w o r k i n g
p a p e r

10 17R

The Ins and Outs of Unemployment in the Long Run: A New Estimate for the Natural Rate?

Murat Tasci

FEDERAL RESERVE BANK OF CLEVELAND

Working Paper 10-17R January 2011*

The Ins and Outs of Unemployment in the Long Run:
A New Estimate for the Natural Rate?

Murat Tasci

In this paper, we present a simple, reduced-form model of comovements in real activity and worker flows and use it to uncover the trend changes in these flows, which determine the trend in the unemployment rate. We argue that this trend rate has several key features that are reminiscent of a “natural rate.” We show that the natural rate, measured this way, has been relatively stable in the last decade, even after the most recent recession. This was due to two opposing trend changes: On the one hand, the trend in the job-finding rate, after being relatively stable for decades, declined by a significant margin after 2000, pushing trend unemployment up. But the trend in the separation rate has somewhat offset that effect, with a continued secular decline since the early 1980s. We also show that, contrary to the business-cycle movements of the unemployment rate, most of the low-frequency variation in the rate can be accounted for by changes in the trend of the separation rate, not the job-finding rate. The notable exception is during the last decade, when the trend changes in the flows that caused the opposing effects on the trend unemployment rate also implied a slower rate of worker reallocation in the U.S. economy. This slow rate of worker reallocation implies a much slower decline in the unemployment rate over the near term than would have been possible with a high degree of churning, which was a feature of U.S. labor markets before the past decade. We also show that the estimated trend for the unemployment rate is very robust to labor force movements and to the use of different filters, as long as one utilizes the information on the underlying worker flow rates.

Key words: Unemployment; Natural Rate; Job-finding Rate; Separation Rate; Labor Market Search.

JEL classification: E24; E32; J64.

*This paper was initially posted in October 2010.

Murat Tasci is a research economist at the Federal Reserve Bank of Cleveland. He can be reached at murat.tasci@clev.frb.org. The author would like to thank Ken Beauchemin, Charles Carlstrom, Tim Dunne, Martin Fukac, Mark Schweitzer, and seminar participants at the Federal Reserve Bank of Cleveland and the Federal Reserve Bank of Boston. He is grateful to Saeed Zaman for his able research assistance.

1 Introduction

This paper lays out a simple, reduced-form model that is useful for exploring the comovement of real economic activity and flows into and out of unemployment. We argue that the large body of literature on the search theory of unemployment makes a compelling case that understanding these job-finding and job-separation rates is the key to understanding the long-run behavior of the unemployment rate.[1] We also show that one can construct an unemployment rate trend based on these flows, and that this rate can be interpreted as the rate of unemployment in the long run, with which the actual unemployment rate would converge if the trends in these flows persisted. The method essentially provides us with a time-varying trend estimate for the unemployment rate. We argue that this trend rate has several key features that are reminiscent of a "natural rate"; hence, we use the terms "natural rate" and "unemployment trend" interchangeably from here onward.

We show that, measured this way, the natural rate has been relatively stable over the past decade, even after the most recent recession. Underlying this stability are two offsetting trends in the flows; the first is the trend in the job-finding rate, which, after being relatively stable for decades, declined significantly since 2000, pushing trend unemployment up. The second is the trend in the separation rate, which has partially offset the effect of the job-finding trend by declining secularly since the early 1980s. We also show that, unlike business-cycle frequency movements of the unemployment rate, a significant fraction of the low-frequency variation in the rate can be explained by changes in the trend of the separation rate rather than the trend of the job-finding rate, especially before 1985. The exception was during the last decade, when the changes in the flows that caused opposing effects on the trend unemployment rate also implied a slower rate of worker reallocation for the U.S. economy. Furthermore, we show via a set of numerical exercises that this slow worker reallocation has important implications for the adjustment process of the unemployment rate in the near term. In particular, our model suggests that because the worker reallocation rate (the sum of the separation and job-finding rates) has slowed, unemployment will decline substantially less in the near term. We also provide a quantitative example of the potential impact of "weaker" output growth during the

[1] For a survey of the labor market search literature, see Mortensen and Pissarides (1999). Pissarides (2000) provides a nice textbook treatment of the subject.

current recovery on this adjustment process. The experiments show the potential usefulness of the model we propose. Moreover, we show that the model has several more desirable features such as precision of the estimates and minor retrospective revisions it implies with additional data. Finally, we find that our results are robust to allowing for labor force fluctuations and flows into and out of inactivity.

The next section discusses the related literature, especially that on the natural rate, and where our empirical approach fits in. Section 3 presents our simple, reduced-form model, which describes the comovement of real GDP and worker flows. This section also includes our description of the data, particularly how we construct worker flow rates and conduct our estimation. Section 4 presents our estimation results and unemployment rate decompositions due to each flow rate, both at the business cycle frequency and over the long run. Section 5 includes our discussion on the Great Recession in light of our model where we address whether the last recession changed the trend of the unemployment rate, and how significant the effects of slow worker reallocation and weak output growth will be on the dynamics of the unemployment rate in the near term. Section 6 compares our methodology to purely statistical filters and a reasonable alternative that is similar in methodology to ours, and shows that our use of worker flow rates rather than the observed unemployment rate is behind our robust result. This section also argues that our model has important desirable statistical features such as precision of the estimates and minor retrospective revisions it requires with additional data. This section also presents some evidence that our results are robust to changes in the labor force. The last section concludes.

2 Related Literature

Our estimate for the long-run trend of the unemployment rate, as we noted earlier, is reminiscent of the natural rate of unemployment. The concept dates back at least to Friedman (1968) and Phelps (1968).[2] It is probably one of the most frequently used, yet most vaguely defined, concepts utilized by macroeconomists. Rogerson criticizes this in his review essay, concluding

[2] For a good discussion on the topic, one can look at a set of papers in two volumes: *Journal of Economic Perspectives (Winter 1997)* and the *American Economic Review, Papers and Proceedings (May 1988)*, as well as a survey by Johnson and Layard (1986).

that "economics would benefit from being deprived of these concepts" and that "we have reached a point where our theories of unemployment are ahead of language" (Rogerson 1997, 74–5). We can trace the origin of the "natural rate of unemployment" concept to Milton Friedman. In his presidential address to the members of the American Economic Association (1968, 8), Friedman spelled out this concept. He did not provide a clear, well-defined characterization of this concept, but rather described some features that it should have:

> The "natural rate of unemployment"... is the level that would be ground out by the Walrasian system of general equilibrium equations, provided there is imbedded in them the actual structural characteristics of the labor and commodity markets, including market imperfections, stochastic variability in demands and supplies, the cost of gathering information about job vacancies and labor availabilities, the cost of mobility, and so on.

Although he further qualified this concept elsewhere, it turned out to be vague enough to make it hard for economists to agree on a clear way to map the concept into a quantitative measure (see Rogerson (1997)). Some economists developed this concept into another concept, the NAIRU (non-accelerating inflation rate of unemployment). This concept assumes an inherent trade off between inflation and the unemployment rate in the sense that when the unemployment rate is above the NAIRU because of slack in the labor market, there will be downward pressure on prices and wages, and inflation will go down. Similarly, a lower measured unemployment rate relative to the NAIRU is assumed to put upward pressure on prices and wages. However, if anything, Friedman (1968, 9) made it clear that he used the term "... 'natural' for the same reason that Wicksell did—to try and separate real forces from monetary forces." Thus, from this perspective, we do not consider the NAIRU concept useful, and it will not be our focus here.[3]

[3] Nevertheless, we should note that the NAIRU has been the focus of a large body of literature, where it is sometimes used synonymously with the natural rate concept we have discussed; see, for example, Ball and Mankiw (2002). A substantial body of literature focuses on estimating the NAIRU, and some of it uses unobserved components methods similar to those employed here or a variant of the Phillips curve; see, for example, Staiger, Stock, and Watson (1997 and 2001), and King and Watson (1994). The usefulness of this concept for policy is a whole different topic; see, for example, Rogerson (1997), David Gordon (1988), Robert Gordon (1997), and Orphanides and Williams (2002), among others.

Another point Friedman emphasized in his address was that the natural rate itself might change over time due to market forces or economic policies. This is very intuitive. For instance, labor market policies such as high unemployment compensation, strict firing rules, and severance policies have been blamed for persistently high unemployment in Europe. It is conceivable that these policies resulted in a higher "natural" rate for Europe, thereby keeping the actual (measured) unemployment rate high during the past three decades as well (Blanchard (2006)).

In our attempt to measure this "natural" rate of unemployment, we follow this guidance and look for a rate that is moving at a relatively low frequency and could potentially change over time, albeit smoothly. For the purposes of this paper, we will call this the long-run trend of the unemployment rate or the natural rate, interchangeably. The unique aspect of our approach is that we estimate the natural rate by first isolating the underlying trends in the job-finding and job-separation rates. We then employ these rates to estimate the long-term trend in unemployment by using the fact that it can be expressed as the ratio of the separation rate to the overall reallocation rate. We think that this exercise gives us a useful empirical concept, which clearly maps into the theory of unemployment in many ways.

In principle, one can use a benchmark search model and estimate it structurally to back out this long-run trend from the model. However, we think that there are at least two reasons why we might do better by pursuing a useful empirical concept instead. First, this class of models is subject to well-known problems that manifest themselves as inability to match many key moments for the labor market variables, including those for unemployment itself. In particular, Hall (2005) and Shimer (2005) argue that standard models of labor market search require implausibly large shocks to generate substantial variation in key variables: unemployment, vacancies, and market tightness (the vacancy-to-unemployment ratio). This quantitative problem makes it harder to use this class of models for a measurement exercise like the one we have in mind here. Secondly, many of the low-frequency changes in the underlying flows represent low-frequency changes in the economic environment, such as labor market policies, demographic changes, and technological advances (in either production or matching technology); incorporating all of these potential driving forces into a parsimonious model would be fairly complicated. By imposing a low-frequency change in these flows, our simple, reduced form model allows for these potential channels to the extent that they affect unemployment flows. If inflow into unemployment turns out to be the main driving force that determines the long-run trend, that is, the "natural rate," as we find, then one can potentially focus on theoretical features in these models, which would manifest themselves as changes in inflows.[4] Hence, we believe that our approach here could also be useful for modelling unemployment in the long run.

