Journal of Econometrics 177 (2013) 233–249

Contents lists available at ScienceDirect

Journal of Econometrics

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

Sequential estimation of shape parameters in multivariate dynamic models ☆

Dante Amengual [a], Gabriele Fiorentini [b,c], Enrique Sentana [a,*]

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

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

[c] RCEA, Rimini, Italy

ARTICLE INFO

Article history:
Available online 27 April 2013

JEL classification:
C13
C32
G01
G11

Keywords:
Confidence intervals
Elliptical distributions
Efficient estimation
Global systematically important banks
Systemic risk
Risk management

ABSTRACT

Sequential maximum likelihood and GMM estimators of distributional parameters obtained from the standardised innovations of multivariate conditionally heteroskedastic dynamic regression models evaluated at Gaussian PML estimators preserve the consistency of mean and variance parameters while allowing for realistic distributions. We assess their efficiency, and obtain moment conditions leading to sequential estimators as efficient as their joint ML counterparts. We also obtain standard errors for VaR and CoVaR, and analyse the effects on these measures of distributional misspecification. Finally, we illustrate the small sample performance of these procedures through simulations and apply them to analyse the risk of large eurozone banks.

1. Introduction

Both academics and financial market participants are often interested in features of the distribution of asset returns beyond its conditional mean and variance. In particular, the Basel Capital Adequacy Accord forced banks and other financial institutions to develop models to quantify all their risks accurately. In practice, most institutions chose the so-called Value at Risk (VaR) framework in order to determine the capital necessary to cover their exposure to market risk. As is well known, the VaR of a portfolio of financial assets is defined as the positive threshold value V such that the probability of the portfolio suffering a reduction in wealth larger than V over some fixed time interval equals some pre-specified level $\lambda < 1/2$. Similarly, the recent financial crisis has highlighted the need for systemic risk measures that assess how an institution is affected when another institution, or indeed the entire financial system, is in distress. Given that the probability of the joint occurrence of several extreme events is regularly underestimated by the multivariate normal distribution, any such measure should definitely take into account the non-linear dependence induced by the non-normality of financial returns.

A rather natural modelling strategy is to specify a parametric leptokurtic distribution for the standardised innovations of the vector of asset returns, such as the multivariate Student t, and to estimate the conditional mean and variance parameters jointly with the parameters characterising the shape of the assumed distribution by maximum likelihood (ML) (see for example Pesaran et al. (2009) and Pesaran and Pesaran (2010)). Elliptical distributions such as the multivariate t are attractive in this context because they relate mean–variance analysis to expected utility maximisation (see e.g. Chamberlain (1983) or Owen and Rabinovitch (1983)). Moreover, they generalise the multivariate normal distribution but retain its analytical tractability irrespective

☆ We would like to thank Manuel Arellano, Christian Bontemps, Antonio Díez de los Ríos, Olivier Faugeras, Javier Mencía, Francisco Peñaranda, Marcos Sanso, David Veredas and audiences at the Bank of Canada, CEMFI, Chicago Booth, CREST, ECARES ULB, Koç, Princeton, Rimini, Toulouse, the Finance Forum (Granada, 2011), the Symposium of the Spanish Economic Association (Málaga, 2011) and the Conference in honour of M. Hashem Pesaran (Cambridge, 2011) for useful comments and suggestions. We also thank the editors and two anonymous referees for valuable feedback. Luca Repetto provided able research assistance for the empirical application. Of course, the usual caveat applies. Amengual and Sentana gratefully acknowledge financial support from the Spanish Ministry of Science and Innovation through grants ECO 2008-00280 and 2011-26342 while Fiorentini acknowledges funding from MIUR PRIN MISURA - Multivariate models for risk assessment.

* Corresponding author.
E-mail addresses: amengual@cemfi.es (D. Amengual), fiorentini@disia.unifi.it (G. Fiorentini), sentana@cemfi.es (E. Sentana).

0304-4076/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jeconom.2013.04.010

of the number of assets. However, non-Gaussian ML estimators often achieve efficiency gains under correct specification at the risk of returning inconsistent parameter estimators under distributional misspecification (see Newey and Steigerwald (1997)). Unfortunately, semiparametric estimators of the joint density of the innovations suffer from the curse of dimensionality, which severely limits their use. Another possibility would be semiparametric methods that impose the assumption of ellipticity, which retain univariate nonparametric rates regardless of the cross-sectional dimension of the data, but asymmetries in the true distribution will again contaminate the resulting estimators of conditional mean and variance parameters.

