2-1-2015

Economic Value of Modeling Covariance Asymmetry for Mixed-Asset Portfolio Diversifications

Jian Zhou
University of Guelph

Joseph Nicholson
Montclair State University, nicholsonj@mail.montclair.edu

Economic Modelling 45 (2015) 14–21

Contents lists available at ScienceDirect

Economic Modelling

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

Economic value of modeling covariance asymmetry for mixed-asset portfolio diversifications

Jian Zhou [a,*], Joseph R. Nicholson [b,1]

[a] Real Estate and Housing, College of Management and Economics, University of Guelph, Guelph, ON, N1G 2W1, Canada
[b] School of Business, Montclair State University, 1 Normal Ave, Partridge Hall 429, Montclair, NJ 07043, United States

ARTICLE INFO

Article history:
Accepted 31 October 2014
Available online 15 November 2014

Keywords:
Covariance asymmetry
Mixed-asset portfolio
Economic value

ABSTRACT

Mounting evidence from the literature points to the existence of covariance asymmetry for financial assets. That is, conditional volatility and correlation of financial returns tend to rise more after negative return shocks than after positive ones of the same size. This paper extends the literature by investigating whether investors could gain significant economic benefits from incorporating the feature into mixed-asset portfolio diversifications. We carry out the investigation for a portfolio consisting of US stock, REITs, and the risk-free asset, and find that covariance asymmetry is indeed a value-added feature for mixed-asset diversifications. This conclusion is robust to different portfolio performance metrics and asset allocation periods. More importantly, we demonstrate that the value added by modeling covariance asymmetry is unlikely to be offset by transaction costs. This leads credence to the implementability of a portfolio strategy which embeds the feature of covariance asymmetry. Our findings have important implications for fund managers and their clientele.

1. Introduction

Understanding the dynamics of volatility and correlations for financial returns is important for many financial tasks (e.g. portfolio diversification, risk management and asset pricing, etc.). This has motivated the development of a large number of econometric models and the associated empirical investigations. Interested readers may refer to Bauwens et al. (2012) for an up-to-date overview of the broad finance literature. From the large literature, one of the most salient findings suggests that both conditional volatility and correlation display asymmetric response to return shocks: they tend to rise more after negative return shocks than after positive ones of the same size (e.g. Nelson, 1991; Glosten et al., 1993; Longin and Solnik, 2001; Cappiello et al., 2006; etc.). This phenomenon is typically referred to as covariance asymmetry, due to the fact that volatility and correlation are the two constituents of covariance and both of them respond asymmetrically to financial innovations.

Given this finding, a natural question arises: what financial implications does covariance asymmetry have for fund managers and their investor clientele? In particular, could they reap tangible economic benefits by accounting for covariance asymmetry in portfolio constructions? And if so, how much the benefits would be? These questions are important. While the literature has widely explored the existence of covariance asymmetry and the econometric modeling of it, few studies have assessed the potential economic value that fund managers and their clientele could gain from incorporating the feature into portfolio decisions. Admittedly, documenting the existence of covariance asymmetry is a good first step, but such analysis *per se* is not particularly informative to investors as it falls short of answering whether there are significant economic gains from modeling covariance asymmetry. This paper takes an asset allocation perspective and aims to contribute to the literature along several dimensions. First, we will investigate whether covariance asymmetry is a value-added feature for mixed-asset portfolio diversifications. This is different than previous studies (e.g. Patton, 2004; Thorp and Milunovich, 2007) which focus on all-equity portfolios. We want to see if the feature of covariance asymmetry would bring different magnitudes of value for a mixed-asset portfolio than for just an all-equity portfolio. Second, to ensure the robustness of our findings, we will use a variety of metrics to evaluate the potential economic value added by considering covariance asymmetry. We will also discuss both the economic and statistical significance of the value. Third, we will examine the impact of transaction costs. This is a critical question: if the value added turned out insufficient to cover the higher transaction costs incurred by modeling covariance asymmetry, then it would be moot for fund managers and their clients to consider the feature. Unfortunately, this question has been consistently neglected in the relevant literature. Finally, we will explore whether our above findings are sensitive to asset allocation periodsanother missing point from the literature.

* Corresponding author. Tel.: +1 519 824 4120x56634; fax: +1 519 823 1964.
E-mail addresses: jian@uoguelph.ca (J. Zhou), nicholsonj@mail.montclair.edu (J.R. Nicholson).
[1] Tel.: +1 973 655 7514; fax: +1 973 655 7629.

http://dx.doi.org/10.1016/j.econmod.2014.10.054

To carry out the investigation, we consider a portfolio consisting of stocks, Real Estate Investment Trusts (REITs), and risk-free assets. We do not include bonds, because they do not display a strong feature of covariance asymmetry (e.g. Cappiello et al., 2006). REITs are included for two reasons: first, REITs, along with stocks, are rich in the feature of covariance asymmetry (e.g. Hung and Glascock, 2010; Liow, 2012; Yang et al., 2012; Zhou and Kang, 2011; etc.); second, they are a distinctive investment alternative to stocks by allowing easy access to real estate investments without directly owning or managing the underlying assets. Over the last two decades, REITs have experienced rapid market expansion and have attracted increasing attention from fund managers (Chandrashekaran, 1999). To model covariance asymmetry for this mixed-asset portfolio, we use GJR-ADCC (Glosten et al.'s (1993) GARCH–Asymmetric Dynamic Conditional Correlation of Cappiello et al., 2006). As is shown later, this multivariate GARCH model captures asymmetry in both volatility and correlation. It also accommodates all stylized facts for financial returns such as volatility clustering, and time-variations in conditional volatility and correlation. In contrast, GARCH-DCC—a nested model of GJR-ADCC (Generalized Autoregressive Conditional Heteroscedasticity–Dynamic Conditional Correlation of Engle, 2002) neglects covariance asymmetry, even though this nested model captures all other stylized facts as mentioned above. By applying both methods to a same asset allocation problem, we expect to evaluate the economic value of modeling covariance asymmetry.