Our reduced form empirical model and the estimation method we employ are closely related to the study of measuring the cyclical component of economic aggregates, as in Clark (1987, 1989) and Kim and Nelson (1999).[5] Our approach—identifying the trend of the unemployment rate over time via long-term trends of the underlying flows into and out of unemployment—is perhaps most closely related to Darby, Haltiwanger, and Plant (1985) and Barro (1988). Darby, Haltiwanger, and Plant (1985) look into the importance of heterogeneity in worker flows for unemployment persistence. Barro (1988) focuses on the same long-run equilibrium condition for unemployment that we focus on here, that is, the separation rate over the sum of the separation rate and the job-finding rate; he emphasizes how worker reallocation determines persistence in unemployment. In this paper, however, we try to tease out the cyclical variation in these flows from the trend changes, in order to estimate the unemployment rate trend. More recently, Dickens (2009) also proposes an empirical model that uses information from the Beveridge curve. Although he incorporates unemployment flows into the model, his main focus is to estimate a time-varying NAIRU.

3 A Simple Model

We are going to write down a simple, reduced form model that incorporates the comovement of flows into and out of unemployment into previous attempts at estimating the natural rate, such as Clark (1987, 1989) and Kim and Nelson (1999). The reduced form model assumes that real GDP has both a stochastic trend and a stationary cyclical component, where only real GDP is observed by the econometrician. We also assume that both flow rates, F_t and S_t, (job-finding and separation rate respectively) have a stochastic trend as well as a stationary component.

[4] See, for instance, Shimer (2007), Elsby, and Michaels and Solon (2009) for the cyclical contributions of flows to unemployment fluctuations in the US.

[5] The idea is similar to the one employed by Laubach and Williams (2003), where they estimate the unobserved natural rate of interest.

Furthermore, the stochastic trend follows a random walk, but the cyclical component in the flow rates depends on the cyclical component of real GDP. More specifically, let Y_t be log real GDP, $\bar{y}_t$ a stochastic trend component, and y_t the stationary cyclical component. Similarly, let F_t (S_t) be the quarterly job finding (separation) rate, $\bar{f}_t$ ($\bar{s}_t$) its stochastic trend component, and f_t (s_t) the stationary cyclical component. Then we consider the following unobserved components model:

$$Y_t = \bar{y}_t + y_t; \quad \bar{y}_t = g_{t-1} + \bar{y}_{t-1} + \varepsilon_t^{yn}; \quad g_t = g_{t-1} + \varepsilon_t^{g}; \quad y_t = \phi_1 y_{t-1} + \phi_2 y_{t-2} + \varepsilon_t^{yc} \tag{1}$$

$$F_t = \bar{f}_t + f_t; \quad \bar{f}_t = \bar{f}_{t-1} + \varepsilon_t^{fn}; \quad f_t = \rho_1 y_t + \rho_2 y_{t-1} + \rho_3 y_{t-2} + \varepsilon_t^{fc} \tag{2}$$

$$S_t = \bar{s}_t + s_t; \quad \bar{s}_t = \bar{s}_{t-1} + \varepsilon_t^{sn}; \; s_t = \theta_1 y_t + \theta_2 y_{t-1} + \theta_3 y_{t-2} + \varepsilon_t^{sc} \tag{3}$$

where g_t is a drift term in the stochastic trend component of output which is also a random walk, following Clark (1987). All the error terms, ε_t^{yn}, ε_t^{g}, ε_t^{yc}, ε_t^{fn}, ε_t^{fc}, ε_t^{sn}, ε_t^{sc}, are independent white-noise processes. There is nothing very controversial about (1), which governs the movement in real output. We impose a stochastic trend, which might be subject to occasional drifts, and a persistent but stationary cyclical component. What is more unconventional is the comovement in the rates of job finding and separations in (2) and (3). We argue that the low-frequency movements in the trends, $\bar{f}_t$ and $\bar{s}_t$, will capture the effects of institutions, demographics, tax structure, labor market rigidities, and the long-run matching efficiency of the labor markets, which will be more important in determining the steady state of unemployment, consistent with our arguments in the preceding section. The cyclical components, f_t and s_t, on the other hand, are moving in response to purely cyclical changes in output. One can easily legitimize this in a simple extension of the textbook search model with endogenous job destruction and shocks to aggregate productivity, as in Mortensen and Pissarides (1994). In this class of models, market tightness—hence the job-finding rate—increases during expansions and declines during recessions. Similarly, when aggregate productivity is temporarily low, there will be a surge of separations, resulting in higher unemployment, because some existing matches cease to be productive enough in the recession. Hence, the assumed relationship of (2) and (3) is in line with the predictions of the search theory of unemployment.

Recall that the trend of the unemployment rate, according to our definition, is pinned down by the stochastic trend components of the job-finding and separation rates. We can estimate our model and use a Kalman filter to back out the underlying trends in order to get an estimate of a time-varying trend. To start, we can write down the system of equations in (1)-(3), in the following state-space representation:

$$\begin{bmatrix} Y_t \\ F_t \\ S_t \end{bmatrix} = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & \rho_1 & \rho_2 & \rho_3 & 0 & 1 & 0 \\ 0 & \theta_1 & \theta_2 & \theta_3 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \bar{y}_t \\ y_t \\ y_{t-1} \\ y_{t-2} \\ g_t \\ \bar{f}_t \\ \bar{s}_t \end{bmatrix} + \begin{bmatrix} 0 \\ \varepsilon_t^{fc} \\ \varepsilon_t^{sc} \end{bmatrix} \tag{4}$$

$$\begin{bmatrix} \bar{y}_t \\ y_t \\ y_{t-1} \\ y_{t-2} \\ g_t \\ \bar{f}_t \\ \bar{s}_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & \phi_1 & \phi_2 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} \bar{y}_{t-1} \\ y_{t-1} \\ y_{t-2} \\ y_{t-3} \\ g_{t-1} \\ \bar{f}_{t-1} \\ \bar{s}_{t-1} \end{bmatrix} + \begin{bmatrix} \varepsilon_t^{yn} \\ \varepsilon_t^{yc} \\ 0 \\ 0 \\ \varepsilon_t^{g} \\ \varepsilon_t^{fn} \\ \varepsilon_t^{sn} \end{bmatrix} \tag{5}$$

where all error terms come from an i.i.d. normal distribution, with zero mean and variance σ_i such that $i = \{yn, g, yc, fn, fc, sn, sc\}$. Once we estimate this model using US data, we can back out our estimate of a time-varying unemployment rate trend by using the estimates of the unobserved trend components. In particular, $\bar{u}_t = \frac{\bar{s}_t}{\bar{s}_t + \bar{f}_t}$ will give us the desired rate of unemployment trend, that the trend in the flows will predict in the long-run. In principle, this methodology can also provide an estimate of the trend output, $\bar{y}_t$. However, two principal problems need to be tackled in this estimation strategy. First, we need data on job-finding and separation rates for the aggregate economy, which are not readily available. Second, the model, as spelled out in equations (4)-(5), is subject to an identification problem. Even though we have only three observables, we are attempting to estimate parameters for seven shocks. We explain in detail how we handle these problems in the following data and estimation subsections.

3.1 Data

Our measure of real output is calculated as quarterly gross domestic output in billions, from the Bureau of Economic Analysis (Department of Commerce) and spans the period 1948:Q1 through 2011:Q2.[6] Flow rates, on the other hand, are not readily available for the aggregate economy. However, recent research on the cyclical features of unemployment, led by Shimer (2005, 2007) and, more recently, by Elsby, Michaels, and Solon (2009) provides us with a simple method to measure these rates using Current Population Survey (CPS) data. The method infers continuous time hazard rates into and out of unemployment by using readily available short-term unemployment, aggregate unemployment, and labor force data. Here we briefly describe the method used to infer these rates, without getting too far into the tedious details. Our presentation will closely follow that of Elsby, Michaels, and Solon (2009).

Let u_t be the number of unemployed in month t of the CPS, u_t^s, the number who are unemployed less than five weeks in month t and l_t the size of the labor force in month t. At the heart of the measurement is a simple equation determining the evolution of unemployment over time in terms of flows into and out of unemployment:

$$\frac{du_t}{dt} = S_t(l_t - u_t) - F_t u_t \tag{6}$$

in terms of flows into and out of unemployment. Given this simple accounting equation, we start with a typical unemployed worker's probability of leaving unemployment. As Shimer (2007) and Elsby, Michaels, and Solon (2009) show, job-finding probability will be given by the following relationship:

$$\hat{F}_t = 1 - \left[\left(u_{t+1} - u_{t+1}^s\right)/u_t\right] \tag{7}$$

which maps into an outflow hazard, job-finding rate, $F_t = -\log(1 - \hat{F}_t)$. This formulation in (7) computes the job-finding probability for the average unemployed person by implicitly assuming that contraction in the pool of unemployed, net of newcomers to the pool (u_{t+1}^s), results from unemployed workers finding jobs. The next step is to estimate the separation rate S_t. This

[6] It is seasonally adjusted at an annual rate and expressed in chained 2005 dollars.

step involves solving the continuous-time equation of motion for unemployment forward to get the following equation, which uniquely identifies S_t.

$$u_{t+1} = \frac{\left(1 - e^{-F_t - S_t}\right) S_t}{F_t + S_t} l_t + e^{-F_t - S_t} u_t \tag{8}$$

Given the outflow hazard, F_t, measured through (7), and data on u_t and l_t, we can solve for S_t numerically for each month t. One potential problem that could bias our estimates is the redesign of the CPS in 1994. As discussed by Shimer (2007) and Elsby, Michaels, and Solon (2009), the CPS redesign deflated the actual number of short-term unemployed by changing the way it computes this for every rotation group except the first and the fifth.[7] To correct for this bias, we follow Elsby, Michaels, and Solon (2009) and use the average fraction of short-term unemployment among the unaffected first and fifth rotation groups to inflate the aggregate short-term unemployment number. This reduces to multiplying every month's u^s_{t+1} by 1.1549 from February 1994 through the end of the sample period. Following this correction finally provides us with the data we need for unemployment flow rates.

