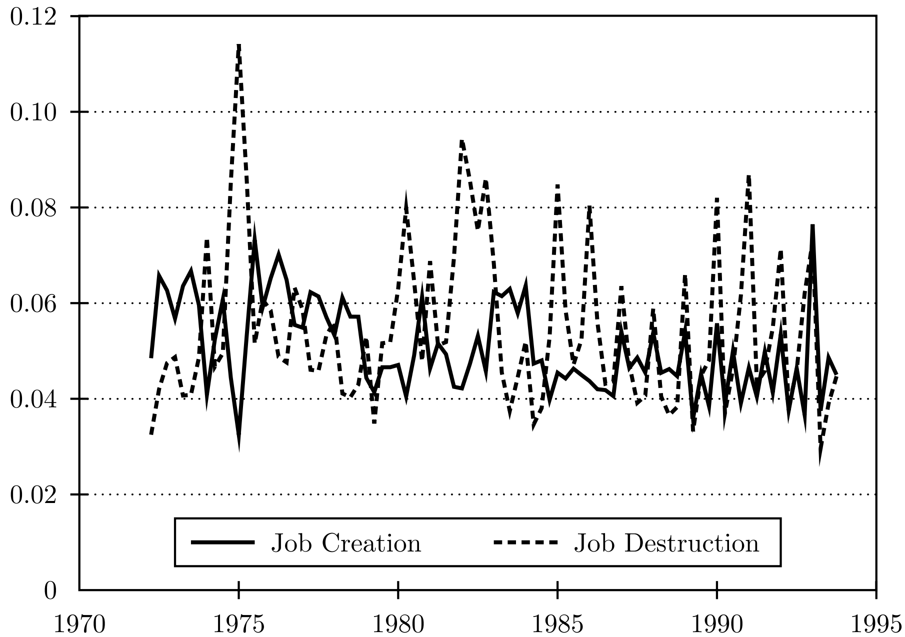

Figure 11: Job Creation and Job Destruction in Manufacturing, United States, 1972Q2–1993Q4. The data are constructed by Davis, Haltiwanger, and Schuh and are available from `http://www.bsos.umd.edu/econ/haltiwanger/download.htm`. They are seasonally adjusted.

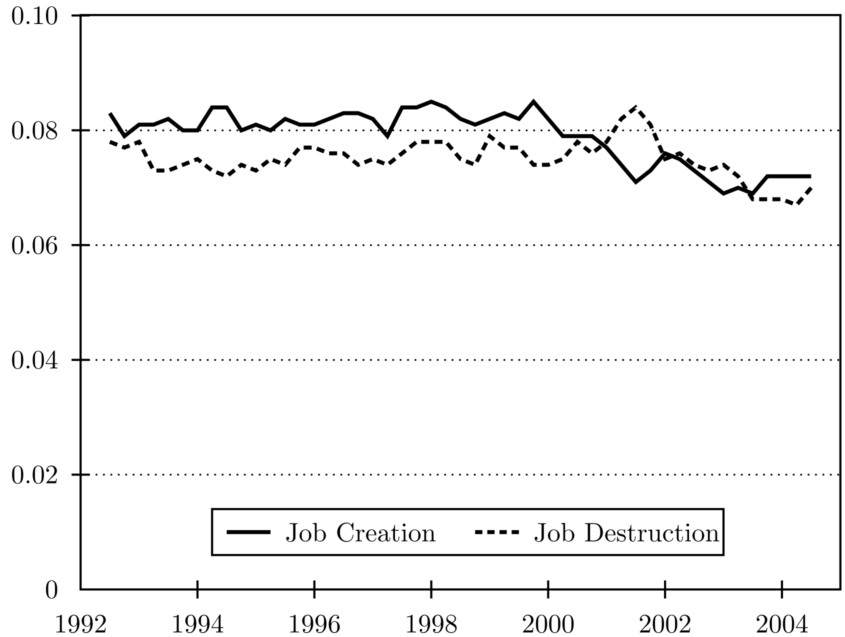

Figure 12: Job Destruction and Job Creation, United States, 1992Q3–2004Q3. The data are constructed by the BLS as part of the BED and are seasonally adjusted.