THE RIMINI CENTRE FOR ECONOMIC ANALYSIS

WP 39_13

Mark J. Holmes
Waikato University, New Zealand

Jesús Otero
Universidad del Rosario, Colombia

Theodore Panagiotidis
University of Macedonia, Greece
The Rimini Centre for Economic Analysis (RCEA), Italy

MODELLING THE BEHAVIOUR OF UNEMPLOYMENT RATES IN THE US OVER TIME AND ACROSS SPACE

The Rimini Centre for Economic Analysis
Legal address: Via Angherà, 22 – Head office: Via Patara, 3 - 47900 Rimini (RN) – Italy
www.rcfea.org - secretary@rcfea.org

Modelling the behaviour of unemployment rates in the US over time and across space*

Mark J. Holmes†
Department of Economics
Waikato University
New Zealand

Jesús Otero‡
Facultad de Economía
Universidad del Rosario
Colombia

Theodore Panagiotidis§
Department of Economics
University of Macedonia
Greece

June 2013

Abstract

This paper provides evidence that unemployment rates across US states are stationary and therefore behave according to the natural rate hypothesis. We provide new insights by considering the effect of key variables on the speed of adjustment associated with unemployment shocks. A highly-dimensional VAR analysis of the half-lives associated with shocks to unemployment rates in pairs of states suggests that distance between states and vacancy rates respectively exert a positive and negative influence. We find that higher homeownership rates do not lead to higher half-lives. When the symmetry assumption is relaxed through quantile regression, support for the Oswald hypothesis through a positive relationship between homeownership rates and half-lives is found at the higher quantiles.

JEL Classification: E240, J600, F150, R100.

Keywords: Unemployment; market integration; speed of adjustment.

---

*We would like to thank three anonymous referees for their useful comments and suggestions on an earlier version of the paper. Any remaining errors are our own responsibility.

†E-mail: holmesmj@waikato.ac.nz
‡E-mail: jesus.otero@urosario.edu.co
§E-mail: tpanag@uom.edu.gr

1 Introduction

The different behaviour of the unemployment rate across regions has been a keen area of research for many years. One important direction of this research has been the examination of the stationarity properties of the unemployment rate. While output variations drive fluctuations in unemployment over the business cycle, two competing hypotheses offer to explain how fluctuations in the unemployment rate might bear out in the long-run. Under the *natural rate hypothesis*, unemployment rate fluctuations are temporary suggesting that unemployment will return back to the natural rate in the long-run (mean reverting process). Under the *hysteresis hypothesis*, these same unemployment fluctuations have permanent effects on unemployment on account of labour market rigidities. Either way, regional unemployment rates reflect the regional economic structure and examining their time series properties is crucial for understanding the impact from shocks, and for analysing the regional impacts resulting from stabilisation policy.

A further dimension to existing research is how regional unemployment rates are related, how quickly they adjust towards equilibrium, and how (potential common) drivers might influence regional speeds of adjustment. Understanding the relative speeds of adjustment of state unemployment rates to shocks is not only an issue of interest from the perspective of relative regional prosperity or for policymakers, but also because it offers insights into labour market mismatch or efficiency. In assessing potential drivers, Oswald [34] has suggested that since home owners are relatively less mobile across geographic locations than renters, regional home ownership rates are positively correlated with the level of regional unemployment rates. This is often referred in the literature as the Oswald hypothesis; for a more up to date discussion see Blanchflower and Oswald [2] for the US states and Laamanen [24] for Finland. Furthermore, there is also potential insight based on the Beveridge curve, according to which high vacancy regions might be characterised by a lower level of unemployment if mismatch is low.

In this paper we address two key questions. First, does the natural rate or hysteresis hypothesis best describe the behaviour of US state unemployment rates? In other words, is unemployment mean reverting? Noting the lack of consensus in the existing literature, we adopt a time-series approach to examine the integration properties of the data. This involves applying two panel data tests, namely the Pesaran [41] test for joint non-stationarity and the Hadri [21] test for joint stationarity. Both tests allow for the presence of cross-sectional dependency, which is a critical issue because of the nature of the data. If unaccounted for, this can lead to severe size distortion and therefore invalid statistical inference. Our results point to the presence of (mean) stationary unemployment rates across US states, and therefore offer support for the natural rate hypothesis. This finding is consistent with unemployment rates that have converged insofar as they are not drifting further apart over the long-run. However, a crucial issue concerns the speed at which unemployment rates adjust towards equilibrium. Indeed, differing states may be characterised by very different speeds of adjustment in response to, say, a large negative demand shock. This context leads to the second question in which we are interested: Can the expected heterogeneity in the speed at which unemployment rates adjust to shocks be explained by (potential common) drivers such as distance, homeownership and vacancy rates? This is a question that is addressed within a cross-section framework. At this stage it is worth noticing that the empirical approach adopted in the paper is in sharp contrast to that of the existing literature, which instead focuses on the effects of the determinants on the level of unemployment indicated above.

This paper seeks to further enlighten our understanding of the interdependencies between state unemployment rates. The distinguishing feature of our empirical modelling exercise is that we combine both time-series and cross-section perspectives. Indeed, we start off by estimating a highly-dimensional vector autoregressive (VAR) model comprising 48 state unemployment rates observed (on a monthly frequency) over a time period of more than 35 years. The resulting VAR model is used to compute generalised impulse-response functions (GIRFs), which in turn are then employed to provide a measure of the corresponding half-lives of shocks to the unemployment rates. Next, the estimated half-lives are analysed within a cross-section framework to consider the factors that determine the speed at which unemployment rate shocks are transmitted across states. Our results point towards positive and negative roles for distance and relative state vacancies in driving half-lives. Despite the predictions of the Oswald hypothesis concerning the level of unemployment, our initial findings (under the assumption of a symmetric response around the conditional mean) do not suggest that higher homeownership leads a reduction in the speed of adjustment in unemployment rates. However, in a further contribution to the existing literature, we provide an exploration of the role of these variables through the use of quantile regression analysis (which allows for an asymmetric response around the conditional moments) which suggests that higher homeownership leads a reduction in the speed of adjustment at the higher quantiles only.

The rest of the paper is organised as follows. Section 2 presents a brief survey of the literature. Section 3 describes the US unemployment data and discusses their time series properties. Section 4 reports the results of a VAR-based modelling exercise of the unemployment rates. Here, the analysis turns around the computation of the half-lives of shocks, from which a (cross-section) model is subsequently estimated with the purpose of finding the main factors that help explain their behaviour. The final section offers some concluding remarks.