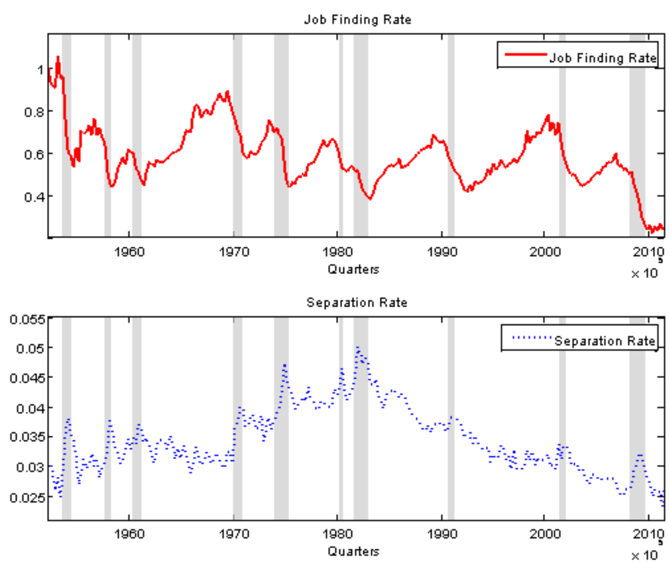

Figure 1: Job-finding and separation rates are constructed using equations (7) and (8) and corrected for CPS redesign. Shaded areas indicate NBER recession periods. Rates are the quarterly averages of the monthly data.

[7] See Polivka and Miller (1998) and Abraham and Shimer (2001) for more detail.

As figure 1 shows, these flows generally follow a pattern in a typical business cycle. As the economy enters a downturn, separations start rising, and job-finding rates start falling. These movements cause the overall unemployment rate to rise. But the separation rate usually stabilizes before the unemployment rate peaks. After the separation rate levels off, most of the subsequent increase in the unemployment rate is caused by a low job-finding rate. Note that this combination implies that the average duration of unemployment gets longer, although the flow of people into the pool of unemployed workers does not increase. The low job-finding rate means that the flow of workers out of the pool of unemployed slows enough to cause an increase in the average duration of unemployment. When the economy finally starts recovering, durations decrease as firms create new jobs and absorb some of the unemployed. The unemployment rate falls. However, this highly stylized description of cyclical movements in these rates ignores the varying degree of importance of one flow or another in accounting for unemployment fluctuations over a particular cycle. For instance, separations seem to have been more responsive to the most recent cycle than during the previous two cyclical downturns. In fact, this relative dominance of the job finding rate before was what led Shimer (2007) to conclude that the job-finding rate is the more important flow, at least for cyclical changes in unemployment; it also spurred a large body of literature that explicitly assumed that separations are not cyclical.[8] Since we have a model which distinguishes between cyclical and trend components of these flows, we can analyze the contributions of each flow to unemployment fluctuations more explicitly. We explain our findings in section 4.1.

Our constructed data cover most of the post–World War II recessions; however, we present the data only since 1952 here, to be consistent with our estimation in the next section. More importantly, figure 1 shows that there are cyclical fluctuations in these flow rates and some general low-frequency movement, which is especially apparent for the separation rates. Hence, we believe that the reduced form model we laid out is a sensible one. Our next task is to estimate the underlying trend in both flow rates, more specifically, $\bar{f}_t$ and $\bar{s}_t$. This is what we discuss in the next section.

[8] For the debate on which flow drives unemployment fluctuations over business cycles, see, for instance, Shimer (2007), Elsby, Michaels, and Solon (2009), and Fujita and Ramey (2009).

3.2 Estimation

We estimate the reduced form model in (1)-(3) via maximum likelihood, and use the state-space representation in (4)-(5). Since the stochastic trend and cyclical components of our variables are not observable, we rely on a Kalman filter to infer them and construct our log-likelihood. One important issue we need to address is the identification problem. This arises from the fact that one observable variable in each equation, (1)-(3) is forced to identify movements in more than one error term. One way to get around this problem is to impose a relative ratio for the standard deviations of trend and cyclical components.[9] For instance, let $\gamma_f = \frac{\sigma_{fn}}{\sigma_{fc}}$ be the relative variance of the error in the trend of the job-finding rate to that in its cycle. This will be a free parameter in our estimation and, in principle, our results might depend on the value of γ_f. Similarly, $\gamma_s = \frac{\sigma_{sn}}{\sigma_{sc}}$, would be a parameter of our estimation with regard to the behavior of the separation rate. The problem is also evident for the real output, since we have three error terms governing movements in the observable output. We start with relative ratios based on those reported in Kim and Nelson (1999) for output. One encouraging fact is that the likelihood function varies in a significant way with the relative ratios, $\gamma_y = \frac{\sigma_{yn}}{\sigma_{yc}}$, $\gamma_g = \frac{\sigma_g}{\sigma_{yc}}$. Hence, we pick the γ_y, $\gamma_g$that yields the highest log-likelihood.[10] Unfortunately, the case for γ_f, γ_s is less obvious. In that case, we estimate our model for various values of γ_f, γ_s and pin down our preferred values by looking at two statistics—the log-likelihood and correlation between the inferred natural rate and the trend of the actual unemployment rate—using a bandpass filter. The idea here is to preserve the likelihood of the model while at the same time inferring a natural rate that is not far from the low-frequency statistical trend of actual unemployment. We discuss our results from this robustness check below. As a result of this exercise, for the benchmark case we choose a parameterization where $\gamma_f = 1$, $\gamma_s = 1.5$. In the next section, we report how our estimation varies with other values for these parameters.

Another minor point in our estimation concerns the random-walk nature of our model. The stochastic trend components are modeled as random walks; hence, we need to initialize the variance–covariance matrix for the Kalman filter with something other than the unconditional

[9] Laubach and Williams (2003) addresses a similar problem in the context of an unobserved components model for the natural rate of interest.

[10] They are 0.85 and 0.027, respectively.

mean. To get around this problem, we start with a diffuse prior, that is, a high initial variance for our unobserved state variables, and remove the first 16 quarters from our actual estimation in order to reduce the impact of this arbitrary initialization. Therefore, we report our estimates starting from 1952:Q1 instead of the beginning of our sample.

4 Results

Here, we present the results of our benchmark estimation, imposing the restrictions $\gamma_f = 1$, $\gamma_s = 1.5$, $\gamma_y = 0.85$, $\gamma_g = 0.027$. This implies that we only estimate 11 parameters. As table 1 shows, all parameters of the reduced form model in (1) - (3) are quite tightly estimated, with the possible exception of θ_3 Given our estimates of the parameters, we can use a Kalman filter to back out the unobserved state variables, namely, $\bar{f}_t$, $\bar{s}_t$ and $\bar{y}_t$. Given these unobserved states, we can compute the implied long-run steady state of the unemployment rate for every quarter with the identity $\bar{u}_t = \frac{\bar{s}_t}{\bar{s}_t + \bar{f}_t}$. Figure 2 shows the trends in the job-finding rate, the job-separation rate, and the unemployment rate using these estimates.

Table 1: Estimation Results: 1952:Q1-2011:Q2

	Estimate			Estimate	
ϕ_1	1.6318	(0.0594)	θ_2	0.1058	(0.0470)
ϕ_1	-0.6805	(0.0583)	θ_3	0.0370	(0.0253)
ρ_1	1.0874	(0.5730)	σ_{yn}	0.0060	(0.0003)
ρ_2	3.8076	(0.9409)	σ_{fn}	0.0188	(0.0012)
ρ_3	-1.1793	(0.6316)	σ_{sn}	0.0006	(0.00005)
θ_1	-0.2016	(0.0313)	L	2466.9	

Standard deviations are in ().

Looking into the underlying trends in unemployment flows gives us considerable insight into the nature of time variation in the trend of the unemployment rate, that is, the natural rate. Both the job-finding and separation rates have trended down over time—the separation rate for almost three decades, the job-finding rate mostly in the last decade. If there were not any significant decline in the trend of the job-finding rate, but only an increase in the trend of the separation rate, our definition of the time-varying unemployment trend would imply an increase in its level. According to our estimates, this was indeed the case throughout the 1970s. The opposite has been happening since then for the separation rate trend; it has shown a secular decline since the early 1980s. Over the course of three decades, the separation rate trended down by almost 50 percent. Over the same period, however, the job-finding rate trend declined but only by a smaller magnitude. Hence, the implied "natural rate" started to decline from its peak levels in the early 1980s. These general patterns seem to be consistent with findings in the literature on the natural rate. Overall, our estimates suggest that over the last four decades, the unemployment rate trend has moved between 5 percent and 7 percent, and currently stands around 6.1 percent.

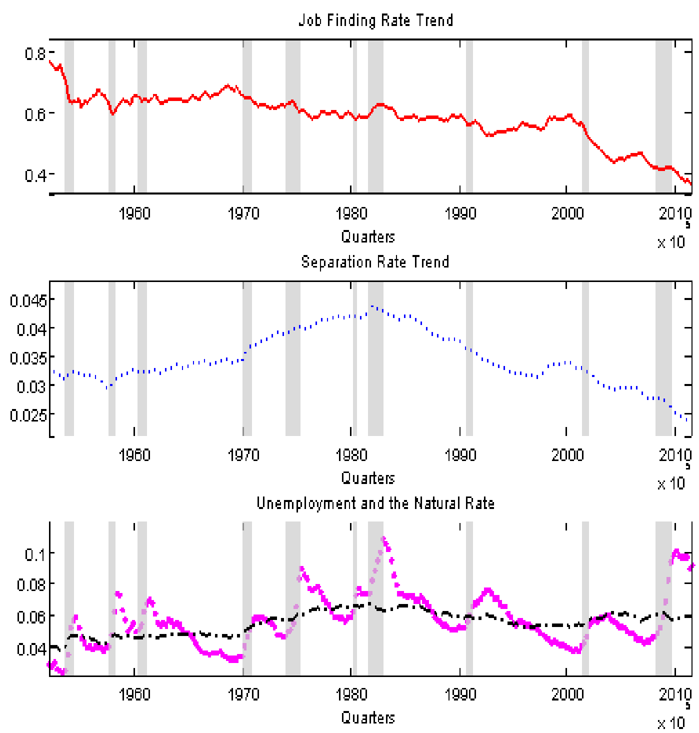

Figure 2: Unobserved trend in all three variables are backed out and smoothed by Kalman filter. Shaded areas indicate NBER recession dates. In the third panel, line (-.) indicates the natural rate as the ratio of separation trend ($\bar{s}_t$) over total worker reallocation ($\bar{s}_t + \bar{f}_t$).

