Volatility-return dynamics across different timescales

Ramazan Gençay[*]   Faruk Selçuk[†]

June 2004

Abstract

This paper proposes a simple yet powerful methodology to analyze the relationship between stock market return and volatility at multiple time scales (horizons). The high frequency DJIA index is studied at intraday (within a trading day) and at interday (any horizon longer than a trading day) horizons. Based on the proposed decomposition, we are able to disentangle the relationship between volatility and return at different timescales. Particularly, we show that the leverage effect (negative correlation between current return and future volatility) is weak at high frequencies, and becomes prominent at lower frequencies (long horizons). On the other hand, the volatility feedback (negative correlation between current volatility and future returns) is a very short lived phenomenon. The positive correlation between the current volatility and future returns becomes dominant at timescales one day and higher, providing evidence that the risk and return are positively correlated.

Key Words: intraday volatility, leverage effect, volatility feedback, wavelet analysis.

JEL No: C2, G0, G1

1 Introduction

The relationship between a stock market index and its volatility has been studied extensively in the literature. A common finding is that innovations to a stock market index and innovations to volatility are negatively related, e.g., a decrease in stock price is associated with an increase in its volatility. Furthermore, the relationship is asymmetric: an absolute change in volatility after a negative shock to the return series is significantly higher than the absolute change in volatility after a positive shock with the same magnitude.[1]

The negative relationship between current returns and future volatilities is labeled as "leverage effect". Early researchers argued that the fall in stock price causes an increase in the debt-equity ratio (financial leverage) of the firm and the risk associated with the firm increases subsequently (Black, 1976; Christie, 1982). More recent research argues that "volatility feedback" is the main source of the negative relationship between return shocks and volatility (Campbell and Hentschel, 1992). According to this approach, an anticipated increase in the risk of a stock induces a high risk premium on the stock and the stock price falls immediately. In other words, if the *expected* stock return increases when its volatility increases, the stock price must fall on impact when volatility increases (Campbell *et al.*, 1997, p. 497). An implication of the volatility feedback hypothesis is that the risk-premium is not constant but time changing, and it may be another contributing factor to the very short-term negative relationship between current volatility and future return. In longer time horizons, this negative relationship between volatility and return must reverse and the current volatility and the future return must be positively correlated because of the required higher risk premium. It should be noted that in the leverage effect hypothesis, the stock return causes volatility while the volatility feedback hypothesis implies that the causality runs the other way around.

Empirical studies on the subject report mixed results. For example, French *et al.* (1987) and Campbell and Hentschel (1992) show that volatility and expected return are positively related which supports the volatility feedback hypothesis. On the other hand, Nelson (1991) and Glosten *et al.* (1993) document a negative relationship between expected stock returns and their conditional volatility. Regarding the asymmetric behavior of the stock market volatility, Campbell and Hentschel (1992) reports that both the volatility feedback effect and the leverage effect play an important role. Harrison and Zhang (1999) also examines the relation over different holding periods and uncovers a significantly positive risk and return relation at long holding intervals, such as one and two years, which does not exist at short holding periods such as one month. A negative relationship between shocks to a stock

---

[1]See Bekaert and Wu (2000) and Wu (2001) and references therein for a recent review of the literature.

market return and shocks to its volatility is not confined to developed markets. An early study by Bekaert and Harvey (1997) at monthly frequency in emerging markets found some evidence on asymmetric volatility, e.g., negative shocks increase volatility by more than positive shocks. Recently, Selçuk (2004) reports a significant negative correlation between shocks to the stock market index and shocks to volatility at daily frequency in ten emerging market economies.[2] Overall, as noted by Bollerslev and Zhou (2003), the current literature reports significantly different results on the relationship between volatility and return, depending on the different definitions of volatility, the length of the return horizon, instruments employed in regression estimations, and conditioning information used in estimating the relationship.

