

Realized Daily Variance of S&P 500 Cash Index: A Revaluation of Stylized Facts*

Shirley J. Huang

*Lee Kong Chian School of Business, Singapore Management University 50
Stamford Road, Singapore 178899
E-mail: shirleyhuang@smu.edu.sg*

Qianqiu Liu

*College of Business Administration, University of Hawaii, 2404 Maile Way
Honolulu, HI 96822
E-mail: qianqiu@hawaii.edu*

and

Jun Yu

*School of Economics and Social Sciences, Singapore Management University
90 Stamford Road, Singapore 178903
E-mail: yujun@smu.edu.sg*

In this paper the realized daily variance is obtained from intraday transaction prices of the S&P 500 cash index over the period from January 1993 to December 2004. When constructing realized daily variance, market microstructure noise is taken into account using a technique proposed by Zhang, Mykland and Ait-Sahalia (2005). The time series properties of realized daily variance are compared with those of variance estimates obtained from parametric GARCH and stochastic volatility models. Unconditional and dynamic properties concerning the realized daily variance are examined, the relationship between realized variance and returns is investigated, and the stylized facts concerning realized daily variance are reevaluated with this long dataset. While many properties are similar to what have been reported based on artificially constructed five-minute returns, three distinct results stand out in our empirical analysis. First, we find evidence that both the realized standard deviation and the realized log variance are not covariance stationary, but nonetheless have memory parameter less than unity. Second, we document a positive and statistically significant risk-return trade-off. Finally, we find a monotonically decreasing news impact function.

Key Words: High frequency data; Integrated variance; Microstructure noise; GARCH; Stochastic volatility; Long range dependence; Intertemporal CAPM.

JEL Classification Numbers: G10; C10; C20.

1. INTRODUCTION

With the availability of ultra high frequency data, there has been growing interest in constructing daily variance using intraday high frequency data. Important contributions in this rapidly expanding literature include Andersen, Bollerslev, Diebold, and Labys (2001, ABDL hereafter), Andersen, Bollerslev, Diebold, Ebens (2001, ABDE hereafter), Barndorff-Nielsen and Shephard (2002), and the survey paper by Andersen, Bollerslev, Diebold (2005) and Bandi and Russell (2006). The main idea is to sum squared intraday returns over a day (the so-called realized volatility or empirical quadratic variation) as an estimate of the integrated daily variance (or theoretical quadratic variation). Compared with parametric approaches based on daily data using GARCH or stochastic volatility models, one major advantage of this approach is that it is model-free.

A common practice suggested in the earlier literature is to use five- or thirty-minute returns even though data may be available at much shorter intervals such as seconds. The longer sampling frequency is chosen in practice because of the tradeoff between the signal and the noise. In particular, it is well known that the observations of efficient price are contaminated by market microstructure noise. The effect of market microstructure noise on variance is well illustrated by volatility signature plots, which depict variance as a function of the sampling interval (see for example, Fang, 1996 and ABDE, 2001). Assuming that the market microstructure noise process is independent and identically distributed (i.i.d.) over time and also independent of the efficient price process, Zhang, Mykland and Ait-Sahalia (2005, ZMA hereafter) and Bandi and Russell (2005) showed that the realized volatility of observed prices goes to infinity as the sampling interval goes to zero. As a result, the noise dominates the signal and market microstructure noise swamps the variance of observed prices at the highest frequency. On the other hand, as the sampling interval increases, the signal/noise ratio goes up. But if the sampling interval is too big, a large amount of data is discarded, leading to inefficient estimation of realized

* The work was initiated when Liu visited Singapore in May 2004. Liu wishes to thank the School of Economics and Social Sciences at Singapore Management University and the Institute of Mathematical Science at National University of Singapore for hospitality during the visit. Huang and Yu gratefully acknowledge financial support from the Research Office at Singapore Management University. We wish to thank Peter Phillips and Anthony Tay for detailed comments on an earlier version of the paper. The time series of realized daily variance for the S&P 500 cash index from January 4, 1993 to December 31, 2004 is available from Jun Yu's web site <http://www.mysmu.edu.sg/faculty/yujun/default.html>.

variance. By sampling at five- or thirty-minute horizons, it is hoped that a compromise is attained whereby information loss is controlled and the observed price is a good approximation to the underlying (unobserved) efficient price. This approach, suggested in ABDL (2001), ABDE (2001) and Ebens (1999), has produced fruitful empirical results. For example, by calculating realized daily variance from five-minute returns, ABDE (2001) examined the statistical properties of realized daily variance for 30 Dow Jones stocks from January 1993 to May 1998, while Ebens (1999) examined the statistical properties of realized daily variance for the Dow Jones index over the same sample period. Both papers documented important empirical regularities concerning realized variance.

More recent work along this line of research approaches the problem by modeling microstructure noise explicitly and hence aims to understand how noise affects the realized variance estimate in a more systematic way. For example, Bandi and Russell (2005), Hansen and Lunde (2006), ZMA (2005), Barndorff-Nielsen, Hansen, Lunde and Shephard (2005) have adopted the assumption of pure noise (i.e., noise is i.i.d and independent with the efficient price). Based on the pure noise assumption, various estimation methods, including nonparametric kernel and sub-sampling techniques, have been developed for reducing the effects of the microstructure noise bias. For example, Bandi and Russell (2005) have found an optimal sampling frequency based on the mean square error criterion. ZMA (2005) suggested using all available data to compute realized volatility. More recently, Ait-Sahalia, Mykland and Zhang (2005a) relaxed the i.i.d. assumption by allowing for stationary temporal dependence in the noise process while maintaining the assumption of independence between the noise and the efficient price.

In this paper we compute the realized daily variance in the S&P 500 cash index using a method of ZMA and analyze the statistical properties of realized daily variance and the risk-return relationship. We make three contributions. First, we revisit the empirical regularities concerning realized daily variance obtained in the earlier literature. We document three new empirical results, namely nonstationarity in realized standard deviation, a significant and positive risk-return tradeoff, and monotonicity in the news impact function. Second, we compare the realized variance obtained from the intraday data with the variance estimated by parametric GARCH and stochastic volatility models. Such a comparison is important to gauge the informational gain in intraday data for the purpose of volatility estimation. Finally, as an important index of the US stock market, the S&P 500 is a dataset to which many parametric volatility models have been fitted (see, for example, Bollerslev et al, 1994 and Jacquier et al, 1994). However, studies on estimating and analyzing daily variance of S&P 500 using intraday data are much less extensive in the literature. Using a long

span of S&P 500 intraday data, we document the time series properties of realized daily variance.

The paper is organized as follows. Section 2 reviews the approach of ZMA. Section 3 describes the data. Section 4 compares the realized daily variance obtained from the intraday data with the estimated variance from parametric models. Section 5 examines the unconditional and the dynamic properties of realized variance while the risk-return relations are investigated in Section 6. Section 7 concludes.