Perhaps the more interesting point about our estimates of these trends is that worker reallocation, as measured by the sum of the job-finding and separation rates, is declining in the U.S. This is a crucial result with important implications for the natural rate as well as how the adjustment in the observed unemployment rate might evolve over time. These results give us considerable insight into the nature of recent changes in unemployment rates. We see that the declining job-finding rate is not temporary, but part of a long-run trend. Along with the more obviously declining trend in separation rates, the declining trend in job-finding rates essentially implies that U.S. labor markets are exhibiting increasingly less worker reallocation. Not only are workers finding jobs at a slower rate on average; independent of the state of the economy, they are also losing (or leaving) their jobs at a slower average rate. This picture of less reallocation also appears to apply to jobs. Several studies show that job reallocation in the US has shown signs of decline over the course of the last two decades; see, for instance Faberman (2008) and Davis et al. (2010). Slower worker reallocation affects the rate of convergence of observed unemployment towards its long-run trend. The sum of these two rates, in essence, determines how fast the economy is able to gravitate to its imputed trend. Hence, one clear implication is that the adjustment from current levels of unemployment towards the level of 6.1 percent will take longer than it would in an economy with more churning but the same implied natural rate.

These results, in principle, could be sensitive to the exact values of γ_f, and γ_s that we use. In our benchmark estimation, we stick to values of 1, and 1.5, respectively. As figure 1 shows, the separation rate has a much clearer low-frequency trend than the job-finding rate. Hence, it is reasonable to have a relatively smoother trend in the separation rate, as our benchmark values of γ_f, and γ_s imply. To pin down the exact numbers, we re-estimate our model over a fine grid for both γ_f, and γ_s; $\gamma_f = \{0.25, 0.375, 0.5...., 3.375, 3.5\}$ and $\gamma_s = \{0.5, 0.625, 0.75...., 3.875, 4\}$. We look at two moments to match: One is the maximum log-likelihood over this combination of points; the other is the correlation between the implied natural rate from our estimation and the trend of the observed unemployment rate, calculated using a bandpass filter. Since we do not use actual unemployment rate in the estimation, we are trying to impose some discipline on our estimation by bringing in these data.[11] The objective here is to maximize the likelihood of the model without getting an implied unemployment trend that is far from a statistical trend. Figure 3 shows how these two moments change across γ_f, and γ_s.

[11] Note that, with the flow rates themselves, the unemployment rate does not give any more information for our reduced form model; hence, it is not part of it.

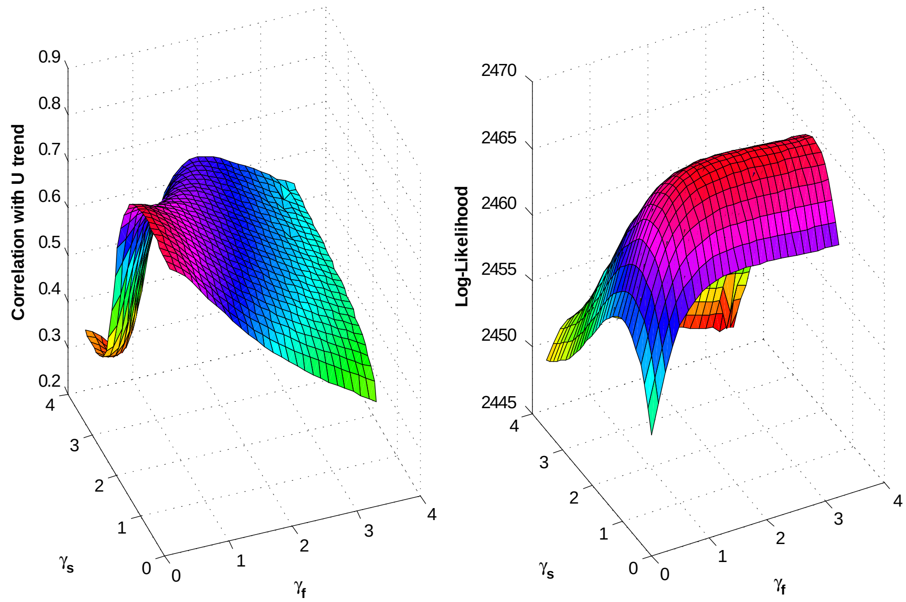

Figure 3: Left panel shows the correlation between the implied natural rate and the statistical trend of the observed unemployment rate computed by bandpass filter, for different values of γ_f, and γ_s. Right panel shows the value of log-likelihood for different γ_f, and γ_s.

Our preferred benchmark values maximize the objective of high log-likelihood and high correlation, as is also clear in figure 3. For instance, we do not improve the likelihood of the model for higher values of γ_f, whereas smaller values result in substantial declines. The likelihood value seems more concave in γ_s, and our preferred value of 1.5 is close to its maximum. As we decrease γ_s, the trend of the separation converges to a straight line; hence, the natural rate will be determined more by the trend of the job-finding rate. The opposite is true when γ_f is small and its trend is close to a straight line. Hence, when one flow has a constant trend imposed (low γ_i), and the other flow has a very small cyclical variation (high $\gamma_j, j \neq i$), we miss the low-frequency movements in the observed unemployment rate by a significant margin. Our objective function determines the optimal trade-off between these two dimensions by putting more weight on the more informative moment, that is, by using the inverse of the covariance matrix as the weighting matrix. Finally, for almost all of the values of γ_f, and γ_s, the natural rate implied by the model varies between 5.5 percent and 6.2 percent at the end of the sample.

4.1 The Ins and Outs of the Natural Rate

Throughout this paper, we have argued that flows provide us with more information about the unemployment rate than unemployment itself could provide. We can distinguish between the forces that affect the duration of unemployment versus those that affect its incidence. Unemployment at any point in time is determined by the magnitude of one flow relative to the other. As we discussed earlier, much of the high-frequency movement in the unemployment rate seems to come from high-frequency variation in the job-finding rate, as shown in figure 1. There is a body of literature focused on teasing out the particular flow that drives unemployment fluctuations over the business cycle.[12] Our model laid out in the previous section gives us the estimates of cyclical and trend components in the underlying flow rates. Hence, in principle, we can use a similar decomposition used in Fujita and Ramey (2009) to analyze the contribution of each flow rate to variations in the unemployment rate, both at the high frequency and the low frequency. In particular, let $\Delta u_t = \log\left(\frac{u_t}{\bar{u}_t}\right) = \log\left(\frac{S_t}{S_t+F_t} / \frac{\bar{s}_t}{\bar{s}_t+\bar{f}_t}\right)$ denote the variation in the unemployment rate in period t from its time-varying trend implied by the model. Similarly, define the variation in the separation and job finding rate from their time-varying trends respectively as $\Delta s_t = \log\left(\frac{S_t}{\bar{s}_t}\right)$ and $\Delta f_t = \log\left(\frac{F_t}{\bar{f}_t}\right)$. Fujita and Ramey (2009) shows that the contributions of each worker flow to high-frequency variation in the unemployment rate are given by factors, $\beta^s = \frac{cov((1-\bar{u}_{t-1})\Delta s_t, \Delta u_t)}{var(\Delta u_t)}$ and $\beta^f = \frac{cov((1-\bar{u}_{t-1})\Delta f_t, \Delta u_t)}{var(\Delta u_t)}$. One can write down a similar decomposition for the low-frequency variation in the unemployment rate's trend, i.e. variations in our estimate of the natural rate, $\bar{u}_t$, relative to its historical mean, $\bar{u}$, by redefining the objects, $\Delta \bar{u}_t = \log\left(\frac{\bar{u}_t}{\bar{u}}\right)$, $\Delta \bar{s}_t = \log\left(\frac{\bar{s}_t}{\bar{s}}\right)$ and $\Delta \bar{f}_t = \log\left(\frac{\bar{f}_t}{\bar{f}}\right)$, where $\bar{u}$, $\bar{s}$ and $\bar{f}$ denote average trend values for the relevant variable. Corresponding factors for the trends are then defined as $\bar{\beta}^s = \frac{cov((1-\bar{u})\Delta \bar{s}_t, \Delta \bar{u}_t)}{var(\Delta \bar{u}_t)}$ and $\bar{\beta}^f = \frac{cov((1-\bar{u})\Delta \bar{f}_t, \Delta \bar{u}_t)}{var(\Delta \bar{u}_t)}$.

Figure 4 shows the respective variation in the cyclical and trend components of both flows. It is clear that most of the variation in cyclical components is driven by the variation in the job finding rate's cyclical component. However, as the lower panel of figure 4 shows that, for most of the sample period, separation rates alone can explain much of the variation in the trend component of the unemployment rate. Until about the beginning of the 2001 recession, the

[12] See for instance, Shimer (2007), Elsby, Michaels, and Solon (2009), Fujita and Ramey (2009) and Barnichon and Figura (2010), as well as earlier work by Darby, Haltiwanger, and Plant (1986).

