

Bayesian Lattice Filters for Time-Varying Autoregression and Time–Frequency Analysis

Wen-Hsi Yang*, Scott H. Holan†, and Christopher K. Wikle‡

Abstract. Modeling nonstationary processes is of paramount importance to many scientific disciplines including environmental science, ecology, and finance, among others. Consequently, flexible methodology that provides accurate estimation across a wide range of processes is a subject of ongoing interest. We propose a novel approach to model-based time–frequency estimation using time-varying autoregressive models. In this context, we take a fully Bayesian approach and allow both the autoregressive coefficients and innovation variance to vary over time. Importantly, our estimation method uses the lattice filter and is cast within the partial autocorrelation domain. The marginal posterior distributions are of standard form and, as a convenient by-product of our estimation method, our approach avoids undesirable matrix inversions. As such, estimation is extremely computationally efficient and stable. To illustrate the effectiveness of our approach, we conduct a comprehensive simulation study that compares our method with other competing methods and find that, in most cases, our approach performs superior in terms of average squared error between the estimated and true time-varying spectral density. Lastly, we demonstrate our methodology through three modeling applications; namely, insect communication signals, environmental data (wind components), and macroeconomic data (US gross domestic product (GDP) and consumption).

Keywords: locally stationary, model selection, nonstationary, partial autocorrelation, piecewise stationary, sequential estimation, time-varying spectral density.

1 Introduction

Recent advances in technology have lead to the extensive collection of complex high-frequency nonstationary signals across a wide array of scientific disciplines. In contrast to the time-domain, the time-varying spectrum may provide better insight into important characteristics of the underlying signal (e.g., Holan et al., 2010, 2012; Rosen et al., 2012; Yang et al., 2013, and the references therein). For example, Holan et al. (2010) demonstrated that features in the time–frequency domain of nonstationary *Enchenopa* treehopper mating signals may describe crucial phenotypes of sexual selection. Other examples of using time–frequency methods to extract salient features from nonstationary time series include Holan et al. (2012), Yang et al. (2013), and Martinez et al. (2013),

*CSIRO Computational Informatics, Ecosciences Precinct, GPO Box 2583, Brisbane QLD 4001, Australia, Wen-Hsi.Yang@csiro.au

†(To whom correspondence should be addressed) Department of Statistics, University of Missouri, 146 Middlebush Hall, Columbia, MO 65211-6100, holans@missouri.edu

‡Department of Statistics, University of Missouri, 146 Middlebush Hall, Columbia, MO 65211-6100, wikle@missouri.edu

among others. The former example, Holan et al. (2012), incorporates time–frequency representations of high-dimensional nonstationary time series of daily NASDAQ log-returns, over several quarters, as a regressor in a predictive model for economic recessions. Subsequently, Yang et al. (2013) use time–frequency methods to predict spawning success of shovelnose sturgeon in the Lower Missouri River. Finally, Martinez et al. (2013) use time–frequency representations to study chirp syllables of free-tailed bats using Bayesian functional mixed models.

In general, time–frequency analyses can either proceed using a nonparametric or model-based (parametric) approach. The most common nonparametric approach is the short-time Fourier transform (i.e., windowed Fourier transform) which produces a time–frequency representation characterizing local signal properties (Gröchenig, 2001; Oppenheim and Schaffer, 2009). Another approach to obtain time–frequency representations proceeds using smoothing splines (Rosen et al., 2009, 2012) or by parameterizing the spectral density to estimate the local spectrum via the Whittle likelihood (Everitt et al., 2013). Similarly, time–frequency representations can be achieved by applying smooth localized complex exponential (SLEX) functions to the observed signal (Ombao et al., 2001). In contrast to window based approaches, the SLEX functions are produced using a projection operator and are, thus, simultaneously orthogonal and localized in both time and frequency. Alternatively, one can use the theory of frames and overcomplete bases to produce a time–frequency representation. For example, continuous wavelet transforms (Vidakovic, 1999; Percival and Walden, 2000; Mallat, 2008) or Gabor frames (Wolfe et al., 2004; Feichtinger and Strohmer, 1998; Fitzgerald et al., 2000) could be used. By introducing redundancy into the basis functions, these representations may provide better simultaneous resolution over both time and frequency.

Model-based approaches typically proceed through the time-domain in order to produce a time–frequency representation for a given nonstationary signal. In this setting common approaches include fitting piecewise autoregressive (AR) models as well as time-varying autoregressive (TVAR) models. The former approach assumes that the nonstationary signal is piecewise stationary. Consequently, the estimation procedure attempts to identify the order of the AR models along with the location of each piecewise stationary series. For example, Davis et al. (2006) propose the AutoPARM method using minimum description length (MDL) in conjunction with a genetic algorithm (GA) to automatically locate the break points and AR model order within each segmentation. In addition to providing a time–frequency representation, this approach also locates changepoints. Wood et al. (2011) propose fitting mixtures of AR models within each segment via Markov chain Monte Carlo (MCMC) methods. Their approach selects a common segment length and then divides the signal into these segments prior to implementation of the fitting procedure. Although such approaches may accommodate signals with several piecewise stationary structures, they lack the capability of capturing momentary shocks to the system (i.e., changes to the evolutionary structure that only occur over relatively few time points).

For many processes, TVAR models may provide superior resolution within the time–frequency domain for both large and small scale features through modeling time-varying parameters. To estimate the TVAR model coefficients, Kitagawa and Gersch (1996) and

Kitagawa (2010) treat the coefficients as a stochastic process and model them using difference equations under the assumption of a maximum fixed order of the TVAR model. Their estimation procedure for the coefficients is then based on state-space models with smoothness priors. In this context, the innovation variances are treated as constant and estimated using a maximum likelihood approach. Subsequently, West et al. (1999) propose a fully Bayesian TVAR framework that simultaneously models the coefficients and the innovation variances using random walk models. Alternatively, by assuming a constant innovation variance, Prado and Huerta (2002) model the coefficients and order of the TVAR model using random walk models. Further, to make the TVAR models stable, the constraint that the roots of the characteristic polynomial lie within the unit circle could be imposed. However, such an added condition makes estimation more complicated and computationally expensive.

To avoid these issues, we instead work with the partial autocorrelation coefficients (i.e., in the partial correlation (PARCOR) coefficient domain) and then use the Levinson recursion to connect the PARCOR coefficients and TVAR model coefficients (Kitagawa and Gersch, 1996; Godsill et al., 2004). Godsill et al. (2004) model the PARCOR coefficients and innovation standard deviations using a truncated normal first-order autoregression and a log-Gaussian first-order autoregression, respectively, with a given constant order. To estimate these values, a sequential Monte Carlo algorithm is used. Alternatively, as previously alluded to, by assuming a constant innovation variance, Kitagawa and Gersch (1996) implement the smoothness prior within a lattice filter to estimate the PARCOR coefficients. After the PARCOR coefficients have been estimated, a constant innovation variance is estimated using a maximum likelihood approach. However, the former approach is computationally expensive and may suffer from the degeneracy problem (i.e., the collapse of approximations of the marginal distributions) when TVAR model order is large. In addition, certain hyperparameter values (i.e., the TVAR coefficients associated with the two latent models) may be sensitive to starting values and may require prior knowledge or expert supplied subjective information to achieve convergence (Godsill et al., 2004). Although the latter takes advantage of the lattice form to estimate the PARCOR coefficients, estimation of the innovation variance is achieved outside of the lattice structure; that is, the estimation procedure is a two-stage method. Importantly, this approach is designed for a constant innovation variance and cannot deal with time-dependent innovation variances. As such, we propose a novel approach that addresses these issues within a fully Bayesian context.

Kitagawa (1988) and Kitagawa and Gersch (1996) constitute the first attempts at utilizing the lattice structure to estimate the coefficients of a TVAR model when the innovations follow a Gaussian distribution with mean zero and constant variance. At each stage of the lattice filter, they assume that the residual at each time between the forward and backward prediction errors follows a Cauchy distribution, and that the PARCOR coefficient is modeled as a Gaussian random walk. This produces a non-Gaussian state space model at each stage and thus, a numeric algorithm is conducted for estimation. Moreover, the assumptions of their approach ignore an implicit connection between the innovation term of the TVAR model and the residual term between the forward and backward prediction errors. That is, the distribution of residual term should be a Gaussian distribution rather than a Cauchy distribution. Consequently, their approach leads

to a TVAR model with innovation terms following a Cauchy distribution; hence the innovation variances do not exist. In contrast, our approach assumes the residual term follows a Gaussian distribution.

We propose a fully Bayesian approach to efficiently estimate the TVAR coefficients and innovation variances within the lattice structure. One novel aspect of our approach is that we model both the PARCOR coefficients and the TVAR innovation variances within the lattice structure and then estimate them simultaneously. This is different from the frequentist two-stage method of Kitagawa (1988) and Kitagawa and Gersch (1996). Another novel aspect is that we take advantage of dynamic linear model (DLM) theory (West and Harrison, 1997; Prado and West, 2010) to regularize the PARCOR coefficients instead of using truncated distributions. Thus, our method provides marginal posterior distributions with standard forms for both the PARCOR coefficients and innovation variances. Since our approach takes advantage of the lattice structure, the computational efficiency of our approach is not affected by the order of the TVAR model; that is, our approach avoids having to calculate higher dimensional inverse matrices. To select the TVAR model order, we provide both a visual and a numerical method. Importantly, the simulation study we provide demonstrates that our approach leads to superior performance in terms of estimating the time–frequency representation of various nonstationary signals, as measured by average squared error. Thus, our approach provides a stable and computationally efficient way to fit TVAR models for time–frequency analysis.

The remainder of this paper is organized as follows. Section 2 briefly introduces the lattice structure and describes our methodology along with prior specification. Section 3 presents a comprehensive simulation study that illustrates the effectiveness of our approach across an expansive array of nonstationary processes. Subsequently, in Section 4, our methodology is demonstrated through three modeling applications; namely, insect communication signals, environmental data (i.e., wind components), and macroeconomic data (i.e., US gross domestic product (GDP) and consumption). Lastly, Section 5 concludes with discussion. For convenience of exposition, details surrounding the estimation algorithms and additional figures are left to a Supplemental Appendix.

2 Methodology

2.1 Time-Varying Coefficient Autoregressive Models

The TVAR model of order P for a nonstationary univariate time series x_t , $t = 1, \dots, T$, can be expressed as

$$x_t = \sum_{m=1}^P a_{t,m}^{(P)} x_{t-m} + \epsilon_t, \quad (1)$$

where $a_{t,m}^{(P)}$ and ϵ_t are the TVAR coefficients associated with time lag m at time t and the innovation at time t , respectively. Typically, the innovations are assumed to be uncorrelated mean-zero Gaussian random variables (i.e., $\epsilon_t \sim N(0, \sigma_t^2)$), with time-varying variance σ_t^2). Therefore, the TVAR model corresponds to a nonstationary AR

model with the AR coefficients and variances evolving through time. In such settings, the model is locally stationary but nonstationary globally. As will be illustrated, the assumption of local stationarity is not required for our approach; that is, the forward and backward partial autocorrelations (defined in Section 2.2) need not be equal. Because this model generally allows both slow and rapid changes in the parameters, it can flexibly model the stochastic pattern changes often exhibited by complex nonstationary signals.