This paper proposes a simple, yet an envelope methodology to investigate the relationship between volatility and return at different time scales. The proposed method is based on a wavelet multi-scaling approach which decomposes the data into its low- and high-frequency components across time. The decomposition is simple to calculate, and does not depend on model-specific parameter choices. It is also translation invariant, and has the ability to decompose an arbitrary length series without boundary adjustments. In addition, the method is associated with a zero-phase filter and is circular. The zero-phase property ensures that there is no phase shift in filtered data while circularity helps to preserve the entire sample unlike other two-sided filters where data loss occurs from the beginning and the end of the studied sample.

We demonstrate the methodology utilizing the Dow Jones Industrial Average (DJIA), recorded at 5-minute intervals during the sample period of September 19, 1994 - October 16, 2002. Our findings indicate that the leverage effect (negative correlation between current return and future volatility) is very weak at high frequencies, and becomes prominent for horizons longer than two days. On the other hand, the negative correlation between current volatility and future returns is a very short lived phenomenon. The positive correlation between the current volatility and past returns becomes dominant at timescales one day and longer, providing evidence that the risk and return are positively correlated at longer horizons.

This paper is structured as follows. The following section introduces the wavelet decomposition and its implications in economics and finance. Section 3 studies the cross correlation between volatility and returns at different timescales for different leads and lags. We conclude afterwards.

---

[2]Bekaert and Harvey (2003) and other articles in a special issue of the Journal of Empirical Finance (2003, v10, 1-2) study several issues in emerging financial markets.

2　Wavelet decomposition

Wavelet methods are rather newer ways of analyzing time series and can be seen as a natural extension of the Fourier analysis. The formal subject matter, in terms of their formal mathematical and statistical foundations go back only to the 1980s. In recent years, there have been several unique applications of wavelet methods to financial problems. For instance, it is well documented that strong intraday seasonalities may induce distortions in the estimation of volatility models.[3] These seasonalities are also the dominant source for the underlying misspecifications of the various volatility models. Gençay *et al.* (2001a) propose a simple method for intraday seasonality extraction that is free of model selection parameters. Their methodology is based on a wavelet multi-scaling approach which decomposes the data into its low and high frequency components through a discrete wavelet transform. Gençay *et al.* (2001a) filtering method is translation invariant, has the ability to decompose an arbitrary length series without boundary adjustments, is associated with a zero-phase filter and is circular. Being circular helps to preserve the entire sample unlike other two-sided filters where data loss occurs from the beginning and the end of the studied sample. Gençay *et al.* (2001c) investigate the scaling properties of foreign exchange volatility by decomposing the variance of a time series and the covariance between two time series on a scale (time horizon) by scale basis through the application of a discrete wavelet transformation. It is shown that foreign exchange rate volatilities follow different scaling laws at different horizons. Particularly, there is a smaller degree of persistence in intra-day volatility as compared to volatility at one day and higher scales. Therefore, a common practice in the risk management industry to convert risk measures calculated at shorter horizons into longer horizons through a global scaling parameter may not be appropriate.

In Gençay *et al.* (2003) and Gençay *et al.* (2004b), a new approach to estimating the systematic risk (the beta of an asset) in a capital asset pricing model (CAPM) have been proposed. At each time scale (horizon), the variance of the market return and the covariance between the market return and a portfolio are calculated to obtain an estimate of the portfolio's beta. The empirical results show that the relationship between the return of a portfolio and its beta becomes stronger for longer time horizons. Therefore, the predictions of the CAPM model are more relevant at medium-long run as compared to short time horizons.

