A local vector autoregressive framework and its applications to multivariate time series monitoring and forecasting

Ying Chen, Bo Li* and Linlin Niu

Our proposed local vector autoregressive (LVAR) model has time-varying parameters that allow it to be safely used in both stationary and non-stationary situations. The estimation is conducted over an interval of local homogeneity where the parameters are approximately constant. The local interval is identified in a sequential testing procedure. Numerical analysis and real data applications are conducted to illustrate the monitoring function and forecast performance of the proposed model.

Keywords and phrases: Adaptive estimation, Multivariate time series, Non-stationarity, Yield curve.

1. INTRODUCTION

Non-stationarity or near non-stationarity is a stylized fact of many macroeconomic and financial time series such as annual GDP growth, inflation rate, interest rates and exchange rates, etc. Even though there is no permanent trend in these series, a high degree of persistence is typical as often detected visually or measured by a slowly decaying autocorrelation function. This feature poses many challenges, not only for theoretical modeling, making inferences and conducting tests, but also for real time economic monitoring and forecasting. Many structural macro or financial models of multiple time series are also defeated by the persistence feature. Although the structural models prescribe very few lags of state variables, empirically, the vectors of a time series require more lags in order to fit the data. To address the non-stationarity issue, the modeling approaches can be broadly classified in two: the long memory approach versus the short memory approach. The long memory view utilizes the data generating process as described by models with constant parameters and innovations with slowly or non-decaying effects, such as the fractionally integrated processes in Granger (1980); Granger and Joyeux (1980) and Hosking (1981), etc. The short memory view considers persistence as spuriously generated by changes in basic modeling parameters, such as heteroscedasticity, structural breaks or regime switching, which are not directly observable and might be inappropriately modeled, see discussions

*Corresponding author.

in Diebold and Inoue (2001) and Granger and Hyung (2004). Technically, both the long memory view and the short memory view have merits in explaining persistence in the data. However, the short memory view often has economic underpinnings to support the various changes corresponding to policy shifts, regime transition, or varying features of exogenous shocks, etc. Conditional on the types of changes, the underlying model can be simple and intuitive. Also, appropriate modeling of the sources of changes per se can be helpful in understanding and monitoring the evolution of the fundamental economic or financial processes.

Recently, an adaptive approach with simple underlying models has emerged from the short memory class, which can incorporate the various sources of changes without explicitly assuming change types and timing. Specifically, the state variables are described by some simple dynamic process such as the autoregressive (AR) process or AR process with exogenous variables (ARX), and the parameters of the processes are time-dependent without explicit functional form. Under the conditions that most of the parameters are slowly evolving with small variations, and that large variations or breaks happen infrequently, the time series can be divided into local intervals such that over each local interval the data generating process can be well approximated by a local AR (LAR) or a local ARX (LARX) model with constant parameters. This approach is particularly useful in real-time monitoring of changes and forecasting under the assumption that the local model holds for the forecast horizon. Chen, Härdle and Pigorsch (2010) propose the LAR detection and estimation methodology and successfully apply it to forecasting realized volatility in a financial time series. Chen and Niu (2013) extend the method to an LARX model and apply it to yield curve modeling and forecasting, where under a traditional dynamic Nelson-Siegel factor model Diebold and Li (2006) each state yield factor is modeled as an LARX with inflation as the exogenous variable. Chen and Niu (2013) demonstrate that the data-driven LARX model brings clear advantages in forecasting the whole yield curve, when compared to a wide spectrum of traditional yield curve models with rolling or recursive estimation strategies using predetermined estimation windows. Moreover, the detected local intervals and the resultant parameter evolutions estimated in real-time are of great help for monitoring the yield curve

dynamics, understanding its changing relationship with the inflation factor, and providing evidence on policy changes and on changing interest rate behavior during the recent financial crisis. This paper generalizes the univariate process of LAR or LARX into a local vector autoregressive (LVAR) framework. The general framework can be used in multivariate structural models for effective modeling and real-time applications in macroeconomics and finance.

Within the VAR modeling framework, our LVAR modeling shares some similar background with the threshold vector autoregressive (TVAR) model and the time-varying VAR model, which feature coefficient changes conditional on states and time, respectively. The threshold autoregressive type model is more suitable for nonlinear time series; see a thorough discussion by Tong (1983) and the development and applications by Tong (1987); Tong and Lim (1980) and Tsay (1989), among others. Tsay (1998) proposes a TVAR model to generalize the model into multivariate settings. The model has been widely applied to business cycle effects, policy regime and transmission mechanism in developed economies, such as Atanasova (2003); Balke (2000); Li and St-Amant (2010), and Afonso, Baxa, and Slavik (2011), etc. The time-varying VAR models are more flexible for studying the changing behaviors of economic systems, which are widely applied to model changes in macroeconomic policies, regimes and exogenuous shocks, such as Cogley and Sargent (2001) and Cogley and Sargent (2005) on a small scale macroeconomic VAR with time-varying coefficients and stochastic volatility, respectively. Besides, Primiceri (2005) studies the causes of the Great Moderation with both time-varying coefficients and stochastic volatility. However, the flexibility of TVAR and of time-varying VAR is achieved at the cost of substantially increasing the dimension of parameter space. Furthermore, specific assumptions on the types and process of the parameter changes are required. Also, the estimation is technically demanding and time consuming. In contrast, the LVAR proposes much simpler underlying models, without specific assumptions on change types and processes, and the estimation is relatively simple.

The rest of the paper is arranged as follows. In Section 2, we discuss the method with modeling assumptions, the local interval detection technique using an empirical testing procedure, and the applications of monitoring and forecasting the vector of variables in a multivariate setting. In Section 3, we demonstrate the properties of the testing procedure and the forecast performance using a Monte Carlo study. Section 4 provides the real data analysis, where we illustrate the use of LVAR with applications to modeling and forecasting US yield curves. Finally, Section 5 concludes.

2. METHOD

In this section, we propose a local vector autoregressive (LVAR) model to estimate the joint dynamics of a vector of variables. The LVAR model allows parameters to be time dependent, without any particular assumptions on the time variation. Time-varying parameters at each point in time are, of course, too flexible to constitute an identifiable dynamic model. We therefore employ a local homogeneity assumption to balance between model flexibility and estimation feasibility. Local homogeneity assumes that at any particular time point there exists a past time interval, over which the local sample can be well approximated by a VAR model with constant parameters. A sequential testing procedure is used to find the longest interval that satisfies the local homogeneity assumption. The identified interval that defines the locally homogeneous sample is called the *interval of local homogeneity*.

