

Contents lists available at [ScienceDirect](#)

European Economic Review

journal homepage: www.elsevier.com/locate/eer

Forecasting unemployment across countries: The ins and outs

Regis Barnichon ^{a,*}, Paula Garda ^{b,1}^a CREI, Universitat Pompeu Fabra and CEPR, Spain^b OECD, France

ARTICLE INFO

Available online 17 November 2015

JEL classification:

E24

E27

J6

Keywords:

Worker flows

Stock-flow model

ABSTRACT

This paper evaluates the flow approach to unemployment forecasting proposed by Barnichon and Nekarda (2012) for a set of OECD countries characterized by very different labor markets. We find that the flow approach yields substantial improvements in forecast accuracy over professional forecasts for all countries, with especially large improvements at longer horizons (one-year ahead forecasts) for European countries. Moreover, the flow approach has the highest predictive ability during recessions and turning points, when unemployment forecasts are most valuable.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Forecasting the unemployment rate is an important and difficult task for policymakers. 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, such as autoregressive moving average (ARMA) models, to forecast unemployment.

Incorporating information from labor force flows in a stock-flow model of unemployment has been recently shown to substantially improve near-term forecasts of the U.S. unemployment rate (Barnichon and Nekarda, 2012).

A big advantage of such a “flow approach” to unemployment forecasting is its small data requirement, which makes the method applicable for a large set of countries. In fact, following Barnichon and Nekarda (2012), the International Labor Organization started using flow-based models to forecast unemployment in G7 countries (International Labour Organization, 2015).

However, whether the improvements that were found for the US also apply to other countries is an open question. The European labor markets are markedly different from the US labor market, in particular with much smaller worker flows (Elsby et al., 2013), and the US results cannot be trivially extrapolated to European countries.

In this paper, we evaluate the flow approach to unemployment forecasting for a set of OECD countries (France, Germany, Spain, the UK, Japan and the US) spanning a broad range of labor market structures. We find that the flow approach yields large improvements over conventional forecasting methods for *all* countries, with the highest predictive ability being achieved during recessions and turning points, precisely when forecasts are most valuable. Moreover, while improvements are largest at short forecast horizons (one- to three-month ahead) in the US, we obtain large improvements at both short and long horizons (one-year ahead forecasts) in European countries.

A simple analogy helps understand how incorporating information from labor force flows can improve unemployment forecasts. The unemployment rate can be thought of as the amount of water in a bathtub, a stock. Given an initial water level, the level of the water in the next period is determined by the rate at which water flows into the tub from the faucet and the rate at which water flows out of the tub through the drain. When the inflow rate equals the outflow rate, the amount of water in the

* Corresponding author. Tel.: +34 93 542 27 40.

E-mail address: rbarnichon@crei.cat (R. Barnichon).¹ The views expressed in this paper are the authors and are not necessarily shared by the OECD or its member countries.




Fig. 1. Average in and outflow rates across countries. *Source:* Authors calculations based on [Elsby et al. \(2013\)](#). *Notes:* Average of monthly in and outflow rates from unemployment. The starting year varies between 1968 (for the U.S.) and 1986 (for New Zealand and Portugal). For all countries, the data ends in 2009.

tub remains constant. But if the inflow rate increases, we know that the water level will be higher in the future. In other words, the inflow rate and the outflow rate provide information about the future level of water – or in this case, unemployment.

The analogy also helps understand why the “flow approach” can offer improvements at long forecast horizons in the case of European countries, while, in the case of US, the improvements were most remarkable at short horizons. The magnitude of the labor market flows governs the speed of convergence to steady state, i.e., the time needed for the level of water to stabilize at a level consistent with the flows, and [Fig. 1](#) reveals substantial cross-country heterogeneity in labor market flows. At one extreme lies the US, with large worker flows, and at the other extreme lie countries from continental Europe, with much smaller worker flows. With large labor market flows, as in US, convergence occurs relatively rapidly (in the order of magnitude of a quarter), and the flows provide information about movements in unemployment in the short run but little information about the value of unemployment in the longer run. In other words, the model can generate good forecasts in the near term. With small flows, as in Europe, convergence occurs much more slowly (in the order of magnitude of a year), so that observing the current worker flows provide information about movements in unemployment in the longer run. However, for this to happen, the flows must be sufficiently persistent. Otherwise, the influence of the current worker flows on unemployment in the longer-run is small, and observing the current worker flows provides little information about the value of longer-run unemployment. In other words, while the large improvements in forecasting accuracy in Europe are intuitive given the smaller worker flows and smaller convergence rates, they were by no means guaranteed. The worker flows help forecast unemployment at longer forecast horizons in Europe, not only because the flows are small but also because they are persistent.

This paper builds on [Barnichon and Nekarda \(2012\)](#) and [Montgomery et al. \(1998\)](#), and extends a growing literature aimed at improving unemployment forecasts.² The paper also draws on the recent literature on labor force flows, which has been overlooked by the forecasting literature, but has been the subject of numerous studies aimed at understanding the determinants of labor market fluctuations.³

The paper is organized as follows. [Section 2](#) presents the flow approach to unemployment forecasting. [Section 3](#) presents the data and construction methods. [Section 4](#) evaluates empirically the forecasting performance of our model and gives the intuition why a flow approach to unemployment forecasting does a good job. Finally, the last section concludes.

2. The flow approach to unemployment forecasting

This section presents our flow approach to unemployment forecasting. First, we present the theory underlying our approach. We use a stock-flow model to show how the unemployment rate – a stock – varies over time because of variations in the rate at which workers flow into and out of unemployment. In particular, we show how the unemployment rate converges to its steady-state rate at a time-varying rate, and both the steady-state unemployment rate and the convergence rate depend on the worker flows into and out of unemployment. The flow approach to unemployment forecasting consists in using information on worker flows to exploit this convergence property of unemployment. We present our “Flow-based unemployment Forecasting model” (FbF) in the second part of this section.

² See, for example, [Rothman \(1998\)](#), [Golan and Perloff \(2004\)](#), [Brown and Moshiri \(2004\)](#), and [Milas and Rothman \(2008\)](#).

³ Some related papers are [Shimer \(2012\)](#); [Petrongolo and Pissarides \(2008\)](#); [Solon et al. \(2009\)](#); [Elsby et al. \(2013\)](#); [Barnichon \(2012\)](#); [Nekarda \(2009\)](#) and [Fujita and Ramey \(2012\)](#).

2.1. A stock-flow model of unemployment

Individuals can only be in one of two labor force states: employed or unemployed.⁴ In addition to providing a simple framework for understanding the basic flow-based accounting of the steady-state unemployment rate, the data requirements from using a two-state approach are relatively benign, so that the approach can be applied to a broad set of countries.

Denote $u_{t+\tau}$ the unemployment rate at instant $t+\tau$ with t indexing the period (e.g., a quarter) and $\tau \in [0, 1]$ a continuous measure of time within the period. Assume that between date t and date $t+1$ all unemployed persons find a job according to a Poisson process with constant arrival rate f_{t+1} , and all employed workers lose their job according to a Poisson process with constant arrival rate s_{t+1} . We adopt this timing convention to reflect data availability, as the hazard rate is only observed at date $t+1$. Indeed, in real time a forecaster does not observe s_{t+1} and f_{t+1} , but only s_t and f_t . This is because at date t one can only observe labor force flows from $t-1$ to t .

The unemployment rate then evolves according to

$$\frac{du_{t+\tau}}{d\tau} = s_{t+1}(1 - u_{t+\tau}) - f_{t+1}u_{t+\tau}, \quad (1)$$

as changes in unemployment are given by the difference between the inflows and the outflows. Solving Eq. (1) yields

$$u_{t+\tau} = \beta_{t+1}(\tau)u_{t+1}^* + [1 - \beta_{t+1}(\tau)]u_t, \quad (2)$$

where

$$u_{t+1}^* \equiv \frac{s_{t+1}}{s_{t+1} + f_{t+1}} \quad (3)$$

denotes the *Steady-State Unemployment Rate* (SSUR), and

$$\beta_{t+1}(\tau) \equiv 1 - e^{-\tau(s_{t+1} + f_{t+1})} \quad (4)$$

is the *rate of convergence* to that steady state.

SSUR is the unemployment rate that would eventually prevail were the flows into and out of unemployment to remain at their current rate forever.

Fig. 2 shows the tight, leading relationship between the steady-state unemployment rate, u^* , and the actual unemployment rate, u for a range of OECD countries,⁵ and Table A1 in the Appendix confirms this visual inspection by showing the cross-correlations between unemployment and steady-state unemployment. The steady-state rate leads the actual unemployment rate by one- to two-quarters, and this leading relationship forms the basis of our approach to unemployment forecasting.⁶

2.2. The flow-based unemployment forecasting model

We now present our Flow-based unemployment Forecasting model, which consists of two stages: (i) a forecast of the worker flows determining the current and future values of the steady-state unemployment rate, and (ii) an iteration on the law of motion of unemployment (Eq. (2)).

By forecasting the worker flow rates and feeding these forecast into the non-linear law of motion of unemployment, our model takes a crucial aspect of the behavior of unemployment into account: the unemployment converges to its *time-varying* steady-state rate at a *time-varying* rate, and both the steady-state rate and the convergence rate are determined by the worker flows taking place in the labor market.

2.2.1. First stage: forecasting the labor force flows

Eq. (2) suggests a simple way to forecast unemployment using information on worker flows. If we assume that the hazard rates remain constant at their last observed value over the forecast horizon, Eq. (2) directly gives us the forecasted value of unemployment at horizon τ .⁷ If the hazard rates are persistent enough, this basic approach may provide reasonable forecasts.

