BAYESIAN TIME SERIES

A (hugely selective) introductory overview

- contacting current research frontiers -

Mike West

Institute of Statistics & Decision Sciences

Duke University

June 5th 2002, Valencia VII - Tenerife

Topics

Dynamic linear models (state space models)

- Sequential context, Bayesian framework
- Standard classes of models, model decompositions

Models and methods in physical science applications

- Time series decompositions, latent structure
- Neurophysiology - climatology - speech processing

Multivariate time series:

- Financial applications - Latent structure, volatility models

Simulation-Based Computation

- MCMC - Sequential simulation methodology

Standard Dynamic Models

Dynamic Linear Models **Linear State Space Models**

$$y_t = x_t + \nu_t \qquad x_t = \mathbf{F}_t'\boldsymbol{\theta}_t \qquad \boldsymbol{\theta}_t = \mathbf{G}_t\boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t$$

- signal x_t, state vector $\boldsymbol{\theta}_t = (\theta_{t1}, \ldots, \theta_{td})'$
- regression vector $\mathbf{F}_t$ and state matrix $\mathbf{G}_t$
- zero mean measurement errors ν_t and state innovations $\boldsymbol{\omega}_t$
  - often zero-mean and normally distributed

Examples

“Slowly varying” level observed with noise:

$$y_t = x_t + \nu_t \qquad\qquad x_t = x_{t-1} + \omega_t$$

Dynamic linear regression:

$$y_t = x_t + \nu_t \qquad\qquad x_t = \mathbf{F}_t'\boldsymbol{\theta}_t \qquad\qquad \boldsymbol{\theta}_t = \boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t$$

Error models for $\nu_t, \boldsymbol{\omega}_t$

- normal distributions
- mixtures of normals: outliers and abrupt “structural” changes

Simple regression example

$$y_t = x_t + \nu_t \qquad\qquad x_t = a_t + b_t X_t$$

$\mathbf{F}_t = (1, X_t)'$ and $\boldsymbol{\theta}_t = (a_t, b_t)'$ "wanders" through time

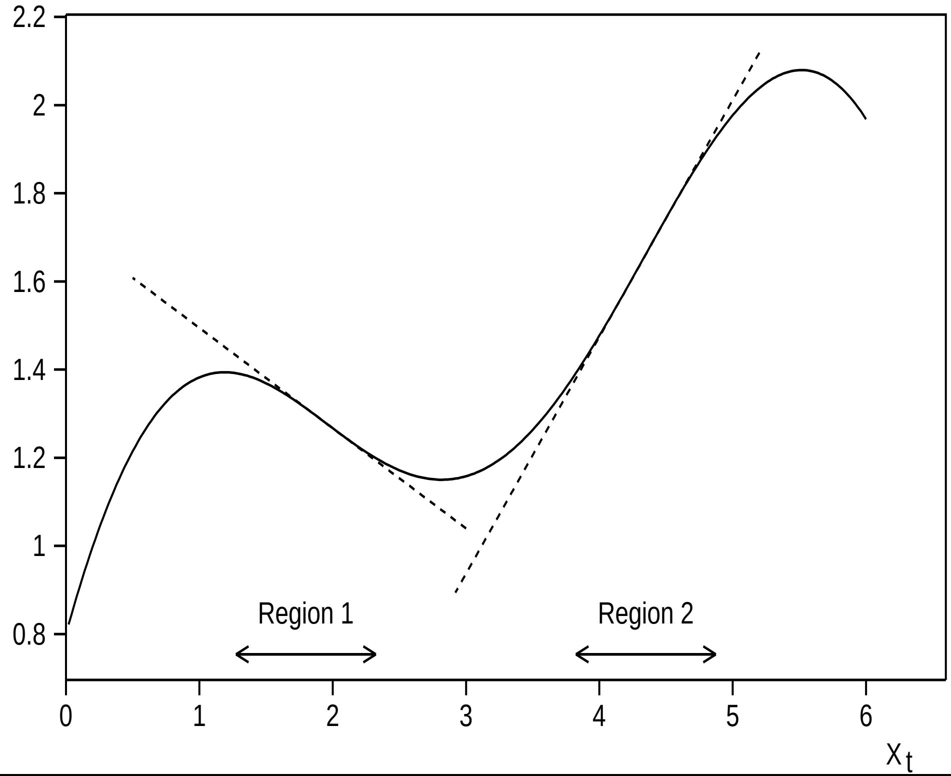

Sales Data Example

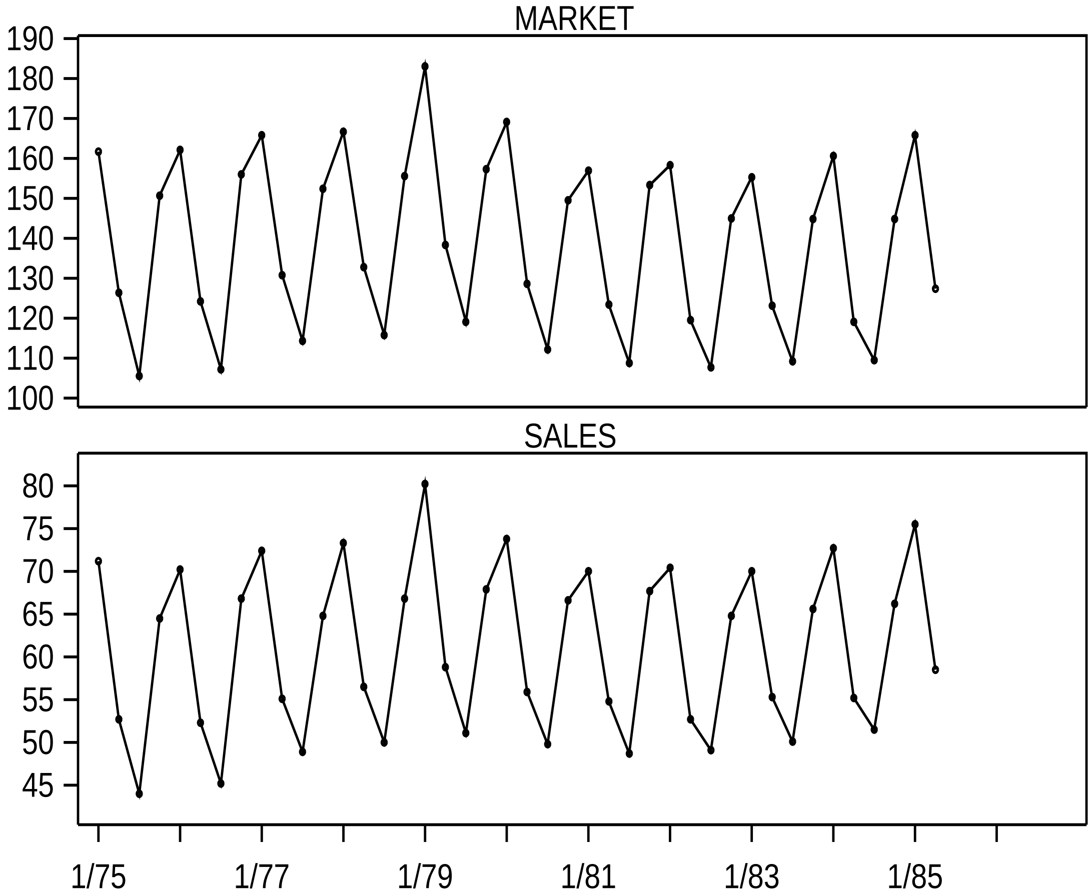

Sales Data Example

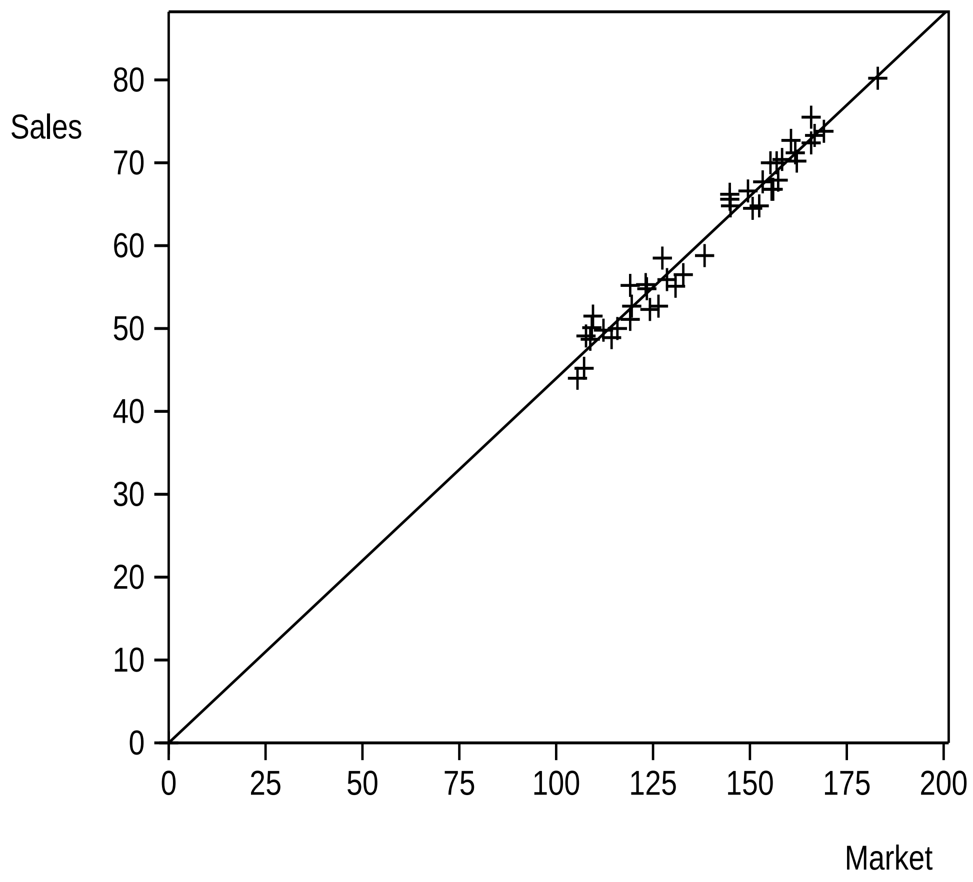

Sales Data Example

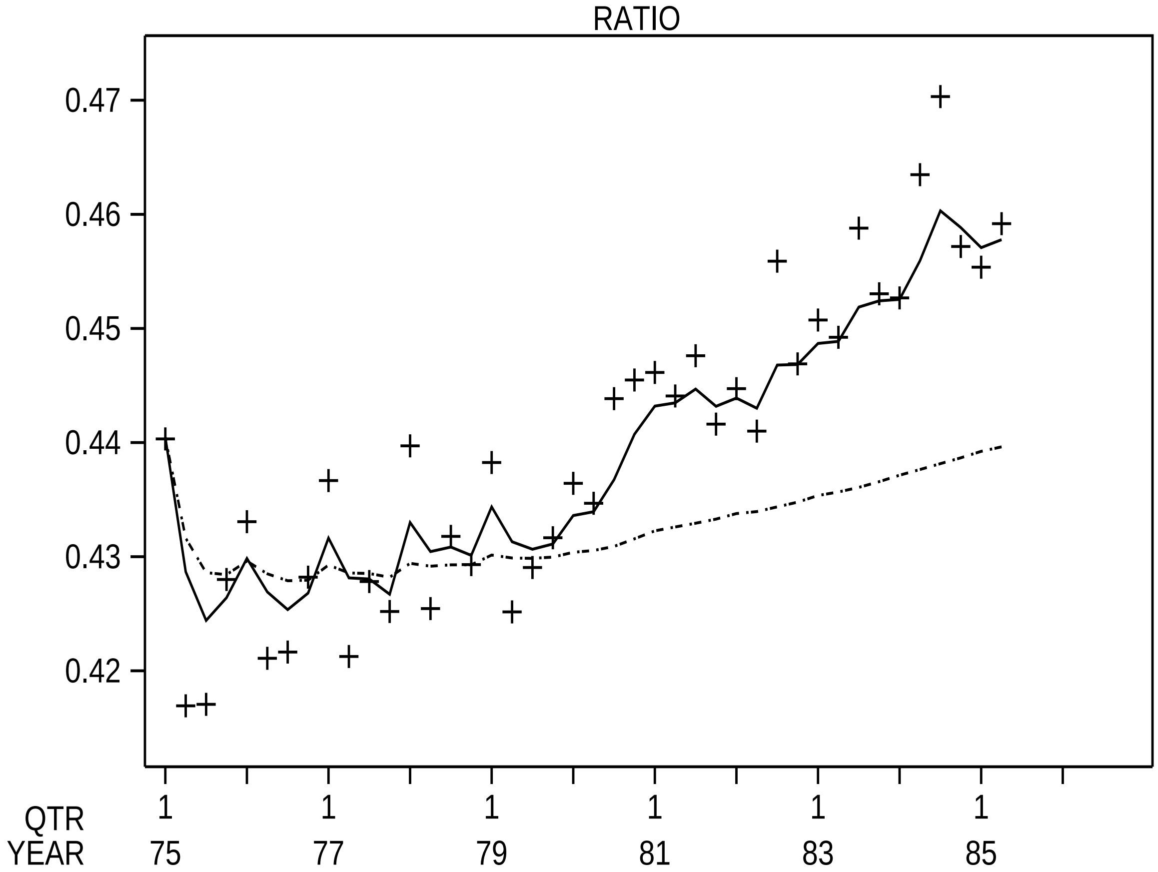

Simple Regression Example

Relative to "static" model, dynamic regression delivers:

- improved estimation via adaptation for "local" regression parameters
- and increased (honest) uncertainty about regression parameters
- adaptability to (small) changes $\rightarrow$ improved point forecasts
- partitions variation: *parameter* vs *observation error*
  $\rightarrow$ increased precision of stated forecasts

  i.e., improved prediction: point forecasts AND precision

General Dynamic Linear Model

$$y_t = x_t + \nu_t \qquad x_t = \mathbf{F}_t' \boldsymbol{\theta}_t \qquad \boldsymbol{\theta}_t = \mathbf{G}_t \boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t$$

$$\begin{array}{ccccccc}
 & y_{t-1} & & y_t & & y_{t+1} & \\
 & \uparrow & & \uparrow & & \uparrow & \\
 & x_{t-1} & & x_t & & x_{t+1} & \\
 & \uparrow & & \uparrow & & \uparrow & \\
\longrightarrow & \boldsymbol{\theta}_{t-1} & \longrightarrow & \boldsymbol{\theta}_t & \longrightarrow & \boldsymbol{\theta}_{t+1} & \longrightarrow
\end{array}$$

- Sequential model definition : Markov evolution structure
- CI structure : $\boldsymbol{\theta}_t$ sufficient for “future” at time t

Bayesian Forecasting

Key concepts:

- Bayesian: modelling & learning is probabilistic
- Time-varying parameter models: often non-stationary
- Sequential view, sequential model definitions
  - encourages interaction, intervention

Statistical framework:

- Forecasting: "What might happen?" and "What if?"
- Data processing and statistical learning from observations
- Updating of models and probabilistic summaries of belief
- Time series analysis ... *Retrospection:* "What happened?"

Bayesian Machinery

- ***Inferences*** based on information $D_t = \{(y_1, \ldots, y_t), I_t\}$

  Find and summarise

  - $p(\boldsymbol{\theta}_t|D_t)$ and $p(\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_t|D_t)$
  - and update as $t \rightarrow t+1 \rightarrow \cdots$

- ***Forecasts:***
  - $p(y_{t+1}, \ldots, y_{t+k}|D_t)$
  - and update as $t \rightarrow t+1 \rightarrow \cdots$

- ***Implementation & computations:***
  - Linear/normal models: neat theory, Kalman filtering
  - Extend to infer variance components, non-normal errors ..
    - ∗ need approximations, simulation methods, MCMC

Commercial Applications

- Short term forecasting of consumer sales and demand
- Monitoring: stocks and inventories of consumer products
- Many items or sectors: Aggregation and multi-level models

Standard models for commercial applications:

$$\begin{array}{ccccccccc} \text{Data} & = & \text{Trend} & + & \text{Seasonal} & + & \text{Regression} & + & \text{Error} \\ \uparrow & & \uparrow & & \uparrow & & \uparrow & & \uparrow \\ y_t & = & x_{1t} & + & x_{2t} & + & x_{3t} & + & \nu_t \end{array}$$

Models in Commercial Applications

Component signals x_{jt} follow individual dynamic models

- *Trends:* e.g., "locally linear" trend

$$x_{1t} = x_{1,t-1} + \beta_t + \partial x_{1t}, \qquad \beta_t = \beta_{t-1} + \partial \beta_t$$

- *Seasonals:*

$$x_{2t} = \sum_j (a_{j,t} \cos(2 * \pi j/p) + b_{j,t} \sin(2 * \pi j/p))$$

where $(a_{j,t}, b_{j,t})$ wander through time

Key concept: Model (De)Composition

- Modelling, prior specification, interventions: **component-wise**
- Posterior inference: detrending, deseasonalisation, etc

Special Classes of DLMs

Time Series DLMs – constant $\mathbf{F}, \mathbf{G}$

$$y_t = x_t + \nu_t \qquad x_t = \mathbf{F}'\boldsymbol{\theta}_t \qquad \boldsymbol{\theta}_t = \mathbf{G}\boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t$$

- Includes all "standard" point-forecasting methods (exponential smoothing, variants, ...)
- Polynomial trend and seasonal components in commercial models
- Includes all practically useful ARIMA models

Multiple representations:

$$\boldsymbol{\phi}_t = \mathbf{E}\boldsymbol{\theta}_t \leftrightarrow \mathbf{G} \rightarrow \mathbf{E}\mathbf{G}\mathbf{E}^{-1}$$

General class: Time-varying $\mathbf{F}_t, \mathbf{G}_t$

Includes non-stationary models, time-varying ARIMA models, etc., ...

Simulation-Based Computation

$$y_t = x_t + \nu_t \qquad x_t = \mathbf{F}_t' \boldsymbol{\theta}_t \qquad \boldsymbol{\theta}_t = \mathbf{G}_t \boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t$$

- Normal error models $p(\nu_t), p(\boldsymbol{\omega}_t)$
- Fixed time window $t = 1, \ldots, n$
- State vector set: $\boldsymbol{\Theta} = \{\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_n\}$
- **Require:** full posterior sample $\boldsymbol{\Theta}^*$ from $p(\boldsymbol{\Theta}|D_n)$

Available via *"Forward-filtering: Backward-sampling"* algorithm

Carter and Kohn (1994) Biometrika

Frühwirth-Schnatter (1994) J Time Series Anal

West and Harrison 1997

FFBS Algorithm

Forward-filtering:

- Standard normal/linear analysis: Kalman filter
- delivers normal $p(\boldsymbol{\theta}_t|D_t)$ at each $t = 1, \ldots, n$

Backward-sampling:

- at $t = n:$ sample $\boldsymbol{\theta}_n^*$ from $p(\boldsymbol{\theta}_n|D_n)$
- for $t = n-1, n-2, \ldots, 1:$ sample $\boldsymbol{\theta}_t^*$ from normal distribution $p(\boldsymbol{\theta}_t|D_t, \boldsymbol{\theta}_{t+1}^*)$

Builds up $\boldsymbol{\Theta}^*$ as $\boldsymbol{\theta}_n^*, \boldsymbol{\theta}_{n-1}^*, \ldots$

Exploits Markovian/CI model structure

MCMC in DLMs

DLM parameters: e.g., constant model

$$y_t = x_t + \nu_t \qquad x_t = \mathbf{F}'\boldsymbol{\theta}_t \qquad \boldsymbol{\theta}_t = \mathbf{G}\boldsymbol{\theta}_{t-1} + \boldsymbol{\omega}_t$$

Parameters:

- Variances (variance matrices) of $\nu_t, \boldsymbol{\omega}_t$
- Elements of $\mathbf{F}, \mathbf{G}$
- Indicators in normal mixture models for errors

Add parameters to analysis: MCMC utilising FFBS

MCMC in DLMs

Parameters $\boldsymbol{\Phi}$ (may depend on sample size)

Gibbs sampling: $p(\boldsymbol{\Theta}, \boldsymbol{\Phi}|D_n)$ iteratively resampled via

- Apply FFBS algorithm to draw $\boldsymbol{\Theta}^*$ from $p(\boldsymbol{\Theta}|\boldsymbol{\Phi}^*, D_n)$
- Draw new value of $\boldsymbol{\Phi}^*$ from $p(\boldsymbol{\Phi}|\boldsymbol{\Theta}^*, D_n)$
- Iterate

"Standard" Gibbs sampling: MCMC

May need "creativity" in sampling $\boldsymbol{\Phi}^*$: Metropolis-Hastings, etc

Often "easy": as in Autoregressive DLM

MCMC: Autoregressive Model Example

Data: $y_t = x_t + \nu_t$

State AR(d) $x_t = \sum_{j=1}^{d} \phi_j x_{t-j} + \epsilon_t$

DLM for for x_t : $x_t = (1, 0, \ldots, 0)\mathbf{x}_t, \quad \mathbf{x}_t = \mathbf{G}\mathbf{x}_{t-1} + \boldsymbol{\omega}_t$

$$\mathbf{G} = \begin{pmatrix} \phi_1 & \phi_2 & \phi_3 & \cdots & \phi_{d-1} & \phi_d \\ 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & & & \ddots & 0 & \vdots \\ 0 & 0 & \cdots & \cdots & 1 & 0 \end{pmatrix}, \quad \boldsymbol{\omega}_t = \begin{pmatrix} \epsilon_t \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

- **Parameters:** $\{\phi_1, \ldots, \phi_d;\ V(\nu_t), V(\epsilon_t)\}$
- Conditional posteriors standard: linear regression parameters

Models in Scientific Applications

Historical interest in biomedical monitoring (change-points), engineering applications (control, tracking), environmental monitoring, ...

Goals:

- Exploratory "discovery" of interpretable latent processes
- Nonstationary time series: "hidden" quasi-periodicities
- Changes over time at different time scales
- Time:frequency structure (in time domain)

State-space models:

- Stationary and/or nonstationary, time-varying parameters
- General decomposition theory for state space-space models
- DLM autoregressions and time-varying autoregressions

Autoregressive DLM

Latent AR(d) process: $x_t = \sum_{j=1}^{d} \phi_j x_{t-j} + \epsilon_t$

Time series: $y_t = x_{0t} + x_t + \nu_t$ with trend, etc in x_{0t}

DLM for for x_t : $x_t = (1, 0, \ldots, 0)\mathbf{x}_t, \qquad \mathbf{x}_t = \mathbf{G}\mathbf{x}_{t-1} + \boldsymbol{\omega}_t$

$$\mathbf{G} = \begin{pmatrix} \phi_1 & \phi_2 & \phi_3 & \cdots & \phi_{d-1} & \phi_d \\ 1 & 0 & 0 & \cdots & 0 & 0 \\ 0 & 1 & 0 & \cdots & 0 & 0 \\ \vdots & & & \ddots & 0 & \vdots \\ 0 & 0 & \cdots & \cdots & 1 & 0 \end{pmatrix}, \qquad \boldsymbol{\omega}_t = \begin{pmatrix} \epsilon_t \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix}$$

– x_t latent, unobserved

– $\mathbf{G} = \mathbf{G}(\boldsymbol{\phi})$ with $\boldsymbol{\phi} = (\phi_1, \ldots, \phi_d)'$ to be estimated

Time Series Decomposition

Eigenstructure: $\mathbf{G} = \mathbf{E}^{-1}\mathbf{A}\mathbf{E}$

Reparametrise to "diagonal" model: $\mathbf{x}_t \rightarrow \mathbf{E}\mathbf{x}_t$

Transforms to

$$x_t = \sum_{j=1}^{d_z} z_{t,j} + \sum_{j=1}^{d_a} a_{t,j}$$

- z terms: one for each pair of complex conjugate eigenvalues
- a terms: one for each real eigenvalue
- underlying *latent* processes $z_{t,j}$ and $a_{t,j}$ follow "simple" models
  - $a_{t,j}$ is AR(1) process - short-term correlations
  - $z_{t,j}$ is *quasi-periodic* ARMA(2,1) - noisy sine wave with randomly time-varying amplitude & phase, fixed frequency

Time-varying Autoregression

TV-AR(d) model: $x_t = \sum_{j=1}^{d} \phi_{t,j} x_{t-j} + \epsilon_t$

- *AR parameter:* $\boldsymbol{\phi}_t = (\phi_{t,1}, \ldots, \phi_{t,d})'$ "wanders" through time:

$$\boldsymbol{\phi}_t = \boldsymbol{\phi}_{t-1} + \partial\boldsymbol{\phi}_t$$

- stochastic "shocks" $\partial\boldsymbol{\phi}_t$
- *Innovations:* $\epsilon_t \sim N(0, \sigma_t^2)$ – time-varying variance σ_t^2

Flexible representations:

- non-stationary process, time-varying spectral properties
- latent component structure
- other evolutions for $\boldsymbol{\phi}_t$ (e.g., Godsill *et al* on speech processing)

DLM Form of TV-AR(d):

$$x_t = (1, 0, \ldots, 0)' \mathbf{x}_t$$

$$\mathbf{x}_t = \mathbf{G}(\boldsymbol{\phi}_t)\mathbf{x}_{t-1} + (\epsilon_t, 0, \ldots, 0)'$$

$$\boldsymbol{\phi}_t = \boldsymbol{\phi}_{t-1} + \partial \boldsymbol{\phi}_t$$

with

$$\mathbf{G}(\boldsymbol{\phi}_t) = \begin{pmatrix} \phi_{t,1} & \phi_{t,2} & \cdots & \phi_{t,d-1} & \phi_{t,d} \\ 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & & \ddots & 0 & \vdots \\ 0 & 0 & \cdots & 1 & 0 \end{pmatrix}$$

Analysis: Posterior distributions for $\{\boldsymbol{\phi}_t, \sigma_t : \quad \forall t\}$

Component models: $y_t = x_{0t} + x_{1t} + \nu_t$

$\rightarrow$ infer latent TV-AR processes too: $\{x_{0t}, x_t : \quad \forall t\}$

TVAR Decomposition

$$x_t = \sum_{j=1}^{d_z} z_{t,j} + \sum_{j=1}^{d_a} a_{t,j}$$

- z terms: one for each pair of complex conjugate eigenvalues
- a terms: one for each real eigenvalue
- underlying *latent* processes:
  - $a_{t,j}$ is **TV-**AR(1) process - short-term correlations
    + time-varying correlation
  - $z_{t,j}$ is **TV-**ARMA(2,1) - noisy sine wave with randomly time-varying amplitude & phase
    + time-varying frequency (2π/wavelength)

Number of complex/real eigenvalues can vary over time too!