In a recent paper, Gençay *et al.* (2004a) argue that conventional time series analysis, focusing exclusively on a time series at a given scale, lacks the ability to explain the nature of the data generating process. A process equation that successfully explains daily price changes, for example, is unable to characterize the

---

[3]Extensive literature on this subject can be found in Dacorogna *et al.* (2001).

nature of hourly price changes. On the other hand, statistical properties of monthly price changes are often not fully covered by a model based on daily price changes. Gençay *et al.* (2004a) simultaneously model regimes of volatilities at multiple time scales through wavelet-domain hidden Markov models. They establish an important stylized property of volatility across different time scales and call this property *asymmetric vertical dependence.* It is asymmetric in the sense that a low volatility state (regime) at a long time horizon is most likely followed by low volatility states at shorter time horizons. On the other hand, a high volatility state at long time horizons does not necessarily imply a high volatility state at shorter time horizons.

Gençay *et al.* (2001b) presents a general framework for the basic premise of wavelets within the context of economic/financial time series. They illustrate that wavelets provide a natural platform to deal with the time-varying characteristics found in most financial time series at multiple time scales. Excellent recent reviews of wavelets from the finance perspective can also be found in Ramsey (1999, 2002).

One major use of wavelets is to study changes in (weighted) averages within and across time scales (intervals). We are often interested in how return, volatility or volume vary across overlapping (or nonoverlapping) time intervals. In particular, we may ask questions such as how 5-minute returns (or volatility) vary across the sample path. Having examined the changes in returns (or volatility) in 5-minute intervals, we may ask, for instance, what does this imply for the magnitude of changes in 30-min, daily or weekly returns (or volatility)? The questions of this sort can be examined without limiting the research to a particular arbitrary time interval and all possible time intervals can be examined simultaneously. The random walk hypothesis, for instance, need not be studied with a weekly or monthly data, but all time intervals from seconds to years can be studied simultaneously within the wavelet methodology. It is possible that a particular model may fit the data at a particular time interval but not otherwise, and therefore we identify those pockets in the data which exhibit different dynamics relative to a model.

A wavelet is a small wave which grows and decays in a limited time period.[4] To formalize the notion of a wavelet, let $\psi(.)$ be a real valued function such that its integral zero,

$$\int_{-\infty}^{\infty} \psi(t)\, dt = 0, \tag{1}$$

and its square integrates to unity,

$$\int_{-\infty}^{\infty} \psi(t)^2\, dt = 1. \tag{2}$$

---

[4]The contrasting notion is a big wave such as the sine function which keeps oscillating indefinitely.

While Equation 2 indicates that $\psi(.)$ has to make some excursions away from zero, any excursions it makes above zero must cancel out excursions below zero due to Equation 1, and hence $\psi(.)$ is a wave, or a wavelet.

Wavelets are, in particular, useful for the study of how weighted averages vary from one averaging period to the next. Let $x(t)$ be real-valued and consider the integral

$$\bar{x}(s,e) \equiv \frac{1}{s-e} \int_s^e x(u)\, du \tag{3}$$

where we assume that $e > s$. $\bar{x}(s,e)$ is the average value of $x(.)$ over the interval $[s,e]$. Instead of treating an average value $\bar{x}(s,e)$ as a function of end points of the interval $[s,e]$, it can be considered as a function of the length of the interval,

$$\lambda \equiv s - e$$

while centering the interval at

$$t = (s+e)/2.$$

λ is referred to as the scale associated with the average, and using λ and t, the average can be redefined such that

$$a(\lambda, t) \equiv \bar{x}(t - \frac{\lambda}{2}, t + \frac{\lambda}{2}) = \frac{1}{\lambda} \int_{t-\frac{\lambda}{2}}^{t+\frac{\lambda}{2}} x(u)\, du$$

where $a(\lambda, t)$ is the average value of $x(.)$ over a scale of λ centered at time t. The change in $a(\lambda, t)$ from one time period to another is measured by

$$w(\lambda, t) \equiv a(\lambda, t + \frac{\lambda}{2}) - a(\lambda, t - \frac{\lambda}{2}) = \tag{4}$$
$$\frac{1}{\lambda} \int_t^{t+\lambda} x(u)\, du - \frac{1}{\lambda} \int_{t-\lambda}^t x(u)\, du.$$

