Bayesian Forecasting and Dynamic Models
– CBMS 2020 Lectures –

Schedule of Lectures

day 1	time 1	1. Dynamic Modeling & Time Series: Structure, Forecasting and Decisions
day 1	time 2	2. Dynamic Factor Models & and Time-Varying Covariance Models
day 2	time 1	3. Time-Varying Vector Autoregressions & Hierarchical Latent Factor Models
day 2	time 2	4. Dynamic Latent Threshold Modeling
day 3	time 1	5. Dynamic Dependency Network Models
day 3	time 2	6. Simultaneous Graphical Dynamic Models
day 4	time 1	7. Dynamic Count Systems: Network Flow
day 4	time 2	8. Dynamic Count Systems: Multiscale Forecasting of Count-Valued Time Series
day 5	time 1	9. Sequential Monte Carlo in Dynamic Models
day 5	time 2	10. Bayesian Predictive Synthesis in Time Series Forecasting

1. Dynamic Modeling & Time Series: Structure, Forecasting and Decisions
Mike West

The introductory lecture begins with discussion of background to the field of Bayesian forecasting and associated univariate and core multivariate dynamic model contexts. This sets perspective on the essentials of the course topics: the development and use of state-space models for forward/sequential learning and forecasting of time series. Examples of univariate dynamic linear models (DLMs) and some essentials of their analyses link to interests in applications in many areas, heavily driven historically by uses in business, economics, signal processing engineering, and others, but nowadays also exploited across many other fields. The context of sequential learning in a Bayesian framework is central, and resulting analysis is implemented via manipulation/coding of changing distributions that deliver posterior and predictive inferences as they change over time. Bayesian state-space modelling stresses interpretation: uncertain, time-varying model parameters are the vehicles that deliver forecast information in terms of Bayesian predictive distributions, so understanding their roles and impact as model components is fundamental to model use and intervention in live forecasting contexts.

Following this introduction, the lecture discusses aspects of multivariate dynamic models for time series, moving to more recent developments, core topics to be covered in later lectures in this series, and connections to open research areas. Recent and currently topical applied contexts include areas such as policy-oriented macro-economic forecasting, forecasting to feed into dynamic financial portfolio investment decisions, biomedical signal processing and monitoring, environmental time series, monitoring and anomaly detection in time-varying systems, among others. This overview sets the stage for the models, methods and applications to follow in lectures 2–10.

Readings:

- Core and basic background in Prado and West (2010) chapters 4 and 10, and West and Harrison (1997) chapters 1-4 and 16.
- West (2013) is a short, easy-to-read overview paper that contacts much of the core material in this lecture at non-technical, overview level. The paper gives historical context and has a long list of references linked to core topics and material.
- West (2020) is a recent overview/review focused heavily on multivariate time series models and forecasting, and that links to a range of topics in dynamic modelling and forecasting presented in later lectures here, again with many– and many more recent– references.

2. Dynamic Factor Models & Time-Varying Covariance Models
Hedibert Lopes

This lecture starts by reviewing the Bayesian approach to the linear Gaussian factor model, focusing on its dimension reduction properties and its constrained covariance structures. Extending factor models by allowing common factors and/or common and idiosyncratic factor variances to be dynamic has become commonplace in many applied areas, including macro-econometrics, financial econometrics and environmetrics. On the one hand, dynamic factor models are usually large-scale state-space models and commonly related to time-varying vector autoregressions (see Lecture 3 below). We will explore the highly flexible class of spatial dynamic factor models, where factor loadings are spatially structured and common factors are temporally evolving. On the other hand, factor stochastic volatility (FSV) and Cholesky stochastic volatility (CSV) models are amongst the more flexible representations of time-varying covariance matrices. In both cases, the factor loadings or the lower-triangular Cholesky matrix are used to induce sparsity in high-dimensional models, while the common factor and idiosyncratic factor variances follow conditionally independent univariate stochastic volatility structures. This lecture will also present computationally efficient methods for posterior inference and forecasting.