Time:Frequency Analysis

- Fit flexible (high order?) AR or TV-AR models
- Estimate latent components and their frequencies, amplitudes over time
- Time domain representation of spectral structure
- Often, some $z_{t,j}$ physically meaningful, some (high frequency) represent noise, model approximation
- $a_{t,j}$ – noise, model approximation and and (possibly) low frequency “trend”

Paleoclimatology Example

- deep ocean cores: relative abundance of δ^{18}O
- δ^{18}O $\downarrow$ as global temperatures $\uparrow$ (smaller ice mass)
- *reverse sign:* higher recent global temperatures
- "well known" periodicities: earth orbital dynamics $\rightarrow$ impact on solar insolation – Milankovitch; Shackleton *et al* since 1976

  eccentricity: 95-120 kyear

  obliquity: 40-42 kyear

  precession: 19-25 kyear *(1 or 2?)*

Oxygen Isotope Data

- Form of time variation in individual cycles ?
- Timing/nature of onset of "ice-age" cycle $\leftrightarrow$ eccentricity component $\sim$ 1000 kyears ago ?
- *Time scale: errors, interpolation, ... measurement, sampling error, etc*

Models: High order TV-AR, $p = 20$, plus smooth trend (outliers?)

Variance components estimated: changing AR parameters

Decomposition: Posterior mean of $x_t, \phi_{t,j}$ at each t

4 dominant quasi-periodic components: order by estimated *amplitude* (innovation variance)

Others: residual structure &/or contaminations

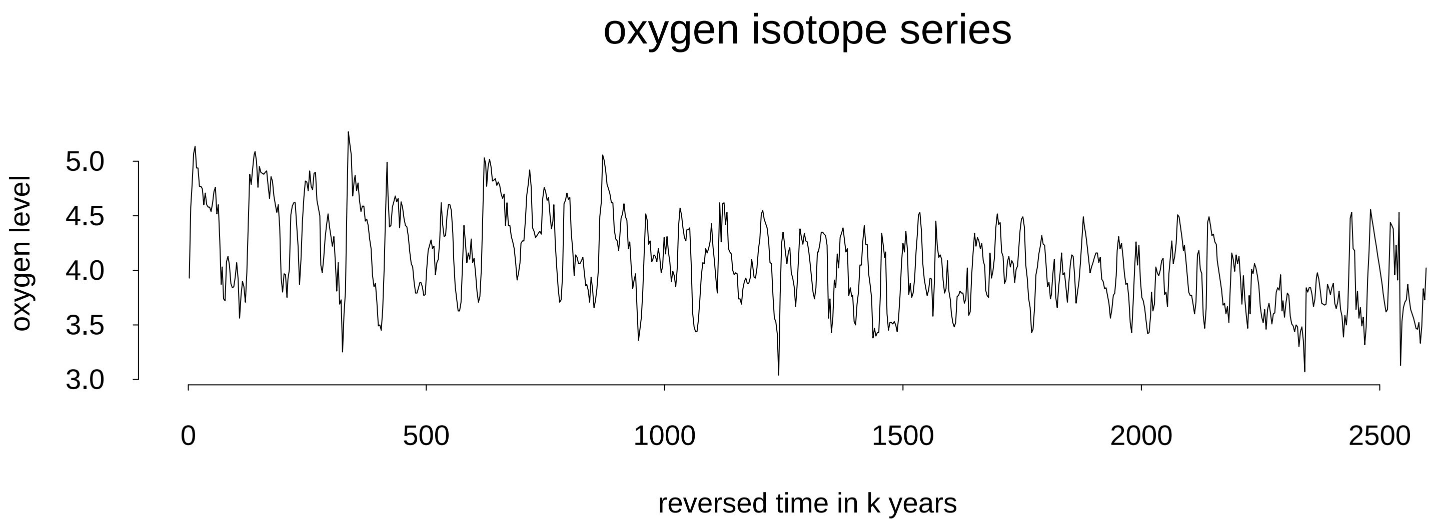
oxygen isotope series
oxygen level
5.0
4.5
4.0
3.5
3.0
0
500
1000
1500
2000
2500
reversed time in k years

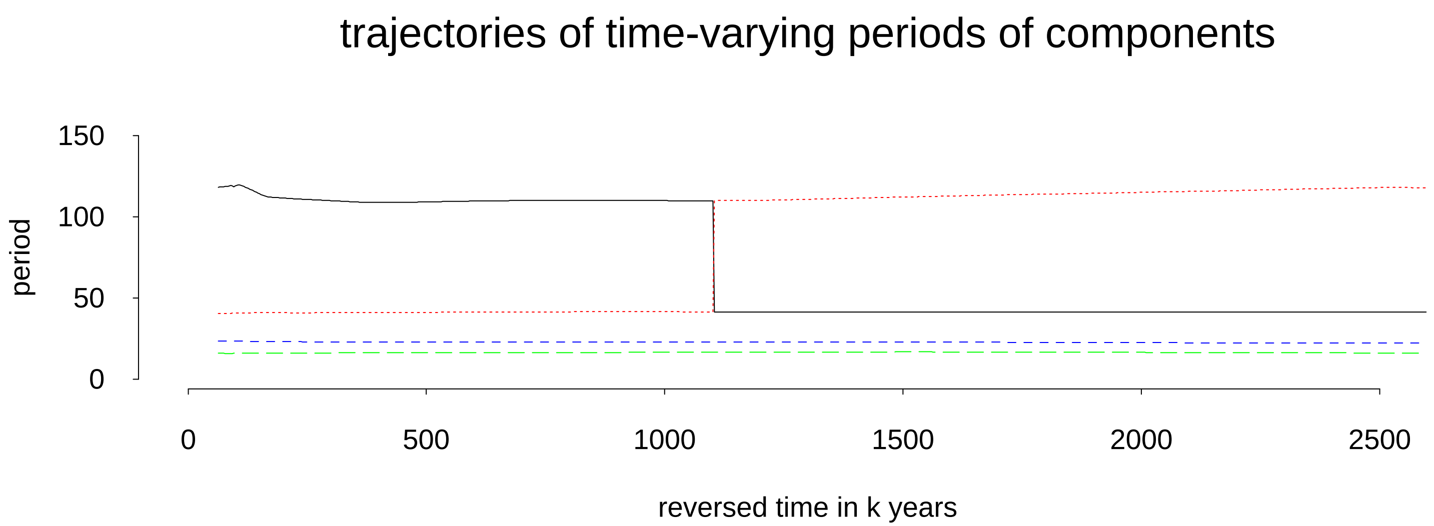
trajectories of time-varying periods of components
period
150
100
50
0
0
500
1000
1500
2000
2500
reversed time in k years

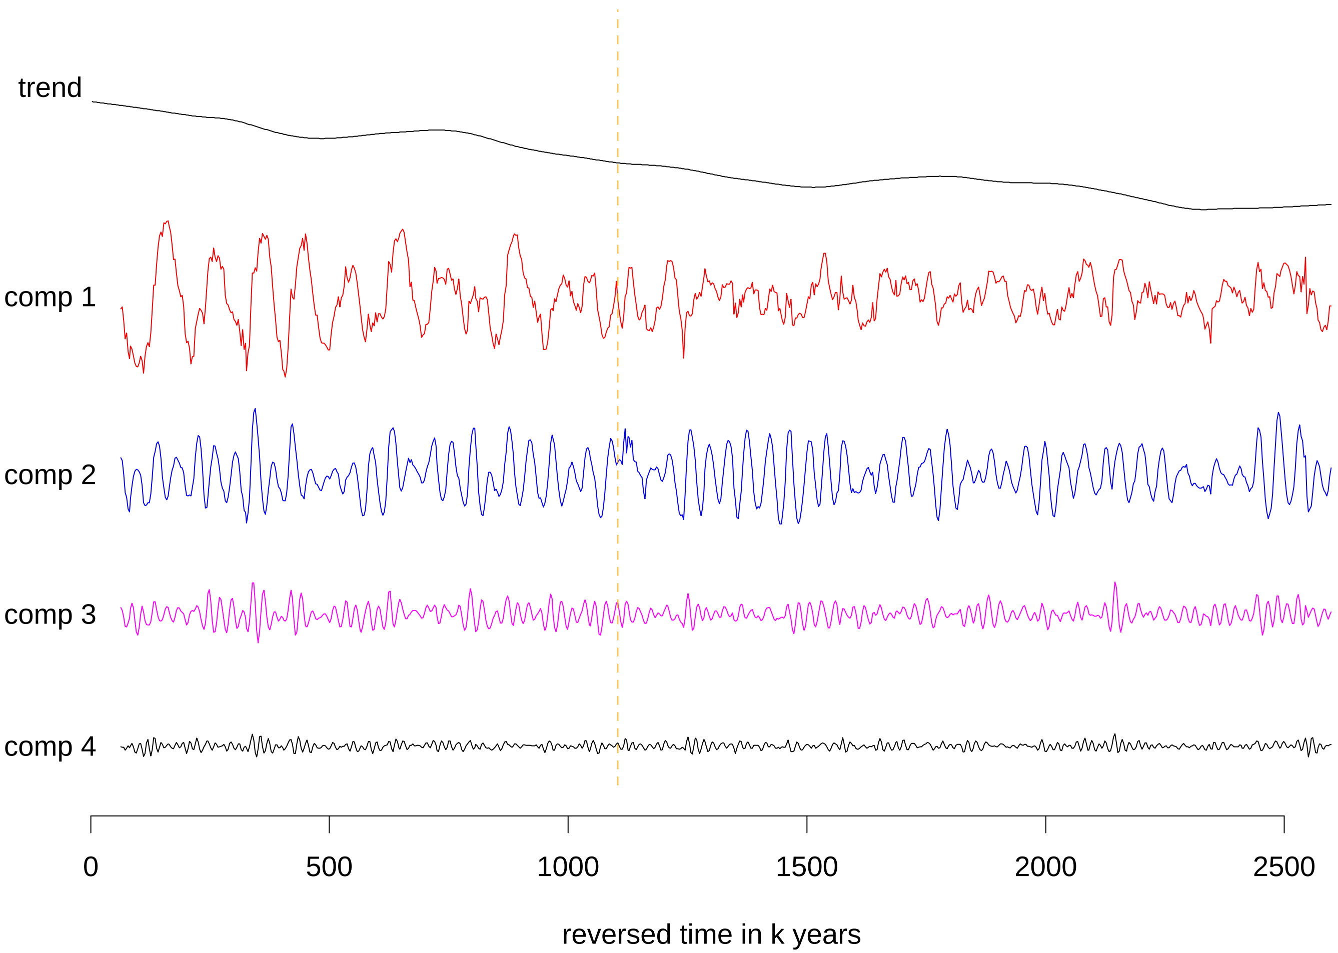

trend
comp 1
comp 2
comp 3
comp 4
0
500
1000
1500
2000
2500
reversed time in k years

Oxygen

- Wavelengths vary only modestly
- Estimated periods/wavelengths consistent with geological determinations

  .... 108–120, peak 110

  .... 40.8–41.6, peak 41.5

  .... 22.2–23, peak 22.8
- "Switch" due to order of estimated amplitude
  - Geological interpretation? Structural climate change $\sim$ 1.1m yrs?