Instead of using all the available past information, as in the recursive estimation approach, we conduct the estimation utilizing the interval of local homogeneity in the adaptive approach; the interval satisfying local homogeneity is time dependent and of possibly varying interval length. It is also different from the conventional rolling window or recursive window technique that adopts a fixed window size or expands the window size throughout the estimation. The fitted adaptive model is then used to monitor the model parameters and stable intervals, and to forecast the variables if the model is sufficiently parsimonious and the average intervals are sufficiently long to include efficient information for accurate prediction.

2.1 Adaptive vector autoregressive model

Let $\boldsymbol{X}_t \in \mathbb{R}^d$ denote d-dimensional autoregressive time series variables, with $t = 1, \ldots, T$. The adaptive vector autoregressive model is defined with time-varying parameters as:

$$\boldsymbol{X}_t = c_t + A_{1t}\boldsymbol{X}_{t-1} + \cdots + A_{p_t}\boldsymbol{X}_{t-p_t} + \epsilon_t, \quad \epsilon_t \sim N(0, \Sigma_t),$$

where $c_t = (c_{1t}, \ldots, c_{dt})^\top$ is the intercept vector at time point t and A_{jt} is a $d \times d$ matrix for $j = 1, \ldots, p$. The stochastic innovation ϵ_t is assumed to be Gaussian distributed satisfying $\mathsf{E}(\epsilon_t) = 0$ and $\mathsf{E}(\epsilon_t \epsilon_t^\top) = \Sigma_t$. Moreover, we assume that there is no serial correlation between any two innovations across time, that is $\mathsf{E}(\epsilon_t \epsilon_{t-k}^\top) = 0$ for $k \neq 0$. The LVAR model, with time-varying parameters, is appropriate in a non-stationary situation where structural changes exist. It also works well in the homogeneous case by fixing the parameters as constant. In either case, the estimation is conducted under the local homogeneity assumption. That is, the multivariate time series are approximated by a parametric model over an interval of local homogeneity.

Although the order of the LVAR model, denoted by p, is allowed to be more than 1, we fix the lag order to 1 for ease of elaboration. In addition, the adoption of the simplest model structure is significantly motivated by tractability of monitoring and good out-of-sample forecast performances. In Section 3, we investigate the model misspecification issue in a simulation study, where the true data generating

process has a higher order than the recommended one. It shows that the LVAR model of order 1 provides stable performance.

The LVAR model of order 1, represented in matrix form, is as follows:

$$\boldsymbol{X}_t = c_t + A_t \boldsymbol{X}_{t-1} + \epsilon_t.$$

For notational simplicity, we denote the unknown parameters by $\Theta_t = (c_t, A_t, \Sigma_t)$.

2.2 Estimation under local homogeneity

Suppose that at time point t, the time series is homogeneous with $\Theta_t = \Theta$ over an interval $I_t = [t - m_t + 1, t]$. Parameter m_t is the interval length, corresponding to the number of observations in the local sample. The local (quasi-) log-likelihood function is defined as:

$$\ell(I_t, \Theta) = -\frac{m_t}{2} \ln|2\pi\Sigma_t| - \frac{1}{2} \sum_{s=t-m_t+1}^{t} \epsilon_s^\top \Sigma_t^{-1} \epsilon_s$$

from which we obtain the local maximum likelihood estimate (MLE):

$$\widetilde{\Theta}_t = \arg\max \ell(I_t, \Theta).$$

Now we relax the local homogeneity assumption such that $\Theta_t \approx \Theta$. Then the modeling bias of the parametric model with constant parameter Θ and the local parametric model with time dependent parameter Θ_t can be measured by:

$$\Delta_t = |\ell(I_t, \Theta) - \ell(I_t, \Theta_t)|^{1/2}$$

which should be small. Therefore, the local MLE, though not unbiased in this situation, can be used.

In practice, the interval of local homogeneity is unknown. The question is how to identify it or equivalently how to select interval length m_t at any particular time point t. With too large a value, there is a high probability of having nontrivial modeling bias, which violates the local homogeneity assumption. On the contrary, a small value of m_t though satisfying a small modeling bias, unnecessarily discards too many observations that are useful for estimation. The goal is to select the longest interval that does not violate the local homogeneity assumption.

Suppose there are K candidate intervals at time point t, which contain the interval of local homogeneity I_t:

$$I_t^{(1)} = \big[t - m_t^{(1)} + 1, t\big], \quad \ldots, \quad I_t^{(K)} = \big[t - m_t^{(K)} + 1, t\big]$$

with $I_t^{(1)} \subset \cdots \subset I_t^{(K)}$. A sequential testing procedure helps to select the longest interval that satisfies local homogeneity. It is worth noting that beyond the selected interval, there is a high probability of structural changes. The adaptive technique proceeds as follows. The procedure starts from the shortest interval $I_t^{(1)}$, over which the local homogeneity assumption probably holds. We accept the interval and denote

the accepted estimator by $\hat{\Theta}_t^{(1)} = \widetilde{\Theta}_t^{(1)}$. Iteratively, the procedure extends to the next interval $I_t^{(k)}$ for $k \geq 2$, where we introduce the test statistic:

$$T_t^{(k)} = \big|\ell\big(I_t^{(k)}, \widetilde{\Theta}_t^{(k)}\big) - \ell\big(I_t^{(k)}, \hat{\Theta}_t^{(k-1)}\big)\big|^{1/2},$$

which measures the divergence of the hypothetical model from the recently accepted local model. If the divergence is significant, where the significance level is controlled by a critical value ζ_k, it indicates that there is a significant structure change larger than the one arising due to sampling changes. In this case, we reject the null hypothesis of local homogeneity and terminate the selection procedure. The last accepted interval would be the final selection and we have $\hat{\Theta}_t^{(j)} = \hat{\Theta}_t^{(k-1)}$ for $j = k, \ldots, K$. Here j denotes the index of the candidate interval where the procedure is terminated. Otherwise, we accept the longer interval $I_t^{(k)}$, and update estimate $\hat{\Theta}_t^{(k)} = \widetilde{\Theta}_t^{(k)}$. The test procedure is continued on the next interval until either a change is detected or the longest candidate interval is reached.

2.3 Calibrate critical values

The success of the adaptive selection procedure depends on the critical values, which are calibrated in Monte Carlo experiments. As the critical values control the significance level under the local homogeneity assumption that requests a small modeling bias, we generate samples with homogeneity and measure the modeling bias using an adaptive estimation. The critical values are selected such that they are capable of providing the prescribed performance of the testing procedure.

