Journal of Econometrics 146 (2008) 10–25

Contents lists available at ScienceDirect

Journal of Econometrics

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

Indirect estimation of large conditionally heteroskedastic factor models, with an application to the Dow 30 stocks

Enrique Sentana [a,*], Giorgio Calzolari [b], Gabriele Fiorentini [b]

[a] CEMFI, Casado del Alisal 5, E-28014 Madrid, Spain

[b] Università di Firenze, Viale Morgagni 59, I-50134 Firenze, Italy

ARTICLE INFO

Article history:
Available online 26 June 2008

JEL classification:
C13
C15
C32

Keywords:
ARCH
Idiosyncratic risk
Inequality constraints
Kalman filter
Sequential estimators
Simulation estimators
Volatility

ABSTRACT

We derive indirect estimators of conditionally heteroskedastic factor models in which the volatilities of common and idiosyncratic factors depend on their past unobserved values by calibrating the score of a Kalman-filter approximation with inequality constraints on the auxiliary model parameters. We also propose alternative indirect estimators for large-scale models, and explain how to apply our procedures to many other dynamic latent variable models. We analyse the small sample behaviour of our indirect estimators and several likelihood-based procedures through an extensive Monte Carlo experiment with empirically realistic designs. Finally, we apply our procedures to weekly returns on the Dow 30 stocks.

1. Introduction

Many issues in finance, including asset pricing, portfolio allocation and risk management, require the analysis of the variances and covariances of a large number of security returns. Traditionally, these issues have been considered in a static framework, but the emphasis has gradually shifted to intertemporal models in which agents' actions are based on the distribution of returns conditional on their information set. Parallel to these theoretical developments, a large family of statistical models for the time variation in conditional variances has grown following Engle's (1982) work on ARCH processes, and numerous applications have already appeared. By and large, though, most applied work in this area has been on univariate series, as the application of these models in a multivariate context has been hampered by the sheer number of parameters involved. In this sense, it is worth mentioning that even with the massive computational power available to the economists nowadays, the multivariate vec analogue of the ubiquitous univariate GARCH(1, 1) model has hardly ever been estimated for more than two series at a time, often with many additional restrictions to ensure that the resulting conditional covariance matrices are positive definite (see Bauwens et al. (2006) for a recent survey).

Given the strong commonality in the volatility movements of different financial assets, it is not surprising that one of the most popular approaches to multivariate dynamic heteroskedasticity employs the same idea as traditional factor analysis to obtain a parsimonious representation of conditional second moments. That is, it is assumed that each of N observed variables is a linear combination of k ($k \ll N$) common factors plus an idiosyncratic term, but with conditional heteroskedasticity in the underlying variables. Such models are particularly appealing in finance, as the concept of factors plays a fundamental role in major asset pricing theories, such as the Arbitrage Pricing Theory of Ross (1976) (see King et al. (1994)). In addition, they automatically guarantee a positive definite conditional covariance matrix for the observed series, once we ensure that the variances of common and specific factors are non-negative.

If the conditional variances of the latent variables are measurable functions of observed variables, as in Engle et al. (1990), maximum likelihood (ML) estimation is time-consuming, but straightforward. However, one has to exercise care in dealing with conditional variances that depend on past values of the common or idiosyncratic factors, as their true values do not necessarily belong to the econometrician's information set (see Harvey et al. (1992)) (HRS). The original latent factor model with ARCH effects on the common factors introduced by Diebold and Nerlove (1989) is the

* Corresponding author. Tel.: +34 91 429 0551; fax: +34 91 429 1056.
E-mail addresses: sentana@cemfi.es (E. Sentana), calzolar@ds.unifi.it (G. Calzolari), fiorentini@ds.unifi.it (G. Fiorentini).

