Journal of Econometrics 180 (2014) 30–48

Contents lists available at ScienceDirect

Journal of Econometrics

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

Detecting big structural breaks in large factor models[☆]

Liang Chen[a], Juan J. Dolado[b,c], Jesús Gonzalo[b,*]

[a] Department of Economics, Nuffield College and Institute of New Economic Thinking, University of Oxford, United Kingdom
[b] Department of Economics, Universidad Carlos III de Madrid, Spain
[c] European University Institute, Italy

ARTICLE INFO

Article history:
Received 9 May 2012
Received in revised form
3 October 2013
Accepted 21 January 2014
Available online 31 January 2014

JEL classification:
C12
C33

Keywords:
Structural break
Large factor model
Factor loadings
Principal components

ABSTRACT

Time invariance of factor loadings is a standard assumption in the analysis of large factor models. Yet, this assumption may be restrictive unless parameter shifts are mild (i.e., local to zero). In this paper we develop a new testing procedure to detect *big* breaks in these loadings at either known or unknown dates. It relies upon testing for parameter breaks in a regression of one of the factors estimated by Principal Components analysis on the remaining estimated factors, where the number of factors is chosen according to Bai and Ng's (2002) information criteria. The test fares well in terms of power relative to other recently proposed tests on this issue, and can be easily implemented to avoid forecasting failures in standard factor-augmented (FAR, FAVAR) models where the number of factors is a priori imposed on the basis of theoretical considerations.

1. Introduction

Despite being well acknowledged that some parameters in economic relationships can be subject to important structural breaks (e.g., those related to technological change, globalization or strong policy reforms), a standard practice in the estimation of large factor models (FM, hereafter) has been to assume time-invariant factor loadings. Possibly, one of the main reasons for this benign neglect of breaks stems from the important results obtained by Stock and Watson (2002, 2009) regarding the consistency of the estimated factors by principal components analysis (PCA hereafter) when the loadings are subject to small (i.e., local-to-zero) instabilities. These authors conclude that the failure of factor-based forecasts is mainly due to the instability of the forecast function, rather than of the different components of the FM. As a result, their advice is to use full-sample factor estimates and subsample forecasting equations to improve forecasts.

However, the main emphasis placed on local-to-zero breaks has been subsequently questioned. For example, by means of a Monte Carlo study, Banerjee et al. (2008) conclude that, in contrast to Stock and Watson's diagnosis, the instability of factor loadings when big (i.e., not local-to-zero) breaks occur is the most likely reason behind the worsening factor-based forecasts, particularly in small samples. Likewise, when discussing Stock and Watson's research on this topic, Giannone (2007) argues that "*…to understand structural changes we should devote more effort in modelling the variables characterized by more severe instabilities…*". In this paper, we pursue this goal by providing a precise characterization of the different conditions under which big and small breaks in the factor loadings may occur, as well as develop a simple test to distinguish between them. We conclude that, in contrast to small breaks, big breaks should not be ignored since they may lead to misleading results in standard econometric practices using FM and in the potential interpretation of the factors themselves.

A forerunner of our paper is Breitung and Eickmeier (2011, BE henceforth) who were the first to propose a proper testing procedure to detect big breaks in the factor loadings. Their test relies on the idea that, under the null of no structural break (plus some additional assumptions), the estimation error of the factors can be ignored and thus the estimated factors can be treated as the true ones. Consequently, a Chow-type test can be implemented

☆ We are grateful to a co-editor, two anonymous referees, Dante Amengual, Manuel Arellano, Stephane Bonhomme and Soren Johansen for insightful comments, and to Atsushi Inoue for helpful discussions on his related but independent research. We also thank seminar participants at Aarhus, CEMFI, Conference in Honor of Sir David F. Hendry (St. Andrews), ESEM 2011 (Oslo), and Workshops on High-Dimensional Econometric Modelling (Cass Business School) and Time Series Econometrics (Zaragoza) for suggestions. Financial support from the Spanish Ministerio de Economía y Competitividad (grant ECO2010-19357), Comunidad de Madrid (grant Excelecon), Bank of Spain (grant ER program), Open Society Foundations and the Oxford Martin School, is gratefully acknowledged.

* Corresponding author. Tel.: +34 916249853.
E-mail addresses: liang.chen@economics.ox.ac.uk (L. Chen), dolado@eco.uc3m.es (J.J. Dolado), jesus.gonzalo@uc3m.es (J. Gonzalo).

http://dx.doi.org/10.1016/j.jeconom.2014.01.006
0304-4076/

by regressing each variable in the data set on both the estimated factors using the whole sample and their truncated versions from the date of the break onwards. Focusing on the joint statistical significance of the estimated coefficients on the truncated factors, their test compares the empirical rejection frequency among the individual regressions to a nominal size of 5% under the null. In our view, this approach suffers from two limitations: (i) the overall limiting distribution of their test remains unknown when testing for the equality of all the elements of the loading matrix in both subsamples[1]; and (ii) it lacks non-trivial power when the number of factors is chosen according to some consistent estimator of r. This last problem can be serious. For example, as explained further below, a FM with r original factors where the loadings of one of them exhibit a big structural break at the same date admits a standard factor representation with $r+1$ factors without a break. Hence, if the number of factors were to be chosen as $r+1$, instead of r, their testing approach would fail to detect any break at all when in fact there is one.