The homogeneous VAR processes are generated with constant parameters $\Theta^* = (c^*, A^*, \Sigma^*)$, such that

$$\boldsymbol{X}_t = c^* + A^* \boldsymbol{X}_{t-1} + \epsilon_t, \quad \epsilon_t \sim N\big(0, \Sigma^*\big).$$

Each process includes T observations and the generation is repeated N times. For each generated process $\boldsymbol{X}_{1:T}^n$, $n = 1, \ldots, N$, the same interval set $I_t^{(k)}$ is used everywhere for $t = m^K, \ldots, T$ and $k = 1, \ldots, K$, where m^K is the longest interval length for interval $I_t^{(K)}$. In the following, for ease of elaboration, we drop the series index n. Under the assumption of homogeneity, the estimation error can be measured by the fitted log-likelihood ratio over each interval:

$$(1) \qquad R_k = E_{\Theta^*}\big|\ell\big(I^{(k)}, \tilde{\Theta}_t^{(k)}\big) - \ell\big(I^{(k)}, \Theta^*\big)\big|^{1/2},$$

where R_k can be computed numerically with knowledge of Θ^*.

Once a set of critical values $\zeta_1, \ldots, \zeta_K$ is given, one can employ the adaptive procedure, by checking the significance of the test statistic $T_t^{(k)}$, to obtain the adaptive estimate $\hat{\Theta}_t^{(k)}$ of the time-dependent parameter Θ_t. Given the MLE $\tilde{\Theta}_t^{(k)}$ of the constant parameter Θ^*, the temporal realized

modeling bias can be measured as:

$$\delta_t^{(k)} = \left| \ell\big(I^{(k)}, \tilde{\Theta}_t^{(k)}\big) - \ell\big(I^{(k)}, \hat{\Theta}_t^{(k)}\big) \right|^{1/2}.$$

The adaptive estimation should behave as well as the true underlying characteristics under the null of time homogeneity, in the sense that the modeling bias is bounded by estimation error R_k, with knowledge of the true model, in Equation (1):

(2)
$$E_{\Theta^*}\big(\delta_t^{(k)}\big) = E_{\Theta^*}\left| \ell\big(I^{(k)}, \tilde{\Theta}_t^{(k)}\big) - \ell\big(I^{(k)}, \hat{\Theta}_t^{(k)}\big) \right|^{1/2} \leq R_k.$$

Clearly, the critical values are the only unknown parameters in the above inequality Equation (2), which can be calibrated.

The computation of critical values relies on two hyperparameters, the interval candidates $(I_t^{(1)}, \ldots, I_t^{(K)})$ and Θ^*. In our study, at each point of time, we consider $K = 19$ intervals for the adaptive estimation, starting with 12 months and a continuous increment of $M = 6$ months between any adjacent intervals, i.e., 120 months (10 years) is the maximal sample size. Ideally, Θ^* should be close to the true parameter underlying the real data series at each point of time, which is actually the target of our estimation. In practice, we approximate Θ^* with the estimate from a sub-sample, for example, the sub-sample before the forecast exercise starts. We find that the adaptive technique is quite robust to the selection of the hyperparameters, as is illustrated in Section 3. There is no significant difference in terms of forecast accuracy for different sets of interval candidates determined by K and M as well as for possible misspecifications of Θ^* with $\pm 20\%$ deviation from the true values.

In the following numerical analysis, we use the MLE of the available real sample in the 1983:1–1997:12 period as Θ^*, i.e., using information before the forecast exercise starts. We then generate a homogeneous VAR series, and calibrate the set of critical values as described above. The same set of calibrated critical values is adopted for every time point throughout the real-time estimation and forecast.

The adaptive estimation algorithm for any particular time point t is as follows:

1. Calibrate critical values given a set of interval candidates $\{I_t^{(k)}\}_{k=1}^K$.
   - Generate homogeneous VAR processes with constant parameter Θ^*. We use the MLE for the time point before the first forecast origin.
   - Compute MLEs $\tilde{\Theta}_t^{(k)}$ and the risk bound R_k over each interval candidate.
   - Given an initial set of critical values, obtain the adaptive estimator $\hat{\Theta}_t^{(k)}$. Compute the realized modeling bias $\delta_t^{(k)}$ and check the risk bound, (2). If it holds, reduce the critical values. Otherwise, increase the critical values. Repeat until the cutting point is found.

2. Given the calibrated critical values, for the data set of interest to be investigated, conduct the sequential testing procedure to identify the interval of local homogeneity and estimate the adaptive estimator over time. Starting from an initial time t_0, for $t \geq t_0$:
   - **Initialization:** We accept the shortest interval and set $\hat{\Theta}_t^{(1)} = \tilde{\Theta}_t^{(1)}$.
   - **Loop:** For $k \geq 2$,
     if $T_t^{(k)} = |\ell(I_t^{(k)}, \tilde{\Theta}_t^{(k)}) - \ell(I_t^{(k)}, \hat{\Theta}_t^{(k-1)})|^{1/2} \leq \zeta_k$, we accept the interval $I_t^{(k)}$ and update the estimate:

     $$\hat{\Theta}_t^{(k)} = \tilde{\Theta}_t^{(k)}.$$

     Otherwise, terminate the procedure and we have:

     $$\hat{\Theta}_t^{(j)} = \hat{\Theta}_t^{(k-1)}, \quad j = k, \ldots, K.$$

   - **Final estimate:** $\hat{\Theta}_t = \hat{\Theta}_t^{(K)}$.

3. We assume that the interval of local homogeneity is extendable over the forecasting horizon, denoted by h. The fitted LVAR model is used for the prediction: $\hat{\boldsymbol{X}}_{t+h} = \hat{c}_t + \hat{A}_t \boldsymbol{X}_t$.

3. SIMULATION

In this section, we conduct a simulation study to demonstrate the performance of the LVAR model. In particular, we evaluate the forecast accuracy of the adaptive procedure, compared to alternative methods with window length selection such as the rolling window technique. Furthermore, the robustness of the forecast performance for the adaptive procedure is investigated with respect to the choice of hyperparameters (Θ^*, K, M). Moreover, we address the model misspecification issue of whether the simplest adaptive model of order 1 is sufficient to handle autoregressive processes of a higher order.

3.1 Simulation design

We consider two kinds of scenarios: a homogeneous scenario with globally constant parameters, denoted as HOM, and a heterogeneous scenario with time-varying parameters shifting from one level to another, i.e., a regime switching scenario, denoted as RS.