2. ZMA ESTIMATORS WITH MICROSTRUCTURE NOISE

Suppose the efficient log-price in high frequency data is $p^*(t)$ and $(0 = t_0, \dots, t_n = 1)$ is a grid of discrete points in a unit interval (say, a day). The asymptotic theory is derived in the literature by requiring $\sup\{t_{i+1} - t_i\} \rightarrow 0$. In the case of equidistant sampling, we have $t_{i+1} - t_i = 1/n \equiv h, \forall i$ and hence $\sup\{t_{i+1} - t_i\} \rightarrow 0$ is equivalent to $n \rightarrow +\infty$ or $h \rightarrow 0$.

Typically $p^*(t)$ is assumed to follow a continuous time stochastic volatility model,

$$dX(t) = \mu(t)dt + \sigma(t)dB(t), \quad (1)$$

where $B(t)$ is the standard Brownian motion, $\mu(t)$ is a stochastic process which is predictable and has a locally bounded sample path, and $\sigma(t)$ is another stochastic process with a càdlàg sample path. This assumption contains many important models as special cases. The quantity of interest is $\int_0^1 \sigma^2(s)ds$ which is termed the integrated variance (IV) in the option pricing literature.

If $p^*(t)$ is observed, the theory of quadratic variation implies that, under (1)

$$\sum_{i=1}^n (p^*(t_{i+1}) - p^*(t_i))^2 \xrightarrow{a.s.} \int_0^1 \sigma^2(s)ds. \quad (2)$$

That is, the realized variance (RV) based on an empirical grid of observations where the maximum grid size tends to zero will produce a strongly consistent estimator of IV. Due to the importance of IV in many financial decisions and the availability of ultra-high frequency data, it is not surprising that this nonparametric approach to estimating the IV has recently received a great deal of attention in the literature. For example, ABDL (2001), ABDE (2001) and Ebens (1999) used it to document the properties of daily exchange rate variance, stock variance and stock index variance, respectively. Jacod (1994) and Barndorff-Nielsen and Shephard (2002) have derived the limiting distribution of the RV.

In spite of the appealing theoretical foundations and the mild assumptions that the RV approach is based on, many researchers tend to caution

against using all the available data (i.e. all the transaction prices or quote prices) to compute the realized variance. This is because the presence of market microstructure noise such as non-synchronous trading, bid-ask spread and price discreteness, precludes a direct observation of the efficient price, $p^*(t)$. To have a nearly continuous record of price on the one hand, and to mitigate the microstructure problems on the other, ABDL (2001), ABDE (2001) and Ebens (1999) used sparsely sampled (such as 5-minute) returns to estimate the IV. More recently, researchers have started examining the impact of market microstructure noise on realized variance and the statistical properties of estimators which use more information from data. Examples include ZMA (2005) and Ait-Sahalia, Mykland and Zhang (2005a, b), Hansen and Lunde (2006), Barndorff-Nielsen, Hansen, Lunde and Shephard (2005), and Bandi and Russell (2005), to name a few.

In this paper we will use one of the approaches suggested by ZMA (2005) to compute realized daily variance. Let $p(t)$ be the logarithmic transactions price, observed at $0 = t_0, \dots, t_n = 1$. ZMA assume that the observed price and the efficient price $p^*(t)$ are related as

$$p(t_i) = p^*(t_i) + \epsilon(t_i) \tag{3}$$

They further make the following assumptions about the noise

$$\epsilon(t_i) \sim iid(0, \sigma_\epsilon^2) \text{ and } \epsilon(t_i) \perp p^*(t_i). \tag{4}$$

Define $RV^{full} = \sum_{i=1}^{n-1} (p(t_{i+1}) - p(t_i))^2$. As a direct consequence of assumptions (4), we have

$$E(RV^{full} | p^*(t) \text{ process}) = IV + 2n\sigma_\epsilon^2. \tag{5}$$

When $n \rightarrow \infty$, the realized variance is of order $O(n)$ and hence diverges instead of converging, invalidating the approach of using the RV calculated from tick-by-tick data to estimate IV. Correspondingly, the volatility signature plot should asymptote as $h \rightarrow 0$.

Define the full grid that contains all the transactions prices by $\mathcal{G} = \{t_0, \dots, t_n\}$. Partition \mathcal{G} into K mutually exclusive sub-grids, called $\mathcal{G}^{(k)}$ with $k = 1, \dots, K$, so that the k^{th} sub-grid $\mathcal{G}^{(k)}$ starts at t_{k-1} and then select every K^{th} sample point after that, until t_n . Typically $\mathcal{G}^{(k)}$ is much more sparse than \mathcal{G} . For example, the average frequency in $\mathcal{G}^{(k)}$ is 5 minutes while the average frequency in \mathcal{G} is 15 seconds. Denote the average size of the sub-samples by \bar{n} .

The most frequently used estimator in the literature is the one based on a subgrid at a modest frequency such as 5-minute, say $\mathcal{G}^{(1)}$ (ABDE, 2001). Another estimate is based on a subgrid at the frequency which minimizes the mean square error of RV (see also Bandi and Russell, 2005).

By sampling sparsely, both these two estimators discard some datapoint in the original sample. One estimator which uses all available information calculates the RV based on all K subgrids (call them $RV^{\mathcal{G}^{(k)}}$ with $k = 1, \dots, K$) and then computes the average of these K RV estimates, that is

$$RV^{(K)} = \frac{1}{K} \sum_{k=1}^K RV^{\mathcal{G}^{(k)}} \quad (6)$$

The bias in $RV^{(K)}$ is of order $O(\bar{n})$ which is smaller than that based on the full grid (ie $O(n)$), as shown in ZMA. To further reduce the bias, ZMA employ the bias property in (5) and propose an estimator obtained by constructing a linear combination of the estimates at two different time scales, defined as

$$RV^{(K)} - \frac{\bar{n}}{n} RV^{full} \quad (7)$$

This two-time-scale method is constructed in the same spirit as the jackknife method of Quenouille (1956). The jackknife estimator is basically a weighted average of the full sample estimator and the sub-sample estimators, just as (7) above. Phillips and Yu (2005) provide another useful application of the jackknife method in the context of asset pricing. The ability of the jackknife estimator to reduce bias can be explained intuitively. Under assumptions (4), the bias in RV^{full} is of order $O(n)$ while the bias in $RV^{(K)}$ is of order $O(\bar{n})$. Therefore, the biases are canceled out in the weighted average (7). It should be emphasized that assumptions in (4) are critical for obtaining the precise weight in equation (7). When (4) is violated, the estimator defined in equation (7) may perform worse than the other estimators.

Although the assumptions about $\epsilon(t_i)$ in (4) substantially simplify econometric treatment, including identification, estimation, and implementation, (7) may not be a realistic representation of microstructure noise. Ait-Sahalia, Mykland and Zhang (2005a) provide a useful generalization of the two-time-scale estimator to the case with dependent noise, but the possible dependence between noise and the efficient price imposes further challenge to identification and estimation of IV. These issues are discussed in Hansen and Lunde (2006) and Phillips and Yu (2006).

3. DATA AND STYLIZED FACTS

In our empirical analysis, we use high frequency data on the S&P 500 cash index for the period January 4, 1993 to December 31, 2004, obtained from the Chicago Mercantile Exchange (CME). Within each day, we consider the transaction record between 8:30am to 15:00pm central standard time

TABLE 1.

