Using Labor Force Flows to Forecast the Labor Market

Regis Barnichon
CREI and Universitat Pompeu Fabra

Christopher J. Nekarda
Federal Reserve Board of Governors

January 2012

Abstract

This paper presents a simple non-linear forecasting model of unemployment based on labor force flows data that, in real time, dramatically outperforms the Survey of Professional Forecasters, the Federal Reserve Board's Greenbook forecast, and basic time series models for short-term forecasts. Our model reduces the mean squared error of the best forecast by 30 to 40 percent. The model also does a good job at identifying turning points several quarters ahead of other forecasters and models. Further, because our model uses information on worker flows typically ignored by other approaches, a combined forecast including our model and the Greenbook forecast yields improvement of about 50 percent for same-quarter forecast, 40 percent for next quarter forecast, and even slight improvements at longer horizons.

The views in this paper are those of the authors and do not necessarily represent the views or policies of the Board of Governors of the Federal Reserve System or its staff. We would like to thank Wouter den Haan, Peter Tulip, Tara Sinclair, Herman Stekler, and Paolo Surico.

1 Introduction

Forecasting the unemployment rate is one of the most important and difficult tasks confronting policymakers. Moreover, forecasting the jobless rate is particularly important surrounding economic downturns. Despite decades of research on the topic, policy makers often rely on Okun's law—the empirical relationship between output growth and unemployment changes—or simple time series models to forecast unemployment.

This paper presents a simple nonlinear forecasting model of unemployment based on labor force flows data. We exploit the tight relationship between the unemployment rate derived from underlying labor force flows, the so-called steady-state unemployment rate (SSUR), and the actual unemployment rate. Because the actual unemployment rate typically converges to its steady-state rate with a short lag, the SSUR provides information about the future unemployment rate. Thus, we can improve our forecast of the unemployment rate by incorporating information from the underlying labor force flows.

Our model dramatically outperforms the Survey of Professional Forecasters (SPF), the Federal Reserve Board's Greenbook forecast, and standard times series models for short-term forecasts, reducing the mean-squared error (MSE) of the best model by 30 to 40 percent. The model also does a good job at identifying turning points several quarters ahead of alternative models. Further, because our model forecasts are based on information not incorporated by other forecasts, a combined forecast including our model and the Greenbook forecast yields an improvement in MSE of about 50 percent for the current-quarter forecast, 40 percent for the next-quarter forecast, and even slight improvements at longer horizons.

Our forecasting model is built on two elements: a (non-linear) law of motion describing how the unemployment rate converges to its steady-state value—the unemployment rate for which the flows of workers in and out of unemployment are balanced—and a forecast of these worker flows. In turn, the model's performance stems from two principles: (i) unemployment converges to its steady-state value in about two quarters, (ii) we emphasize the flows into and out of unemployment rather than the stock.

First, our model exploits the relationship between the unemployment rate and the steady-state value implied by the flows. When the flow of workers into and out of unemployment are not balanced, the actual unemployment rate diverges from its steady-state and adjusts until balance among the flows is restored. As a result, the actual unemployment rate converges towards its steady-state value. Because the unemployment rate usually converges to its steady state with a lag, the SSUR provides information about the actual unemployment rate in the near future. As shown in Figure 1, the SSUR leads actual unemployment by about one quarter. By using information on current labor force flows, we can exploit this convergence mechanism and improve near-term forecasts of the unemployment rate.

However, relying solely on current labor force flows constrains our approach to near-term forecasts, because the steady state to which the actual unemployment rate converges also evolves over time when worker flows change. Thus, we forecast worker flows using a time-series model and feed those forecasts into our law of motion to generate unemployment forecasts at longer horizons. Directly forecasting the flows into and out of unemployment rather than the unemployment stock itself, as is customary, is the second reason why our model outperforms other approaches. It allows the model to better capture the dynamics of unemployment, because the unemployment stock is driven by flows with different time series properties, and because the contribution of the different flows changes throughout the cycle.[1]

In particular, the outflows are the main driving force of unemployment in normal times, but the inflows are responsible for the rare but large fluctuations in unemployment at the onset of recessions.[2] Moreover, the long-run trend in unemployment is driven mainly by a trend in the inflow rate, whereas the outflow rate has no obvious trend.[3]

While a model of the stock can capture the average time series properties of the stock, it cannot allow for different time series properties at different stages of the cycle. Pushing further in that direction, we also show that disaggregating worker flows by sex and age generate even better forecasts (reducing the MSE by another 10 percent for next-quarter forecasts), notably because women and men flows can exhibit very different properties.[4]

A second advantage of focusing on the flows is that it allows us to capture the asymmetric nature of unemployment movements. Inflows to unemployment are responsible for the steepness asymmetry of unemployment—increases are steeper than decreases.[5] By incorporating information on inflows, our model can better capture the asymmetric nature of unemployment than standard models and address one of the criticisms raised by Montgomery et al. (1998) of standard (symmetric) unemployment forecasting models.[6] While our model is not explicitly asymmetric, it relies on the underlying labor market flows that are responsible for the asymmetry of unemployment. By using such information as inputs in the forecasts, our model can incorporate the impulses responsible for the asymmetry of unemployment.[7]

One final benefit on focusing on the flows is the fact that the inflows lead the outflows by about a quarter.[8] As a result, our model does a good job at identifying turning points several quarters ahead of other forecasters and models. Indeed, because a turning point in the inflow

---

1. Barnichon (2012).

2. At a quarterly frequency, the autocorrelation of the outflow rate is 0.91, but the inflow rate is only 0.73 (Shimer, 2007). Further, while the distribution of the (detrended) inflow rate is positively skewed and highly kurtotic, the distribution of the (detrended) outflow rate exhibits no skewness and low kurtosis (Barnichon, 2012).

3. Barnichon and Figura (2010).

4. Elsby, Hobijn and Şahin (2010) and Barnichon and Figura (2010).

5. Barnichon (2010). See Mitchell (1927), Neftçi (1984), DeLong and Summers (1986), and Sichel (2007) for discussion of the steepness asymmetry of unemployment.

6. Because of unemployment asymmetry, forecast errors from linear models are particularly large during economic downturns and around business cycle turning points, when accurate forecasts are extremely valuable for central banks and policy makers. This observation led Montgomery et al. (1998) to advocate nonlinear models of the unemployment rate.

7. And unlike standard time series models used to capture asymmetries (such as threshold autoregressive models), our model does not rely on an arbitrary threshold to introduce asymmetry.

8. Fujita and Ramey (2009).

rate typically signals a turning point in unemployment several quarters in advance, our model can better predict turning points. This is a significant improvement compared to the consensus forecast, the Greenbook, or time series models—all of which typically miss turning points during contractionary periods.[9]

Our model is a useful addition to the set of forecasting models because our approach uses information on worker flows typically ignored by standard approaches. A new combined forecast between the model's forecasts and the most successful historical forecasts (the Greenbook forecasts) yields a reduction in MSE of about 40 percent for same to next quarter forecasts, 20 percent for two-quarter ahead forecasts, and even slight improvements at longer horizons.

Finally, by considering labor flows into and out of the labor force, our method can be easily extended to forecast the labor force participation rate, and also bring considerable improvements to forecasting models of labor force participation.

This paper builds on the influential work of Montgomery et al. (1998) and extend the growing literature aimed at improving the performances of unemployment forecasting models.[10] We particularly draw on the recent literature on labor market flows, which has typically been overlooked by the forecasting literature, but has been the subject of numerous studies aimed at understanding the determinants of labor market fluctuations.[11] In this context, an important advantage of our forecasting method is that it can be easily implemented by forecasters and policy makers for a wide range of OECD countries flow which worker flows data can easily be constructed.[12]

The paper is organized as follows. Section 2 presents our forecasting model, first in the simpler case where there are only two labor market outcomes: employed and unemployed. We then present in section 2.2 a more general model that allows for nonparticipation. Section 3 describes the data we use to estimate our model and generate forecasts. We evaluate the forecasting performance of our model in section 4. Section 5 evaluates the gains of combining forecasts, and section 6 evaluates a forecasting model with flows disaggregated by age and sex. Section 8 concludes.