In the HOM scenario, we calibrate the VAR coefficients from a three factor VAR(1) model constructed by Nelson-Siegel yield factors with US yield curve data from 1983 to 1997. The Nelson-Siegel (NS) model is parameterized according to Diebold and Li (2006), and the data set is a fifteen yield series as used in Chen and Niu (2013). We denote the underlying parameters as $\Theta_0 = (c_0, A_0, \Sigma_0)$ and keep them constant throughout the whole sample. In the RS scenario, we design two experiments and label them as: RS-A, where A denotes the vector autoregressive coefficient matrix, and RS-C, where C denotes the intercept vector.

Table 1. Parameters in the simulation scenarios. HOM refers to the homogeneous scenario; RS refers to the regime-switching (structural change) scenario. In each of the RS scenarios, only the labeled parameter is changed in Phase 2. The other parameters remain the same as in the original set-up

Scenarios	Parameters		
HOM	$c_0 = \begin{pmatrix} 0.093 \\ 0.111 \\ -0.314 \end{pmatrix}$		$A_0 = \begin{pmatrix} 0.989 & 0.011 & -0.005 \\ -0.031 & 0.933 & 0.054 \\ 0.062 & 0.090 & 0.853 \end{pmatrix}$
RS	Phase 1 $t \in [1, 200]$		Phase 2 $t \in [201, 400]$
RS-A	$A_t = \begin{pmatrix} 0.989 & 0.011 & -0.005 \\ -0.031 & 0.933 & 0.054 \\ 0.062 & 0.090 & 0.853 \end{pmatrix}$		$A_t = \begin{pmatrix} 0.493 & -0.167 & 0.177 \\ 0.259 & 0.952 & -0.082 \\ 0.523 & 0.511 & 0.462 \end{pmatrix}$
RS-C	$c_t = (0.093, 0.111, -0.314)^\top$		$c_t = (2.789, -1.974, -3.503)^\top$

In the RS scenarios, only the labeled parameters shift from the original level to a new set of parameter values estimated using a different subsample of the NS factors, during the recent financial crisis from 2008 to 2010. The other parameters remain the same as in the original set up.

For each scenario and experiment we simulate 200 data series, each with 400 observations. In the RS scenarios, each of the regimes lasts for 200 time points. The details of the scenario designs regarding timing and parameters are described in Table 1.

We employ both the adaptive procedure and the rolling window strategies to compute the one-step ahead forecasts for the same forecasting period from $t = 122$ to 400 for each of the simulated samples. In the adaptive case, the critical values are calibrated using the true underlying parameters, i.e., $\Theta^* = \Theta_0$. The candidate intervals start from 12 months (1 year) and end at 120 months (10 years), with $K = 19$ and $M = 6$. The interval lengths are 12, 18, 24, 30, 36, 42, 48, 54, 60, 66, 72, 78, 84, 90, 96, 102, 108, 114, 120.

Figure 1 displays the resulting critical values that are used in each scenario. At the same time, we consider 19 alternative window sizes in the prediction using the rolling window technique, i.e., ranging from I^1 (12 months) to I^K (10 years), which correspond to our interval candidates in the adaptive procedure. Forecast accuracy is determined by the forecast root mean squared error (RMSE) values.

3.2 Forecast accuracy

Table 2 presents the RMSE values of the adaptive technique and the alternatives. For ease of exposition we do not report all forecasting results of the rolling windows. Instead, we only list those window sizes yielding the best forecast accuracy (with minimal RMSE values) and the worst accuracy (with maximum RMSE values). The respective rolling window sizes are indicated in the parentheses. The number of times that the LVAR is superior to the 19 alternative window choices is highlighted in the column of "No. of Winning".

The numerical results reveal that the adaptive approach with varying local window sizes introduces more flexibility
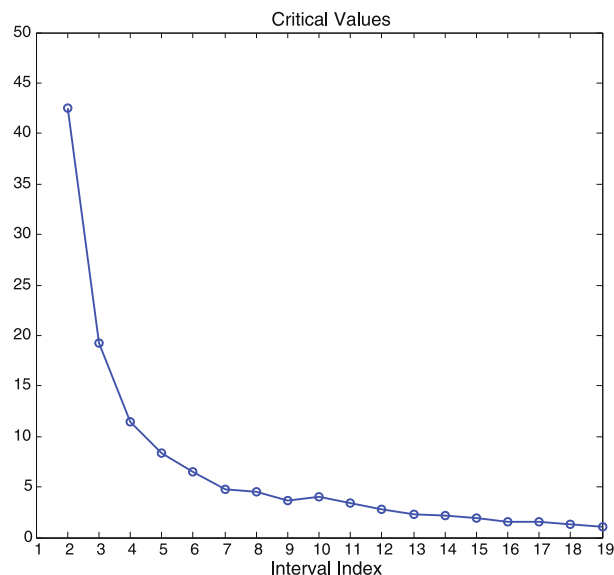

Figure 1. Critical values. The hyperparameters are $M = 6$, $K = 19$ and $\Theta^* = \Theta_0$.

into the procedure, leading to a comparable performance to the optimal sample under the homogeneous scenario and a generally better performance under the scenarios with structural changes. More specifically, in the homogeneous scenario, the adaptive technique, though with a misspecified assumption of time-varying coefficients, still provides reasonable accuracy. In the structural change scenarios with time-varying parameters, our technique is superior to all 19 alternative rolling window estimations.

The adaptive interval selection procedure contributes to the improvement of forecast accuracy and simultaneously provides stable performances. In the structural change scenarios with a parameter shift at $t = 201$, the average values of the selected intervals drop quickly after that point, see Figure 2. As the sample following the new data generating process extends, the lengths of the selected intervals

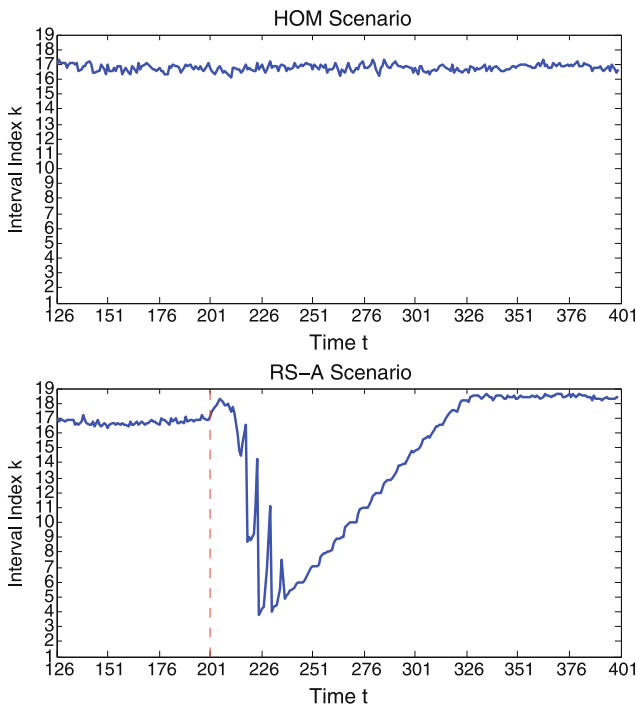| | Scenario | Rolling window (window size) | | Adaptive | No. of Winning |
		Best	Worst		
NS1	HOM	0.336 (120)	0.424 (12)	0.340	13/19
	RS-A	0.400 (78)	0.517 (12)	0.395	19/19
	RS-C	0.501 (54)	0.606 (12)	0.489	19/19
NS2	HOM	0.370 (120)	0.474 (12)	0.374	14/19
	RS-A	0.399 (120)	0.514 (12)	0.398	19/19
	RS-C	0.472 (72)	0.588 (12)	0.465	19/19
NS3	HOM	0.816 (120)	1.049 (12)	0.826	13/19
	RS-A	0.851 (120)	1.084 (12)	0.855	16/19
	RS-C	0.953 (96)	1.195 (12)	0.954	16/19