Readings: Lopes and Carvalho (2007); Lopes et al. (2008, 2011b); Lopes (2014); Kastner et al. (2017); Shirota et al. (2017); Lopes et al. (2018); Carvalho et al. (2018)

3. Time-Varying Vector Autoregressions & Hierarchical Latent Factor Models
Raquel Prado

This lecture presents state-space representations, decompositions and interpretation of vector autoregressions (VARs) and time-varying vector autoregressions (TV-VARs), viewing them as special cases of multivariate DLMs or conditionally Gaussian DLMs introduced in Lecture 1. Time-frequency and partial autocorrelation (PARCOR) representations of these models will also be discussed in detail. Furthermore, connections with latent factor models and hierarchical factor models will be explored and illustrated. This lecture will also present computationally efficient methods for posterior inference and forecasting. The models and methods will be illustrated in several application studies, with a focus on analysis of multi-dimensional data arising from neuroscience studies including, among others, the analysis of multi-location EEG data and high-resolution brain imaging data.

Readings: Prado and West (2010) chapters 5 and 9; Prado et al. (2001); Prado (2009, 2010)

4. Dynamic Latent Threshold Modeling
Mike West

This lecture links to Bayesian sparsity modelling in multivariate time series analysis and forecasting, and discusses the relevance and role of dynamic sparsity in a range of model contexts: dynamic factor models, time-varying vector autoregressions, multivariate volatility models, and more elaborate models based on combinations of these main forms. The core methodological focus is on recent innovations for dynamic sparsity modelling using the concept of dynamic latent thresholding.

The latent threshold concept defines a rather general framework for dynamic sparsity in multiparameter dynamic models: time-varying parameters are linked to latent processes that are probabilistically thresholded to induce zero/non-zero values adaptively over time. This induces natural mechanisms for dynamic variable inclusion/selection that can apply to any set of state-space model parameters. Resulting broad classes of multivariate latent threshold models (LTMs) provide the opportunity for time- and data-adaptive thresholding, with the practically effective parameters being exactly zero over periods of time and then non zero, but typically time-varying, in other periods. This can induce parsimony in multivariate dynamic models, reduce estimation and prediction uncertainties as result, improve predictive model fit and out-of-sample forecasts, and identify time-varying network relationships linked to dynamically evolving feed-forward and contemporaneous relationships among series.

The lecture discusses aspects of Bayesian model specification, analysis and prediction in dynamic regressions, time-varying vector autoregressions (TV-VAR), dynamic latent factor models, and Cholesky-style multivariate volatility models using latent thresholding. Elements of computation for model fitting and forecasting are also noted. Highlights of the utility of LTMs draw on examples in applications in several areas. This includes a series of econometric and financial studies where the LTM approach has had demonstrable success, and in biomedical signal processing as a statistical approach to identifying dynamics in network structures.

Readings: Nakajima and West (2013a,b); Zhou et al. (2014); Nakajima and West (2015, 2017)

5. Dynamic Dependency Network Models
Mike West

In recent years the structured modelling of multivariate time series using decouple/recouple ideas in the context of so-called dynamic dependency network models (DDNMs) has become standard technology, especially in stochastic dynamic models in the financial industry, core areas of macroeconomics and allied areas of business. DDNMs are Bayesian state-space models that characterize sparse patterns of dependence among multiple time series via extensions of traditional subset time-varying vector autoregressions, dynamic regressions and multivariate volatility models. The key features of DDNMs are (i) structural sparsity in representing cross-series linkages, and (ii) that they enable scaling to higher numbers of individual time series–

scaling of computations that is linear in the number of time series. The theory of DDNMs shows how the individual series can be *decoupled* for sequential analysis, and then *recoupled* for applied forecasting and decision analysis. Decoupling allows fast, efficient analysis of each of the series in individual univariate models that are linked– for later recoupling– through a theoretical multivariate volatility structure defined by a sparse underlying graphical model. Computational advances are especially significant in connection with model uncertainty about the sparsity patterns among series that define these graphical models; Bayesian model averaging using discounting of historical information can build substantially on this computational advance. An extensive, detailed case study showcases the use of these models, and the improvements in forecasting and financial portfolio investment decisions that are achievable. Using a long series of daily international currency, stock indices and commodity prices, the case study includes evaluations of multi-day forecasts and Bayesian portfolio analysis with a variety of practical utility functions, as well as comparisons against commodity trading advisor benchmarks.