2 The Steady-State Unemployment Rate

Our forecast is built on two elements: (i) a law of motion describing how the unemployment rate converges to its steady-state value, and (ii) a forecast of the worker flows determining steady-state unemployment and the speed at which actual unemployment converges to steady state. We first present a model with only two labor market states and then expand it to the more

---

[9] Montgomery et al. (1998).

[10] See, for example, Rothman (1998), Golan and Perloff (2004), Brown and Moshiri (2004), and Milas and Rothman (2008).

[11] See, e.g., Shimer (2007); Elsby, Michaels and Solon (2009); Nekarda (2009); Barnichon (2012); and Elsby, Hobijn and Şahin (2011), among others.

[12] Elsby, Hobijn and Şahin (2011).

general case with three labor market states.

2.1 The Labor Market with Two States

We first develop a model with only two labor market states: employed and unemployed. That is, we explicitly assume that there are no movements into and out of the labor force. In addition to providing a simpler framework for understanding the basic flow-based accounting of the SSUR, the two-state model can also be estimated over a long period using duration data. Moreover, because the data used in the two-state version of model are less noisy than in the three-state model described in section 2.2, the two-state model's forecasts are better.

2.1.1 The Law of Motion for Unemployment

Denote $u_{t+\tau}$ the unemployment rate at instant $t+\tau \in \mathbb{R}_+$ with $t \in \mathbb{N}$ and $\tau \in [0,1]$. Assume that during a "period t" of one month—i.e. $\tau \in [0,1]$—all unemployed workers find a job according to a Poisson process with constant arrival rate f_t and all employed workers lose their job according to a Poisson process with constant arrival rate s_t. The unemployment rate then evolves according to

$$\frac{\mathrm{d}u_{t+\tau}}{\mathrm{d}\tau} = s_t\left(1-u_{t+\tau}\right) - f_t u_{t+\tau}. \tag{1}$$

Assuming that the flow hazard rates are constant within time periods, solving equation 1 yields

$$u_{t+\tau} = \beta_t(\tau)u_t^* + \left[1-\beta_t(\tau)\right]u_t \quad \forall \quad \tau \in \mathbb{R}^+, \tag{2}$$

where

$$u_t^* \equiv \frac{s_t}{s_t+f_t} \tag{3}$$

denotes the steady-state unemployment rate, and

$$\beta_t(\tau) \equiv 1 - \mathrm{e}^{-\tau\left(s_t+f_t\right)} \tag{4}$$

is the rate of convergence to steady-state. Since the unemployment outflow rate is much larger than the unemployment inflow rate, we can approximate the convergence speed by $\beta_t(\tau) \approx 1-\mathrm{e}^{-\tau f_t}$.

Using equation 2, we can forecast the unemployment rate one month hence from

$$u_{t+1} = \beta_t u_t^* + (1-\beta_t)u_t, \tag{5}$$

where $\beta_t = 1-\mathrm{e}^{-f_t}$.

Figure 1. Unemployment Rate and Steady-State Unemployment Rate

Percent

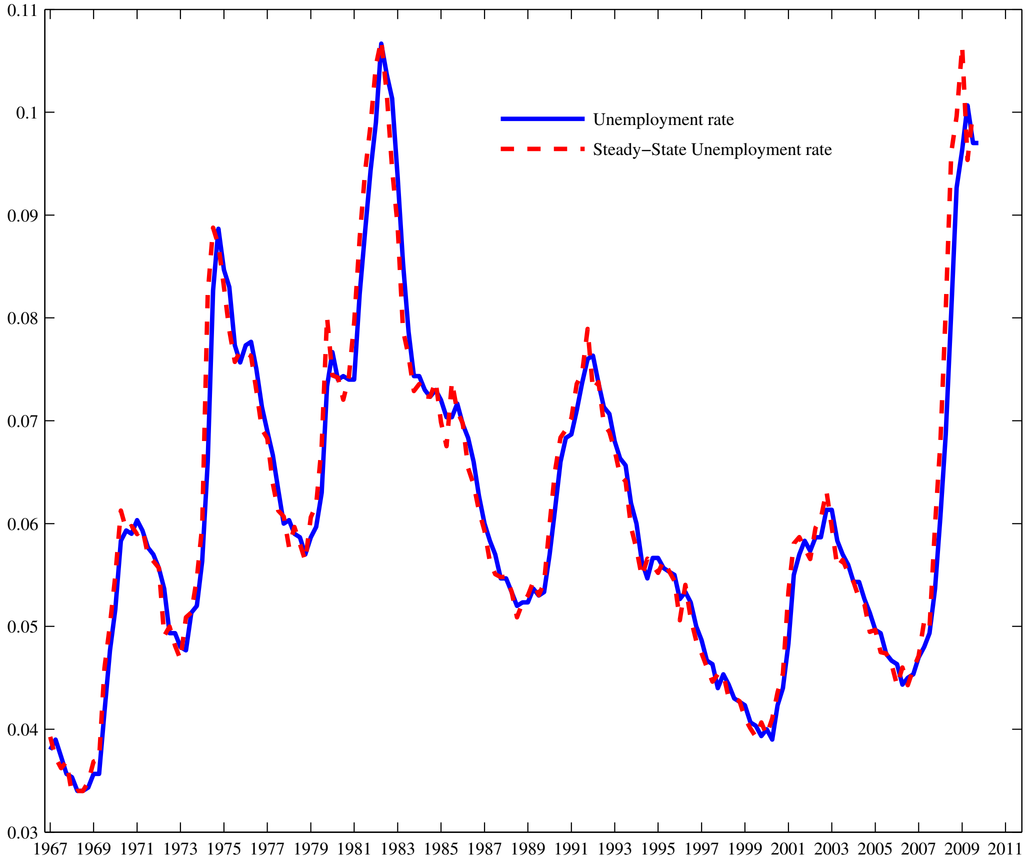

Source: BLS data and authors' calculations.

Notes: Plots unemployment rate and steady-state unemployment rate given by equation 3.

Since the monthly unemployment outflow rate (f) averaged 0.45 over 1948–2007, the half-life deviation of unemployment from its steady-state is about a month and a half.[13] As a result, unemployment converges to its steady state in about two quarters. Figure 1 confirms this observation and shows that the steady-state unemployment rate leads the actual unemployment rate by about one quarter. This observation forms the basis of our approach to forecasting unemployment with labor market flows.

Because the unemployment rate converges quickly to its steady-state value, using labor market flows improve unemployment rate forecasts in the short-term (current quarter through about two quarters hence) over other models. Moreover, the improvement in forecast per-

---

[13] $(\ln 2)/0.45 \approx 1.5$.

formance does not only come from including labor flows data in the information set; rather, the specific nonlinear law of motion (equation 5) through which those labor flows affect the unemployment rate matters.[14]

2.1.2 Forecasting Labor Market Flows

Since equation 5 only forecasts the unemployment rate one month ahead given current values of s and f, forecasts of s_{t+j} and f_{t+j} are needed to generate unemployment rate forecasts at longer horizons.

A simple approach consists in assuming that the hazard rates change little between months so that $\left(s_{t+1}, f_{t+1}\right) \approx \left(s_t, f_t\right)$. In that case, we can use equation 5 and time t values of s_t and f_t to forecast unemployment T-months ahead from[15]

$$u_{t+T} = \left(1 - \mathrm{e}^{-f_t T}\right) u_t^* + \mathrm{e}^{-f_t T} u_t. \tag{6}$$

If hazard rates movements are persistent enough, this procedure may provide reasonable near-term forecasts. In practice however, this assumption is only likely to be valid in the near-term, thereby limiting the applicability of our approach to only near-term forecasts. Moreover, *in real time*, we do not observe s_t and f_t but only s_{t-1} and f_{t-1}, since, at time t, we can only measure the flows of workers over the period $[t-1, t]$. Thus, to forecast unemployment from equation 5, we need forecasts of s_t and f_t.