We consider a risk-averse investor who forms portfolios by minimizing variance subject to a target return. We use S&P 500 Index, FTSE/NAREIT All Equity REITs Index, and the 3-month Treasury bill rates to respectively represent the three asset classes. We obtain data from January 2, 2007 through December 31, 2012. The whole sample is then divided into two periods: an estimation period (January 2, 2007 to December 31, 2010; 1000 observations) and a testing or asset allocation period (January 3, 2011 to December 31, 2012; 500 observations).We then use a recursive procedure to construct portfolios over the testing period. Overall we find that modeling covariance asymmetry yields significant economic value for mixed-asset portfolio diversifications, and the added value seems to be greater than what has been previously reported for an all-equity portfolio. More importantly, we show that the added value is unlikely to be offset by transaction costs. These results are found to hold for a different testing period (i.e. the year of 2012). This implies that covariance asymmetry is indeed an implementable value-added feature for portfolio. Our findings should benefit both fund managers and their investor clientele.

The reminder of this paper is organized as follows. Section 2 outlines the econometric methodologies. Section 3 discusses the data. Section 4 presents the empirical findings. Section 5 concludes.

2. Econometric methodologies

2.1. The asset allocation strategy

We consider an investor who allocates funds across assets by minimizing portfolio variance subject to a target return constraint. Let $\boldsymbol{\mu}$ be a vector of expected excess returns $(\mathbf{r}_{t+1} - r_f\mathbf{1})$, where $\mathbf{r}_{t+1}$ is a vector of expected returns of k risky assets, r_f is there turn of risk-free asset, and $\mathbf{1}$ is a vector of ones. Then the asset allocation problem can be written as

$$\min_{\mathbf{w}_t} \mathbf{w}_t'\boldsymbol{\Sigma}_t\mathbf{w}_t \tag{1}$$

$$s.t. \mathbf{w}_t'\boldsymbol{\mu} = \mu_p \tag{2}$$

where $\boldsymbol{\Sigma}_t \equiv E_t[(\mathbf{r}_{t+1} - \boldsymbol{\mu})(\mathbf{r}_{t+1} - \boldsymbol{\mu})']$ is the expected covariance matrix, μ_p is the target return, and $\mathbf{w}_t$ is a $k \times 1$ vector of weights on the risky assets. The solution to this optimization problem is

$$\mathbf{w}_t = \frac{\mu_p \boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}}{\boldsymbol{\mu}'\boldsymbol{\Sigma}_t^{-1}\boldsymbol{\mu}}. \tag{3}$$

Note that we do not impose short-sales constraints so that any wealth not accounted for by $\mathbf{w}_t$ is implicitly invested in the risk-free asset, which has a weight of $(1 - \mathbf{w}_t'\mathbf{1})$.

2.2. Forecasting the conditional covariance

Implementing the above asset allocation strategy requires estimating $\boldsymbol{\mu}$ and $\boldsymbol{\Sigma}_t$. To simplify our analysis, we follow Fleming et al. (2001) by using in-sample mean return to estimate $\boldsymbol{\mu}$.Doing so allows us to focus solely on the impact of covariance structures. Another reason is that expected returns are typically estimated with far less precision than expected covariance matrices (Merton, 1980). So in what follows we mainly discuss how to estimate $\boldsymbol{\Sigma}_t$.

As a benchmark model for $\boldsymbol{\Sigma}_t$, we use the GARCH-DCC (Dynamic Conditional Correlation) model of Engle (2002). This model has been widely used. It is capable of capturing certain stylized facts of covariance structure such as volatility clustering and dynamic correlations but ignores the feature of covariance asymmetry. As an alternative that can model covariance asymmetry, we resort to a model named GJR-ADCC. On the one hand, GJR, representing the GARCH model of Glosten et al. (1993), can capture volatility asymmetry—the first layer of covariance asymmetry. A generic GJR model has the following specification:

$$h_{i,t} = \varpi_i + \alpha_i\varepsilon_{i,t-1}^2 + \gamma_i I_{t-1}\varepsilon_{i,t-1}^2 + \beta_i h_{i,t-1} \tag{4}$$

where $h_{i,t}$ is the conditional volatility for each asset of the portfolio, ε is the demeaned returns, $I_{t-1} = I(\varepsilon_{t-1} < 0)(I(\cdot)$ is an indicator function which takes on value 1 if the argument is true and 0 otherwise), and ϖ, α, γ, & β are parameters. It is easy to see that volatility asymmetry is modeled through γ, as a positive value of γ would indicate that negative return shocks generate higher volatility than positive ones of the same magnitude. γ is thus the parameter of volatility asymmetry. Setting $\gamma = 0$ reduces GJR to the standard GARCH, which neglects volatility asymmetry. On the other hand, ADCC, representing Asymmetric Dynamic Conditional Correlation model, captures correlation asymmetry—the second layer of covariance asymmetry. A generic ADDC model can be written as:

$$\boldsymbol{\Sigma}_t = \mathbf{D}_t\mathbf{R}_t\mathbf{D}_t \tag{5}$$

where $\mathbf{D}_t$ is a $k \times k$ diagonal matrix with $\sqrt{h_{i,t}}$ on the ith diagonal, and $\mathbf{R}_t$ is the correlation matrix to be estimated. According to Cappiello et al. (2006), $\mathbf{R}_t$ can be formulated as follows:

$$\mathbf{R}_t = diag(\mathbf{Q}_t)^{-1}\mathbf{Q}_t diag(\mathbf{Q}_t)^{-1} \tag{6}$$

$$\mathbf{Q}_t = (1-a-b)\overline{\mathbf{Q}} - \phi\overline{\mathbf{N}} + a(\mathbf{z}_{t-1}\mathbf{z}'_{t-1}) + b\mathbf{Q}_{t-1} + \phi(\boldsymbol{\eta}_{t-1}\boldsymbol{\eta}'_{t-1}) \tag{7}$$

where $diag(\mathbf{Q}_t) = \left[\sqrt{q_{i,i,t}}\right]$ is a diagonal matrix containing the square root of the diagonal elements of matrix $\mathbf{Q}_t$, $\mathbf{z}_t$ is the vector of standardized residuals (i.e. $z_{i,t} = \varepsilon_{i,t}/\sqrt{h_{i,t}}$), $\overline{\mathbf{Q}} = E(\mathbf{z}_t\mathbf{z}'_t)$, $\boldsymbol{\eta}_t = I(\boldsymbol{\varepsilon}_t < 0) \circ \boldsymbol{\varepsilon}_t$ ($\circ$ denotes the Hadamard product), $\overline{\mathbf{N}} = E(\boldsymbol{\eta}_t\boldsymbol{\eta}'_t)$ and a, b, and ϕ are nonnegative scalar parameters. It is worth noting that ϕ captures correlation asymmetry—the second layer of covariance asymmetry, as a positive value of ϕ would suggest that cross-market correlations between equity returns increase more during market downturns than during market upturns. ϕ is thus the parameter of correlation asymmetry. Setting $\phi = 0$ reduces ADCC to the standard DCC, which ignores correlation asymmetry.