⁴ We assume that the contribution of movements in-and-out of the labor force to unemployment fluctuations is negligible, consistent with recent literature (Solon et al., 2009). Although a three state model with unemployment, employment and inactivity (which allows for movements in-and-out of the labor force) is theoretically possible, the data requirements are strong (relying on household survey micro data), so that such a model is very difficult to implement for most countries. Moreover, micro data are typically available with a significant delay, making them generally ill-suited for use in forecasting models. In contrast, a two-state model is easy to implement for many countries. In the case of US, Barnichon and Nekarda (2012) show that the performances of the 2-state and 3-state models are comparable. The 3-state model is a more realistic characterization of the labor market, but this advantage is compensated by the stronger data requirements, which lead to a higher noise-to-signal ratio in the data.

⁵ Section 3 will describe the procedure to construct the worker flow series.

⁶ Note that the leading relationship between u^* and u differs across countries. For the US, steady-state unemployment leads unemployment by one quarter, so that the series are almost indistinguishable. In contrast, for Germany, steady-state unemployment rate leads unemployment by two- to three quarters. We will see that this difference shows up again in the different performances of FbF across countries.

⁷ Despite its extreme simplicity, we will see that even this basic approach forecasts as well, or even better, than standard stock models, which shows how powerful a flow-based approach to unemployment forecasting can be.

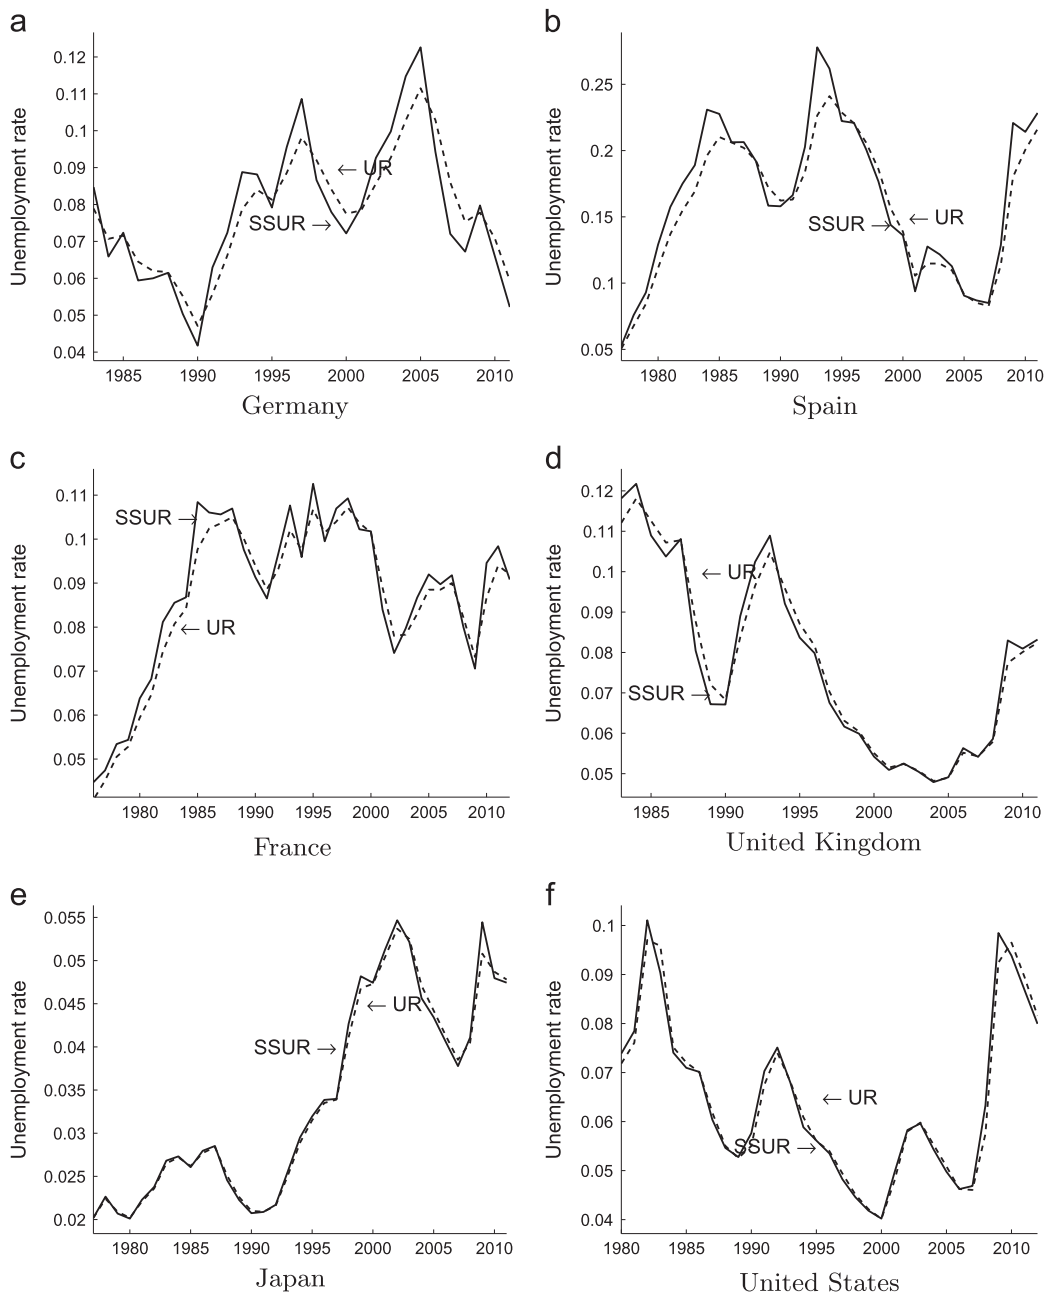


Fig. 2. Unemployment rate (UR) and steady state unemployment rate (SSUR).

Source: Authors calculations based on [Elsby et al. \(2013\)](#) and National Statistics Institutes data. Notes: The dashed line is the unemployment rate, and the continuous line is the steady state unemployment rate $u^* = s/s+f$. Average annual data.

However, because the hazard rates change over time, a better approach consists in properly forecasting the flow rates. Thus, the first stage of our baseline FbF model consists in producing forecasts of the worker flows. To generate such forecasts, we use a VAR, where we include leading indicators of labor force flows, such as vacancy posting v_t , claims for unemployment insurance ui_t , and GDP gdp_t . Specifically, we consider a vector of the form

$$\mathbf{y}_t = (\ln s_t, \ln f_t, \Delta \ln u_t, \Delta \ln v_t, \Delta \ln gdp_t, \dots)' \quad (5)$$

and we estimate the VAR

$$\mathbf{y}_t = C + \phi_1 \mathbf{y}_{t-1} + \phi_2 \mathbf{y}_{t-2} + \dots + \phi_n \mathbf{y}_{t-n} + \varepsilon_t \quad (6)$$

Table 1
Data sources.

Data	France	Spain	UK
Quarterly duration data	Q1.1992–Q4.2011	Q1.1987–Q4.2011	Q2.1992–Q4.2011
Source	INSEE-Pôle emploi	INE-LFS	ONS-LFS
Gross worker flow data			Q1.1989–Q2.2008
Source			Smith (2011)
Annual duration data	1977–2011	1977–2011	1982–2011
Source	OECD, Eurostat-LFS, Elsby et al. (2013)		
	Germany	Japan	US
Quarterly duration data	Q1.2005–Q3.2011	Q1.2002–Q3.2011	M1.1968–M9.2011
Source	DeStatis-LFS	Stats Bureau-LFS	BLS-CPS
Gross worker flow data	M1.1984–M6.2009	Q1.1978–Q4.2009	
Source	Hertweck and Sigris (2012)	Lin and Miyamoto (2012)	
Annual duration data	1985–2011	1977–2011	
Source	OECD, Eurostat-LFS, Elsby et al. (2013)		

over a ten-year rolling window.⁸ Since many VAR specifications are possible and since the best-performing specification depends on the country of interest, specification (5) is illustrative, and [Table B2](#) in [Appendix B](#) reports the VAR specification and the number of lags n for each country. In each case, we chose the VAR specification (varying the lag length from 1 to 4 quarters and choosing between (log) variables in levels or in first-difference) that generated the smallest average Mean-Square Errors over the different forecast horizons. The results change little with alternative specifications.

2.2.2. Second stage: iterating using unemployment's non-linear law of motion

Given a set of worker flows forecasts, the second stage of a FbF forecast then consists in iterating on [Eq. \(2\)](#), i.e., iterating on unemployment's law of motion. Specifically, given forecasts of the flow rates $\hat{f}_{t+j|t}$ and $\hat{s}_{t+j|t}$ with $j \in \mathbb{N}$, a j -period-ahead forecast of the unemployment rate, $\hat{u}_{t+j|t}$, can be constructed recursively from

$$\hat{u}_{t+j|t} = \hat{\beta}_{t+j|t} \hat{u}_{t+j|t}^* + (1 - \hat{\beta}_{t+j|t}) \hat{u}_{t+j-1|t}, \quad (7)$$

with

$$\hat{u}_{t+j|t}^* = \frac{\hat{s}_{t+j|t}}{\hat{s}_{t+j|t} + \hat{f}_{t+j|t}} \quad (8)$$

and

$$\hat{\beta}_{t+j|t} = 1 - e^{-(\hat{s}_{t+j|t} + \hat{f}_{t+j|t})}. \quad (9)$$

In other words, the forecasted value of unemployment at date $t+j$ is obtained by taking a weighted average of the previous-period ($t+j-1$) unemployment forecast (or actual unemployment rate when $j=1$) and the time ($t+j$) steady-state unemployment rate, with the weights determined by the speed of convergence to steady state. Importantly, the speed of convergence and thus the weights are also time-varying, so that the law of motion of unemployment is non-linear, a point to which we will return in the performance evaluation section.