A concrete example helps clarify this point. The January employment report (received in early February) provides information on the stock of unemployment in January and the average unemployment inflow and outflow rates between December and January (s_{t-1} and f_{t-1}). Looking back at equation 5, the January employment report allows to measure β_{t-1}, u_{t-1}^* and u_t. Thus, to forecast u_{t+1} (the unemployment rate in February), we need forecasts of f_t and s_t (from January to February).

More generally, consider forecasting $\left(s_{t+j}, f_{t+j}\right)$ for $j \geq 0$. For this, we use a vector autoregression (VAR) with ($\ln s_t$, $\ln f_t$) and two leading indicators of labor market flows: initial claims for unemployment insurance and vacancy posting. In particular, initial claims is a high-frequency indicator of layoffs and has been shown to improve labor market forecasts.[16] Importantly, this approach for forecasting worker flows is relatively simple; a more elaborate model for forecasting the worker flows could further improve our model's forecasts.

After generating forecasts of $\left\{f_{t+j}, s_{t+j}\right\}_{j\geq 0}$, we iterate on equation 5 to obtain j-period-

---

[14] See Table 1 in section 4.

[15] This law of motion forms the basis of Elsby, Hobijn and Şahin (2011) strategy to generalize to out of steady-state dynamics the unemployment decomposition by Shimer (2007); Elsby, Michaels and Solon (2009); and Fujita and Ramey (2009).

[16] Montgomery et al. (1998).

ahead forecasts of unemployment with

$$u_{t+1+j} = \beta_{t+j} u^*_{t+j} + (1 - \beta_{t+j}) u_{t+j} \quad \forall \quad j > 0 \tag{7}$$

with $\beta_{t+j} = 1 - \mathrm{e}^{-(s_{t+j} + f_{t+j})}$.

2.2 The Labor Market with Three States

So far, we have only considered two states of the labor market: inside the unemployment pool or outside. However, individuals outside the unemployment pool can either be employed or out of the labor force. In this section, we consider the behavior of the employed and inactive workers separately.

An important advantage of considering three states is the additional possibility of forecasting the employment-population ratio and the labor force participation rate. In addition, a three-state model can more accurately capture the worker flows taking place in the labor market. For instance, the unemployment inflow rate comprises both job losers/leavers as well as entrants to the labor force. Since these two flows (and fact all six flows) display different time-series properties, a three-state model may produce better forecasts than a two-state model.[17]

To generalize our two-state framework to three states, we need to specify and solve the system of differential equations governing the number of people in unemployment U, employment E or out of the labor force N. Denote λ_t^{ab} the hazard rate of transiting from state $a \in \{E, U, N\}$ to state $b \in \{E, U, N\}$, the number of unemployed, employed and out of the labor force workers satisfies the system

$$\begin{cases} \dot{U}_{t+\tau} = \lambda_t^{EU} E_{t+\tau} + \lambda_t^{NU} N_{t+\tau} - (\lambda_t^{UE} + \lambda_t^{UN}) U_{t+\tau} \\ \dot{E}_{t+\tau} = \lambda_t^{UE} U_{t+\tau} + \lambda_t^{NE} N_{t+\tau} - (\lambda_t^{EU} + \lambda_t^{EN}) E_{t+\tau} \\ \dot{N}_{t+\tau} = \lambda_t^{EN} E_{t+\tau} + \lambda_t^{UN} U_{t+\tau} - (\lambda_t^{NE} + \lambda_t^{NU}) N_{t+\tau}. \end{cases} \tag{8}$$

For instance, as shown in the first line, changes in unemployment are given by the difference between the inflows to unemployment (workers losing or quitting their jobs and persons joining the labor force) and the outflows from unemployment (unemployed persons finding a job or dropping out of the labor force).

The details of the solution for the 3-state system are shown in the appendix. Using the initial and terminal conditions, the one-step ahead forecasts of unemployment, employment and nonparticipation can be solved as functions of the hazard rates (λ^{ab}s) and the steady-state numbers of unemployed (U_t^*), employed (E_t^*), and nonparticipants (N_t^*). These steady-state

---

[17] See Barnichon and Figura (2010).

stocks are given by

$$\begin{aligned} U_t^* &= \frac{P_t}{s_t + f_t + o_t} s_t \\ E_t^* &= \frac{P_t}{s_t + f_t + o_t} f_t \\ N_t^* &= \frac{P_t}{s_t + f_t + o_t} o_t, \end{aligned}$$

where P_t is the working age population at time t and with s_t, f_t and o_t defined by

$$\begin{cases} s_t = \lambda_t^{EN}\lambda_t^{NU} + \lambda_t^{NE}\lambda_t^{EU} + \lambda_t^{NU}\lambda_t^{EU} \\ f_t = \lambda_t^{UN}\lambda_t^{NE} + \lambda_t^{NU}\lambda_t^{UE} + \lambda_t^{NE}\lambda_t^{UE} \\ o_t = \lambda_t^{EU}\lambda_t^{UN} + \lambda_t^{UE}\lambda_t^{EN} + \lambda_t^{UN}\lambda_t^{EN}. \end{cases}$$

We use the solution to these equations (presented in the appendix) to generate one-period-ahead forecasts of the unemployment rate and labor force participation rate from

$$\begin{aligned} u_{t+1} &= \frac{U_{t+1}}{U_{t+1} + E_{t+1}} \\ lfpr_{t+1} &= \frac{U_{t+1} + E_{t+1}}{P_{t+1}}. \end{aligned}$$

Similar to the two-state case, we forecast the six hazard rates from a VAR with the 6 hazard rates and include the leading indicators unemployment insurance claims and vacancy posting. To forecast the population growth rate, we use the middle forecast from the U.S. Census Bureau.

3 Data

Our approach rests on obtaining good measures of worker flows across labor market states.

When we consider the case of only two labor market states, measures of s_t and f_t, the unemployment inflow and outflow rates, are obtained following Shimer (2007) over 1951–2009. Specifically, the unemployment outflow rate is estimated from the share of short-term (less than 5 weeks) unemployed with $f_t = -\ln(1 - F_t)$ and $F_t = 1 - \frac{u_{t+1} - u_{t+1}^{<5wks}}{u_t}$. The separation rate is then estimated by solving equation 1 over $[t, t+1]$ and finding s_t such that the solution $u_{t+\tau}$ equals u_{t+1} for $\tau = 1$.

When we consider all three labor market states, we need measures of the six transition rates. We construct month-to-month labor force transitions from Current Population Survey (CPS) microdata by matching workers across consecutive months.[18] Because the CPS microdata begin only in 1976, we use Joe Ritter's tabulation of the gross flows data from June 1967 to

---

[18] Shimer (2007); Nekarda (2009).

December 1975. Thus, our CPS flows data cover 1967–2009. More recently, the Bureau of Labor Statistics (BLS) started publishing gross flows data on worker transitions covering 1990–2011. Compared to the transitions inferred from the raw CPS microdata, the published series has the advantage to be corrected for margin error and thus of being more reliable and less noisy.[19] However, using the published flows data or the microdata from 1990 onwards makes little difference to the forecasting performances of our model. We thus use the published flows after 1990.

Compared to microdata on worker flows, these measures have the advantage of being obtained from published macro data (unemployment duration). While the matched flows data provide a direct measure of worker flows, the published macro data (unemployment duration) have the advantage of being less noisy, available over a longer time period and easily accessible to any professional forecaster.