2 Brief literature review

Nelson and Plosser [31] argued that most macroeconomic series (including the unemployment rate) become stationary after first differencing. This seminal paper has produced a vast literature that has focused on the stationarity properties of the series. Numerous studies also address the natural rate versus hysteresis hypothesis for the US include Song and Wu [46], Payne et al. [38], León-Ledesma [27], Clemente et al. [8] and Cheng et al. [7] among others. These studies can only offer mixed evidence on the nature of fluctuations in US state unemployment rates. For example, Payne et al. [38] find that all state unemployment rates are non-stationary based on univariate augmented Dickey and Fuller (ADF) [14] unit root and variance ratio testing. Further support for hysteresis is provided by León-Ledesma [27], who applies the (subsequently published) panel unit root test by Im et al. [25] to test for unemployment hysteresis in the US states and the EU countries, against the alternative of a natural rate. The results show that hysteresis for the EU and the natural rate for the US states are the most plausible hypotheses. In contrast to this, Song and Wu [46] reject joint non-stationarity in favour of the natural rate hypothesis once they move from univariate to panel data unit root testing based on Levin and Lin [28]. Clemente et al. [8] offer results that are, in general, favourable to the rejection of the presence of a unit root. However, they observe that this conclusion is clearly qualified by the level of disaggregation employed. Most recently, Cheng et al. [7] investigate the stochastic nature of the unemployment rate allowing for cross-section dependence from a panel of US state-level data. They find significant evidence of a non-stationary common component when the data from the most recent recession are included. Even when stationarity is empirically supported, the bias-corrected half-life of the common component appears very long, casting doubt on the usefulness of the natural rate hypothesis. Cover and Mallick [12] employ a battery of univariate procedures that include the familiar ADF tests and the Phillips-Perron (PP) [39] tests and show that the national US unemployment rate contains a unit root. This finding is further supported by the more powerful Elliott et al. (ERS) [15] test as well as Kwiatkowski et al. (KPSS) [26] stationarity test. In further comparing with the EU, Chang et al. [6] employ ADF unit root testing within a seemingly unrelated regression framework and conclude that the hysteresis hypothesis is supported for all ten European countries in their sample except for Belgium and the Netherlands. In a recent paper, Abadir et al. [1] revisit Nelson and Plosser [31] using an autocorrelation-function-based approach, and find that most macroeconomic variables (including the unemployment rate) are not integrated processes.

The literature that addresses the dispersion between US state unemployment rates includes Chalmers and Greenwood [5], Bronars and Jansen [3], Partridge and Rickman [36] and [37], Vedder and Gallaway [48], Payne et al. [38], Nissan and Carter [32], Conley and Topa [9] and Nistor [33] among others. A variety of perspectives and conclusions are offered by these studies. For example, Bronars and Jansen [3] present estimates of the time series and spatial pattern of unemployment rate fluctuations in the United States. They consider the impact of a temporary unemployment rate shock on unemployment rates in adjacent areas. Partridge and Rickman [36] and [37] find that unemployment between states is persistent on account of local economic and political factors that influence state behaviour. These influences include industrial, regional, demographic and non-demographic factors. Using bivariate Engle-Granger cointegration tests, Payne et al. [38] find evidence of cointegration between state and national unemployment rates in only two out of fifty cases. Conley and Topa [9] find some evidence that there is a positive and statistically significant degree of spatial dependence in the distribution of raw unemployment rates. Lastly, Chalmers and Greenwood [5] and Nistor [33] find evidence of asymmetries across states in unemployment behaviour.

In their survey, Havet and Penot [23] conclude that Oswald's hypothesis finds little support. This is continued in recent studies such as Farber [17] who, analysing the US labour market in the Great Recession and its aftermath, notes the great difficulty in unemployment falling. Looking at mobility rates for unemployed homeowners and renters, there is no support for the idea that the housing market crisis has reduced mobility of the unemployed. In earlier work, Coulson and Fisher [11] conclude that while individual homeowners may have inferior labour market outcomes as compared to renters, from the viewpoint of society, higher homeownership rates may actually result in greater job creation and overall production, among other benefits. Munch et al. [30] investigate the effects of homeownership on labour mobility and unemployment duration. In distinguishing between finding employment locally or being geographically mobile, they find that homeownership hampers the propensity to move for job reasons, but improves the chances of finding local jobs. Their empirical findings thus contradict the so-called Oswald hypothesis, even if support is found for the main mechanism behind the hypothesis, namely that homeownership hampers mobility. Coulson and Fisher [10] argue that homeowners, conditionally or unconditionally, have better labour market outcomes than renters. More recently, Blanchflower and Oswald [2] revisit the relationship between homeownership and unemployment using panel data from US states. They conclude that high levels of home-ownership tend to destroy jobs with a one year lag. Laamanen [24] employs a unique dataset from Finland, a country that undertook a rental housing market deregulation reform; this policy reform is a feature that allows the author to overcome potential endogeneity issues. The main conclusion of this study is that an increase in home-ownership is associated with higher unemployment.

3 US unemployment data and their time-series properties

We employ monthly seasonally adjusted data on unemployment rates for 48 US states. The data, which are freely available from the Federal Reserve Economic Data (FRED) of the Federal Reserve Bank of St. Louis, cover the period between 1976m1 and 2012m7, which yields a total of 439 time observations for each state. Alaska and Hawaii are excluded from our analysis on the grounds that these two states are not geographically contiguous with any other state in the US, so that some of the mechanisms that may underpin the relationship of unemployment rates across states within the US may not operate in these cases. The range of variation (i.e. the difference between the maximum and minimum values) of the unemployment rates of the 48 states under consideration are plotted in Figure 1, which for the purposes of the empirical analysis are considered as a fraction of the state labour force. The study period includes periods of strong and steady growth. It also includes the recessions experienced in the early 1980s, 1990s and 2000s as well as the most recent great recession and financial crisis episodes. This figure supports the view that discrepancies in state unemployment rates have been somewhat noticeable over the sample period. Indeed, the range of variation of the unemployment rates reveals that the extent of dispersion has fluctuated between 14.2 and 13.5 per cent for West Virginia (WV) and Michigan (MI), respectively, and 4.2 and 3.5 per cent for North Dakota (ND) and South Dakota (SD), respectively. Clearly, this constitutes a sizable spread in relation to the mean unemployment rate across states. In addition, it is the case that many states exhibit consistently high or low unemployment rates. This might be symptomatic of variations in the natural rate of unemployment across states. Against this background, states might experience similar unemployment fluctuations around their own high or low natural rates.

Figure 2 displays the evolution through time of the cross-sectional range of variation of the unemployment rate, and its relation to the business cycle, where the latter is represented in the figure (in the shadowed area) by means of the National Bureau of Economic Research (NBER) official peak to trough periods. As can be seen from this plot, while the state unemployment rates have generally moved inversely with economic performance, they also appear to become more dispersed during the course of a recession, and generally less dispersed during the stronger economic periods. Since dispersion rises during the course of a recession, this is consistent with state unemployment rates rising at different speeds. The great recession sees the largest increase in range which could be consistent with the greatest speed variation compared to the earlier recessions; interestingly, the spikes that are observed in the plot in September, October and November 2005 are due to (short-run) effects of Hurricane Katrina in the Louisiana (LA) and Mississippi (MS) unemployment rates.