With the setup of Eqs. (4)–(7), we first use a maximum likelihood estimator to estimate the parameters and then proceed to generate the expected value of the covariance matrix (i.e. one-step-ahead forecast of $\mathbf{\Sigma}_t$). We need to forecast covariance matrix because it is the expected value of the covariance matrix that is ultimately used for asset allocation (see Eq. (3)). To forecast $\mathbf{\Sigma}_t$, we follow Eq. (5) by first obtaining a forecast of $\mathbf{D}_t$ through Eq. (4) and then having a forecast of $\mathbf{R}_t$ based on Eqs. (7) and (6). It is worth noting that during this process, we generate two sets of covariance forecasts: one by using GARCH-DCC, which ignores covariance asymmetry and the other by using GJR-ADCC, which considers covariance asymmetry.

2.3. Measuring portfolio performances

To compare the performances of the portfolios constructed based on GARCH-DCC and GJR-ADCC, we resort to a variety of metrics. The first one is the standard Sharpe ratio (SR), which can be computed by applying the estimated optimal weights (Eq. (3)) to the actual asset returns. Denote μ_p and σ_p respectively the mean and volatility of the realized portfolio returns. We have $SR = (\mu_p - r_f)/\sigma_p$. By comparing SRs from the two portfolio-forming strategies, we can tell if one outperforms the other. However, such a comparison may underestimate the outperformance of one strategy over the other, because the comparison does not consider the potentially different risk levels of the two competing strategies. This issue can be addressed by using the modified Sharpe ratio (mSR) of Graham and Harvey (1997). mSR measures the abnormal return that a strategy would have earned if it had the same risk as the benchmark strategy. This measure is directly related to the estimated Sharpe ratios. Assume GARCH-DCC is our benchmark strategy and GJR-ADCC is the alternative. mSR is then written as $mSR = \sigma_p^b(SR_p^a - SR_p^b)$, where the superscript 'a' and 'b' denote respectively the alternative and benchmark strategies.

A third measure we consider is the performance fee (Δ) proposed by West et al. (1993) and Fleming et al. (2001). As a popular measure of the economic value of different asset allocations, Δ estimates the management fee that an investor would be willing to pay to switch from a benchmark portfolio strategy (GARCH-DCC) to an alternative strategy (GJR-ADCC) without being worse off in terms of utility. Δ can be calculated as the average return that can be subtracted from the portfolio return of the alternative strategy, such that the investor becomes indifferent when choosing between the two competing strategies. By assuming quadratic utility, Δ can be estimated by solving the following equation (Fleming et al., 2001):

$$\sum_{t=0}^{T-1}\left[\left(r_{p,t+1}^a-\Delta\right)-\frac{\gamma}{2(1+\gamma)}\left(r_{p,t+1}^a-\Delta\right)^2\right] = \sum_{t=0}^{T-1}\left[r_{p,t+1}^b-\frac{\gamma}{2(1+\gamma)}\left(r_{p,t+1}^b\right)^2\right] \tag{8}$$

where T is the sample size, r_p is portfolio return, and γ is the degree of relative risk aversion.

An important factor for assessing portfolio performances is transaction costs. Suppose we find the GJR-ADCC strategy leads to positive economic value over the GARCH-DCC one. However, if that value turned out insufficient to cover the possibly higher transaction costs incurred by the GJR-ADCC strategy, then the GJR-ADCC strategy would not be implementable. As such, we need to check into transaction costs. Transaction costs (TC) can be assumed to be a fixed fraction (τ) of the value traded (VT) for all assets in the portfolio, namely $TC = \tau * VT$. According to Han (2006),

$$VT = \sum_{t=1}^{T}\sum_{i=1}^{k}\left(w_{i,t}-\frac{w_{i,t-1}\left(1+r_{i,t}\right)}{1+r_{p,t}}\right). \tag{9}$$

Note that the investors' initial wealth is normalized to be 1. Given Eq. (9), if the alternative strategy outperforms the benchmark but at the same time enjoys a lower VT (i.e. lower turnover of assets), then we would say that we can reap the gains caused by the alternative strategy without necessarily paying more transaction costs. This directly leads credence to the implementability of the GJR-ADCC strategy. On the other hand, if the alternative has a higher VT than the benchmark, we then compute the breakeven transaction cost τ^{be}, which measures the level of transaction costs required to make the investor indifferent when choosing between the alternative and benchmark strategy. τ^{be} is calculated as the proportional cost that cancels out the utility advantage (and hence the positive performance fee) of the outperforming alternative strategy. If τ^{be} is found to be greater than actual transaction costs, it implies that transaction costs cannot offset value. It is thus overall beneficial to consider covariance asymmetry in portfolio constructions. Otherwise, it does not make economic sense to consider this feature.

3. Data

We consider portfolios consisting of stock, equity REITs and a risk-free asset. We use S&P 500 Index to represent stock, and FTSE/NAREIT All Equity REITs Index to represent REITs. Both indices are collected from DataStream. We use the 3-month Treasury bill rates from the Federal Reserve Bank of St. Louis to represent the risk-free asset. The data are daily and cover the period from January 2, 2007 through December 31, 2012 for a total of 1500 observations. Table 1 reports the summary statistics of the asset returns. Stock has a positive mean return while REITs have a negative mean return. At the same time, the REIT market appears to be more volatile than the stock market. Other interesting points suggest that all returns are skewed, and exhibit significant excess kurtosis (i.e. fat-tailness). As such, normal distribution is found inappropriate to describe returns. Unit-root tests indicate that all return series are stationary. Fig. 1 presents the time-series plots of the stock and REIT returns. There are large swings which periodically occur over the sample period. The most pronounced ones took place from mid-2008 to mid-2009. This apparently corresponds to the recent global financial crisis.