Example: EEG Study

- Clinical uses of electroconvulsive therapy
- Measure seizure treatment outcomes via long, multiple EEG (electroencephalogram) series – electrical potential fluctuations on scalp
- Many multiple series: one seizure – 19 channels, 256/sec

Models to:

- Characterise seizure "waveforms" ... time varying amplitudes at ranges of frequencies (alpha waves, etc)
- Superimposed on "normal" waveform, noise, ...
- Identify/extract latent components: infer **seizure effects**
- Spatial connectivities: related multiple series

Why TVAR Models?

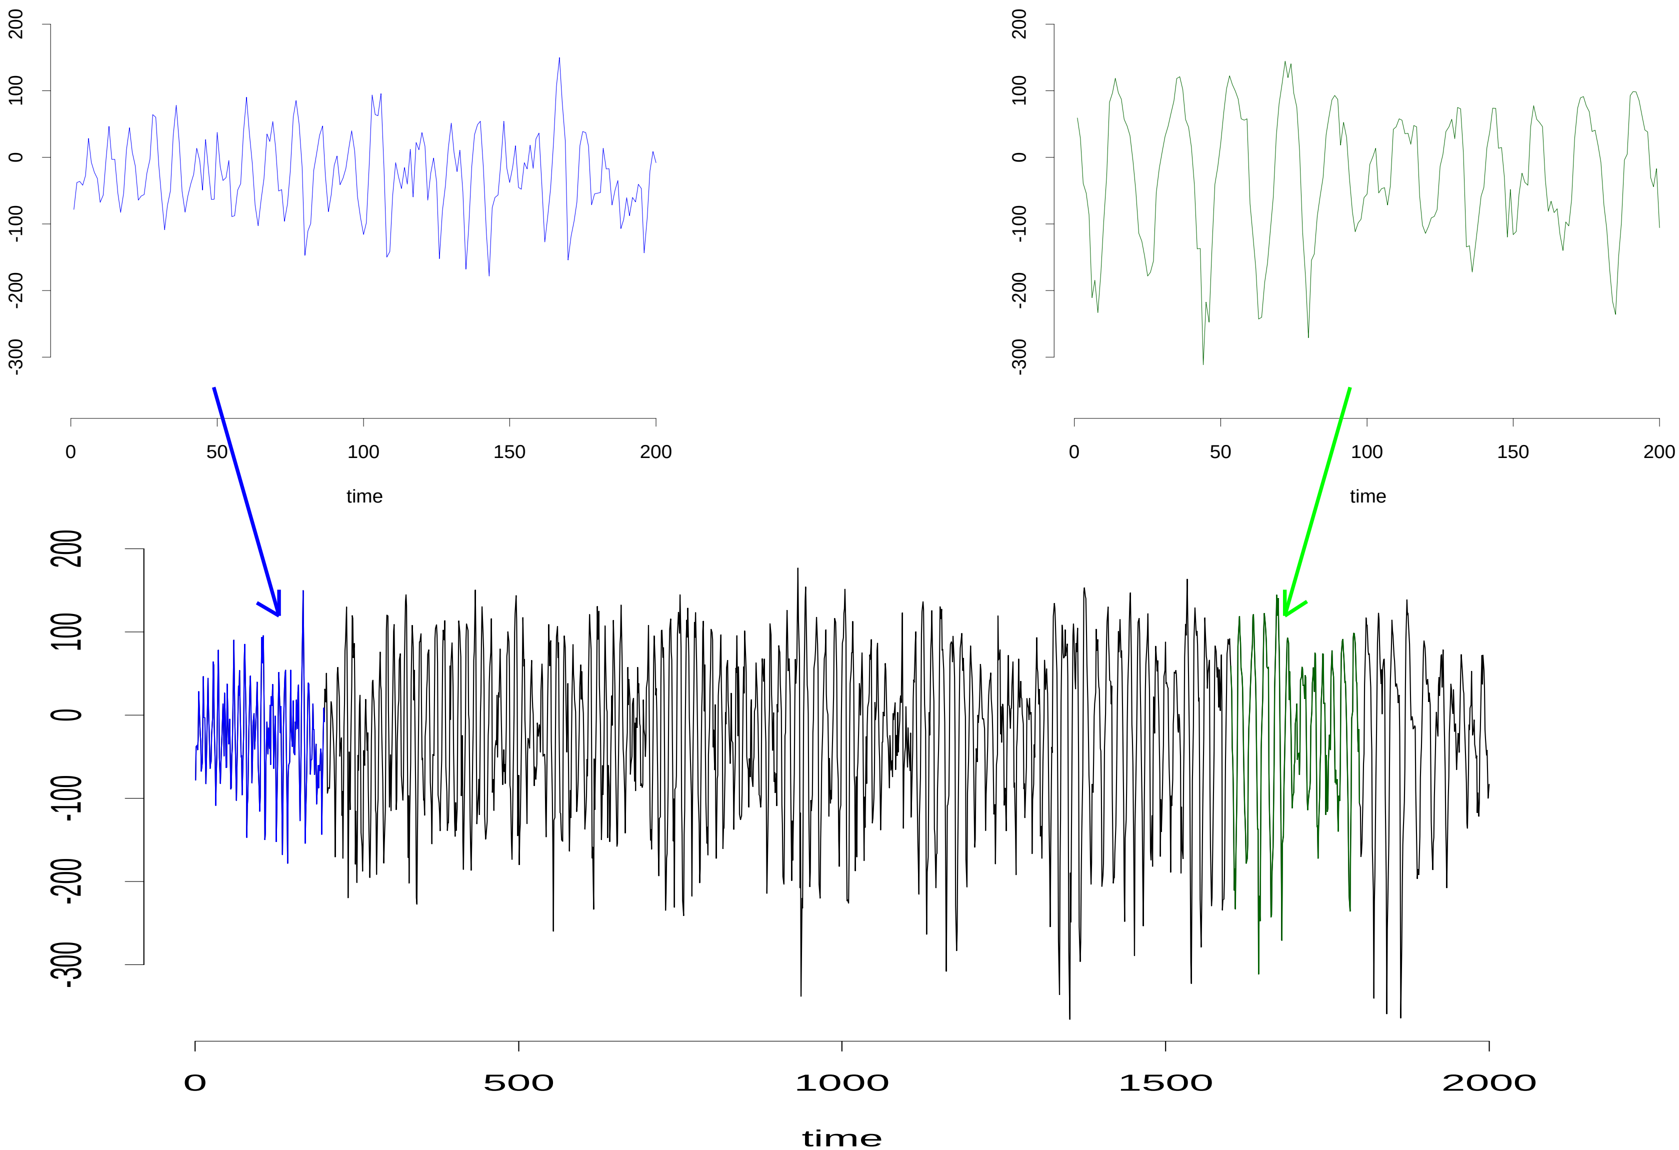

EEG Decomposition Examples

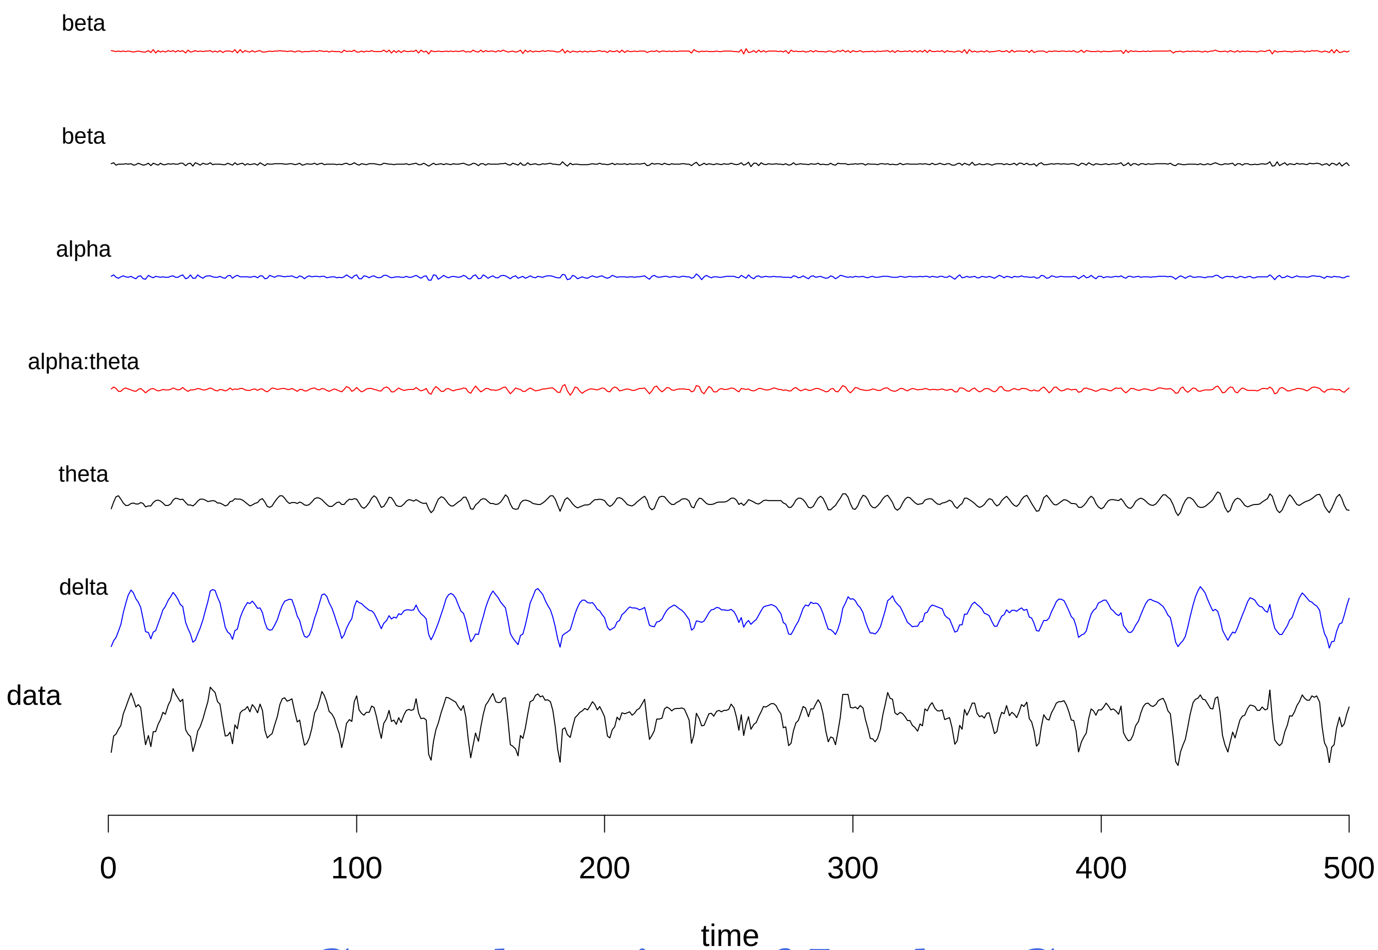

Central section of Ictal19-Cz

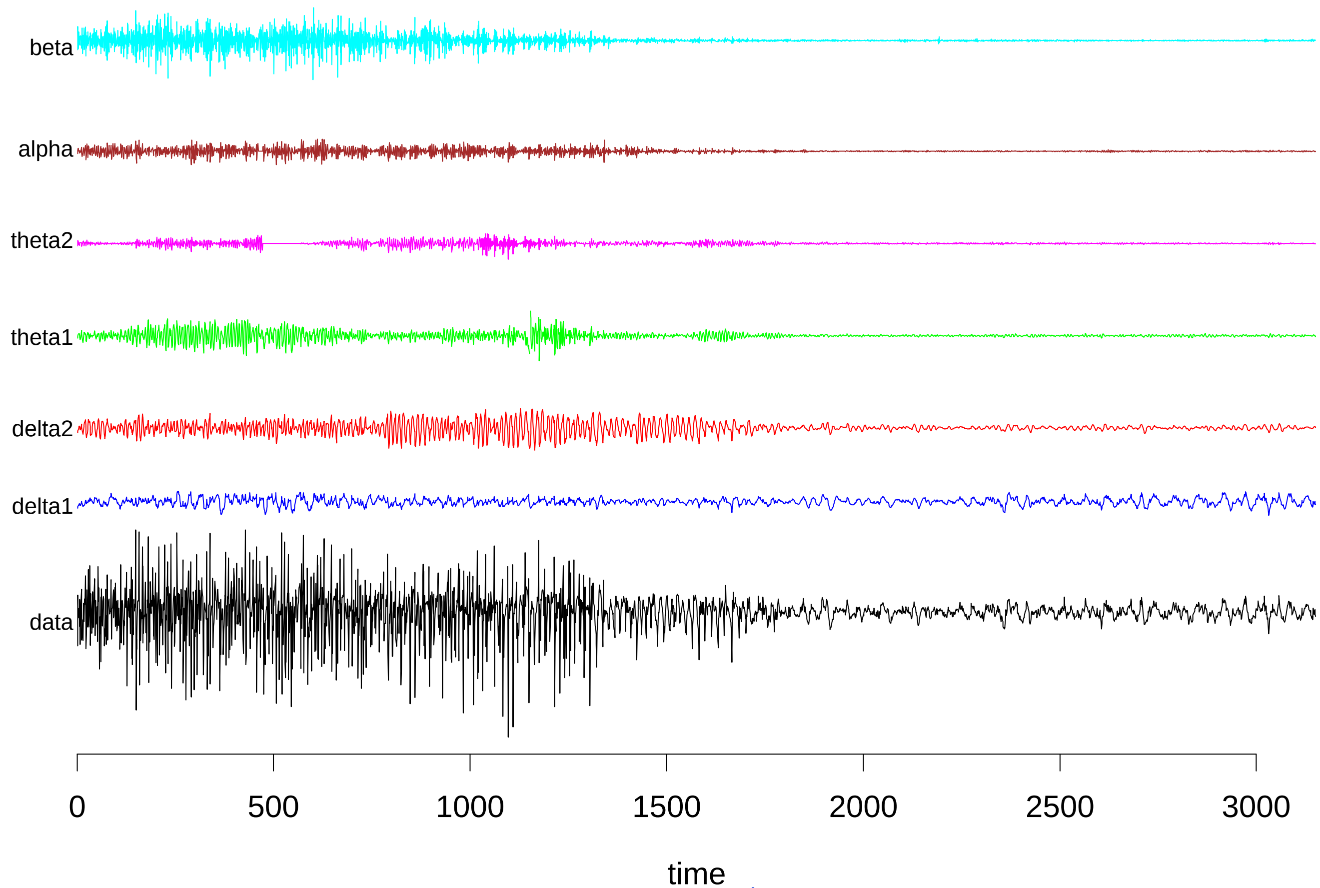

S26.Low Patient/Treatment

EEG Treatment Comparison

Two EEG series on one individual: S26.Low cf S26.Mod

- Repeat seizures with varying **ECT treatment**
- TVAR(20) with time-varying σ_t^2
- Evident high frequency structure and "spiky" traces
  $\rightarrow$ higher order models
- Several frequency bands influenced by seizure – several components
- **Identification issues**

EEG Comparisons

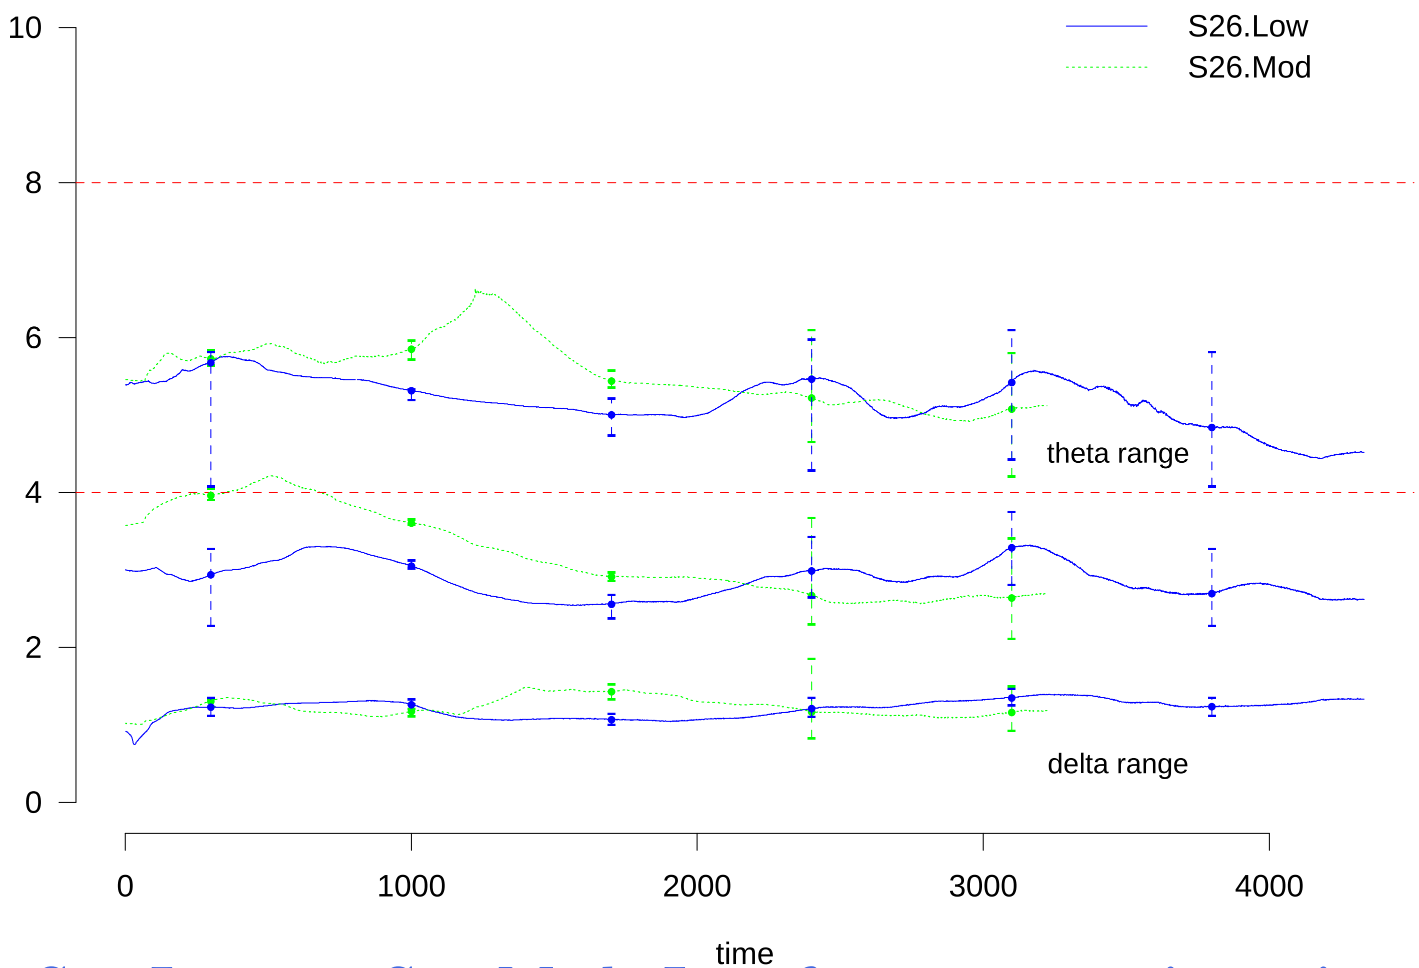

S26.Low vs. S26.Med: Low frequency trajectories

Sequential Simulation Analysis

Time $t:$

- States $\boldsymbol{\Theta}_t = \{\boldsymbol{\theta}_{t-k}, \ldots, \boldsymbol{\theta}_t\}$ of interest
- Summarised via set of posterior samples: $p(\boldsymbol{\Theta}_t|D_t)$

Time $t+1:$

- Observe $y_t \sim p(y_t|\boldsymbol{\theta}_t)$
- Require updated summary, posterior samples: $p(\boldsymbol{\Theta}_{t+1}|D_t)$

Issues:

- Expanding/Changing state space and dimension
- Simulation-based summaries: discrete approximations
- Inference on parameters as well as state vectors
- New data may "conflict" with prior/predictions

Particle Filtering

Key Goal: sequentially update posteriors

$$\cdots \rightarrow \; p(\boldsymbol{\theta}_t|D_t) \; \rightarrow \; p(\boldsymbol{\theta}_{t+1}|D_{t+1}) \; \rightarrow \cdots$$

Numerical Approximations (points/weights):