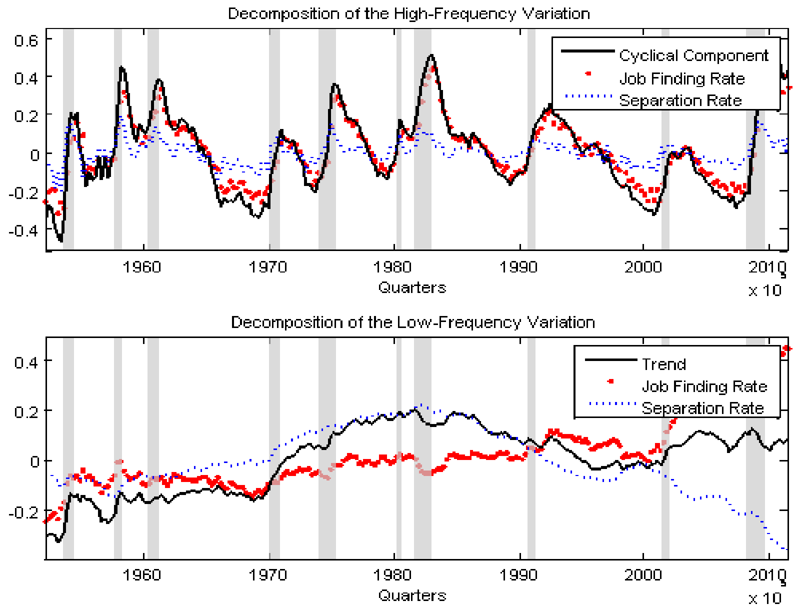

Figure 4: Upper panel plots Δu_t, Δs_t and Δf_t over time. Lower panel shows $\Delta \bar{u}_t$, $\Delta \bar{s}_t$ and $\Delta \bar{f}_t$.

separation rate trend can account for most of the behavior of the natural rate. In a sense, this is not very surprising, given the small variation in the job-finding rate trend over this period relative to the last 10 years in the sample (figure 2). The picture for the last decade is starkly different. It is clear that neither of the flow rate trends by themselves can generate the observed variation in the estimated natural rate in figure 4. There are offsetting effects of the trend changes in both flows.

Table 2 summarizes the information in figure 4 in a different way by providing the variance decomposition factors at different frequencies and sample-periods. Throughout the whole sample period, It seems though the job finding rate consistently explains more than three fourths of the variation in the cyclical component of the unemployment rate. The dominant role for the job finding rate, however, is mostly present for the variation in the unemployment rate trend after 1985. For the period before 1985, the separation rate trend explains more than 60 percent of the variation in the natural rate. This changes in the rest of the sample, with the job finding rate explaining almost two thirds of the variation in the trend unemployment. Hence, our analysis in this paper not only confirms the dominant role of the job-finding rate for unemployment fluctuations at the business cycle frequency, but also underscores the importance of the separation rate for the long-run trend in unemployment for most of our sample period.

Table 2: Variance Decomposition for Unemployment Rate

	Cyclical Component				Trend Component		
	pre - 1985	post - 1985	1952- 2011		pre - 1985	post - 1985	1952- 2011
β^f	0.7263	0.7979	0.7550	$\bar{\beta}^f$	0.3614	0.6714	0.6372
β^s	0.2774	0.2095	02500	$\bar{\beta}^s$	0.6397	0.3284	0.3612

5 Discussion: The Great Recession

Between December 2007 and June 2009, the U.S. economy experienced one of the worst recessions since the Great Depression. Over the course of the latest recession, the U.S. economy shrank by 5.1 percent. This large aggregate shock had correspondingly large effects on the labor market. A total of 8.7 million jobs were lost (through February 2010), and the unemployment rate rose from 4.7 percent to a peak of 10.1 percent in late 2009. Currently, more than 14 million people are officially unemployed, and many are underemployed. More striking is the length of time people remain unemployed. Unemployed workers stay jobless for 39 weeks on average now, about twice as long than at previous cyclical peaks. These large effects of the aggregate shock on the labor market raise some obvious questions: Has the recession changed the long-run trend for the unemployment rate?

5.1 Has the Great Recession Changed the Long-Run Trend?

Given the accompanying substantial decline in employment in some sectors (construction, finance, manufacturing), it might be natural to expect a change in the trend after the longest recession since World War II. It is conceivable that sectoral reallocation, lower matching efficiency, and longer durations of unemployment insurance compensation might lead to changes in the natural rate. To the extent that these changes are reflected in the measured flow rates, our framework can capture this change in the trend. One obvious way to answer this question is to look at our estimates of the natural rate before and after the recession. Our estimate in 2007:Q3, just before the recession started, was approximately 6.1 percent. Even though the natural rate, estimated using our method, hit around 6.3 percent in the midst of the recession, it is back around 6.1 percent at the end of the sample.[13] Most of the intervening slight increase over the recession resulted from a sharp increase in the separation rate, which represented a temporary slowdown in the declining secular trend in the separation rate. The Kalman filter seems to have identified the surge in separations partly as a trend slowdown.

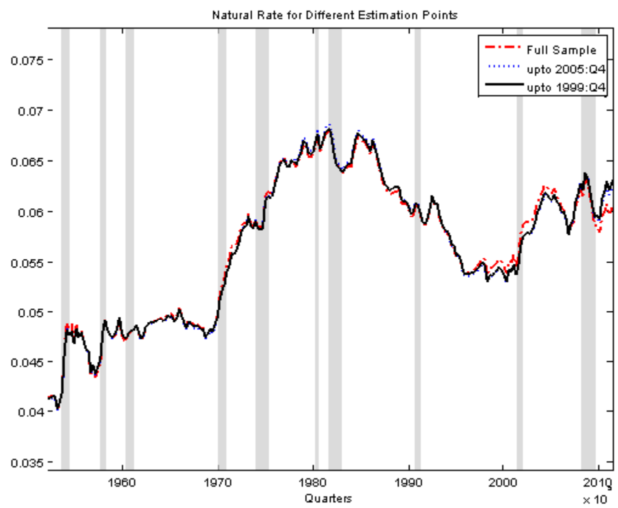

Figure 5: Shaded areas indicate NBER recession dates.

Another issue that has been raised about the effects of the last recession is that the comovement of unemployment with output has changed substantially.[14] Our framework provides a nice testing ground for this. Obviously, since we do not have a structural model, it is impossible for us to distinguish among potential reasons. However, in a reduced form sense we can see whether the last recession in fact changed the underlying nature of the comovement between output and flows into and out of unemployment. We conduct this test by estimating our model for different sample periods during which we think that these "structural" changes may have happened, and then letting the Kalman filter back out the unobserved states with the full-sample data. If there is any substantial difference between the implied natural rates, that difference will be due to the changing structure of the relationship between unemployment

[13]For a slightly different view on this see Weidner and Williams (2011). Even though they argue that the 'natural rate' might have increased as much as 1.7 percentage point to 6.7 percent, their conclusion about the prospect of short-term adjustment is similar to our conclusion here.

[14]See, for instance, Daly and Hobjin (2010) and Gordon (2010a and 2010b).

flows and output. This is obviously not a test for a regime change in the usual sense; however, it is a relatively simple way to address the question within the scope of this paper.

We re-estimate our model with two more subsamples, before 2006 and before 2000. The first subsample, which includes data through 2005:Q4, excludes data for the last business cycle and most of the recovery after the previous recession. However, the second subsample, which includes data until 1999:Q4, excludes data on the previous recession, that is, the last jobless recovery episode. We present our results in figure 5 for both subsamples and the full sample. Note that, regardless of where we end our estimation, the implied natural rate is very close to the estimated one from the full sample. The differences between the three reported estimates are very small. Hence, this simple test shows that the last recessionary episode did not significantly change the natural rate through its effects on the parameters of the model.

5.2 Why Is the Decline in the Unemployment Rate So Slow?

Even though we contend that we probably have not seen a significant increase in the natural rate over the last several years, we can safely predict that convergence to the estimated natural rate will be slow for two reasons: The first is the sheer extent of the gap between the current unemployment rate and its estimated trend level. This gap reflects the size of the aggregate shock that hit the economy. When the U.S. economy experienced a similarly sized shock after the 1981–82 recession, it took several years for the observed unemployment rate to drop to levels closer to the trend, even though the rebound in output growth was exceptionally strong. Second, as we argued earlier, slower worker reallocation will itself imply slower adjustment because the adjustment rate depends on how fast workers are reallocated between unemployment and employment.

We present two numerical exercises in this section to show the quantitative significance of these implications. The first exercise compares the behavior of labor market aggregates since 2009:Q3 with a hypothetical scenario in which output growth rate experiences the same shocks as it did after the 1982 recession. The second exercise, on the other hand, compares simulations that use current reallocation rates with the counterfactual, in which labor markets have much more churning.

Clearly, our simple empirical model implies that strong output growth will lead to a faster

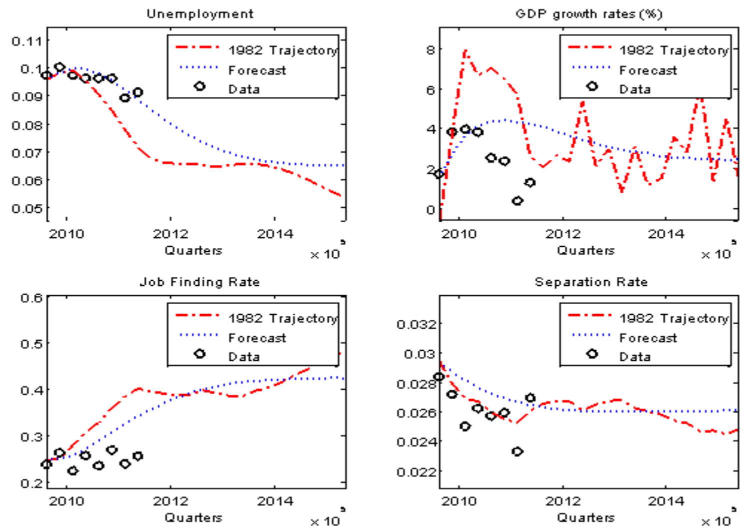

Figure 6: The line labeled '1982 Trajectory (-.)' plots the results from model simualtions with $\varepsilon_t^g, \varepsilon_t^{yn}, \varepsilon_t^{yc}$ set to their realizations during the 24 quarters after 1982:Q3, when the recovery started according to the NBER. 'Forecast (...)' presents the unconditional forecast from the model. They are both expressed as averages of 10,000 simulations of 24 quarters starting from 2009:Q3, when the current recovery started. GDP growth rates are annualized.

