ARTICLE IN PRESS

Journal of Econometrics ▮ (▮▮▮▮) ▮▮▮–▮▮▮

Contents lists available at ScienceDirect

Journal of Econometrics

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

A spectral EM algorithm for dynamic factor models☆

Gabriele Fiorentini [a,*], Alessandro Galesi [b], Enrique Sentana [c]

[a] Università di Firenze and RCEA, Viale Morgagni 59, I-50134 Firenze, Italy
[b] Banco de España, Calle Alcalá 48, E-28014 Madrid, Spain
[c] CEMFI, Casado del Alisal 5, E-28014 Madrid, Spain

ARTICLE INFO

Article history:
Available online xxxx

JEL classification:
C32
C38
C51

Keywords:
Indirect inference
Kalman filter
Sectoral employment
Spectral maximum likelihood
Wiener–Kolmogorov filter

ABSTRACT

We make two complementary contributions to efficiently estimate dynamic factor models: a frequency domain EM algorithm and a swift iterated indirect inference procedure for ARMA models with no asymptotic efficiency loss for any finite number of iterations. Although our procedures can estimate such models with many series without good initial values, near the optimum we recommend switching to a gradient method that analytically computes spectral scores using the EM principle. We successfully employ our methods to construct an index that captures the common movements of US sectoral employment growth rates, which we compare to the indices obtained by semiparametric methods.

1. Introduction

Dynamic factor models have been extensively used in macroeconomics and finance since their introduction by Geweke (1977) and Sargent and Sims (1977) as a way of capturing the cross-sectional and dynamic correlations between multiple series in a parsimonious way. A far from comprehensive list of early and more recent applications include not only business cycle analysis (see Litterman and Sargent (1979), Stock and Watson (1989, 1993)), Diebold and Rudebusch (1996) or Gregory et al. (1997) and bond yields (Singleton (1981), Jegadeesh and Pennacchi (1996), Dungey et al. (2000) or Diebold et al. (2006)), but also wages (Engle and Watson, 1981), employment (Quah and Sargent, 1993), commodity prices (Peña and Box, 1987) and financial contagion (Mody and Taylor, 2007).

An expanding, influential body of literature has shown that one may accurately recover the unobserved series by using the frequency domain version of principal components put forward by Brillinger (1981, ch. 9) and further extended by Forni et al. (2000) (FHLR), which is based on a non-parametric estimate of the spectral density matrix of the observed series (see Forni et al. (2015) for more recent developments). In fact, it might even be possible to use static principal components if certain additional assumptions hold (see Bai and Ng (2008)). Aside from avoiding the numerical optimisation of a criterion

☆ We are grateful to Dante Amengual and Domenico Giannone, as well as to audiences at the Conference on Indirect Estimation Methods in Finance and Economics (Konstanz, 2014), the Econometric Methods for Banking and Finance Conference (Bank of Portugal, 2014), the NBER-NSF Time Series Conference (Federal Reserve Bank of St. Louis, 2014), the Advances in Econometrics Conference on Dynamic Factor Models (Aarhus, 2014), the EC^2 Advances in Forecasting Conference (UPF, 2014), the Italian Congress of Econometrics and Empirical Economics (Salerno, 2015) and the V Workshop in Time Series Econometrics (Zaragoza, 2015) for helpful comments, discussions and suggestions. The requests of the editors and two anonymous referees have also substantially improved the paper. Of course, the usual caveat applies. The views expressed herein are those of the authors and not necessarily reflect the opinion of the Bank of Spain or the Eurosystem. Financial support from the project MIUR PRIN MISURA 2013-15 (Fiorentini) and the Spanish Ministries of Science & Innovation and Economy & Competitiveness through grants ECO 2008-00280 and ECO 2014-59262 (Sentana) is also gratefully acknowledged.

* Corresponding author.
E-mail addresses: fiorentini@disia.unifi.it (G. Fiorentini), alessandro.galesi@bde.es (A. Galesi), sentana@cemfi.es (E. Sentana).

https://doi.org/10.1016/j.jeconom.2018.03.013

function, the main advantage of such methods is that they remain valid in the presence of some mild contemporaneous and dynamic correlation between idiosyncratic terms when the cross-sectional dimension, N, is commensurate with the time series dimension, T.

There are two closely related issues, though. First, the cross-sectional asymptotic boundedness conditions on the eigenvalues of the autocovariance matrices of the idiosyncratic terms underlying the approximate factor structures originally suggested by Chamberlain and Rothschild (1983) are largely meaningless in empirical situations in which N is small relative to T. And second, although the factors could be regarded as a set of parameters in any given realisation, efficiency considerations indicate that a signal extraction approach which treats the underlying latent variables as stochastic processes would be more appropriate for such data sets. In addition, Doz et al. (2012) have recently closed the gap between the two strands of the literature by proving the NT-consistency of the Gaussian pseudo ML estimators (MLE) of exact versions of dynamic factor models even when not all the maintained assumptions hold.