$$\{\boldsymbol{\theta}_t^{(j)}, \omega_t^{(j)} : \; j = 1, \ldots, N_t\}$$

Theoretical update:

$$p(\boldsymbol{\theta}_{t+1}|D_{t+1}) \propto p(\mathbf{y}_{t+1}|\boldsymbol{\theta}_{t+1}) p(\boldsymbol{\theta}_{t+1}|D_t)$$

MC approximation to "prior":

$$p(\boldsymbol{\theta}_{t+1}|D_t) \approx \sum_{k=1}^{N_t} \omega_t^{(k)} p(\boldsymbol{\theta}_{t+1}|\boldsymbol{\theta}_t^{(k)})$$

- Mixture prior: sample and accept/reject ideas natural

Example: Auxilliary Particle Filter

APF state update from $t \rightarrow t+1$:

- for each k,
  - "estimates" $\boldsymbol{\mu}_{t+1}^{(k)} = E(\boldsymbol{\theta}_{t+1}|\boldsymbol{\theta}_t^{(k)})$
  - and weights $g_{t+1}^{(k)} \propto \omega_t^{(k)} p(\mathbf{y}_{t+1}|\boldsymbol{\mu}_{t+1}^{(k)})$
- sample (aux) indicators j with probs $g_{t+1}^{(j)}$
- **time $t+1$ samples**: $\boldsymbol{\theta}_{t+1}^{(j)} \sim p(\boldsymbol{\theta}_{t+1}|\boldsymbol{\theta}_t^{(j)})$
- **and weights**:

$$\omega_{t+1}^{(j)} = \frac{p(\mathbf{y}_{t+1}|\boldsymbol{\theta}_{t+1}^{(j)})}{p(\mathbf{y}_{t+1}|\boldsymbol{\mu}_{t+1}^{(j)})}$$

Multivariate Models in Finance

(*Quintana* et al *invited talk, Valencia VII*)

- Futures markets, exchange rates, portfolio selection
- Multiple time series: time-varying covariance patterns
- Econometric/dynamic regressions/hierarchical models
- Latent factors in hierarchical, dynamic models
- Common time-varying structure in multiple series
- Bayesian multivariate stochastic volatility

$$\mathbf{y}_t = (y_{1t}, \ldots, y_{pt})'$$

e.g., y_{it} is $p-$vector of *returns* on investment i
(exchange rate futures, etc.)

Dynamic Factor Models

Exchange rate modelling for dynamic asset allocation

- Monthly (daily) currency exchange rates
- Dynamic regression/econometric predictors
- Residual structure and residual stochastic volatility
  - Time-varying variances and covariances
  - Dynamic factor models
- Dynamic asset allocation & risk management: portfolio studies
- Bayesian analysis: model fitting, sequential analysis, forecasting, decision analysis

Excess Returns on Exchange Rates (monthly)

