Journal of Econometrics 178 (2014) 22–44

Contents lists available at ScienceDirect

Journal of Econometrics

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

A unified approach to validating univariate and multivariate conditional distribution models in time series

Bin Chen[a,*], Yongmiao Hong[b,c,d,e,**]

[a] Department of Economics, University of Rochester, Rochester, NY, 14627, United States
[b] Department of Economics, Cornell University, Ithaca, NY 14850, United States
[c] Department of Statistical Science, Cornell University, Ithaca, NY 14850, United States
[d] Wang Yanan Institute for Studies in Economics (WISE), Xiamen University, Xiamen 361005, China
[e] MOE Key Laboratory in Econometrics, Xiamen University, Xiamen 361005, China

ARTICLE INFO

Article history:
Available online 12 August 2013

JEL classification:
C4
G0

Keywords:
Diagnostic procedure
Empirical distribution function
Frequency domain
Generalized Cramer–von Mises test
Kernel method
Non-Markovian process
Time series conditional distribution model

ABSTRACT

Modeling conditional distributions in time series has attracted increasing attention in economics and finance. We develop a new class of generalized Cramer–von Mises (GCM) specification tests for time series conditional distribution models using a novel approach, which embeds the empirical distribution function in a spectral framework. Our tests check a large number of lags and are therefore expected to be powerful against neglected dynamics at higher order lags, which is particularly useful for non-Markovian processes. Despite using a large number of lags, our tests do not suffer much from loss of a large number of degrees of freedom, because our approach naturally downweights higher order lags, which is consistent with the stylized fact that economic or financial markets are more affected by recent past events than by remote past events. Unlike the existing methods in the literature, the proposed GCM tests cover both univariate and multivariate conditional distribution models in a unified framework. They exploit the information in the joint conditional distribution of underlying economic processes. Moreover, a class of easy-to-interpret diagnostic procedures are supplemented to gauge possible sources of model misspecifications. Distinct from conventional CM and Kolmogorov–Smirnov (KS) tests, which are also based on the empirical distribution function, our GCM test statistics follow a convenient asymptotic $N(0, 1)$ distribution and enjoy the appealing "nuisance parameter free" property that parameter estimation uncertainty has no impact on the asymptotic distribution of the test statistics. Simulation studies show that the tests provide reliable inference for sample sizes often encountered in economics and finance.

1. Introduction

The modeling of conditional distributions in time series has been advancing rapidly, with a wide range of applications in economics and finance (e.g. Granger, 1999; Corradi and Swanson, 2006b). Enormous empirical evidences document that economic and financial variables are typically nonlinear and nonnormally distributed, and have asymmetric comovements.[1] Consequently, one has to go beyond the conditional mean and conditional variance to obtain a complete picture for the dynamics of time series of interest. The conditional distribution characterizes the full dynamics of economic variables. As pointed out by Granger (2003), the knowledge of the conditional distribution is essential in performing various economic policy evaluations, financial forecasts, derivative pricing and risk management.[2]

In economics and econometrics, effort has been devoted to using higher moments and the entire distribution. Rothschild and Stiglitz's (1971, 1972) seminal works have demonstrated that the risk or uncertainty should be characterized by the distribution function, rather than the first two moments. In particular, the

* Corresponding author.
** Corresponding author at: Department of Economics, Cornell University, Ithaca, NY 14850, United States.
E-mail addresses: bchen8@mail.rochester.edu (B. Chen), yh20@cornell.edu (Y. Hong).

1 Empirical evidences against normality can be dated back to Mills (1927) and continue through today, see, e.g., Ang and Chen (2002), Bollerslev (1986), Longin and Solnik (2001).

2 A prominent example is in the option pricing context, where the price is determined by not just the conditional mean and variance, but functions of conditional distribution. Another example is to calculate value-at-risk (VaR), where the key step is to accurately estimate the conditional distribution of asset returns and the preassumed normal distribution can significantly underestimate the downward risk.

0304-4076/$ – see front matter
http://dx.doi.org/10.1016/j.jeconom.2013.08.004

ranking of the cumulative distribution function (CDF) by certain rules always coincides with that of the risk-averter's preference,[3] while the mean–variance analysis is only applicable to the restricted family of utility functions or distribution functions. Granger (1999), in a model evaluation context, suggests that the predictive conditional distribution should be provided, since forecasts based on conditional means are optimal only for a very limited class of loss functions.[4]

In time series analysis, the most popular models are ARMA models for conditional mean and GARCH models for conditional variance. However, as Hansen (1994) points out "there is no reason to assume, in general, that the only features of the conditional distribution which depend upon the conditional information are the mean and variance". Although still in an early stage, some time series models have been developed to study skewness, kurtosis and even the entire distribution. Hansen (1994) develops a general model for autoregressive conditional density (ARCD), which allows for time-varying first four conditional moments via a generalized skewed t-distribution. Harvey and Siddique (1999) propose a generalized autoregressive conditional skewness model (GARCHS) in a conditional non-central t-distribution framework by explicitly modeling the conditional second and third moments jointly. Brooks et al. (2005) develop a generalized autoregressive conditional heteroscedasticity and kurtosis (GARCHK) model via a central t distribution with time-varying degrees of freedom. Other examples of distribution models include Engle and Russell's (1998, 2005) autoregressive conditional duration (ACD) and autoregressive conditional multinomial (ACM) models, Bowsher's (2007) vector conditional intensity model, Hamilton's (1989, 1990) Markov regime switching models and Geweke and Amisano's (2007) compound Markov mixture models.