2.2 Lattice Structures

The Levinson–Durbin algorithm yields a unique correspondence between the PARCOR coefficients and the AR coefficients (Shumway and Stoffer, 2006; Kitagawa, 2010). Therefore, a disciplined approach to fitting AR models can be achieved through estimation of the PARCOR coefficients. The lattice structure described below provides a direct way of associating the PARCOR coefficients with the observed time series (see Hayes (1996, p. 225) and the Supplementary Appendix (Yang et al., 2015) for additional discussion). As such, the lattice structure provides an effective path to AR model estimation. In fact, the Levinson–Durbin algorithm for a stationary time series can be derived using the lattice structure (see Kitagawa, 2010, Appendix B).

Let $f_t^{(P)}$ and $b_t^{(P)}$ denote the prediction error at time t for a forward and backward AR(P) model, respectively, where

$$f_t^{(P)} = x_t - \sum_{m=1}^P a_m^{(P)} x_{t-m} \quad \text{and} \quad b_t^{(P)} = x_t - \sum_{m=1}^P d_m^{(P)} x_{t+m}.$$

Then, the m th stage of the lattice filter can be characterized by the pair of input–output relations between the forward and backward predictions,

$$f_t^{(m-1)} = \alpha_m^{(m)} b_{t-m}^{(m-1)} + f_t^{(m)}, \tag{2}$$

$$b_t^{(m-1)} = \beta_m^{(m)} f_{t+m}^{(m-1)} + b_t^{(m)}, \quad m = 1, 2, \dots, P, \tag{3}$$

with the initial condition, $f_t^{(0)} = b_t^{(0)} = x_t$, and where $\alpha_m^{(m)}$ and $\beta_m^{(m)}$ are the lag m forward and backward PARCOR coefficients, respectively. Equation (2) shows that the forward PARCOR coefficient of lag m is a regression coefficient of the forward prediction error $f_t^{(m-1)}$ regressed on the backward prediction error $b_{t-m}^{(m-1)}$ and the residual term $f_t^{(m)}$ is the forward prediction error of the forward AR(m) model. Similarly, (3) shows that the backward PARCOR coefficient of lag m is a regression coefficient of the backward prediction error $b_t^{(m-1)}$ regressed on the forward prediction error $f_{t+m}^{(m-1)}$ and the residual term $b_t^{(m)}$ is the backward prediction error of the backward AR(m) model. Using (2) and (3) recursively, we can derive the PARCOR coefficients for a given lag. In the stationary case, the forward and backward PARCOR coefficients are equivalent, i.e., $\alpha_m^{(m)} = \beta_m^{(m)}$.

The PARCOR coefficients $\alpha_P^{(P)}$ are equal to the last component of the coefficients of the forward AR(P) model, i.e., $\alpha_P^{(P)} = a_P^{(P)}$. Using the Levinson–Durbin algorithm,

the remainder of the AR coefficients and the innovation variance can be obtained as follows:

$$a_m^{(P)} = a_m^{(P-1)} - a_P^{(P)} a_{P-m}^{(P-1)}, \quad m = 1, \dots, P-1. \quad (4)$$

This equation implies that once the PARCOR coefficient $\alpha_P^{(P)}$ is estimated, then all of the other coefficients are immediately determined as well.

2.3 The Lattice Structure of the TVAR Model

Given the assumption of second-order stationarity, the forward and backward PARCOR coefficients are constant over time (i.e., shift-invariant). However, since most real-world signals are nonstationary, the shift-invariant PARCOR coefficients are typically inappropriate. In such cases we can modify (2) and (3) as follows:

$$f_t^{(m-1)} = \alpha_{t,m}^{(m)} b_{t-m}^{(m-1)} + f_t^{(m)}, \quad (5)$$

$$b_t^{(m-1)} = \beta_{t,m}^{(m)} f_{t+m}^{(m-1)} + b_t^{(m)}, \quad m = 1, 2, \dots, P, \quad (6)$$

with both the forward and backward PARCOR coefficients $\alpha_{t,m}^{(m)}$ and $\beta_{t,m}^{(m)}$ now time dependent. Note that for notational simplicity, $f_t^{(m-1)}$ and $b_t^{(m-1)}$ here denote the prediction error at time t of the forward and backward TVAR($m-1$). For locally stationary signals, we may impose the constraint that $\alpha_{t,m}^{(m)} = \beta_{t,m}^{(m)}$ at each time t . However, for general nonstationary cases, $\alpha_{t,m}^{(m)}$ and $\beta_{t,m}^{(m)}$ may not be identical at each time t . Therefore, our approach will proceed without this constraint. Also, the residual terms, $f_t^{(m)}$ and $b_t^{(m)}$, are assumed to follow zero-mean Gaussian distributions, $N(0, \sigma_{f,m,t}^2)$ and $N(0, \sigma_{b,m,t}^2)$, respectively. Importantly, when the true process is TVAR(P), the variance $\sigma_{f,P,t}^2$ is equal to the innovation variance σ_ϵ^2 . To verify this statement, we use the fact that PARCOR coefficient of AR(P) is equal to zero when the lag is larger than P . This property can be applied to TVAR models since TVAR models correspond to AR models at each time t . Using such property, $\alpha_{t,m}^{(m)}$ of (5) is equal to zero for $m > P$. Consequently, $f_t^{(m)} = f_t^{(m+1)}$ for $m \geq P$. The forward prediction error $f_t^{(P)}$ is identical to ϵ_t , i.e., $f_t^{(P)}$ and ϵ_t are identically distributed. Therefore, as mentioned in Section 1, the Gaussian distribution provides a more reasonable assumption for the target model (1) than the Cauchy distributions used by Kitagawa (1988) and Kitagawa and Gersch (1996).

In principle, it would be possible to choose distributions for $f_t^{(m)}$ and $b_t^{(m)}$ other than $N(0, \sigma_{f,t,m}^2)$ and $N(0, \sigma_{b,t,m}^2)$, respectively. However, by assuming $f_t^{(m)}$ and $b_t^{(m)}$ follow non-Gaussian distributions, we will lose the computational efficiency of the current model specification. Specifically, under a non-Gaussian distribution assumption, the algorithm in Supplemental Appendix A will no longer be applicable due to lack of conjugacy. Consequently, more computationally expensive algorithms such as sequential Monte Carlo methods would be needed (see Godsill et al., 2004).

For each stage m of the lattice structure, we construct the following equations to obtain the coefficients, $a_{t,k}^{(m)}$ and $d_{t,k}^{(m)}$, of the forward and backward TVAR models (Hayes, 1996; Kitagawa and Gersch, 1996; Haykin, 2002):

$$a_{t,k}^{(m)} = a_{t,k}^{(m-1)} - a_{t,m}^{(m)} d_{t,m-k}^{(m-1)}, \tag{7}$$

$$d_{t,k}^{(m)} = d_{t,k}^{(m-1)} - d_{t,k}^{(m)} a_{t,m-k}^{(m-1)}, \quad k = 1, 2, \dots, m - 1, \tag{8}$$

with $a_{t,m}^{(m)} = \alpha_{t,m}^{(m)}$ and $d_{t,m}^{(m)} = \beta_{t,m}^{(m)}$. Equations (7) and (8) describe the relationship between the coefficients of the forward and backward TVAR models. In particular, these relations illustrate that the forward coefficients at the current stage are a linear combination of the forward and backward coefficients of the previous stage, with the weights equal to the PARCOR coefficients. Importantly, such a combination also includes the stationary and locally stationary cases. For the stationary case, since $\alpha_{t,m}^{(m)} = \beta_{t,m}^{(m)}$ are constant over time, the general equations (7) and (8) can be reduced to (4). For locally stationary cases, since $\alpha_{t,m}^{(m)} = \beta_{t,m}^{(m)}$ at time t , (7) and (8) are identical.

2.4 Model Specification and Bayesian Inference

Since both the forward and backward PARCOR coefficients of (5) and (6) as well as the corresponding innovation variances require time-varying structures, we consider random walk models for their evolutions. In such cases, the following two hierarchical components are added to (5) and (6). The evolution of the forward and backward PARCOR coefficients are modeled, respectively, as follows:

$$\alpha_{t,m}^{(m)} = \alpha_{t-1,m}^{(m)} + \epsilon_{\alpha,m,t}, \quad \epsilon_{\alpha,m,t} \sim N(0, w_{\alpha,m,t}), \tag{9}$$

$$\beta_{t,m}^{(m)} = \beta_{t-1,m}^{(m)} + \epsilon_{\beta,m,t}, \quad \epsilon_{\beta,m,t} \sim N(0, w_{\beta,m,t}), \tag{10}$$

where $w_{\alpha,m,t}$ and $w_{\beta,m,t}$ are time dependent system variances. These system variances are then defined in terms of hyperparameters $\gamma_{f,m}$ and $\gamma_{b,m}$, so called *discount factors* with range (0, 1), respectively (West and Harrison, 1997); see the Supplemental Appendix for further details. Usually, we treat $\gamma_{f,m} = \gamma_{b,m} = \gamma_m$ at each stage m . These two equations also imply a sequential update form that the PARCOR coefficient at time $t + 1$ is equal to the sum of the PARCOR coefficient at time t plus a correction.

Similarly, both the evolution innovation variances, $\sigma_{f,m,t}^2$ and $\sigma_{b,m,t}^2$, are modeled through multiplicative random walks as follows

$$\sigma_{f,m,t}^2 = \sigma_{f,m,t-1}^2 (\delta_{f,m} / \eta_{f,m,t}), \quad \eta_{f,m,t} \sim \text{Beta}(g_{f,m,t}, h_{f,m,t}), \tag{11}$$

$$\sigma_{b,m,t}^2 = \sigma_{b,m,t-1}^2 (\delta_{b,m} / \eta_{b,m,t}), \quad \eta_{b,m,t} \sim \text{Beta}(g_{b,m,t}, h_{b,m,t}), \tag{12}$$

where $\delta_{f,m}$ and $\delta_{b,m}$ are hyperparameters (i.e., discount factors on the range (0,1)), and the multiplicative innovations, $\eta_{f,t,m}$ and $\eta_{b,t,m}$ follow beta distributions with parameters, $(g_{f,m,t}, h_{f,m,t})$ and $(g_{b,m,t}, h_{b,m,t})$, respectively (West et al., 1999). These parameters are defined at each time t by the discount factors, $\delta_{f,m}$ and $\delta_{b,m}$, as detailed in the Supplemental Appendix. In many cases, we also assume that $\delta_{f,m} = \delta_{b,m} = \delta_m$ at each stage m . The series of stochastic error terms $\epsilon_{\alpha,m,t}$, $\epsilon_{\beta,m,t}$, $\eta_{f,m,t}$, and $\eta_{b,m,t}$ are mutually independent, and independent of the forward and backward innovations, $f_t^{(m)}$ and $b_t^{(m)}$ of (5) and (6). It is worth noting that in the context of (9)–(12) that the level of smoothness of the PARCOR coefficients (and hence the TVAR coefficients) is

related to the magnitude of γ whereas the smoothness and deviation of the estimated time-varying innovation variance is related to both γ and δ .