Readings: Zhao et al. (2016); Nakajima and West (2013a, 2017)

6. Simultaneous Graphical Dynamic Models
Mike West

Building on the success and broader application of DDNMs, the major recent developments of a more general class of simultaneous graphical dynamic linear models (SGDLMs) has substantially advanced the ability to model and analyze much larger-scale systems of time series. A key and critical issue faced by DDNMs is the requirement to define an ordering of the individual, univariate series in a multivariate system. While in some contexts this is both natural and constructive– in terms of enabling incorporation of context and theory into models– it is more generally regarded as a challenging and potentially limiting requirement. Defining collectives of coupled systems of simultaneous dynamic models in the more recent SGDLM framework obviates these issues, and opens up the scope for major expansion in terms of application area and– critically– scalability of analysis. Essentially, SGDLMs define an over-arching framework for sensitive and flexible modelling of dynamics and relationships at the level of individual univariate series coupled with similarly flexible structures for representing cross-series/multivariate stochastic and dynamic volatility. As with the subclass of DDNMs, the computational demands of the resulting Bayesian analysis scale linearly in the number of series.

This lecture discusses the framework and decouple/recouple methodology of SGDLMs. The scope includes issues of variable selection, Bayesian computation for scalability, and a case study in exploring the resulting potential for improved short-term forecasting of large-scale volatility matrices. The case study concerns financial forecasting and portfolio optimization for decisions with a 400-dimensional series of daily stock prices. Analysis shows that the SGDLM forecasts volatilities and co-volatilities well, making it ideally suited to contributing to quantitative investment strategies to improve portfolio returns. Analysis also identifies performance metrics linked to the sequential Bayesian filtering analysis that turn out to define a leading indicator of increased financial market stresses, comparable to but leading standard financial risk measures. Parallel computation using GPU implementations substantially advance the ability to fit and use these models.

Readings: Gruber and West (2016, 2017)

7. Dynamic Count Systems: Network Flow
Mike West

Traffic flow count data in networks arise in many applications, such as automobile or aviation transportation, certain directed social network contexts, and Internet studies. With a central example of Internet browser traffic flow through site-segments of an international news website, this lecture highlights advances in modelling of dynamic network flows based on the decouple/recouple concept. Bayesian analyses of two linked classes of models, in tandem, allow fast, scalable and interpretable Bayesian inference. A novel class of flexible but decoupled state-space models for streaming count data is able to adaptively characterize and quantify network dynamics efficiently in real-time. These models are used as *emulators* of more structured, time-varying gravity models that allow formal dissection of network dynamics. This yields interpretable inferences on traffic flow characteristics, and on dynamics in interactions among network nodes. Bayesian monitoring theory defines a strategy for sequential model assessment and adaptation in cases when network flow data deviates from model-based predictions. Exploratory and sequential monitoring analyses of evolving traffic on a network of web site-segments in e-commerce demonstrate the utility of this coupled Bayesian emulation approach to analysis of streaming network count data.

For large-scale networks, second-stage development involves customized dynamic generalized linear models (DGLMs) integrated into the network context via the decouple/recouple approach. Development is anchored in an expanded, higher-dimensional case-study of flows of visitors to the commercial news web site defining a long time series of node-node counts on over 56,000 node pairs. Characterizing inherent stochasticity in traffic patterns, understanding node-node interactions, adapting to dynamic changes in flows and allowing for sensitive monitoring to flag anomalies are central questions. The methodology of dynamic network DGLMs will be of interest and utility in broad ranges of dynamic network flow studies, and the underlying dynamic models will apply in studies of integer flows in dynamical systems more broadly.

Readings: Chen et al. (2018, 2019)

8. Dynamic Count Systems: Multiscale Forecasting of Count-Valued Time Series
Mike West

Problems of forecasting many related time series of counts arise in many areas, such as consumer behaviour in a range of socio-economic contexts, various natural and biological systems, and commercial and economic problems of analysis and forecasting of flows and discrete outcomes. Such data are particularly prevalent in consumer demand and sales contexts. With one motivating applied context of multi-step ahead forecasting of daily sales of many supermarket items across a system of outlets, this lecture discusses new classes of dynamic models for complex series of counts. The models address efficiency, efficacy and scalability of dynamic models based on the concept of decouple/recouple applied to multiple series that are individually represented via novel univariate state-space models for non-negative counts. The latter involve dynamic generalized linear models for binary and conditionally Poisson time series, with dynamic random effects for over-dispersion, allowing use of dynamic covariates in both binary and non-zero count components. Sequential Bayesian analysis allows fast, parallel analysis of sets of decoupled time series. New multivariate models then enable information sharing in contexts when data at a more highly aggregated level provide more incisive inferences on shared patterns such as trends and seasonality. A novel multi-scale approach– a further example of the concept of decouple/recouple in time series– enables information sharing across series. This incorporates cross-series linkages while insulating parallel estimation of univariate models, hence enables scalability in the number of series.