Equation 4 measures how much the average changes between two adjacent nonoverlapping time intervals, from $t - \lambda$ to $t + \lambda$, each with a length of λ. Because the two integrals in Equation 4 involve adjacent nonoverlapping intervals, they can be combined into a single integral over the real axis to obtain

5

$$w(\lambda, t) = \int_{-\infty}^{\infty} \tilde{\psi}(t) x(u) \, du \tag{5}$$

where

$$\tilde{\psi}(t) = \begin{cases} -1/\lambda, & t - \lambda < u < t, \\ 1/\lambda, & t < u < t + \lambda, \\ 0, & \text{otherwise.} \end{cases}$$

$w(\lambda, t)$'s are the wavelet coefficients and they are essentially the changes in averages across adjacent (weighted) averages.

2.1 Discrete wavelet transformation

In principle, wavelet analysis can be carried out in all arbitrary time scales. This may not be necessary if only key features of the data are in question, and if so, discrete wavelet transformation (DWT) is an efficient and parsimonious route as compared to the continuous wavelet transformation (CWT). The DWT is a subsampling of $w(\lambda, t)$ with only dyadic scales, i.e., λ is of the form 2^{j-1}, $j = 1, 2, 3, \ldots$ and, within a given dyadic scale 2^{j-1}, t's are separated by multiples of 2^j.

Let $\mathbf{x}$ be a dyadic length vector $(N = 2^J)$ of observations. The length N vector of discrete wavelet coefficients $\mathbf{w}$ is obtained by

$$\mathbf{w} = \mathcal{W}\mathbf{x},$$

where $\mathcal{W}$ is an $N \times N$ real-valued orthonormal matrix defining the DWT which satisfies $\mathcal{W}^T \mathcal{W} = I_N$ ($n \times n$ identity matrix).[5] The nth wavelet coefficient w_n is associated with a particular scale and with a particular set of times. The vector of wavelet coefficients may be organized into $J + 1$ vectors,

$$\mathbf{w} = [\mathbf{w}_1, \mathbf{w}_2, \ldots, \mathbf{w}_J, \mathbf{v}_J]^T,$$

where $\mathbf{w}_j$ is a length $N/2^j$ vector of wavelet coefficients associated with changes on a scale of length $\lambda_j = 2^{j-1}$ and $\mathbf{v}_J$ is a length $N/2^J$ vector of scaling coefficients associated with averages on a scale of length $2^J = 2\lambda_J$.

Using the DWT, we may formulate an additive decomposition of $\mathbf{x}$ by reconstructing the wavelet coefficients at each scale independently. Let $\mathbf{d}_j = \mathcal{W}_j^T \mathbf{w}_j$

---

[5]Since DWT is an orthonormal transform, orthonormality implies that $\mathbf{x} = \mathcal{W}^T \mathbf{w}$ and $||w||^2 = ||x||^2$.

define the jth level *wavelet detail* associated with changes in $\mathbf{x}$ at the scale λ_j (for $j = 1, \ldots, J$). The wavelet coefficients $\mathbf{w}_j = \mathcal{W}_j\mathbf{x}$ represent the portion of the wavelet analysis (decomposition) attributable to scale λ_j, while $\mathcal{W}_j^T\mathbf{w}_j$ is the portion of the wavelet synthesis (reconstruction) attributable to scale λ_j. For a length $N = 2^J$ vector of observations, the vector $\mathbf{d}_{J+1}$ is equal to the sample mean of the observations.