Finally, initial unemployment insurance claims are taken from the Department of Labor, and vacancy posting is proxied with the composite help-wanted index presented in Barnichon (2010).[20]

4 Forecasting Performance

We evaluate the performance of our SSUR models by comparing its unemployment rate forecasts with alternative methods along two dimensions: MSE of out-of-sample forecasts and ability to forecast business cycle turning points. The first dimension is a quantitative exercise; the second is more qualitative in nature.

4.1 Evaluation Method

As in Montgomery et al. (1998), we quantify our models' performance by the MSE of out-of-sample forecasts. We consider four alternative forecasts of the unemployment rate. The first is a basic univariate time-series model. Following Montgomery et al., we use an ARIMA(2,0,1) model for the unemployment rate.[21] The second alternative is the median forecast from the SPF. The third alternative is the Greenbook forecast.

---

[19] Owing to sample rotation and temporary absence of individuals, transition information is unavailable for a subset of the sample. This failure to match individual workers across months is referred to as margin error, and it leads to omission of possible transitions from the survey data. To alleviate this problem, the BLS reweights the flow measures in a manner that minimizes the discrepancies between officially reported stocks of workers and the stocks that would be imputed from the flow measures.

[20] This composite index uses the Conference Board print help-wanted index until 1994 to proxy for vacancy posting. After 1994, the composite index controls for the emergence of online advertising (at the expense of print advertising) by combining information from the Conference Board print and online help-wanted advertising indexes with the BLS's Job Openings and Labor Turnover survey.

[21] Montgomery et al. (1998) consider a range of time series models and show that, given the same information set, they are always dominated by the median SPF forecast. We report only one time series model for illustration.

The fourth alternative forecast is a VAR that incorporates the same information set as the SSUR model (unemployment inflow and outflow rates, vacancy posting, initial claims for unemployment insurance, and the unemployment rate). This allows us to evaluate the nonlinear relationship directed by the theory compared to an atheoretical linear time series model using the same information set.

To reflect the actual environment within which forecasters must operate, we estimate our models and generate forecasts using real time data. For example, because worker flows lag the unemployment rate by a period, we forecast the current-period flows. This puts the models on similar footing as the SPF and the Greenbook, which are obviously made in real time.

The SPF sends out the survey questionnaire sometime in the first month of a quarter, and the survey participants are asked to mail back the completed questionnaire by the middle of the second month of the quarter, implying that forecasts included in the SPF incorporate data from the first month of the quarter. To make forecasts comparable, the SSUR model forecasts incorporate data up until (and including) the first month of the current quarter. The ARIMA model is estimated at a quarterly frequency, and the forecast jumps off from the last quarter.[22]

The forecasts are generating from rolling model estimates. That is, for a given forecast origin ($t+0$), the model is estimated using all observations that a forecaster would have had at the origin, and then the model forecasts are generated for $t+0$ through $t+4$. To obtain stable parameters, we estimate the models over 1967–78 and start the rolling forecasts in 1979. We generate 4 forecasts per year when comparing with SPF forecasts or the time series model, or 8 forecasts per year to mimic the forecasting dates of the Greenbook.[23] The VARs used to forecast the hazard rates include 6 lags.

4.2 Forecast Errors

Table 1 reports the mean squared forecast error of quarterly forecasts over a one-year horizon (including a forecast of the current quarter, $t+0$) for several models. Since the SPF has been shown to outperform many time series model forecasts (e.g., Montgomery et al., 1998), we use the SPF forecast as a benchmark, and report the MSE relative to those of the SPF. Because the Federal Reserve Board produces eight Greenbook forecasts per year (two each quarter), to compare with the SPF and the other quarterly forecasts, we take the projection prepared closest to the middle of each quarter.[24]

The SSUR-2 model outperforms the SPF model dramatically for short-term forecasts. As shown in row 2 of table The MSE is reduced by 40 percent for same quarter forecast and by

---

[22] Using instead a monthly ARIMA model jumping off from the first month of the current quarter (consistent with the SPF and SSUR model) generates a relative MSE of 1.59 for same-quarter forecast ($t+0$) and forecasts that are noticeably worse than the quarterly ARIMA model at longer horizons.

[23] The data set of Greenbook forecast errors covers 1979 to 2006 and was compiled by Tulip (2009). The Greenbook forecasts are made public with a five-year lag; thus our comparison with the Greenbook forecast only covers through 2006.

[24] Table 2 provides a comparison to the Greenbook forecast based on 8 forecasts per year.

Table 1. Unemployment Rate Forecast Errors Relative to SPF, 1979–2006

MSE relative to SPF

	Forecast horizon (quarters)				
Model	$t+0$	$t+1$	$t+2$	$t+3$	$t+4$
1. SPF	1.00	1.00	1.00	1.00	1.00
2. SSUR-2	0.60	0.72	0.90	1.01	1.04
3. Greenbook	0.90	0.66	0.62	0.61	0.63
4. ARIMA	–	1.28	1.98	2.55	2.93
5. VAR	0.65	0.87	1.04	1.14	1.15
Memo:					
6. MSE of SPF	0.02	0.11	0.24	0.42	0.63

Source: Authors' calculations using data from BLS, Federal Reserve Board, and Federal Reserve Bank of Philadelphia.

Notes: MSE calculated from 4 forecasts per year. $t+0$ denotes current-quarter forecast. SSUR-2 denotes the steady-state model with 2 labor market states. ARIMA model is a quarterly ARIMA(2,0,1). The VAR includes the unemployment rate, the inflow and outflow rates, initial claims for unemployment insurance and vacancy posting; all variables have four lags.

about 30 percent for one-quarter ahead forecast.[25] At longer forecast horizons, the difference is not as stark or the consensus forecast performs slightly better. This is not surprising given that, unlike the Greenbook or SPF forecasts based on an array of economic data and models of the U.S. economy and its different sector, SSUR-2 is a simple statistical model that incorporates no information other than labor market data. The univariate ARIMA model performs markedly worse.

Row 5 of table 1 reports the performance of a VAR that incorporates the same information set as SSUR-2 (the inflow and outflow rates, vacancy posting, unemployment insurance claims and the unemployment rate). The results show that nonlinearity is an important feature of our model, as the VAR performs worse than SSUR-2 at all forecast horizons.

The real test of our forecasting models, however, comes from with a comparison with the Greenbook forecasts, that generate substantially smaller MSE than SPF forecasts (Table 1). Since there are 8 Greenbook forecasts a year that are performed at different stages of the quarter—some forecasts are performed with no monthly data for the current quarter, while others with two months of data for the current quarter, we simulate a series forecasts over 1979–2006 that mimic data availability for each Greenbook date, i.e., taking as inputs *only* the data that were available at the time of the forecast.

Table 2 confirms the results of Table 1, with SSUR-2 reducing the MSE dramatically for same

---

[25] The improvements are statistically significant at the 1 percent margin. To formally test that the variance of the new forecasts is significantly less that of our benchmark models, we followed Montgomery et al. (1998) and calculated the p-values of the Meese-Rogoff test.

Table 2. Unemployment Rate Forecast Errors Relative to Greenbook, 1979–2006

MSE relative to Greenbook

	Forecast horizon (quarters)				
Model	$t+0$	$t+1$	$t+2$	$t+3$	$t+4$
1. Greenbook	1.00	1.00	1.00	1.00	1.00
2. SSUR-2	0.52	0.81	1.12	1.28	1.36
3. SSUR-3	0.69	0.95	1.25	1.40	1.45
Memo:					
4. MSE of Greenbook	0.03	0.13	0.30	0.54	0.84

Source: Authors' calculations using data from BLS, Federal Reserve Board, and Federal Reserve Bank of Philadelphia.