3. Data

We are interested in producing quarterly unemployment forecasts for six large OECD countries – France, Germany, Japan, Spain, the UK and the US. Quarterly data are available for these countries since 2000, so that the FbF model can currently be easily used to forecast unemployment in these countries.

However, in order to first estimate and evaluate the forecasting performances of our model, we need longer quarterly time series of the inflow and outflow rates.

Worker flow series can be constructed from data on the stocks of unemployment and short-term unemployment following [Shimer \(2012\)](#) and [Elsby et al. \(2013\)](#). Since quarterly unemployment duration data are not always available before 2000, we construct quarterly flow rates series over the last 30 years by combining yearly OECD duration data (as [Elsby et al., 2013](#)) with quarterly transition rates measured from micro household survey data. Specifically, we proceed in two steps.

First, we construct yearly outflow rates series as in [Elsby et al. \(2013\)](#) by using information on the number of persons unemployed, U_t , and on the number of unemployed of less than d months, $U_t^{<d}$. Specifically, the probability that an unemployed worker

⁸ A rolling window (in which the model is estimated over the previous K periods) yielded more accurate forecasts than a recursive window (in which the model is estimated over the entire observed history). The size of the window was restricted by data availability.

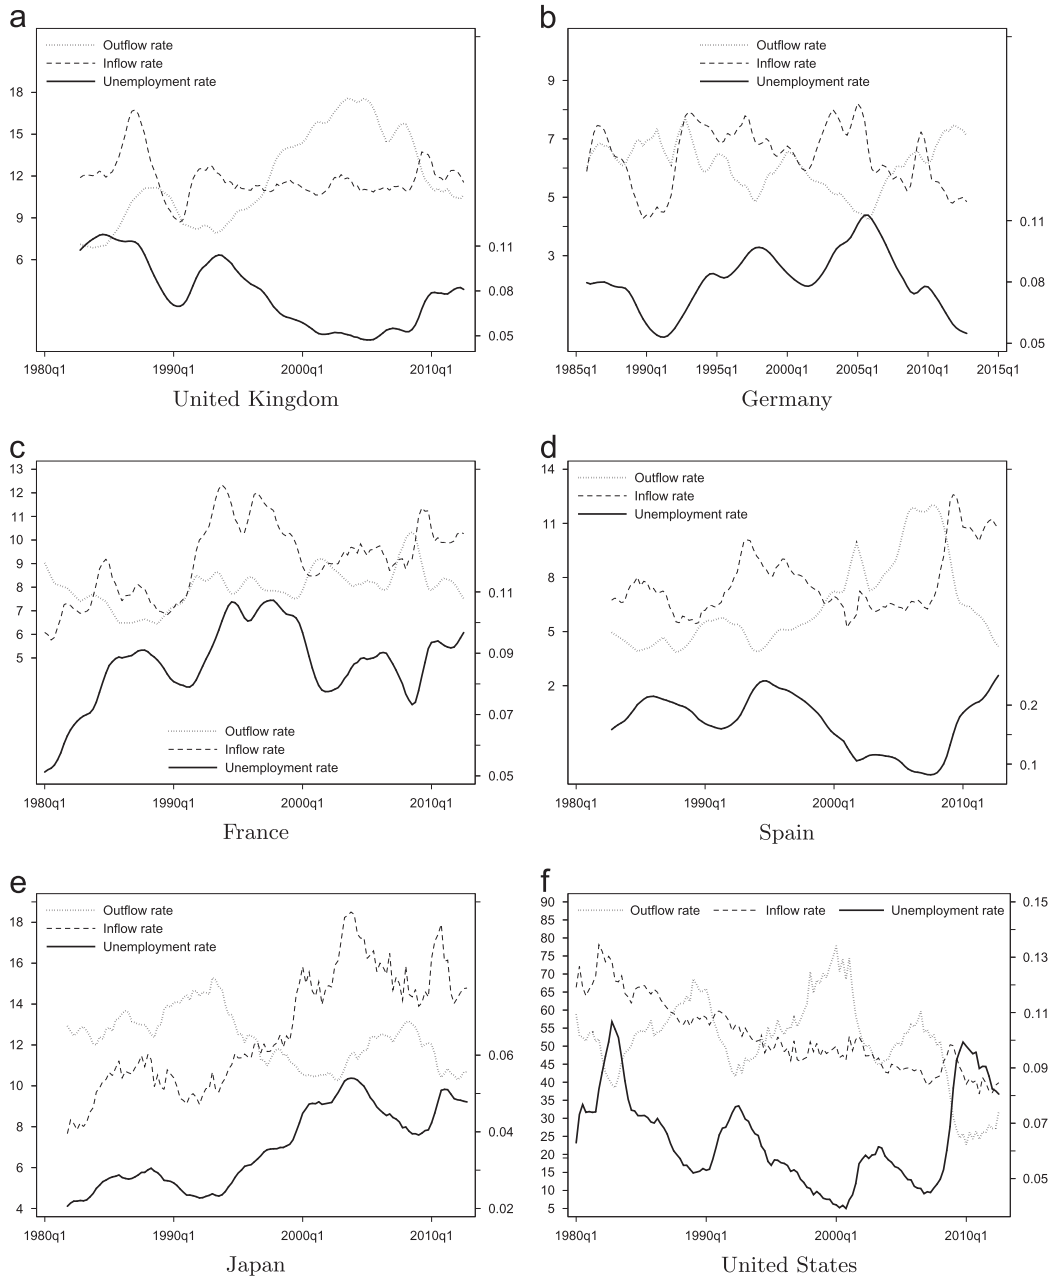


Fig. 3. Unemployment Inflows and Outflows. *Source:* Authors calculations based on Hertweck and Sigris (2012), Lin and Miyamoto (2012), Smith (2011), Elsby et al. (2013) and National Statistics Institutes data. *Notes:* Quarterly data smoothed by a four-quarter moving average. The unemployment rate is on the right axis, and the flow rates are on the left axis. For clarity, the inflow rate is rescaled $s_t/E(u_t)$. s_t and f_t are quarterly averages of the monthly inflow and outflow rates and expressed in percent.

exits unemployment within d months, F , can be calculated from

$$F_{t+1} = 1 - \frac{U_{t+1} - U_{t+1}^{<d}}{U_t}$$

with $f_{t+1} = -\ln(1 - F_{t+1})/d$ the monthly hazard rate associated with the probability that an unemployed worker at time t completes his spell within the subsequent d months.⁹ The estimated outflow rates are very close to the ones reported by Elsby et al. (2013).

⁹ We use $d=12$ months for Spain, France and Germany, $d=6$ months for the UK and Japan, and $d=1$ for the US. An alternative approach is to combine information on the share of workers in different unemployment duration bins, as described in Elsby et al. (2013). Estimated outflow rates are little changed.

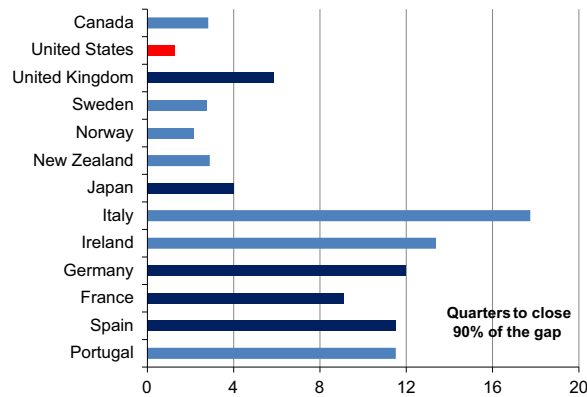


Fig. 4. Time to convergence to steady-state unemployment. *Source:* Authors calculations based on [Elsby et al. \(2013\)](#) and National Statistics Institutes data, and [Barnichon and Nekarda \(2012\)](#). *Notes:* Time in quarters needed to close 90 percent of the gap with steady-state unemployment rate $u = s/s + f$. Averages over the period 1985–2012.

Second, we construct quarterly outflow rates in two alternative ways, depending on data availability. When quarterly duration data are available (see [Table 1](#) for data availability by country), we use the stock of unemployed and short term unemployed as described above. When, quarterly duration data are not available, we spliced the annual duration data with the quarterly series to expand the time coverage of the quarterly series. To convert the annual duration data to a quarterly frequency, we impute the inter-year variations from quarterly gross worker flows constructed from household survey data whenever possible ([Hertweck and Sigrist, 2012](#) for Germany, [Lin and Miyamoto, 2012](#) for Japan, [Smith, 2011](#) for the UK), or, if not possible, we keep the quarterly flow rate constant at its beginning of the year value.¹⁰

The inflow rate, s , is then obtained by solving Eq. (1) forward over $[t, t + 1]$ and finding the value of s_{t+1} that solves

$$U_{t+1} = \frac{[1 - e^{-(f_{t+1} + s_{t+1})}] s_{t+1}}{f_{t+1} + s_{t+1}} (U_t + E_t) + e^{-(f_{t+1} + s_{t+1})} U_t.$$

Note that in this accounting, given a value for the unemployment outflow rate (which also captures movements out of the labor force) and the stock of unemployed persons, the inflow rate is the rate that explains the observed stock of unemployed persons in the next month. As a result, the inflow rate incorporates all movements in unemployment not accounted for by the unemployment outflow rate.