recovery in the labor market, as the cyclical components of the job finding and separation rates disappear sooner. There is some concern among economists that the current pace of the economic recovery is relatively weak compared to historical norms, especially before the mid-80s. The upper-right panel of figure 6 provides some evidence that this may indeed be the case. According to our model, the growth rate of real GDP, at this point in the recovery, is predicted to be twice the rate it is in the data, at least during the last two quarters of the sample. These predictions are based on average of 10,000 simulations of the model, each one for 24 quarters, starting from the third quarter of 2009. Based on the parameter estimates, average GDP growth rates at this point in the cycle would have been around 4.3 percent. One can compare the path of unemployment under this scenario with a particular realization of shocks, $\varepsilon_t^g, \varepsilon_t^{yn}, \varepsilon_t^{yc}$, in a specific episode. Our benchmark here is the recovery after the 1982 recession. To do this comparison, we back out the realization of these shocks from 1982:Q3 and feed them into the model simulations, generating a forecast for four variables conditional on a particular output growth path. Comparing this conditional forecast, which follows a post-1982 trajectory in terms of output growth, with the unconditional forecast from the model

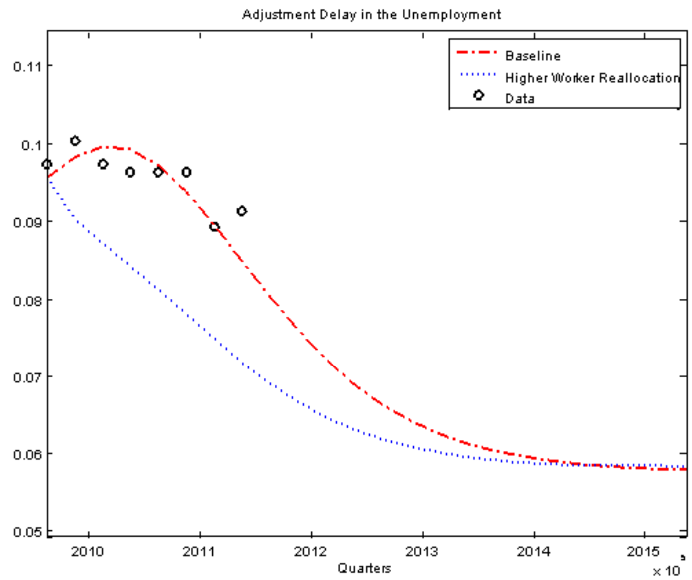

Figure 7: Average of 10,000 simulations. Unemployment is imputed based on the simple equation of motion for unemployment, (6), and predicted values of flow rates. Baseline refers to the benchmark case where worker reallocation rates are consistent with current estimates.

shows that, along the transition path, the decline in the observed unemployment rate could be significantly lower with a weaker recovery, as large as 1.75 percentage points. Figure 6 also shows that the model overestimates the job finding rate and underestimates the separation rate in the near term, providing us with a relatively accurate forecast of unemployment for the past seven quarters. Overall, the results of this exercise suggest that some of the persistence in the current level of the unemployment rate could be explained by the weakness of output growth, both relative to historical averages predicted by the model, and the particular recovery episode following the 1982 recession.

Next we try to quantify the effect of slower worker reallocation on the unemployment rate's convergence towards a long-run trend. The effect of slower worker reallocation on the pace of the adjustment process might be as strong as that of output growth, according to our second numerical experiment. This experiment involves comparing the path of unemployment under two different assumptions about worker reallocation. First we generate a set of simulations using the levels of job finding, and separation rate trends at the end of 2009:Q2, 0.42 and 0.026, respectively. Using the equation of motion for unemployment, eq. (6), and an initial rate of unemployment, one can generate a forecast path for unemployment from 2009:Q3 onward. We label this path as the baseline in figure 7. Our counterfactual is from a period where trend worker reallocation was very high, as measured by the sum of job finding and separation rate trends. More specifically, we set the job finding rate trend, $\bar{f}_t$, in the last period of our sample to the level it was in 1982:Q4. This amounts to $\bar{f}_t = 0.62$, implying a large shock to the job finding rate trend at the end of the sample. Note that this is very close to the sample average of this rate, which is 0.59. Since trend flow rates follow a random walk, this will have permanent effects. In order to be consistent, we also set $\bar{s}_t$ to a higher level at the end of the sample so that the unemployment rate converges to the same level in the long-run. This requires setting $\bar{s}_t = 0.039$, which is very close to the separation rate trend in 1982:Q4. As figure 7 shows, higher worker reallocation clearly implies a faster decline in the observed unemployment rate. The difference could be as large as 1.6 percentage points along the transition path, even though both economies ultimately converge to the same long-run level.

As both of these experiments suggest, having a relatively unchanged unemployment rate trend even after the last recession does not necessarily imply an optimistic picture for the unemployment rate in the near term. The strength of the growth in real output and the effects of slower worker reallocation in the US labor market will be among the crucial factors determining this adjustment process. The significance of the latter factor is a novel feature of the framework we use in this paper, and it suggests that structural reasons behind slow worker reallocation might have important implications for unemployment dynamics over business cycles. Understanding these structural factors requires going beyond our reduced-form framework, and it is clearly beyond the scope of this paper.

6 Robustness

The method we used to uncover a time-varying unemployment trend is ultimately a filtering method. One can reasonably ask whether our estimate of this unemployment trend moving at a low-frequency is robust to different filters. Our framework also implicitly argues for uncovering trends in the underlying flow rates as opposed to the unemployment rate itself. In this section, we present some evidence that shows that what is really important is not the specific filtering method per se, but the inclusion of worker flow information instead of the unemployment rate alone. Moreover, we show that our model provides estimates for the natural rate that are at least as precise as those one can get from a reasonable alternative that uses a similar methodology but ignores the additional information in worker flows. Our method also requires, on average, much less revision to these estimates as we incorporate more data. We also present some evidence that our results are very robust to changes in the labor force participation rate over time and allowing for worker flows into and out of inactivity.

6.1 Alternative Methods and Filters

6.1.1 Statistical filters and the use of worker flow rates

Using a Kalman filter helps us to put some structure on the empirical relationship between worker flows and aggregate economic activity. However, it comes with a reduced form we imposed on the data. One might argue that if the objective is to derive an empirically useful unemployment rate trend, a pure statistical trend of the unemployment rate might be more practical, if worker flow information does not seem to provide us with any additional information. Having already established in the previous section that worker flows are in fact crucial to understanding the long-run behavior of the unemployment rate, we focus here on different statistical filtering methods with and without worker flows to distinguish the role they play.

Taking an HP-filter of the unemployment rate itself has been one approach used in the literature to identify a trend for the unemployment rate in the context of the natural rate debate (see Rogerson (1997)). We compare our estimate of the long-run trend for the unemployment rate with those that could be obtained using an HP or a bandpass filter. Figure 8 presents the results of this exercise. When we omit the information on worker flows and filter the quarterly unemployment rate, we find a lot of variation in the trend and significant diversion across different filters. For instance, applying an HP-filter with a high smoothing parameter gives a relatively smooth trend that moves closely with our preferred trend from the model. However, a bandpass filter or an HP-filter with a smaller smoothing parameter produces much more variation in the trend. The lower panel also shows the well-known problem of overemphasizing the end points of the sample.

A strikingly different picture emerges if we include information on worker flows and impute

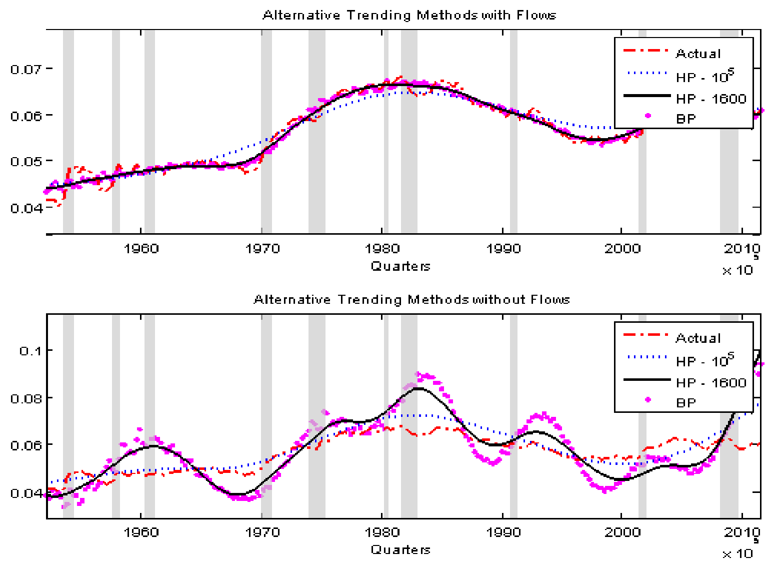

Figure 8: The upper panel presents unemployment rate trends imputed by different statistical filters on worker flow rates. The lower panel presents pure statistical trends based solely on unemployment rate data. The line labeled actual - displays our preferred version that is based on our model. We also use an HP-filter (with smoothing parameters 1600, and 10^5) as well as a bandpass filter (with parameters $(6, 32)$).

an unemployment rate trend, as we did in our model, based on the trends of these underlying flows. As the upper panel of figure 8 shows, unemployment trends imputed this way do not vary much across different filters and are much smoother than the trend estimates based solely on unemployment rate information. Moreover, our model, which puts a lot more structure on the comovement of worker flows and real output, produces a trend that moves closely with these other filters. We interpret this result as evidence of the importance of worker flows in understanding the unemployment rate trend over the long run. The obvious discrepancy between various estimates of the trend with different filters when worker flows data are ignored makes it harder to get an empirically consistent, useful measure.

