International Journal of Forecasting 30 (2014) 963–980

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Bayesian forecasting and portfolio decisions using dynamic dependent sparse factor models

Xiaocong Zhou, Jouchi Nakajima*, Mike West

Duke University, Durham NC, USA

ARTICLE INFO

Keywords:
Bayesian forecasting
Benchmark neutral portfolio
Dynamic factor models
Latent threshold dynamic models
Multivariate stochastic volatility
Portfolio optimization
Sparse time-varying loadings

ABSTRACT

We extend the recently introduced *latent threshold dynamic models* to include dependencies among the dynamic latent factors which underlie multivariate volatility. With an ability to induce *time-varying sparsity* in factor loadings, these models now also allow time-varying correlations among factors, which may be exploited in order to improve volatility forecasts. We couple multi-period, out-of-sample forecasting with portfolio analysis using standard and novel *benchmark neutral* portfolios. Detailed studies of stock index and FX time series include: multi-period, out-of-sample forecasting, statistical model comparisons, and portfolio performance testing using raw returns, risk-adjusted returns and portfolio volatility. We find uniform improvements on all measures relative to standard dynamic factor models. This is due to the parsimony of latent threshold models and their ability to exploit between-factor correlations so as to improve the characterization and prediction of volatility. These advances will be of interest to financial analysts, investors and practitioners, as well as to modeling researchers.

1. Introduction

Since the early Bayesian approaches to factor volatility modeling (e.g., Aguilar, Prado, Huerta, & West, 1999; Aguilar & West, 2000; Pitt & Shephard, 1999), there has been an increasing interest in refined models, based on their practical benefits in financial studies in particular (e.g., Quintana, Carvalho, Scott, & Costigliola, 2010; Quintana, Lourdes, Aguilar, & Liu, 2003). While the original approaches assumed constant factor loadings and no time dependence of the latent factors for financial returns series, recent extensions have introduced short-term time series models for factor loadings (e.g., Carvalho, Lopes, & Aguilar, 2011). To date, little has been discussed about dependencies among factor processes, due primarily to the adoption of identifying constraints under which independent factor processes are mandated. With the increasing interest in sparse factor models – models in which multiple factor loadings are zero over some periods of time – this has changed: such models now allow for dependencies among latent factor processes, and our main modeling goal here is to develop and exploit this in forecasting and portfolio decisions.

We achieve these developments in analyses of dynamic factor models using *latent thresholding*, an idea and methodology which was recently introduced and developed theoretically by Nakajima and West (2013a), with applications to dynamic regression and time-varying VAR models. A follow-up application (see Nakajima & West, 2013b) added dynamic sparsity to traditional factor models; the current paper extends this with the development of *dependent factor model* structures, novel portfolio constructions, and their embedding in a complete analysis and forecasting system. In the applied studies of the paper, this is shown to be of quite substantial benefit in terms

* Corresponding author.
E-mail address: jouchi.nakajima@boj.or.jp (J. Nakajima).

http://dx.doi.org/10.1016/j.ijforecast.2014.03.017

of improved forecasting performances and portfolio decision outcomes, as well as in improved model fits on purely statistical criteria. We develop the applied examples using a range of stylized Bayesian portfolio decision constructs, and, as part of this, introduce a novel strategy that explicitly integrates a *benchmark neutral* strategy into a more or less standard portfolio optimization. In addition to demonstrating the ability of sparse, dependent factor models to outperform standard models under this and other portfolio rules, this development will also be of interest to forecasters and financial decision makers in other contexts.

This work contributes modeling, forecasting and decision analytic advances to the growing body of literature on dynamic factor approaches to time series analysis. Beginning with earlier developments of dynamic factor models in time series (e.g., Geweke & Zhou, 1996; Peña & Box, 1987; Stock & Watson, 1989), these approaches have become popular in macroeconomic applications (e.g., Aruoba, Diebold, & Scotti, 2009; Bai & Ng, 2006; Del Negro & Otrok, 2008; Forni & Gambetti, 2010; Forni, Hallin, Lippi, & Reichlin, 2000; Koop & Potter, 2004; Stock & Watson, 2002), as well as in financial applications in which multivariate volatility is represented in factor structures and other forms (e.g., Aguilar & West, 2000; Asai, McAleer, & Yu, 2006; Chib, Nardari, & Shephard, 2006; Doz & Renault, 2006; Fan, Fan, & Lv, 2008; Han, 2005; Harvey, Ruiz, & Shephard, 1994; Philipov & Glickman, 2006; Pitt & Shephard, 1999; Yu & Meyer, 2006). Recent developments in time-varying factor loadings models that provide part of the foundation for our work here have been noted particularly for the forecasting and statistical improvements they can generate (e.g., Del Negro & Otrok, 2008; Lopes & Carvalho, 2007). Our work builds on structural and dynamic model concepts from these areas, introducing dynamic, sparse factor models with dependencies among latent factor processes that are shown to be able to provide substantial additional improvements in model fit, forecasting and portfolio decisions.

