Journal of Econometrics 193 (2016) 17–34

Contents lists available at ScienceDirect

Journal of Econometrics

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

Inference theory for volatility functional dependencies[☆]

Jia Li[a], Viktor Todorov[b], George Tauchen[a,*]

[a] Department of Economics, Duke University, Durham, NC 27708, United States

[b] Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, United States

ARTICLE INFO

Article history:
Received 2 September 2014
Received in revised form
26 May 2015
Accepted 23 January 2016
Available online 10 February 2016

JEL classification:
C51
C52
G12

Keywords:
High-frequency data
Occupation measure
Semimartingale
Specification test
Stochastic volatility

ABSTRACT

We develop inference theory for models involving possibly nonlinear transforms of the elements of the spot covariance matrix of a multivariate continuous-time process observed at high frequency. The framework can be used to study the relationship among the elements of the latent spot covariance matrix and processes defined on the basis of it such as systematic and idiosyncratic variances, factor betas and correlations on a fixed interval of time. The estimation is based on matching model-implied moment conditions under the occupation measure induced by the spot covariance process. We prove consistency and asymptotic mixed normality of our estimator of the (random) coefficients in the volatility model and further develop model specification tests. We apply our inference methods to study variance and correlation risks in nine sector portfolios comprising the S&P 500 index. We document sector-specific variance risks in addition to that of the market and time-varying heterogeneous correlation risk among the market-neutral components of the sector portfolio returns.

1. Introduction

Economic models often place exact restrictions across the realizations of a set of random variables. One case in point is an affine term structure model for bond prices that constrains all bond yields to lie along a very low dimensional manifold; see, for example, Singleton (2006) and citations therein. More generally, factor models specify the pricing kernel as a function of a low-dimensional factor process. Combining this structure with a model for aggregate asset payoffs implies low-dimensional factor structure for the conditional distribution, and thereby the first and second conditional moments, of asset returns and derivative prices.

In a broad sense, dimension reduction is also often imposed for the purposes of parsimony in modeling high-dimensional objects so as to mitigate the statistical and/or computational complexity of econometric models. One example is the use of diffusion index in macroeconomic forecasting (see Stock and Watson (2002) and references therein). Other examples arise in models of the stochastic covariance matrix of a multivariate process, where dimension-reduction restrictions may take the form of time-invariant correlations (Bollerslev (1990)) or equal stochastic correlations among multiple time series (Engle and Kelly (2012)). Similar strategies have also been used in moderately high-dimensional models for nonlinear dependence, such as copula models (Oh and Patton (2013)).

The primary focus of this paper is the estimation and testing for such model restrictions among the elements of the spot covariance matrix of a multivariate process of asset returns. The estimation is based on high-frequency (intraday) observations of a multivariate Itô semimartingale process on a fixed time interval with mesh of the observation grid shrinking to zero. For this process, we are interested in estimating and testing pathwise models that impose time-invariant (over the observation interval) relations among possibly nonlinear transforms (e.g. beta, correlation and idiosyncratic variance) of the spot covariance matrix.

The statistical uncertainty in our setting arises from the fact that the spot covariance matrix is not directly observed and needs to be estimated, or “measured”, from the discrete observations. The measurement error makes the functional relationship among

☆ We would like to thank the editor, an associate editor and two anonymous referees for many useful comments and suggestions. We also thank Tim Bollerslev, Federico Bugni, Marine Carrasco, Silvia Goncalves, Wolfgang Hardle, Jean Jacod, Ilze Kalnina, Matthew Masten, Andrew Patton and Ya'acov Ritov for helpful discussions. Financial support from the National Science Foundatation (NSF) under grants SES-1227448 and SES-1326819 (Li) and SES-0957330 (Todorov) is gratefully acknowledged.

* Corresponding author.
E-mail addresses: jl410@duke.edu (J. Li), v-todorov@northwestern.edu (V. Todorov), george.tauchen@duke.edu (G. Tauchen).

http://dx.doi.org/10.1016/j.jeconom.2016.01.004
0304-4076/

latent risks hold only approximately for their estimated counterparts. Nevertheless, as the sampling frequency increases, the error vanishes asymptotically, so that we can uncover and rigorously test model restrictions for the latent risks. To the best of our knowledge, our test is the first general method for testing such model restrictions in the high-frequency setting, while allowing for general forms of nonstationarity and dependence in the data.

In the first part of our theoretical analysis, we consider estimation based on forming moment conditions under the covariance occupation measure which are implied by our pathwise volatility model. We then construct sample analogues of these moment conditions by plugging in local nonparametric estimators of volatility formed over blocks of high-frequency price increments with asymptotically decreasing length of each of the blocks. This is similar to Jacod and Rosenbaum (2013) who use block volatility estimates to construct estimators for integrated nonlinear functions of volatility. Finally, we weight the moment conditions using a feasible weight matrix and form a quadratic-form objective function that our estimator minimizes. We derive the limit behavior of our estimator not only in the case when the model is correctly specified but also in the case of model misspecification, and further provide feasible estimates for the standard errors of the parameters in the model.