Number of trading days in each year from 1993 to 2004

Year	1993	1994	1995	1996	1997	1998
# of days	253	252	251	254	253	252
Year	1999	2000	2001	2002	2003	2004
# of days	252	252	248	251	250	252

(CST). In total, there are 3020 trading days in the sampling period.¹ Table 1 shows the number of trading days in each year. Being an important index of US stock market behavior, the S&P 500 has been modeled with a large number of parametric volatility specifications. On the other hand, few studies have focused on the estimation and analysis of the variance of S&P 500 using intraday data.

The S&P 500 is a value-weighted index of 500 prominent common stocks listed on the New York Stock Exchange (NYSE), American Stock Exchange (AMEX), and the National Association of Security Dealers Automated Quotation system (NASDAQ). Whenever there is a new transaction recorded in any one of the 500 stocks, the S&P 500 cash index is updated. Since the S&P 500 stocks are the most actively traded equities, the S&P 500 cash index should be highly liquid. Naturally the aggregation would mitigate some microstructure effects such as the bid ask spread and the price discreteness. However, the effect of non-synchronous trading is retained in the S&P 500 index. While in theory one should expect one update for the S&P 500 at each time stamp, the data provided by CME is much sparser. Indeed, the first, second and third quartiles of inter-price durations are all 15 seconds in the S&P 500 cash index. This observation seems to suggest that CME only update the cash index every 15 seconds. The total number of transaction prices is around 1800 in a typical trading day.

Since we deal with a large amount of empirical data with several million time series observations, there are some obvious data errors on the price and mis-recording of time stamps. We therefore first removed these data entry errors in a cleaning operation. We then remove price bouncebacks, defined by Ait-Sahalia, Mykland and Zhang (2005a) as a price jump of size greater than a cutoff of 1%, immediately followed by a jump of similar magnitude but opposite sign.

To decide on which method to use for estimating realized daily variance, we first inspect the volatility signature plots for the “cleaned” data. To do so, we construct RVs based on 15, 18, \dots , 1800 seconds using the method

¹June 27, 1995 S&P 500 cash index data is missing in the CME data files. In addition, April 1, 1994 and September 11, 2001 are removed from the sample due to insufficient transactions prices recorded for the day.

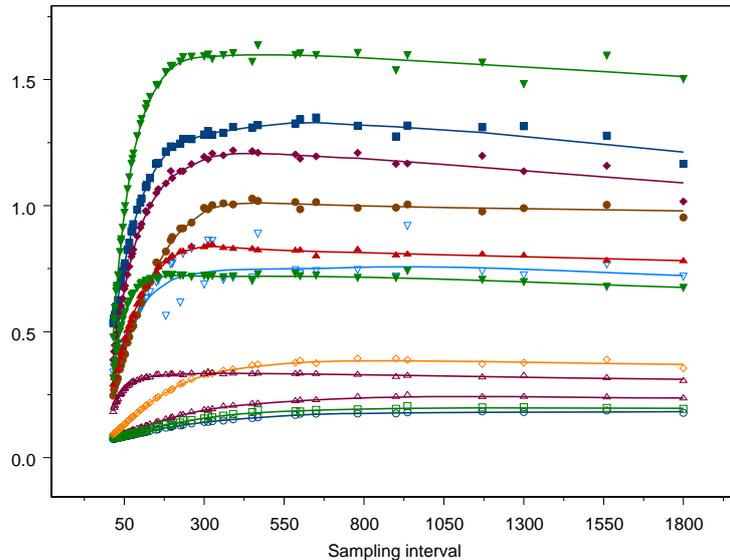


FIG. 1. Volatility signature plots. The vertical axis is the RV estimator ($\times 10000$), averaged over all trading days in each year. The horizontal axis is the sampling interval, taking the values of 15, 16, 18, \dots , 1000, 1800 seconds. Each curve corresponds to each year in the sample.

of Andersen et al (2001). In Figure 1 we plot the RV estimator, averaging over all trading days within each year, as a function of the sampling interval. When the price at particular time stamp is not available, we use the previous tick method to approximate it (see Hansen and Lunde (2006) for further details about the previous tick method). In contrast to signature plots that are typically found in the literature, such as Aït-Sahalia, Mykland and Zhang (2005a) and Bandi and Russell (2005), none of the signature plots for S&P 500 blows up as the sampling interval gets smaller. Indeed all the signature plots slope downward at the highest frequencies. According to Hansen and Lunde (2006), this is evidence of negative correlation between the microstructure noise and the efficient price, violating the two assumptions in (4). As a result, the two-time-scale RV estimator of ZMA may not be the best estimator.

However, all the signature plots seem to stabilize around $h = 600$ seconds. This observation suggests that one way to estimate the IV is to construct RV based on a subgrid of 10-minute returns. By doing so, of course, only 39 out of about 1800 observations are used in each day and hence a large amount of data is discarded. In this paper, we use $RV^{(K)}$, defined by Equation (7), to estimate the IV. In particular, we first apply the logarithmic transformation to the transactions prices, which defines

$p(t_i)$. Equation (6) with $K = 600$ is then used to obtain the realized daily variance over the sample period. The time series sequence of realized daily variance is defined by $\{RV_t\}_{t=1}^{3020}$. Due to the averaging effect, this estimator has much smaller variance than one based on a single grid (ZMA, 2005).

It would be interesting to know if there is any difference between the nonparametric estimate of variance and parametric counterparts, and if the properties of realized daily variance found in ABDE (2001) and Ebens (1999) remain qualitatively unchanged. It is worthwhile, therefore, to first review the following stylized facts about realized daily variance that have been documented in the literature:

1. Volatilities are time varying and clustering.
2. Although the return distribution is non-Gaussian and leptokurtic, the standardized return (the ratio of return to realized standard deviation) distribution conforms well with the normal distribution.
3. While neither realized daily variance nor realized daily standard deviation follow the normal distribution, the distribution of realized daily log-variance is closer to the normal distribution.
4. Long range dependence and covariance stationarity are found in realized daily variance, featured by an estimate of the memory parameter significantly larger than 0 but generally less than 0.5.
5. Standard unit root tests often reject the presence of a unit root in realized daily variance.
6. The news impact function (NIF) defined by Engle and Ng (1993) is asymmetrically and V-shaped.
7. The evidence about the risk-return tradeoff is mixed and often statistically insignificant.

4. COMPARING PARAMETRIC AND NONPARAMETRIC VARIANCE ESTIMATES

Two extensive literatures, seeking to estimate daily variance, have developed since the introduction of the ARCH model by Engle (1982) and the introduction of the stochastic volatility (SV) model by Taylor (1982). See Bollerslev et al (1994) and Shephard (2005) for the review of the GARCH literature and the SV literature, respectively. Andersen et al (2005) provide an interesting discussion on the relationship between the conditional variance of GARCH and SV models and the RV implied by intraday data. In this section, based on the same dataset, we examine the relationship of the three sequences of variance estimates, obtained, respectively, from the intraday data, the GARCH(1,1) model fitted to the open-to-close returns, and the basic SV model fitted to the open-to-close returns.