We specify conjugate initial priors for $\alpha_{0,m}^{(m)}$ and $\sigma_{f,m,0}^2$ at each stage m as follows:

$$p(\alpha_{0,m}^{(m)} | D_{f,m,0}, \sigma_{f,m,0}^2) \sim N(\mu_{f,m,0}, c_{f,m,0}), \quad (13)$$

$$p(\sigma_{f,m,0}^{-2} | D_{f,m,0}) \sim G(v_{f,m,0}/2, \kappa_{f,m,0}/2), \quad (14)$$

where $D_{f,m,0}$ denotes the information set at the initial time $t = 0$, $G(\cdot, \cdot)$ is the gamma distribution, $\mu_{f,m,0}$ and $c_{f,m,0}$ are the mean and variance for a normal distribution, and $v_{f,m,0}/2$ and $\kappa_{f,m,0}/2$ are the shape and rate parameters for a gamma distribution. Usually, we treat the starting values $\mu_{f,m,0}$, $c_{f,m,0}$, $v_{f,m,0}$, and $\kappa_{f,m,0}$ as common constants over stage m . Typically, we choose $\mu_{f,m,0}$ and $c_{f,m,0}$ to be zero and one, respectively. In addition, to set $v_{f,m,0}$, we first fix $v_{f,m,0} = 1$ and calculate the sample variance of the initial components of the signal. Given these two values, we can obtain $\kappa_{f,m,0}$ through the formula for the expectation of the gamma distribution. In such prior settings, the DLM sequential filtering and smoothing algorithms provide the necessary components for the marginal posterior distributions (West and Harrison, 1997). Specifically, for each stage $m = 1, \dots, P$ and times $t = 1, \dots, T$, with the information set $D_{f,m,T}$ up to time T , the marginal posterior distributions $p(\alpha_{t,m}^{(m)} | D_{f,m,T})$ is a t -distribution and $p(\sigma_{f,m,t}^{-2} | D_{f,m,T})$ is approximated by a gamma distribution. Analogous to $\alpha_{0,m}^{(m)}$ and $\sigma_{f,m,0}^2$, the same conjugate initial priors for $\beta_{0,m}^{(m)}$ and $\sigma_{b,m,0}^2$ are specified at each stage m . In general, based on a sensitivity analysis, the simulation studies provided in Section 3 and the case studies provided in Section 4 are not overly sensitive to the prior specification, with the chosen priors able to glean the important features contained in the time–frequency representation. Details of the sequential filtering and smoothing for the PARCOR coefficients and innovation variances for each stage m are discussed in the Supplemental Appendix.

2.5 Model Selection

Selection of the model order and set of discount factors $\{P, \gamma_m, \delta_m; m = 1, \dots, P\}$ is essential for our approach. First, one can assume $\gamma_m = \gamma$ and $\delta_m = \delta$, for $m = 1, \dots, P$. Then, the analysis will proceed using a set of various pre-specified combinations of (P, γ, δ) . Since γ is related to the variability of the PARCOR coefficients, it also affects the variability of the TVAR coefficients. Hence, one can model the variance of the time-varying coefficients and the innovation variances of the TVAR models using discount factors γ and δ , respectively (West et al., 1999). Note that, in our context, the discount factors (γ_m, δ_m) are a function of m (the lattice filter stage), whereas the West et al. (1999) setting does not make use of the lattice filter and, thus, there is only one set of discount factors (γ, δ) that need to be estimated. However, estimation of (P, γ, δ) using the approach of West et al. (1999) entails repeatedly having to calculate inverse matrices in the sequential filtering process. Our approach allows γ_m and δ_m to vary by stage. Thus, we first specify a potential maximum value of P and a set of combinations of $\{\gamma_m, \delta_m\}$ for each stage m . Given a value of P , we search for the combination of $\{\gamma_1, \delta_1\}$ maximizing the log-likelihood of (5) at stage one. Using the selected γ_1 and

δ_1 , we can obtain the corresponding series $\{f_t^{(2)}\}$ and $\{b_t^{(2)}\}$, for $t = 1, \dots, T$, as well as the value, \mathcal{L}_1 , of log-maximum-likelihood of (5). We then, repeat the above search procedure for stage two using the output $\{f_t^{(2)}\}$ and $\{b_t^{(2)}\}$ obtained from implementing the selected hyperparameters γ_1 and δ_1 . In turn, this produces a new series of $\{f_t^{(3)}\}$ and $\{b_t^{(3)}\}$, for $t = 1, \dots, T$, as well as a value \mathcal{L}_2 . We repeat the procedure until the set of $\{\gamma_m, \delta_m, \mathcal{L}_m\}$, $m = 1, \dots, P$, has been selected.

Here, we provide both a visual and numerical method to select the order. Similar to the scree plot widely used in multivariate analysis (Rencher, 2002), we can plot \mathcal{L}_m against the order m . When the observed series follows an AR or TVAR model, the values of \mathcal{L}_m will stop increasing after a specific lag, this lag can be chosen as the order for the estimated model. Henceforth, this plot is referred to as “BLF-scree.” This type of visual order determination can be directly quantified through the relative change of \mathcal{L}_m . Specifically, we provide a numerical method of order selection based on calculating the percent change in going from \mathcal{L}_{m-1} to \mathcal{L}_m with respect to m ,

$$|(\mathcal{L}_m - \mathcal{L}_{m-1})/\mathcal{L}_{m-1}| * 100 < \tau. \tag{15}$$

Based on simulation of various TVAR models, we choose $\tau = 0.5$ with $m - 1$ reflecting the “best” value for the order. That is, we have found that 0.5 provides an effective cut-off for choosing the order. Although this approach provides a good guide to order selection, other model selection methods could be considered (e.g., shrinkage through Bayesian variable selection, reversible jump MCMC, or by minimizing an information criteria). Development of alternative model selection approaches in this setting constitutes an area of future research.

We now summarize our approach for fitting TVAR models. Given a set of hyperparameters $\{P, \gamma_m, \delta_m; m = 1, \dots, P\}$, the procedure starts by setting $f_t^{(0)} = b_t^{(0)} = x_t$, for $t = 1, \dots, T$. Next, plugging $\{f_t^{(0)}\}$ and $\{b_t^{(0)}\}$ into (5) and (6) and using sequential filtering and smoothing algorithms, we obtain a series of estimated parameters $\{\hat{\alpha}_{t,1}^{(1)}\}$, $\{\hat{\beta}_{t,1}^{(1)}\}$, $\{\hat{\sigma}_{f,1,t}^2\}$, and $\{\hat{\sigma}_{b,1,t}^2\}$, as well as the new series of forward and backward prediction errors, $\{f_t^{(1)}\}$ and $\{b_t^{(1)}\}$, for $t = 1, \dots, T$. We then repeat the above procedure until $\{\hat{\alpha}_{t,P}^{(P)}\}$, $\{\hat{\beta}_{t,P}^{(P)}\}$, $\{\hat{\sigma}_{f,P,t}^2\}$, and $\{\hat{\sigma}_{b,P,t}^2\}$ have been obtained. Then, recursively plugging the estimates of $\{\alpha_{t,m}^{(m)}\}$ and $\{\beta_{t,m}^{(m)}\}$, from $m = 1, \dots, P$ into (7) and (8), we obtain the estimated time-varying coefficients of (1). As part of this algorithm, the series of estimated innovation variances are equal to $\{\hat{\sigma}_{f,P,t}^2\}$. Finally, for $t = 1, \dots, T$, the time-frequency representation associated with the TVAR(P) model can be obtained by the following equation:

$$S(t, \omega) = \frac{\sigma_t^2}{\left|1 - \sum_{m=1}^P a_{t,m}^{(P)} \exp(-2\pi i m \omega)\right|^2}, \quad -1/2 \leq \omega \leq 1/2, \tag{16}$$

where $i = \sqrt{-1}$ (Kitagawa and Gersch, 1996). Plugging the estimated values $\hat{a}_{t,m}^{(P)}$, $m = 1, \dots, P$, and $\hat{\sigma}_{f,P,t}^2$ into (16) yields the estimated time-varying AR(P) time-frequency representation $\hat{S}(t, \omega)$. See the Supplemental Appendix for further discussion.

3 Simulation Studies

In this section, we simulate various nonstationary time series in order to compare the performance of our approach with five other approaches used to estimate the time–frequency representation. The first approach is AdaptSPEC proposed by Rosen et al. (2012). This approach adaptively segments the signal into finite pieces and then estimates the time–frequency representation using smoothing splines to fit local spectra via the Whittle likelihood approximation. The size of a segment and the number of spline basis functions are two essential parameters for this approach. To reduce any subjectivity in our comparisons, we choose settings for these two parameters similar to those considered in Rosen et al. (2012) (with their $t_{min} = 40$), as well as the same settings for MCMC iterations and burn-in. However, rather than using 10 spline basis functions we use 15, as this provides slightly better results along with superior computational stability. We note that the approach of Everitt et al. (2013) is not considered here due to the fact that the parameterization of the spectral density through the Wittle likelihood requires subjective knowledge.

The second approach is the AutoPARM method (Davis et al., 2006). Although this approach combines the GA and MDL to automatically search for potential break points along with the AR orders for each segment, four parameters are crucial for the GA: the number of islands, the number of chromosomes in each island, the number of generations for migration, and the number of chromosomes replaced in a migration; see Davis et al. (2006) for a comprehensive discussion. All of these parameters were chosen identical to those used in Davis et al. (2006).

The third approach is the AutoSLEX method (Ombao et al., 2001). Given a fixed value for the complexity penalty parameter of the cost function, AutoSLEX can automatically segment a given signal and choose a smoothing parameter. Following the suggestion of Ombao et al. (2001), we set this parameter equal to one. The fourth approach we consider is the approach of Kitagawa and Gersch (1996), referred to as KG1996. The last method we consider is the approach of West et al. (1999), referred to as WPK1999. This approach requires specification of three parameters: the TVAR order and two discount factors – one associated with the variance of the time-varying coefficients and the other with the innovation variances. In general, the discount factor values are in the range 0.9 – 0.999 (West et al., 1999). Therefore, for our simulations, we give each discount factor a set of values from 0.8 to 1 (with equal spacing of 0.02) and, further, a set of values for the TVAR order from 1 to 15. Given these values, we choose the combination that achieves the maximum likelihood (West et al., 1999).

Our approach uses the two selection methods discussed in Section 2.5 to search for the TVAR order with appropriate discount factor values. The selected combination of (P, γ, δ) with γ and δ held fixed over all stages of the Bayesian lattice filter is referred to as BLFFix. The selected combination of (P, γ_m, δ_m) , for $m = 1, \dots, P$, is referred to as BLFDyn. Again, the candidate space of parameters for both discount factors are from 0.8 to 1 (with equal spacing of 0.02), along with orders from 1 to 15.

We consider five types of nonstationary signals: (i) TVAR of order 6 with constant innovation variance; (ii) a piecewise AR process with constant innovation variance;

(iii) simulated signals based on an *Enchenopa* treehopper communication signal (Holan et al., 2010) – see Section 4.1; (iv) TVAR of order 2 with time-varying innovation variance; and (v) TVMA of order 2 with time-varying innovation variance. Note that these simulations were specifically chosen based on previous simulations in the literature coupled with the necessity of having a known time–frequency representation. These simulations also are designed to demonstrate the performance in the time-varying innovation variance context. Each simulation consists of 200 realizations. To evaluate the performance in estimating the various time–frequency representations, we calculate the average squared error (ASE) for each realization as follows (Ombao et al., 2001):

$$ASE_n = (TL)^{-1} \sum_{t=1}^T \sum_{l=1}^L \left(\log \widehat{S}(t, \omega_l) - \log S(t, \omega_l) \right)^2, \quad (17)$$

where $n = 1, \dots, 200$, $\omega_l = 0, 0.005, \dots, 0.5$, and T denotes the length of the simulated series. Lastly, we denote $\overline{ASE} = (1/200) \sum_{n=1}^{200} ASE_n$. For each simulation study, Table 1 summarizes the mean values and standard deviations for ASE_n .