0304-4076/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2008.06.001

best known example. In such models, the distribution and moments of the observed variables conditional on their past values alone are unknown. To some degree, this has prompted interest in other parameter driven models (see Andersen (1994) or Shephard (1996)), in which the volatility of the latent factors evolves according to discrete-time stochastic volatility processes (see Pitt and Shephard (1999), Aguilar and West (2000), Meddahi and Renault (2004), Chib et al. (2006) and Doz and Renault (2006)), or discrete-state Markov chains (see Carter and Kohn (1994), Shephard (1994) and Kim and Nelson (1999), and the references therein).

Despite the attractive features of these alternative models, one should not necessarily abandon the use of GARCH processes in latent variable models, especially if one considers that many macro and finance theories are often specified using conditional moments, although those moments are defined with respect to the agents' information set, not the econometrician's. For that reason, Fiorentini et al. (2004) develop computationally feasible Markov Chain Monte Carlo (MCMC) simulation algorithms that can be used to obtain exact likelihood-based estimators of factor models with GARCH structures in the common factors, thereby avoiding the inconsistencies associated with the Kalman filter approximations to the log-likelihood function proposed by Diebold and Nerlove (1989) and HRS. However, they maintain the assumption of constant idiosyncratic variances, which seems at odds with the empirical evidence.

In this paper we drop that assumption and analyse alternative simulation-based estimation methods belonging to the class of indirect estimation procedures proposed by Gallant and Tauchen (1996), Gouriéroux et al. (1993) and Smith (1993). In fact, Gouriéroux et al. (1993) explicitly considered conditionally heteroskedastic factor models as one of their examples, and suggested a first order, discrete-state Markov chain as auxiliary model for the case of ARCH(1) dynamics (see also Billio and Monfort (2003), who use non-parametric auxiliary models when N is small). Our approach is more closely related to Dungey et al. (2000), who also developed indirect estimators for such models. Specifically, they considered two auxiliary parametric models: a "dual" VAR model for the levels and squares (but not cross products) of the observed series (see also Zhumabekova and Dungey (2001)), and the Kalman filter-based approximation to the log-likelihood function used by Diebold and Nerlove (1989). Although Dungey et al. (2000) found in a limited Monte Carlo exercise that the latter yields indirect estimators with substantially smaller root mean square errors than the former, they did not use it in their empirical application. Another problem with their VAR approach is that the number of parameters of the auxiliary model increases with the square of the number of asset returns, which rules out its application to large models.

In this context, our methodological contribution is twofold. First, we show that the HRS approximation is an ideal auxiliary model because (a) it has exactly the same number of parameters as the model of interest, and with the same meaning, and it is also easy to estimate; and (b) it spans the score of the model of interest in some important limiting cases, providing a very good approximation to it in more general situations. Second, we derive an alternative joint indirect estimator on the basis of the sequential estimators of the HRS approximation in Sentana and Fiorentini (2001) (SF), which can be applied to situations with rather large cross-sectional dimensions if we add the empirically plausible assumption that the dynamic variance coefficients of the idiosyncratic terms are common. In addition, we conduct an empirically realistic Monte Carlo experiment to assess the finite sample performance of our two proposed indirect estimators relative to the approximate methods of HRS and SF. We also compare them to the Bayesian estimators of Fiorentini et al. (2004) in their restricted case.

Importantly, we explain how our proposed estimators can easily be adapted to handle any state space model with GARCH disturbances, which includes many examples that have been used in the empirical economic and finance literatures, such as structural time series models, or regression models with time-varying coefficients (see Harvey (1989) and Kim and Nelson (1999)).

Finally, we apply our estimators to weekly excess returns on the 30 components of the Dow Jones Industrial Average. We also gauge the importance of allowing for time-variation in conditional correlations and idiosyncratic volatilities, as well as for non-normality of returns. In addition, our empirical results shed some light on whether the increase in idiosyncratic risk documented by Campbell et al. (2001) continued after the dot-com bubble burst.

In Section 2, we define the model of interest and the HRS approximation, and study their relationship in detail. Then in Section 3 we introduce our two indirect estimators, and explain how they can be extended to deal with more general models. Our Monte Carlo evidence is included in Section 4. Finally, the results of the empirical application are presented in Section 5, and our conclusions in Section 6. Proofs and auxiliary results can be found in appendices.