A multiresolution analysis (MRA) may now be defined via

$$x_t = \sum_{j=1}^{J+1} \mathbf{d}_{j,t} \quad t = 1, \ldots, N. \tag{6}$$

That is, each observation x_t is a linear combination of wavelet detail coefficients at time t. Let $\mathbf{s}_j = \sum_{k=j+1}^{J+1} \mathbf{d}_k$ define the jth level *wavelet smooth*. Whereas the wavelet detail $\mathbf{d}_j$ is associated with variations at a particular scale, $\mathbf{s}_j$ is a cumulative sum of these variations and will be smoother and smoother as j increases. In fact, $\mathbf{x} - \mathbf{s}_j = \sum_{k=1}^{j} \mathbf{d}_k$ so that only lower-scale details (high-frequency features) from the original series remain. The jth level *wavelet rough* characterizes the remaining lower-scale details through

$$\mathbf{r}_j = \sum_{k=1}^{j} \mathbf{d}_k, \quad 1 \le j \le J + 1.$$

The wavelet rough $\mathbf{r}_j$ is what remains after removing the wavelet smooth from the vector of observations. A vector of observations may thus be decomposed through a wavelet smooth and rough via

$$\mathbf{x} = \mathbf{s}_j + \mathbf{r}_j,$$

for all j.

A variation of the DWT is called the maximum overlap DWT (MODWT). Similar to the DWT, the MODWT is a subsampling at dyadic scales, but in contrast to the DWT, the analysis involves all times t rather than the multiples of 2^j. Retainment of all possible times eliminates alignment effects of DWT and leads to more efficient time series representation at multiple time scales. In this paper, we use the MODWT in our disentangling of the intraday from the interday dynamics.

2.2 Analysis of variance

The orthonormality of the matrix $\mathcal{W}$ implies that the DWT is a variance preserving transformation where

$$\|\mathbf{w}\|^2 = \sum_{j=1}^{J} \sum_{t=0}^{N/2^j - 1} w_{j,t}^2 + v_{J,0}^2 = \sum_{t=0}^{N-1} x_t^2 = \|\mathbf{x}\|^2 .$$

This can be easily proven through basic matrix manipulation via

$$\|\mathbf{x}\|^2 = \mathbf{x}^T\mathbf{x} \quad = \quad (\mathcal{W}\mathbf{w})^T\mathcal{W}\mathbf{w}$$
$$= \quad \mathbf{w}^T\mathcal{W}^T\mathcal{W}\mathbf{w} = \mathbf{w}^T\mathbf{w} = \|\mathbf{w}\|^2 \,.$$

Given the structure of the wavelet coefficients, $\|\mathbf{x}\|^2$ is decomposed on a scale-by-scale basis via

$$\|\mathbf{x}\|^2 = \sum_{j=1}^{J} \|\mathbf{w}_j\|^2 + \|\mathbf{v}_J\|^2 \,, \tag{7}$$

where $\|\mathbf{w}_j\|^2$ is the sum of squared variation of $\mathbf{x}$ due to changes at scale λ_j and $\|\mathbf{v}_J\|^2$ is the information due to changes at scales λ_J and higher. An alternative decomposition of $\|\mathbf{x}\|^2$ to Equation 7 is

$$\|\mathbf{x}\|^2 = \sum_{j=1}^{J} \|\mathbf{d}_j\|^2 + \|\mathbf{s}_J\|^2$$

which decomposes the variations in $\mathbf{x}$ across the variations in details and the smooth.

Percival and Mofjeld (1997) proved that the MODWT is an energy (variance) preserving transform such that the variance of the original time series is perfectly captured by the variance of the coefficients from the MODWT. Specifically, the total variance of a time series can be partitioned using the MODWT wavelet and scaling coefficient vectors by

$$\|\mathbf{x}\|^2 = \sum_{j=1}^{J} \|\widetilde{\mathbf{w}}_j\|^2 + \|\widetilde{\mathbf{v}}_J\|^2 \tag{8}$$

where $\widetilde{\mathbf{w}}_j$ is a length $N/2^j$ vector of MODWT wavelet coefficients associated with changes on a scale of length $\lambda_j = 2^{j-1}$ and $\widetilde{\mathbf{v}}_J$ is a length $N/2^J$ vector of MODWT scaling coefficients associated with averages on a scale of length $2^J = 2\lambda_J$.[6] This will allow us to construct MODWT versions of the wavelet variance.