In the case of the AutoSLEX method the number of frequencies in (17) differs from the other approaches considered. In particular, the AutoSLEX approach dyadically segments the signal up to a given maximum scale J such that 2^J is less than signal length, T . Subsequently, AutoSLEX automatically determines whether a segment at a particular scale will be included in final segmentation. Once this has been completed, the frequency resolution for the AutoSLEX approach is equal to $T/2^{(j+1)}$ where j is the scale of the largest segment included in the final segmentation.

	TVAR6	PeiceAR	ASE(sd _{ASE})		TVAR2Tvar	TVMA2Tvar
			SimBugs			
AdaptSPEC	0.2195 (0.0351)	0.1070 (0.0227)	1.4760 (0.2237)		3.2913 (0.9571)	0.6000 (0.0837)
AutoPARM	0.2233 (0.0439)	0.0702 (0.0392)	0.6711 (0.0725)		7.0674 (0.0253)	0.6227 (0.0713)
AutoSLEX	0.2885 (0.0796)	0.1141 (0.0296)	2.3993 (0.6078)		3.2434 (0.6721)	1.5568 (0.4414)
WPK1999	0.0771 (0.0171)	0.0931 (0.0218)	0.4627 (0.0796)		1.8253 (0.6049)	0.3231 (0.0370)
BLFFix	0.0841 (0.0191)	0.0921 (0.0205)	0.3444 (0.0659)		1.6530 (0.5701)	0.3117 (0.0364)
BLFDyn	0.0543 (0.0276)	0.1607 (0.1107)	0.3237 (0.1075)		1.0496 (0.3096)	0.3151 (0.0383)
KG1996	0.1914 (0.2955)	0.1204 (0.0513)	17.5660 (0.3336)		5.7748 (0.1631)	7.1856 (0.0315)

Table 1: The mean, \overline{ASE} and standard deviation, sd_{ASE} , of the ASE values for the simulations. Note that the bold values represent the approach having minimum \overline{ASE} . TVAR2Tvar and TVMA2Tvar denote the TVAR(2) and TVMA(2) with time-varying innovation variance, respectively.

3.1 Time-Varying AR(6) Process

We consider signals from the same time-varying AR(6) process of order six (TVAR6) used in Rosen et al. (2009). This time-varying AR(6) process can be compactly expressed as $\phi_t(B)x_t = \epsilon_t$, $t = 1, \dots, T$, in terms of a characteristic polynomial function $\phi_t(B)$, with $\epsilon_t \stackrel{iid}{\sim} N(0, 1)$ and B the backshift operator (i.e., $B^p x_t = x_{t-p}$). The characteristic polynomial function for this process can be factorized as

$$\phi_t(B) = (1 - a_{t,1}B)(1 - a_{t,1}^*B)(1 - a_{t,2}B)(1 - a_{t,2}^*B)(1 - a_{t,3}B)(1 - a_{t,3}^*B),$$

where the superscript $*$ denotes the complex conjugate. Also, for $p = 1, 2, 3$, let $a_{t,p}^{-1} = A_p \exp(2\pi i \theta_{t,p})$, where the $\theta_{t,p}$ s are defined by $\theta_{t,1} = 0.05 + (0.1/(T-1))t$, $\theta_{t,2} = 0.25$, and $\theta_{t,3} = 0.45 - (0.1/(T-1))t$, with $T = 1024$ and the values of A_1 , A_2 , and A_3 equal to 1.1, 1.12, and 1.1, respectively.

The BLF-sree plot (not shown) suggests that order six is the appropriate choice for all 200 realizations. The TVAR(6) contains three pairs of time-varying conjugate complex roots. Figure 1 illustrates that the TVAR(6) has a time-varying spectrum with three peaks. For this case, the TVAR-based models outperform the group of non-TVAR-based models, with BLFDyn performing superior to the others. Moreover, the percent reduction in $\overline{\text{ASE}}$ is significant relative to the other approaches considered (Table 1).

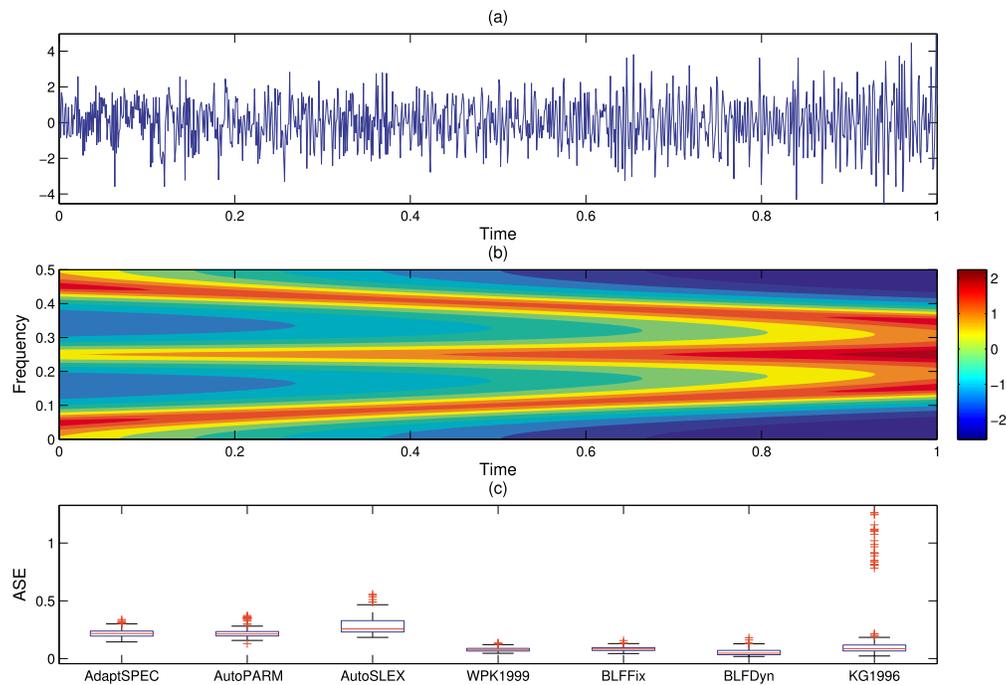


Figure 1: Plots (a) and (b) depict one realization along with the true time–frequency representation of the time-varying AR(6) process (TVAR6), respectively (Section 3.1). Plot (c) illustrates the box-plots of the average squared error (ASE) values corresponding to the time–frequency representation of the TVAR6 for all of the approaches considered.

3.2 Piecewise Stationary AR Process

The signals simulated here are based on the same piecewise stationary AR process, used by Davis et al. (2006) and Rosen et al. (2009, 2012) and is defined as follows:

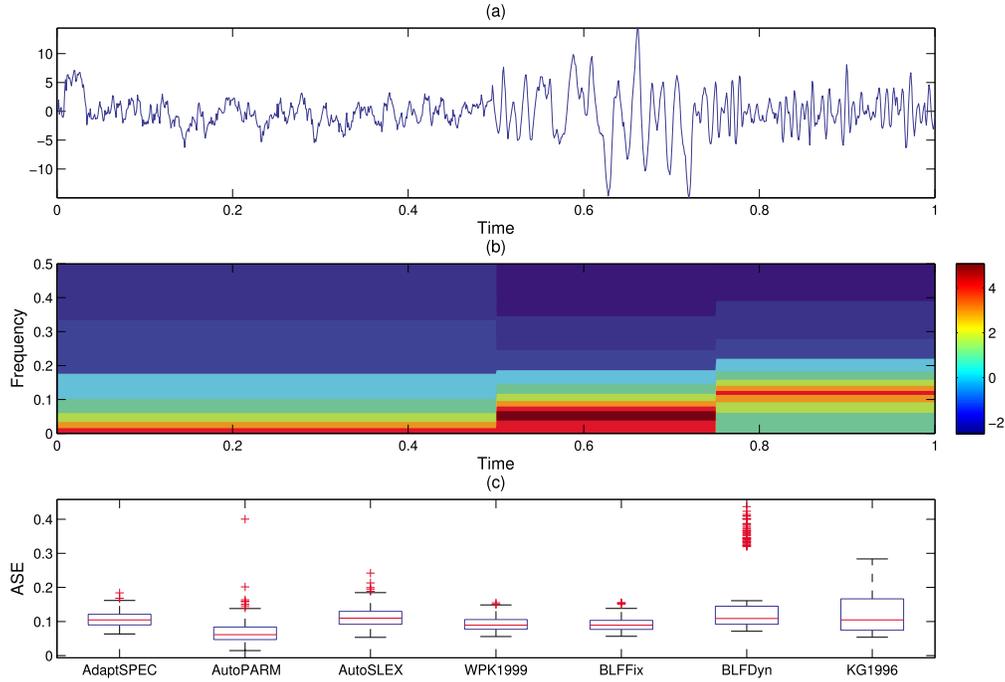


Figure 2: Plots (a) and (b) depict one realization along with the true time–frequency representation of the piecewise AR process (PieceAR), respectively (Section 3.2). Plot (c) illustrates the box-plots of the average squared error (ASE) values corresponding to the time–frequency representation of the PieceAR for all of the approaches considered.

$$x_t = \begin{cases} 0.9x_{t-1} + \epsilon_t; & \text{if } 1 \leq t \leq 512, \\ 1.69x_{t-1} - 0.81x_{t-2} + \epsilon_t; & \text{if } 513 \leq t \leq 768, \\ 1.32x_{t-1} - 0.81x_{t-2} + \epsilon_t; & \text{if } 769 \leq t \leq 1024, \end{cases}$$

where $\epsilon_t \stackrel{\text{iid}}{\sim} N(0, 1)$. These generated signals are referred to as PieceAR. Since it is difficult to choose the order for some realizations visually using the BLF-sctree plot, we use (15), with $\tau = 0.5$, to choose the order. The numerical method suggests order two for some realizations and order three for the others. The true process includes three segments, with each of the segments mutually independent. The piecewise nature of this process is clearly depicted by its time-varying spectrum (Figure 2(b)). The box-plot (Figure 2(c)) shows that AutoPARG exhibits superior performance in terms of the smallest median ASE. However, WPK1999 and BLFFix may perform more robustly (i.e., less outlying ASE values). Although we see a 23.78% reduction in ASE for AutoPARG versus BLFFix, we find that BLFFix performs superior to the remainder of the approaches and has the smallest standard deviation across the 200 simulations. Table 1 summarizes the mean values and standard deviations for this simulation. The findings here are not surprising as AutoPARG is ideally suited toward identifying and estimat-