[Fig. 3](#) plots the unemployment rate and the constructed flow rates series for the 6 OECD countries. Two observations are worth noting. First, while the flow rates move over time, they also display substantial persistence (see Appendix A for summary statistics of the flow rate series). This indicates that the contemporaneous value of the flow rates contains information about the future values of the steady-state unemployment rate and thus about the future unemployment rate. We will see that the persistence in the flow rates is an important reason behind FbF forecasting performance. Second, the level of the flow rates varies substantially across countries (as first highlighted by [Elsby et al., 2013](#)) with the US displaying flow rates about 10 times larger than those in Europe. Since the flow rates affect β_t , the convergence rate to steady-state, variations in the level of flow rates have important implications for the dynamics of unemployment. [Fig. 4](#) shows the time needed for unemployment to close 90 percent of the gap with its steady state value. In the US, unemployment closes 90 percent of the gap in about four months, but in Germany, it takes almost three years. We will see that this difference in convergence rate shows up again in the different performances of FbF across countries.

4. Empirical forecasting performance

We can now evaluate the empirical performances of FbF by comparing its unemployment rate forecasts with alternative forecasts along two dimensions. First, we assess the relative performances of each model by using the Root-Mean-Squared-Error (RMSE) of out-of-sample forecasts, considering forecast horizons ranging from one-quarter-ahead to two-year-ahead. Second, because it is harder, but especially valuable, to forecast the unemployment rate around recessions, we assess the model's performance over the business cycle.

¹⁰ To infer the quarterly movements in f from movements in the quarterly transition rates calculated from micro data, we posit that the outflow rate f behaves similarly to the unemployment outflow rate \hat{f} derived in a three labor market state model where $\hat{f} = \lambda^{UE} + \lambda^{UI} * \lambda^{IE} / \lambda^{IE} + \lambda^{IU}$ ([Barnichon and Figura, 2013](#)) with λ^{AB} the transition rate from A to B, with U unemployment, E employment, and I inactivity. Using \hat{f} as a proxy for the outflow rate, we can convert the annual series f to a quarterly frequency whenever quarterly duration data are unavailable. When both data sources overlap, f and \hat{f} are highly correlated.

Table 2

Forecast accuracy: RMSE of FbF and Consensus Forecasts. RMSE (in ppt) of FbB and Relative RMSE.

	t+1	t+2	t+4	t+6
FbF				
UK	0.33	0.63	0.72	1.55
Germany	0.69	0.64	0.86	1.45
France	0.55	0.57	0.92	1.13
Spain	–	–	2.21	–
Japan	0.17	0.24	0.47	0.49
United States	0.30	0.55	1.09	–
Consensus Forecasts relative to FbF				
UK	1.91** (0.05)	1.19 (0.34)	1.14 (0.23)	0.86 (0.40)
Germany	0.73 (0.43)	1.33 (0.15)	1.49* (0.08)	0.95 (0.89)
France	1.63*** (0.05)	1.64* (0.09)	1.20 (0.32)	1.07 (0.79)
Spain	–	–	1.79*** (0.01)	–
Japan	1.54** (0.03)	1.29** (0.05)	0.90 (0.43)	1.16 (0.43)
United States	1.66* (0.09)	0.91 (0.38)	0.73 (0.13)	0.70 (0.26)

Source: Authors' calculations. Notes: First panel is Root Mean Square in Error of FbF in percentage points. Second panel shows the RMSE of Consensus forecasts relative to FbF. Consensus Forecasts are from Consensus Economics in the US, UK, Germany, France and Japan, OECD forecasts for Spain. Forecasts evaluations cover 1993–2011, except for Germany 1996–2011, and Spain 1997–2011. *p*-values of Giacomini–White test statistic are reported in parentheses.

*** Indicates statistically different from FbF at 1/percent.

** Indicates statistically different from FbF at 5/percent.

* Indicates statistically different from FbF at 10 percent.

This section starts by introducing the alternative models used for comparisons, and then reports and discusses the relative performances of FbF.

4.1. Alternative forecasts

We evaluate FbF forecasts against two sets of alternatives: (i) professional forecasts, and (ii) forecasts from standard time series models.

4.1.1. Professional forecasts

Professional forecasts were obtained from Consensus Economics¹¹ for all countries in our sample except Spain. Consensus Economics surveys the forecasts of a large number of private and public forecasters (investment banks, large international corporations, economic research institutes, and universities). Consensus Economics conducts forecast surveys on a monthly basis, and professional forecasters are surveyed in the middle of each month about their forecast for the current year and the next.¹² We use the mean forecast of the survey as the professional forecast, and we only use consensus forecasts published in the last month of each quarter.¹³ This allowed us to construct series of professional forecasts over 1993–2011 for one-quarter ahead ($t+1$), two-quarter ahead ($t+2$), one-year ahead ($t+4$) and six-quarter ahead ($t+6$) forecasts.

Since historical forecasts for Spain are not available from Consensus Economics, we use instead OECD forecasts.¹⁴ The OECD releases forecasts in December (the last month of Q4) for next year unemployment rate, using data available as of November, which allows us to construct a series of one-year ahead ($t+4$) forecasts over 1997–2011.¹⁵

When comparing FbF forecasts to professional forecasts, it is important that FbF does not have a larger information set than the real-time forecasters. Labor force surveys are generally conducted at a quarterly or monthly frequency, and the survey's release date differs across countries. To reflect data availability, we only allow FbF to have access to the latest published unemployment report at the time of the forecast. For France and Germany, the reports are released with

¹¹ <http://www.consensuseconomics.com/>.

¹² For instance, in June 2011, Consensus Economics published forecasts for the level of unemployment in 2011-Q4 (i.e., a two-quarter ahead forecast, or $t+2$) and 2012-Q4 (i.e., a six-quarter ahead forecast, or $t+6$).

¹³ This is done to ensure that professional forecasts have as much information as possible about the current quarter (and stack the cards against FbF).

¹⁴ OECD unemployment forecasts have been shown to be as good as Consensus Economics forecasts (Batchelor, 2010), and we verified that the performances of OECD forecasts for the US, UK, Germany, and France were very similar to the performances of Consensus Economics forecasts.

¹⁵ For instance, in December 2011, the OECD published forecasts for the level of unemployment in 2012-Q4 (i.e., a four-quarter ahead forecast, or $t+4$).

Table 3

Forecast accuracy: RMSE of FbF and RMSE of alternative models relative to FbF.

	Forecast horizon					
	t+0	t+1	t+2	t+3	t+4	t+8
UK						
FbF	0.16	0.29	0.43	0.59	0.77	1.47
<i>VAR no flows</i>	1.41* (0.09)	1.40*** (0.00)	1.25* (0.06)	1.21 (0.19)	1.23 (0.21)	1.38 (0.24)
<i>VAR</i>	1.34*** (0.00)	1.28*** (0.01)	1.14* (0.09)	1.09 (0.23)	1.09 (0.21)	1.07 (0.36)
\bar{u}^*	1.22*** (0.00)	1.19*** (0.00)	1.14*** (0.01)	1.08 (0.15)	1.00 (0.76)	0.82 (0.06)
$\bar{\beta}$	1.10* (0.07)	1.15 (0.15)	1.20 (0.16)	1.22 (0.12)	1.22* (0.09)	1.18** (0.02)
<i>discrete-FbF</i>	1.28 (0.15)	1.23 (0.2)	1.18 (0.14)	1.11 (0.11)	1.05 (0.10)	0.97 (0.05)
Germany						
FbF	0.13	0.33	0.56	0.76	0.94	1.64
<i>VAR no flows</i>	1.91*** (0.00)	1.74*** (0.00)	1.68** (0.04)	1.63** (0.03)	1.56** (0.02)	1.35** (0.03)
<i>VAR</i>	1.45*** (0.01)	1.52*** (0.01)	1.52*** (0.01)	1.48*** (0.01)	1.43*** (0.01)	1.33*** (0.01)
\bar{u}^*	1.34 (0.35)	1.13 (0.91)	1.03 (0.42)	1.00 (0.43)	0.99 (0.58)	0.97 (0.48)
$\bar{\beta}$	1.10 (0.25)	1.06 (0.55)	1.04 (0.65)	1.05 (0.6)	1.08 (0.53)	1.06 (0.57)
<i>discrete-FbF</i>	1.45 (0.37)	1.13 (0.96)	0.99 (0.65)	0.95 (0.82)	0.94 (0.79)	0.91 (0.68)
France						
FbF	0.19	0.35	0.51	0.68	0.83	1.23
<i>VAR no flows</i>	1.25 (0.18)	1.30 (0.26)	1.45 (0.3)	1.59 (0.38)	1.71 (0.33)	1.74 (0.21)
<i>VAR</i>	1.15 (0.22)	1.19 (0.29)	1.21 (0.33)	1.23 (0.33)	1.24 (0.28)	1.18 (0.26)
\bar{u}^*	0.99 (0.21)	1.02 (0.42)	1.05 (0.56)	1.05 (0.59)	1.08 (0.5)	1.07 (0.68)
$\bar{\beta}$	1.05 (0.62)	1.05 (0.56)	1.05 (0.84)	1.05 (0.42)	1.04 (0.23)	1.03 (0.18)
<i>discrete-FbF</i>	1.05 (0.65)	1.01 (0.64)	0.98 (0.57)	0.96 (0.47)	0.95 (0.41)	0.95 (0.37)
Spain						
FbF	0.49	1.08	1.78	2.49	3.17	5.44
<i>VAR no flows</i>	1.17 (0.38)	1.12 (0.68)	1.10 (0.52)	1.14 (0.41)	1.17 (0.36)	1.51 (0.30)
<i>VAR</i>	1.11 (0.43)	1.06 (0.94)	1.05 (0.81)	1.09 (0.57)	1.13 (0.44)	1.39 (0.30)
\bar{u}^*	0.78** (0.04)	0.77** (0.02)	0.76** (0.04)	0.76* (0.07)	0.76* (0.08)	0.77 (0.14)
$\bar{\beta}$	1.10 (0.21)	1.04 (0.21)	0.99 (0.22)	0.98 (0.37)	0.97 (0.48)	0.98 (0.88)
<i>discrete-FbF</i>	1.24 (0.17)	1.09 (0.22)	0.98 (0.25)	0.92 (0.28)	0.89 (0.28)	0.86 (0.34)
Japan						
FbF	0.16	0.22	0.30	0.38	0.46	0.88
<i>VAR no flows</i>	1.09** (0.04)	1.18 (0.13)	1.19 (0.16)	1.20 (0.23)	1.17 (0.41)	0.95 (0.54)
<i>VAR</i>	1.05** (0.05)	1.08 (0.15)	1.07 (0.17)	1.09* (0.10)	1.11* (0.08)	1.09* (0.10)
\bar{u}^*	1.14** (0.04)	1.30*** (0.00)	1.25** (0.02)	1.28*** (0.03)	1.18 (0.23)	0.89 (0.46)
$\bar{\beta}$	1.01 (0.19)	1.03 (0.11)	1.04* (0.07)	1.05* (0.08)	1.05* (0.08)	1.04* (0.10)
<i>discrete-FbF</i>	1.04 (0.50)	1.08 (0.25)	1.04 (0.62)	1.02 (0.92)	0.97 (0.72)	0.80* (0.09)