2. Conditionally heteroskedastic factor models

2.1. Definition

Consider the following multivariate model:

$$\mathbf{x}_t = \mathbf{B}\mathbf{f}_t + \mathbf{u}_t, \tag{1}$$

$$\begin{pmatrix} \mathbf{f}_t \\ \mathbf{u}_t \end{pmatrix} | I_{t-1} \sim N\left[\begin{pmatrix} \mathbf{0} \\ \mathbf{0} \end{pmatrix}, \begin{pmatrix} \mathbf{\Delta}_t & \mathbf{0} \\ \mathbf{0} & \mathbf{\Psi}_t \end{pmatrix}\right], \tag{2}$$

where $\mathbf{x}_t$ is a $N \times 1$ vector of observable random variables, $\mathbf{f}_t$ is a $k \times 1$ vector of unobserved common factors, $\mathbf{B}$ is the $N \times k$ matrix of factor loadings, with $N \geq k$ (typically $N \gg k$) and rank $(\mathbf{B}) = k$, $\mathbf{u}_t$ is a $N \times 1$ vector of idiosyncratic noises, which are conditionally orthogonal to $\mathbf{f}_t$, $\mathbf{\Delta}_t$ is a $k \times k$ diagonal positive definite (p.d.) matrix of time-varying factor variances, $\mathbf{\Psi}_t$ is a $N \times N$ diagonal positive semidefinite (p.s.d.) matrix of time-varying idiosyncratic variances, and I_{t-1} is an information set that contains the values of $\mathbf{x}_t$ and $\mathbf{f}_t$ up to $t-1$. Thus, the distribution of $\mathbf{x}_t$ conditional on I_{t-1} is $N(\mathbf{0}, \mathbf{\Sigma}_t)$, where $\mathbf{\Sigma}_t = \mathbf{B}\mathbf{\Delta}_t\mathbf{B}' + \mathbf{\Psi}_t$ has the usual exact factor structure. For this reason, we shall refer to the data generation process specified by (1) and (2) as a multivariate conditionally heteroskedastic exact factor model. To simplify the exposition, we maintain the normality assumption until the empirical application in Section 5.

Such a formulation nests several models for asset returns widely used in the empirical finance literatures on asset pricing, portfolio selection, hedging and risk management. Those applications typically assume that $\mathbf{f}_t$ and $\mathbf{u}_t$ follow dynamic heteroskedastic processes, but differ in the particular specification of their conditional variances $\delta_{jt} = V(f_{jt}|I_{t-1})$ $(j = 1, \ldots, k)$ and $\psi_{it} = V(u_{it}|I_{t-1})$ $(i = 1, \ldots, N)$. For instance, Diebold and Nerlove (1989) parametrised the common factors as univariate *strong* ARCH models, in the sense of Drost and Nijman (1993). For the sake of concreteness, in this paper we study in detail the more realistic covariance stationary GARCH(1, 1) formulation:

$$\delta_{jt} = \varsigma_j + \phi_j f_{jt-1}^2 + \rho_j \delta_{jt-1}, \quad (j = 1, \ldots, k) \tag{3}$$

$$\psi_{it} = \varsigma_i^* + \phi_i^* u_{it-1}^2 + \rho_i^* \psi_{it-1}, \quad (i = 1, \ldots, N) \tag{4}$$

with $\varsigma_j = (1 - \phi_j - \rho_j)\delta_j$, $\delta_j = E(\delta_{jt}|\varrho) = V(f_{jt}|\varrho)$, $\varsigma_i^* = (1 - \phi_i^* - \rho_i^*)\psi_i$ and $\psi_i = E(\psi_{it}|\varrho) = V(u_{it}|\varrho)$, where $E(.|\varrho)$ denotes expected values taken with respect to model (1)–(4)