4. Empirical findings

4.1. Preliminary analysis

To avoid in-sample overfitting, we divide the whole sample into two subsamples: an estimation period (January 2, 2007 to December 31,

Table 1
Summary statistics.

Panel A:	Mean	Std. Dev.	Skewness	Kurtosis	Normality	Unit-root
Stock	0.001	1.575	−0.343*	7.385*	3.435 *	−22.712*
REITs	−0.010	2.898	−0.116**	8.097*	4.098 *	−24.175*
Risk-free	0.004	0.006	1.527*	0.779*	0.621 *	−3.047*

Notes: This table reports the summary statistics of daily returns for stock, REITs and risk-free asset over the period January 2007 to December 2012. Stock is represented by S&P 500 Index, REITs by FTSE/NAREIT All Equity REITs Index, risk-free by 3-month Treasury bill. Mean and Std. Dev. (standard deviation) are expressed in the percentage form. Normality is tested for using the Jarque–Bera test and the test statistics are expressed in the unit of 1000. Unit-root is tested for using the ADF (Augmented Dickey–Fuller) test, under which the null hypothesis H_0 is that the series of interest is $I(1)$.

* Significant at 5% level.
** Significant at 10% level.

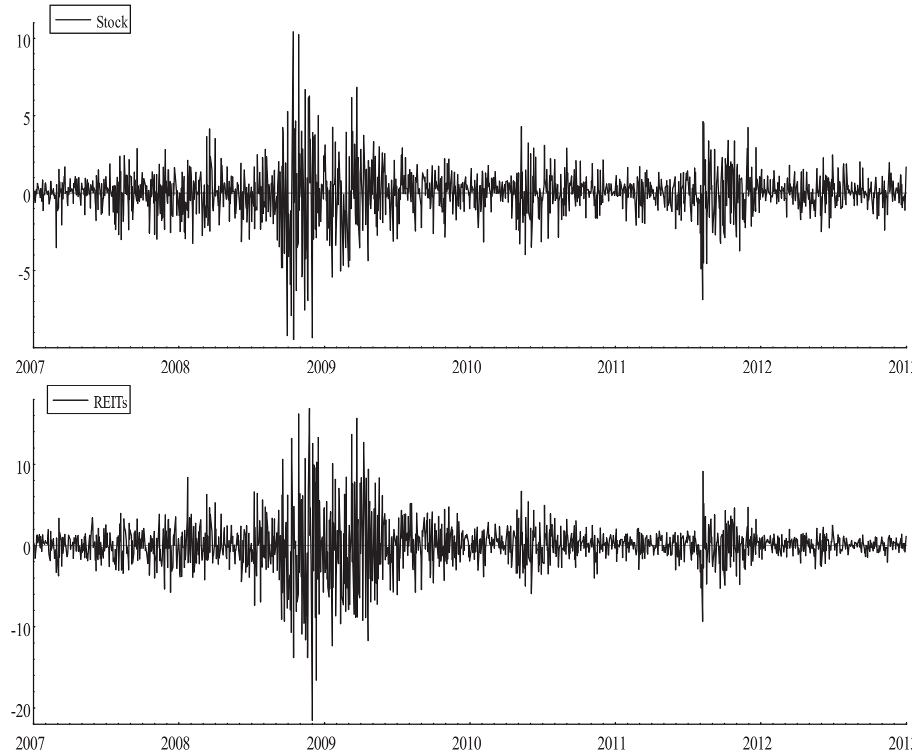

Fig. 1. Time series plots of daily asset returns. The sample period is from January 2007 to December 2012.

2010; 1000 observations) and a testing or asset allocation period (January 3, 2011 to December 31, 2012; 500 observations). To generate one-step-ahead forecasts of the conditional covariance, we adopt a recursive estimation procedure: the initial estimation period January 2, 2007 to December 31, 2010 is used to produce the first forecast (i.e. January 3, 2011) for the testing period, and then the models are re-estimated

