BAYESIAN DYNAMIC MODELS: TIME SERIES ANALYSIS & FORECASTING

1-day Short Course | Sunday 3rd August, 2014 - 08:30-17:00 | JSM 2014, Boston MA

This short-course covers basic principles and methods of Bayesian dynamic modeling in time series analysis and forecasting, with methodological details of central model classes explored in a range of examples. A main focus is on dynamic linear models and related methods of inference and forecasting, including multivariate time series analysis. Links between time and frequency domain, and stationary time series models, will be covered, as well as selected developments in nonlinear and non-Gaussian dynamic models and associated Monte Carlo Markov chain simulation methods for analysis. The course will conclude by contacting some recent modeling and applied developments in multivariate time series and forecasting. The course draws on a range of examples and case studies from business, finance, signal processing and the biomedical sciences.

Course participants will gain:

- exposure to the basic ideas and approaches of Bayesian model-based time series analysis using key classes of dynamic models;
- an appreciation of the roles of computation— analytic- as well as simulation-based methods— for time series analysis in dynamic models, including filtering, parameter learning and smoothing;
- awareness of texts and software that will enable follow-on explorations and analysis;
- an appreciation of some of the breadth of application Bayesian dynamic modelling has had, and can have, in various applied fields.

The course material will be accessible to advanced students, academics and/or professionals with strong statistical modelling backgrounds and prior exposure to essentials of Bayesian analysis. Familiarity with— and working facility in— multivariate distribution theory and statistical inference are prerequisites. Prior exposure to some areas of time series analysis will be useful though is not necessary. Prospective participants can get a focused flavour of the level and nature of the material from the Reading web page. The course uses software (R and Matlab) of the authors, linked to the Software web page.

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

$$y_{t,j} = \log(\text{Price}_{t,j}), \quad \text{month } t, \text{asset } j$$

$$\begin{aligned} \mathbf{y}_t' &= \mathbf{F}_t'\Theta_t + \boldsymbol{\nu}_t', \\ \Theta_t &= \mathbf{G}_t\Theta_{t-1} + \boldsymbol{\Omega}_t \end{aligned}$$

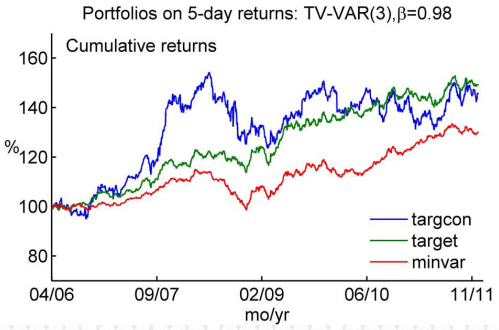

ASA Section on Bayesian Statistical Science

ISBA INTERNATIONAL SOCIETY FOR BAYESIAN ANALYSIS

BAYESIAN DYNAMIC MODELS: TIME SERIES ANALYSIS & FORECASTING

1-day Short Course | Sunday 3rd August, 2014 - 08:30-17:00 | JSM 2014, Boston MA

Schedule (with sessions annotated by sections of Prado & West— P&W)

The outline schedule links to the slides. We will be selective and responsive to course participants' interests, so that the slides here should be regarded as encompassing; some of the examples and detailed technical developments in some of the slides may be covered relatively briefly compared to others, to ensure effective coverage of core material leavened with engaging and topical examples.