Notes: MSE calculated from 8 forecasts per year. $t+0$ denotes the current quarter forecast. SSUR-2 is the steady-state model with employment and unemployment; SSUR-3 is the steady-state model with employment, unemployment, and nonparticipation. "Greenbook + SSUR" is optimal combination; weight on Greenbook forecast shown in parentheses.

quarter—about 50 percent—and one quarter ahead—about 20 percent—forecasts. Table 2 also evaluates the performance of SSUR-3, based on a three labor market state representation of worker flows and using matched CPS microdata. While SSUR-3 also significantly outperforms the Greenbook (and thus the other models considered in Table 1) at short term horizon, it performs less well than SSUR-2. This is perhaps not surprising given that SSUR-3 forecasts rely on microdata that are noisier than the aggregate duration data used by SSUR-2.

4.3 Business Cycle Turning Points

To study the performances of the SSUR model during contractionary periods and before turning points, Figures 2 and 3 plot the forecasts from SSUR-2 and the Greenbook over the last four recessions. The figures illustrate the evolution of the forecasts as the recessions gain in momentum. Specifically, for each recession, we consider two different jumping-off points. The first point, shown in Figure 2, is at the nascent onset of a recession—roughly before any significant increase in unemployment. The second point, shown in Figure 3, corresponds to the peak of the unemployment inflow rate, which is about half-way through the increase in the unemployment rate.

Jumping off from the very early stages of a recession, the SSUR model does not perform significantly better than the Greenbook at horizons of a year or more. Except for the early 1980s recession, the impetus from the inflow rate is too small, and the SSUR model understates the increases in unemployment.

However, once the unemployment rate has risen, the SSUR model is better at identifying turning points. When jumping off roughly mid-way through the increase in unemployment, the

Figure 2. Unemployment Rate Forecasts at the Onset of Recessions

Percent

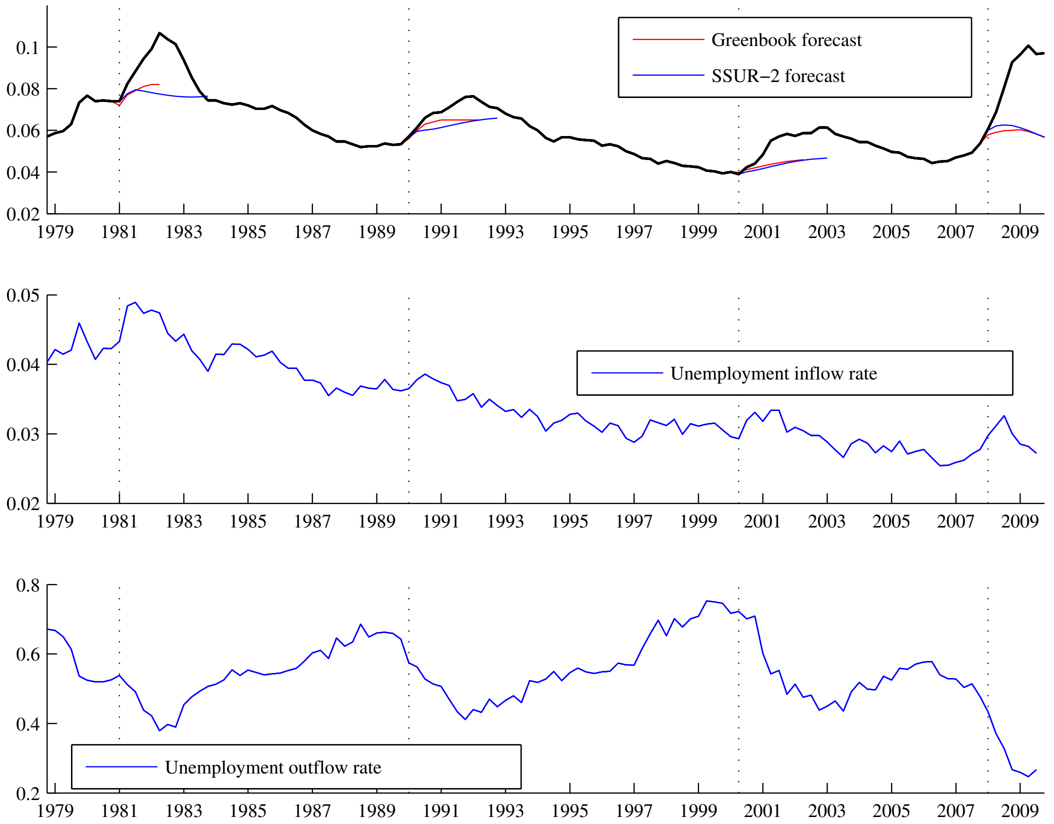

Source: BLS and Federal Reserve Bank of Philadelphia data and authors' calculations.

Notes: Dashed vertical lines indicate forecast jump-off dates: 1981m9, 1990m9, 2000m11, and 2008m8. Because Greenbook forecasts after 2006 are not yet public, the 2008m8 forecast is from the SPF.

SSUR model clearly outperforms the Greenbook forecast and is able to predict the turning point in unemployment for the last four recessions as far a year in advance. Indeed, at each of the dates picked (and especially in 1982 and 1990), the Federal Reserve Board staff had projected that the unemployment rate was near its peak for the cycle, when it would continue rising for some time. In contrast, the SSUR model predicted that the turning point would occur much later, often up to a year later. And for all four recessions, the model's predicted turning point was close to the actual turning point. The SSUR forecast was also much closer to the actual path of unemployment than any of the Greenbook forecasts.

To get some intuition for the model's performance, the middle and lower panels of Figures 2 and 3 plot the behavior of the unemployment inflow and outflow rates around the jumping-off

Figure 3. Unemployment Rate Forecasts Mid-Way through Recessions

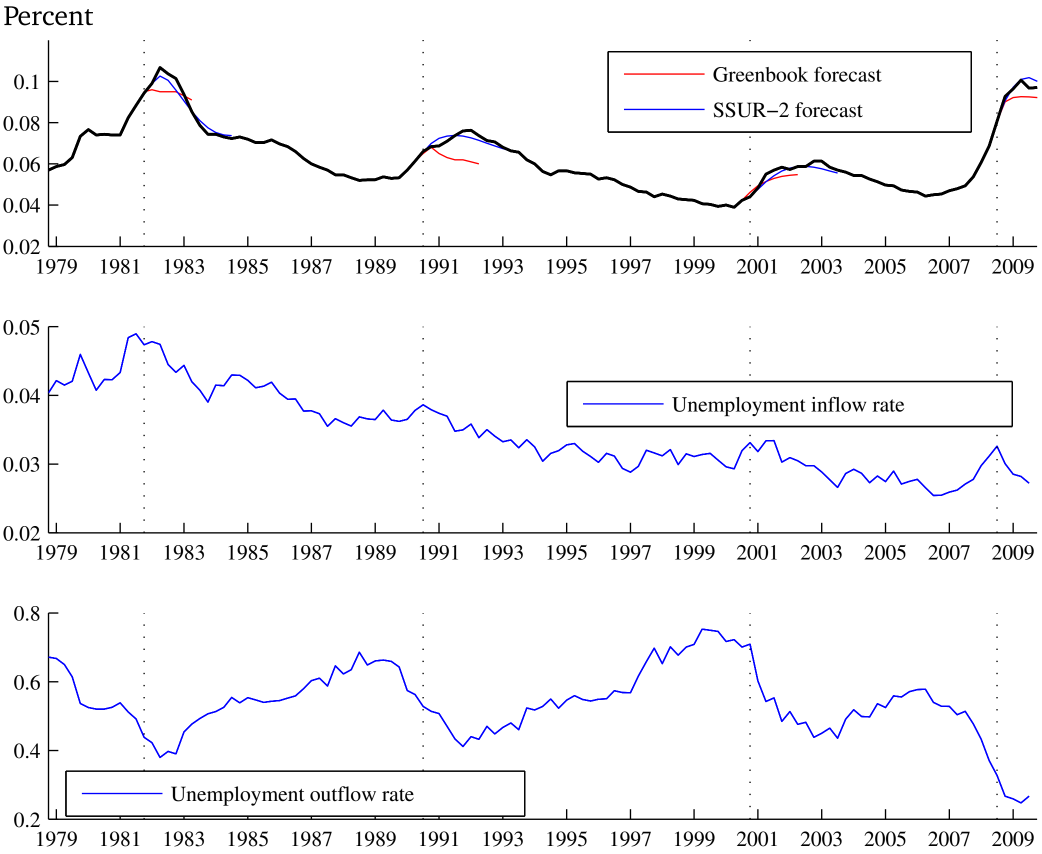

Source: BLS and Federal Reserve Bank of Philadelphia data and authors' calculations.