Figure 2. The average value of the selected intervals from time index 122 to 400 over the 200 generated processes in the HOM and RS-A scenarios.

increase. The conventional rolling window technique with a fixed window size, on the other hand, does not have such flexibility. Moreover, in the homogeneous scenario where there is no structural change, the optimal interval selection should be the longest one of 120 months. The average values of the selected intervals are quite reasonable, with

values around 108 months ($k = 17$) for each time point. A direct comparison confirms that the simple yet flexible LVAR model can be safely applied to both stationary and non-stationary situations.

3.3 Robustness check

The forecast accuracy of the adaptive model depends on the critical values which themselves depend on the underlying hyperparameters K, M and Θ^*. As an illustration, we analyze the robustness of the hyperparameter choices under the RS-A scenario. The default values are $K = 19$, $M = 6$ and $\Theta^* = \Theta_0$. We report the forecast accuracy with different hyperparameters and also analyze the impact on the forecast performance when Θ^* is misspecified.

We first consider four alternative interval sets: given $M = 6$, taking fewer or more candidates with $K = 10$ or $K = 30$; given $K = 19$ with the first interval being 12 months, taking shorter or longer steps with $M = 3$ or $M = 12$ between two adjacent intervals. With these alternative interval sets, the critical values are re-calibrated. In order to match the longest possible interval length for $I^{(K)}$, the initial forecasting points for cases with $K = 10$ and $K = 30$ are different, which are $t = 122$ and $t = 188$, respectively.

Moreover, we consider the cases where parameter Θ^* is misspecified. Instead of using the true underlying parameter values, i.e., $\Theta^* = \Theta_0$, we compute critical values under each of the two misspecified hypothetical parameters, i.e., $\pm 20\%$ deviation on the VAR coefficients. More specifically, we decompose matrix A and shift its eigenvalues by 20% to have $0.8 \times \text{EV}$ and $1.2 \times \text{EV}$, and denote the scenarios as mis08 and mis12, respectively. We use these parameter sets to generate Monte Carlo experiments and to calibrate the critical values, although the series actually follow the VAR model with Θ_0. A potential problem occurs in scenario mis12, where stationarity is not valid even in the shortest interval. To guarantee the existence of local homogeneity, we artificially select a matrix that satisfies the stationarity condition. The forecast accuracy of scenario RS-A with the new parameters is computed and compared again with the alternative rolling window forecasts.

Table 3 presents the forecast accuracy under the alternative or misspecified hyperparameters. The results confirm the robustness of the adaptive technique, with RMSE values very close to those in the default case where true parameter values are used to calibrate the critical values.

3.4 Model misspecification

In the following experiment, we investigate the stability of the proposed adaptive model in terms of model misspecification. It is necessary to answer whether the adaptive model of order 1 provides reasonable forecast accuracy, if the true data generating process has a higher lag order. As an illustration, we consider the true data generating process to be a LVAR process with lag order 5. Among other scenarios, we conduct the simulation under the RS-A scenario where

Table 3. Robustness testing (scenario RS-A): RMSE values. We compare the default case of $M = 6$, $K = 19$ and $\Theta^* = \Theta_0$ to several cases of alternative hyperparameters of $M = 3$ or 12, $K = 10$ or 30 and misspecified parameter Θ^* in the critical value calibration

	default	$K = 10$	$K = 30^{*1}$	$M = 3$	$M = 12$	default^{*2}	mis08^{*2}	mis12^{*2}
NS1	0.395	0.395	0.409	0.396	0.397	0.348	0.349	0.347
NS2	0.398	0.402	0.403	0.400	0.399	0.378	0.381	0.376
NS3	0.855	0.862	0.862	0.858	0.857	0.844	0.849	0.841

[*1] The first forecast is at time index 188 (instead of 122 as for others) relative to the longest possible interval length

[*2] An artificial VAR coefficient matrix is used to guarantee the existence of local homogeneity after being multiplied by 120% in the mis12 scenario

Table 4. Model misspecification with the true data generating process of LVAR(5) and time-varying VAR parameters. In the table, $p = 1$ and $p = 5$ refer to the misspecified and correct lag orders in the adaptive estimation, respectively. Only the best and worst results of all the rolling window approaches (with the corresponding window sizes) are reported. The last two columns contain the LVAR results and the number of cases where LVAR is better than the rolling window approaches in terms of RMSE values

		Rolling window (window size)		Adaptive	No. of
	Scenario	Best	Worst		Winning
NS1	$p = 1$	0.385 (60)	0.406 (24)	0.378	17/17
	$p = 5$	0.403 (120)	0.706 (24)	0.409	11/17
NS2	$p = 1$	0.388 (84)	0.412 (24)	0.383	17/17
	$p = 5$	0.408 (120)	0.695 (24)	0.416	11/17
NS3	$p = 1$	0.822 (84)	0.878 (24)	0.814	17/17
	$p = 5$	0.862 (120)	1.458 (24)	0.876	11/17

US Yield Curves

Figure 3. Time evolution of US yield curves.

the VAR autoregressive coefficient matrix has regime shifts. In the first regime, the underlying LVAR(5) parameter set is again computed from a VAR(5) model using the three NS factors extracted from the US yield curve from 1983 to 1997. The autoregressive coefficient matrix changes to a new value estimated by using a different sample from 1983 to 2010. Forecasting is conducted using both a LVAR(5) model and a misspecified model, the LVAR of order 1. The set of interval candidates is the same as before, starting from 12 months and ending at 10 years.