Our contribution here is to propose a simple testing procedure to detect big breaks in FMs stemming from a single source which does not suffer from the previous shortcomings. We focus on breaks in the loadings though we also provide a procedure to detect whether the breaks originate from the loadings or from factors themselves. In particular, we first derive some asymptotic results finding that, in contrast to small breaks, the number of factors is overestimated under big breaks, a result which was also used by BE (2011). We argue that neglecting these breaks can have serious consequences on the forecasting performance of some popular regression-based models using factors, where their number is a priori imposed. Likewise, under big breaks, it may be difficult to provide a structural interpretation of the estimated factors when they are chosen according to some consistent information criteria (see Bai and Ng, 2006b; Chen, 2012). Our proposal relies upon a very simple regression-based testing procedure. As sketched earlier, the insight is that a FM with big breaks in the loadings can be re-parameterized as another FM with constant loadings but a larger set of factors, where the number and the space spanned by the latter can be consistently estimated by PCA under fairly standard assumptions. Hence, rather than directly testing for whether all the elements of the loadings matrix are stable, which will suffer from an infinite-dimensionality problem as the number of variables in the panel data set grows, one can test if the relationships among the larger finite-dimensional set of estimated factors are stable.

Specifically, our procedure consists of two steps. First, the number of factors for the whole sample period is chosen as $\bar{r}$ according to Bai and Ng's (2002, BN henceforth) information criteria, and then $\bar{r}$ factors are estimated by PCA. Next, one of the estimated factors (e.g., the first one) is regressed on the remaining $\bar{r}-1$ factors, to next apply the standard Chow Test or the Sup-type Test of Andrews (1993) to this regression, depending on whether the date of the break is treated as known or unknown. If the null of no breaks is rejected in the second-step regression, we conclude that there are big breaks and, otherwise, that either no breaks exist at all or that only small breaks occur. Further, on the basis of the rank properties of the covariance matrix of the estimated factors in different subsamples, we also provide some guidance on how to distinguish between breaks stemming either from the loadings or from the data generating process (DGP hereafter) of the factors. This difference is important since the latter may lead to reject the null of constant loadings when it is true, implying a misleading interpretation of the source of the break.

After completing the first draft of this paper, we became aware of a closely related paper by Han and Inoue (2012, HI hereafter) who, in an independent research, adopt a similar approach to ours in testing for big breaks in FM. The two approaches, however, differ in some relevant respects. In effect, rather than using a simple regression-based approach to avoid the infinite-dimensionality problem as we do here, HI (2012) test directly for differences before and after the break in all the elements of the covariance matrix of the estimated factors. We will argue below that, despite the fact that the HI tests use more information than ours, both tests generally exhibit similar power. Indeed, our regression-based test based on the Wald principle, which behaves much better in general than the Lagrange multiplier (LM hereafter) tests for detecting structural breaks, is even more powerful than the corresponding HI's test for small sample sizes, such as $N = T = 50$. One additional advantage of our simple linear-regression setup is that it amenable to use many other existing methods for testing breaks, including multiple ones (see, e.g. Perron, 2006, for an extensive review of these tests).

The rest of the paper is organized as follows. In Section 2, we present the basic notation, assumptions and the definitions of *small* and *big* breaks. In Section 3, we analyze the consequences of big breaks on the choice of the number of factors and their estimation, as well as their effects on standard econometric practices with factor-augmented regressions. In Section 4, we first derive the asymptotic distributions of our tests and next discuss, when a big break is detected, how to identify its sources. Section 5 deals with the finite sample performance of our test relative to the competing tests using Monte-Carlo simulations. Section 6 provides an empirical application. Finally, Section 7 concludes. An appendix contains detailed proofs of the main results.

2. Notation and preliminaries

We consider FM that can be rewritten in the static canonical form:

$$X_t = AF_t + e_t$$

where X_t is the $N \times 1$ vector of observed variables, $A = (\alpha_1, \dots, \alpha_N)'$ is the $N \times r$ matrix of factor loadings, r is the number of common factors which is finite, $F_t = (F_{1t}, \dots, F_{rt})'$ is the $r \times 1$ vector of common factors, and e_t is the $N \times 1$ vector of idiosyncratic errors. In the case of dynamic FMs, all the common factors f_t and their lags are stacked into F_t. Thus, a dynamic FM with r dynamic factors and p lags of these factors can be written as a static FM with $r \times (p+1)$ static factors. Further, given the assumptions we make about e_t, the case analyzed by BE (2011) where the e_{it} disturbances are generated by individual specific autoregressive (AR hereafter) processes is also considered.[2]

We assume that there is a single structural break in the factor loadings of all factors at the same date τ:

$$X_t = AF_t + e_t \quad t = 1, 2, \dots, \tau, \tag{1}$$

$$X_t = BF_t + e_t \quad t = \tau + 1, \dots, T \tag{2}$$

where $B = (\beta_1, \dots, \beta_N)'$ is the new factor loadings after the break. By defining the matrix $C = B - A$, which captures the size of the

1 With the notation used below in (1)–(2), the limiting distribution of the rejection frequencies for the joint hypothesis $A = B$ is not known, although the individual tests for the hypothesis $\alpha_i = \beta_i$ have known limiting distributions.

2 Notice, however, that our setup excludes the generalized dynamic FM considered by Forni and Lippi (2001), where the polynomial distributed lag possibly tends to infinity.