Session 1 (P&W: 4.1-4.3)

- Introduction - Dynamic linear models (DLMs)
- Trend, seasonal & regression models; Composition; Discount factors
- Sequential learning & forecasting; Retrospective analysis/smoothing

Session 2 (P&W: 2.1, 4.1-4.2, 3.1, 4.5)

- DLMs for AR models: Decompositions, ties to frequency analysis
- Examples of MCMC in dynamic models: FFBS, parameters
- More on sequential learning, forecasting, retrospective analysis/smoothing

Session 3 (P&W: 7.5, 6.1-6.2)

- More examples of MCMC in DLMs if desired (mixtures of DLMs, etc)
- Non-stationary models: TVAR(p) models
- Decompositions and time-varying frequency analysis

Session 4 (P&W: 8.1, 9.1, 10.1-10.4)

- Selected examples of multivariate dynamic models:
- Exchangeable DLMs; Elements of VAR and TV-VAR models
- Multivariate volatility: Variance matrix learning & discounting
- Forecasting & decisions: Financial time series examples

**breaks also provide some opportunity to explore code and examples*

BAYESIAN DYNAMIC MODELS: TIME SERIES ANALYSIS & FORECASTING

1-day Short Course | Sunday 3rd August, 2014 - 08:30-17:00 | JSM 2014, Boston MA

Books and Paper

- Time Series: Modeling, Computation, and Inference, by Raquel Prado & Mike West, 2010, Chapman Hall/CRC Press Taylor & Francis Group.

  This is the course text. Registeredcourse participants have a one-time offer of a 25% discount, obtained by pasting the discount code ...TBA... into the *Promotion Code* field at checkout at the CRC web site.

  The 2010 P&W text covers all course material at a detailed level, and much more. A number of the examples and data sets in the course are taken from this text. The course sessions on the course schedule are annotated by by relevant sections of the book. Reviewing at least introductory material in the book in advance is highly recommended. A cursory read over the sections annotated in the course Schedule is an ideal way of ensuring and understanding of the level/prequisite knowledge in order to find the course immediately accessible.

- Bayesian Forecasting & Dynamic Models, by Mike West & Jeff Harrison, 1997 (2nd edition), Springer-Verlag.

  Some participants may already have— or will likely find useful— this standard text. W&H covers the core theory and methodology of dynamic models, Bayesian forecasting and time series analysis in extensive and foundational detail. By virtue of vintage, W&H is not so advanced as the P&W text on more recent applications and, especially, computation and multivariate dynamic modelling. A number of the examples and data sets in the course are taken from W&H.

- Bayesian Dynamic Modelling, by Mike West.

  Mike's overview paper (written in late 2011, published in early 2013) contacts much of the course at a gentle, non-technical overview level. Importantly, the paper has a long list of references linked to topics and material covered in this course, so participants can quickly and easily find connections to the literature to follow-up and explore more deeply after the course. This overview paper can be considered a companion to the course-linked video.

BAYESIAN DYNAMIC MODELS: TIME SERIES ANALYSIS & FORECASTING

1-day Short Course | Sunday 3rd August, 2014 - 08:30-17:00 | JSM 2014, Boston MA

Software

Course attendees may bring laptops to explore modelling examples (Matlab & R code). The course will be based on the instructors working through slide presentation material and interactively exploring several examples during the sessions throughout the day. Attendees will benefit most by exploring code and working through examples, so prior experience with Matlab & R will be most beneficial.

The course will explore a range of examples using:

- R code:
  - Raquel's R code for several course examples.
  - Dynamic Linear Models with R, by Giovanni Petris, Sonia Petrone and Patrizia Campagnoli.

  Raquel's R code covers a lot of the basic models of this course, including core univariate DLM analysis and forecasting, MCMC methods, and more. Some of this also uses the DLM R package that comes with Giovanni's 2009 Springer book; this is a free-standing package with a lot of the basic dynamic model facilities, and more, and is fully documented and annotated by the text.

- Mike's Matlab code for several course examples. Mike's Matlab code covers of lot of basics as well, including analysis and forecasting in a range of univariate and multivariate DLMs.

This software is made available to course participants on the understanding that they will not advertise or distribute it, nor use it for anything but educational and research purposes. Raquel & Mike can provide no support nor assistance with implementations beyond the course examples, nor extensions of the code for other purposes. It is understood by the user that neither the authors, nor the ASA, the University of California or Duke University, bear any responsibility nor assume any liability for any end-use of this software. It is expected that appropriate credit/acknowledgement be given should the software be included as an element in other software development or in publications.