We begin our empirical investigation with an analysis of cross-sectional dependence of innovations (shocks) in unemployment rates. To do this, we calculate the Pesaran [40] general diagnostic test for cross section dependence in panels, denoted as the CD statistic. To implement this test, we first fit an ADF regression for each cross section unit i separately, using p lags of the dependent variable, and denote the resulting residuals as individual series $\hat{e}_{it}$. The importance of estimating an ADF(p) model is that it allows us to remove any serial correlation pattern in the individual time series i. Second, we compute the cross-correlation coefficient between the residuals of cross section units i and j as:

$$\hat{\rho}_{ij} = \frac{\sum_{t=1}^{T} \hat{e}_{it}\hat{e}_{jt}}{\left(\sum_{t=1}^{T} \hat{e}_{it}^2\right)^{1/2} \left(\sum_{t=1}^{T} \hat{e}_{jt}^2\right)^{1/2}}. \tag{1}$$

Finally, we calculate the CD statistic as:

$$CD = \sqrt{\frac{2T}{N\left(N-1\right)}} \left(\sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \hat{\rho}_{ij} \right) \sim N\left(0,1\right). \tag{2}$$

Table 1 summarises the results of applying the CD statistic to the 48 unemployment rates using $p = 1, 2 \ldots, 6$ lags in the ADF(p) regressions. As can be seen from the table, the null hypothesis that the series innovations are cross sectionally independent is strongly rejected for all augmentation orders.

Next, the order of integration of the unemployment rates is investigated using panel unit root (stationarity) tests. Univariate tests (such as the ADF) suffer from disappointing power properties, as reported, for instance, by DeJong et al. [13]. In an attempt to overcome this deficiency, testing for unit roots (stationarity) in panel data has been considered as a possible way to achieve power gains over unit root (stationarity) tests applied to a single time series. This is because panel data combine information from both the time-series and the cross-section dimension and thus require fewer time observations for the tests to exhibit power.

Among the panel unit root tests that are available in the literature, the one proposed by Im et al. [25], commonly referred to as the IPS test, has proved to be popular probably due to the fact that it is straightforward to compute, being based on the simple average of ADF(p) statistics obtained across the N members of the panel. However, a critical assumption underlying the IPS test is that of cross section independence among the individual time series in the panel. Indeed, the IPS test suffers from severe size distortions in the presence of cross section dependence, as shown in the Monte Carlo simulations carried out by Pesaran [41]. To account for cross section dependence, the latter suggests modifying the IPS test by augmenting the standard ADF(p)-type regression with the cross section averages of lagged level and lagged first-differences of the individual series in the panel. Thus, the test of the unit root hypothesis would be based on the following regression model:

$$\Delta y_{it} = a_i + b_i y_{i,t-1} + \sum_{r=1}^{p_i} c_{ir} \Delta y_{i,t-r} + d_i \bar{y}_{t-1} + \sum_{r=0}^{p_i} f_{ir} \Delta \bar{y}_{t-r} + \varepsilon_{it}, \tag{3}$$

where $\bar{y}_t = (N)^{-1} \sum_{i=1}^{N} y_{it}$, $i = 1, ..., N$ and $t = 1, ..., T$. The corresponding cross-sectionally augmented version of the IPS test statistic, denoted CIPS, is given by:

$$\text{CIPS} = (N)^{-1} \sum_{i=1}^{N} t_i, \tag{4}$$

where t_i is the ADF statistic based on the regression t-statistic for testing $Ho : b_i = 0$ in equation (3). The CIPS statistic is similar to the IPS in that they both test the joint null of a unit root, against the alternative of at least one stationary series in the panel. Pesaran [41] tabulates critical values for the most commonly used specifications of the deterministic components in the test regressions (namely no constant and no trend; constant and no trend; and constant and trend), but only for a limited range of time observations (T) and a limited number of cross-sectional units (N) in the panel; Otero and Smith [35] estimate response surfaces for the CIPS tests that can also be used to calculate approximate p–values for these tests.

In addition to the CIPS test, we also apply the Hadri [21] test for the null hypothesis that all individual series are stationary (either around a level or a deterministic time trend), against the alternative of at least a single unit root in the panel. This test, which is the panel version of the well-known univariate KPSS stationarity test, offers the advantage that if the null hypothesis is not rejected, then one may conclude that all the unemployment rates in the panel are stationary.

To compute the Hadri test, we start off by calculating the KPSS statistic for which a critical issue is how to approximate the so called long-run variance. For the latter, we follow Monte Carlo simulation results reported by Carrión-i-Silvestre and Sansó [4], which recommend estimation of this key parameter applying the boundary condition rule advocated by Sul et al. [47]. An additional important issue is related to the fact that, similar to the IPS test, the Hadri [21] test suffers from severe size distortions in the presence of cross-sectional dependence, the magnitude of which increases as the strength of the cross-sectional dependence increases; see Giulietti et al. [19]. To account for cross section dependence, we perform an autoregressive-based bootstrap, implemented as described in Hadri and Rao [22].

The overall picture that emerges from applying the panel unit root and stationarity tests is that the unemployment series under consideration can be best described as stationary processes. On the one hand, the CIPS test, incorporating individual-specific intercepts, rejects the joint null hypothesis of a unit root at least at a 10% significance level, a conclusion that appears robust to the number of lags of the dependent variable that are included in the test regression (see Table 2). The panel stationarity tests on the other, are presented in Table 3 (Hadri test), along with their corresponding bootstrap p-values in brackets, which in turn are based on 10,000 replications used to derive the empirical distributions of the test statistics. As can be seen, the evidence reported in Table 3 does not reject the null hypothesis of stationarity, supporting the view that the unemployment rates under consideration are mean reverting. [1] This finding is important not only because it means that shocks have only temporary effect on unemployment rates, but also because it enables us to have an estimate of the mean reverting level for the different unemployment rates, and see how close they are to the 6.5% level set by the federal Reserve as an intermediate target of monetary policy. This comparison is made in Figure 3, which shows that there are 30 states within $\pm 1\%$ of the unemployment intermediate target rate set by the Federal Reserve. The overall stationarity finding is also in line with Abadir et al. [1].

4 VAR-based modelling of US unemployment rates and the determinants of their speed of adjustment

This section examines the dispersion of unemployment rates over time and across markets. The specific question we aim to answer is whether the geographical separation of states constitutes a factor that helps explain the dynamic behaviour of unemployment rates. Our empirical investigation is based on the estimation of a

[1]Perhaps it is interesting to observe that if one incorrectly assumes cross-sectional independence, and therefore compares the calculated Hadri statistic against critical values derived from a standard normal distribution, then the null hypothesis would be incorrectly rejected.

VAR model consisting of the unemployment rates in the 48 US states under consideration. The use of a VAR-based modelling approach can be justified on two grounds. First, the results of the preliminary analysis of the integration properties of the data supported the view that the unemployment series under investigation could be best characterised as stationary over the study period. Indeed, the finding of stationarity precludes the possibility that cointegrating relationships may exist among the unemployment rates, and therefore makes the series suitable for modelling within a standard VAR framework. Second, it is unlikely to be able to identify a specific state, let us say k, that is dominant in the sense that shocks to it propagate to the other states, while shocks to the remaining states have little effect on k. Thus, the adoption of a VAR-based approach appears appealing as it offers the advantage that all unemployment rates may be treated as potentially endogenous variables; see Sims [45].