Table 4 displays the forecast performance based on both the correct and the misspecified models of LVAR(5) and LVAR(1), respectively. The adaptive models are in most cases superior to the alternatives using various rolling window approaches. Misspecified model LVAR(1) provides even better accuracy with smaller RMSE values than LVAR(5). This implies that the proposed LVAR(1) model is capable of providing reasonable forecast accuracy even when it is misspecified. The simple structure is beneficial for out-of-sample forecasts.

4. REAL DATA ANALYSIS

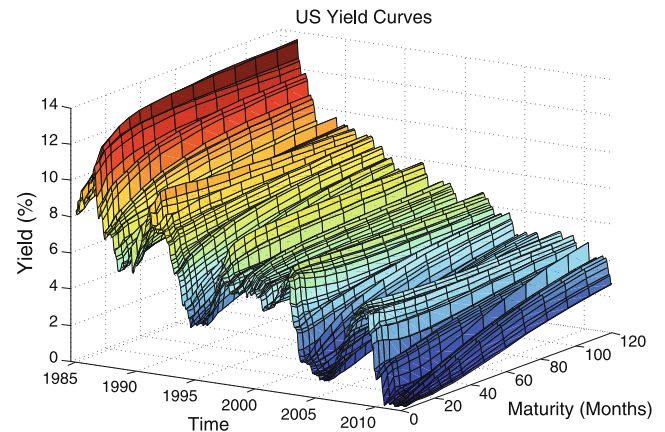In this section, the proposed LVAR model is fitted to monitor and forecast the US Treasuries, spanning January 1983 to September 2010. Each month has a curve for interest rates of 15 maturities (3, 6, 9, 12, 18, 24, 30, 36, 48, 60, 72, 84, 96, 108 and 120 months). The short-term yields on 3 and 6 months are converted from the 3- and 6-month Treasury Bill rates on a discount basis, available from the Federal Reserve's H.15 release of selected interest rates. The remaining yields with maturities of integer years are taken from publicly available research data of the Federal Reserve Board, as released by Gürkaynak, Sack and Wright (2007). We add the 9-, 18-, and 30-month yields interpolated according to the parameters provided in their data file to emphasize the fit for medium-term yields. The time evolution of the yield curves is displayed in Figure 3.

The three Nelson-Siegel (NS) factors are extracted in the framework of Nelson and Siegel (1987):

(3)
$$y_t(\tau) = \beta_{1t} + \beta_{2t}\left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau}\right) + \beta_{3t}\left(\frac{1 - e^{-\lambda\tau}}{\lambda\tau} - e^{-\lambda\tau}\right) + \epsilon_t(\tau),$$

where $\epsilon_t(\tau) \sim N(0, \sigma_\epsilon^2)$ and $y_t(\tau)$ denotes the yield curve with maturity τ (in months) at time t. We follow Diebold and Li (2006) to set $\lambda = 0.0609$ which maximizes the curvature loading at a medium maturity of 30 months. The three factors, β_{1t}, β_{2t} and β_{3t}, represent level, slope and curva-
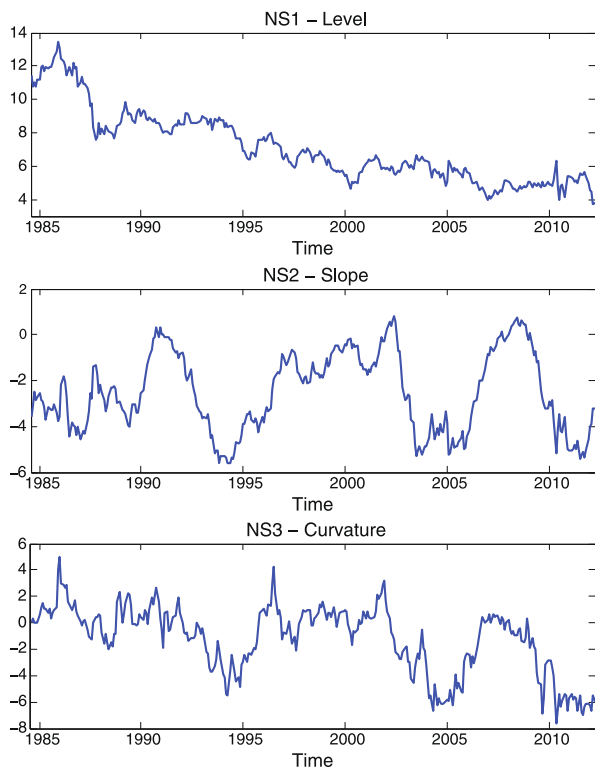
Figure 4. Time evolution of Nelson-Siegel factors extracted from US yield curves.

ture, respectively. The three factors' dynamics are displayed in Figure 4.

Under the NS framework with a fixed value of λ, if there exists any form of non-stationarity in the yield curves y_t, then it is solely attributed to changes in the sequences of the state factors, denoted as $\boldsymbol{X}_t = [\beta_{1t}, \beta_{2t}, \beta_{3t}]'$. The sample autocorrelation and cross-correlation plots of the factors are presented in Figure 5. It shows that the NS factors are cross-dependent not only on its own lagged values but also on the other factors. Moreover, these factors are persistent, with slowly decaying and significant autocorrelations and cross-correlations up to high lag orders, which cannot be easily captured by a VAR with a low order. The persistency feature provides motivation to employ the LVAR model to describe and forecast the factors.

At any forecast time, we use the adaptive technique to identify an interval of local homogeneity, over which we estimate the parameters. The fitted model is then used to iteratively compute multistep ahead forecasts of the NS factors, which are further used to obtain the forecast of interest rates at various maturities in the NS framework. The first forecast originates from December 1997, where we iteratively obtain the 1-month to 12-month ahead forecasts for January 1998 to December 1998. We then move to the next period to identify the interval of local homogeneity and fit the LVAR model for another set of forecasts up to a 12-month ahead

forecast. This estimation and interval selection exercise is executed until September 2009 for the last set of forecasts between October 2009 and September 2010. A total of 142 forecasts are generated for each forecast horizon.

The identified intervals of local homogeneity from December 1997 to September 2009 are shown in Figure 6. The plot shows that as the estimation moves forward, there are some commonly identified ending periods of homogenous intervals, such as 1990, 1993 and 2000–2001, of which the timing of 1990 and 2001 coincide with US economic recessions.