Table 3 (continued)

	Forecast horizon					
	t+0	t+1	t+2	t+3	t+4	t+8
US						
FbF	0.07	0.30	0.55	0.81	1.09	1.90
VAR no flows	2.23*** (0.01)	1.34** (0.02)	1.22* (0.07)	1.21* (0.09)	1.19 (0.13)	1.31 (0.13)
VAR	1.50*** (0.00)	1.13 (0.16)	1.06 (0.38)	1.03 (0.53)	0.99 (0.65)	0.97 (0.25)
\bar{u}^*	1.36*** (0.00)	1.05 (0.54)	0.95 (0.29)	0.90 (0.21)	0.90 (0.22)	0.91 (0.34)
$\bar{\beta}$	1.50* (0.08)	1.36 (0.13)	1.17 (0.11)	1.12 (0.12)	1.09 (0.14)	1.02 (0.24)
discrete-FbF	1.47*** (0.01)	1.19 (0.22)	1.06 (0.67)	0.99 (0.81)	0.93 (0.28)	0.88 (0.12)

Source: Authors' calculations. Notes: Rows starting with (FbF) report the Root Mean Square in Error of (FbF) forecasts in percentage points. All the other rows report the relative RMSEs of the VAR, \bar{u}^* , $\bar{\beta}$ and discrete-FbF forecasts relative to FbF. The evaluation of the models' forecasts is calculated from 77 forecasts over 1992q1–2011q4 (except for Germany 1995q2–2011q4). $t+0$ denotes current quarter forecast. p -values of Giacomini–White test statistic for the comparison with (FbF) are reported in parentheses.

*** indicates statistically different from (FbF) at 1 percent.

** indicates statistically different from (FbF) at 5 percent.

* indicates statistically different from (FbF) at 10 percent.

considerable delay, and we consider that FbF only has access to data published two quarters ago. For instance, for a forecast as of June 2011, FbF information set ends in 2010-Q4.¹⁶ For the other countries, US, UK, Japan and Spain, the labor force surveys are published faster, and we consider that FbF has access to the employment report published last quarter.¹⁷ For instance, for a forecast as of June 2011, FbF information set ends in 2011-Q1. Note that our approach is very conservative, because, for European countries, forecasters also have access, unlike FbF, to registered unemployment data that are published at a monthly frequency. Thus, the information set of FbF is certainly smaller than that of professional forecasters.

4.1.2. Alternative time series models

We now considered a number of alternative time-series forecasts of the unemployment rate. While using professional forecasts as a benchmark is important to establish the usefulness of our method in practice, considering different time series models is also useful for two reasons. First, it allows us to more clearly establish the superiority of the flow-based approach over standard stock-based models. Second, it allows us to better understand the key elements underlying the superior performances of FbF. In particular, we will use different time-series models to illustrate how the performances of FbF come from (i) the use of worker flows data, (ii) non-linearities in the law of motion of unemployment, and in particular the time-varying nature of both the steady-state rate u_t^* and the convergence rate β_t .

First, we consider two VAR models. The first VAR model does not use data on worker flows and thus can be seen as a benchmark time series model that does not include information from worker flows, which is the standard approach in unemployment forecasting and the starting point of our method. The second VAR model uses information from worker flows. By comparing the performances of these two VARs, we can evaluate the value-added of using worker flow information when forecasting unemployment. Moreover, by comparing FbF forecasts with those of a VAR with worker flow data, we can evaluate the extent to which the nonlinear relationship implied by the theory is quantitatively important to forecast unemployment. Except for the use of worker flows data, both VARs include the same specification as the VAR used to predict the worker flows in the FbF model (including leading labor market indicators).¹⁸ Like FbF, the VAR models are estimated using a ten-year rolling window.

Second, in order to better understand the origins of the superior performance of FbF, we consider successive modifications of FbF.

¹⁶ For instance, in France, the 2013-Q1 employment report was only published on June 6. We take a conservative approach and consider the 2013-Q1 employment report *unavailable* to FbF for a forecast as of June.

¹⁷ This is clear for the US since it is a monthly survey released on the first Friday of the next month. The Japan also conducts a monthly survey. The 2013-M3 was released on April 29th, so that 2013-Q1 unemployment data were available for forecasters as of June.

The UK and Spain conduct quarterly surveys. Since 1992 UK reports monthly estimates of the unemployment rate. However, they are not designed to represent national statistics. The LFS sample is designed so that the data collected for any three consecutive monthly reference periods (or rolling quarters) are representative of the UK population. However, the data for any given single month is unlikely to be representative of the UK. Because, these sampling effects can cause movements in the single month that are a consequence of the survey nature of the LFS and are not a true reflection of change in the wider economy we use quarterly data. Moreover, a monthly exercise did not provide better forecasting performance. For the UK, the 2013-Q1 unemployment rate date was released on May 15th, and was thus available to forecasters as of June. For Spain, the 2013-Q1 unemployment rate date was released on April 25th, and was thus available for forecasters as of June.

¹⁸ Using alternative specifications (varying the lag length from 1 to 4 or using variables in levels or in first difference) give very similar conclusions.

First, we hold the inflow and outflow rates constant at their last known values and simply let the model converge to its current steady-state u_t^* at the constant rate β_t , as predicted by the law of motion of unemployment. In other words, we forecast unemployment at horizon j by iterating on a simplified version of Eq. (7):

$$\hat{u}_{t+j|t} = \beta_t u_t^* + (1 - \beta_t) \hat{u}_{t+j-1|t}.$$

We refer to this model as the (\bar{u}^*) model. Shutting down the evolution of the hazard rates isolates the contribution of the current steady-state unemployment rate to the forecasting performances of FbF.

Second, we let the steady-state evolve as predicted by the forecasted flows, but we keep the speed of convergence fixed at its last known value β_t . Specifically, we forecast unemployment at horizon j by iterating on

$$\hat{u}_{t+j|t} = \beta_t \hat{u}_{t+j|t}^* + (1 - \beta_t) \hat{u}_{t+j-1|t}.$$

We refer to this model as the $(\bar{\beta})$ model. This exercise will allow us to evaluate how important are time-variations in β_t , the convergence rate to the steady-state, to forecasting accuracy.

Finally, we consider a flow-based model that is simpler to implement than FbF. If the job finding rate and job separation rate are small, Eq. (2) with $\tau=1$ gives that the law of motion of unemployment is approximately given by

$$u_{t+1} \simeq S_{t+1}(1 - u_t) + (1 - F_{t+1})u_t$$

with $S_{t+1} = 1 - e^{-s_{t+1}}$ the job separation probability between $t+1$ and t , and $F_{t+1} = 1 - e^{-f_{t+1}}$ the job finding probability over the same period.¹⁹ Using this simpler law of motion, we can forecast unemployment at horizon j by iterating on

$$\hat{u}_{t+j|t} = \hat{S}_{t+j|t} - (\hat{S}_{t+j|t} + \hat{F}_{t+j|t}) \hat{u}_{t+j-1|t}.$$

Intuitively, if the period of observation is small enough (such that s_t and f_t are small enough), the law of motion of unemployment is approximately a linear first-order difference equation. Since this model (thereafter referred to as “discrete-FbF”) is arguably simpler to implement, we will evaluate its performance and contrast it with our baseline FbF model.

4.2. Forecast errors

Tables 2 and 3 report the RMSE of forecasts for quarterly unemployment rates from the FbF model over a two-year horizon and the relative RMSE of alternative forecasts. To evaluate the statistical significance of our results, we report the p -values of the unconditional Giacomini-White (2006) predictive ability test statistic of equal predictive ability between our FbF forecast and the comparison forecast.²⁰

4.2.1. Professional forecasts

Table 2 reports the performance of FbF against professional forecasts and shows that FbF dramatically outperforms the consensus forecasts: FbF improves upon professional forecasts in most cases with all the significant differences corresponding to improvements offered by FbF. More specifically, FbF’s RMSE for one-year ahead forecasts ($t+4$) are lower than professional forecasts by 80% for Spain and 50% for Germany and 20% for France. For Japan, the UK and the US, the RMSE of one-quarter-ahead forecasts is reduced by respectively 50%, 90% and 65%.