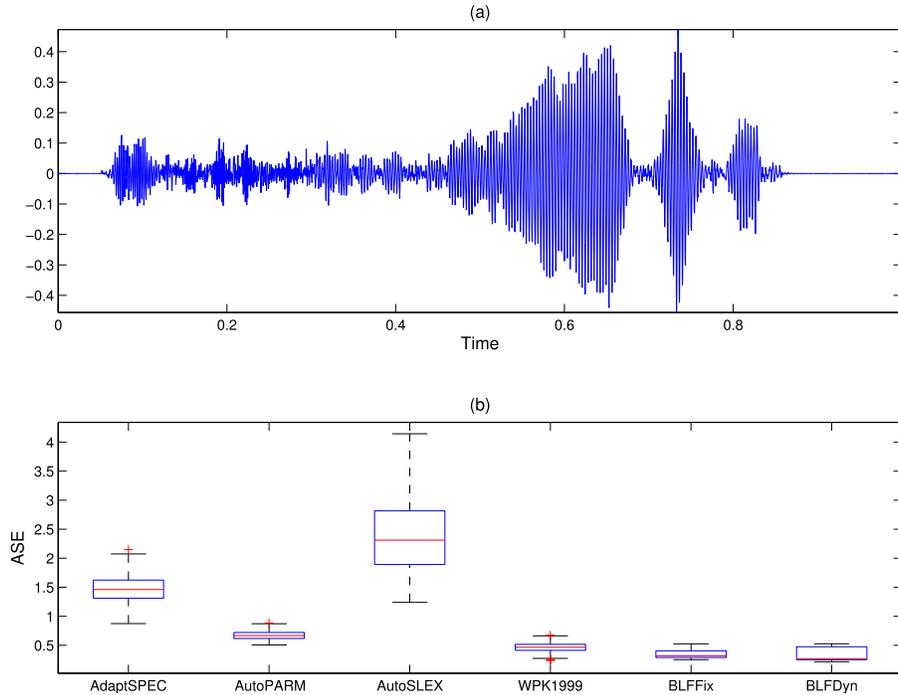


Figure 3: Plot (a) depicts one realization of the simulated insect communication signals (SimBugs), Section 3.3. Plot (b) illustrates the box-plots of the average squared error (ASE) values corresponding to the time–frequency representation of the SimBugs for all of the approaches considered. Note that the minimum ASE for the KG1996 approach is 16.49.

ing piecewise AR processes. For our approach, taking (γ, δ) fixed is advantageous for processes that are not slowly-varying.

3.3 Simulated Insect Communication Signals

The signals considered in this simulation are formulated such that they exhibit the same properties as an *Enchenopa* treehopper mating signal; see Section 4.1 for a complete discussion. Specifically, we fit a TVAR(6) model to the signal x_t , $t = 1, \dots, 4096$, to obtain time-varying AR coefficients and innovation variances. Typically, with these type of nonstationary signals, the innovation variances are time dependent, which is markedly different from the previous examples where the innovation variance was constant. The signals generated by these parameters are referred to as SimBugs. As expected, the BLF-screep plot (not shown) suggests that order six may be an appropriate choice for all 200 realizations. Figure 3(a) illustrates one realization of the SimBugs, whereas Figure 3(b) provides box-plots that characterize the distribution of ASE_n over the 200 simulated signals. Specifically, from Figure 3(b), we see that BLFFix and BLFDyn perform better

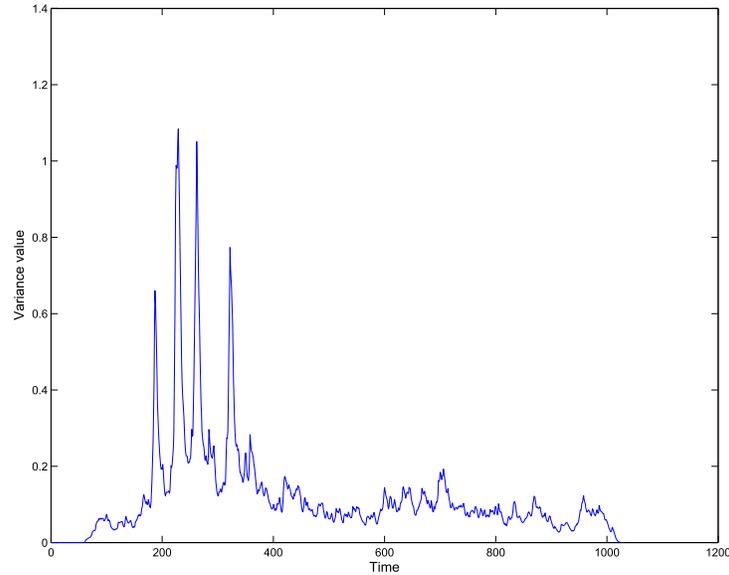


Figure 4: Time-varying innovation variance used for the TVAR2Tvar and the TVMA2Tvar simulations in Sections 3.4 and 3.5, respectively.

than the other approaches, in terms of median ASE. Further, we find that BLFFix and BLFDyn are similar, in terms of $\overline{\text{ASE}}$, although the median of BLFDyn (0.2675) is smaller than that of BLFFix (0.3176).

Table 1 summarizes the mean values and standard deviations for this simulation. From this table we see that the reduction in $\overline{\text{ASE}}$ for BLFDyn is 6.01% over BLFFix and that both exhibit a substantial percent reduction in $\overline{\text{ASE}}$ over the other methods considered.

3.4 Time-Varying AR(2) Process with Time-Varying Variance

We simulate 200 signals from a time-varying AR(2) process with time-varying variance (TVAR2Tvar), which is defined as follows:

$$\begin{aligned}x_t &= a_t x_{t-1} - 0.81x_{t-2} + \epsilon_t, \\ a_t &= 0.8(1 - 0.5\cos(\pi t/1024)),\end{aligned}$$

where $\epsilon_t \stackrel{\text{iid}}{\sim} N(0, \sigma_t^2)$ and $t = 1, \dots, 1024$. Figure 4 illustrates the time-varying innovation, σ_t^2 , used to simulate the TVAR2Tvar signals. Aside from the time-varying innovation variance, the time-varying AR(2) (TVAR2) process used for this simulation is identical to the TVAR2 used in Davis et al. (2006) and Rosen et al. (2009, 2012). Specifically, Davis et al. (2006) and Rosen et al. (2009, 2012) used these equations to generate time-varying AR(2) processes with constant variance. Since the time-varying

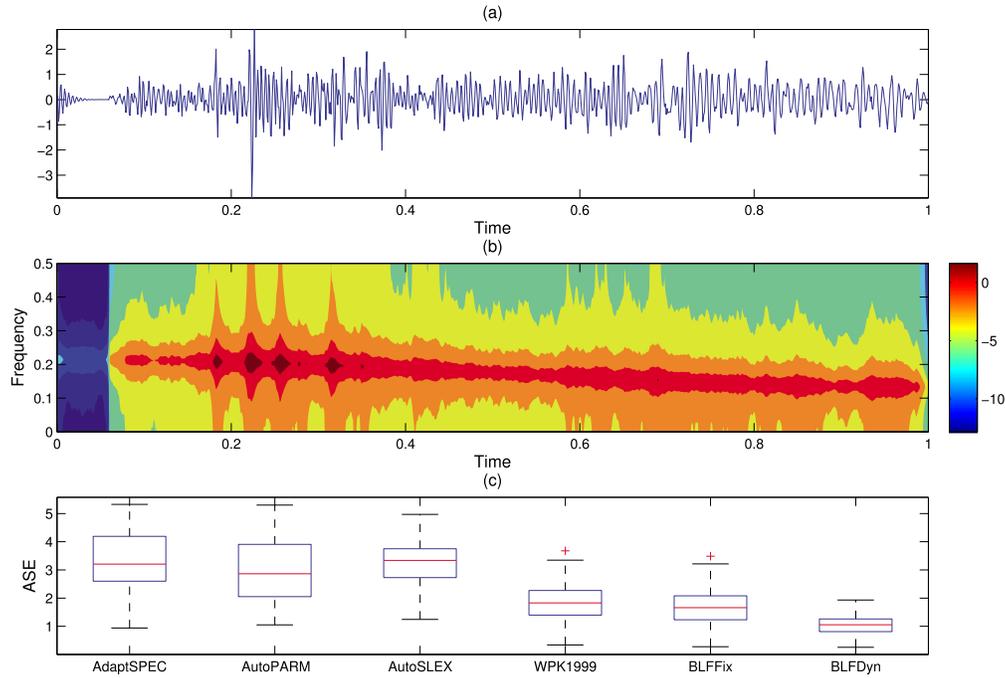


Figure 5: Plots (a) and (b) depict one realization along with the true time–frequency representation of TVAR2Tvar, respectively. Plot (c) illustrates the box-plots of the average squared error (ASE) values corresponding to the time–frequency representation of the TVAR2Tvar for all of the approaches considered. Note that the minimum ASE the KG1996 approach is 7.0413.

coefficient a_t varies slowly with time, this process naturally exhibits a slowly evolving time-varying spectrum (Figure 5). In contrast, the time-varying innovation variance produces several regions of higher spectral energy at the beginning of the time series. For this simulation, the order is chosen using (15), with $\tau = 0.5$. The box-plots of the ASE values in Figure 5 show that the group of TVAR-based models with time-varying innovations (i.e., WPK1999, BLFFix, and BLFDyn) perform superior to the group of non-TVAR-based models (i.e., AdaptSPEc, AutoPARM, and AutoSLEX), with BLFDyn performing the best. In particular, there is a significant percent reduction in $\overline{\text{ASE}}$ for BLFDyn relative to the other methods considered (Table 1).

3.5 Time-Varying MA(2) Process with Time-Varying Variance

We also simulate 200 signals from a time-varying MA(2) process with time-varying variance (TVMA2Tvar), which is defined as follows:

$$\begin{aligned} x_t &= 0.5\epsilon_{t-2} + a_t\epsilon_{t-1} + \epsilon_t, \\ a_t &= 1.122\{1 - 1.781\sin(\pi t/2048)\}, \end{aligned}$$

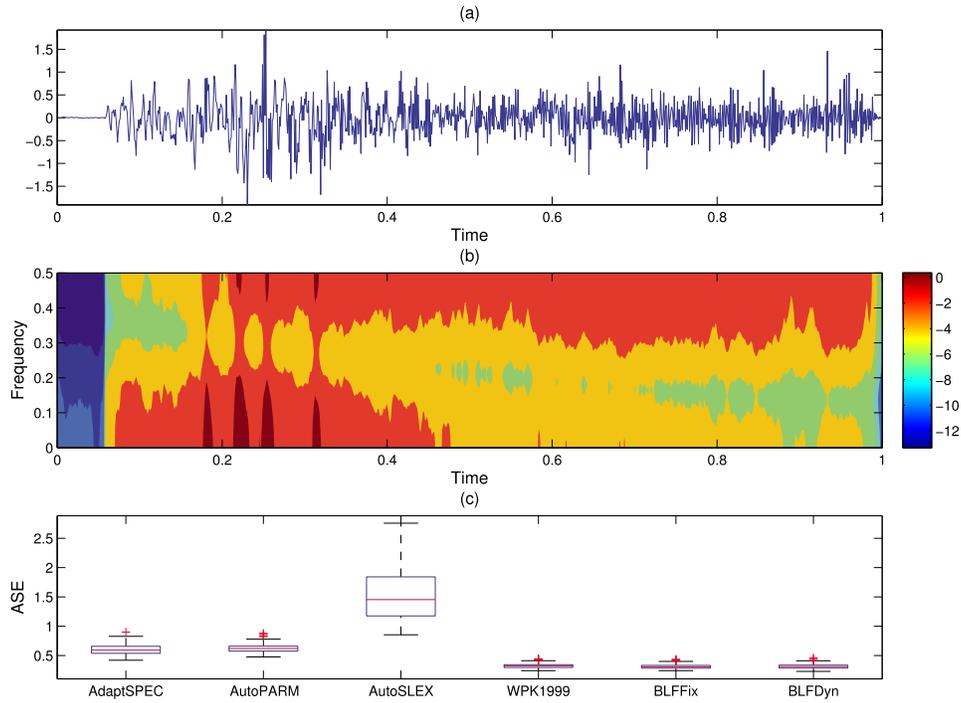


Figure 6: Plots (a) and (b) depict one realization along with the true time–frequency representation of TVMA2Tvar, respectively. Plot (c) illustrates the box-plots of the average squared error (ASE) values corresponding to the time–frequency representation of the TVMA2Tvar for all of the approaches considered. Note that the minimum ASE for the KG1996 approach is 7.1236.

where $\epsilon_t \stackrel{iid}{\sim} N(0, \sigma_t^2)$ and $t = 1, \dots, 1024$. Figure 4 illustrates the usage of σ_t^2 for simulating the TVMA2Tvar signals. Davis et al. (2006) also used these equations to simulate time-varying MA(2) processes but with constant variance. Similar to the TVAR2Tvar signals, the order is chosen using (15), with $\tau = 0.5$. Again, the box-plots of the ASE values in Figure 6 show that the group of TVAR-based models with time-varying innovations (i.e., WPK1999, BLFFix, and BLFDyn) perform superior to the group of non-TVAR-based models (i.e., AdaptSPEC, AutoPARM, and AutoSLEX), with BLFFix performing the best. In particular, there is a significant percent reduction in \overline{ASE} for BLFFix relative to AdaptSPEC, AutoPARM, AutoSLEX, and the KG1996 approach (Table 1).