3 Empirical Results

3.1 Data and preliminary analysis

Our data set is the Dow Jones Industrial Average (DJIA), recorded at 5-minute intervals during the sample period of September 19, 1994 - October 16, 2002.[7] The

---

[6]More information on the MODWT transformation can be found in Percival and T. (2000).

[7]Statistical properties of the same data set are studied extensively by Gençay and Selçuk (2004).

New York Stock Exchange opens at 9:30 a.m. (EST) (13:30 GMT) and the first record of the DJIA index for that day is registered at 9:35 a.m. The market closes at 4:00 p.m. (EST) (20:05 GMT) and the last record of the day is registered at 4:05 p.m.[8] Therefore, there are 79 index records at 5-minute intervals during one business day. We eliminated weekends, holidays, and other days that the market was closed.

The 5-minute stock market return is defined as

$$r_t = \log x_t - \log x_{t-1}, \quad t = 1, 2, \ldots, 151,443, \tag{9}$$

where x_t is the DJIA level at time t. The volatility is defined as squared return, r_t^2.

In addition, we eliminated the days in which there were at least 12 consecutive 5-minute zero returns.[9] The opening 5-minute return each day is not a "true" 5-minute return since it is the log difference between the previous days's close and the first record of the index at 9:35 a.m. during the day. Therefore, we eliminated the opening returns from each day. As a result, there are 78 5-minute returns each day, resulting in 149,526 sample points and covering 1,917 business days.

The estimated autocorrelation coefficients of returns at 5-minute intervals are plotted against their lags along with 95 percent confidence intervals in Figure 1(b). There is a significant autocorrelation at the first two lags ((0.038 and -0.02, respectively). The positive autocorrelation for intraday and daily returns on stock portfolios is well known, see, for example, O'Hara (1995), Campbell *et al.* (1997) and Ahn *et al.* (2002) and references therein. More recently, Bouchaud and Potters (2000) has reported significant autocorrelations (up to four lags) of the S&P 500 increments measured at 5-minute intervals. One possible explanation for the positive autocorrelation at the first lag is nonsynchronous trading. That is, one group of stocks in an index may react to new information slower than others which results in a strong autocorrelation for return on the index (Ahn *et al.*, 2002). On the other hand, it is well known that bid-ask bounce leads to negative autocorrelation in stock returns (Roll, 1984). Therefore, one may argue that the true first-order autocorrelation coefficient of intraday returns may be higher than the one reported here. Figure 1(a) also reports average returns at 5-minute intervals during the day. The first (9:35 a.m. to 9:40 a.m.) and the last (4:00 p.m. to 4:05 p.m.) 5-minute average returns are significantly higher than any other 5-minute average returns. Although the opening effect may have some influence on the first 5-minute return in our sample, there is no intuitive explanation for the relatively very high return at the last 5-minute.

---

[8]See Harris (2003) for extensive coverage of the market mechanism in the NYSE.