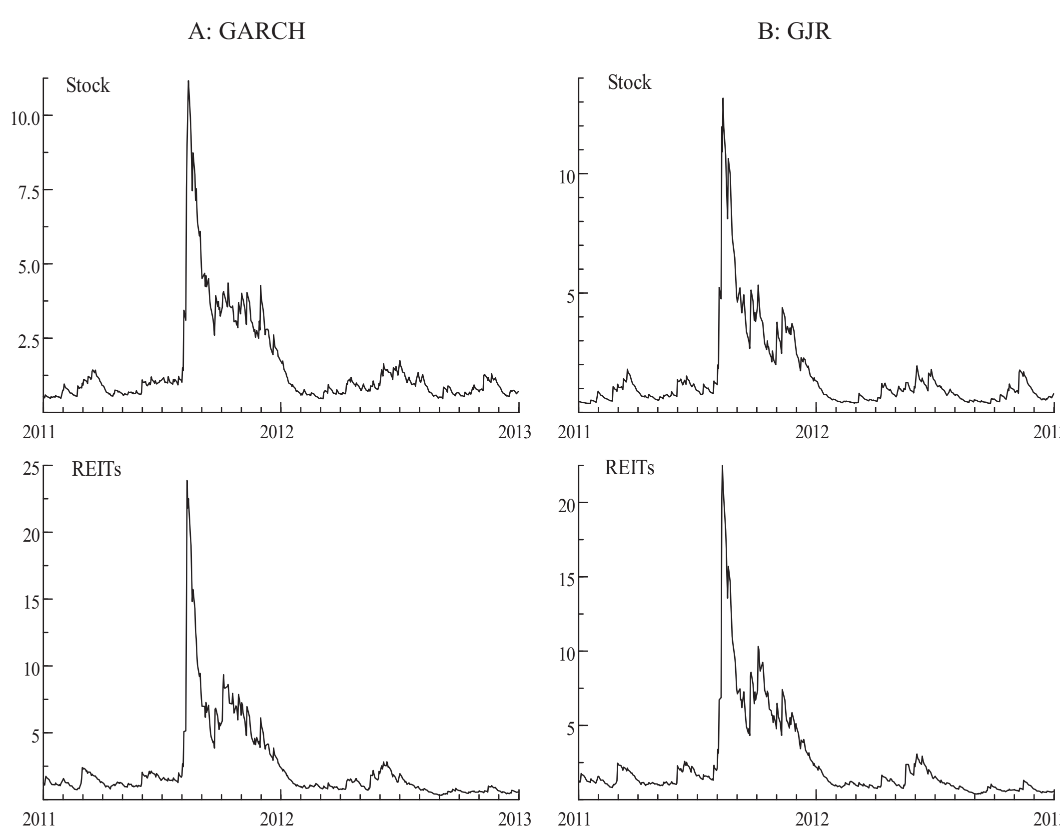

Fig. 2. Time series plots of forecasted conditional volatilities over the testing period January 2011 to December 2012. The forecasting is conducted based on Eq. (4): $h_{i,t} = \varpi_i + \alpha_i\varepsilon^2_{i,t-1} + \gamma_i I_{t-1}\varepsilon^2_{i,t-1} + \beta_i h_{i,t-1}$, which is the general specification of GJR. GARCH is a special case ($\gamma = 0$).

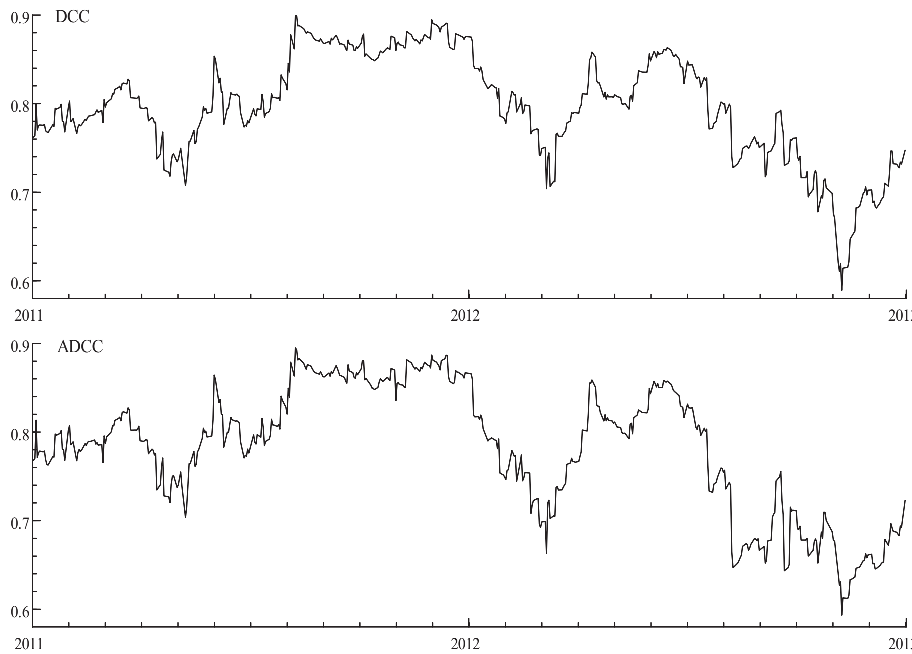

Fig. 3. Time series plots of forecasted conditional correlations over the testing period January 2011 to December 2012. The forecasting is conducted based on Eq. (6): $\mathbf{R}_t = diag(\mathbf{Q}_t)^{-1}\mathbf{Q}_t diag(\mathbf{Q}_t)^{-1}$ and Eq. (7): $\mathbf{Q}_t = (1-a-b)\overline{\mathbf{Q}} - \phi\overline{\mathbf{N}} + a(\mathbf{z}_{t-1}\mathbf{z'}_{t-1}) + b\mathbf{Q}_{t-1} + \phi(\mathbf{\eta}_{t-1}\mathbf{\eta'}_{t-1})$. The two equations consist of the general specification of ADCC. DCC is a special case ($\phi = 0$).

using the sample period January 2, 2007 to January 3, 2011 to produce the next forecast for January 4, 2011, and so on throughout the remainder of the testing period.

Figs. 2 and 3 plot the forecasts of conditional volatilities and correlations—the two decomposition elements of the covariance matrix (see Eq. (5)). We observe a large spike in the volatilities over the second half of 2011. This corresponds well to the tumultuous nature of the year due to fears of recession and a currency crisis in Europe. Interestingly, GJR, thanks to its ability to capture volatility asymmetry, predicts higher levels of volatility than GARCH during those abnormal periods (especially true for stock). Regarding conditional correlations, both DCC and ADCC predict positive correlations across the markets, even

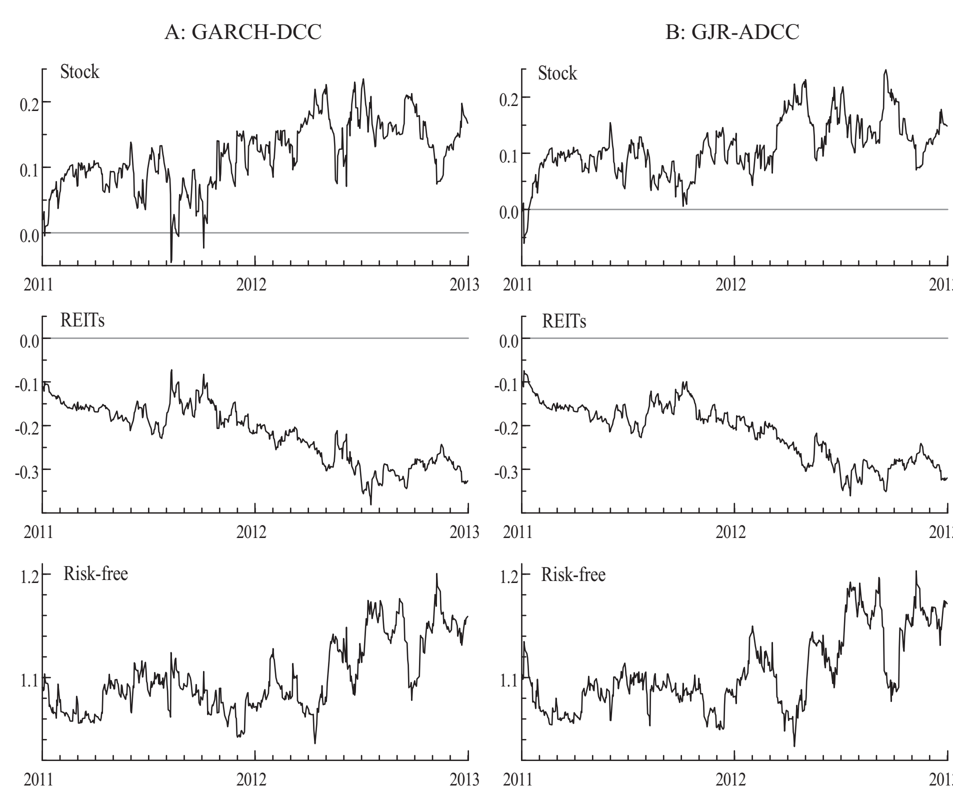

Fig. 4. Time series plots of optimal weights given a 6% target return over the testing period January 2011 to December 2012.