Notes: Dashed vertical lines indicate forecast jump-off dates: 1982m6, 1991m3, 2001m5, 2009m1. Because Greenbook forecasts after 2006 are not yet public, the 2009m1 forecast is from the SPF.

points. The inflow rate is most responsible for SSUR model's superior performance. Because the inflow rate leads the unemployment rate and because bursts of separation are responsible for the sharp increases in unemployment at the onset of recessions, incorporating information from the inflow rate allows the model to capture the fast increase in unemployment during recessions.[26] Thus, SSUR correctly predicts a period of increasing unemployment following all four jumping-off points. In contrast, the Greenbook missed or understated the increase in unemployment in all four recessions.[27] In addition, because the turning point in the inflow rate

---

[26] Fujita and Ramey (2009) and Barnichon (2012).

[27] As discussed in Montgomery et al. (1998) and Baghestani (2008), this tendency to understate increases in unemployment during recessions is not only a property of the Greenbook forecast, but is an undesirable feature of the SPF and most models.

signals the turning point in the unemployment rate several quarters (and sometimes as far as a year) in advance, the SSUR model is able to predict the turning point in the unemployment rate with relatively high confidence several quarters in advance.

5 Combining Forecasts

An important question is whether any additional improvements can be made to the forecasts by combining the best available forecast (in this case, the Greenbook forecast) with our SSUR model. To answer these questions, we construct a new combined forecast, which exploits the differences in correlation among the forecast errors.[28]

This combined forecast is constructed by taking a weighted average of the forecasts from two main models. The weights are determined through ordinary least squares regression, with a constant included to account for any systematic biases in the estimate. Given the SSUR models' comparative advantage in near-term forecasts, we estimate weights separately for each forecast horizon. These weights allow us to evaluate the marginal contributions of each model over the Greenbook forecast. If the SSUR model forecast had no incremental benefit over the Greenbook forecast, then the weight on the Greenbook forecast would be 1.

As shown in Table 3, this is not the case, and combining the SSUR model with an Okun's law–type forecast improves forecasting performance significantly from zero to two quarters forecast horizon. Compared to the baseline Greenbook forecast, the reduction in MSE achieved by the combined forecast amounts to about 50 percent for current-quarter forecasts, 25 percent for one-quarter ahead forecasts, 12 percent for two-quarter ahead forecasts, and even 5 percent for three-quarter ahead forecasts. Compared to the SSUR-2 forecast alone, the combined forecast performs significantly better for one- to three-quarter ahead forecasts.

The optimal weights reflect the contribution of the SSUR model for short-term forecasting, and the combined forecast puts a lot more weight on the SSUR model at short-term horizon. Importantly, that the combined forecast performs significantly better than any forecast alone indicates that the SSUR model brings relevant information not contained in the Greenbook forecast. In other words, because the forecast errors of the two models are not strongly correlated, the combined forecast performs substantially better.

6 Disaggregated Models by Sex and Age

Table 1 showed that even a simple linear VAR with labor market flows improves unemployment forecasts compared to other methods that only use information on the unemployment stock. This improvement comes from the fact that the different time series properties and behavior of the flows cannot be captured by stock variables.

---

[28] Granger and Newbold (1986).

Table 3. Unemployment Rate Forecast Errors, Optimal Combined Forecasts, 1979–2006

MSE relative to Greenbook

	Forecast horizon (quarters)				
Model	$t+0$	$t+1$	$t+2$	$t+3$	$t+4$
1. Greenbook	1.00	1.00	1.00	1.00	1.00
2. SSUR-2	0.52	0.81	1.12	1.28	1.36
3. SSUR-2 + Greenbook	0.51	0.74	0.88	0.96	0.98
	(0.18)	(0.33)	(0.58)	(0.74)	(0.83)
4. SSUR-3	0.69	0.95	1.25	1.40	1.45
Memo:					
5. MSE of Greenbook	0.03	0.13	0.30	0.54	0.84

Source: Authors' calculations using data from BLS, Federal Reserve Board, and Federal Reserve Bank of Philadelphia.

Notes: MSE calculated from 8 forecasts per year. $t+0$ denotes the current quarter forecast. SSUR-2 is the steady-state model with employment and unemployment; SSUR-3 is the steady-state model with employment, unemployment, and nonparticipation. "SSUR-n + Greenbook" is optimal combined forecast; weight on Greenbook forecast shown in parentheses.

In this spirit, an additional aspect that may further improve the forecasting performances of the SSUR model is the fact that women and men exhibit flows with very different properties. As discussed in Elsby, Hobijn and Şahin (2010), the unemployment inflow rate is a lot more cyclical for men than for women, with men experiencing much starker increases in job separation during recessions.[29] In addition, while men's inflow rate is relatively stationary, women's inflow rate displays a downward trend that is largely responsible for the trend in unemployment since 1976.[30]

Treating separately men and women may thus allow us to improve forecasting performances. In essence, this amounts to adding degrees of freedom in our model since there will be two sets of hazard rates, one for each sex. Importantly however, the number of observation would also double since it is possible to observe the hazard rates of both sexes.[31]

Considering the model with two labor market states, disaggregating by sex yields 2 laws of motion for unemployment similar to equation 7:

$$u_{t+1+j}^{sex} = \beta_{t+j}^{sex} u_{nt+j}^{*,sex} + (1-\beta_{t+j}^{sex}) u_{t+j}^{sex} \ \forall \ j > 0, sex \in \{m, f\}. \tag{9}$$

---

[29] In contrast, the outflow rates behave similarly. However, inflow rates are the key driver of the SSUR model.

[30] Barnichon and Figura (2010).

[31] The limiting factor of this approach is thus the sample size of the CPS. As we disaggregate hazard rates further, the sampling variance increases and movements in individuals' hazard rates become less informative.

Table 4. Comparison of Disaggregated SSUR Model Forecast Errors, 1989–2009

MSE relative to SSUR-2

	Forecast horizon (quarters)				
Model	$t+0$	$t+1$	$t+2$	$t+3$	$t+4$
1. SSUR-2	1.00	1.00	1.00	1.00	1.00
2. SSUR-2-sex	0.98	0.91	0.89	0.92	0.95
3. SSUR-2-age	1.00	0.98	0.90	0.92	0.95

Source: Authors' calculations using data from BLS, Federal Reserve Board, and Federal Reserve Bank of Philadelphia.

Notes: MSE calculated from 8 forecasts per year. $t+0$ denotes the current quarter forecast. SSUR-2 is the steady-state model with 2 labor market states; SSUR-2-sex denotes the SSUR-2 model used separately by sex. SSUR-2-age denotes the SSUR-2 model used separately for three age groups (young, prime-aged, and older).

The aggregate unemployment rate can be projected by

$$u_{t+1+j} = \omega_{t+1+j} u^{m}_{t+1+j} + (1-\omega_{t+1+j}) u^{f}_{t+1+j} \ \forall\, j > 0, \tag{10}$$

with ω the share of men in the labor force. We use separate VARs including the sex-specific hazard rates and the share of that group in the labor force to forecast β^{sex}_{t+j}, $u^{*,sex}_{t+j}$ and ω_{t+1+j}.