4 Case Studies

4.1 Animal Communication Signals

Understanding the dynamics of populations is an important component of evolutionary biology. Many organisms exhibit complex characteristics that intricately relate to fitness.

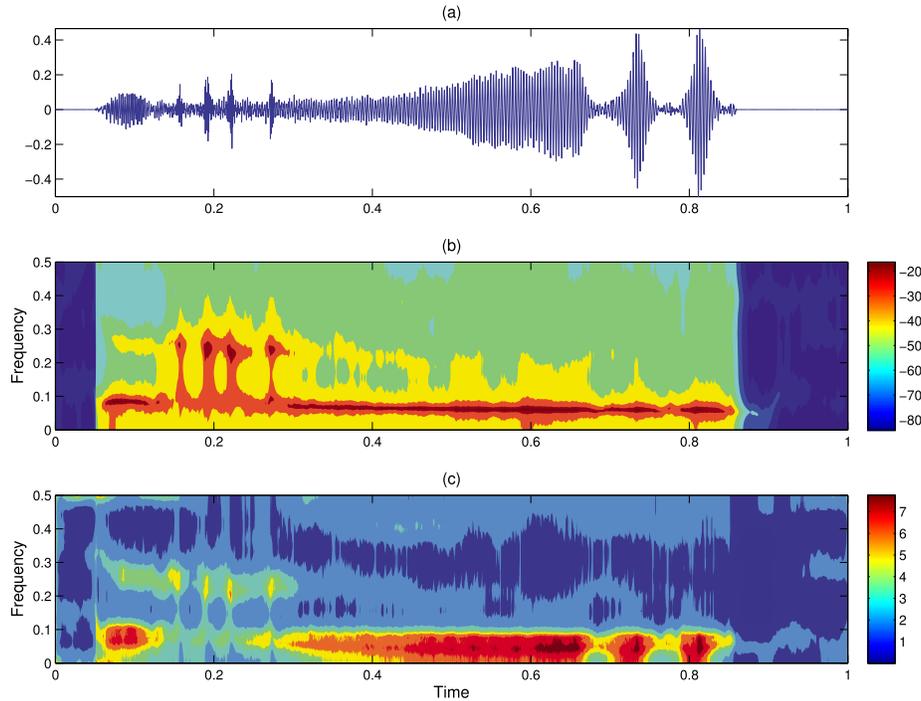


Figure 7: (a) An example of typical signal corresponding to a successful mated (Section 4.1). Plots (b) and (c) present posterior mean and standard deviation of the TVAR(6) spectral representation of the signal in plot (a).

For example, the mating signal of the *Enchenopa* treehopper represents a phenotype of the insect that is used in mate selection. During the mating season, males in competition deliver their vibrational signals through stems of plants to females (see Cocroft and McNett, 2006, and the references therein). The data considered here comes from an experiment that was previously analyzed in Cocroft and McNett (2006) and Holan et al. (2010). The experiment was designed with the goal of reducing potential confounding effects between environmental and phenotypical variation. In this experiment, males signals were recorded one week prior to the start of mating. Figure 7(a) displays a typical signal of from a successful mated, with length 4,739 downsampled from registered signals of length 37,912. Justification for the appropriateness of downsampling the original signal, in this context, can be found in Holan et al. (2010). Also, as discussed in Holan et al. (2010), this signal shows a series of broadband clicks preceding a frequency-modulated sinusoidal component, followed by a series of pulses.

For this analysis, we used the BLFDyn approach to search for a model having both discount factors in the range of 0.8 to 1 (with equal spacing of 0.02) and an order between 1 and 25. Figure 8(a) shows an increase in \mathcal{L}_m along with the order, which is different from the simulated TVAR(6) model in Section 3.3. Hence, we use (15) with the

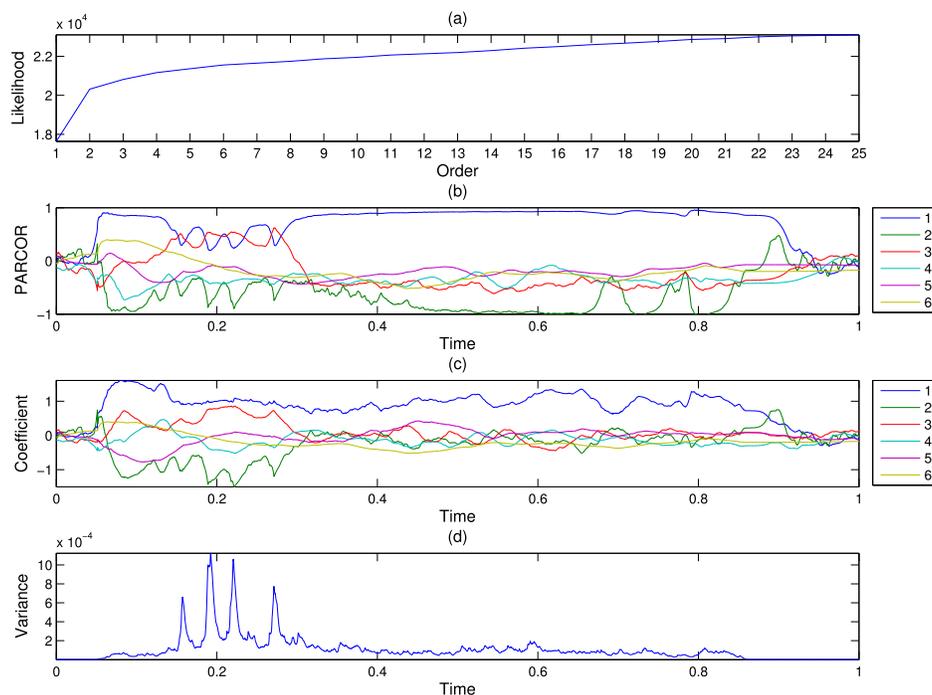


Figure 8: Plot (a) shows the BLF-snee plot of the treehopper communication signal. Plot (b) depicts the first six time-varying estimated PARCOR coefficients. Plots (c) and (d) show the estimated time-varying coefficients and innovation variances of the TVAR(6) model.

model order chosen by $\tau = 0.5$. This rule yields a TVAR(6) model. In Figure 8(b), the PARCOR coefficients of lag larger than two are close to zero following time around 0.3 (where the time axis has been normalized such that $t \in (0, 1)$). Thus, the last four TVAR coefficients after time 0.3 are close to zero (Figure 8(c)). Such phenomena suggests that the period before time 0.3 has a more complex dependence structure. Figure 8(d) illustrates that the innovation variance exhibits higher volatility at the beginning signal. These bursts in the innovation variance are related to the series of broadband clicks at the beginning of the signal. Finally, Figure 7(b) presents the time–frequency representation of the treehopper signal using the TVAR(6) model and corroborates the significance of the broadband clicks at the beginning of the signal. Figure 7(c) shows the posterior standard deviation of the time–frequency representation using 2,000 MCMC samples (in order to obtain a posterior distribution of the time–frequency surface for uncertainty quantification) from 210,000 MCMC iterations by discarding the first 10,000 iterations as burn-in and keeping every 100th iteration of the remainder.

In general, goodness-of-fit is an important component of model assessment. Although the fitted model produced residuals that are normally distributed, as determined

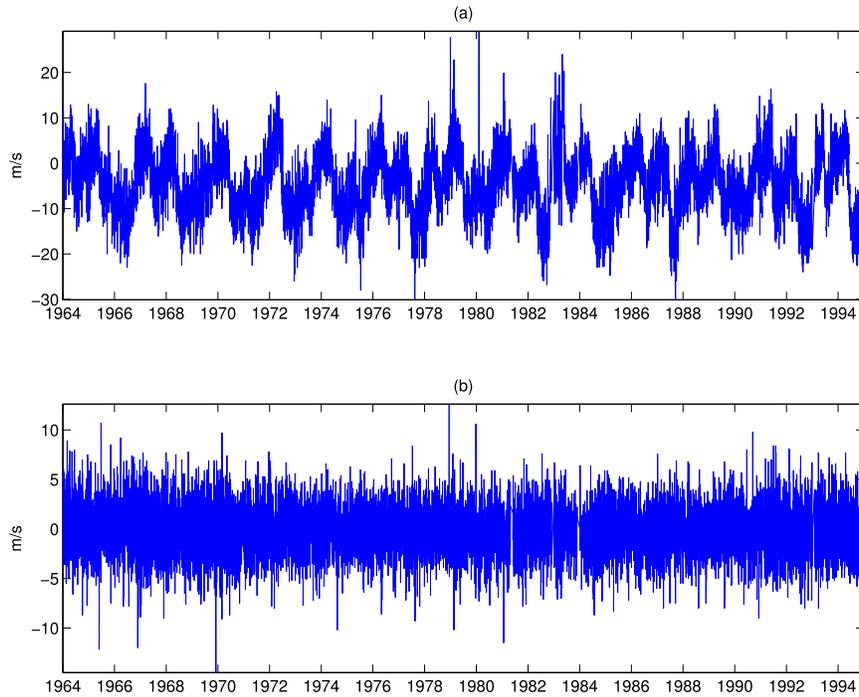


Figure 9: Plots (a) and (b) show daily time series (1964–1994) of east/west and north/south components of wind, respectively. Both components are measured in meters per second (m/s).

through visual inspection of the histogram, QQ plot, and Shapiro–Wilk test (results not shown), they do not behave like white-noise. Additionally, increasing the model order did not remedy this issue. Nevertheless, it is important to note that the estimated time–frequency representation presented for the animal communication signal corroborates previous analyses (e.g., Holan et al. (2010)) and subject matter expertise.