6.1.2 Precision of the estimates and revisions

An important issue in the empirical literature that tries to estimate the natural rate (of either unemployment or interest) is the precision of the estimates and the significant revisions observed with the inclusion of subsequent data. Here, we briefly discuss how the empirical model we proposed in this paper performs on these two fronts. We find that, in terms of precision

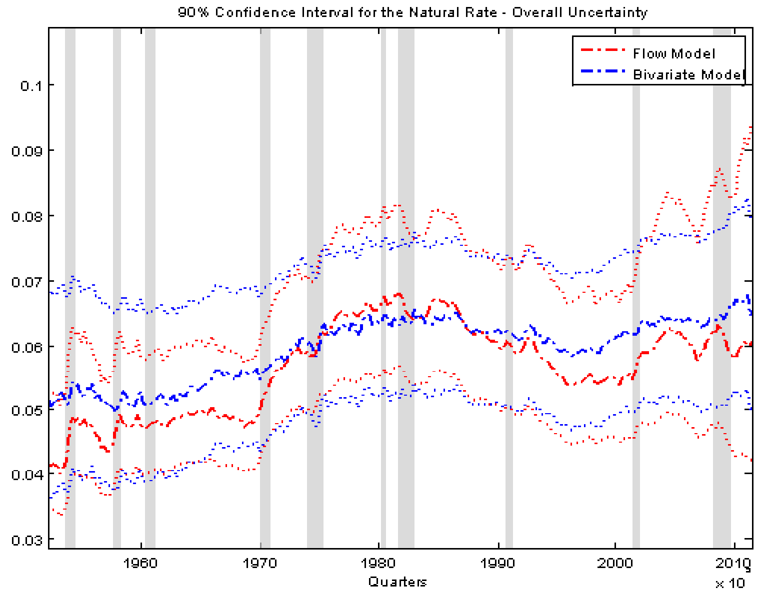

Figure 9: Flow model refers to the model expressed in equations (1)-(3). Bivariate model is the benchmark bivariate model without worker flows, as in Kim and Nelson (1999, 38-9).

of estimates, our model with worker flows performs as well as a benchmark bivariate model that uses data only on unemployment rate and real output but also relies on Kalman filter to infer the unobserved trend in the unemployment rate. Moreover, the model with worker flows does not imply significant revisions to previous estimates of the unobserved trend as the benchmark bivariate model does, thereby making it a useful method to estimate a natural rate more consistently over time.

It is well known that the estimated state vector of an unobserved components model, such as our model here, is subject to both parameter and filtering uncertainty. Using a standard Monte Carlo procedure, we compute the 90 percent confidence bands around the estimates of the unobserved state (unemployment's trend) in our model.[15] We compare the precision of our estimates with those estimated from a benchmark bivariate model, that resembles our model in structure, but does not involve any worker flows data. More specifically, we estimate a bivariate model presented in Kim and Nelson (1999, pp. 38-39), which assumes an unobserved trend in unemployment rate that is a random walk and the cyclical component of the observed unemployment rate depends on the cyclical component of the aggregate output.

[15] Details are available upon request.

Figure 9 plots the overall uncertainty around the estimate of the unemployment rate trend in both cases. Even though it looks like the bivariate model has a narrower confidence band towards the end of the sample, on average the flow model does not perform exceptionally bad. For instance, the width of the 90 percent confidence band implies, on average, a range of 2.35 percentage points around the mean estimate over time in the flow model (-0.97 and 1.38). The benchmark bivariate model performs slightly worse, with a range of 2.47 percentage points around the mean estimate over time (-1.18 and 1.29). The standard deviation of the error band is also slightly smaller for the flow model, 0.75 vs. 0.83. Hence, our empirical model does as well as, if not better than, a reasonable alternative that uses a similar methodology but ignores the additional information in worker flows.

Another desirable feature of our framework is its robustness to additional data. Since we use Kalman smoothing to back out the unobserved states, as additional information becomes available estimates of the unobserved state might change, in principle, all the way back to the beginning of the sample. In this respect, the model with worker flows performs remarkable well relative to the benchmark bivariate model we used to compare the precision of estimates. The particular numerical exercise we conduct is very simple: We re-estimate both models with two more subsamples, before 2006 and before 2000, and compare the estimates of the unemployment rate trend for each case until 1999:Q4. Ideally, if we have a robust approach, the addition of a small set of new data should not change the estimates of the unobserved state, i.e. the natural rate, prior to 1999:Q4.

Table 3: Revisions to the Natural Rate (% points) before 1999:Q4

	Bivariate Model		Flow Model	
Excluded Data	2005:Q4-current	1999:Q4-current	2005:Q4-current	1999:Q4-current
Avg.	0.0775	0.0678	0.0338	0.0259
Std.	0.0856	0.0757	0.0411	0.0323

Table 3 presents some interesting moments from this numerical exercise. On average, our model revises its estimate of the natural rate by a much smaller margin than the benchmark bivariate model does. For instance, the natural rate estimate is revised by on average 0.034 percentage points in the flow model when we include additional data covering the period after

2005:Q4, as opposed to 0.078 percentage points in the benchmark bivariate model. As columns three and five in table 3 show, this result is robust to the inclusion/exclusion of the entire last decade. The variation in the required magnitude of revisions is almost twice as large in the benchmark bivariate model. Hence, we conclude that our framework based on worker flows is superior, not only to purely statistical filters, but also to similar approaches that ignore worker flows data.

6.2 Labor Force Participation

Starting in the early 1970s, labor force participation began to rise significantly in the U.S. This well-known secular trend was mostly due to higher female labor force participation. In principle, our measurement of the job finding and separation rates could be affected by changes in the labor force. The overall size of the labor force affects this measurement through the equation of motion for unemployment, (6). We do not expect to see any effect on the overall job finding rate, since it uses only information on the stock of unemployed by duration (see eq. (7)). However, separation rates are implicitly a function of the overall labor force, as eq. (8) shows. Hence, both high, and low-frequency movements in the size of the labor force will directly affect this measurement and our estimate of the unemployment trend. By taking the size of the labor force exogenously over time, we are to some extent controlling for these effects in our estimation. However, variations in the labor force participation rate turns out to have very minor quantitative effects on our results. To see this point, we compare worker flow rates from our benchmark measurement with the rates that are computed assuming a constant labor force participation rate over the whole sample. More specifically, we create an alternative set of job finding and separation rates where the labor force is assumed to grow such that labor force participation is constant over time. If movements in the labor force, such as the secular rise in female labor force participation, drives our results, this comparison will provide us with a quantitative measure of the impact.

Figure 10 plots the results of this experiment. Not surprisingly, imposing a constant level of labor force participation does not alter the job finding rate. However, separation rates are affected by this assumption to some extent. The impact on the separation rates is so small that it does not change our estimates of the unemployment rate trend over time (not shown in

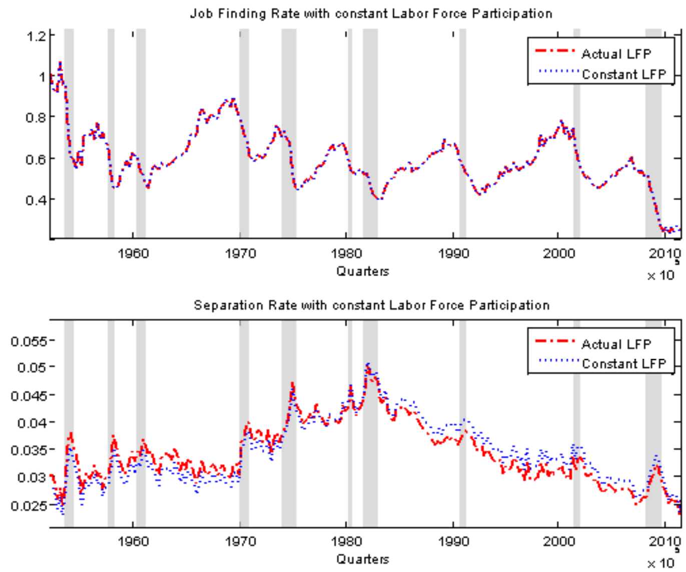

Figure 10: Job finding and separation rates measured in two different ways: Actual LFP uses the labor force in the data to measure worker flows in (6) - (8). Constant LFP assumes a constant labor force participation (the mean) when measuring worker flows in (6) - (8).

figure 10). This result confirms that our methodology of imputing an unemployment rate trend based on worker flows is very robust to changes in the aggregate labor force participation rate over time in the U.S.

Unfortunately, there are other ways in which our results could change because of movements in the labor force. The entire methodology we have used for measuring worker flows has been standard since Shimer (2005). However, it does not allow for any separations into inactivity and flows into employment from out of the labor force. When these flows are taken into consideration, measures of job finding and separation rates will change. To the extent that these flows have non-negligable effects on the labor force participation rate, or more precisely flows into and out of the labor force, it potentially could affect our estimation. To extend our methodology this way will require incorporating additional flows using the large micro data from the CPS and will be more cumbersome. An advantage of our methodology is that it requires only macro data that is publicly available at quarterly frequency as far back as 1948.[16] Moreover, it is not clear whether we would learn more about the driving forces behind the unemployment rate from such an experiment. Shimer (2007) presents evidence that the aggregate job finding rate is

[16] Using CPS micro files to redo this exercise is not possible for pre-1967, at least to our knowledge.