Once the VAR model has been estimated, we examine the speed at which unemployment rates adjust to exogenous shocks or innovations, for which we employ half-life estimates based on impulse response functions. In the case of a simple AR(1), the half-life of a shock can be estimated using the formula $-\ln(2)/\ln(\delta)$, where δ indicates the value of the autoregressive coefficient. However, Goldberg and Verboven [20], in footnote 11, observe that in case of more complicated processes, such as a higher order AR(p) process or an ARMA(p,q) process, the previous formula is no longer valid, and thus impulse response functions should be preferred (see also Seong et al. [44]). Bearing this aspect in mind, we employ the GIRFs proposed by Pesaran and Shin [42], which are invariant to the way shocks in the underlying VAR model are orthogonalised. Thus, GIRFs offer an extension to the traditional impulse response analysis, which is sensitive to the ordering of the variables included in the VAR; see Lütkepohl [29].[2]

An important initial stage in the analysis of VAR models is the selection of the optimal lag order. In this paper, the underlying VAR model is estimated by ordinary

[2] More insight into the application of this methodology can be found in studies such as Ewing and Thompson [16].

least squares (OLS), including a constant as deterministic component, for $p_{max} = 4$ lags, and then the appropriate lag length is determined using the AIC, the SIC, and the GTS algorithm. As expected, the lag orders selected by the AIC and GTS algorithm (i.e. 2 and 4, respectively) are larger than that selected by the SIC (i.e. 1). However, when using the SIC some of the fitted AR equations in the VAR model suffer from residual serial correlation, and so it seems prudent to choose the higher lag order specifications favoured by either the AIC or the GTS algorithm. Taking into consideration that the number of time observations that are available, that is 439, might easily turn out to be too small relative to the number of variables in the VAR model, that is 48, we opt for estimating the model using $p = 3$ lags as a compromise; it is worth pointing out that the results obtained for alternative lag orders are qualitatively the same.[3]

Having selected the optimal order p , the estimated model can be used to identify which state unemployment rate(s) Granger-causes the remaining ones. The appealing feature of this is that one might argue that Federal Reserve policy-makers could target a particular (state) unemployment rate. To perform the Granger causality tests, given the large number of variables included in the VAR model, we focus on the unemployment rates of the four largest states, as judged by their corresponding average share on national GDP over the period 2008 to 2011. With this criterion in mind, it turns out that these states are California (CA), Texas (TX), New York (NY) and Florida (FL), which account for about 13.2%, 8.5%, 7.7% and 5.2%, respectively. The results reported in Table 4 indicate that the four largest states Granger-cause (at the 10% significance level) the unemployment rate in 25 out of 44 states. In addition, when one examines whether the individual largest states Granger-cause the remaining ones, in some instances there seems to be evidence of spatial association. Indeed, notice for instance that although CA, TX, NY and FL do not Granger-cause Washington (WA), New Mexico (NM) and Minnesota (MN), CA Granger-causes WA, TX Granger-causes NM, and NY Granger-causes MN.

[3] All estimations were performed using the econometric software RATS version 8.10.

Next, we compute the associated GIRF that describes the time profile of the effect of a shock observed in the respective state, as well as that of shocks that originate in a different state (where shocks are measured by one standard deviation).[4] For each state the resulting lag weights are then normalised so that they add up to one, and the half-life of a shock is calculated as the number of months required for 50 per cent (or the first half) of the adjustment to take place.[5] Notice that there is no need for half-lives to be symmetric; that is, the response of unemployment in state i to a shock in unemployment in state j, hl_{ij}, is not necessarily the same as the response of unemployment in state j to a shock in unemployment in state i, denoted hl_{ji}. This can be motivated by the presence of leader and follower states. For example, one might expect New York or California to be a leader state.

Perhaps it is interesting to start the description of the results by comparing the average of the half-life estimates from GIRF for shocks observed in the respective market (own shock) with the average of the half-life estimates of the shocks that originate in a different market. Our results (available upon request) uncover an interesting spatial dimension on the speed of adjustment of unemployment rates to shocks, as the average half-life estimate of an own shock is approximately 50 months, while the average half-life estimate of the same shock on other markets takes much longer to dissipate, that is approximately 68 months. The difference between the latter and the former is positive and statistically different from zero. At this point, it is worth examining the effect (if any) that the quantitative easing (QE) policies that have been implemented by the US monetary authorities since November 2008 may have had on the half-lives of shocks. Bearing in mind that the high dimensionality of the VAR model under consideration prevents us to perform estimation after this date, we estimate the model up to October 2008, and compute the resulting half-lives of shocks. Our results indicate that in the absence of QE policies the average half-life estimate of an own shock is approximately 60 months, while for shocks on

[4]The dynamic responses of each endogenous variable in the VAR to a shock to the system are computed for 120 periods, that is, 10 years.

[5]Even though in few instances the resulting lag weights turn out to be negative, the cumulative lag weights are positive and thus the resulting half-lives can be interpreted in a meaningful way.

other markets it is about 71 months. In other words, in the absence of QE policies US unemployment rates would have to take much more to adjust to shocks (which is particularly noticeable in the case of own shocks). Of course, the persistence we have identified here should not be confused with hysteresis. The effects of the shocks are long-lasting, but not permanent.

We now consider a range of potential factors that might influence the speed of adjustment between state unemployment rates. We begin with spatial considerations in terms of whether states pairs that are closest in terms of distance are characterised by a faster speed of adjustment or shorter half-life in response to a shock to equilibrium. To examine in a more formal way the role of distance on the half-life of shocks, we estimate a regression of hl_{ij} on an intercept, and the distance between states i and j, d_{ij}, which we consider in logarithms and denote ld_{ij}. The data for this regressor corresponds to the Euclidian distance between the population centres of any two states, based on the geographic coordinates (latitude and longitude) obtained from the Census Bureau for the year 2000.[6] It ought to be noticed that in the case of the response of unemployment to own shocks, that is when $i = j$, we set $d_{ij} = 1$ so that $ld_{ij} = 0$.

The second driver of half-lives that we consider is the state job vacancy rate. The Beveridge curve represents a graphical inverse relationship between the unemployment and job vacancy rates. The matching process will determine how efficiently workers find new jobs. Improvements in the matching system would shift the curve towards the origin, because an efficient matching process will find jobs faster- filling vacancies and employing the unemployed. Depending on the degree of mismatch, we might expect higher vacancy rates to be associated with shorter half-lives. Another factor used to explain hl_{ij} is the average vacancy rate between states i and j denoted as $avac_{ij}$.[7] More vacancies are expected to facilitate a faster speed of ad-

[6] We thank Gary Wagner for kindly providing these data, which were used in Garrett et al. [18].

[7] The vacancy data were taken from (several issues of) the News Release of The Conference Board, as downloaded from *www.conference-board.org*. Restrictions on data availability limited access to every month over the full sample period. Therefore, state vacancy rates are computed as the mean values observed in December 2007 and May 2009. The choice of these two specific months

justment when it comes to considering unemployment responses within states, which implicitly assumes that state vacancies are the "right" vacancies insofar as these are jobs that the unemployed can fill readily. While a higher rate of unemployment normally occurs with a lower rate of vacancies, inefficient labour markets can be due to mismatches between available jobs and the unemployed and an immobile labour force.

