Journal of Econometrics 194 (2016) 153–186

Contents lists available at ScienceDirect

Journal of Econometrics

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

A nonparametric test of a strong leverage hypothesis ☆

Oliver Linton [a], Yoon-Jae Whang [b], Yu-Min Yen [c,*]

[a] Department of Economics, University of Cambridge, Austin Robinson Building, Sidgwick Avenue, Cambridge CB3 9DD, United Kingdom

[b] Department of Economics, Seoul National University, Seoul 151-742, Republic of Korea

[c] Department of International Business, National Chengchi University, 64, Sec. 2, Zhi-nan Rd., Wenshan, Taipei 116, Taiwan

ARTICLE INFO

Article history:
Received 22 July 2013
Received in revised form
30 January 2016
Accepted 10 February 2016
Available online 9 June 2016

JEL classification:
C14
C15

Keywords:
Distribution function
Leverage effect
Gaussian process

ABSTRACT

The so-called leverage hypothesis is that negative shocks to prices/returns affect volatility more than equal positive shocks. Whether this is attributable to changing financial leverage is still subject to dispute but the terminology is in wide use. There are many tests of the leverage hypothesis using discrete time data. These typically involve fitting of a general parametric or semiparametric model to conditional volatility and then testing the implied restrictions on parameters or curves. We propose an alternative way of testing this hypothesis using realized volatility as an alternative direct nonparametric measure. Our null hypothesis is of conditional distributional dominance and so is much stronger than the usual hypotheses considered previously. We implement our test on individual stocks and a stock index using intraday data over a long span. We find only very weak evidence against our hypothesis.

1. Introduction

The so-called leverage hypothesis, Black (1976) and Christie (1982), is essentially that negative shocks to stock prices affect their volatility more than equal magnitude positive shocks. Whether this is attributable to changing financial leverage or is a result of the volatility feedback effect (French et al., 1987; Campbell and Hentschel, 1992), is still subject to dispute (Engle and Ng, 1993; Figlewski and Wang, 2000; Bekaert and Wu, 2000; Bollerslev et al., 2006 and Dufour et al., 2012), but the terminology is in wide use. There are many statistical tests of the leverage hypothesis using discrete time data. These typically involve fitting of a general parametric or semiparametric model to conditional volatility and then testing the implied restrictions on parameters or curves, see for example Nelson (1991), Engle and Ng (1993), Linton and Mammen (2005), and Rodriguez and Ruiz (2012). Most authors have found the parameters governing asymmetric volatility response in daily individual stock returns and in indexes to be statistically significant.

A theoretical justification of the leverage effect is given in Christie (1982) inside a continuous time model, and recently there has been an important literature on measuring leverage effects in high frequency data. Aït-Sahalia et al. (2013) investigate the leverage effect "puzzle" within the continuous time framework. The puzzle is that natural estimators of the leverage effect based on high frequency data are usually very small and insignificant. They take apart the sources of this finding and interpret it as bias due to microstructure noise issues, and they propose a solution to this based on a bias correction. Empirically their method seems to uncover a stronger leverage effect. Wang and Mykland (2014) propose a nonparametric estimator of a class of leverage parameters inside a very general class of continuous time stochastic processes. They propose an estimator that is quite simple and easily studied and provide its limiting properties. They extend the theory to allow for measurement error and more sophisticated estimators of volatility and leverage. Their modified procedure is consistent and asymptotically mixed normal in this case too, although the rate of convergence is slower. They provide the means to conduct inference about the leverage parameter, although their application is more toward prediction of volatility. They demonstrate the value added that their leverage effect has in this purpose.

Bandi and Renò (2012) propose a nonparametric method for estimating the leverage effect in a continuous time stochastic volatility with jumps model. They use a flexible function of the

☆ We thank Valentina Corradi, Jean-Marie Dufour, and Haim Levy for helpful comments. R software for carrying out the conditional dominance test is available from the web site www.oliverlinton.me.uk.

* Corresponding author.
E-mail addresses: obl20@cam.ac.uk (O. Linton), whang@snu.ac.kr (Y.-J. Whang), yyu_min@nccu.edu.tw (Y.-M. Yen).

http://dx.doi.org/10.1016/j.jeconom.2016.02.018
0304-4076/

state of the firm, which is associated with intraday returns and spot variances to measure the leverage effect. They prove consistency of the functional estimates as the number of observations diverges to infinity and the interval for estimating the intradaily spot variances approaches zero. Asymptotic properties of their estimators also depend on behaviors of the jump components in the price process. Using the proposed estimators with intradaily asset returns and estimated spot variances as inputs, they find that the leverage effect is time varying and its magnitude increases with the variance level.