Extension of these models are dynamic count mixture models that apply to forecasting individual customer transactions, coupled with a novel probabilistic model for predicting counts of items per transaction. The resulting transactions-sales models allow use of dynamic covariates in both transaction and sales levels components, and can incorporate a diverse range of trend, seasonal, price, promotion, random effects and other outlet-specific predictors at the level of individual items. Decouple/recouple enabling effective multi-scale analysis is again central. The motivating case study context of many-item, multi-period, multi-step ahead supermarket sales forecasting provides examples that demonstrate improved forecast accuracy in a range of traditional and statistical metrics, while also illustrating the benefits of full probabilistic models for forecast accuracy evaluation and comparison.

Readings: Berry and West (2019); Berry et al. (2020)

9. Sequential Monte Carlo in Dynamic Models
Hedibert Lopes

Linear and Gaussian dynamic models represent a vast and well known class of time series models. Its notoriety is arguably due to its closed form sequential/online Bayesian learning and updating, which allows sequential model comparison, model averaging and out-of-sample forecasting as observations arrive. Despite such generality, when Gaussianity and/or linearly are given up, closed-form sequential learning becomes computational infeasible. Markov chain Monte Carlo sampling schemes revolutionised the Bayesian approach to non-linear and non-Gaussian dynamic models by allowing accurate approximations of the smoothed distribution of the latent variables. Nonetheless, sequential learning remained highly costly since MCMC algorithms would have to be re-run when a new observation become available. Roughly speaking, Sequential Monte Carlo (SMC) methods, also commonly known as particle filters, perform online sampling importance resampling to approximate sequential Bayesian learning. We will revisit some of the most popular SMC schemes, such as the bootstrap filter, the auxiliary particle filter, the Liu-West filter and the particle learning filter.

Readings: Liu and West (2001); Carvalho and Lopes (2007); Carvalho et al. (2010); Lopes et al. (2011a); Lopes and Tsay (2011); Lopes and Prado (2013)

10. Bayesian Predictive Synthesis in Time Series Forecasting
Mike West

This lectures concerns Bayesian reasoning and modelling for comparison, calibration, and combination of forecast distributions in the context of multiple of models or forecasters. These issues are central to applied forecasting, linking back to 1980s/90s literatures cutting across statistics, economics, management science and other fields. Then, these questions are increasingly central in research as the ability to fit multiple models– and to access forecast information from multiple sources– escalates.

This lecture focuses on Bayesian predictive synthesis (BPS) that defines a methodological framework for evaluation, calibration, comparison, and context-and data-informed combination of multiple forecast distributions arising from several "competing" models or sources. BPS subsumes and explains– in a formal Bayesian framework– existing density forecasts combination methods, helping to identify strengths and weaknesses of specific approaches, while providing an over-arching framework for forecast synthesis whether density forecasts represent predictions from sets of models, individual forecasters, agencies, or other sources. BPS also defines, as special cases, standard Bayesian model uncertainty analysis using Bayesian model probabilities (and Bayesian model averaging a.k.a. BMA), and admits extension to problems of analysing partial forecast information as well as full forecast distributions.

Following background and discussion of foundational aspects in using multiple dynamic models for forecasting univariate time series via BPS, the lecture focuses on a broad and flexible class of dynamic BPS models for multivariate time series forecast synthesis. Inherently, the theory underlying BPS defines classes of multivariate, dynamic latent factor models in which latent factor processes represent individual models or forecasters. The framework naturally allows modelling and estimation– sequentially and adaptively over time– of varying forecast biases and facets of miscalibration of individual forecast densities, and critically of time-varying inter-dependencies among models or forecasters over multiple series. BPS analysis is developed in one subset of the implied BPS model class– a class of dynamic, seemingly-unrelated, dynamic factor and regression models (DFSURE) in which each forecast model is linked to one set of multivariate dynamic latent factor processes. Bayesian simulation-based computation enables implementation.