Row 2 of Table 4 evaluates the forecasting performance of the SSUR-2 model with hazard rates disaggregated by sex. Since disaggregated hazard rates are only available since 1976, we estimate the model over 1976–89 and then generate rolling forecasts from 1989 to 2009. Allowing for different flows for men and women reduces the MSE relative to the baseline SSUR-2 model by about 10 percent for one to two-quarters ahead forecasts. Compared to the Greenbook forecast (Table 2), the reduction amounts to about 30 percent (from 0.91 to 0.81) for next-quarter forecasts. Disaggregating by sex also improves the performance for two-quarter ahead forecasts and somewhat at longer horizons. Same-quarter forecasts are only marginally improved, perhaps a negative consequence of the higher noise in the disaggregated hazard rates.

Because the changing age composition of the workforce accounts for a large share of the trend in unemployment in the United States.[32], we also explore disaggregating flows by age groups. We construct worker flows by three broad age groups: young (16–25), prime-age (25–55) and older (55–85). Row 3 of table 4 shows the results of the SSUR model disaggregated by age. As with sex, disaggregation improves the forecasting performance of the SSUR model. The improvement for same- and next-quarter forecasts is less important than in the case of age-specific disaggregation. This is probably a result of the fact that, unlike the sex-specific hazard

---

[32] Barnichon and Figura (2010).

rates, the cyclical behavior of the age-specific hazard rates are similar.[33] At longer forecast horizons (two- to four-quarters ahead), age-disaggregation improves performance significantly, because demographics is an important factor behind unemployment's trend and incorporating age-specific information allows the model to better capture the low-frequency movements in unemployment.

However, a disadvantage of disaggregating by demographics is that it increases the noise (due to the much smaller sample sizes of sex- and age-based flows), thus limited the benefits of disaggregation. Nonetheless, the behaviors of the sex-specific hazard rates are sufficiently different to compensate the loss of information due to higher sampling variance and allow the sex-specific SSUR model to perform even better than the baseline SSUR model.

7 Forecasting Labor Force Participation

Unlike with the unemployment rate, there is less of a systematic aggregate relationship between labor force participation and output growth. In fact, aggregate participation was largely thought to be acyclical over 1960–2006, where changes in the labor force participation rate were only weakly related to output growth.[34] As a result, forecasting the labor force participation rate was often seen as subordinate to forecasting the unemployment rate.

The large and unexpected decline in labor force participation in the 2008–09 recession challenged that conventional wisdom and highlighted the importance of forecasting labor force participation. However, given the absence of a strong relation between output and labor force participation prior to the recession, forecasters have few models to turn to to forecast labor force participation.

An important advantage of our three-state model is that it also generates forecasts of the labor force participation rate (and, by extension, the employment-to-population ratio). Table 5 evaluates the performance of the SSUR-3 model compared to the Greenbook forecasts.[35] SSUR-3 improves upon the Greenbook forecast for same quarter forecast, reducing the MSE by about 20 percent. At longer forecast horizons, SSUR-3 performs less well than the Greenbook.

However, SSUR-3 forecast errors need not be correlated with Greenbook forecast errors, and a combined forecast may generate significant improvements. Row 3 of Table 5 confirms this intuition. An optimal combined forecast between the Greenbook and SSUR-3 performs significantly better at all forecasting horizons considered, reducing the MSE by 40 percent for same-quarter forecasts, and by about 20 percent for one- to four-quarter ahead forecasts. The

---

[33] Elsby, Hobijn and Şahin (2010).

[34] Nonetheless several papers identified the importance of demographics for the aggregate participation rate. Aaronson et al. (2006) and Fallick and Pingle (2006) use cohort-based models to help isolate demographic and other structural factors from cyclical variation in the participation rate. They find that the apparent acyclicality of aggregate participation is the result of moderately cyclical participation among certain demographic groups that roughly offsets when aggregated.

[35] The historical Greenbook forecasts contain quarterly forecasts for the participation rate beginning only in 2000.

Table 5. LFPR Forecast Errors Relative to Greenbook, 2000–06

MSE relative to Greenbook

	Forecast horizon (quarters)				
Model	$t+0$	$t+1$	$t+2$	$t+3$	$t+4$
1. Greenbook	1.00	1.00	1.00	1.00	1.00
2. SSUR-3	0.82	1.48	1.83	1.93	1.75
3. SSUR-3 + Greenbook	0.59	0.76	0.82	0.83	0.78
	(0.42)	(0.64)	(0.71)	(0.73)	(0.68)
Memo:					
4. MSE of Greenbook	0.02	0.04	0.07	0.10	0.14

Source: Authors' calculations using data from BLS, Federal Reserve Board, and Federal Reserve Bank of Philadelphia.

Notes: MSE calculated from 8 forecasts per year. $t+0$ denotes the current quarter forecast. SSUR-3 is the steady-state model with employment, unemployment, and nonparticipation. "Greenbook + SSUR-3" is optimal combination; weight on Greenbook forecast shown in parentheses.

higher weight on SSUR-3 in the near-term reflects the superior performances of SSUR-3 for same quarter forecast, but even at longer horizons, the weights on SSUR-3 are nontrivial at 30 to 40 percent and reflect the contribution of SSUR-3 to the improved performances.

8 Conclusion

Although the unemployment rate is typically not considered a leading or coincident indicator, increases in the unemployment rate have preceded the last three recessions.[36] And recent research by Fleischman and Roberts (2011) find that the unemployment rate provides the best single signal about the state of the business cycle in real time. Nevertheless, despite extensive research on the topic, policy makers often rely on Okun's law or simple time series models to forecast the unemployment rate.

This paper presents a nonlinear model for forecasting the unemployment rate based on labor force flows that, in real time, dramatically outperforms basic times series models, the SPF, and the Federal Reserve Board's Greenbook forecast at short horizons. The model, and its performance, derive from two principles: (i) the unemployment rate converges to its steady-state value in about two quarters according to a nonlinear law of motion, and (ii) individual labor market flows have different time series properties.

Empirically, the model has a mean squared forecast error about 40 percent less than the next-best forecast for same-quarter and next-quarter forecasts. Our model also does a good job at identifying turning points several quarters ahead of others models and forecasters. In

---

[36] Stekler (2010).

addition, because the model brings new information to the forecast, a combined forecast including our model and the Greenbook forecast yields improvement of about to 50 percent for same-quarter forecast, 40 percent for next quarter forecast, 20 percent for two-quarter ahead forecasts, and even slight improvements at longer horizons.

In the case with two labor market states, the forecasting model is easy to implement, and the data are readily available. In this respect, the recent work by Elsby, Hobijn and Şahin (2011) is of great importance for forecasters and policy makers interested in forecasting the labor market outside of the United States. The authors constructed unemployment hazard rates across a 14 OECD countries. While the data are compiled at an annual frequency, the flow hazard rates necessary for our model can also be measured at a higher frequency for many of those countries. Thus, leveraging the Elsby, Hobijn and Şahin data, our model could yield better labor market forecasts in those countries as well.

References

Aaronson, Stephanie, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William Wascher. 2006. "The Recent Decline in the Labor Force Participation Rate and Its Implications for Potential Labor Supply." *Brookings Papers on Economic Activity*, 37(1): 169–54.

Baghestani, Hamid. 2008. "Federal Reserve versus Private Information: Who is the Best Unemployment Rate Predictor?" *Journal of Policy Modeling*, 30(1): 101–10.

Barnichon, Regis. 2010. "Building a Composite Help-Wanted Index." *Economics Letters*, 109(3): 175–78.

Barnichon, Regis. 2012. "Vacancy Posting, Job Separation, and Unemployment Fluctuations." *Journal of Economic Dynamics and Control*, 36(3): 315–30.

Barnichon, Regis, and Andrew Figura. 2010. "What Drives Movements in the Unemployment Rate?" Federal Reserve Board of Governors Finance and Economics Discussion Series paper 2010-48.

Brown, Laura, and Saeed Moshiri. 2004. "Unemployment Variation over the Business Cycles: A Comparison of Forecasting Models." *Journal of Forecasting,* 23(7): 497–511.