Sequential estimators of shape parameters that use the Gaussian Pseudo ML estimators of the mean and variance parameters as first step estimators offer an attractive compromise because they preserve the consistency of the first two conditional moments under distributional misspecification as long as those moments are correctly specified and the fourth moments are bounded (see Bollerslev and Wooldridge (1992)), while allowing for more realistic conditional distributions. From a more practical point of view, they also simplify the computations by reducing the dimensionality of the optimisation problem at each stage, thereby increasing the researcher's confidence that she has not found a local minimum. In this regard, it is worth bearing in mind that most commercially available econometric packages have been fine tuned to the Gaussian case, which even leads to closed-form estimators in commonly used models.

The focus of our paper is precisely the econometric analysis of sequential estimators obtained from the standardised innovations evaluated at the Gaussian PML estimators. Specifically, we consider not only sequential ML estimators, but also sequential generalised method of moments (GMM) estimators based on certain functions of the standardised innovations.

To keep the exposition simple we focus on elliptical distributions in the text, and relegate more general cases to the supplemental Appendix. We illustrate our results with several examples that nest the normal, including the Student t and some rather flexible families such as scale mixtures of normals and polynomial expansions of the multivariate normal density, both of which could form the basis for a proper nonparametric procedure. We explain how to compute asymptotically valid standard errors of sequential estimators, assess their efficiency, and obtain the optimal moment conditions that lead to sequential MM estimators as efficient as their joint ML counterparts. Although we consider multivariate conditionally heteroskedastic dynamic regression models, our results apply in univariate contexts as well as in static ones.

We then analyse the use of our sequential estimators in the computation of commonly used risk management measures such as VaR, and recently proposed systemic risk measures such as Conditional Value at Risk (CoVaR) (see Adrian and Brunnermeier (2011)). In particular, we compare our sequential estimators to nonparametric estimators, both when the parametric conditional distribution is correctly specified and also when it is misspecified. Our analytical and simulation results indicate that sequential ML estimators of flexible parametric families of distributions offer substantial efficiency gains, while incurring in small biases.

Finally, we illustrate our results with data for four Global Systematically Important Banks from the eurozone. As expected, we find that their stock returns display considerable non-normality even after controlling for time-varying volatilities and correlations, which in turn gives rise to the type of non-linear dependence that is relevant for systemic risk measurement.

The rest of the paper is as follows. In Section 2, we describe the model, present the elliptical distributions we use as examples and introduce a convenient reparametrisation satisfied by most static and dynamic models. Then, in Section 3 we discuss the sequential ML and GMM estimators, and compare their efficiency. In Section 4, we study the effect of those estimators on risk measures under both correct specification and misspecification, and derive asymptotically valid standard errors. A Monte Carlo evaluation of the different parameter estimators and risk measures can be found in Section 5, and the empirical application in Section 6. Finally, we present our conclusions in Section 7. Proofs and auxiliary results are gathered in appendices.

2. Theoretical background

2.1. The dynamic econometric model

Discrete time models for financial time series are usually characterised by a parametric dynamic regression model with time-varying variances and covariances. Typically, the N dependent variables, $\mathbf{y}_t$, are assumed to be generated as:

$$\mathbf{y}_t = \boldsymbol{\mu}_t(\boldsymbol{\theta}_0) + \boldsymbol{\Sigma}_t^{1/2}(\boldsymbol{\theta}_0)\boldsymbol{\varepsilon}_t^*,$$
$$\boldsymbol{\mu}_t(\boldsymbol{\theta}) = \boldsymbol{\mu}(\mathbf{z}_t, I_{t-1}; \boldsymbol{\theta}), \qquad \boldsymbol{\Sigma}_t(\boldsymbol{\theta}) = \boldsymbol{\Sigma}(\mathbf{z}_t, I_{t-1}; \boldsymbol{\theta}),$$