In addition, the forecast results of LVAR are compared with the alternative forecasts using the rolling window technique and the recursive approach. The rolling window technique adopts a fixed window size, 60-month and 120-month in our study, while the recursive approach uses all the past available information. Table 5 summarizes the forecast performance of the three models for NS factors NS1, NS2 and NS3. Table 6 shows the multi-step ahead forecast results of the US yields at 3-month, 12-month, 36-month, 60-month and 120-month maturities. Similar interpretations could be drawn from the two forecast results tables. In several cases, LVAR performs better than the rolling window technique with fixed window size of 60 months, showing smaller RMSE values. However, generally speaking, further improvement of the LVAR model is needed to beat the rolling window with 120-month maturity and the recursive approach.

In fact, as a specific case of this general setting, Chen and Niu (2013) show that restricting the state dynamics to an AR(1) model for each NS factor greatly improves forecast accuracy, and that the resulting performance beats the alternative rolling or recursive forecast uniformly. The reason may be due to the off-diagonal elements in the VAR autoregressive coefficient matrix being typically sparse and close to zeros, which does not contribute much to the improvement of forecast accuracy, but deteriorates the information efficiency. However, as a general illustration of LVAR application, we believe that this example provides a good scenario for monitoring and forecasting economic and financial time series.

5. CONCLUSION

We have proposed a local vector autoregressive (LVAR) model that is capable of providing reasonable forecast accuracy under both homogeneity and structural changes. Compared to the conventional dynamic time series models using either a rolling window with a globally fixed window size or a recursive technique using all past information, the adaptive procedure carefully selects an interval of local homogeneity at any particular time point. With this flexibility on interval selection, the LVAR model provides stable performance both in a simulated homogeneous situation and under regime shift scenarios.

The real data analysis provides an example of real-time monitoring and forecasting of the yield curves. The selected
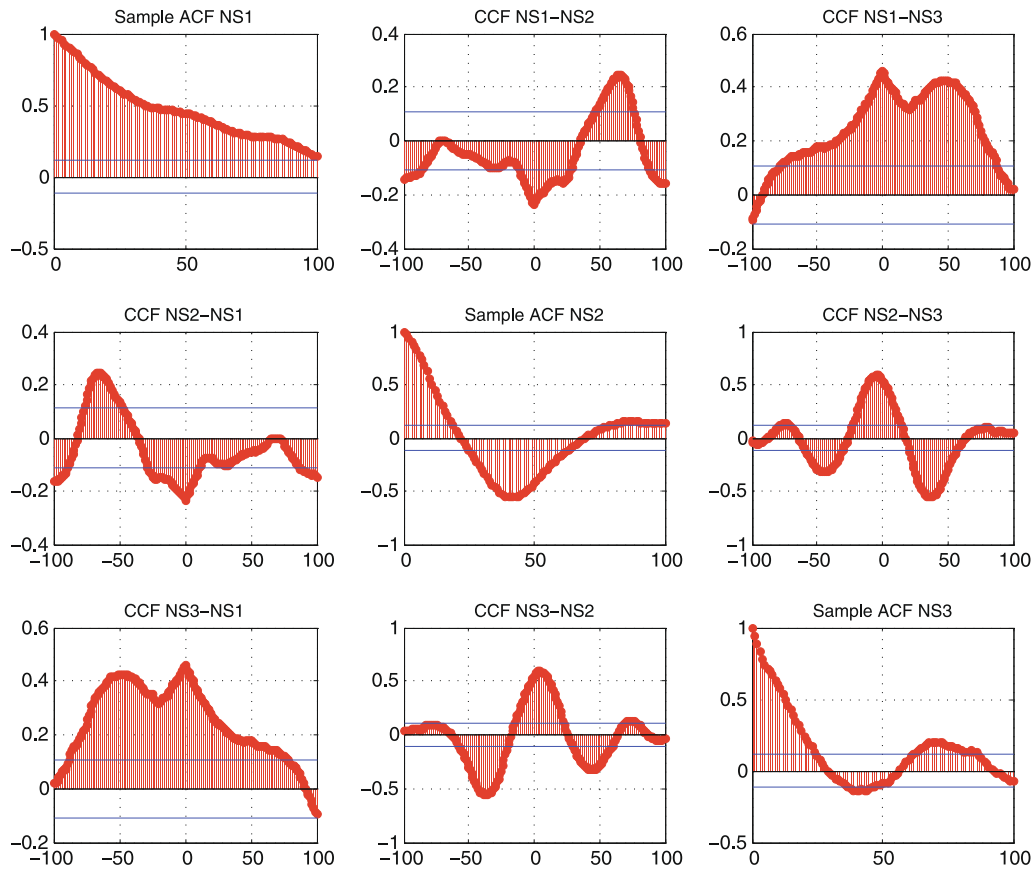
Figure 5. Sample autocorrelations and cross-correlations of the three NS factors.
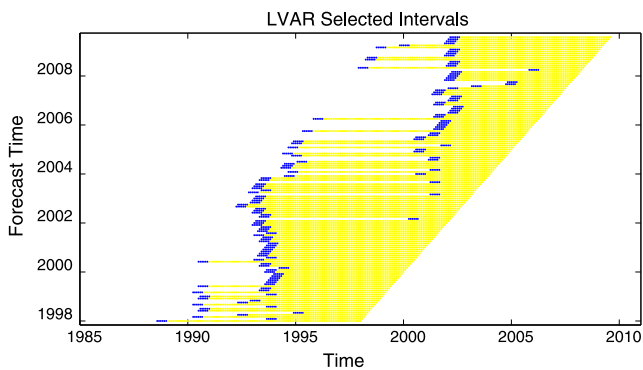

Figure 6. Selected intervals of local homogeneity from December 1997 to September 2009, over which the parameters are estimated and the fitted model is used to obtain the iterative forecasts. The vertical axis represents the time when the estimation is made. The selected interval is marked horizontally as a light yellow line. The dark blue line represents the interval during which the most recent break is detected.

