Journal of Econometrics 177 (2013) 289–304

Contents lists available at ScienceDirect

Journal of Econometrics

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

Consistent factor estimation in dynamic factor models with structural instability☆

Brandon J. Bates [a], Mikkel Plagborg-Møller [b], James H. Stock [b], Mark W. Watson [c,*]

[a] BlackRock, Inc., United States
[b] Harvard University, United States
[c] Princeton University, United States

ARTICLE INFO

Article history:
Available online 17 April 2013

ABSTRACT

This paper considers the estimation of approximate dynamic factor models when there is temporal instability in the factor loadings. We characterize the type and magnitude of instabilities under which the principal components estimator of the factors is consistent and find that these instabilities can be larger than earlier theoretical calculations suggest. We also discuss implications of our results for the robustness of regressions based on the estimated factors and of estimates of the number of factors in the presence of parameter instability. Simulations calibrated to an empirical application indicate that instability in the factor loadings has a limited impact on estimation of the factor space and diffusion index forecasting, whereas estimation of the number of factors is more substantially affected.

1. Introduction

Dynamic factor models (DFMs) provide a flexible framework for simultaneously modeling a large number of macroeconomic time series.[1] In a DFM, a potentially large number of observed time series variables are modeled as depending on a small number of unobserved factors, which account for the widespread co-movements of the observed series. Although there is now a large body of theory for the analysis of high-dimensional DFMs, nearly all of this theory has been developed for the case in which the DFM parameters are stable, in particular, in which there are no changes in the factor loadings (the coefficients on the factors); among the few exceptions are Stock and Watson (2002, 2009) and Breitung and Eickmeier (2011). This assumption of parameter stability is at odds with broad evidence of time variation in many macroeconomic forecasting relations. Recently, a number of empirical DFM papers have explicitly allowed for structural instability, e.g., Banerjee et al. (2008), Stock and Watson (2009), Eickmeier et al. (2011) and Korobilis (2013). However, theoretical guidance remains scant.

The goal of this paper is to characterize the type and magnitude of parameter instability that can be tolerated by a standard estimator of the factors, the principal components estimator, in a DFM when the coefficients of the model are unstable. In so doing, this paper contributes to a larger debate about how best to handle the instability that is widespread in macroeconomic forecasting relations. On the one hand, the conventional wisdom is that time series forecasts deteriorate when there are undetected structural breaks or unmodeled time-varying parameters, see for example Clements and Hendry (1998). This view underlies the large literatures on the detection of breaks and on models that incorporate breaks and time variation, for example by modeling the breaks as following a Markov process (Hamilton, 1989; Pesaran et al., 2006). In the context of DFMs, Breitung and Eickmeier (2011) show that a one-time structural break in the factor loadings has the effect of introducing new factors, so that estimation of the factors ignoring the break leads to estimating too many factors.

On the other hand, a few recent papers have provided evidence that sometimes it can be better to ignore parameter instability when forecasting. Pesaran and Timmermann (2005) point out that whether to use pre-break data for estimating an autoregression trades off an increase in bias against a reduction in estimator variance, and they supply empirical evidence supporting the use of

☆ We thank Gary Chamberlain, Herman van Dijk, Anna Mikusheva, Allan Timmermann and two anonymous referees for helpful comments.

* Corresponding author.
E-mail addresses: brandon.bates@blackrock.com (B.J. Bates), plagborg@fas.harvard.edu (M. Plagborg-Møller), james_stock@harvard.edu (J.H. Stock), mwatson@princeton.edu (M.W. Watson).

[1] The early work on DFMs considered a small number of time series. DFMs were introduced by Geweke (1977), and early low-dimensional applications include Sargent and Sims (1977), Engle and Watson (1981), Watson and Engle (1983), Sargent (1989) and Stock and Watson (1989). Work over the past fifteen years has focused on methods that facilitate the analysis of a large number of time series, see Forni et al. (2000) and Stock and Watson (2002) for early contributions. For recent contributions and discussions of this large literature see Bai and Ng (2008), Eickmeier and Ziegler (2008), Chudik and Pesaran (2011) and Stock and Watson (2011).