The third driver of half-lives we consider is homeownership across US states. The Oswald hypothesis implies a positive coefficient on homeownership insofar as slowing down the speed of adjustment or increasing the half-life but, as indicated earlier, the empirical literature is not very supportive when it comes to looking at the unemployment rate. To examine the validity of this hypothesis, we compute the mean homeownership rate for each state over the period 1990 to 2010. Interestingly, we observe that the resulting state homeownership rates over the last two decades or so exhibit quite a noticeable degree of cross-sectional variation, with West Virginia and Michigan in the upper part of the spectrum (77.1 and 75.2, respectively), while California and New York are at the lower end (57.2 and 54.5, respectively). Next, with this cross-sectional information, we compute the average homeownership rate between states i and j, which is denoted $hown_{ij}$.[8]

Finally, consideration is also made for the potential effect of state-to-state migration flows (as a percentage of the labour force in the recipient state). For this variable, we use the 2010 American Community Survey produced by the US Census Bureau, which provides estimates of this year state-to-state migration flows for the one-year-old and older population. However, results not reported here (for brevity) indicate that this variable does not turn out to be statistically significant (at least

are the closest data match to the respective peak (December 2007) and trough (June 2009) of the latest US business cycle, as dated by the National Bureau of Economic Research (NBER).

[8]The source of these data is the US Department of Commerce: Census Bureau, as taken from *www.census.gov/housing/hvs/*.

at the 10% significance level), and is therefore omitted from the analysis.[9] [10]

OLS estimation of a regression model in which hl_{ij} is regressed against ld_{ij}, $avac_{ij}$ and $hown_{ij}$ yields the following results:

$$\begin{aligned} hl_{ij} \quad = \quad & \underset{(15.315)}{149.259} + \underset{(0.602)}{1.879}\ ld_{ij} - \underset{(1.580)}{15.313}\ avac_{ij} \\ & - \underset{(0.195)}{0.810}\ hown_{ij}, \qquad R^2 = 0.041,\ \text{Obs.} = 2304, \end{aligned} \tag{5}$$

where the numbers in parentheses denote White's heteroskedastic-consistent standard errors (the White F-test for heteroskedasticity of an unknown form, including cross terms, is clearly rejected at 8.296). As can be seen, the estimated coefficients on distance and vacancy rate have the expected positive and negative signs, respectively, and are statistically significant. These results support the view that shocks to unemployment rates take longer to disappear for regions farther apart, while adjustment is quicker the greater the average vacancy rate. On the other hand, the estimated coefficient on homeownership is negative and statistically significant. Whereas the prediction of the Oswald hypothesis is that higher homeownership is associated with a higher unemployment rate, we find that higher homeownership is associated with a shorter half-life. This finding points to better labour market outcomes for homeowners and is more consistent with studies such as Coulson and Fisher [10] and [11] and Munch et al. [30].[11] Moreover, the positive OLS coefficient offers a qualification to the Oswald hypothesis. While it is conceivable that high homeownership is accompanied by high unemployment, our results suggest that labour market adjustment in response to a regional shock may in fact be conducted relatively more quickly by homeowners over renters. Taking into account the abovementioned studies, our

[9] This lack of significance of this variable is perhaps explained by the fact that for some states there appears to be a great deal of uncertainty surrounding the estimates of the migration flows, as judged by the associated margins of error that are reported based on 90% confidence intervals.

[10] Rowthorn and Glyn [43] find little or no mean-reversion in state-level *employment* rates. They find that the speed of convergence is very slow. They conclude that US regional labour markets are not highly flexible insofar as their analysis suggests that migration and other forces behind adjustment to regional shocks are now quite weak in the US.

[11] As an alternative model specification, we also included an interaction variable based on distance and homeownership. However, this variable did not turn out to be statistically significant, and is therefore omitted from the analysis.

findings support the idea that high homeownership states have a better facility for job creation and overall production and that homeowners have better prospects of securing local employment.[12]

However, focusing on this last finding, and in what appears to be an interesting twist of results, Table 5 reports estimates of the regression coefficients at selected extreme quantiles, τ, of the conditional distribution of the half-life. Interestingly, while at τ = 0.05, 0.10, 0.15, 0.20 and 0.25 the estimated coefficient on $hown_{ij}$ is negative and significant, the quantile regression analysis also indicates a positive and statistically significant coefficient at the higher quantiles, that is when τ = 0.75, 0.80, 0.85, 0.90 and 0.95. In other words, state pairs characterised by an already sluggish unemployment rate adjustment appear to be those state pairs for which the Oswald effect is present. In contrast, moving along the conditional distribution of the half-life, the estimated coefficients on the other two covariates, that is ld_{ij} and $avac_{ij}$, do not change qualitatively.

5 Concluding remarks

In this paper, we first set out to assess whether the natural rate or the hysteresis hypothesis best describes the behaviour of US state unemployment rates. Our empirical analysis, which employs panel unit root and stationarity testing that account for cross sectional dependency, offers support for the natural rate hypothesis, i.e. unemployment rates across US states behaving as stationary processes. We also set out to identify the drivers of the half-lives or speeds of adjustment of unemployment rates towards their long-run equilibrium. This is in contrast to an existing literature more focussed on modelling the level of the unemployment rate. With no tendency to drift apart in the long-run, a multi-dimensional VAR analysis of the bivariate state pairs of unemployment rates suggests that the half-life of a shock is positively affected by distance and negatively affected by vacancy rates. Whereas

[12]While our prime interest here is the accuracy or significance of the covariate slope coefficients, the low R-squared is nonetheless significantly different from zero.

the Oswald hypothesis points to a positive association between homeownership and the unemployment rate, we obtain an initial result which suggests that the half-life of unemployment rate shocks is negatively associated with homeownership (based on a mean response). This finding is consistent with a literature that points towards homeowners having better labour market outcomes. However, further insight provided by an analysis based on quantile regression (where the symmetry assumption is relaxed) suggests that a positive relationship occurs at the higher quantiles. In other words, the Oswald hypothesis is more likely where adjustment of the unemployment rate is already relatively sluggish.

References

[1] K. Abadir, G. Caggiano and G. Talmain. Nelson-Plosser revisited: The ACF approach. *Journal of Econometrics*, 175:22–34, 2013.

[2] D. G. Blanchflower and A. J. Oswald. Does home-ownership impair the labor market?. NBER Working Paper Series 19079, 2013.

[3] S. G. Bronars and D. W. Jansen. The geographical distribution of unemployment rates in the US. *Journal of Econometrics*, 36:251–79, 1987.

[4] J. Carrión-i-Silvestre and A. Sansó. A guide to the computation of stationarity tests. *Empirical Economics*, 31:433–448, 2006.

[5] J. A. Chalmers and M. J. Greenwood. The regional labor market adjustment process: determinants of changes in rates of labor force participation, unemployment, and migration. *Annals of Regional Science*, 19:1–17, 1985.