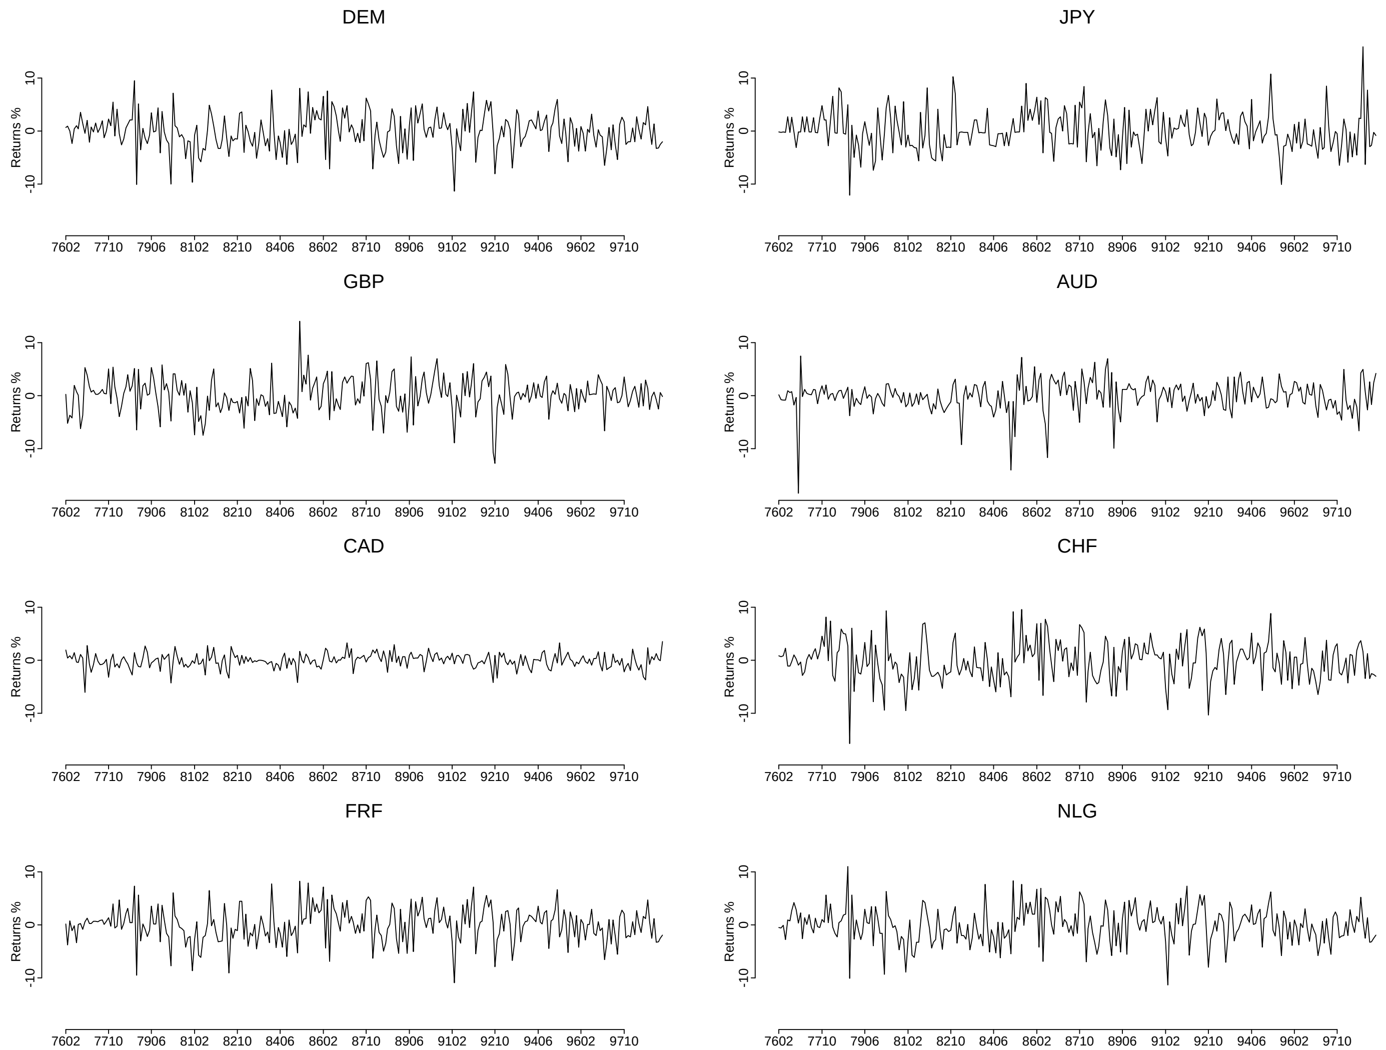

Short Interest Rates (monthly)

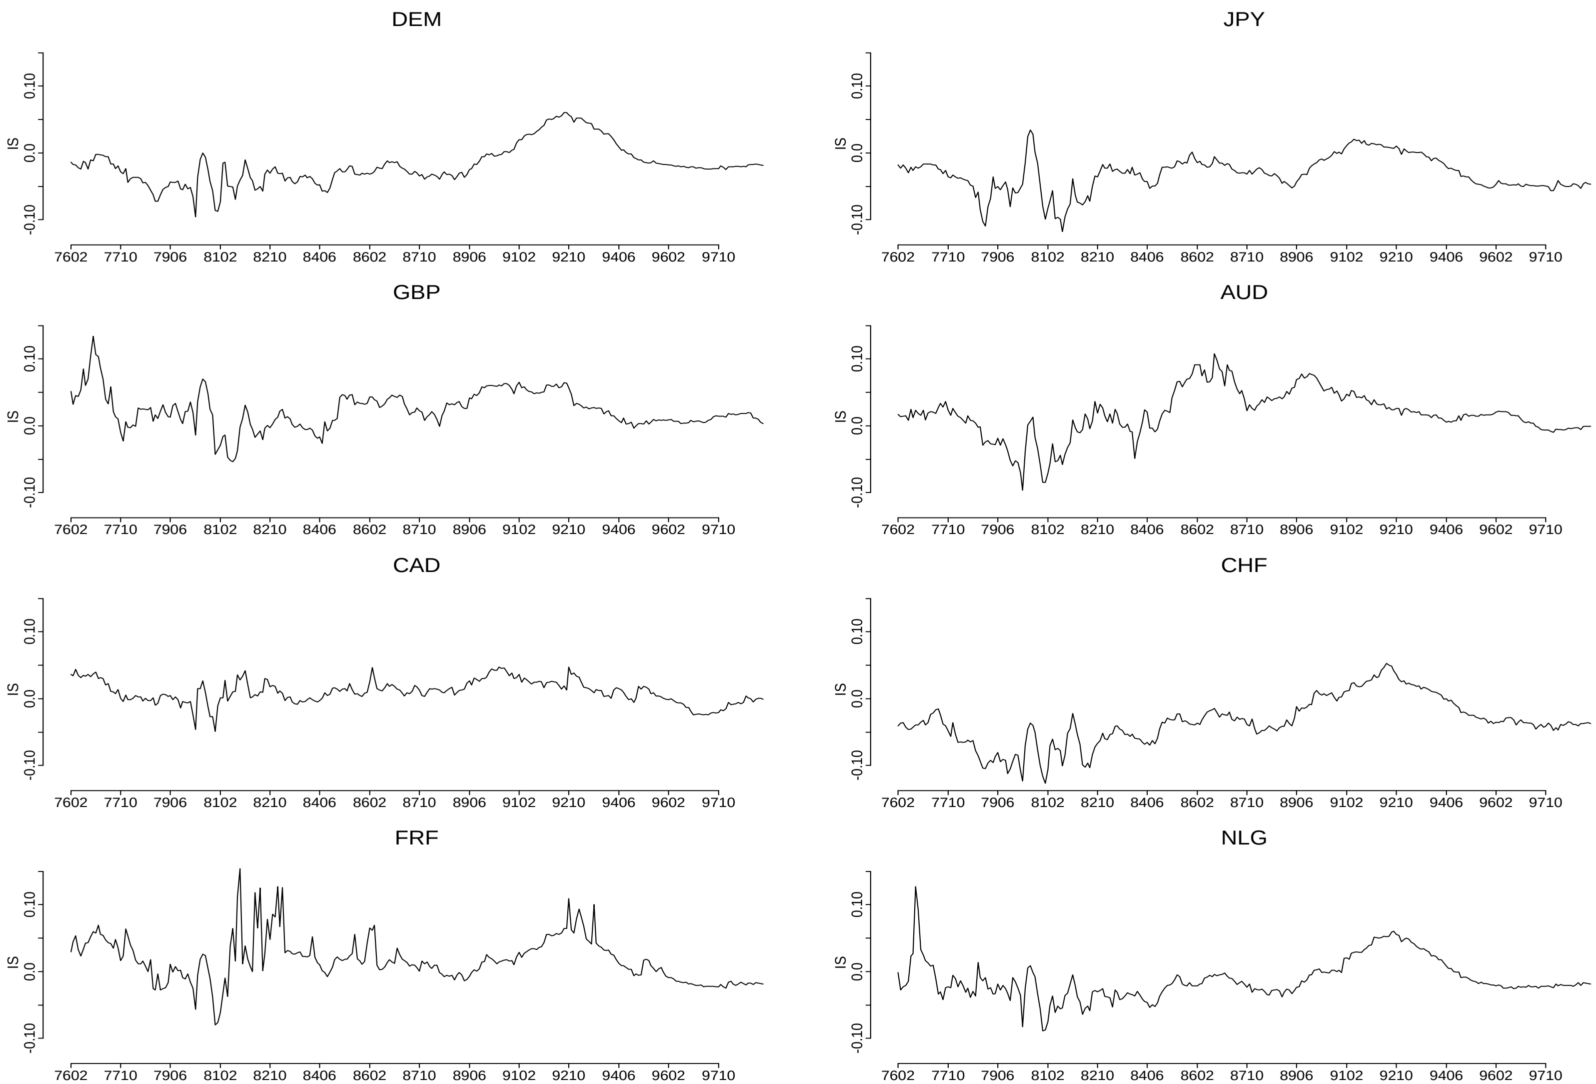

Stochastic Volatility Factor Models

$\mathbf{y}_t = \text{dynamic regression}_t + \text{residual}_t$ – $\text{residual}_t \sim N(\mathbf{0}, \mathbf{\Sigma}_t)$

$$\begin{aligned} \text{residual}_t &= \mathbf{X}_t \mathbf{f}_t + \mathbf{e}_t \\ \mathbf{f}_t &\sim N(\mathbf{f}_t | \mathbf{0}, \mathbf{F}_t) \\ \mathbf{e}_t &\sim N(\mathbf{e}_t | \mathbf{0}, \mathbf{E}_t) \end{aligned}$$

$\mathbf{f}_t$ ... $k-$vector of latent factors

$\mathbf{F}_t = \text{diag}(\exp(\lambda_{1,t}), \ldots, \exp(\lambda_{k,t}))$

$\mathbf{e}_t$ ... $q-$vector of "idiosyncracies"

$\mathbf{E}_t = \text{diag}(\exp(\lambda_{k+1,t}), \ldots, \exp(\lambda_{k+q,t}))$

$$\mathbf{\Sigma}_t = \mathbf{X}_t \mathbf{F}_t \mathbf{X}_t' + \mathbf{E}_t$$

Factor and idiosyncratic latent log volatilities: $\lambda_t = (\lambda_{1,t}, \ldots, \lambda_{k+q,t})'$

Dynamic Factor Model: Factor Loadings Structure

- Constant factor loadings matrix $\mathbf{X}_t = \mathbf{X}$

  – or "slowly varying" over time –

- Identification constraints on $\mathbf{X}$:

$$\mathbf{X} = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ x_{2,1} & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \cdots & 0 \\ x_{k,1} & x_{k,2} & x_{k,3} & \cdots & 1 \\ x_{k+1,1} & x_{k+1,2} & x_{k+1,3} & \cdots & x_{k+1,k} \\ \vdots & \vdots & \vdots & \cdots & \vdots \\ x_{q,1} & x_{q,2} & x_{q,3} & \cdots & x_{q,k} \end{pmatrix}$$

- Series order *defines* interpretation of factors

Stochastic Volatility Models: Factors & Idiosyncracies

Multivariate SV models

- vector AR(1) model for latent log volatilities λ_t
- volatility persistence: AR(1) coefficients $\mathbf{\Phi}$ (diagonal)
- marginal volatility levels: time-varying μ_t
- dependent innovations: shocks to volatilities related across series
  global/sector "effects" represented through correlations

$$\lambda_t = \mu_t + \mathbf{\Phi}(\lambda_{t-1} - \mu_{t-1}) + \omega_t$$

$$\omega_t \sim N(\mathbf{0}, \mathbf{U})$$

$$\mu_t \sim \text{random walk}$$

Dynamic Regressions and Shrinkage Models

- several predictors (e.g., interest rates)
- regression coefficients time-varying
- shrinkage models: coefficients related across series

e.g., sensitivity to short-term interest rates "similar" across currencies

Series j, time t :

$$\beta_{jt} = \phi_{jt}\beta_{j,t-1} + (1 - \phi_{jt})\gamma_t + \text{innovation}_{jt}$$

- global or sector "average" γ_t
- time-varying degrees of "shrinkage" ϕ_{jt}
- multiple series, several predictors: state-space model formulations

Model Fitting & Analysis: MCMC

Inference based on a fixed sample over $t = 1, \ldots, T$

- Bayesian analysis via posterior simulations
- Monte Carlo samples from *joint posterior* for
  - model parameters, dynamic regression parameters, AND
  - *latent processes: factors & volatilities*

$$\{ \mathbf{f}_t, \boldsymbol{\lambda}_t : \; t = 1, \ldots, T \}$$

- examples of highly structured, hierarchical models with many latent variables and parameters
- custom MCMC built of standard "modules"
- simple MC evaluation of forecast/predictive distributions

Model Fitting & Analysis: Sequential Analyses

Sequential particle filtering over $t = T+1, T+2, , \ldots$

- "particulate" approximation to posterior distributions
- **prior**: cloud of particles, associated importance weights
- **posterior**: sampling/importance resampling to *revise* posterior samples

$$\cdots \to p(\cdot|D_{t-1}) \to p(\cdot|D_t) \to \cdots$$

- sample "regeneration" by local smoothing of particulate prior
- time-varying latent variables *and* model parameters