The GARCH(1,1) model of Bollerslev (1982) takes the form of

$$y_t = \mu_y + \sigma_t \epsilon_t = \mu_y + \exp(h_t/2) \epsilon_t, \quad \epsilon_t \sim \text{i.i.d. } N(0, 1) \quad (8)$$

and

$$\sigma_{t+1}^2 = \alpha_0 + \alpha_1 \sigma_t^2 + \beta y_t^2, \quad (9)$$

where y_t represents the open-to-close return on day t . The model is estimated by the maximum likelihood (ML) method. The normality assumption of ϵ_t can be relaxed, in which case ML estimation becomes quasi maximum likelihood. From the ML estimates of α_0 , α_1 and β , the variance can be easily estimated by

$$\hat{\sigma}_{t+1}^2 = \hat{\alpha}_0 + \hat{\alpha}_1 \sigma_t^2 + \hat{\beta} y_t^2,$$

with appropriate initializations on σ_0^2 and y_0^2 .

The basic SV model of Taylor (1982) takes the form of

$$y_t = \mu_y + \sigma_t \epsilon_t = \mu_y + \exp(h_t/2) \epsilon_t, \quad \epsilon_t \sim \text{i.i.d. } N(0, 1) \quad (10)$$

and

$$h_{t+1} = \alpha + \phi h_t + \sigma \eta_t, \quad \eta_t \sim \text{i.i.d. } N(0, 1). \quad (11)$$

where y_t represents the open-to-close return on day t . The model was estimated by various techniques in Mahieu and Schotman (1998). In the present paper, we estimate the model using a Bayesian Markov chain Monte Carlo (MCMC) method of Meyer and Yu (2000).

To estimate the variance from an SV model, two quantities have been used in the literature, namely, the smoothed variance ($E(\sigma_{t+1}^2 | I_T)$) and the filtered variance ($E(\sigma_t^2 | I_t)$), where $I_t = \sigma(y_1, \dots, y_t)$. The smoothed estimate of variance is a by-product of the MCMC algorithm as σ_{t+1}^2 is in the augmented parameter space (Jacquier et al, 1994) and hence readily available once the MCMC output is obtained (Tsay, 2005). The filtered estimate is generally more difficult to calculate in the context and may be obtained by means of the particle filter (Pitt and Shephard, 1999 and Kitagawa, 1996). Berg et al (2004) discuss how to use Kitagawa's algorithm to obtain the filtered variance of the SV model given by (10) and (11). While the Kalman filter works for linear Gaussian state-space models, particle filter is applicable to nonlinear non-Gaussian state models. The basic idea of particle filter is to draw "particles" from the filtered density ($F(\sigma_t^2 | I_t)$), advance the particles by drawing from $F(\sigma_{t+1}^2 | \sigma_t^2, I_t)$, and then do resampling to take the new information implied by y_{t+1} into account. Although the smoothed variance is easy to compute in a SV model, the filtered variance is a more reasonable quantity to use in the context because it has the

same information set as the GARCH model. Hence it is used in the present paper and the number of particles is chosen to be 10,000.

Figure 2 shows the time series plots for the three sequences of variance estimates. Several results are evident from Fig. 2. First, all the three sequences have a very similar pattern. For example, both at the beginning of the sample and at the end of the sample, variance is low while in the middle variance is high. Second, while October 28, 1997 (corresponding to the Asian financial crisis) has the highest realized daily variance in the sample, the variance estimated from the two parametric models suggests that it reached a peak in July 2002. Thirdly, both the GARCH and the SV models produce much smoother estimates of the variance than realized variance. Underestimation of variance is especially serious during the volatile period. These results remain qualitatively unchanged when more flexible GARCH and SV models are used.

To quantify the underestimation problem in the parametric volatility models, we fit the following empirical regression

$$RV_t - \hat{\sigma}_t^{2,parametric} = \hat{\beta}_0 + \hat{\beta}_1 RV_t + \hat{e}_t,$$

where RV_t is the RV at day t and $\hat{\sigma}_t^{2,parametric}$ is the estimate of variance from the GARCH(1,1) model or the basic SV model at day t . The OLS estimates, the associated t statistics based on Newey-West standard errors, and R^2 are reported in Table 2. Figures 3-4 depict the scatter plots of $RV_t - \hat{\sigma}_t^{2,parametric}$ against RV_t . Superimposed is a nonparametric curve obtained by locally weighted least square regression of Cleveland and Devlin (1988). The underestimation is evident in both cases and appears to be statistically significant. Relative to the GARCH(1,1) model, however, the variance estimates implied by the basic SV model are closer to the RVs, featured by a much smaller value for R^2 and a smaller value for the t statistic for β_1 in the regression model.

5. UNCONDITIONAL AND DYNAMIC PROPERTIES

5.1. Unconditional Properties

Table 3 reports some basic summary statistics for RV. The feature of variance clustering is manifest in the first autocorrelation function (ACF) (see also Panel 3 of Figure 2), confirming the first stylized fact.

Figures 5-6 plot the unconditional distribution of the open-to-close returns (y_t), and the unconditional distribution of the standardized returns (defined as y_t/RV_t). We also report some summary statistics of the two series, including the skewness, the kurtosis, and the Jarque-Bera test statistic and the associated p -value. It is clear that the return distribution has fatter tails than the normal distribution, featured by the high kurtosis (6.87).

TABLE 2.

Variance estimation by parametric volatility models

GARCH			
Parameter	β_0	β_1	R^2
Estimate	-7.15E-05	0.502712	0.331454
t-statistic	-14.02493	7.841215	
SV			
Parameter	β_0	β_1	R^2
Estimate	-8.86E-05	0.12768	0.128089
t-statistic	-13.97981	3.940623	

Note: This table reports the OLS estimates, associated t statistics based on the Newey-West standard errors, and R^2 in the following empirical model,

$$RV_t - \hat{\sigma}_t^{2,parametric} = \hat{\beta}_0 + \hat{\beta}_1 RV_t + \hat{e}_t,$$

where RV_t is the RV at day t and $\hat{\sigma}_t^{2,parametric}$ is the estimate of variance from the GARCH(1,1) model or the basic SV model at day t .

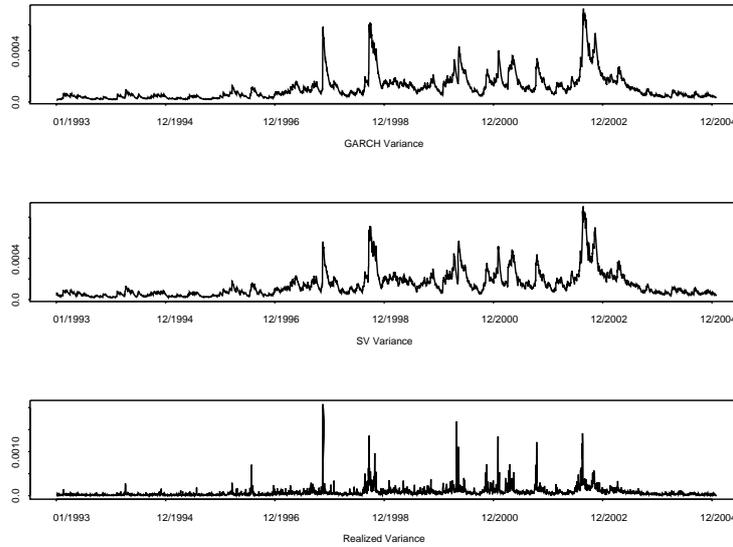


FIG. 2. Time series plots of the conditional variance estimate for the GARCH(1,1) model, the filtered variance of the basic SV model, and the realized daily variance from January 4, 1993 to December 31, 2004.