In addition to the univariate time series distribution modeling, the recent literature has documented a rapid growth of multivariate conditional distribution models, due to an increasing need to capture the joint dynamics of multivariate processes, such as in macroeconomic control, pricing, hedging and risk management.[5] For example, CAPM studies the relationship between individual asset returns and the market return, which has motivated the development of multivariate GARCH models (e.g., Bollerslev et al., 1988, Engle, 2002b). Among multivariate distribution models, copula-based models have become increasingly popular in characterizing the comovement between markets, risk factors and other relevant variables (e.g., Patton, 2004, Hu, 2006, Lee and Long, 2009). Another example is the extension of Markov regime switching models to a multivariate framework (e.g., Clements and Krolzig, 2003, Chauvet and Hamilton, 2006). Markov regime switching models can capture the asymmetry, nonlinearity and persistence of extreme observations of time series.

Efficient parameter estimation, optimal distribution forecast, valid hypothesis testing and economic interpretation all require correct model specification. The work on testing distributional assumptions at least dates back to the Kolmogorov–Smirnov (KS) test. One undesired feature of this test is that it is not distribution free when parameters are estimated. Andrews (1997) extends the KS test to conditional distribution models for independent observations, where a bootstrap procedure is used to obtain critical values. Meanwhile, Zheng (2000) proposes a nonparametric test for conditional distribution functions based on the Kullback–Leibler information criterion and the kernel estimation of the underlying distributions. Fan et al. (2006) extend Zheng's (2000) test to allow for discrete dependent variables and for mixed discrete and continuous conditional variables. However, a limitation of the above tests is that the data must be independently and identically distributed, therefore ruling out time series applications especially when the underlying time series is non-Markovian.

Observing the fact that when a dynamic distribution model is correctly specified, the probability integral transform of observed data via the model-implied conditional density is i.i.d. $U[0, 1]$, Bai (2003) proposes a KS type test with Khmaladze's (1981) martingale transformation, whose asymptotic distribution is free of impact of parameter estimation. However, Bai's (2003) test only checks uniformity rather than the joint i.i.d. $U[0, 1]$ hypothesis. It will have no asymptotic unit power if the transformed data is uniform but not i.i.d. Moreover, in a multivariate context, the probability integral transform of data with respect to a model-implied multivariate conditional density is no longer i.i.d. $U[0, 1]$, even if the model is correctly specified. Bai and Chen (2008) evaluate the marginal distribution of both independent and serially dependent multivariate data by using the probability integral transform for each individual component. This test is legitimate, but it may miss important information on the joint distribution of a multivariate model. In particular, when applied to each component of multivariate time series data, Bai and Chen's (2008) test may fail to detect misspecification in the joint dynamics. For example, the test may easily overlook misspecification in the conditional correlations between individual time series.

Corradi and Swanson (2006a) propose bootstrap conditional distribution tests in the presence of dynamic misspecifications. However, they consider a finite dimensional information set and thus may not have good power against non-Markovian models. Their tests are designed for univariate time series. When extended to multivariate time series, their tests are not consistent against all alternatives to the null. Moreover, their critical values are data dependent and cannot be tabulated. Bierens and Wang (2012) propose a weighted integrated conditional moment (ICM) test of the validity of parametric specifications of conditional distribution models for stationary time series, extending Bierens' (1984) test. Their ICM test is consistent against all stationary alternatives, but its asymptotic distribution is case dependent and a bootstrap method has to be applied to obtain critical values, which is computationally intensive.

In a continuous-time diffusion framework, Ait-Sahalia et al. (2009) and Li and Tkacz (2006) propose tests by comparing the model-implied distribution function with its nonparametric counterpart. Both tests maintain the Markov assumption for the DGP, and only check one lag dependence, therefore are not suitable for non-Markovian models like GARCH or MA type models. Another undesired feature of these tests is that they have severe size distortion in finite samples and bootstrap must be used to approximate the distribution of the test statistics. Bhardwaj et al. (2008) consider a simulation-based test, which is an extension of Andrews' (1997) conditional KS test, for multivariate diffusion models. The limit distribution of their test is not nuisance parameter free and asymptotic critical values must be obtained via a block bootstrap.

In this paper, we shall propose a new class of generalized Cramer–von Mises (GCM) tests of the adequacy of univariate and multivariate conditional distribution models, without requiring prior knowledge of possible alternatives (including both functional forms and lag structures). Compared with the existing tests for conditional distribution models in the literature, our approach has several main advantages.

First, our GCM tests are constructed using a new approach, which embeds the empirical distribution function in a spectral framework. Thus it can detect misspecification in both marginal distribution and dynamics of a time series. Thanks to the use

[3] A closely related concept is second-order stochastic dominance, which ranks any pair of distributions with the same mean in terms of comparative risk.

[4] See also Christoffersen and Diebold (1997) for more discussion.

[5] Geweke and Amisano (2007) argue that "while univariate models are a first step, there is an urgent need to move on to multivariate modeling of the time-varying distribution of asset returns".