Volatilities (SD) of Currency Returns (monthly)

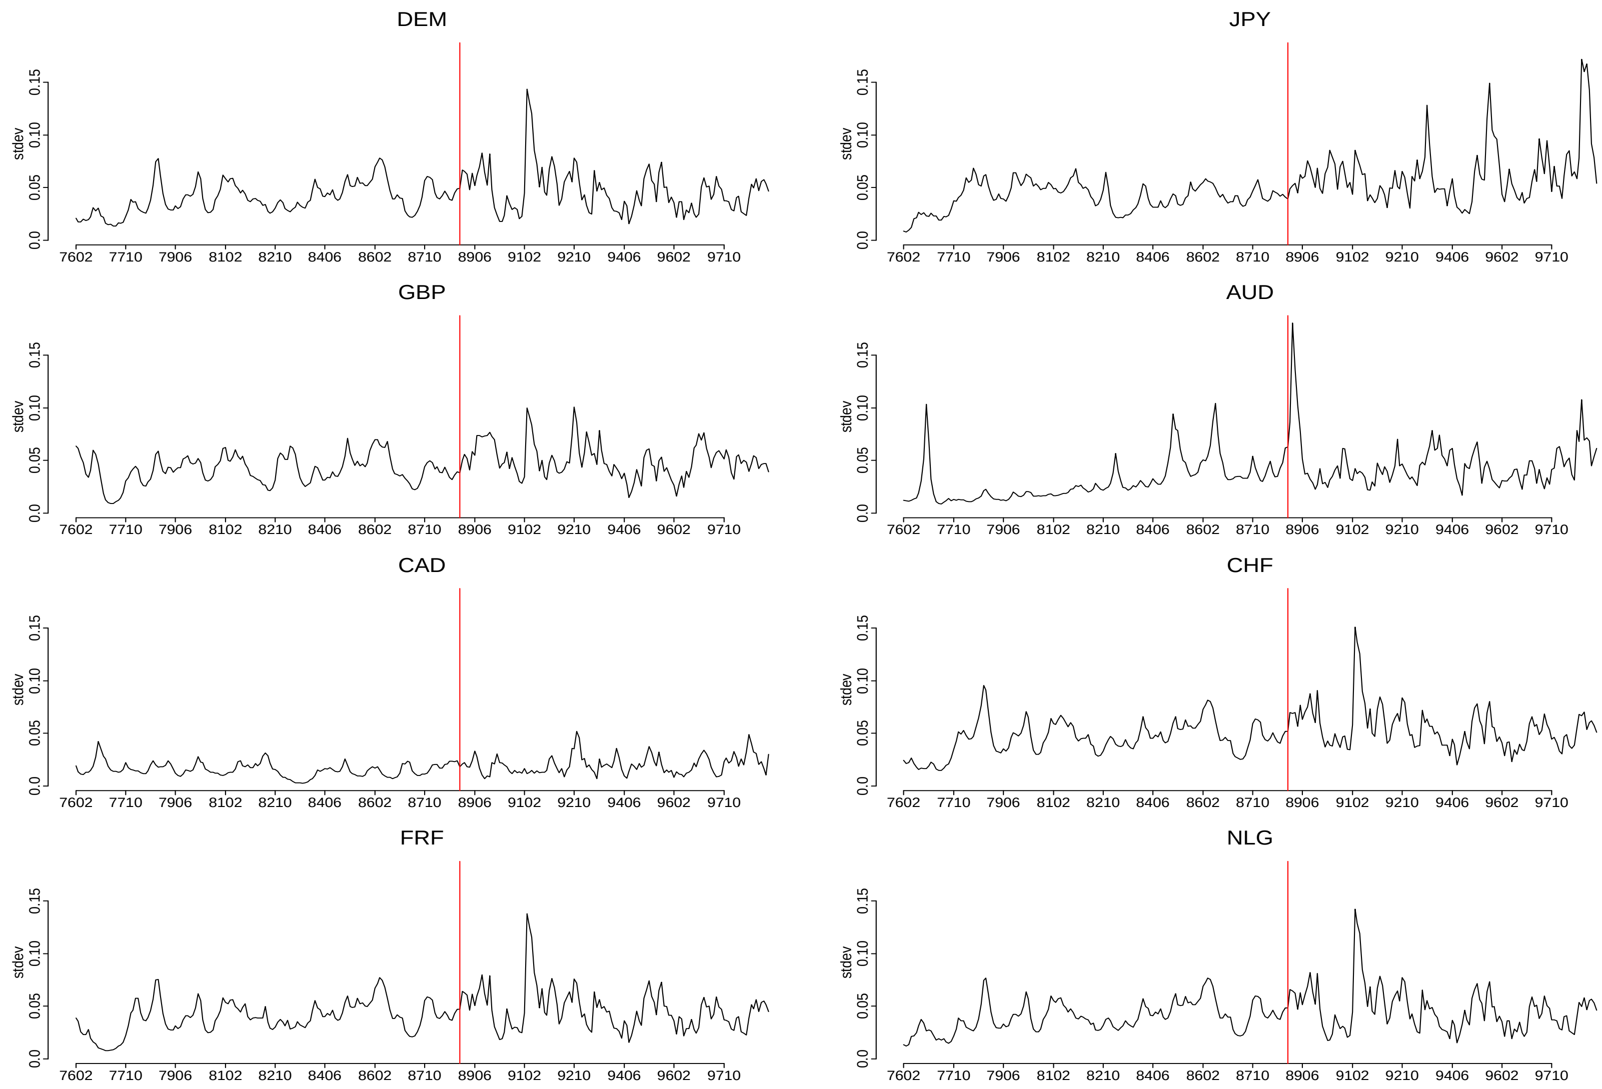

2 Latent Factors and their Volatilities (SD) (monthly)

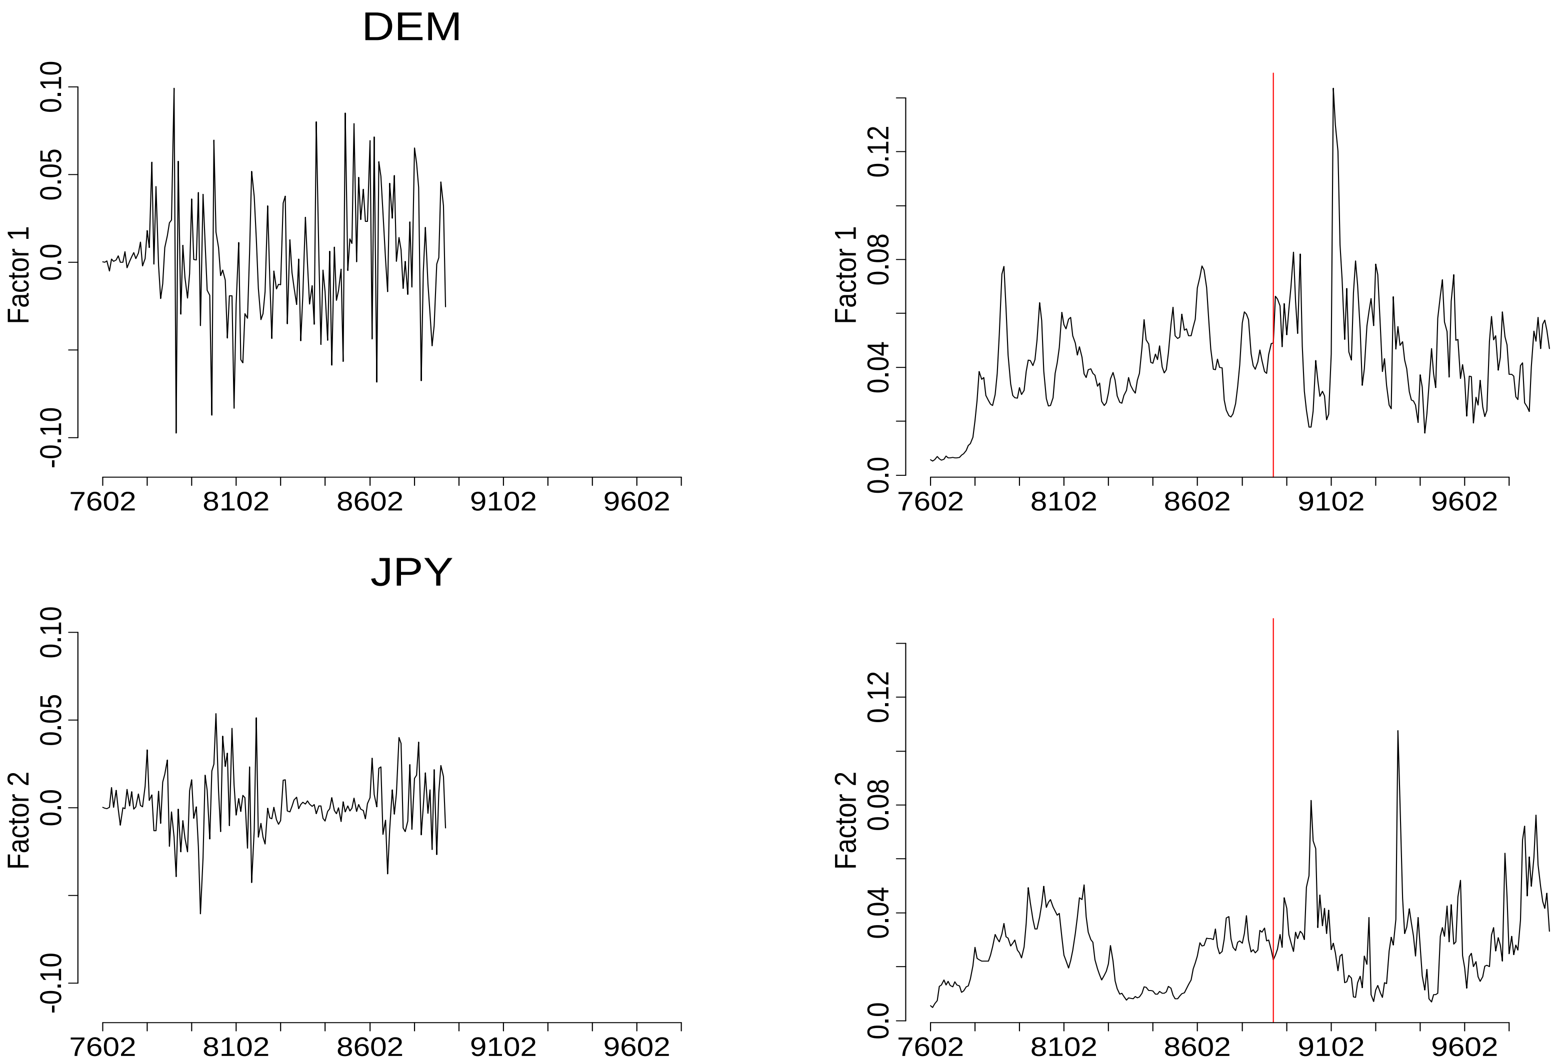

3 Latent Factors and their Volatilities (SD) (monthly)

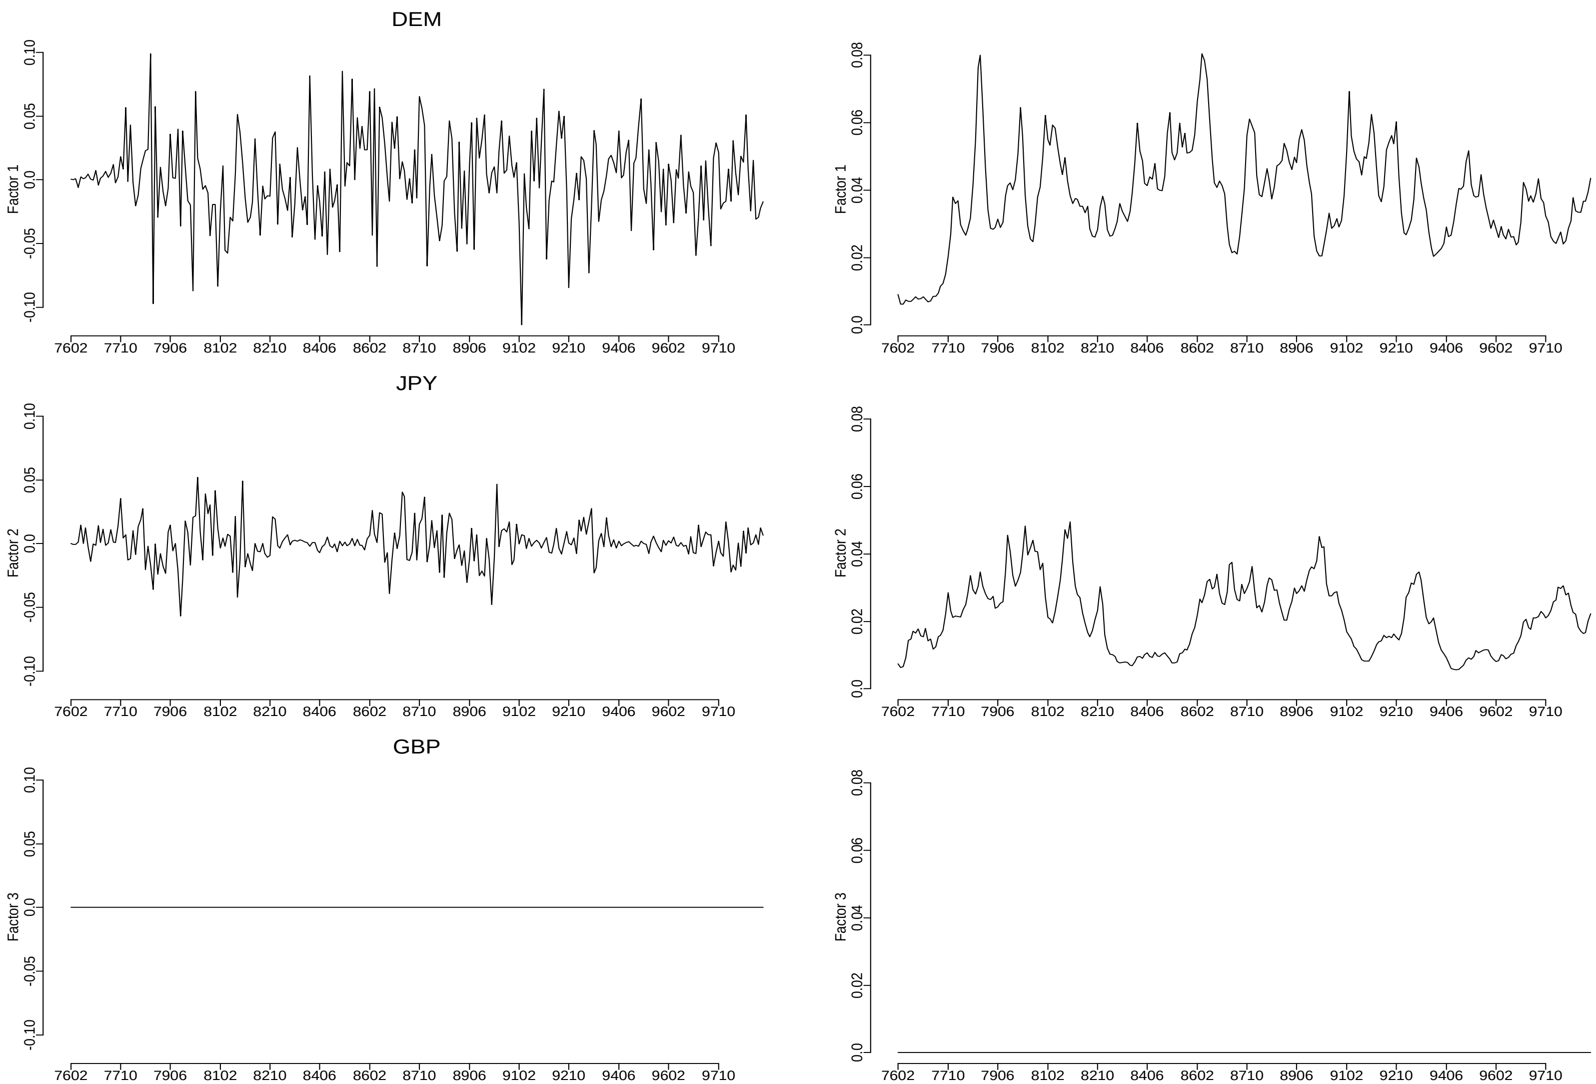

Factor Loading Matrix X (monthly)

	Factor 1	Factor 2
DEM	**1.00 (0.00)**	**0.00 (0.00)**
JPY	**0.55 (0.06)**	**1.00 (0.00)**
GBP	**0.68 (0.05)**	**0.48 (0.14)**
AUD	**0.23 (0.03)**	**0.35 (0.07)**
CAD	**0.04 (0.02)**	**0.03 (0.08)**
CHF	**1.04 (0.03)**	**0.23 (0.07)**
FRF	**0.96 (0.01)**	**-0.01 (0.02)**
ESP	**0.99 (0.01)**	**-0.01 (0.01)**

Factor Loading Matrix X (monthly)

	Factor 1	Factor 2
DEM	**1**	·
FRF	**1**	·
ESP	**1**	·
CHF	**1**	**0.2**
GBP	**0.7**	**0.5**
JPY	**0.5**	**1**
AUD	**0.2**	**0.3**
CAD	·	·

Idiosyncratic Volatilities (SD) (monthly)

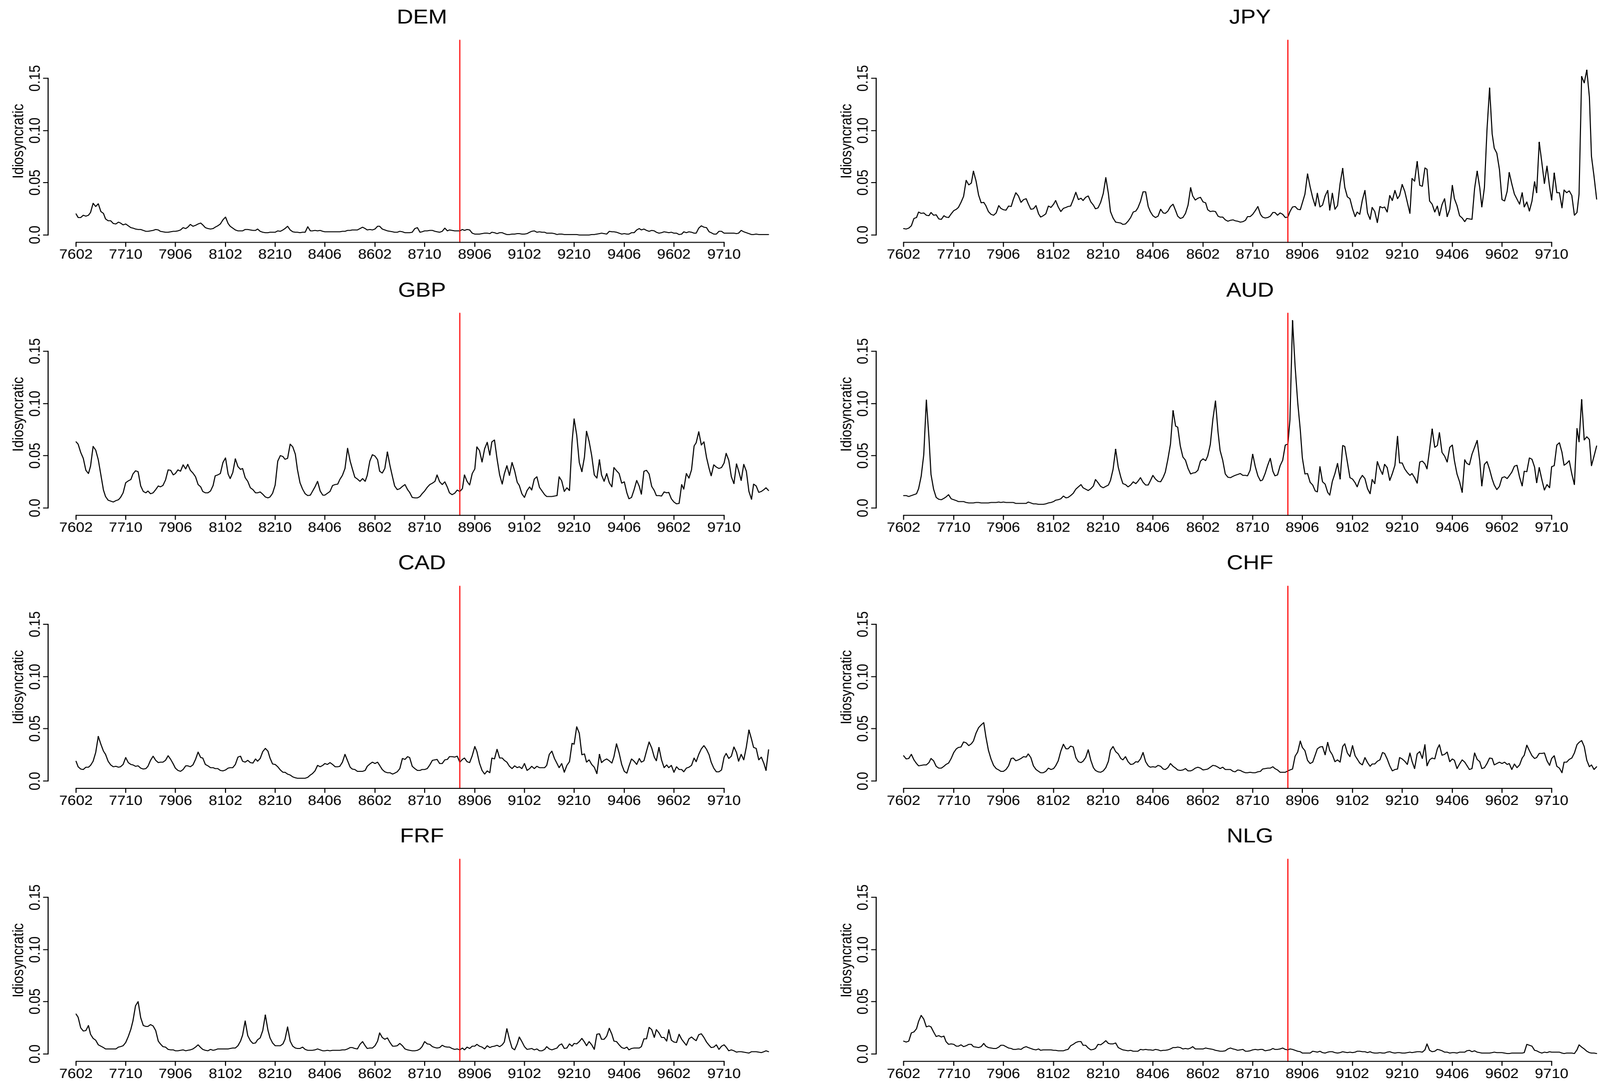

Financial Time Series

- Models/forecasts feed into portfolio decision management
  - *Quintana* et al *invited talk, Valencia VII*
- Live/Operational assessments of dynamic factor models
- Improved sequential particle filtering: parameters
- Time-varying loadings matrices $\mathbf{X}_t$, parameters
- Uncertainty about number of factors

Bayesian Time Series, Currently

Applied aspects:

- Financial modelling and forecasting
- Natural/engineering sciences: signal processing
- Spatial time series: epidemiology, environment, ecology

Models and methods:

- Highly structured multiple time series