TABLE 3.

Summary statistics of realized variance

	Mean	St. dev.	Skewness	Kurtosis	ACF1	ACF100	ACF200
<i>RV</i>	7.74e-6	1.15e-4	6.80	77.60	0.594	0.100	0.113

Note: This table reports summary statistics for the realized daily variance of S&P 500 cash index. ACF1, ACF100, ACF200 represent the ACF of orders 1, 100, 200, respectively.

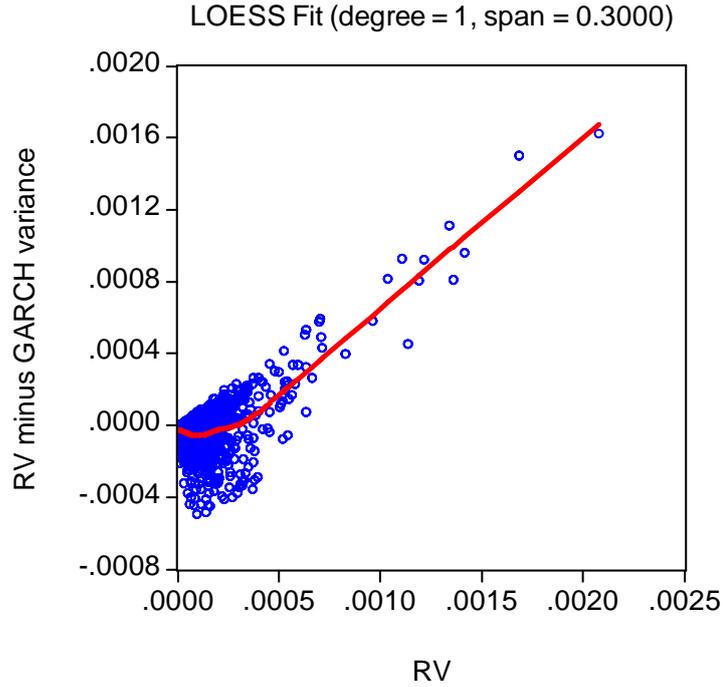


FIG. 3. Empirical relationship between realized variance and the variance estimate from GARCH(1,1). The solid curve is a nonparametric regression smoother.

Moreover, it is skewed to the left (-0.151). The Jarque-Bera test is 1895.85 and has a p -value of 0, rejecting the null hypothesis of normality. However, the unconditional distributions of the standardized return series conform better with the normal distribution. While the Jarque-Bera test still rejects normality, the statistic reduces to 34.33 from 1895.85. Hence the second stylized fact is partially confirmed. Interestingly, the skewness becomes positive in the standardized return.

Figure 7 plots the unconditional distributions of the realized variance, realized standard deviation, and logarithmic realized variance. For com-

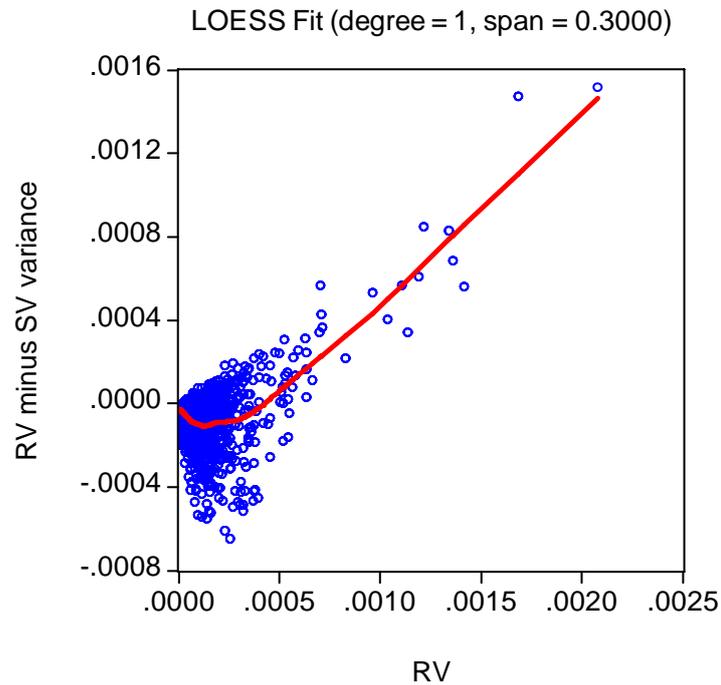


FIG. 4. Empirical relationship between realized variance and the variance estimate from the basic SV model. The solid curve is a nonparametric regression smoother.

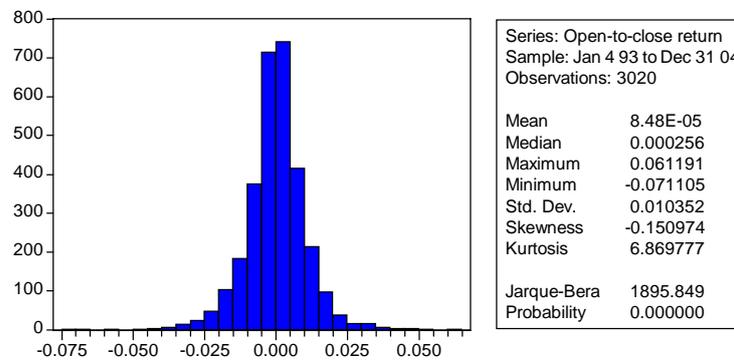


FIG. 5. Density estimate of unconditional distribution and summary statistics of daily returns. The figure shows the unconditional distribution of the daily open-to-close returns. The numbers are the summary statistics of the same data. The sample period is from January 4, 1993 to December 31, 2004.

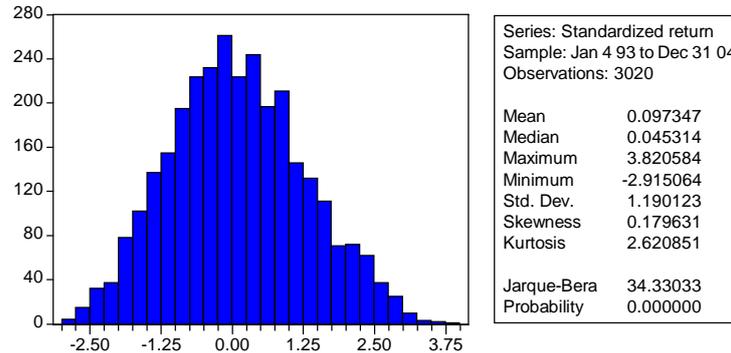


FIG. 6. Density estimate of unconditional distribution and summary statistics of standardized returns, measured by r_t/RV_t . The figure shows the unconditional distribution of the daily open-to-close returns. The numbers are the summary statistics of the same data. The sample period is from January 4, 1993 to December 31, 2004.

parison, we also depict the normal density function with the same mean and the same variance. Neither the variance nor the standard deviation is normally distributed. Compared with the variance and standard deviation, however, the log-variance has a better normal approximation. In Table 4 we report the summary statistics for the standard deviation and log-variance. Comparing Table 4 with Table 3, the better normal approximation is manifest in the smaller value of the skewness, kurtosis, and Jarque-Bera statistic in the log-variances although normality is still rejected by the Jarque-Bera test. Consequently, we confirm the third stylized fact.