[6] T. Chang, K. Lee, C. Nieh and C. Wei. An empirical note on testing hysteresis in unemployment for ten European countries: Panel SURADF approach. *Applied Economics Letters*, 12:881–886, 2005.

[7] K. M. Cheng, N. Durmaz, H. Kim, and M. L. Stern. Hysteresis vs. natural rate of US unemployment. *Economic Modelling*, 29:428–434, 2012.

[8] J. Clemente, L. Lanaspa, and A. Montañés. The unemployment structure of the US states. *Quarterly Review of Economics and Finance*, 45:848–868, 2005.

[9] T. G. Conley and G. Topa. Socio-economic distance and spatial patterns in unemployment. *Journal of Applied Econometrics*, 17:303–327, 2002.

[10] N. E. Coulson and L. M. Fisher. Labor market outcomes of homeownership. *Housing Studies*, 17:35–50, 2002.

[11] N. E. Coulson and L. M. Fisher. Housing tenure and labor market impacts: The search goes on. *Journal of Urban Economics*, 65:252–264, 2009.

[12] J. P. Cover and S. K. Mallick. Identifying sources of macroeconomic and exchange rate fluctuations in the UK. *Journal of International Money and Finance*, 31:1627–1648, 2012.

[13] D. N. DeJong, J. C. Nankervis, N. E. Savin, and C. H. Whiteman. The power problems of unit root tests in time series with autoregressive errors. *Journal of Econometrics*, 53:323–343, 1992.

[14] D. A. Dickey and W. A. Fuller. Distribution of the estimators for autoregressive time series with a unit root. *Journal of the American Statistical Association*, 74:427–431, 1979.

[15] G. Elliot, T .J . Rothenberg and J. H. Stock. Efficient tests for an autoregressive unit root. *Econometrica*, 64:813–836, 1996.

[16] B. T. Ewing and M. A. Thompson. VAR and generalized impulse response analysis of manufacturing unit labor cost. *Physica A*, 387:2575–2583, 2008.

[17] H. Farber. Unemployment in the Great Recession: Did the housing market crisis prevent the unemployed from moving to take jobs? *American Economic Review*, 102:520–525, 2012.

[18] T. A. Garrett, G. A. Wagner, and D. C. Wheelock. Regional disparities in the spatial correlation of state income growth, 1977-2002. *Annals of Regional Science*, 41:601–618, 2007.

[19] M. Giulietti, J. Otero, and J. Smith. Testing for unit roots in three-dimensional heterogeneous panels in the presence of cross-sectional dependence. *Economics Letters*, 101:188–192, 2008.

[20] P. Goldberg and F. Verboven. Market integration and convergence to the law of one price: Evidence from the european car market. *Journal of International Economics*, 65:49–73, 2005.

[21] K. Hadri. Testing for stationarity in heterogeneous panels. *Econometrics Journal*, 3:148–161, 2000.

[22] K. Hadri and Y. Rao. Panel stationarity test with structural break. *Oxford Bulletin of Economics and Statistics*, 70:245–269, 2008.

[23] N. Havet and A. Penot. Does homeownership harm labor market performances? a survey. Groupe d'Analyse et de Théorie Économique Lyon - St Étienne GATE Working Paper 1012, 2010.

[24] J. P. Laamanen. Home-ownership and the Labour Market: Evidence from Rental Housing Market Deregulation. University of Tampere, Working Paper 89.

[25] K. Im, M. H. Pesaran, and Y. Shin. Testing for unit roots in heterogeneous panels. *Journal of Econometrics*, 115:53–74, 2003.

[26] D. Kwiatkowski, P. C. B. Phillips, P. Schmidt, and Y. Shin. Testing the null hypothesis of stationarity against the alternative of a unit root. *Journal of Econometrics*, 54:159–178, 1992.

[27] M. A. León-Ledesma. Unemployment hysteresis in the US states and the EU: A panel approach. *Bulletin of Economic Research*, 54:95–103, 2002.

[28] A. Levin and C. F. Lin. Unit root tests in panel data: Asymptotic and finite-sample properties, 1992. Unpublished manuscript, University of California, San Diego, Discussion Paper No. 92-93.

[29] H. Lütkepohl. *New Introduction to Multiple Time Series Analysis*. Springer Verlag, Heidelberg, 2005.

[30] J. Munch, M. Rosholm, and M. Svarer. Are homeowners really more unemployed? *Economic Journal*, 116:991–1013, 2006.

[31] C. R. Nelson and C. I. Plosser. Trends and random walks in macroeconmic time series: Some evidence and implications. *Journal of Monetary Economics*, 10:139–62, 1982.

[32] E. Nissan and G. H. Carter. Convergence-divergence of U.S. state unemployment rates. *Journal of Economics and Finance*, 25:172–80, 2001.

[33] A. Nistor. Assessing the effectiveness of human capital investments on the regional unemployment rate in the United States: 1990 and 2000. *International Regional Science Review*, 32:65–91, 2009.

[34] A. Oswald. The housing market and Europe's unemployment: A nontechnical paper, 1999. University of Warwick, mimeo.

[35] J. Otero and J. Smith. Response surface estimates of the cross-sectionally augmented IPS tests for panel unit roots. *Computational Economics*, 41:1–9, 2013.

[36] M. D. Partridge and D. S. Rickman. Differences in state unemployment rates: the role of labor and product market structural shifts. *Southern Economic Journal*, 62:89–106, 1995.

[37] M. D. Partridge and D. S. Rickman. The dispersion of the US state unemployment rates: The role of the market and non-market. *Regional Studies*, 31:593–606, 1997.

[38] J. B. Payne, B. T. Ewing, and E. P. George. Time series dynamics of US state unemployment rates. *Applied Economics*, 31:1503–1510, 1999.

[39] P. C. B Phillips and P. Perron. Testing for a unit root in time series regression. *Biometrika*, 75:334–346, 1988.

[40] M. H. Pesaran. General diagnostic tests for cross section dependence in panels, 2004. CESifo Working Papers, No.1233.

[41] M. H. Pesaran. A simple panel unit root test in the presence of cross section dependence. *Journal of Applied Econometrics*, 22:265–312, 2007.

[42] M. H. Pesaran and Y. Shin. Generalized impulse response analysis in linear multivariate models. *Economics Letters*, 58:17–29, 1998.

[43] R. Rowthorn and A. J. Glyn. Convergence and stability in U.S. employment rates. *Contributions to Macroeconomics*, Vol. 6, Iss. 1:Art. 4, 2006.

[44] B. Seong, A. M. Morshed, and S. K. Ahn. Additional sources of bias in half-life estimation. *Computational Statistics and Data Analysis*, 51:2056–2064, 2006.

[45] C. A. Sims. Macroeconomics and reality. *Econometrica*, 48:1–48, 1980.

[46] F. M. Song and Y. Wu. Hysteresis in unemployment: Evidence from 48 US states. *Economic Inquiry*, 35:235–243, 1997.

[47] D. Sul, P. Phillips, and C. Choi. Prewhitening bias in HAC estimation. *Oxford Bulletin of Economics and Statistics*, 67:517–546, 2005.