0304-4076/$ – see front matter
http://dx.doi.org/10.1016/j.jeconom.2013.04.014

pre-break data for forecasting. Pesaran and Timmermann (2007) develop tools to help ascertain in practice whether pre-break data should be used for estimation of single-equation time series forecasting models. In DFMs, Stock and Watson (2009) provide an empirical example using US macroeconomic data from 1960–2007 in which full-sample estimates of the factors are preferable to subsample estimates, despite clear evidence of a break in many factor loadings around the beginning of the Great Moderation in 1984.

We therefore seek a precise theoretical understanding of the effect of instability in the factor loadings on the performance of principal components estimators of the factors. Specifically, we consider a DFM with N variables observed for T time periods and $r \ll N$ factors, where the $N \times r$ matrix of dynamic factor loadings Λ can vary over time. We write this time variation so that Λ at date t equals its value at date 0, plus a deviation; that is, $\Lambda_t = \Lambda_0 + h_{NT}\xi_t$. The term ξ_t is a possibly random disturbance, and h_{NT} is a deterministic scalar sequence in N and T which governs the scale of the deviation. Using this framework and standard assumptions in the literature (Bai and Ng, 2002, 2006a), we obtain general conditions on h_{NT} under which the principal components estimates are mean square consistent for the space spanned by the true factors. We then specialize these general results to three leading cases: i.i.d. deviations of Λ_t from Λ_0, random walk deviations that are independent across series, and an arbitrary one-time break that affects some or all of the series.

For the case in which Λ_t is a vector of independent random walks, Stock and Watson (2002) showed that the factor estimates are consistent if $h_{NT} = O(T^{-1})$. By using a different method of proof (which builds on Bai and Ng, 2002), we are able to weaken this result considerably and show that the estimated factors are consistent if $h_{NT} = o(T^{-1/2})$. We further show that, if $h_{NT} = O(1/\min\{N^{1/4}T^{1/2}, T^{3/4}\})$, the estimated factors achieve the mean square consistency rate of $1/\min\{N, T\}$, a rate initially established by Bai and Ng (2002) in the case of no time variation. Because the elements of ξ_t in the random walk case are themselves $O_p(t^{1/2})$, this means that deviations in the factor loadings on the order of $o_p(1)$ do not break the consistency of the principal components estimator. These rates are remarkable: as a comparison, if the factors were observed so an efficient test for time variation could be performed, the test would have nontrivial power against random walk deviations in a $h_{NT} \propto T^{-1}$ neighborhood of zero (e.g., Stock and Watson, 1998b) and would have power of one against parameter deviations of the magnitude tolerated by the principal components estimator. Intuitively, the reason that the principal components estimator can handle such large changes in the coefficients is that, if these shifts have limited dependence across series, their effect can be reduced, and eliminated asymptotically, by averaging across series.

We further provide the rate of mean square consistency as a function of h_{NT}, both in general and specialized to the random walk case. The resulting consistency rate function is nonlinear and reflects the tradeoff between the magnitude of the instability and, through the relative rate N/T as T increases, the amount of cross-sectional information that can be used to"average out" this instability. To elaborate on the practical implications of the theory, we conduct a simulation study calibrated to the Stock and Watson (2009) dataset. The results confirm that the principal components estimator and derived diffusion index forecasts are robust to empirically relevant degrees of temporal instability in the factor loadings, although the precise quantitative conclusions depend on the assumed type of structural instability and the persistence of the factors. Interestingly, the robustness obtains even though the Bai and Ng (2002) information criterion estimator of the rank of the factor space appears to be asymptotically biased for some of our parametrizations.

The rest of the paper proceeds as follows. Section 2 lays out the model, the assumptions, and the three special cases. Our main result on consistency of the principal components estimator is presented in Section 3. Rank selection and diffusion index forecasting are discussed in Section 4. Section 5 provides Monte Carlo results, and Section 6 concludes.