4.2 Wind Components

We study the time–frequency representations of the east/west and north/south wind components, recorded daily at Chuuk Island in the tropical Pacific during the period of 1964 to 1994 (see Cressie and Wikle, 2011, Sections 3.5.3 and 3.5.4). The data studied are at the level of 70 hPa, which is important scientifically due to the likely presence of westward and eastward propagating tropical waves, and the presence of the quasi-biennial oscillation (QBO) (Wikle et al., 1997). Figure 9 shows the two wind component time series from which we can discern visually that the east/west series clearly exhibits the QBO signal, but no discernible smaller-scale oscillations are present in either series. Our interest is then whether the time-varying spectra for these series suggest the presence of time-varying oscillations, which are theorized to be present.

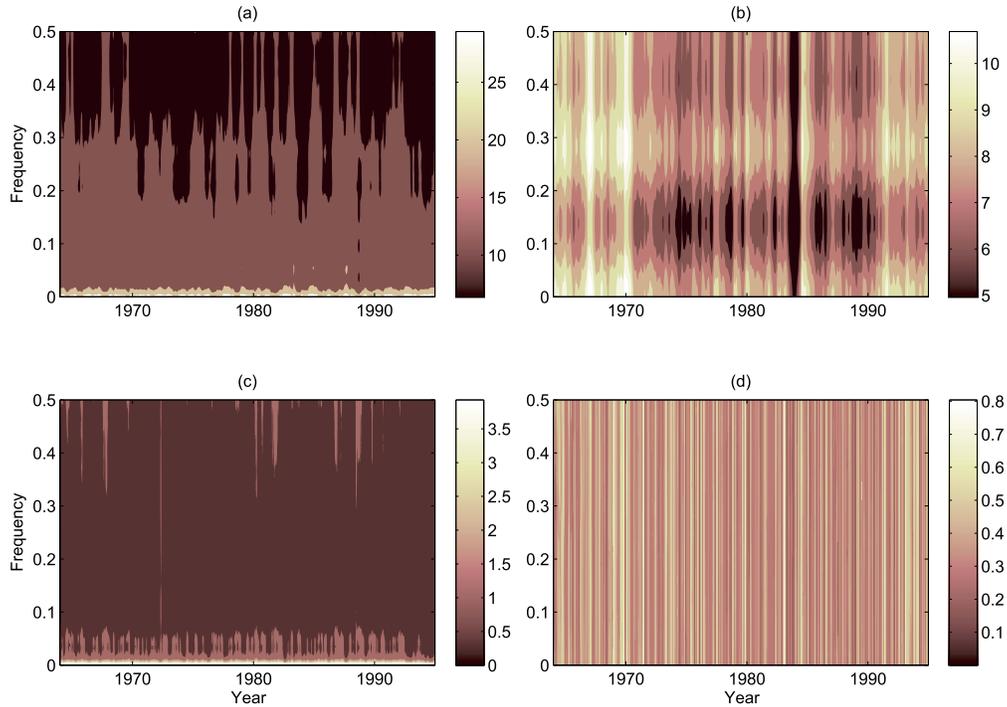


Figure 10: Plots (a) and (c) display the posterior mean and standard deviation of time–frequency representations of the wind east/west component by fitting a TVAR(18) model. Plots (b) and (d) display the posterior mean and standard deviation of time–frequency representations of the wind north/south component by fitting a TVAR(4) model.

We consider the same search space for the model order and discount factors as that used for the treehopper communication signal (Section 4.1). The BLF-sctree plot for the east/west component shows an increase of \mathcal{L}_m along with the order. Therefore, we use (15) with $\tau = 0.5$ along with goodness-of-fit testing results (i.e., white-noise residuals) and choose the order equal to 18. Additionally, the lag-one PARCOR coefficient is time dependent whereas the remaining 17 PARCOR coefficients appear to be constant. The innovation variances of TVAR(18) model are also time dependent. On the other hand, the BLF-sctree plot for the north/south component shows a turning point at order four so that we choose the order equal to four. The PARCOR coefficients are time independent but the innovation variances of the TVAR(4) model are time dependent. The preceding findings are illustrated in Figures 6 and 7 of the Supplemental Appendix. Figures 10(a) and 10(b) show the estimated time–frequency representations of the wind east/west and north/south components. Figures 10(c) and 10(d) present the associated posterior standard deviations of the time–frequency representation using 2,000 MCMC samples (in order to obtain a posterior distribution of the time–frequency surface for uncertainty quantification) from 210,000

MCMC iterations by discarding the first 10,000 iterations as burn-in and keeping every 100th iteration of the remainder. Goodness-of-fit for the models presented is assessed through white-noise residuals, as determined through acf plots and the Ljung–Box test (not shown). Further, through visual inspection of the histogram, QQ plot, and Shapiro–Wilk test all of the residuals are seen to be normally distributed (not shown).

The east/west component time-varying spectrum does suggest that the QBO intensity varies considerably as evidenced by the power in the low-frequencies. Perhaps more interesting is the suggestion of time-varying equatorial waves in the north/south wind component time-varying spectrum. In particular, the lower-frequency (Kelvin and Rossby) waves with frequencies between 0.1 and 0.2 show considerable variation in duration of wave activity, as well as intensity. One also sees time-variation in the likely mixed-Rossby gravity waves in the frequency band between 0.2 and 0.35. Interestingly, in some cases these are in phase with the lower-frequency wave activity but more often act in opposition. We also note the almost complete collapse of the equatorial wave activity centered on 1984.

4.3 Economic Index

Koopman and Wong (2011) study the time–frequency representation of the log-difference of the US gross domestic product (GDP), consumption, and investment from the first quarter of 1947 to the first quarter of 2010. For comparison, we obtained the series from the Federal Reserve Bank of St. Louis (<http://research.stlouisfed.org/fred2>). Since the investment series is unavailable from the website, we only consider the GDP and consumption in our study. Figure 11(a) illustrates the trend of the logarithm of the GDP and consumption series. Similar to Koopman and Wong (2011), we analyzed the log-difference of these two series (Figures 11(b) and 11(c)).

We consider the same search space for the model order and discount factors as that used for the treehopper communication signal (Section 4.1). By using (15) with $\tau = 0.5$, we choose TVAR(1) for the log-difference of GDP and TVAR(2) for the log-difference of consumption. Goodness-of-fit for the models presented is assessed through white-noise residuals, as determined through acf plots and the Ljung–Box test (not shown). In addition, through visual inspection of the histogram, QQ plot, and Shapiro–Wilk test all of the residuals are normally distributed (not shown). Figures 12(a) and 12(b) show the estimated time–frequency representations and the associated standard deviations of the log-difference of GDP and consumption series, respectively. Similar to the results of Koopman and Wong (2011), both the US macroeconomic indices exhibit relatively larger spectra (or fluctuations) for the early period around 1950. Then, due to the oil crisis, another fluctuation comes out in the period of 1970–1980. The fluctuation at the end of the sample reflects the 2008–2009 worldwide financial crisis. See Koopman and Wong (2011) for further discussion. Using the same number of MCMC iterations as Section 4.1 (in order to obtain a posterior distribution of the time–frequency surface for uncertainty quantification), Figures 12(c) and 12(d) present the associated posterior standard deviations.

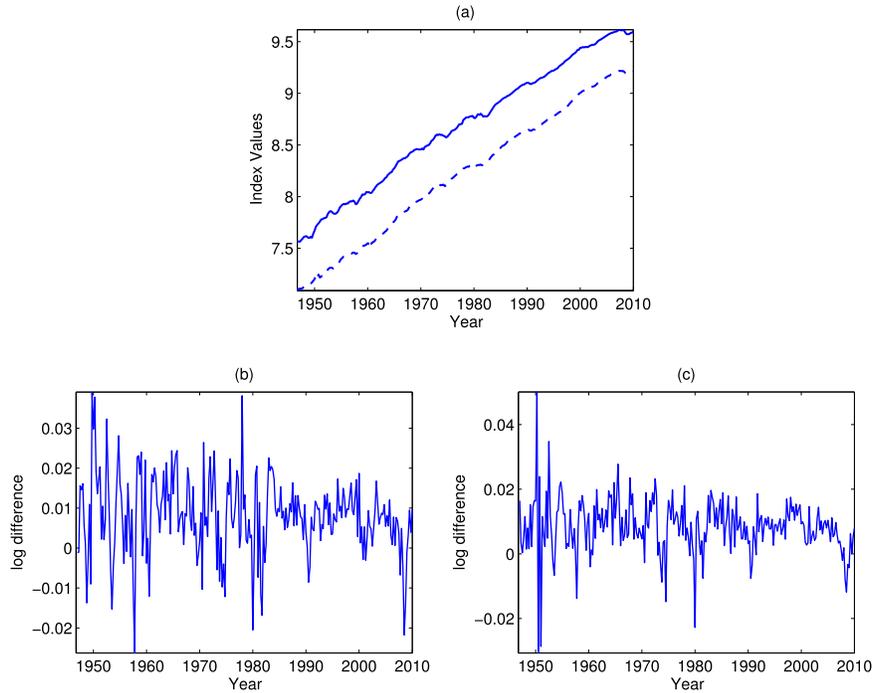


Figure 11: Plot (a) displays the logarithm of GDP (solid line) and consumption (dashed line) in from the first quarter of 1947 to the first quarter of 2010. Plots (b) and (c) present the log-difference of GDP and the log-difference of consumption.

5 Discussion

This paper develops a computationally efficient method for model-based time–frequency analysis. Specifically, we consider a fully Bayesian lattice filter approach to estimating time-varying autoregressions. By taking advantage of the partial autocorrelation domain, our approach is extremely stable. That is, the PARCOR coefficients and the TVAR innovation variances are specified within the lattice structure and then estimated simultaneously. Notably, the marginal distributions arising from our approach are all of standard form and, thus, facilitate easy estimation.

The framework we propose extends the current model-based approaches to time–frequency analysis and, in most cases, provides superior performance, as measured by the average squared error between the true and estimated time-varying spectral density. In fact, for slowly-varying processes we have demonstrated significant estimation improvements from using our approach. In contrast, when the true process comes from a piecewise AR model the approach of Davis et al. (2006) performed best, with our approach a close competitor and performing second best. This is not unexpected as the AutoPARM method is a model-based segmented approach and more closely mimics the behavior of a piecewise AR.

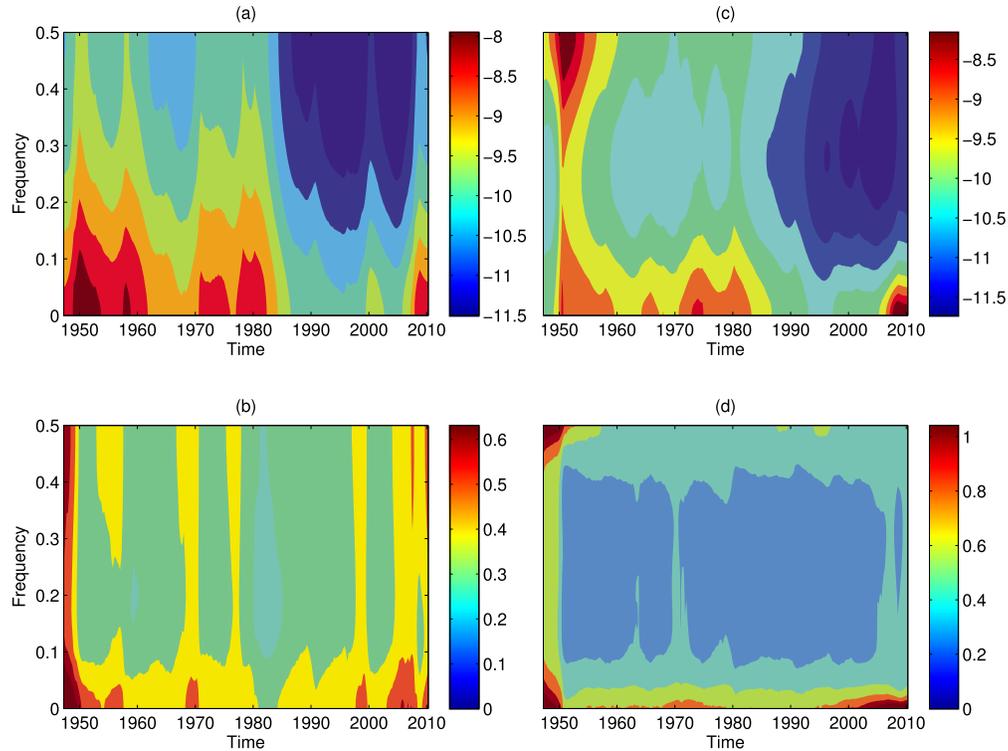


Figure 12: Plots (a) and (b) display the time–frequency representations of the log-difference of GDP and consumption series, respectively. Plots (c) and (d) show the standard deviations of the time–frequency representations of the log-difference of GDP and consumption series, respectively.