[48] R. Vedder and L. Gallaway. Spatial variations in U.S. unemployment. *Journal of Labor Research*, 17:445–61, 1996.

Table 1: CD cross-section dependence test on unemployment rates

Lags	CD test	p-value
1	105.345	[0.000]
2	121.006	[0.000]
3	111.566	[0.000]
4	110.653	[0.000]
5	109.497	[0.000]
6	105.457	[0.000]

Note: The Pesaran [40] CD statistic is based on the cross-correlation of the residuals that result from estimating *pth*-order ADF type regressions (including a constant) for each of the individual unemployment rates that conform the panel. This statistic follows a standard normal distribution under the null hypothesis of cross-sectional independence.

Table 2: CIPS test for panel unit root on US states unemployment rates

Lags	CIPS test	p-value
1	-2.083	[0.079]
2	-2.508	[0.000]
3	-2.437	[0.000]
4	-2.216	[0.014]
5	-2.119	[0.053]
6	-2.116	[0.055]

Note: The test regression includes a constant as deterministic component. The number of lags of Δy_{it} and of $\Delta \bar{y}_t$ that are included in the test regression is given in the first column. The p-values of the test statistics are based on the response surfaces estimated in Otero and Smith [35].

Table 3: Hadri tests for mean stationarity of unemployment rates

Criteria to determine lag length	Hadri test	p-value
SIC	374.478	[0.686]
GTS	394.383	[0.545]

Note: To construct the individual KPSS statistics, the estimate of the long-run variance of the tests is computed employing the boundary condition rule advocated by Sul et al. [47]. SIC and GTS indicate that the optimal number of lags required to compute the test statistic is determined using the Schwarz information criterion and the General To Specific algorithm, respectively, based on $p_{max} = 12$ lags. The idea of the GTS algorithm is to estimate an autoregressive model with $p_{max} = p$ lags, and test the significance of the coefficient associated to the last lag. If it is significant using a significance level of, say, 10%, one selects $p_{max} = p$ lags. Otherwise, the order of the autoregression is reduced by one until the coefficient on the last included lag is found to be statistically different from zero. Bootstrap p-values are based on 10000 replications.

Table 4: Granger causality F-tests

	Excluded state unemployment rate:				
State	CA	FL	NY	TX	All four
AL	3.622⋆	0.292	0.472	1.922	1.333
AR	3.143⋆	1.952	3.268⋆	0.114	2.290⋆
AZ	0.187	0.666	2.584⋆	3.854⋆	1.752⋆
CO	4.991⋆	1.576	0.482	1.730	2.395⋆
CT	0.967	0.089	0.678	1.291	0.804
DE	2.076	1.461	4.191⋆	1.401	2.187⋆
GA	5.664⋆	1.145	1.740	1.247	3.150⋆
IA	0.142	0.251	1.485	0.396	0.596
ID	4.934⋆	2.427⋆	4.717⋆	1.004	3.057⋆
IL	3.929⋆	2.481⋆	0.536	0.668	1.830⋆
IN	2.196⋆	1.530	2.765⋆	0.358	1.737⋆
KS	0.839	0.197	2.806⋆	2.991⋆	1.866⋆
KY	0.770	0.877	6.329⋆	2.258⋆	3.088⋆
LA	0.235	0.412	0.086	0.479	0.320
MA	4.838⋆	1.890	2.305⋆	3.652⋆	2.718⋆
MD	2.181⋆	0.745	0.677	1.259	1.308
ME	2.996⋆	0.689	2.039	1.955	2.483⋆
MI	0.393	0.348	1.186	0.985	0.950
MN	0.621	1.252	2.805⋆	0.385	1.075
MO	0.812	1.162	1.572	1.181	1.186
MS	2.297⋆	0.234	0.498	0.772	1.022
MT	5.554⋆	0.032	1.108	1.091	1.868⋆
NC	0.176	0.500	1.655	0.467	0.831
ND	0.809	0.620	0.195	1.176	0.850
NE	0.449	4.071⋆	2.460⋆	0.498	1.876⋆
NH	2.867⋆	0.630	0.282	0.466	0.943
NJ	1.210	0.363	1.630	1.179	1.230
NM	0.609	1.755	0.691	2.682⋆	1.279
NV	1.162	2.497⋆	0.890	2.620⋆	1.630⋆
OH	6.682⋆	3.152⋆	0.413	0.917	2.387⋆
OK	2.904⋆	4.219⋆	1.595	1.968	2.138⋆
OR	1.520	0.311	0.629	1.625	1.213
PA	1.471	0.687	1.211	1.842	1.190
RI	0.963	1.557	1.577	2.007	1.320
SC	4.609⋆	3.100⋆	0.388	0.565	1.760⋆
SD	0.842	0.512	2.380⋆	1.954	1.804⋆
TN	3.198⋆	1.074	7.761⋆	0.134	3.594⋆
UT	1.150	2.971⋆	2.020	0.546	1.955⋆
VA	2.482⋆	3.276⋆	1.825	1.998	2.414⋆
VT	1.613	0.288	1.864	3.377⋆	2.179⋆
WA	2.514⋆	0.421	2.095	0.926	1.496
WI	3.759⋆	3.167⋆	4.626⋆	1.775	3.542⋆
WV	4.388⋆	0.439	2.662⋆	0.917	2.171⋆
WY	0.026	1.801	0.599	1.179	1.259

Note: ⋆ indicates that the null hypothesis of no Granger-causality is rejected at the 10% level.

Table 5: Quantile process estimates

Quantile	Intercept		ld_{ij}		$avac_{ij}$		$hown_{ij}$	
	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.	Coeff.	S.E.
0.05	46.755	(31.216)	0.947	(0.322)	-2.620	(2.661)	-0.527	(0.363)
0.10	165.345	(57.801)	2.083	(0.519)	-14.280	(4.359)	-1.765	(0.713)
0.15	241.821	(35.041)	2.696	(0.605)	-21.345	(4.531)	-2.521	(0.430)
0.20	232.866	(32.350)	2.713	(0.694)	-22.171	(4.439)	-2.249	(0.422)
0.25	253.390	(24.071)	2.585	(0.699)	-22.387	(3.651)	-2.418	(0.304)
0.75	98.063	(14.146)	1.045	(0.564)	-13.549	(1.613)	0.284	(0.170)
0.80	92.934	(13.571)	1.611	(0.518)	-11.123	(1.519)	0.281	(0.168)
0.85	89.112	(14.039)	1.770	(0.458)	-10.600	(1.448)	0.366	(0.182)
0.90	78.967	(13.701)	1.908	(0.457)	-9.577	(1.391)	0.524	(0.178)
0.95	98.298	(19.827)	1.075	(0.856)	-7.895	(1.438)	0.333	(0.236)

Note: Standard errors calculated using a Huber Sandwich method.