Interestingly, improvements vary substantially across countries and show an interesting pattern: FbF performs best at long forecast horizons in countries with small workers flows (Germany and Spain, Fig. 4). In contrast, improvements are only observed for short forecast horizons in countries with large worker flows (the US). Moreover, note the non-monotonicity of the relative performances of FbF for France and Germany with performances initially increasing and then decreasing with the forecast horizon.

4.2.2. Alternative time series models

Table 3 reports the performance of FbF against alternative time-series models.

A VAR model without worker flows, a popular starting point for professional forecasters, performs substantially worse than FbF at all horizons, consistent with our previous result that FbF improves substantially upon professional forecasts. Interestingly, adding worker flows substantially improves the performance of the VAR, indicating that including labor force flows in the information set already provides large gains in forecast accuracy compared to stock-based approaches. However, the VAR with flows still performs worse than FbF, indicating that taking into account the non-linear nature of unemployment’s law of motion is important to produce good forecasts.

The importance of forecasting the flows is clear from comparing the forecast accuracy of FbF and (\bar{u}^*) . Recall that (\bar{u}^*) forecasts are based solely on the convergence property of unemployment to the current value of steady-state unemployment. (\bar{u}^*) performs worse than FbF at all horizons (with the exception of Spain, a point that we discuss below), indicating that time variation in the flow rates is, indeed, an important element of our model. Nonetheless, it is remarkable that the

¹⁹ To show this, use the fact that $u_{t+1}^* = s_{t+1}/s_{t+1} + f_{t+1}$ and $1 - e^{-s_{t+1} - f_{t+1}} \simeq s_{t+1} + f_{t+1}$.

²⁰ We use the Giacomini and White (2006) predictive ability test, because it is robust to both non-nested and nested models (as are the VAR, \bar{u}^* and FbF models), unlike the Diebold and Mariano (1995) test.

non-estimated model (\bar{u}^*) performs as well or better than the *estimated* VAR model. This indicates how powerful the flow approach can be compared to the standard stock approach.

Looking at the performance of (β^*), we can see that the time-variation in convergence to steady-state is also an important aspect of a good forecast. In recessions, the speed of convergence declines, and capturing this fact improves forecasting performances, sometimes dramatically as in the case of the US, sometimes by about 5–10% as in the case of European countries. The reason for these large differences across countries is that β_t shows relatively small variations in Europe, but shows large deviations in the US. This is simply a result of the differences in the levels of the flow rates in the US and in Europe, with the US having much larger flows: since $\beta_t = 1 - e^{-(f_t + s_t)}$ depends on the levels of the flow rates, the convergence rate is much more volatile in the US than in Europe (see Fig. A1 in Appendix).

Interestingly, note that for Spain, Japan or US, a model with fixed hazard rates (or only fixed convergence rate) can sometimes perform better in the long-run than our baseline FbF (particularly for 2 year ahead forecasts). This result points to the general difficulty in forecasting long-run behavior in the labor market. These three countries are characterized by strong trends in their inflow rate (Fig. 3), and a VAR without trend (as we use to forecast the worker flows) will mean-revert and may not capture these long-run trends. We see this result as highlighting the fact that our results are conservative with respect to the potential of a flow-based approach, and that there is still substantial room for improvement (in this case by a proper modeling of low-frequency movements).

Finally, Table 3 reports the performances of a simpler discrete-time version of FbF. The improvements compared to a standard stock-based VAR are still large, but we can see that the discrete-time approximation does penalize performance, indicating that the baseline FbF model is generally preferable.

4.3. Forecasting performance over the business cycle

Accurately forecasting unemployment is especially valuable during recessions and turning points. In this section, we show that FbF performs especially well, i.e., its performances relative to other models are even higher, during turning points and recessions, i.e., times of higher volatility.

To evaluate whether FbF performs differently over the course of the business cycle, we use the Giacomini–Rossi predictive ability test in unstable environments (Giacomini and Rossi, 2010). The test develops a measure of the relative local forecasting performance of two models and is ideal for testing whether the performance of our model varies over the cycle (compared to a benchmark model). We use as a benchmark an ARIMA model, estimated, like FbF, over a ten-year rolling window.²¹ We evaluate the local forecasting performance over a five-year window from quarterly forecasts.²²

Fig. 5 plots the Giacomini–Rossi fluctuation test for the one year-ahead forecast, along with the corresponding 5 percent critical value. The bold line shows the (standardized) rolling difference in mean-squared-error between the two models. This is measure of the relative performance; a positive value indicates a superior performance of FbF. The thin lines show the standard deviation of the inflow rate, ($s_t/E(u_t)$, dashed thin line), and outflow rate, (f_t , plain thin line).

We can see that FbF performs especially well around recessions, when the volatility of f and s is high. It does particularly well during the deep last recession in 2007–08 in all countries, and during times of large and swift movements in the inflow rate. In other words, FbF yields the greatest improvement over a univariate model around turning points, precisely when accurate unemployment forecasts are the most valuable.

4.4. Intuition for the model's performance

Our model's performance is particularly striking in two respects: (1) while improvements are largest at short forecast horizons in the US, we obtain large improvements at long horizons in European countries, and (2) the model performs especially well during turning points, i.e., during times of higher volatility.

While a theoretical exploration of the forecasting properties of FbF is behind the scope of this paper, in this section, we discuss the intuition behind the performances of FbF. We argue that the forecasting accuracy of FbF depends on three parameters: (i) the level of the flows, (ii) the persistence of the flows, and (iii) whether different flows have different time series properties.

4.4.1. Average performance

To better understand how the use of worker flow data can help (or not) improve forecasting performances, we will consider a simple illustrative experiment. Through this exercise, we will see that the value-added of using worker flow data depends on *both* the level and the persistence of the worker flows. In particular, FbF performs well over long horizons for European countries, because the flows in Europe are (i) small *and* (ii) sufficiently persistent.

We start at time $t=0$ with an unemployment rate u_0 at its steady-state value and with the inflow and outflow rates at their mean value. That is, we have $s_0 = Es_t$, $f_0 = Ef_t$ and $u_0 = Es_t/Es_t + Ef_t$. To simplify notations, denote $s = Es_t$ and $f = Ef_t$.

²¹ Appendix B Table B1 presents a table with the ARIMA models estimated for each country. Again, for each country, we selected the ARIMA model with the best average performance across forecast horizons.

²² Although the professional forecast would be an interesting benchmark, we use an ARIMA model because the Giacomini–Rossi test is only valid for models estimated over rolling-windows. (Both models are estimated over a ten-year rolling window.)




Fig. 5. Giacomini-Rossi fluctuation test: (FbF) vs ARIMA. (a) United Kingdom. (b) Germany. (c) France. (d) Spain. (e) Japan. (f) Unites States. Authors' calculations. *Notes:* “One-year ahead” is the relative performance of one-year ahead forecasts on the left axis, and “SD Outflow rate” and “SD Inflow rate” are the standard deviation of f and s on the right axis. Relative performance is the five-year rolling difference in MSE between forecasts from the FbF and ARIMA models. Standard deviations are calculated over five-year rolling windows. For clarity, the inflow rate was rescaled as $s_t/E(u_t)$. Dashed horizontal line indicates 5 percent critical value.

At time $t=1$, consider a one-time shock to the separation rate such that $s_1 = s + \varepsilon_1$ and then assume that the separation rate mean-reverts at some rate $\rho_s < 1$, i.e., the separation rate s follows an AR(1) process $s_{h+1} - s = \rho_s(s_h - s)$, $h > 0$. The job finding rate f_h is assumed to remain constant throughout the experiment.²³

The question is then the following: how is unemployment going to respond following this shock? If the response is short-lived, it means that the current separation rate contains information about unemployment in the near-term but little information about unemployment in the long-term. In contrast, if the response is persistent, the current value of the separation contains information about unemployment in the longer-term.

²³ Proceeding similarly, we could have considered a shock to the job finding rate.




Fig. 6. Impulse response function of unemployment to a separation rate shock. *Notes:* Impulse response function of unemployment to a unit shock to the inflow rate. “DE” refers to Germany ($f=0.06$) and “US” to the United States ($f=0.6$).

For a small shock (i.e., to a first-order), it is easy to show that $\psi(h) = d \ln u_{h+1} / d \ln s_1$, the impulse response of the unemployment rate to a shock to the separation rate, is given by

$$\frac{d \ln u_{h+1}}{d \ln s_1} \simeq \rho_s^{-h} (1 - e^{-f \cdot h}), \quad h \geq 0 \tag{10}$$

with $d \ln x_t = x_t - x/x$ and using $s \ll f$, which holds for all OECD countries (cf Table A2).

Starting from $\psi(0) = 1 - e^{-f}$, the impulse-response function is first increasing, reaches a maximum at h^* and then decreases to 0 with $\psi(h) \rightarrow 0$. Fig. 6 illustrates this impulse response function for two countries: the US with $f = 0.6$ and Germany with $f = 0.06$ (taking an autocorrelation $\rho_s = 0.85$ in both cases).

The maximum of the impulse response occurs at

$$h^* = 1 + \frac{1}{\beta} \ln \left(1 - \frac{\beta}{\ln \rho_s} \right) \tag{11}$$

with $\beta \approx f$ the (unemployment’s) rate of convergence to steady-state (the same β we encountered before).