[9]We eliminated the entire day, not just consecutive zero-return periods during the day, to keep the frequency characteristics of the data set intact. We thank Olsen Data Inc. (www.olsendata.com) for providing the raw data set.
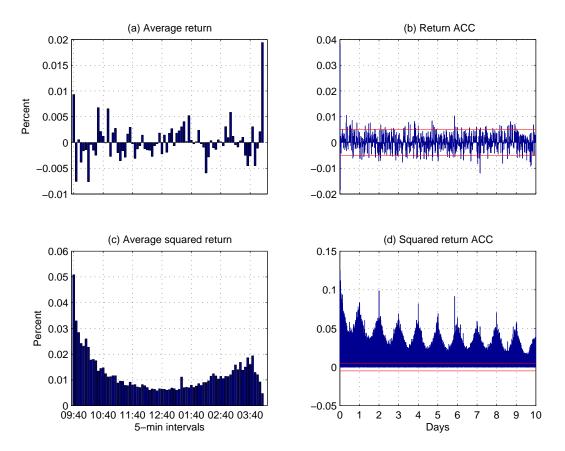
Figure 1: Autocorrelation coefficients (ACC) and intraday averages of the 5-minute DJIA return and volatility. (a) Average return during the day in sample period (in percent). (b) Autocorrelation coefficients of 5-minute returns at 5-minute lags up to 10 days. The first two lag autocorrelation coefficients are statistically significant (0.038 and -0.02, respectively). 95 percent confidence intervals are plotted as dash-dot lines. (c) Average volatility (average squared return) during the day. (d) Autocorrelation coefficients of 5-minute squared returns at 5-minute lags up to 10 days. 95 percent confidence intervals are plotted as dash-dot lines. Sample period is September 19, 1994 - October 16, 2002 (149,526 5-minutes, 1,917 days). Data source: Olsen Data Inc. (www.olsendata.com).

The volatility clustering is evident in Figures 1(c) and 1(d). Figure 1(d) studies the sample autocorrelation coefficients of 5-minute volatility, defined as squared return, at 5-minute lags up to 10 days. There is a significant peak at lag 78 which indicates that there is a strong seasonal cycle which completes itself in one day. Regarding the weekly seasonality, we do not observe a strong peak at 5 days nor at an

integer multiple of 5 days. However, this observation should be interpreted with caution since the presence of very strong short-period seasonalities (periodicities) may obscure relatively weak long-period seasonal dynamics. Figure 1(c) plots average volatility at 5-minute intervals. The highest average volatilities are observed at the opening of the session, especially during the first 45-minutes. The average volatility drops afterwards, reaches a minimum during lunch hours and slightly increases again. Except for the last 20 minutes, the volatility has a U shape during the day, similar to the U shape of volatility autocorrelations. The minimum 5-minute average volatility is observed at the last 5-minute (closing). An early study by Amihud and Mendelson (1986) reported that trading at the opening exposes traders to a greater variance than in the close. They attribute this difference to the trading mechanism in the NYSE. Recently, Barclay and Hendershott (2003) have reported that relatively low after-hours trading can generate significant price discovery. When the market opens, this information advantage is probably reflected on prices causing a larger volatility during the opening than the rest of the day. Unlike the the average return series, there is no significant increase in volatility during the last 5-minutes. In fact, the volatility starts to decline at 3:45 p.m. and reaches its daily minimum in the last 5-minutes.

3.2 Interday and intraday decomposition

The wavelet additive decomposition of 5-minute returns into intraday and interday scales is performed utilizing a maximum overlap discrete wavelet transform (MODWT) multiresolution analysis (MRA) with $j = 6$. Particularly, the Daubechies least asymmetric family of wavelets, LA(8), is utilized in this decomposition. The highest level detail d_{6t} (level 6 detail) captures frequencies $\frac{1}{128} \leq f \leq \frac{1}{64}$ while the 6th level smooth r_t^s contains oscillations with period length of 128 and higher. Since there are 78 5-minute returns per day, details from 1 to 6 contain all intraday and daily dynamics while the smooth component represents interday (over 1.5 days) dynamics. The intraday returns r_t^r are defined as

$$r_t^r = \sum_{j=1}^{6} d_{jt} \tag{10}$$

where d_{jt} is the jth level wavelet detail of returns associated with changes in r_t at scale 2^{j-1}. The original 5-minute return series r_t are therefore

$$r_t = r_t^r + r_t^s \tag{11}$$

Figure 2 plots both intraday and interday volatility series along with the original 5-minute squared returns. Intraday volatility is defined as the vertical sum of the