Figure 1. Range of variation of the unemployment rate in 48 US states (1976m1 - 2012m7)

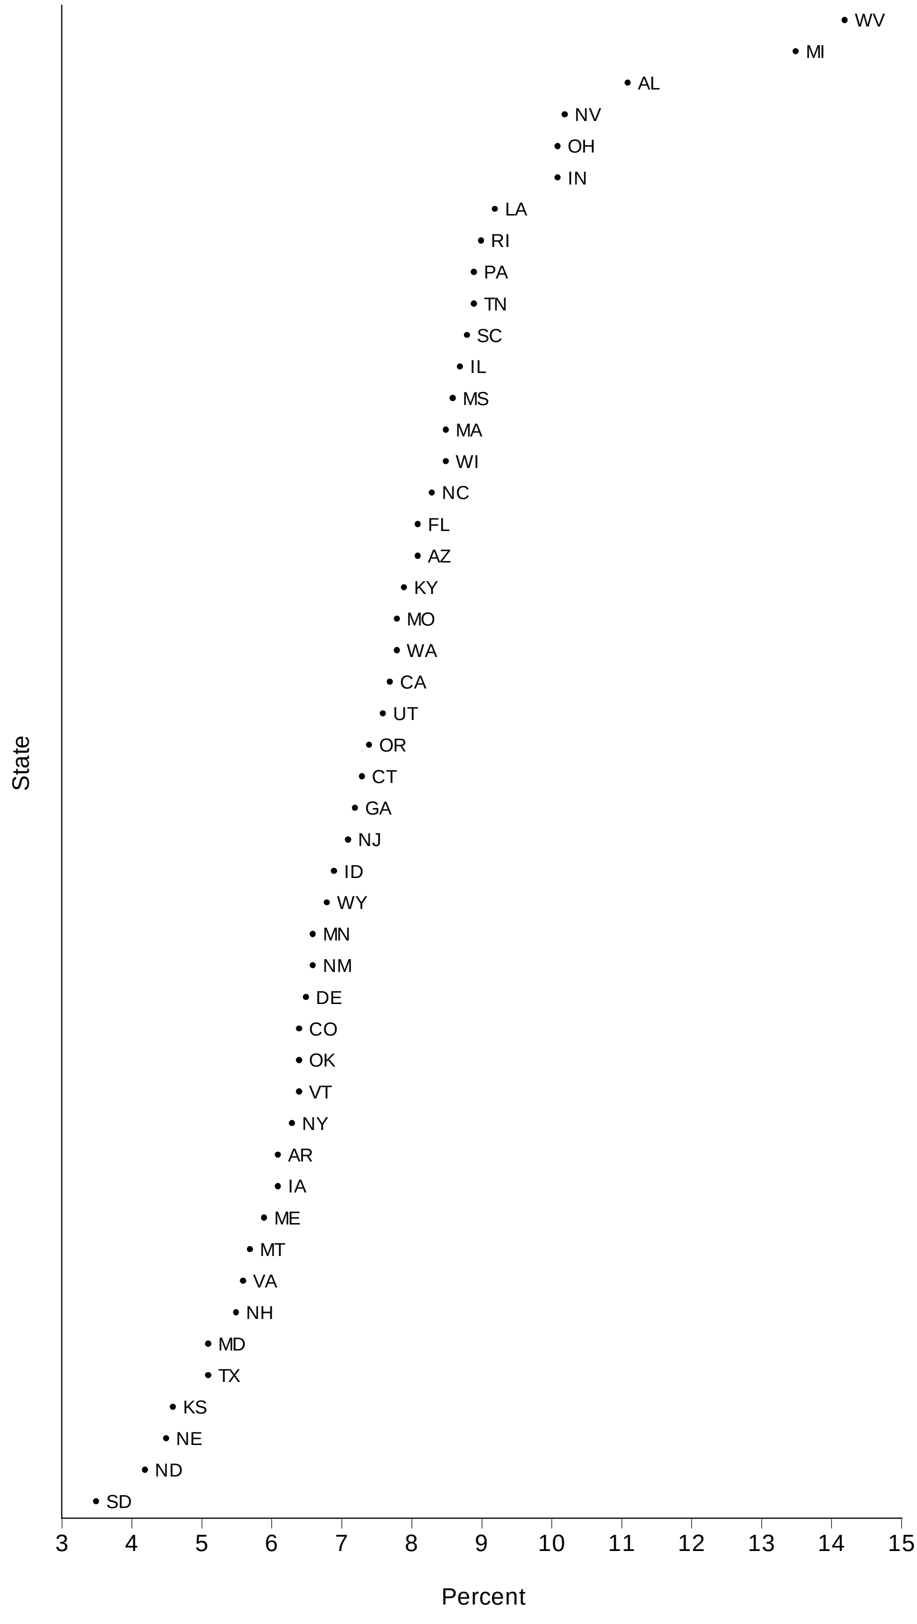

Note: The source of the data is the US Department of Labor (Bureau of Labor Statistics), as downloaded from the Federal Reserve Economic Data (FRED) of the Federal Reserve Bank of St. Louis.

Figure 2. Cross-sectional range of variation in 48 US states through time

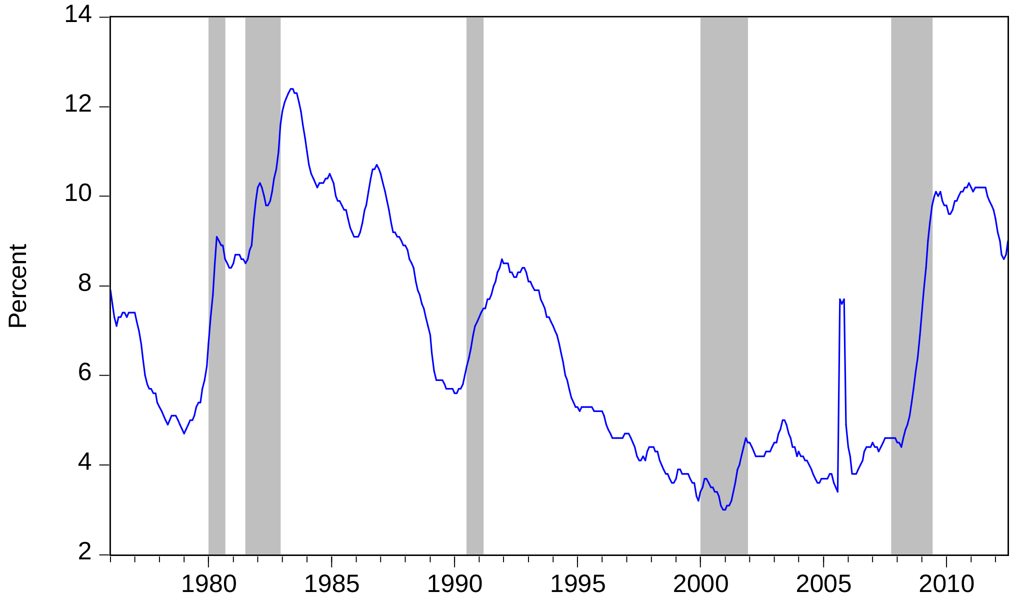

Note: The shadowed areas in the figure indicate the National Bureau of Economic Research peak to trough periods.

Figure 3. Mean reverting level of state unemployment and intermediate target