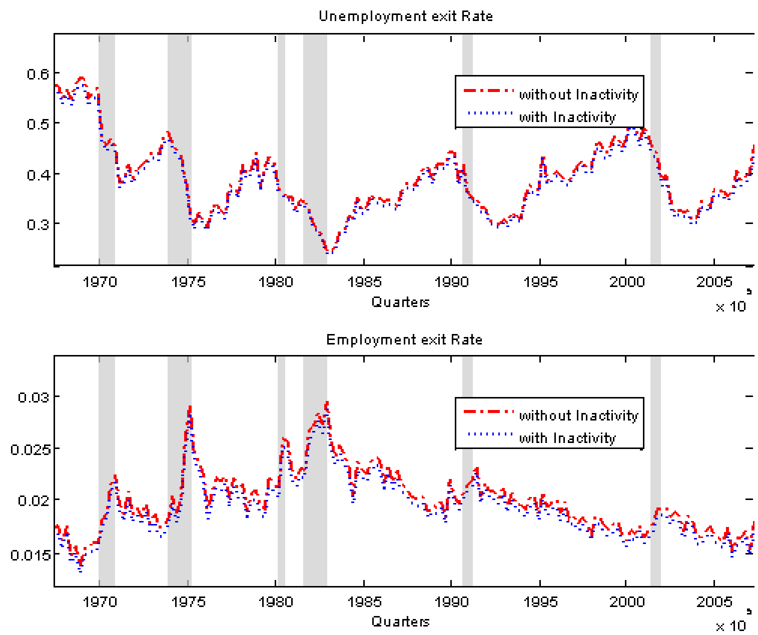

Figure 11: Data is available at http://sites.google.com/site/robertshimer/research/flows.

almost entirely driven by flows from unemployment to employment (at least in the aggregate). Similarly, separation rates closely follow the flow rate from employment to unemployment. Figure 11 presents employment and unemployment exit hazards calculated by Shimer (2007) using CPS. Unemployment exit hazard without inactivity corresponds to the estimate of the job finding rate we used in our model. Similarly, employment exit rate without inactivity is conceptually same as the separation rate we use. Figure 11 confirms our conjecture that, allowing for transitions in and out of inactivity will not change our results, as the implied worker flow hazard rates we need will match the longer time-series we used in our estimation.

7 Conclusion

We present a simple reduced-form model of comovements in real activity and worker flows in this paper and use it to uncover the trend changes in these flows, which determine the trend in the unemployment rate. We argue that this approach provides us with an empirically useful measure of the unemployment rate trend. We use this framework to show that the unemployment rate trend has been relatively stable in the last decade, even after the most recent recession. We also present several numerical exercises where one can address interesting questions regarding unemployment rate dynamics both at low and high frequency.

Our results also suggest that worker reallocation, measured by sum of the job-finding rate and the separation rate, has experienced a trend decline since 2000. This slow worker reallocation has important implications about the dynamics of the unemployment rate, predicting a much slower decline in the near term than would have been possible with high churning, which was a feature of U.S. labor markets before. The estimated trend for the unemployment rate is very robust to labor force movements and use of different filters as long as one utilizes the information on the underlying worker flow rates. It also appears to be at least as precise as an estimate from a similar unobserved components model, where one uses only unemployment rate data, but requires significantly less revisions to past estimates with additional data. Hence, our reduced form approach has many desirable features of an empirically useful measure of the unemployment rate trend and incorporates worker flow rates to impute this rate, as the standard labor market search models would predict. Understanding the actual structural changes that might have led to the observed changes in the trends of worker flows, thereby the implied unemployment rate trend, should be the logical next step for future research. Without an understanding of these structural forces, any policy conclusions based on the estimates from our reduced form model would be misleading and premature.[17]

[17] See, for example, Lucas (1978).

References

[1] Abraham, Katharine and Robert Shimer. 2001. "Changes in Unemployment Duration and Labor-Force Attachment," in *The Roaring '90s: Can Full Employment be Sustained?* ed. Alan B. Krueger and Robert Solow, 3–60. New York: Russell Sage Foundation and Century Foundation.

[2] Ball, Laurence and N. Gregory Mankiw. 2002. "The NAIRU in Theory and Practice." *Journal of Economic Perspectives*, 16(4): 115–36

[3] Barnichon, Regis and Andrew Figura. 2010. "What Drives Movements in the Unemployment Rate? A Decomposition of the Beveridge Curve" *Finance and Economics Discussion Series*, 2010-48. Board of Governors of the Federal Reserve System.

[4] Barro, Robert J. 1988. "The Persistence of Unemployment." *American Economic Review: Papers and Proceedings*, 78(2): 32–7.

[5] Blanchard, Olivier. 2006. "European Unemployment: The Evolution of Facts and Ideas." *Economic Policy*, Centre for Economic Policy Research, 21(45): 5–69.

[6] Clark, Peter K. 1987. "The Cyclical Component of U.S. Economic Activity." *Quarterly Journal of Economics*, 102(4): 797–814.

[7] Clark, Peter K. 1989. "Trend Reversal in Real Output and Unemployment." *Journal of Econometrics*, 40: 15–32.

[8] Daly, Mary and Bart Hobjin. "Okun's Law and the Unemployment Surprise of 2009." *Economic Letter*, Federal Reserve Bank of San Francisco, No. 2010–07.

[9] Darby, Michael R., John C. Haltiwanger, and Mark W. Plant. 1985. "Unemployment Rate Dynamics and Persistent Unemployment under Rational Expectations." *American Economic Review*, 75(4):614–37.

[10] Darby, Michael R., John C. Haltiwanger, and Mark W. Plant. 1986. "The Ins and Outs of Unemployment: The Ins Win." Working Paper 1997, National Bureau of Economic Research.

[11] Davis, Steven, Jason Faberman, John Haltiwanger, Ron Jarmin, and Javier Miranda. 2010. "Business Volatility, Job Destruction, and Unemployment." *American Economic Journal: Macroeconomics*, forthcoming.

[12] Dickens, William T. 2009. "A New Method for Estimating Time-Variation in the NAIRU," in *Understanding Inflation and the Implications for Monetary Policy: A Phillips Curve Retrospective*, ed. Jeff Fuhrer, Jane Sneddon Little, Yolanda K. Kodrzycki, and Giovanni P. Olivei, 207–30. Cambridge, MA: MIT Press.

[13] Elsby, Michael, Ryan Michaels, and Gary Solon. 2009. "The Ins and Outs of Cyclical Unemployment." *American Economic Journal: Macroeconomics*, 1(1): 84–110.

[14] Faberman, Jason. 2008. "Job Flows, Jobless Recoveries, and the Great Moderation." Working Paper 08-11, Federal Reserve Bank of Philadelphia.

[15] Friedman, Milton. 1968. "The Role of Monetary Policy." *American Economic Review*, 58(1): 1–17.

[16] Fujita, Shigeru and Gary Ramey. 2009. "The Cyclicality of Separation and Job Finding Rates." *International Economic Review*, 50(2): 415–30.

[17] Gordon, David M. 1988. "The Un-Natural Rate of Unemployment: An Econometric Critique of the NAIRU Hypothesis." *American Economic Review: Papers and Proceedings*, 78(2): 117–23.

[18] Gordon, Robert. 1997. "The Time-Varying NAIRU and its Implications for Economic Policy." *Journal of Economic Perspectives*, Winter, 11(1): 11–32.

[19] Gordon, Robert. 2010a. "Okun's Law, Productivity Innovations, and Conundrums in Business-cycle Dating." *American Economic Review: Papers and Proceedings*, 100(2): 1–9.

[20] Gordon, Robert. 2010b. "The Demise of Okun's Law and of Procyclical Fluctuations in Conventional and Unconventional Measures of Productivity." Working Paper, Northwestern University.

[21] Hall, Robert E. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review*, 95(1): 50–65.

[22] Johnson, G.E., and P. R. G. Layard. 1986. "The Natural Rate of Unemployment: Explanation and Policy," in *Handbook of Labor Economics*, vol. 2, ed. Orley Ashenfelter and Richard Layard, 92–1000. Amsterdam: North-Holland.

[23] Kim, Chang-Jin and Charles R. Nelson. 1999. *State-Space Models with Regime Switching.* Cambridge, MA: MIT Press.

[24] King, Robert G. and Mark W. Watson. 1994. "The Post-War U.S. Phillips Curve: A Revisionist Econometric History." *Carnegie-Rochester Conference Series on Public Policy*, 41: 157–219.

[25] Laubach, Thomas and John C. Williams. 2003. "Measuring the Natural Rate of Interest." *Review of Economics and Statistics*, 85(4): 1063–70.

[26] Lucas, Robert E. 1978. "Unemployment Policy." *American Economic Review: Papers and Proceedings*, 68(2): 353–357.

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

[28] Mortensen, Dale and Christopher A. Pissarides. 1999. "New Developments in Models of Search in the Labor Market," in *Handbook of Labor Economics*, vol. 3B, ed. Orley Ashenfelter and David Card, 256–2627. Amsterdam: North-Holland.

[29] Orphanides, Athanasios and John C. Williams. 2002. "Robust Monetary Policy Rules with Unknown Natural Rates," *Brookings Papers on Economic Activity*, 2: 63–118.

[30] Phelps, Edmund. 1968. "Money-Wage Dynamics and Labor Market Equilibrium." *Journal of Political Economy*, 76(2): 678–711.

[31] Pissarides, Christopher A. 2000. *Equilibrium Unemployment Theory.* Cambridge, MA: MIT Press.

[32] Polivka, Anne E. and Stephen M. Miller. 1998. "The CPS after the Redesign: Refocusing the Economic Lens," in *Labor Statistics Measurement Issues*, ed. John Haltiwanger, Marilyn E. Manser, and Robert Topel, 249–86. Chicago, IL: University of Chicago Press.

[33] Rogerson, Richard. 1997. "Theory Ahead of Language in the Economics of Unemployment." *Journal of Economic Perspectives*, Winter 11(1): 73–92.

[34] Shimer, Robert. 2005. "The Cyclical Behavior of Unemployment and Vacancies: Evidence and Theory." *American Economic Review*, 95(1): 25–49.

[35] Shimer, Robert. 2007. "Reassessing the Ins and Outs of Unemployment." Working Paper, University of Chicago.

[36] Staiger, Douglas, James H. Stock, and Mark W. Watson. 1997. "The NAIRU, Unemployment, and Monetary Policy." *Journal of Economic Perspectives*, Winter 11(1): 33–49.

[37] Staiger, Douglas, James H. Stock, and Mark W. Watson. 2001. "Prices, Wages and the U.S. NAIRU in the 1990s," in *The Roaring '90s: Can Full Employment be Sustained?* ed. Alan B. Krueger and Robert Solow, 3–60. New York: Russell Sage Foundation and Century Foundation.

[38] Weidner, Justin and John C. Williams. 2011. "What is the New Normal Unemployment Rate?" *Economic Letter*, Federal Reserve Bank of San Francisco, No. 2011–05.