2. Model and assumptions

2.1. Basic model and intuition

The model and notation follow Bai and Ng (2002) closely. Denote the observed data by X_{it} for $i = 1, \ldots, N$, $t = 1, \ldots, T$. It is assumed that the observed series are driven by a small, fixed number r of unobserved common factors F_{pt}, $p = 1, \ldots, r$, such that

$$X_{it} = \lambda_{it}' F_t + e_{it}.$$

Here $\lambda_{it} \in \mathbb{R}^r$ is the possibly time-varying factor loading of series i at time t, $F_t = (F_{1t}, \ldots, F_{rt})'$, and e_{it} is an idiosyncratic error. Define vectors $X_t = (X_{1t}, \ldots, X_{Nt})'$, $e_t = (e_{1t}, \ldots, e_{Nt})'$, $\Lambda_t = (\lambda_{1t}, \ldots, \lambda_{Nt})'$ and data matrices $X = (X_1, \ldots, X_T)'$, $F = (F_1, \ldots, F_T)'$. The initial factor loadings Λ_0 are fixed. We write the cumulative drift in the parameter loadings as

$$\Lambda_t - \Lambda_0 = h_{NT}\xi_t,$$

where h_{NT} is a deterministic scalar that may depend on N and T, while $\{\xi_t\}$ is a possibly degenerate random process of dimension $N \times r$, $\xi_t = (\xi_{1t}, \ldots, \xi_{Nt})'$ (in fact, it will be allowed to be a triangular array). Observe that

$$X_t = \Lambda_t F_t + e_t = \Lambda_0 F_t + e_t + w_t, \quad (1)$$

where $w_t = h_{NT}\xi_t F_t$. Our proof technique will be to treat w_t as another error term in the factor model.[2]

To establish some intuition for why estimation of the factors is possible despite structural instability, let the number of factors be $r = 1$ and consider an independent random walk model for the time variation in the factor loadings, so that $\xi_{it} = \xi_{i,t-1} + \zeta_{it}$, where ζ_{it} is i.i.d. across i and t with mean 0 and variance σ_ζ^2, and suppose that Λ_0 is known. In addition, we look ahead to Assumption 2 and assume that $\Lambda_0'\Lambda_0/N \to D > 0$. Because Λ_0 is known, we can consider the estimator $\hat{F}_t(\Lambda_0) = (\Lambda_0'\Lambda_0)^{-1}\Lambda_0' X_t$. From (1),

$$\hat{F}_t(\Lambda_0) = F_t + (\Lambda_0'\Lambda_0)^{-1}\Lambda_0' e_t + (\Lambda_0'\Lambda_0)^{-1}\Lambda_0' w_t,$$

so

$$\hat{F}_t(\Lambda_0) - F_t \approx D^{-1}N^{-1}\sum_{i=1}^{N}\lambda_{i0}e_{it} + D^{-1}N^{-1}\sum_{i=1}^{N}\lambda_{i0}w_{it}.$$

The first term does not involve time-varying factor loadings and under limited cross-sectional dependence it is $O_p(N^{-1/2})$. Using the definition of w_t, the second term can be written

$$D^{-1}N^{-1}\sum_{i=1}^{N}\lambda_{i0}w_{it} = D^{-1}\left(h_{NT}N^{-1}\sum_{i=1}^{N}\lambda_{i0}\xi_{it}\right)F_t.$$

Since F_t is $O_p(1)$, this second term is the same order as the first, $O_p(N^{-1/2})$, if $h_{NT}N^{-1}\sum_{i=1}^{N}\lambda_{i0}\xi_{it}$ is $O_p(N^{-1/2})$. Under the independent random walk model, $\xi_{it} = O_p(T^{1/2})$, so

$$h_{NT}N^{-1}\sum_{i=1}^{N}\lambda_{i0}\xi_{it} = O_p(h_{NT}(T/N)^{1/2}),$$

[2] As pointed out by our referees, a straightforward approach would be to treat $e_t^* = e_t + w_t$ as a catch-all error term and provide conditions on h_{NT} and ξ_t such that e_t^* satisfies Assumption C in Bai and Ng (2002). Some of the examples below could be handled this way. However, in the case of random walk factor loadings, applying the Bai and Ng assumption to e_t^* would restrict the temporal dependence of ξ_t more severely than required by our Theorem 1 (cf. Assumption 3.2 below).