Table 5. RMSE values of the iterative forecasts for NS factors NS1, NS2 and NS3. Three types of models are employed: the LVAR model with a time-dependent interval of local homogeneity, a VAR rolling model with window sizes of 60 months and 120 months, and a recursive VAR model

NS1	h=1	h=3	h=6	h=12
LVAR	0.337	0.487	0.680	0.932
VAR Rolling 60m	0.337	0.475	0.636	0.783
VAR Rolling 120m	0.332	0.461	0.641	0.884
VAR Recursive	0.331	0.437	0.573	0.731
NS2	h=1	h=3	h=6	h=12
LVAR	0.416	0.717	1.178	2.294
VAR Rolling 60m	0.427	0.742	1.219	2.257
VAR Rolling 120m	0.417	0.687	1.069	1.844
VAR Recursive	0.412	0.662	1.003	1.747
NS3	h=1	h=3	h=6	h=12
LVAR	0.976	1.720	2.476	3.832
VAR Rolling 60m	0.927	1.635	2.375	3.288
VAR Rolling 120m	0.886	1.511	2.160	2.864
VAR Recursive	0.875	1.453	2.013	2.638

Table 6. RMSE values of the iterative forecasts for yields at 3-month, 12-month, 36-month, 60-month and 120-month maturities. Three types of models are employed: the LVAR model with a time-dependent interval of local homogeneity, the VAR rolling model with window sizes of 60 months and 120 months, and the recursive VAR model

y(3)	h=1	h=3	h=6	h=12
LVAR	0.285	0.588	1.033	2.061
VAR Rolling 60m	0.281	0.586	1.044	1.896
VAR Rolling 120m	0.273	0.531	0.900	1.613
VAR Recursive	0.268	0.512	0.858	1.545
y(12)	h=1	h=3	h=6	h=12
LVAR	0.286	0.660	1.153	2.147
VAR Rolling 60m	0.276	0.631	1.103	1.882
VAR Rolling 120m	0.261	0.578	0.987	1.665
VAR Recursive	0.262	0.562	0.930	1.559
y(36)	h=1	h=3	h=6	h=12
LVAR	0.344	0.707	1.115	1.861
VAR Rolling 60m	0.326	0.664	1.031	1.557
VAR Rolling 120m	0.322	0.644	0.998	1.509
VAR Recursive	0.318	0.615	0.916	1.351
y(60)	h=1	h=3	h=6	h=12
LVAR	0.350	0.650	0.977	1.510
VAR Rolling 60m	0.336	0.609	0.888	1.240
VAR Rolling 120m	0.333	0.605	0.901	1.302
VAR Recursive	0.326	0.570	0.814	1.132
y(120)	h=1	h=3	h=6	h=12
LVAR	0.302	0.486	0.706	0.994
VAR Rolling 60m	0.299	0.460	0.626	0.769
VAR Rolling 120m	0.297	0.455	0.652	0.910
VAR Recursive	0.304	0.446	0.601	0.778

intervals indicate that there are some underlying economic interpretations behind the detection results, thus LVAR is useful for monitoring purposes. Although the forecast results are not satisfying, which may be due to frequent changes that are not easily detected by a high-dimensional model setting without penalizing parameter uncertainty, we know from existing work that imposing specification structures could improve predictability. We relegate discussion of further improvement within the general LVAR framework to future work.

ACKNOWLEDGEMENT

Received 15 June 2013

REFERENCES

AFONSO, A., BAXA, J. and SLAVIK, M. (2011). Fiscal development and financial stress: A threshold VAR analysis. *Technical report*, European Central Bank. Working paper series, No. 1319.

ATANASOVA, C. (2003). Credit market imperfections and business cycle dynamics: A nonlinear approach. *Studies in Nonlinear Dynamics and Econometrics* **7** 1–22.

BALKE, N. (2000). Credit and economic activity: Credit regimes and nonlinear propagation of shocks. *Review of Economics and Statistics* **82** 344–349.

CHEN, Y., HÄRDLE, W. and PIGORSCH, U. (2010). Localized realized volatility modelling. *Journal of the American Statistical Association* **105** 1376–1393. MR2796557

CHEN, Y. and NIU, L. (2013). Adaptive dynamics Nelson-Siegel term structure model with applications. *Technical report*. Revision submitted.

COGLEY, T. and SARGENT, T. J. (2001). Evolving Post-World War II U.S. inflation dynamics. *NBER Macroeconomics Annual* **16** 331–373.

COGLEY, T. and SARGENT, T. J. (2005). Drifts and volatilities: Monetary policies and outcomes in the Post-WWII U.S. *Review of Economic Dynamics* **8** 262–302.

DIEBOLD, F. X. and INOUE, A. (2001). Long memory and regime switching. *Journal of Econometrics* **105** 131–159. MR1864215

DIEBOLD, F. X. and LI, C. (2006). Forecasting the term structure of government bond yields. *Journal of Econometrics* **130** 337–364. MR2211798

GRANGER, C. W. J. (1980). Long memory relationships and the aggregation of dynamic models. *Journal of Econometrics* **14** 227–238. MR0597259

GRANGER, C. W. J. and HYUNG, N. (2004). Occasional structural breaks and long memory with an application to the S&P 500 absolute stock returns. *Journal of Empirical Finance* **11** 399–421.

GRANGER, C. W. and JOYEUX, R. (1980). An introduction to long memory time series models and fractional differencing. *Journal of Time Series Analysis* **1** 5–39. MR0605572

GÜRKAYNAK, R. S., SACK, B. and WRIGHT, J. H. (2007). The U.S. Treasury yield curve: 1961 to the present. *Journal of Monetary Economics* **54** 2291–2304.

HOSKING, J. R. M. (1981). Fractional differencing. *Biometrika* **68** 165–176. MR0614953

LI, F. and ST-AMANT, P. (2010). Financial stress, monetary policy, and economic activity. *Bank of Canada Review* **2010** 9–18.

NELSON, C. and SIEGEL, A. (1987). Parsimonious modeling of yield curve. *Journal of Business* **60** 473–489.

PRIMICERI, G. E. (2005). Time varying structural vector autoregressions and monetary policy. *The Review of Economic Studies* **72**(3) 821–852. MR2148143

TONG, H. (1983). *Threshold Models in Non-linear Time Series Analysis*. Lecture Notes in Statistics, vol. 21. Springer-Verlag, New York. MR0717388

TONG, H. (1987). *Non-linear Time Series Models of Regularly Sampled Data: A Review*, vol. 2. VNU Science Press. MR1092478

TONG, H. and LIM, K. S. (1980). Threshold autoregression, limit cycles and cyclical data. *Journal of the Royal Statistical Society B* 245–292.

TSAY, R. S. (1989). Testing and modeling threshold autoregressive processes. *Journal of the American Statistical Association* **84** 231–240. MR0999683

TSAY, R. S. (1998). Testing and modeling multivariate threshold models. *Journal of the American Statistical Association* **93** 1188–1202. MR1649212

Ying Chen
Department of Statistics & Applied Probability
National University of Singapore
6 Science Drive 2, 117546
Singapore
E-mail address: stacheny@nus.edu.sg

Bo Li
Department of Statistics & Applied Probability
National University of Singapore
6 Science Drive 2, 117546
Singapore
E-mail address: lbemily86@gmail.com

Linlin Niu
Wang Yanan Institute for Studies in Economics (WISE)
Xiamen University, Xiamen 361005
China
E-mail address: linlin.niu@gmail.com