In principle, Gaussian (P)MLEs of the parameters can be obtained from the usual time domain version of the log-likelihood function computed as a by-product of the Kalman filter prediction equations or from Whittle's (1962) frequency domain asymptotic approximation. Further, once the parameters have been estimated the Kalman smoother and its Wiener–Kolmogorov counterpart provide optimally filtered estimates of the latent factors. These estimation and filtering issues are well understood (see e.g. Harvey (1989)), and the same can be said of their numerical implementation (see Jungbacker and Koopman (2015)). In practice, though, researchers avoid ML except in relatively small models because of the heavy computational burden involved, which is disproportionately larger as the number of series considered increases.

To ameliorate this problem, Shumway and Stoffer (1982), Watson and Engle (1983) and Quah and Sargent (1993) applied the *Expectation–Maximisation* (EM) algorithm in Dempster et al. (1977) to the time domain versions of these models, thereby avoiding the computation of the likelihood function and its score. This iterative algorithm has been popular in various areas of statistics and econometrics when the data set is incomplete or contains missing values, or the model can be posed in a similar form, such as in the finite mixture models studied by Arcidiacono and Jones (2003) or the dynamic Markov switching models considered by Hamilton (1990) and Fruhwirth-Schnatter (2007). Its popularity can be attributed mainly to the efficiency of the procedure, as measured by its speed, and also to the generality of the approach and its convergence properties (see Ruud (1991) for an elegant review of this method and McLachlan and Krishnan (1996) for a more thorough analysis).

However, the time domain version of the EM algorithm has only been derived for dynamic factor models in which the latent variables follow pure Ar processes (see again Doz et al. (2012)), and works best when the effects of the common factors on the observed variables are contemporaneous, which substantially limits the class of models to which it can be successfully applied. In particular, it excludes models in which either common or idiosyncratic factors follow Arma processes. As is well known, such processes combine autoregressive and moving average components in a rather parsimonious way, and for that reason they are by far the most common approximations used to (Wold) represent univariate stationary series. However, while Ar processes often arise as difference equation-type representations of natural phenomena, the presence of Ma components is sometimes justified as a result of contemporaneous aggregation of several underlying components. On this basis, one might argue that there is no need to introduce Ma terms in dynamic factor models.

Nevertheless, there are at least three compelling reasons for considering Arma processes for common or idiosyncratic factors. First, the temporal aggregation results in Bergstrom (1984) imply that discrete time observations will often contain Ma components even if the underlying continuous time processes follow first-order stochastic difference equations. Obviously, the same applies to the increasingly popular continuous time versions of Arma models (see Chambers and Thorton (2012)). Second, the usual deseasonalisation procedures employed by the national statistical offices imply a transfer function that substantially dampens the spectral density at high frequencies, which pure Ar models struggle to capture (see Maravall (1993)). Given that those filters tend to be very similar for closely related homogeneous series, they are likely to introduce Ma terms not only in the specific component of each series but also in the common one. Finally, recent macroeconomic applications of dynamic factor models have often considered specifications in which the lagged latent variables appear as additional factors (see again Bai and Ng (2008) and the references therein). In those circumstances, plausible cross-sectional restrictions on the dynamic factor loadings made for parsimony reasons also introduce Ma terms in the common factors.

In the context of general dynamic factor models with latent Arma processes, we make two independent but complementary contributions. First, we introduce a frequency domain version of the EM algorithm, which exploits the heteroskedastic factor structure of the spectral density matrix of the model to carry out the *expectation* stage very quickly. Nevertheless, a standard implementation of our algorithm would still require $O(N)$ numerical optimisations at each *maximisation* stage in models with idiosyncratic Ma components. For that reason, our second contribution is a very fast iterated indirect inference procedure for estimating the parameters of univariate Arma models, which is based on a sequence of simple auxiliary OLS regressions of certain filtered series. Importantly, unlike existing indirect inference procedures for those models, our proposed estimator entails no asymptotic efficiency loss for any finite number of iterations. Further, it will generally coincide with the ML estimator in the limit.

The complementary between our proposals is twofold: (i) our iterated indirect inference method can be implemented far more quickly in the frequency domain than in the time domain; and (ii) it can be very easily adapted to deal with latent variables with only minor modifications.

Our combined proposal, differs from more standard applications of indirect estimation methods to factor models, which typically rely on a simpler approximating model as auxiliary model. For example, in the case of a single factor with static loadings we could fit univariate Arma models to the first principal component of the observed series, as well as to the