Our focus is on the low frequency (daily) volatility and return relationship. We propose a way of testing the leverage hypothesis nonparametrically without requiring a specific parametric or semiparametric model. Consequently our test statistics do not need an estimated quantity for measuring the leverage effect as an input. This is a major difference between our approach and the aforementioned methods proposed by Aït-Sahalia et al. (2013), Bandi and Renò (2012) and Wang and Mykland (2014), which all rely on using the estimated leverage effect parameter for statistical inference. Our inference is robust to the model choices that many previous studies have adopted. In fact, we test a "strong leverage" hypothesis. Our null hypothesis is that the conditional distribution of volatility given negative returns and past volatility stochastically dominates in the first order sense the distribution of volatility given positive returns and past volatility. This hypothesis is stronger in some sense than those considered previously since we refer to the distribution rather than just the mean of the outcome.[1] If our null hypothesis is satisfied then any investor who values volatility negatively would prefer the distribution of volatility that arises after positive shocks to returns to the distribution that arises after negative shocks (Levy, 2006). A further advantage of formulating our hypothesis in terms of distributions is that the tests are less sensitive to the existence of moments. A lot of informal evidence around the leverage effect is reported based on cross correlations between squared returns and lags and leads of returns, see for example Bouchaud et al. (2001). As Mikosch and Starica (2000) have shown, the asymptotic behavior of sample correlograms can be badly affected by heavy tails, which themselves have been widely documented in daily stock returns. Therefore, confidence intervals and hypothesis tests under these circumstances need to be evaluated with care. Our distribution theory builds on the work of Linton et al. (2005) who considered tests of unconditional stochastic dominance for time series data. Linton et al. (2010) consider conditional dominance tests but inside specific semiparametric models. We allow for a general stationary and mixing process for both returns and volatility and impose some smoothness conditions needed for our asymptotic approximations, but otherwise our test is model-free. We obtain the limiting distribution of our test statistic: it is a functional of a Gaussian process. Since the limit distribution depends in a complicated way on nuisance parameters, we propose an inference method based on subsampling (Politis and Romano, 1994). Our test is consistent against a general class of alternatives.

A key part of our methodology is volatility, and we work with ex post volatility that is estimated from high frequency data. Our asymptotic framework requires $n \to \infty$ and $T \to \infty$, where n denotes the number of high-frequency intra-period returns used to compute the realized variance in every period, and T denotes the number of low-frequency time-periods used in the estimation of the test statistic. We derive the limiting distribution of the estimated coefficients under this double asymptotic framework.[2] We find that under fairly strong conditions on n and T, the estimates are $\sqrt{T}$-consistent and have the standard distribution as when there is no measurement-error. However, if the above condition is not satisfied, there is an asymptotic bias that would invalidate this approximation. In that case, we find that under weaker conditions on n and T, a bias-corrected estimator has the standard limiting distribution. This improvement is particularly relevant in the empirical case we examine where n is quite modest. The above is an important methodological contribution to the extant literature on high-frequency volatility estimation. Most work has currently been about just estimating that quantity itself and using it to compare discrete time models in settings where the noise is small. Our approach is concerned with small sample issues when using estimated realized volatility as regressors in the estimation of parameters associated with the unobserved quadratic variation. This involves a useful extension of the existing asymptotic results for realized volatility[3] concerned with the uniformity of the estimation error. We establish the properties of the parameter estimates and propose a bias correction in the case where the estimation error is large. Our methodology sits between discrete time econometrics and continuous time econometrics, since we use concepts from both literatures. If the volatility measure we use can be interpreted as an unbiased estimator of ex ante volatility, then our hypothesis can be interpreted inside the typical discrete time framework.

We apply our testing methodology to stock returns. We focus on whether there is a leverage effect between daily volatility and daily lagged returns on the S&P500 (cash) index and on individual stocks. The stocks we consider are five constituents of the Dow Jones Industrial Average. The sample period covers 1993 to the end of 2009, which includes several very volatile episodes as well as some more tranquil ones. In our main empirical analysis, we measure daily volatility using realized volatility (computed from one minute and five minute intraday transactions data) and a realized intraday range estimator, which only requires daily high and low prices. These data are widely available both for indexes and individual stocks. Dufour et al. (2012) in their study of S&P500 futures data used also the VIX index of implied volatility but this type of traded volatility instruments are not available for individual stocks for the long time span we consider. We find little evidence against the strong leverage effect in these data. We also carry out several robustness checks, including using different volatility estimators, different sample periods, different conditioning values, and both with and without an explicit bias correction method. Our main conclusions survive in all cases. In addition, we also conduct intensive simulations to investigate how our testing methodology performs. We find our proposed test statistics work well on detecting the conditional leverage effect under various situations (see Appendix E). Finally, we compare our results with those obtained from three alternative approaches, which include: (1) Estimating HAR-RV type models with the leverage effect, and two newly developed methods for estimating the leverage effect parameter proposed by (2) Wang and Mykland (2014) and (3) Aït-Sahalia et al. (2013). Due to the generality of our approach we can cannnot explicitly quantify the magnitude of the leverage effect,

[1] Although Wang and Mykland (2014) also allow for the leverage effect to be defined through any (given) function F of volatility.

[2] Corradi and Distaso (2006) use realized variance estimators to test for the correct specification of the functional form of the volatility process within the class of eigenfunction stochastic volatility models. The procedure is based on the comparison of the moments of realized volatility measures with the corresponding ones of integrated volatility implied by the model under the null hypothesis. They allow for measurement error in the realized variance and consider an asymptotic framework similar to ours.

[3] See Barndorff-Nielsen and Shephard (2002).