DeLong, J. Bradford, and Lawrence H. Summers. 1986. "Are Business Cycles Symmetric?" In *The American Business Cycle, Continuity and Changes.* , ed. Robert J. Gordon. University of Chicago Press.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2010. "The Labor Market in the Great Recession." *Brookings Papers on Economic Activity*, 41(1): 1–69.

Elsby, Michael W. L., Bart Hobijn, and Ayşegül Şahin. 2011. "Unemployment Dynamics in the OECD." Federal Reserve Bank of San Francisco Working paper.

Elsby, Michael W. L., Ryan Michaels, and Gary Solon. 2009. “The Ins and Outs of Cyclical Unemployment.” *American Economic Journal: Macroeconomics*, 1(1): 84–110.

Fallick, Bruce, and Jonathan Pingle. 2006. “A Cohort-Based Model of Labor Force Participation.” Board of Governors of the Federal Reserve System Finance and Economics Discussion Series paper 2007-09.

Fleischman, Charles A., and John M. Roberts. 2011. “From Many Series, One Cycle: Improved Estimates of the Business Cycle from a Multivariate Unobserved Components Model.” Board of Governors of the Federal Reserve System Finance and Economics Discussion Series paper 2011-46.

Fujita, Shigeru, and Garey Ramey. 2009. “The Cyclicality Of Separation And Job Finding Rates.” *International Economic Review*, 50(2): 415–30.

Golan, Amos, and Jeffrey M. Perloff. 2004. “Superior Forecasts of the U.S. Unemployment Rate Using a Nonparametric Method.” *The Review of Economics and Statistics*, 86(1): 433–38.

Granger, Clive W., and Paul Newbold. 1986. *Forecasting Economic Time Series.* New York:Academic Press.

Milas, Costas, and Philip Rothman. 2008. “Out-of-Sample Forecasting of Unemployment Rates with Pooled STVECM Forecasts.” *International Journal of Forecasting*, 24(1): 101–21.

Mitchell, Wesley C. 1927. *Business Cycles: The Problem of Its Setting.* Cambridge, MA:National Bureau of Economic Research.

Montgomery, Alan L., Victor Zarnowitz, Ruey S. Tsay, and George C. Tiao. 1998. “Forecasting the U.S. Unemployment Rate.” *Journal of the American Statistical Association*, 93(442): 478–93.

Neftçi, Salih N. 1984. “Are Economic Time Series Asymmetric Over the Business Cycle?” *Journal of Political Economy*, 92(2): 307–28.

Nekarda, Christopher J. 2009. “Understanding Unemployment Dynamics: The Role of Time Aggregation.” Board of Governors of the Federal Reserve System Unpublished paper.

Rothman, Philip. 1998. “Forecasting Asymmetric Unemployment Rates.” *The Review of Economics and Statistics*, 80(1): 164–168.

Shimer, Robert. 2007. “Reassessing the Ins and Outs of Unemployment.” National Bureau of Economic Research Working Paper 13421.

Sichel, Daniel E. 2007. “Business Cycle Asymmetry: A Deeper Look.” *Economic Inquiry*, 31(2): 224–36.

Stekler, Herman O. 2010. “Perspectives on Evaluating Macroeconomic Forecasts.” George Washington University RPF Working Paper 2010-002.

Tulip, Peter. 2009. “Has the Economy Become More Predictable? Changes in Greenbook Forecast Accuracy.” *Journal of Money, Credit and Banking*, 41(6): 1217–231.

Appendix

Solution to Three-State Model

Denoting $Y_{t+\tau} = \begin{pmatrix} U_{t+\tau} \\ E_{t+\tau} \\ N_{t+\tau} \end{pmatrix}$, we can rewrite equation 8 as $\dot{Y}_{t+\tau} = A_t Y_{t+\tau}$ with

$$A_t = \begin{pmatrix} -\lambda_t^{UE} - \lambda_t^{UN} & \lambda_t^{EU} & \lambda_t^{NU} \\ \lambda_t^{UE} & -\lambda_t^{EU} - \lambda_t^{EN} & \lambda_t^{NE} \\ \lambda_t^{UN} & \lambda_t^{EN} & -\lambda_t^{NE} - \lambda_t^{NU} \end{pmatrix}.$$

Since the columns of A_t sum to zero, A_t has one eigenvalue equal to zero. Denoting P_t the matrix of eigenvectors of A_t corresponding to the eigenvalues $\{r_{1t}, r_{2t}, 0\}$, a solution to equation 8 is given by

$$Y_{t+\tau} = P_t \begin{pmatrix} c_1 \mathrm{e}^{\tau r_{1t}} \\ c_2 \mathrm{e}^{\tau r_{2t}} \\ c_3 \end{pmatrix}, \tag{A.1}$$

or

$$\begin{cases} U_{t+\tau} = p_{11} c_1 \mathrm{e}^{\tau r_{1t}} + p_{12} c_2 \mathrm{e}^{\tau r_{2t}} + p_{13} c_3 \\ E_{t+\tau} = p_{21} c_1 \mathrm{e}^{\tau r_{1t}} + p_{22} c_2 \mathrm{e}^{\tau r_{2t}} + p_{23} c_3 \\ N_{t+\tau} = p_{31} c_1 \mathrm{e}^{\tau r_{1t}} + p_{32} c_2 \mathrm{e}^{\tau r_{2t}} + p_{33} c_3. \end{cases} \tag{A.1$'$}$$

Using the initial and terminal conditions, the one-step ahead forecasts of unemployment, employment and nonparticipation are given by

$$\begin{cases} U_{t+1} = p_{11} c_1 \mathrm{e}^{-3\beta_{1,t}} + p_{12} c_2 \mathrm{e}^{-3\beta_{2,t}} + U_t^* \\ E_{t+1} = p_{21} c_1 \mathrm{e}^{-3\beta_{1,t}} + p_{22} c_2 \mathrm{e}^{-3\beta_{2,t}} + E_t^* \\ N_{t+1} = p_{31} c_1 \mathrm{e}^{-3\beta_{1,t}} + p_{32} c_2 \mathrm{e}^{-3\beta_{2,t}} + N_t^*, \end{cases} \tag{A.2}$$

with

$$\begin{aligned} \beta_1 &\approx \lambda_t^{UE} + \lambda_t^{UN} \\ \beta_2 &\approx \lambda_t^{EU} + \lambda_t^{EN} + \lambda_t^{NE} + \lambda_t^{NU}, \end{aligned}$$

and c_1 and c_2 given by

$$\begin{pmatrix} c_1 \\ c_2 \end{pmatrix} = \begin{pmatrix} p_{11} & p_{12} \\ p_{21} & p_{22} \end{pmatrix}^{-1} \times \begin{pmatrix} U_t \\ E_t \end{pmatrix}.$$

Using the hazard rate forecasts, we can obtain forecasts of unemployment and labor force participation at horizon j from

$$u_{t+j+1} = \frac{U_{t+j+1}}{U_{t+j+1} + E_{t+j+1}} \tag{A.3}$$

$$lfpr_{t+j+1} = \frac{U_{t+j+1} + E_{t+j+1}}{P_{t+j+1}} \tag{A.4}$$

$$U_{t+j+1} = p_{11} c_1 \mathrm{e}^{-\beta_{1,t+j}} + p_{12} c_2 \mathrm{e}^{-\beta_{2,t+j}} + U^*_{t+j} \tag{A.5}$$

$$E_{t+j+1} = p_{21} c_1 \mathrm{e}^{-\beta_{1,t+j}} + p_{22} c_2 \mathrm{e}^{-\beta_{2,t+j}} + E^*_{t+j} \tag{A.6}$$

$$N_{t+j+1} = p_{31} c_1 \mathrm{e}^{-\beta_{1,t+j}} + p_{32} c_2 \mathrm{e}^{-\beta_{2,t+j}} + N^*_{t+j}. \tag{A.7}$$