TABLE 4.

Summary statistics of realized standard deviation and logarithmic realized variance

	Mean	St. dev.	Skewness	Kurtosis	ACF1	ACF100	ACF200
<i>SD</i>	7.69e-3	4.29e-3	2.147	12.00	0.740	0.297	0.278
<i>log - RV</i>	-10.00	1.012	0.103	2.917	0.772	0.463	0.405

Note: This table reports summary statistics for the realized standard deviation and logarithmic realized variance of the S&P 500 cash index. ACF1, ACF100, ACF200 represent the ACF of orders 1, 100, 200, respectively.

5.2. Dynamic Properties

Figure 8 plots the autocorrelation function (ACF) for the variance, standard deviation and log variance. In all cases, the ACF decays very slowly, suggesting evidence of long range dependence.

Motivated by these plots, we fit the following fractionally integrated model to each series,

$$(1 - L)^d X_t = \mu + \epsilon_t,$$

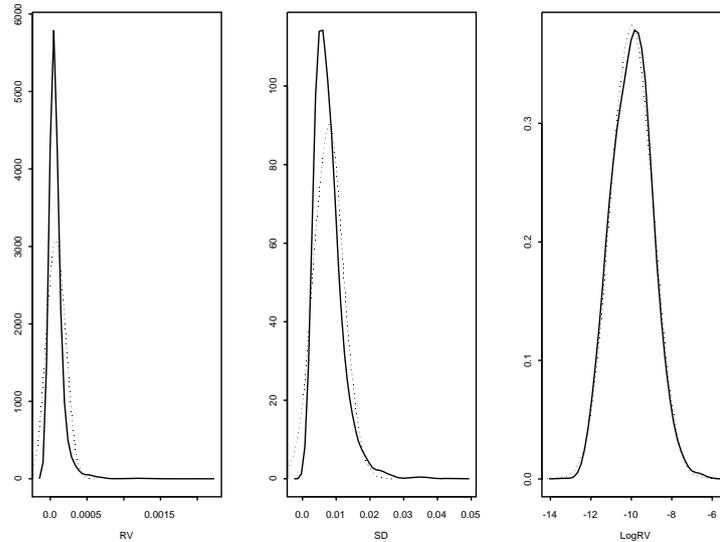


FIG. 7. Density estimate of unconditional distribution of the realized variance (left), realized standard deviation (middle), and logarithmic realized variance (right). The dotted line refers to the normal distribution with the same mean and the same standard error. The sample period is from January 4, 1993 to December 31, 2004.

where X_t is one of the three series, and $\epsilon_t \sim I(0)$. See Baillie (1996) for the review of fractionally integrated processes. As we do not assume a particular range for the memory parameter d , we make use of a feasible exact local whittle method developed recently by Shimotsu and Phillips (2005) and Shimotsu (2006) to estimate d . This method allows d to take a much wider range of possible values than many alternative methods and produces valid confidence intervals and asymptotic standard errors for both stationary and nonstationary values of d . These advantages are practically important because a priori one normally does not know the range of possible values of d . In Table 5 we report these estimates of d together with the corresponding asymptotic standard errors in each of the three cases. In all cases, we find strong evidence of fractional integration. The point estimates of d range between 0.43 and 0.63, all significantly larger than zero. This confirms the stylized fact concerning the long range dependence in realized variance. However, in no case can we reject the null hypothesis of $d = 0.5$ in favor of $d < 0.5$. In two cases the point estimate of d is even larger than 0.5. In one case we reject the null hypothesis of $d = 0.5$ in favor of $d > 0.5$. As a result, we find evidence of non-stationarity in all three realized quantities, especially in the standard deviation and the log variance. This result is

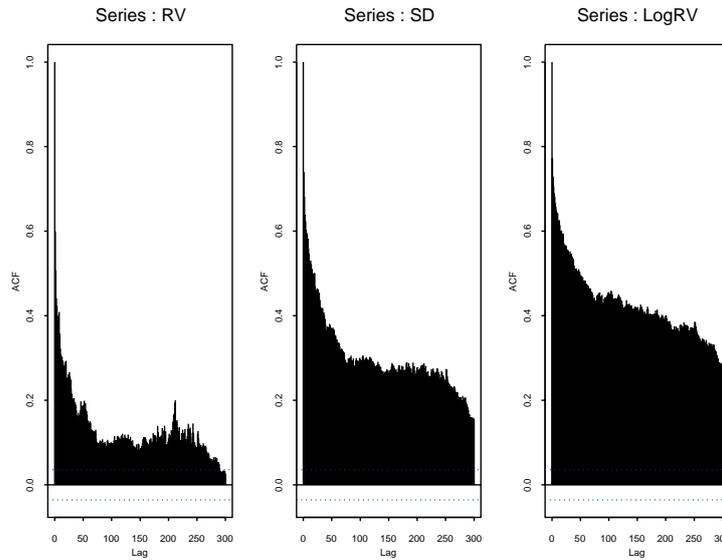


FIG. 8. The autocorrelation functions (ACF). The graphs shown the first 200 ACF of the variance (left), standard deviation (middle), and log variance (right). The dotted line refers to the upper limit of the 95% confidence interval. The sample period is from January 4, 1993 to December 31, 2004.

in sharp contrast to those obtained in ABDE and Ebens (1999) where estimates of d are all significantly less than 0.5 and suggest stationarity.

TABLE 5.

Estimates of the memory parameter for the realized variance, realized standard deviation and logarithmic realized variance

	<i>RV</i>	<i>SD</i>	<i>log - RV</i>
d	0.432	0.574	0.625
asy. std. err.	0.045	0.045	0.045

Note: This table reports the feasible exact local whittle estimate of d and the corresponding asymptotic standard error for the realized variance, realized standard deviation and log realized variance of the S&P 500 cash index. The model fitted is $(1 - L)^d X_t = \mu + \epsilon_t$, where $\epsilon_t \sim I(0)$.

Furthermore, since the estimates of d are always significantly less than 1, we find strong evidence against the unit root (or $I(1)$) hypothesis. Alternatively, one can test for the presence of a unit root in each of the three series using a unit root test. Table 6 contains the results of ADF unit root

test, again suggesting evidence against the I(1) hypothesis. Therefore, we confirm the fifth stylized fact.

TABLE 6.