Table 2
Comparisons of ex post Sharpe ratios.

Target	GARCH-DCC			GJR-ADCC			Difference in *SR*	Testing
Return	μ_p	σ_p	*SR*	μ_p	σ_p	*SR*		
5%	5.797	22.731	0.258	6.463	22.487	0.291	0.033	0.874 (0.402)
6%	6.971	27.277	0.258	7.770	26.984	0.291	0.033	0.901 (0.378)
7%	8.145	31.823	0.258	9.076	31.482	0.291	0.033	0.929 (0.388)
8%	9.318	36.369	0.256	10.383	35.979	0.289	0.033	0.954 (0.353)
9%	10.492	40.915	0.256	11.690	40.476	0.289	0.033	0.975 (0.349)
10%	11.665	45.462	0.257	12.996	44.974	0.289	0.032	0.986 (0.342)

Notes: This table reports the results for ex post Sharpe ratios. μ_p is the annualized mean ex post portfolio returns, σ_p is the annualized standard deviation of portfolio returns, and *SR* is the ex post Sharp ratio. The last column reports the test statistics of Ledoit and Wolf's (2008) robust test for difference in *SR*. In parentheses is the *p*-value for the null hypothesis of zero difference in *SR*. The results reported are for the testing period of January 2011 to December 2012.

though the strengths of the correlations appear to go up and down over time. Applying the forecasted covariance to Eq. (3) would yield the optimal weights given a target return. Fig. 4 plots the estimated weights for a 6% target return.[2] As expected, the sign and magnitude of the weights depend on the estimates of expected returns, conditional volatility and correlation. Because the mean return of REITs is mostly negative over the estimation period (see Table 1), REITs have a negative weight (i.e. being shorted). By the same reasoning, the weight of stock is mostly positive. We find that the estimated weights display marked time variations. For example, volatilities in both markets appeared to be lower in 2012 than in 2011. Given our minimum-volatility portfolio construction strategy (Eq. (1)), this caused the weights of stock (long) and REITs (short) to both increase over time. Note that the risk-free asset is treated in the portfolio as a 'residual' asset. Its weight is simply one minus the sum of the weights of the other two assets. We thus find it is being longed (i.e. investor lends) over the full testing period. Finally, it is worth noting that even though the weight plots are synchronized across the two panels, there still exist noticeable disparities, as a careful visual check can reveal. The differences in weight paths are expected to have important implications for asset allocations, as is discussed below.

4.2. Economic evaluation results

We multiply the above estimated weights by the observed asset returns over the testing period to compute the ex post (realized) portfolio returns. Table 2 reports the annualized mean ex post portfolio return (μ_p), and the standard deviation (σ_p), and the Sharpe ratio (*SR*). Compared with the GARCH-DCC strategy, the GJR-ADCC strategy is found to deliver consistently higher portfolio returns with lower standard deviations. As a result, it yields higher ex post *SR*s than the benchmark GARCH-DCC. But the differences in *SR*s seem pretty small—around 3 basis points per year. Given this, it is natural to ask if the differences are significant. To check out, we employ the test of Ledoit and Wolf (2008). This test is developed to test for the statistical significance of the difference in the Sharpe ratios of two competing strategies. Compared with the traditional Sharpe ratio tests (e.g. Jobson and Korkie, 1981), this test is augmented in two ways: one, it can be applied to returns with heavier tails than the normal distribution. This feature is important, given that fat-tailness is found to be a stylized fact of financial returns and it happens to be the case here; two, this test achieves robust finite sample performances by using a time series bootstrap to construct the standard errors for the difference in the Sharpe ratios. The last column of Table 2 reports the statistics of the Ledoit and Wolf (2008) test. It appears that the differences in *SR*s are statistically insignificant. So in terms of Sharpe ratios, GJR-ADCC seems not to outperform GARCH-DCC.

However, a caveat should be issued because a direct comparison of the Sharpe ratios does not take into account the potentially different levels of risks of the two competing strategies. As discussed in Section 2.3, a more appropriate measure to use is the modified Sharpe ratio (*mSR*), which measures the abnormal return that a strategy would have earned if it had the same risk as the benchmark. Table 3 reports the results for *mSR*. Also presented are the performance fees (Δ), which measure the management fees that the investor would be willing to pay to switch from the benchmark strategy to an alternative strategy without being worse off in terms of utility levels. We calculate Δ for three levels of $\gamma = 1$, 5 or 10, where γ is investors' degree of relative risk aversion. As shown in Table 3, the estimated *mSR* and Δ are all positive and fairly large. This lends preliminary evidence to the outperformance of the alternative over the benchmark.

Before we make any further interpretation of the results, it is important to examine whether the estimated performance measures are statistically positive. This issue has been often neglected in the finance literature (e.g. Fleming et al., 2001; Han, 2006; etc.). We formulate a null hypothesis that the performance measure is negative (i.e. $mSR < 0$ or $\Delta < 0$). To test the null, the empirical distribution of the performance measures needs to be constructed. We turn to the bootstrap technique. Because the bootstrap resampling is based on the stationary time series of the realized portfolio returns,[3] we adopt the stationary bootstrap procedure of Politis and Romano (1994). This procedure has the advantage of retaining the stationarity property for the resampled series.[4] The bootstrapped *p*-values for the null are presented in the parentheses of Table 3. It is easy to see that the null hypothesis of the relevant performance measure being negative is rejected in every case, thus indicating that the estimated performance measures are statistically positive. This adds further credibility to the economic benefits of including covariance asymmetry for portfolio constructions. For example, for investors with a 5% target return, using the GJR-ADCC strategy would earn them an annualized abnormal return of 73.634 basis points over using the GARCH-DCC strategy. At the same time, those investors would pay an annual management fee of around 42.083 basis points to switch from the benchmark strategy to the alternative strategy, assume that they have a low degree of risk aversion ($\gamma = 1$). Also as γ rises, the estimated Δ rises. This implies that as investors become more risk averse, they would value more of the benefits brought by the GJR-ADCC strategy. The results of Table 3 also indicate that the estimated *mSR* and Δ both increase as investors become more aggressive in investment objectives (i.e. higher target return). This suggests that the economic gains offered by the inclusion of covariance asymmetry increase with the perceived riskiness of the investment objectives.