Applied highlights come from development and evaluation of sequential BPS analysis in a multiple macroeconomic time series study with US data. The study involves consideration of several dynamic models forecasting monthly data of six key series (inflation, wages, unemployment, consumption, investment and nominal short-term interest rates). Examples highlight questions of forecast synthesis methodology with respect to forecasting goals: interest in 12 or 24 month-ahead forecasting demands– from a formal Bayesian perspective– analysis customized to the horizon. Analysis results sharply highlight this in defining substantially improved forecasting within each horizon based on customized BPS models. The analyses bear out the potential to define fully Bayesian, interpretable models that can (i) adapt to time-varying biases and miscalibration of multiple models or forecasters, (ii) adaptively and practically account for– while generating useful insights into– patterns of time-varying relationships and dependencies among sets of models or forecasters, while (iii) improving forecast accuracy– in some cases, most substantially– for each of several multiple macroeconomic series together, at multiple horizons.

Readings: McAlinn and West (2019); McAlinn et al. (2020); Johnson and West (2018); McAlinn et al. (2018)

References

L. R. Berry and M. West. Bayesian forecasting of many count-valued time series. *Journal of Business and Economic Statistics*, 2019. doi: 10.1080/07350015.2019.1604372. arXiv:1805.05232.

L. R. Berry, P. Helman, and M. West. Probabilistic forecasting of heterogeneous consumer transaction-sales time series. *International Journal of Forecasting*, 2020. doi: 10.1016/j.ijforecast.2019.07.007. arXiv:1808.04698.

C. Carvalho and H. Lopes. Simulation-based sequential analysis of markov switching stochastic volatility models. *Journal of Statistical Planning and Inference*, 51:4526–42, 2007.

C. Carvalho, M. Johannes, H. Lopes, and N. Polson. Particle learning and smoothing. *Statistical Science*, 25: 88–106, 2010.

C. Carvalho, H. Lopes, and R. McCulloch. On the long run volatility of stocks: time-varying predictive systems. *Journal of the American Statistical Association*, 113:1050–1069, 2018.

X. Chen, K. Irie, D. Banks, R. Haslinger, J. Thomas, and M. West. Scalable Bayesian modeling, monitoring and analysis of dynamic network flow data. *Journal of the American Statistical Association*, 113:519–533, 2018. doi: 10.1080/01621459.2017.1345742. arXiv:1607.02655.

X. Chen, D. Banks, and M. West. Bayesian dynamic modeling and monitoring of network flows. *Network Science*, 7:292–318, 2019. doi: 10.1017/nws.2019.10. arXiv:1805.04667.

L. F. Gruber and M. West. GPU-accelerated Bayesian learning in simultaneous graphical dynamic linear models. *Bayesian Analysis*, 11:125–149, 2016. doi: 10.1214/15-BA946.

L. F. Gruber and M. West. Bayesian forecasting and scalable multivariate volatility analysis using simultaneous graphical dynamic linear models. *Econometrics and Statistics*, 3:3–22, 2017. doi: 10.1016/j.ecosta.2017.03.003. arXiv:1606.08291.

M. C. Johnson and M. West. Bayesian predictive synthesis: Forecast calibration and combination. *Technical Report, Department of Statistical Science, Duke University*, 2018. arXiv:1803.01984.

G. Kastner, S. Frühwirth-Schnatter, and H. F. Lopes. Efficient Bayesian inference for multivariate factor stochastic volatility models. *Journal of Computational and Graphical Statistics*, 26:905–917, 2017.

J. Liu and M. West. Combined parameter and state estimation in simulation-based filtering. In A. Doucet, J. F. G. De Freitas, and N. J. Gordon, eds., *Sequential Monte Carlo Methods in Practice*, pp 197–217. Springer, 2001.

H. Lopes and C. Carvalho. Factor stochastic volatility with time varying loadings and markov switching regimes. *Journal of Statistical Planning and Inference*, 137:3082–91, 2007.

H. Lopes and R. Prado. Sequential parameter learning and filtering in structured ar models. *Statistics and Computing*, 23:43–57, 2013.

H. Lopes and R. Tsay. Particle learning for sequential Bayesian computation (with discussion). *Journal of Forecasting*, 30:168–209, 2011.

H. Lopes, E. Salazar, and D. Gamerman. Spatial dynamic factor models. *Bayesian Analysis*, 3:759–92, 2008.

H. Lopes, C. Carvalho, M. Johannes, and N. Polson. Particle learning for sequential Bayesian computation (with discussion). *Bayesian Statistics*, 9:317–360, 2011.

H. Lopes, D. Gamerman, and E. Salazar. Generalized spatial dynamic factor models. *Computational Statistics and Data Analysis*, 55:1319–30, 2011.