Unit root test statistics for the realized variance, realized standard deviation and logarithmic realized variance

	<i>RV</i>	<i>SD</i>	<i>log - RV</i>
ADF	-8.82	-6.77	-5.08
<i>p</i> -value	0.0	0.0	0.0

Note: This table reports the augmented Dickey-Fuller test and its *p*-value for the realized variance, realized standard deviation and log realized variance of the S&P 500 cash index.

6. RISK-RETURN RELATIONS

In this section we are concerned with two kinds of relationships between variance and returns, namely, the intertemporal relation between returns and variance (i.e., inter-temporal capital asset pricing model or ICAPM), and the news impact relation between variance and lagged returns. Uncovering these two relations has been a topic that has received extensive investigation in the empirical literature.

Regarding the ICAPM intertemporal relation between returns and variance, while Merton (1973) demonstrated a positive risk-return relation from a theoretical perspective, the empirical literature has been unable to document strong statistical evidence to support it. (However, see Ghysels et al (2005) for a counterexample). Following much of the literature, we estimate the following linear regression model,

$$y_t = \alpha + \beta E_{t-1}(RV_t) + \epsilon_t,$$

where y_t represents the open-to-close return on day t and $E_{t-1}(RV_t)$ is measured by the lagged realized daily variance RV_{t-1} . β is the so-called risk aversion parameter.

Figure 9 displays the scatter plot for y_t against RV_{t-1} . Superimposed is the estimated linear regression line which clearly slopes upwards. Table 7 reports the OLS estimates, standard errors, and t-statistics of α and β . The estimated risk aversion parameter is 4.209 which is statistically significant at the 5% level.

Regarding the relation between variance and lagged returns, which is also termed the news impact function (NIF) in the volatility literature, two hypotheses coexist in the literature, namely, the leverage effect and the volatility feedback effect. Although both effects predict a negative

TABLE 7.

Regression of daily return on lagged daily realized variance

	α	β
estimate	-0.00024	4.209
std error	0.00023	1.635
t-statistic	-1.06	2.57

Note: This table reports the OLS regression results from the following model

$$y_t = \alpha + \beta RV_{t-1} + \epsilon_t,$$

where y_t represents the open-to-close return on day t .

correlation between the current variance and the lagged return, the leverage hypothesis predicts a monotonically decreasing and hence asymmetric NIF while the other hypothesis predicts a more flexible asymmetry in NIF. Most studies in the GARCH literature have documented an asymmetrically V-shaped NIF.

Figure 10 displays the scatter plot for logarithmic realized variance against $y_{t-1}/\sqrt{RV_{t-1}}$. Superimposed is a nonparametric regression curve (namely NIF) obtained by a Gaussian kernel and the optimal bandwidth. Comparing our Fig. 10 with Fig. 10 in ABDE (2001), we draw two conclusions. First, there is an important difference between the two NIFs. In ABDE, the NIF is asymmetrically V-shaped when two regression lines are fitted. In our Fig. 10, however, we find new empirical evidence in the high frequency S&P 500 cash index that NIF monotonically decreases. An implication of asymmetrically V-shaped NIF is that variance tends to rise when news (either good or bad) arrives, consistent with the implication of the asymmetric ARCH models. An implication of a monotonic NIF is that variance tends to increase (decrease) when bad (good) news arrives, consistent with the leverage hypothesis of Black (1976), as shown in Yu (2005). Second, there is an important similarity between our Fig. 10 and Fig. 10 in ABDE, that is, the NIFs are all very flat and rather poorly determined.

7. CONCLUSIONS

Following ABDE (2001) and Ebens (1999), this paper examines the properties of realized daily volatilities and return-variance relationships, differing from existing work in two aspects. First, unlike ABDE (2001) which focused on individual stocks and Ebens (1999) which focused on the DJIA index, we use a new dataset, the S&P 500 cash index, provided by CME. Noting that the TAQ database only includes intraday observations on in-

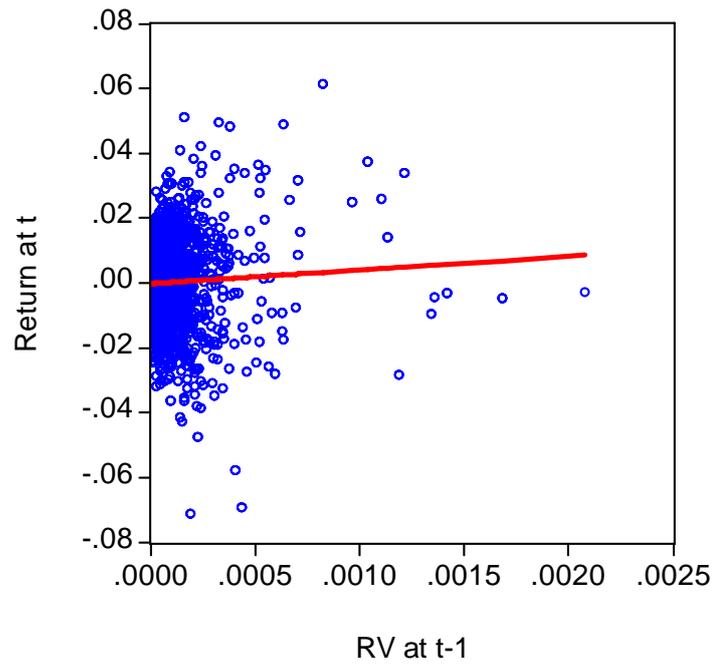


FIG. 9. Risk aversion. The figure shows the scatter plot of the current return against the lagged realized daily variance for S&P500. The solid curve is the estimated regression line.

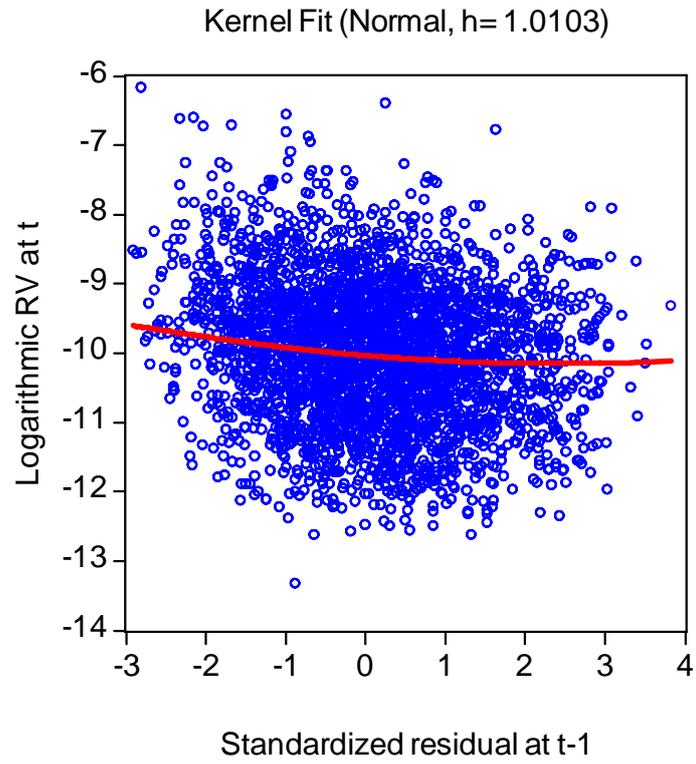


FIG. 10. News impact function. The figure shows the scatter plot of the realized daily variance against the lagged standardized return shock for S&P500. The solid curve is a kernel regression smoother.