Finally, an important factor in portfolio construction is transaction costs. We need to examine whether the positive economic gains found

[2] The weight plots for other levels of target returns show similar patterns. To save space, they are not presented but available upon request.

[3] We apply the ADF (Augmented Dickey–Fuller) test to the series of realized portfolio returns estimated using the two models. The results strongly indicate that all seriesare unit-root-free. Testing results can be obtained from us.

[4] The stationary bootstrap is a block bootstrap with block lengths that are distributed as a geometric (q) random variable. The average block length is $1/q$. We choose q by following the procedure of Politis and White (2004). We bootstrap 10,000 times. For more details of the stationary bootstrap and the determination of its block length, refer to Politis and Romano (1994) and Politis and White (2004).

Table 3
Estimates of the modified Sharpe ratios and performance fees.

Target return	*mSR*	Δ_1	Δ_5	Δ_{10}
5%	73.634*	42.083*	112.801*	128.663*
	(0.036)	(0.045)	(0.036)	(0.027)
6%	88.361*	76.137*	177.095*	199.365*
	(0.039)	(0.046)	(0.039)	(0.025)
7%	103.087*	118.530*	254.735*	285.064*
	(0.037)	(0.043)	(0.044)	(0.027)
8%	117.814*	169.246*	345.454*	384.544*
	(0.037)	(0.043)	(0.044)	(0.020)
9%	132.540*	228.129*	449.023*	498.866*
	(0.038)	(0.047)	(0.040)	(0.034)
10%	147.267*	295.113*	565.014*	624.485*
	(0.040)	(0.047)	(0.040)	(0.025)

Notes: This table reports the estimates of the modified Sharpe ratios (*mSR*) and the performance fees (Δ_γ) ($\gamma = 1, 5$ or 10 is the degree of relative risk aversion). *mSR* and Δ are expressed in annualized basis points(*bps*). In parentheses are the bootstrapped *p*-values for the null hypothesis that the relevant performance measure is negative. The bootstrapping is carried out using the stationary bootstrap procedure of Politis and Romano (1994). The results reported are for the testing period of January 2011 to December 2012.
* Significant at the 5% level.

above can be offset by transaction costs. Table 4 presents the results. Note that the break-even transaction cost τ_γ^{be} is calculated based on the value traded (*VT*) of all assets in the portfolio. If *VT* of the GJR-ADCC strategy is found to fall below that of the GARCH-DCC strategy, it implies that we can reap the gains of considering asymmetry without incurring more transaction costs. A few cases fall into this category (indicated by '–'). For other cases, we calculate τ_γ^{be}. Recall that τ_γ^{be} is calculated as the daily proportional cost (in *bps*) that cancels outs the performance fee. As such, τ^{be} can be interpreted as a measure of the cushion of trading costs. According to Han (2006), a double-digit cushion should be enough to cover the trading costs investors reasonably incur on the equity market while a single-digit cushion raises red flags. As shown, τ^{be} is double-digit in most cases, suggesting that the gains by modeling asymmetry are unlikely to be offset by the actual transaction costs. This leads credence to the implementability of the GJR-GARCH strategy.

4.3. Robustness analysis

As a robustness check, we experiment with a different sample division. We save the last year of our sample (i.e. 2012) as the testing or asset allocation period while using the period January 3, 2011 to December 31, 2011 as the estimation period. Following the aforementioned recursive estimation procedure, we repeat the above analyses. To preserve space, Table 5 only shows the estimation results of the performance fees and break-even transaction costs for a few selected levels of target returns. Other results are available upon request. Based on Table 5, we find that our conclusions still hold: the economic values of covariance asymmetry are all positive and mostly statistically significant; they increase with the investors' degree of risk aversion and their aggressiveness of investment objectives, and transaction costs appear to be a non-issue in most cases. Finally, as an interesting note, we find that the magnitude of the estimated performance fees seems to be slightly affected. Choosing a shorter testing period seems to lead to smaller (but still significant) performance fees, compared with those estimated before for a longer testing period January 2011 to December 2012.

Table 4
Estimates of the break-even transaction costs τ_γ^{be}.

Target return	τ_1^{be}	τ_5^{be}	τ_{10}^{be}
5%	9.300	25.067	28.591
6%	–	–	–
7%	–	–	–
8%	–	–	–
9%	13.419	26.413	29.345
10%	12.832	24.565	27.151

Notes: This table reports the estimates of the break-even transaction costs τ_γ^{be}, where $\gamma = 1, 5$ or 10 is the degree of relative risk aversion. The calculation of τ_γ^{be} is based on the value traded (*VT*) of all assets in the portfolio. When *VT* of the GJR-ADCC strategy is lower than that from the GARCH-DCC strategy, it implies that the GJR-ADCC strategy can be implemented without incurring more transaction costs. For these cases (indicated by '–'), there is no need to calculate τ_γ^{be}. τ_γ^{be} is only calculated when *VT* of the GJR-ADCC strategy exceeds that from the GARCH-DCC strategy. τ^{be} is then defined as the daily proportional cost (in *bps*) required to make the investor indifferent between the two competing strategies. The results reported are for the testing period of January 2011 to December 2012.

Table 5
Robustness check—estimates of Δ_γ and τ_γ^{be} over an alternative testing period.