Expression (11) is key to understand how both the level and the persistence of the flow affect the performance of a FbF model.

Since FbF uses the current worker flows as its main input, the “information content” of the current flow will determine the forecasting performance of FbF. With a shock to the separation rate having its largest effect on the unemployment rate after a time h^* , h^* indicates how much “information” the current worker flow rate s_1 contains about the value of unemployment h periods ahead. s_1 has a substantial effect on u_{h+1} for h close to h^* , that is, s_1 contains information about the value of unemployment h periods ahead. In contrast, as h becomes larger than h^* , the current separation rate contains less and less information about future unemployment u_{h+1} .

Looking at (11), the value of h^* depends on (i) the rate of convergence to steady-state (which itself depends on the average level of the outflow rate f) and (ii) the persistence of the inflow rate. h^* is larger if unemployment converges to steady-state more slowly ($\partial h^* / \partial \beta < 0$) or if the separation rate is more persistent ($\partial h^* / \partial \rho_s > 0$).

This thought experiment can thus help us understand why FbF offers improvements at long forecast horizons in the case of European countries, but only at short horizons in the case of US: the US labor market displays large worker flows while European labor markets have small flows.²⁴ Because of this large difference in average worker flows, $\partial h^* / \partial \beta < 0$ ensures that the impulse response of unemployment is more drawn out in the case of Germany (Fig. 6), and the current separation rate s_1 contains *more information* about one-year ahead unemployment for Germany than for the US.

Intuitively, in the case of European countries where the flows are small, convergence to u^* occurs slowly, and the current transition rates contain a lot of information about unemployment in the longer run, so that information on worker flows can help forecast unemployment in the longer run. In contrast, in the US, the current worker flow rates contain mostly information about the value of unemployment in the shorter-run.

²⁴ For the US, $f=0.6$, which implies $h_{US}^* \simeq 2$ quarters, but for Germany $f=0.06$, which implies $h_{DE}^* \simeq 4$ quarters (taking again an autocorrelation $e^{-\rho} = 0.85$ in both cases).

Moreover, Eq. (11) also shows that the persistence of the separation rate is important for the performance of FbF. With $\partial h^* / \partial \rho_s > 0$, the information content on the current flow increases with the persistence of the flow. Considering again the case of Germany, if the autocorrelation of the separation rate were only 0.6 (instead of 0.85 as in Fig. 6), we would get $h_{DE}^* = 1.8$ quarters, and the current flows would contain much less information about the value of one-year ahead unemployment.

Thus, the performances of FbF depend on the interaction between the speed of convergence (the levels of the flows) and the persistence of the flows, and FbF can help forecast unemployment in the longer-run *only if* the flows are (i) sufficiently small and (ii) sufficiently persistent. In other words, the fact that FbF offers such large improvements at long horizons for European countries was by no means guaranteed and owes to the fact that the flows in Europe are both small and sufficiently persistent.

4.4.2. Performance over the business cycle

The third important characteristic behind the performances of FbF stems from our focus on forecasting the flows rather than the stock. A model of the stock (such as a stock-based VAR model as evaluated previously) cannot perform as well as FbF, because the flows have different time series properties,²⁵ and because the contribution of the different flows changes throughout the cycle (Barnichon, 2012). 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.

A nice illustration of this property can be seen with the asymmetric nature of unemployment fluctuations, and the fact that FbF performs especially well (relative to a stock-based model) during recessions and around turning points.




Fig. 7. Impulse responses to a one standard deviation shock to the unemployment inflow rate^a. Authors calculations.^aImpulse response functions to a 1-standard deviation shock to the inflow rate, calculated from a VAR of $y_t = \ln(f_t, s_t, ur_t)'$ with one lag estimated using quarterly data over the sample with available data. See Table 1.

²⁵ For instance, for virtually all countries in our sample, the autocorrelation of the outflow rate is much higher than that of the inflow rate (Table A2 in the Appendix). There are other differences. For instance, in the case of the US, while the distribution of the (detrended) inflow rate is positively skewed and highly kurtotic, the distribution of the (detrended) outflow rate exhibits low kurtosis and no skewness (see Barnichon, 2012).

The unemployment rate displays steepness asymmetry – that is, increases are steeper than decreases. This asymmetry manifests itself most forcefully during recessions, but a stock-based model such as a (linear) VAR model cannot capture it. While FbF is not explicitly asymmetric, it relies on the worker flows that are responsible for the asymmetry of unemployment (Barnichon, 2012).

Indeed, the beginning of a recession is typically marked by a sharp increase in the inflow rate, and Fig. 7 plots the impulse responses from our estimated VAR to a shock to the inflow rate in 3 representative countries: the UK, Germany and France. While the inflow rate displays a sharp increase with relatively quick mean reversion, the outflow rate displays a delayed hump-shaped response with much slower mean reversion. These different impulse responses are behind the steepness asymmetry of unemployment. Following the initial shock, the inflow rate reverts relatively quickly to its mean. However, the outflow rate takes a lot longer to mean-revert and thus prevents the unemployment rate from decreasing as fast as it increased and generating asymmetry in steepness. By relying on a VAR forecast of the flow rates, following an initial disturbance to the inflow rate at the onset of a recession, a flow-based model like FbF can propagate the cyclical behavior of the flows and thus capture the steepness asymmetry of unemployment. In contrast, a stock-based model cannot capture the asymmetric nature of unemployment fluctuations and will perform worse in recessions.

5. Conclusion

This paper evaluates a novel “flow approach” to unemployment forecasting for six large OECD countries with very different labor market dynamics. We find that the “flow approach” yields very large improvements in forecast accuracy, with large reductions in the mean-squared errors of the best alternative forecasting models. Improvements occur mainly at large horizons (one-year ahead forecasts) for countries characterized by small labor market flows (e.g., Spain or Germany), whereas improvements occur at short horizons (one-quarter ahead forecasts) for countries characterized by small labor market flows (the US). For all countries, improvements are especially large during recessions and turning points, when unemployment forecasts are most valuable.

An important advantage of the “flow approach” to unemployment forecasting is its small data requirements, which make the method applicable for a very large range of countries. In fact, the unemployment flow data underlying our approach have already been calculated for as many as 70 developed and developing countries (Viegelahn and Wieser, 2013). Thus, provided that the flow data are available in real time with reasonable time lags, our method could in principle be applied to as many as 70 countries.

Finally, the large improvements in forecasting performances were obtained with simple VAR-based forecasts of the worker flows. More elaborate techniques may produce better flow forecasts and open the door to even better unemployment forecasts. Exploring such possible improvements would be a fruitful avenue for future research.

Acknowledgements

We thank Bart Hobijn, Juan Dolado, Kris Nimark, Barbara Petrongolo, Barbara Rossi, Yanos Zylberberg and seminar participants for helpful comments, and we gratefully acknowledge Matthias Hertweck, Hiro Miyamoto, Jennifer Smith and Lucile Viel for sharing their transition rates data. Any errors are our own. Barnichon acknowledges financial support from the Spanish Ministerio de Economía y Competitividad (grant ECO2011-23188), the Generalitat de Catalunya (grant 2009SGR1157) and the Barcelona GSE Research Network.

Appendix A. Time series properties

Table A1

Cross-correlation between unemployment rate and steady-state unemployment rate.

i	$u, u^*(-i)$						
	Germany	Spain	France	UK	Japan	US	
0	0.839	0.905	0.819	0.958	0.765	0.959	
1	0.890	0.935	0.862	0.972	0.691	0.982	
2	0.909	0.918	0.860	0.972	0.675	0.955	
3	0.900	0.895	0.824	0.960	0.608	0.905	
4	0.875	0.857	0.748	0.938	0.448	0.841	

Note: Cross-correlation $u, u^*(-i)$ using quarterly data over 1992q1–2011q4. Source: Authors calculations.

Table A2
Time series properties.

	UK	Germany	France	Spain	Japan	US
f_t						
$E(X)$	0.13	0.06	0.08	0.07	0.12	0.58
$sd(X)$	0.03	0.01	0.01	0.02	0.01	0.15
$\rho(X)$	0.97	0.91	0.94	0.96	0.74	0.89
R^2	0.97	0.80	0.93	0.95	0.55	0.82
$s_t/E(u_t)$						
$E(X)$	0.13	0.06	0.09	0.08	0.13	0.58
$sd(X)$	0.01	0.01	0.02	0.02	0.03	0.097
$\rho(X)$	0.41	0.84	0.90	0.85	0.70	0.91
R^2	0.65	0.84	0.83	0.73	0.50	0.83

Notes: Quarterly data over 1992q1–2011q4. Source: Authors calculations.