where $\boldsymbol{\mu}()$ and vech $[\boldsymbol{\Sigma}()]$ are $N \times 1$ and $N(N+1)/2 \times 1$ vector functions known up to the $p \times 1$ vector of true parameter values $\boldsymbol{\theta}_0$, $\mathbf{z}_t$ are k contemporaneous conditioning variables, I_{t-1} denotes the information set available at $t-1$, which contains past values of $\mathbf{y}_t$ and $\mathbf{z}_t$, $\boldsymbol{\Sigma}_t^{1/2}(\boldsymbol{\theta})$ is some particular "square root" matrix such that $\boldsymbol{\Sigma}_t^{1/2}(\boldsymbol{\theta})\boldsymbol{\Sigma}_t^{1/2\prime}(\boldsymbol{\theta}) = \boldsymbol{\Sigma}_t(\boldsymbol{\theta})$, and $\boldsymbol{\varepsilon}_t^*$ is a martingale difference sequence satisfying $E(\boldsymbol{\varepsilon}_t^*|\mathbf{z}_t, I_{t-1}; \boldsymbol{\theta}_0) = \mathbf{0}$ and $V(\boldsymbol{\varepsilon}_t^*|\mathbf{z}_t, I_{t-1}; \boldsymbol{\theta}_0) = \mathbf{I}_N$. Hence,

$$E(\mathbf{y}_t|\mathbf{z}_t, I_{t-1}; \boldsymbol{\theta}_0) = \boldsymbol{\mu}_t(\boldsymbol{\theta}_0), \qquad V(\mathbf{y}_t|\mathbf{z}_t, I_{t-1}; \boldsymbol{\theta}_0) = \boldsymbol{\Sigma}_t(\boldsymbol{\theta}_0). \quad (1)$$

To complete the model, we need to specify the conditional distribution of $\boldsymbol{\varepsilon}_t^*$. We shall initially assume that, conditional on $\mathbf{z}_t$ and I_{t-1}, $\boldsymbol{\varepsilon}_t^*$ is independent and identically distributed as some particular member of the spherical family with a well defined density, or $\boldsymbol{\varepsilon}_t^*|\mathbf{z}_t, I_{t-1}; \boldsymbol{\theta}_0, \boldsymbol{\eta}_0 \sim i.i.d.\, s(\mathbf{0}, \mathbf{I}_N, \boldsymbol{\eta}_0)$ for short, where $\boldsymbol{\eta}$ are q additional shape parameters.

2.2. Elliptical distributions

A spherically symmetric random vector of dimension N, $\boldsymbol{\varepsilon}_t^*$, is fully characterised in Theorem 2.5 of Fang et al. (1990) as $\boldsymbol{\varepsilon}_t^* = e_t\mathbf{u}_t$, where $\mathbf{u}_t$ is uniformly distributed on the unit sphere surface in $\mathbb{R}^N$, and e_t is a non-negative random variable independent of $\mathbf{u}_t$. The variables e_t and $\mathbf{u}_t$ are referred to as the generating variate and the uniform base of the spherical distribution. Often, we shall also use $\varsigma_t = \boldsymbol{\varepsilon}_t^{*\prime}\boldsymbol{\varepsilon}_t^*$, which trivially coincides with e_t^2. Assuming that $E(e_t^2) < \infty$, we can standardise $\boldsymbol{\varepsilon}_t^*$ by setting $E(e_t^2) = N$, so that $E(\boldsymbol{\varepsilon}_t^*) = \mathbf{0}$ and $V(\boldsymbol{\varepsilon}_t^*) = \mathbf{I}_N$. If we further assume that $E(e_t^4) < \infty$, then Mardia's (1970) coefficient of multivariate excess kurtosis

$$\kappa = E(\varsigma_t^2)/[N(N+2)] - 1 \quad (2)$$

will also be bounded. The most prominent examples are the standardised multivariate Student t, in which ς_t is proportional to an F random variable with N and ν degrees of freedom, and the limiting Gaussian case, when ς_t becomes a χ_N^2. Since this involves no additional parameters, we identify the normal distribution with $\boldsymbol{\eta}_0 = \mathbf{0}$, while for the Student t we define η as $1/\nu$, which will always remain in the finite range $[0, 1/2)$ under our assumptions. Normality is thus achieved as $\eta \to 0$ (see Fiorentini et al. (2003)). Other more flexible families of spherical distributions that we will also use to illustrate our general results are:

Discrete scale mixture of normals: $\boldsymbol{\varepsilon}_t^* = \sqrt{\varsigma_t}\mathbf{u}_t$ is distributed as a DSMN if and only if

$$\varsigma_t = [s_t + (1-s_t)\varkappa]/[\alpha + (1-\alpha)\varkappa] \cdot \zeta_t$$