	$\gamma = 1$	$\gamma = 5$	$\gamma = 10$
Target return $= 5\%$			
Δ_γ	29.101	71.108*	94.460*
	(0.067)	(0.043)	(0.034)
τ_γ^{be}	–	–	–
Target return $= 6\%$			
Δ_γ	45.321*	123.070*	153.063*
	(0.047)	(0.040)	(0.036)
τ_γ^{be}	–	–	–
Target return $= 7\%$			
Δ_γ	65.310*	184.732*	222.048*
	(0.045)	(0.042)	(0.033)
τ_γ^{be}	7.257	20.526	24.672

See notes in Tables 3 and 4. The results reported are for an alternative testing period of year 2012.
* Significant at the 5% level.

4.4. Comparison with prior literature and economic implications

Our results, to some degree, are consistent with those reported for all-equity portfolios (e.g. Patton, 2004; Thorp and Milunovich, 2007): modeling covariance asymmetry yields significant economic value. However, a discrepancy exists in that the economic benefits seem to be greater for mixed-asset portfolio than for an all-equity portfolio. For instance, Thorp and Milunovich (2007), using a much longer testing period (nearly four years) than here, showed that the average management fee does not exceed 100 basis points, whereas Patton (2004) estimated the average management fee to be merely 41.5 basis points using a 10-year testing period. These numbers are much smaller than what we find here: if we average the estimated management fee (Δ) across each row (i.e. across the three levels of $\gamma = 1, 5$ or 10) in either Table 3 or Table 5, in most cases the result easily surpasses 100 basis points. A plausible reason for this discrepancy is that mixed assets like REITs and stock have imperfect covariance (Chandrashekaran, 1999) while strong covariance tends to exist between stocks.

Our paper should appeal to both fund managers and their investor clientele. For investors, modeling covariance asymmetry leads to higher risk-adjusted returns, as demonstrated through the positive and significant modified Sharpe ratio (*mSR*). For managers, their clients would like to pay higher management fees (Δ) because doing so does not cause a drop in the utility level. Moreover, the covariance-asymmetry-embedded portfolio strategy is found implementable, as the added economic value is sufficient to cover transaction costs. Finally, the value added appears robust to the asset allocation period.

5. Conclusions

This paper investigates whether one can significantly improve portfolio performances by modeling the feature of covariance asymmetry, which says that conditional volatility and correlation tend to rise more after negative return shocks than after positive ones of the same magnitude. To carry out the investigation, we use GARCH-DCC and GJR-ADCC to construct portfolios and then compare their performances. GJR-ADCC nests GARCH-DCC. The former incorporates the feature of covariance asymmetry while the latter ignores it. We consider an investor who forms portfolios by minimizing variance subject to a target return. We apply both models to a mixed-asset portfolio consisting of US stock, REITs, and a risk-free asset. Overall our findings suggest that portfolios constructed based on GJR-ADCC significantly outperform those constructed based on GARCH-DCC. Both fund managers and their investor clientele can benefit by considering covariance asymmetry in their portfolio decisions. We reach this conclusion by carefully examining a number of portfolio performance metrics. More importantly, we demonstrate that the incremental value of modeling covariance asymmetry is unlikely to be offset by transaction costs. This lends credence to the implementability of the covariance-asymmetry-embedded portfolio strategy. Lastly, our findings are shown to be robust to different asset allocation periods.

References

Bauwens, Luc, Hafner, C.M., Laurent, S., 2012. Handbook of Volatility Models and Their Applications. Wiley.

Cappiello, L., Engle, R.F., Sheppard, K., 2006. Asymmetric dynamics and the correlations of global equity and bond returns. J. Financ. Econ. 4, 537–594.

Chandrashekaran, V., 1999. Time-series properties and diversification benefits of REIT returns. J. Real Estate Res. 17, 91–112.

Engle, R.F., 2002. Dynamic conditional correlation: a simple class of multivariate GARCH models. J. Bus. Econ. Stat. 20, 339–350.

Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing. J. Financ. 56, 329–352.

Glosten, L., Jagannathan, R., Runkle, D., 1993. On the relationship between the expected value and the volatility of the nominal excess return on stocks. J. Financ. 48, 1779–1801.

Graham, J.R., Harvey, C.R., 1997. Grading the performance of market-timing newsletters. Financ. Anal. J. 53 (6), 54–68.

Han, Y., 2006. Asset allocation with a high dimensional latent factor stochastic volatility model. Rev. Financ. Stud. 19, 237–271.

Hung, S.Y.K., Glascock, J.L., 2010. Volatilities and momentum returns in real estate investment trusts. J. Real Estate Financ. Econ. 41, 126–149.

Jobson, J.D., Korkie, B.M., 1981. Performance hypothesis testing with the Sharpe and Treynor measures. J. Financ. 36, 889–908.

Ledoit, Oliver, Wolf, M., 2008. Robust performance hypothesis testing with the Sharpe ratio. J. Empir. Financ. 15, 850–859.

Liow, K.H., 2012. Co-movements and correlations across asian securitized real estate and stock markets. Real Estate Econ. 40, 97–129.

Longin, F., Solnik, B., 2001. Extreme correlation of international equity markets. J. Financ. 56, 649–676.

Merton, R.C., 1980. On estimating the expected return on the market: an exploratory investigation. J. Financ. Econ. 8, 323–361.

Nelson, D., 1991. Conditional heteroskedasticity in asset returns: a new approach. Econometrica 59, 347–370.

Patton, A., 2004. On the out-of-sample importance of skewness and asymmetric dependence for asset allocation. J. Financ. Econ. 2 (1), 130–168.

Politis, D.N., Romano, J.P., 1994. The stationary bootstrap. J. Am. Stat. Assoc. 89, 1303–1313.

Politis, D.N., White, H., 2004. Automatic block-length selection for the dependent bootstrap. Econ. Rev. 23, 53–70.

Thorp, S., Milunovich, G., 2007. Symmetric versus asymmetric conditional covarianceforecasts: does it pay to switch? J. Financ. Res. 30 (3), 355–377.

West, K.D., Edison, H.J., Cho, D., 1993. A utility-based comparison of some models of exchange rate volatility. J. Int. Econ. 35, 23–45.

Yang, J., Zhou, Y., Leung, W.K., 2012. Asymmetric correlation and volatility dynamics among stock, bond, and securitized real estate markets. J. Real Estate Financ. Econ. 45, 491–521.

Zhou, J., Kang, Z.X., 2011. A comparison of alternative forecast models of REIT volatility. J. Real Estate Financ. Econ. 42, 275–294.