- Spatial time series
- Computational methods: Sequential simulation methods

Links & Materials

- Books and papers: `www.isds.duke.edu/~mw`
  - *Copious references to broad literatures*
- 1997 Tutorial – extensive and historical – at website
- Teaching materials, notes, from past time series courses
- Software: `www.isds.duke.edu/~mw`
  - *TVAR (matlab, fortran), AR component models, BATS, links*
- Encyclopedia of Statistical Sciences (1998): *Bayesian Forecasting* (eds: S. Kotz, C.B. Read, and D.L. Banks), Wiley.
- 1997 2nd Edition: West & Harrison/Springer book
  Bayesian Forecasting & Dynamic Models
- Aguilar, Prado, Huerta & West (1999), Valencia 6 invited paper

Key Recent (since 1995) & Current Authors

– many/most are here! –

Omar Aguilar (Lehman Bros)

Chris Carter (Hong Kong)

Sid Chib (Washington Univ/St Louis)

Sylvia Frühwirth-Schnatter (Vienna)

Simon Godsill & Co (Cambridge)

Gabriel Huerta (Univ New Mexico)

Genshiro Kitagawa (ISM Tokyo)

Robert Kohn (Sydney)

Jane Liu (UBS NY)

Hedibert Lopes (Rio)

Viridiana Lourdes (ITAM Mexico)

Giovanni Petris (Univ Arkansas)

Michael Pitt (Warwick)

Nicolas Polson (Chicago)

Raquel Prado (Santa Cruz)

José M Quintana (Nikko NY)

Peter Rossi (Chicago)

Neil Shephard (Oxford)

... and, of course, ...

Valencia VII Invited Papers

- Quintana, Lourdes, Aguilar & Liu
  - Global gambling: multivariate financial time series
- Davy & Godsill
  - Signal processing, latent structure, MCMC

... plus a number of contributed talks and posters