Other code of possible interest on both core and related models, though not part of the course:

- DLM-GASP models
- Several other pieces of Bayesian software including time series related code.

BAYESIAN DYNAMIC MODELS: TIME SERIES ANALYSIS & FORECASTING

1-day Short Course | Sunday 3rd August, 2014 - 08:30-17:00 | JSM 2014, Boston MA

Video(s)

Bayesian dynamic modelling

Mike's *Lecture on Bayesian Foundations* at the ISBA World Meeting (Kyoto, June 2012)

This tutorial/overview presentation contacts key foundational concepts and models (among other things). Course participants may find it useful (and easy viewing).

The above video is one of several of Mike's video presentations on topics in dynamic modelling, most at more advanced/research levels.

BAYESIAN DYNAMIC MODELS: TIME SERIES ANALYSIS & FORECASTING

1-day Short Course | Sunday 3rd August, 2014 - 08:30-17:00 | JSM 2014, Boston MA

Brief Biographies (@ August 2014)

Raquel Prado is Professor of Statistics in the Department of Applied Mathematics and Statistics at the Baskin School of Engineering of the University of California Santa Cruz, USA. Her research areas include time series analysis - with a focus on Bayesian modeling of large-dimensional nonstationary time series data - biomedical signal processing and statistical genetics. An elected ASA Fellow, Raquel has been a recipient of the Outstanding Statistical Application Award from the American Statistical Association and is currently a Board Member of the International Society for Bayesian Analysis (ISBA) and Associate Editor for JASA/TAS Reviews. She is a co-author of the book Time Series: Modeling, Computation, and Inference, and has published papers in mainstream statistical journals (e.g., JASA, Series C, Journal of Time Series Analysis and JSPI among others) and disciplinary journals. Professor Prado has supervised five graduate students, has been a member of the Savage Award Committee and the ISBA Prize Committee, and Chair of the ISBA Program Council responsible for, in particular, the 2014 ISBA World Meeting. Raquel is/has been PI or Co-PI of multiple projects funded by NSF and NIH. In addition to active teaching at all levels, Raquel is an experienced short-course presenter, having run courses on Bayesian analysis, time series and forecasting, in particular, for diverse audiences.

Mike West holds a distinguished chair as The Arts & Sciences Professor of Statistics & Decision Sciences in the Department of Statistical Science at Duke University, where he headed the development of the department (earlier the Institute of Statistics & Decision Sciences) from 1990-2002. A past president of the International Society for Bayesian Analysis and past SBSS Chair, Mike is an elected Fellow of the ASA, IMS, RSS and a founding elected Fellow of ISBA. Mike has served the international statistics communities on founding boards of both the National Institute of Statistical Sciences and the Statistical & Applied Mathematical Sciences Institute in the USA, as well as on boards of other national research institutes including the Institute of Statistical Mathematics in Japan, and the Institute for Mathematical Sciences in UK. Mike has received a number of awards and accolades for research, including the 1994, 1997 and 2012 Mitchell Prize for research in applied Bayesian statistics. Mike's research interests are in Bayesian methodology and applications, with specific emphases on dynamic models for time series and forecasting, latent structure and sparsity in complex statistical modelling, and computation. Interdisciplinary applications have spanned many areas, with current emphases in financial econometrics, and "big data", monitoring and fast analysis in dynamic networks, among others. Mike's professional "kicks" come from a variety of sources, but teaching and working with smart, young, emerging statisticians from all over the world is what he prizes most. In addition to working with and advising many undergraduates and Master's students, Mike has advised more than 56 primary PhD students and postdoctoral associates. Mike is also an experienced short-course presenter, having given multiple short courses on Bayesian dynamic modelling, time series and forecasting, to diverse audiences.