dividual stocks, one advantage with the S&P 500 cash index is that all the intraday observations are available directly from this dataset. However, CME only updates the price every 15 seconds or so, although transactions in the underlying stocks are much more active. Second, instead of using artificially constructed 5-minute returns which discard a lot of observations in the high frequency data, we make use of all the available observations to estimate the realized daily variance. In particular, we employ the method proposed by Zhang, Mykland and Aït-Sahalia (2005) to deal with the microstructure noise problem. Based on the new estimation methodology, we re-examined the stylized facts about realized daily variance reported in ABDE (2001) and Ebens (1999). We found that most of the stylized facts documented in the literature continued to hold, with the exception that the standard deviation and the log variance are no longer stationary. We also compared the realized variance with the variance estimates from the GARCH(1,1) model and the basic SV models. We found that all the three variance series have the same pattern, but the variance estimates implied by the parametric volatility models based on the daily data are too smooth. Moreover, we found evidence of superiority of the SV model relative to the GARCH model in terms of approximating the RV. Finally, based on the realized variance, we examined two relations about returns and volatility. While the literature has found difficult to document a significant ICAPM relation, we find a positive and statistically significant relation. We also found evidence of a monotonically decreasing NIF.

REFERENCES

- Aït-Sahalia, Y., Mykland, P., and L. Zhang, 2005a, Ultra high-frequency volatility estimation with dependent microstructure noise. Working Paper, Princeton University.
- Aït-Sahalia, Y., Mykland, P., and L. Zhang, 2005b, How often to sample a continuous-time process in the presence of microstructure noise. *Review of Financial Studies* **18**, 351-416.
- Andersen, T.G., Bollerslev, T., and F.X. Diebold, 2005, Parametric and nonparametric volatility measurement: In: L.P. Hansen and Y. Aït-Sahalia (eds.), *Handbook of Financial Econometrics*. Amsterdam: North-Holland, forthcoming.
- Andersen, T., T. Bollerslev, F.X. Diebold, and H. Ebens, 2001, The distribution of realized stock return volatility. *Journal of Financial Economics* **61**, 43–76.
- Andersen, T., T. Bollerslev, F.X. Diebold, and P. Labys, 2001, The distribution of realized exchange rate volatility. *Journal of the American Statistical Association* **96**, 42–55.
- Andersen, T., T. Bollerslev, F.X. Diebold, and P. Labys, 2000, Great realization. *Risk* **13**, 105–108.
- Baillie, R., 1996, Long memory processes and fractional integration in econometrics. *Journal of Econometrics* **73**, 5–59.
- Bandi, F. M. and J. Russell, 2005, Microstructure noise, realized volatility, and optimal sampling. Working paper, Graduate School of Business, University of Chicago.

- Bandi, F. M. and J. Russell, 2006, Volatility, *Handbook of Financial Engineering*, forthcoming.
- Barndorff-Nielsen, O., P. Hansen, A. Lunde, and N. Shephard, 2005, Regular and modified kernel-based estimators of integrated variance: The case with independent noise, Working paper, Nuffield College.
- Barndorff-Nielsen, O. and N. Shephard, 2000, Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B* **64**, 253-280.
- Berg, A., R. Meyer, and J. Yu, 2004, Deviance information criterion for comparing stochastic volatility models. *Journal of Business and Economic Statistics* **22**, 107-120.
- Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* **31**, 307-327.
- Bollerslev, T., R. Engle, and D. Nelson, 1994, ARCH models. *Handbook of Econometrics* **4**, 2959-3038.
- Cleveland, W.S. and S.J. Devlin, 1988, Locally-weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association* **83**, 596-610.
- Ebens, H., 1999, Realized stock volatility. Working Paper, Johns Hopkins University.
- Engle, R. and V. Ng, 1993, Measuring and testing the impact of news in volatility. *Journal of Finance* **43**, 1749-1778.
- Fang, Y., 1996, Volatility modeling and estimation of high-frequency data with Gaussian noise. Unpublished PhD thesis, MIT, Sloan School of Management.
- Ghysels, E., Santa-Clara, P., and R. Valkanov, 2005, There is a risk-return tradeoff after all. *Journal of Financial Economics* **76**, 509-548.
- Hansen, P. and A. Lunde, 2006, Realized variance and market microstructure noise (with discussions). *Journal of Business and Economic Statistics* **24**, 127-161.
- Jacod, J., 1994, Limit of random measures associated with the increments of a Brownian motion martingale. Working Paper, University Pierre et Marie Curie.
- Jacquier, E., N.G. Polson, and P.E. Rossi, 1994, Bayesian analysis of stochastic volatility models. *Journal of Business and Economic Statistics* **12**, 371-389.
- Kitagawa, G., 1996, Monte Carlo filter and smoother for Gaussian nonlinear state space models. *Journal of Computational and Graphical Statistics* **5**, 1-25.
- Mahieu, R.J. and P.C. Schotman, 1998, An empirical application of stochastic volatility models. *Journal of Applied Econometrics* **13**, 330-360.
- Merton, R.C., 1973, An intertemporal capital asset pricing model. *Econometrica* **41**, 867-887.
- Meyer, R. and J. Yu, 2000, BUGS for a Bayesian analysis of stochastic volatility models. *Econometrics Journal* **3**, 198-215.
- Phillips, P. C. B. and J. Yu, 2005, Jackknifing bond option prices. *Review of Financial Studies* **18**, 707-742.
- Phillips, P. C. B. and J. Yu, 2006, Discussion on realized variance and market microstructure noise. *Journal of Business and Economic Statistics* **24**, 202-208.
- Pitt, M. and N. Shephard, 1999, Filtering via simulation: Auxiliary particle filter. *The Journal of the American Statistical Association* **94**, 590-599.
- Quenouille, M. H., 1956, Notes on bias in estimation. *Biometrika* **43**, 353-360.

- Shephard, N., 2005, *Stochastic Volatility: Selected Readings*. Oxford University Press.
- Shimotsu, K. and P.C.B. Phillips, 2005, Exact local whittle estimation of fractional integration. *Annals of Statistics* **33**, 1890-1933.
- Shimotsu, K., 2006, Exact local whittle estimation of fractional integration with unknown mean and time trend. Working paper, Queens University.
- Taylor, S.J, 1982, Financial returns modeled by the product of two stochastic processes — a study of the daily sugar prices 1961-75. In: Anderson, O.D., Editor, *Time Series Analysis: Theory and Practice*, 203–226. North-Holland, Amsterdam.
- Tsay, R., 2005, *Analysis of Financial Time Series*, Wiley.
- Yu, J., 2005, On leverage in a stochastic volatility model. *Journal of Econometrics* **128**, 165-178.
- Zhang, L., Mykland, P.A., and Y. Aït-Sahalia, 2005, A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association* **100**, 1394-1411.