Our estimator of the parameters of the volatility model can be viewed as an analogue to the classical minimum distance type estimators with several important differences. First, the moment conditions in our case are formed under the occupation measure and, hence, they hold for the observed path but not necessarily for the invariant distribution of the volatility process; indeed, the invariant distribution is not even required to exist. The strategy of framing inference procedures in terms of moments under the occupation measure opens the possibility of systematically reincarnating many classical moment-based econometric procedures (e.g. Hansen (1982)), which are framed under the probability measure, for conducting inference for multivariate volatility models. Second, the asymptotic behavior of the estimator is equivalent to that generated by observing the moment condition with the true value of the spot covariance matrix plus a Gaussian martingale defined on an extension of the original probability space. This Gaussian martingale has quadratic variation that is adapted to the original filtration and shrinks asymptotically at order Δ_n, where Δ_n is the length of the high-frequency interval. Our estimation problem is thus similar to the problem of estimating a signal with asymptotically shrinking Gaussian noise, see, for example, section VII.4 in Ibragimov and Has'minskii (1981). Third, the limit law of our estimator is mixed Gaussian which means that the precision of estimation will typically vary depending on the particular realization.

The second, and perhaps more important, part of our analysis is specification testing for the pathwise volatility model. Since the model holds almost everywhere in time over the fixed time interval, designing a test based on the distance from zero of the model-based moment conditions under the covariance occupation measure is not sufficient. The reason is fairly intuitive: the covariance occupation measure does not preserve the information about the value of the spot covariance matrix at a particular point in time. For this reason, we introduce the concept of the weighted covariance occupation measure which, unlike the original occupation measure, allows to weigh differently the values of the spot covariance matrix at different points in time. We derive an empirical-process-type theory for an estimator of the weighted occupation measure. We use the latter to design a specification test for our pathwise volatility model by comparing the distance from zero of a set of moment conditions under a family of weighted occupation measures. We show that if the family of weight functions is chosen appropriately, our test statistic produces an asymptotically valid test.

Finally, we apply our inference theory to study the stochastic covariance structure on the industry level using S&P 500 sector index exchange-traded funds (ETFs). We specify and test models for the spot covariance matrix of the components of the industry returns that are orthogonal to the market portfolio. Our results show that not all variations in the stochastic variances of the industry portfolios can be accounted for by their sensitivity to market returns and the market variance. Some sectors like the Financials and Energy have independent sources of variance shocks in addition to that of the market. We further document nontrivial temporal variation in correlation in the market-neutral industry portfolio returns with nontrivial cross-sectional differences. The temporal variation in the market-neutral industry portfolio correlations suggests the presence of additional factors in the industry portfolio returns to span their risks.

The inference methods developed in the current paper are related with several strands of literature. First, our work is closely related to the literature on volatility estimation using high-frequency data. Early work mainly focuses on the estimation of the integrated variance (Barndorff-Nielsen and Shephard (2002) and Andersen et al. (2003)) and covariance (Barndorff-Nielsen and Shephard (2004b)), which can be considered as the mean of the covariance occupation measure. The estimation of nonlinear transforms of the volatility has been considered by Barndorff-Nielsen et al. (2005), Jacod (2008), Mykland and Zhang (2009), Todorov and Tauchen (2012), Jacod and Rosenbaum (2013), Li et al. (2013), Kalnina and Xiu (2014) and Aït-Sahalia and Xiu (2015) among others. Similar to Jacod and Rosenbaum (2013) our estimation is based on local estimates of volatility, which are local versions of the truncated variation of Mancini (2001), over blocks of decreasing length. Unlike the above cited literature on the estimation of volatility functionals, our focus here is on the estimation and specification of pathwise models for the spot covariance matrix. On the technical level, this requires a derivation of empirical-process-type limit results for a family of weighted covariance measures which is new.[1] Second, in Li et al. (2013) we advocate the volatility occupation measure as a unifying framework for high-frequency based volatility estimation, but in Li et al. (2013) we focus only on the estimation of the volatility occupation time in a univariate setting, without weighting the observations, and importantly without deriving feasible central limit theorems. Third, our inference can be compared with the literature on estimating parametric volatility models using realized measures under joint in-fill and long span asymptotics, see, for example, Bollerslev and Zhou (2002), Barndorff-Nielsen and Shephard (2002), Corradi and Distaso (2006), Todorov (2009) and Todorov and Tauchen (2012). Unlike our setup here, the estimation in these papers is always parametric (at least about the stochastic volatility part of the observed process) and relies crucially on the error due to the discrete sampling being dominated by the empirical process type error due to time aggregation. In contrast, our estimation here is performed on a fixed span and does not involve full parametric specification of the volatility process. The pathwise volatility models of interest here hold for whole families of parametric models. The fixed span setting also allows us to accommodate general forms of nonstationarity and dependence in the data.

[1] In a concurrent work, Li and Xiu (2015) study regression-type problems for noisy semimartingales using the stochastic variance as an explanatory variable in conditional moment equality models that arise from derivative pricing and market microstructure models. In contrast, the current paper focuses on the stochastic covariance matrix itself and conducts inference for pathwise restrictions on its behavior.