The remainder of the paper is organized as follows. Section 2 summarizes the standard framework of dynamic factor models. Section 3 discusses model identification, sparse dynamic factors and the key rationale for dependent factor models. Section 4 discusses the latent thresholding concept and its application to factor models. Section 5 summarizes the new class of dynamic sparse factor models, with time-varying volatility matrices allowing correlated factors. Sections 6 and 7 discuss analysis, model comparison, forecasting, and portfolio decisions in two case studies: a 10-dimensional stock price index time series, and a 20-dimensional FX time series. Some summary comments appear in Section 8. An Appendix briefly outlines the Bayesian Markov chain Monte Carlo computational method for model fitting; this links to more extensive technical details in prior publications for interested readers, as well as to software.

Some notation. We use the distributional notation $\boldsymbol{y} \sim N(\boldsymbol{a}, \boldsymbol{A}), d \sim U(a, b), p \sim B(a, b), v \sim G(a, b)$, for the normal, uniform, beta, and gamma distributions, respectively. We use $\text{diag}(a_1, \ldots, a_k)$ to refer to a diagonal matrix whose diagonal elements are the arguments. We also use $s : t$ to denote $s, s+1, \ldots, t$ when $s < t$, for succinct subscripting; e.g., $\boldsymbol{y}_{1:T}$ denotes $\{\boldsymbol{y}_1, \ldots, \boldsymbol{y}_T\}$.

2. Basic setting and background: traditional dynamic factor models

2.1. Basic model context

We begin with traditional dynamic factor models with time-varying factor loadings and volatility components, as follows. The $m \times 1$ vector response time series $\boldsymbol{y}_t$ $(t = 1, 2, \ldots)$ follows

$$\boldsymbol{y}_t = \boldsymbol{c}_t + \boldsymbol{B}_t \boldsymbol{f}_t + \boldsymbol{\nu}_t, \quad \boldsymbol{\nu}_t \sim N(\boldsymbol{0}, \boldsymbol{\Sigma}_t), \tag{1}$$

$$\boldsymbol{f}_t = \boldsymbol{G}\boldsymbol{f}_{t-1} + \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim N(\boldsymbol{0}, \boldsymbol{\Upsilon}_t), \tag{2}$$

where

- $\boldsymbol{c}_t = (c_{1t}, \ldots, c_{mt})'$ is the $m \times 1$ time-varying local mean at time t;
- $\boldsymbol{f}_t = (f_{1t}, \ldots, f_{kt})'$ is a $k \times 1$ vector of latent factors evolving according to a VAR(1) model with a diagonal $(k \times k)$ AR coefficient matrix $\boldsymbol{G} = \text{diag}(\gamma_1, \ldots, \gamma_k)$;
- $\boldsymbol{\varepsilon}_t = (\varepsilon_{1t}, \ldots, \varepsilon_{kt})'$ is a $k \times 1$ vector of factor innovations with time-varying variance matrix $\boldsymbol{\Upsilon}_t$ containing elements υ_{ijt};
- $\boldsymbol{B}_t$ is the $m \times k$ time-varying factor loadings matrix; and
- $\boldsymbol{\nu}_t = (\nu_{1t}, \ldots, \nu_{mt})'$ is a $m \times 1$ residual vector with a diagonal time-varying volatility matrix $\boldsymbol{\Sigma}_t = \text{diag}(\sigma^2_{1t}, \ldots, \sigma^2_{mt})$.

The $\boldsymbol{\varepsilon}_s$ and $\boldsymbol{\nu}_t$ sequences are independent and mutually independent. Eqs. (1) and (2) define a broad class of dynamic latent factor models, and variants of the models are routinely applied to financial time series. The AR coefficient matrix $\boldsymbol{G}$ in the factor evolution model is assumed to be constant and diagonal; this could also be relaxed for other applications in which a persistent but time-varying matrix may be of interest, although the applications here do not suggest such an extension for the current analyses.

To ensure the mathematical identification of factor models, and as a matter of modeling choice, we use the traditional lower triangular constraint on the dynamic factor loadings matrix process $\boldsymbol{B}_t$ (e.g., Aguilar & West, 2000; Geweke & Zhou, 1996; Lopes & West, 2004). Noting that $\boldsymbol{B}_t$ is "tall and skinny" – that is, the number of factors k will typically be far less than the number of series m – the upper triangular elements are $b_{ijt} = 0$ for $k \geq j > i \geq 1$, and the main diagonal elements are $b_{iit} = 1$ for $i = 1 : k$. Importantly, the traditional model has a *diagonal* factor innovations volatility matrix $\boldsymbol{\Upsilon}_t$, while admitting volatility models for the diagonal elements.

2.2. AR models for dynamic parameter processes

Complete model specification requires specific structures for the time-varying parameter processes $\boldsymbol{c}_t$ and $\boldsymbol{B}_t$, and the diagonal $\boldsymbol{\Sigma}_t$ and $\boldsymbol{\Upsilon}_t$. The simplest and most widely used are basic AR(1) models for univariate parameters, as follows.

Factor loadings $\boldsymbol{B}_t$. Univariate AR(1) models for univariate factor loadings have become increasingly popular in the literature (e.g., Del Negro & Otrok, 2008; Lopes & Carvalho, 2007). For each $i = 2 : m$ and $j < i$, denote by β_{ijt} the loading relating series (row) i to factor (column) j in $\boldsymbol{B}_t$, recalling that the upper right triangle elements are zero