In addition to a comprehensive simulation study we have provided three real-data examples, one from animal (insect) communication, one from environmental science, and one from macro-economics. In all cases, the exceptional time–frequency resolution obtained using our approach helps identify salient features in the time–frequency surface. This is especially important when the primary interest is in producing a time–frequency representation to use as a functional predictor in a hierarchical model (e.g., see Holan et al., 2010, 2012, among others). Finally, as a by-product of taking a fully Bayesian approach, we are naturally able to quantify uncertainty and, thus, use our approach to draw inference.

Supplementary Material

Supplemental Materials: Bayesian Lattice Filters for Time-Varying Autoregression and Time–Frequency Analysis (DOI: [10.1214/15-BA978SUPP](https://doi.org/10.1214/15-BA978SUPP); .pdf). Supplementary Mate-

rials: Details surrounding the estimation algorithms, computation times associated with the applications in Section 4, and additional figures for Section 4

References

- Cocroft, R. B. and McNett, G. D. (2006). “Vibrational communication in treehoppers (Hemiptera: Membracidae).” In: *Insect Sounds and Communication: Physiology, Ecology and Evolution*, Drosopoulos, S. and Claridge, M. F. (Eds.), 305–317. Taylor & Francis. 994
- Cressie, N. and Wikle, C. K. (2011). *Statistics for Spatio-Temporal Data*. John Wiley & Sons. MR2848400. 996
- Davis, R. A., Lee, T. C. M., and Rodriguez-Yam, G. A. (2006). “Structural break estimation for nonstationary time series models.” *Journal of the American Statistical Association*, 101(473): 223–239. MR2268041. doi: <http://dx.doi.org/10.1198/016214505000000745>. 978, 986, 988, 991, 993, 999
- Everitt, R. G., Andrieu, C., and Davy, M. (2013). “Online Bayesian inference in some time–frequency representations of non-stationary processes.” *IEEE Transactions of Signal Processing*, 61: 5755–5766. MR3130040. doi: <http://dx.doi.org/10.1109/TSP.2013.2280128>. 978, 986
- Feichtinger, H. G. and Strohmer, T. (1998). *Gabor Analysis and Algorithms: Theory and Applications*. Birkhäuser. MR1601119. doi: <http://dx.doi.org/10.1007/978-1-4612-2016-9>. 978
- Fitzgerald, W. J., Smith, R. L., Walden, A. T., and Young, P. C. (2000). *Nonlinear and Nonstationary Signal Processing*. Cambridge University Press. MR1831319. 978
- Godsill, S. J., Doucet, A., and West, M. (2004). “Monte Carlo smoothing for nonlinear time series.” *Journal of the American Statistical Association*, 99(465): 156–168. MR2054295. doi: <http://dx.doi.org/10.1198/016214504000000151>. 979, 982
- Gröchenig, K. (2001). *Foundations of Time–Frequency Analysis*. Birkhäuser. MR1843717. doi: <http://dx.doi.org/10.1007/978-1-4612-0003-1>. 978
- Hayes, M. H. (1996). *Statistical Digital Signal Processing and Modeling*. John Wiley & Sons. 981, 982
- Haykin, S. (2002). *Adaptive Filter Theory (4th ed)*. Pearson. 982
- Holan, S. H., Wikle, C. K., Sullivan-Beckers, L. E., and Cocroft, R. B. (2010). “Modeling complex phenotypes: generalized linear models using spectrogram predictors of animal communication signals.” *Biometrics*, 66(3): 914–924. MR2758228. doi: <http://dx.doi.org/10.1111/j.1541-0420.2009.01331.x>. 977, 987, 994, 996, 1000
- Holan, S. H., Yang, W.-H., Matteson, D. S., and Wikle, C. K. (2012). “An approach for identifying and predicting economic recessions in real-time using time–frequency functional models.” *Applied Stochastic Models in Business and Industry*, 28: 485–499. MR3008939. doi: <http://dx.doi.org/10.1002/asmb.1954>. 977, 978, 1000

- Kitagawa, G. (1988). “Numerical approach to non-Gaussian smoothing and its applications.” In: *Proceedings of the 20th Symposium on the Interface*, Wegman, E. J., Miller, J., and Gantz, D. T. (Eds.). Interface Foundation of North America. 979, 980, 982
- Kitagawa, G. (2010). *Introduction to Time Series Modeling*. Chapman & Hall/CRC. MR2647741. doi: <http://dx.doi.org/10.1201/9781584889229-c13>. 979, 981
- Kitagawa, G. and Gersch, W. (1996). *Smoothness Priors Analysis of Time Series*. Springer. MR1441074. doi: <http://dx.doi.org/10.1007/978-1-4612-0761-0>. 978, 979, 980, 982, 985, 986
- Koopman, S. and Wong, S. (2011). “Kalman filtering and smoothing for model-based signal extraction that depend on time-varying spectra.” *Journal of Forecasting*, 30: 147–167. MR2758808. doi: <http://dx.doi.org/10.1002/for.1203>. 998
- Mallat, S. G. (2008). *A Wavelet Tour of Signal Processing: The Sparse Way*. Academic Press. MR2479996. 978
- Martinez, J. G., Bohn, K. M., Carroll, R. J., and Morris, J. S. (2013). “A study of Mexican free-tailed bat chirp syllables: Bayesian functional mixed models for nonstationary acoustic time series.” *Journal of the American Statistical Association*, 108(502): 514–526. MR3174638. doi: <http://dx.doi.org/10.1080/01621459.2013.793118>. 977, 978
- Ombao, H., Raz, J., Von Sachs, R., and Malow, B. (2001). “Automatic statistical analysis of bivariate nonstationary time series.” *Journal of the American Statistical Association*, 96(454): 543–560. MR1946424. doi: <http://dx.doi.org/10.1198/016214501753168244>. 978, 986, 987
- Oppenheim, A. V. and Schaffer, R. W. (2009). *Discrete-Time Signal Processing*. Prentice Hall Signal Processing. 978
- Percival, D. B. and Walden, A. T. (2000). *Wavelet Methods for Time Series Analysis*. Cambridge University Press. MR1770693. doi: <http://dx.doi.org/10.1017/CB09780511841040>. 978
- Prado, R. and Huerta, G. (2002). “Time-varying autogressions with model order uncertainty.” *Journal of Time Series*, 23(5): 599–618. MR1925266. doi: <http://dx.doi.org/10.1111/1467-9892.00280>. 979
- Prado, R. and West, M. (2010). *Time Series: Modeling, Computation, and Inference*. CRC Press. MR2655202. 980
- Rencher, A. C. (2002). *Methods of Multivariate Analysis (2nd ed)*. John Wiley & Son. MR1885894. doi: <http://dx.doi.org/10.1002/0471271357>. 985
- Rosen, O., Stoffer, D. S., and Wood, S. (2009). “Local spectral analysis via a Bayesian mixture of smoothing splines.” *Journal of the American Statistical Association*, 104(485): 249–262. MR2504376. doi: <http://dx.doi.org/10.1198/jasa.2009.0118>. 978, 987, 988, 991

- Rosen, O., Wood, S., and Stoffer, D. S. (2012). “AdaptSPEC: Adaptive spectral estimation for nonstationary time series.” *Journal of the American Statistical Association*, 107(500): 1575–1589. MR3036417. doi: <http://dx.doi.org/10.1080/01621459.2012.716340>. 977, 978, 986, 988, 991
- Shumway, R. H. and Stoffer, D. S. (2006). *Time Series Analysis and Its Applications (2nd ed)*. Springer. MR2228626. 981
- Vidakovic, B. (1999). *Statistical Modeling by Wavelets*. John Wiley & Son. MR1681904. doi: <http://dx.doi.org/10.1002/9780470317020>. 978
- West, M. and Harrison, J. (1997). *Bayesian Forecasting and Dynamic Models (2nd ed)*. Springer. MR1482232. 980, 983, 984
- West, M., Prado, R., and Krystal, A. D. (1999). “Evaluation and comparison of EEG traces: Latent structure in nonstationary time series.” *Journal of the American Statistical Association*, 94(446): 375–387. 979, 983, 984, 986
- Wikle, C. K., Madden, R. A., and Chen, T.-C. (1997). “Seasonal variation of upper tropospheric lower stratospheric equatorial waves over the tropical Pacific.” *Journal of the Atmospheric Sciences*, 54: 1895–1909. 996
- Wolfe, P. J., Godsill, S. J., and Ng, W. J. (2004). “Bayesian variable selection and regularization for time–frequency surface estimation.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3): 575–589. MR2088291. doi: <http://dx.doi.org/10.1111/j.1467-9868.2004.02052.x>. 978
- Wood, S., Rosen, O., and Kohn, R. (2011). “Bayesian mixtures of autoregressive models.” *Journal of Computational and Graphical Statistics*, 20(1): 174–195. MR2816544. doi: <http://dx.doi.org/10.1198/jcgs.2010.09174>. 978
- Yang, W.-H., Holan, S. H., and Wikle, C. K. (2015). “Supplemental materials: Bayesian lattice filters for time-varying autoregression and time–frequency analysis.” *Bayesian Analysis*. doi: <http://dx.doi.org/10.1214/15-BA978SUPP>. 981
- Yang, W.-H., Wikle, C. K., Holan, S. H., and Wildhaber, M. L. (2013). “Ecological prediction with nonlinear multivariate time–frequency functional data models.” *Journal of Agricultural, Biological, and Environmental Statistics*, 18(3): 450–474. MR3110902. doi: <http://dx.doi.org/10.1007/s13253-013-0142-1>. 977, 978

Acknowledgments

This research was partially supported by the U.S. National Science Foundation (NSF) and the U.S. Census Bureau under NSF grant SES-1132031, funded through the NSF-Census Research Network (NCRN) program. The authors would like to thank the Reginald Cocroft lab for use of the insect communication data analyzed in Section 4.1. The authors would like to thank Drs. Richard Davis, Hernando C. Ombao, and Ori Rosen for providing their codes. We also thank Dr. Mike West for generously providing his code online. Finally, we thank the Editor, Associate Editor, and an anonymous Referee for providing valuable comments that helped strengthen this manuscript.