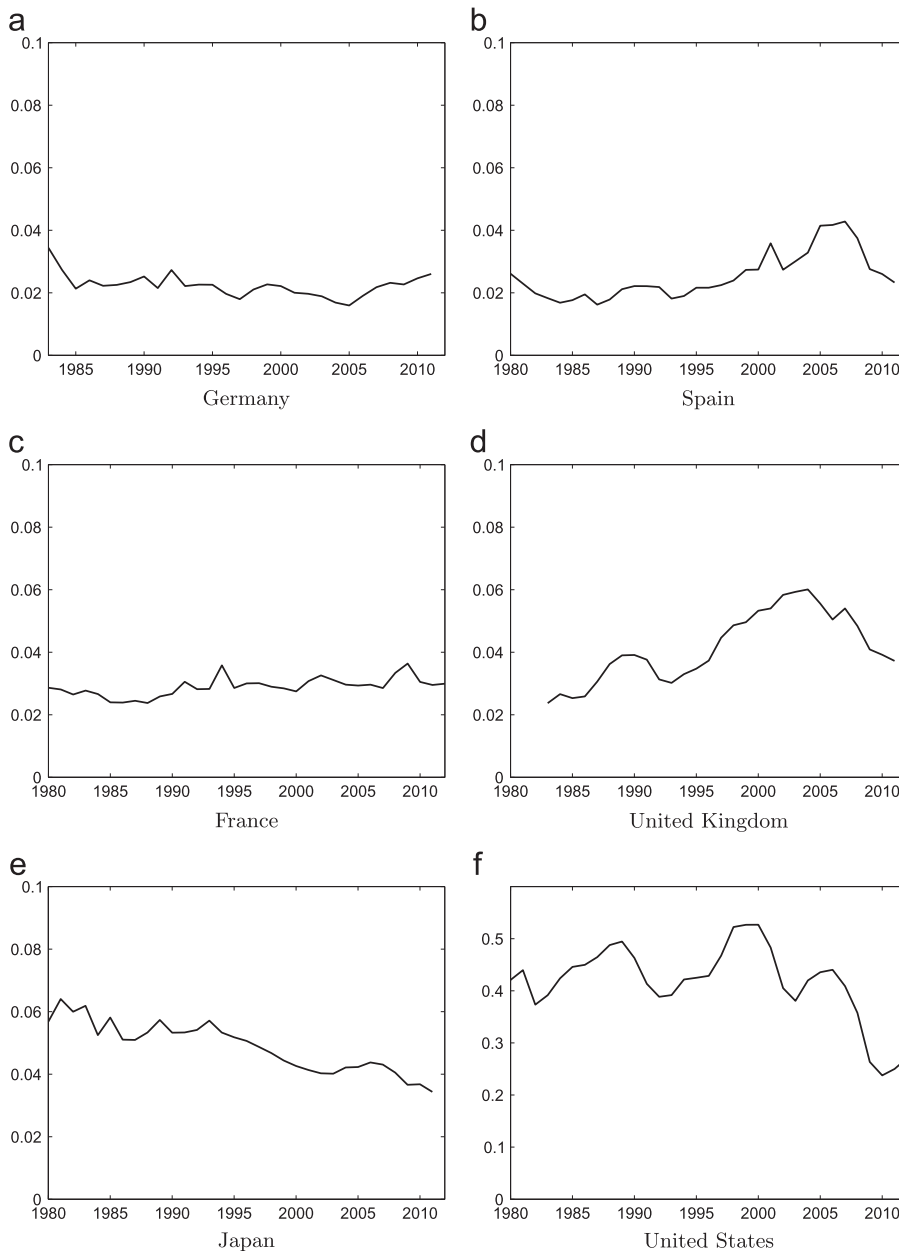


Fig. A1. Convergence rate by country. Note: Convergence rate is $\beta_t \equiv 1 - e^{-(s_t + f_t)}$. The y-scale is substantially larger for US data. Source: Authors calculations based on [Elsby et al. \(2013\)](#) and National Statistics Institutes data.

Appendix B. Model specifications

This section lists the different specifications used in the ARIMA model (Section 4.4.2) and in the VAR model (Sections 2.2.1 and 4.2.2). For the VAR, in addition to the unemployment rate and worker flows, we included leading indicators of the labor market: ui , the number of claims for unemployment insurance each period, ν the job openings or vacancies in each period, and $\Delta \ln GDP$ the growth gross domestic output. The inclusion of variables is restricted by data availability. For the US we also use uic , the monthly average of weekly initial claims for unemployment insurance and hwi , the composite help-wanted index constructed by Barnichon (2010).²⁶

Table B1
ARIMA models for the UR.

Country	Model
UK	ARIMA(2,0,0)
Germany	ARIMA(1,0,1)
France	ARIMA(1,0,1)
Spain	ARIMA(2,0,1)
Japan	ARIMA(2,0,1)
US	ARIMA(2,0,1)

Table B2
VAR specifications.

Country	Variables	Lags	Estimation period
UK	$\ln f_t, \ln s_t, \Delta \ln u_t, \Delta \ln \nu_t,$ $\Delta \ln gdp_t, \Delta \ln ui_t$	1 lag	1982Q2–2011Q4
Germany	$\ln f_t, \ln s_t, \Delta \ln u_t, \Delta \ln \nu_t,$ $\Delta \ln gdp_t, \Delta \ln ui_t$	1 lag	1985Q1–2011Q4
France	$\ln f_t, \ln s_t, \Delta \ln u_t, \Delta \ln gdp_t$	1 lag	1977Q1–2011Q4
Spain	$\Delta \ln f_t, \ln s_t, \Delta \ln u_t, \Delta \ln \nu_t,$ $\Delta \ln gdp_t, \Delta \ln const_t, \Delta \ln ere_t$	1 lag	1982Q1–2011Q4
Japan	$\ln f_t, \ln s_t, \Delta \ln u_t, \Delta \ln \nu_t$	2 lags	1980Q1–2011Q4
US	$\ln f_t, \ln s_t, \Delta \ln u_t, \ln uic_t, \Delta \ln hwi_t$	2 lags	1977M1–2011M12

Appendix C. Supplementary data

Supplementary data associated with this article can be found in the online version at <http://dx.doi.org/10.1016/j.euroecorev.2015.10.006>.

References

- Barnichon, Regis, 2010. Building a composite help-wanted index. *Econ. Lett.* 109 (3), 175–178.
- Barnichon, Regis, 2012. Vacancy posting, job separation and unemployment fluctuations. *J. Econ. Dyn. Control* 36 (3), 315–330.
- Barnichon, Regis, Andrew Figura, 2013. Declining labor force attachment and downward trends in unemployment and participation.
- Barnichon, Regis, Christopher, Nekarda, 2012. The ins and outs of forecasting unemployment: using labor force flows to forecast the labor market. *Brookings Papers on Economic Activity*.
- Batchelor, Roy, 2010. The IMF and OECD versus consensus forecasts. *Appl. Econ.* 33 (2), 225–235.
- Brown, Laura, Moshiri, Saeed, 2004. Unemployment variation over the business cycles: a comparison of forecasting models. *J. Forecast.* 23 (7), 497–511.
- Diebold, Francis X, Mariano, Roberto S., 1995. Comparing predictive accuracy. *J. Bus. Econ. Stat.* 13 (3), 253–263.
- Elsby, Michael W.L., Bart, Hobijn, Ayşegül, Şahin, May 2013. Unemployment dynamics in the OECD, *The Review of Economics and Statistics*, MIT Press, vol. 95(2), pages 530–548, May.
- Fujita, S., Ramey, G., 2012. Exogenous vs. endogenous separation. *Am. Econ. J.: Macroecon.* 4 (4), 68–93.
- Giacomini, Raffaella, White, Halbert, 2006. Tests of conditional predictive ability. *Econometrica* 74 (6), 1545–1578.
- Giacomini, Raffaella, Rossi, Barbara, 2010. Forecast comparisons in unstable environments. *J. Appl. Econom.* 25 (4), 595–620.
- Golan, Amos, Perloff, Jeffrey M., 2004. Superior forecasts of the U.S. unemployment rate using a nonparametric method. *Rev. Econ. Stat.* 86 (1), 433–438.
- Hertweck, Matthias S., Oliver Sigrist, 2012. The Aggregate Effects of the Hartz Reforms in Germany. Department of Economics, University of Konstanz Working Paper Series of the Department of Economics, University of Konstanz 2012-38.
- International Labour Organization, ILO, 2015. World Employment and Social Outlook: Trends 2015.

²⁶ In the case of Spain we also added output in the construction sector (labelled *const*) and a “Record of Employment Regulation” (labelled *ere*) to account for changes in labor market regulations in Spain.

- Lin, Ching-Yang, Miyamoto, Hiroaki, 2012. Gross worker flows and unemployment dynamics in Japan. *J. Jpn Int. Econ.* 26 (1), 44–61.
- Milas, Costas, Rothman, Philip, 2008. Out-of-sample forecasting of unemployment rates with pooled STVECM forecasts. *Int. J. Forecast.* 24 (1), 101–121.
- Montgomery, Alan L., Zarnowitz, Victor, Tsay, Ruey S., Tiao, George C., 1998. Forecasting the U.S. unemployment rate. *J. Am. Stat. Assoc.* 93 (442), 478–493.
- Nekarda, C., 2009. Understanding Unemployment Dynamics: The Role of Time Aggregation, working paper, available at www.chrisnekarda.com. (<http://www.chrisnekarda.com>).
- Petrongolo, B., Pissarides, C.A., 2008. The ins and outs of European unemployment. *Am. Econ. Rev.* 98 (2), 256–262.
- Rothman, Philip, 1998. Forecasting asymmetric unemployment rates. *Rev. Econ. Stat.* 80 (1), 164–168.
- Shimer, Robert, 2012. Reassessing the ins and outs of unemployment. *Rev. Econ. Dyn.* 15 (2), 127–148.
- Smith, Jennifer C., 2011. The ins and outs of UK unemployment. *Econ. J.* 121 (552), 402–444.
- Solon, Gary, Michaels, Ryan, Elsby, Michael.W.L., 2009. The ins and outs of cyclical unemployment. *Am. Econ. J.: Macroecon.* 1 (1), 84–110.
- Viegelahn, Christian, Christina Wieser, 2013. Key Indicators of the Labour Market: The ILO Estimates on Unemployment Flows.