H. F. Lopes. Modern Bayesian factor analysis. In Jeliazkov and Yang, eds., *Bayesian Inference in the Social Sciences*, pp 117–158. Wiley, New York, 2014.

H. F. Lopes, R. E. McCulloch, and R. S. Tsay. Parsimony inducing priors for large scale state-space models. Technical report, Booth School of Business, University of Chicago, 2018.

K. McAlinn and M. West. Dynamic Bayesian predictive synthesis in time series forecasting. *Journal of Econometrics*, 210:155–169, 2019. doi: 10.1016/j.jeconom.2018.11.010. arXiv:1601.07463.

K. McAlinn, K. A. Aastveit, and M. West. Bayesian predictive synthesis– discussion of: Using stacking to average Bayesian predictive distributions, by Y. Yao *et al*. *Bayesian Analysis*, 13:971–973, 2018. doi: 10.1214/17-BA1091.

K. McAlinn, K. A. Aastveit, J. Nakajima, and M. West. Multivariate Bayesian predictive synthesis in macroeconomic forecasting. *Journal of the American Statistical Association*, 2020. doi: 10.1080/01621459.2019.1660171. arXiv:1711.01667.

J. Nakajima and M. West. Bayesian analysis of latent threshold dynamic models. *Journal of Business & Economic Statistics*, 31:151–164, 2013. doi: 10.1080/07350015.2012.747847.

J. Nakajima and M. West. Bayesian dynamic factor models: Latent threshold approach. *Journal of Financial Econometrics*, 11:116–153, 2013. doi: 10.1093/jjfinec/nbs013.

J. Nakajima and M. West. Dynamic network signal processing using latent threshold models. *Digital Signal Processing*, 47:6–15, 2015. doi: 10.1016/j.dsp.2015.04.008.

J. Nakajima and M. West. Dynamics and sparsity in latent threshold factor models: A study in multivariate EEG signal processing. *Brazilian Journal of Probability and Statistics*, 31:701–731, 2017. doi: 10.1214/17-BJPS364. arXiv:1606.08292.

R. Prado. Characterization of latent structure in brain signals. In S. Chow, E. Ferrer, and F. Hsieh, eds., *Statistical Methods for Modeling Human Dynamics: An Interdisciplinary Dialogue*, pp 123–153. Psychology Press, Taylor and Francis Group, New York, 2009. ISBN 9781848728257.

R. Prado. Multi-state models for mental fatigue. In A. O'Hagan and M. West, eds., *The Handbook of Applied Bayesian Analysis*, pp 845–874. Oxford University Press, Oxford, 2010.

R. Prado and M. West. *Time Series: Modelling, Computation & Inference*. Chapman & Hall/CRC Press, 2010.

R. Prado, M. West, and A. D. Krystal. Multi-channel EEG analyses via dynamic regression models with time-varying lag/lead structure. *Journal of the Royal Statistical Society, Series C*, 50:95–109, 2001.

S. Shirota, Y. Omori, H. Lopes, and H. Piao. Cholesky realized stochastic volatility model. *Econometrics and Statistics*, 3:34–59, 2017.

M. West. Bayesian dynamic modelling. In P. Damien, P. Dellaportes, N. G. Polson, and D. A. Stephens, eds., *Bayesian Theory and Applications*, pp 145–166. Clarendon: Oxford University Press, 2013.

M. West. Bayesian forecasting of multivariate time series: Scalability, structure uncertainty and decisions (with discussion). *Annals of the Institute of Statistical Mathematics*, 72:1–44, 2020. doi: 10.1007/s10463-019-00741-3. arXiv:1911.09656. Published online Dec 9 2019.

M. West and P. J. Harrison. *Bayesian Forecasting & Dynamic Models*. Springer, 2nd edition, 1997.

Z. Y. Zhao, M. Xie, and M. West. Dynamic dependence networks: Financial time series forecasting & portfolio decisions (with discussion). *Applied Stochastic Models in Business and Industry*, 32:311–339, 2016. doi: 10.1002/asmb.2161. arXiv:1606.08339.

X. Zhou, J. Nakajima, and M. West. Bayesian forecasting and portfolio decisions using dynamic dependent sparse factor models. *International Journal of Forecasting*, 30:963–980, 2014. doi: http://dx.doi.org/10.1016/j.ijforecast.2014.03.017.