11

squared wavelet detail coefficients at the first six details while interday volatility is the squared wavelet smooth coefficient. Figure 2 reveals that there is an increase in both intraday and interday volatility after the second half of 1997 (following the Asian financial crisis). However, the relative increase in volatility at the interday scale seems to be higher than the relative increase in the intraday scale. We are also able to identify occasional jumps in volatility series in different scales at different times. For example, there is a relative increase in interday volatility as compared to previous days while there is not that much relative increase in intraday volatility at around March 8, 2000 (see Figure 2(c) and (e)). Figure 2 also plots the sample autocorrelation coefficients for all three series. It shows that both intraday and interday volatility series have much stronger autocorrelation structure than what it appears with autocorrelation structure based on 5-minute squared return series. Overall, this examination indicates that the wavelet decomposition is a very useful tool in disentangling complex dynamics of any given time series in terms of different time scales.

We now proceed with cross-correlations between the return and the volatility series at different time scales. Figure 3(a) plots the sample cross-correlation between 5-minute return r_t against future squared return r_{t+i}^2. Although there is a persistent negative cross-correlation between r_t and r_t^2 for several days, the magnitude of the correlation coefficients seems to be very low (less than -2 percent). When we decompose the volatility and the return, there is no significant cross-correlation between intraday return and future volatility (Figure 3.b), except for the first four lags (20 minutes). Figure 3(c), however, reveals that the current return and future volatility have a much stronger correlation than what it appears with the original 5-minute return series. A small but significant positive correlation is observed at one day lag. The correlation reverses itself after one day and becomes negative. Contrary to the raw sample cross-correlation coefficients, the negative relationship between the current return and the future volatility is high and persistent. This result provides evidence for the leverage hypothesis which states that the current return and future volatility are negatively correlated.

Figure 3(d) plots the sample cross-correlation between 5-minute squared return r_t^2 against future return r_{t+i}. Although there are some significant cross-correlation coefficients (both positive and negative) at certain leads, there is no systematic pattern. There is also no significant cross-correlation between intraday volatility and future intraday return (Figure 3.e). Figure 3(f), however, reveals that the current interday volatility and future interday return have a strong correlation. Although it is small in magnitude (around 5 percent) as compared to the cross-correlation between the current return and the future volatility, it remains positive for several days. Notice that the volatility feedback hypothesis implies a very short term negative relationship between the current volatility and the future return. As
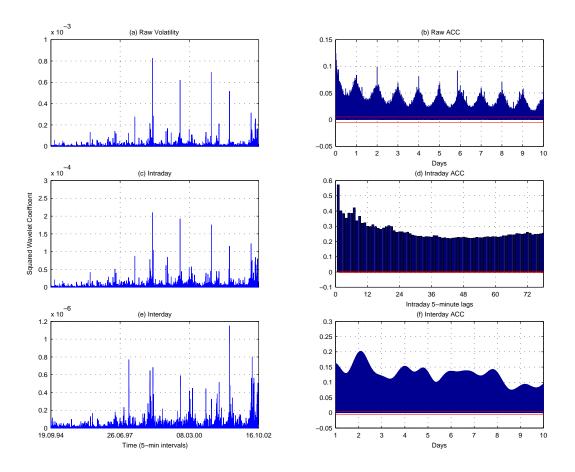
Figure 2: Row (top left), intraday (middle left) and interday (bottom left) volatilities and corresponding sample autocorrelation coefficients (right column). 95 percent confidence intervals are plotted as solid horizontal lines. The raw 5-minute squared return is the sum of intraday and interday squared returns at every 5-minutes. Sample period is September 19, 1994 - October 16, 2002 (149,525 5-minutes, 1,917 days). Data source: Olsen Data Inc. (www.olsendata.com).

time passes, this negative relationship between volatility and return must reverse and the current volatility and the future return must be positively correlated because of the required higher risk premium. Figure 3(f) shows that this is indeed the case for the DJIA return and volatility. Overall, we conclude that both the leverage effect and the volatility feedback play an important role in determining the return-volatility dynamics in stock markets.

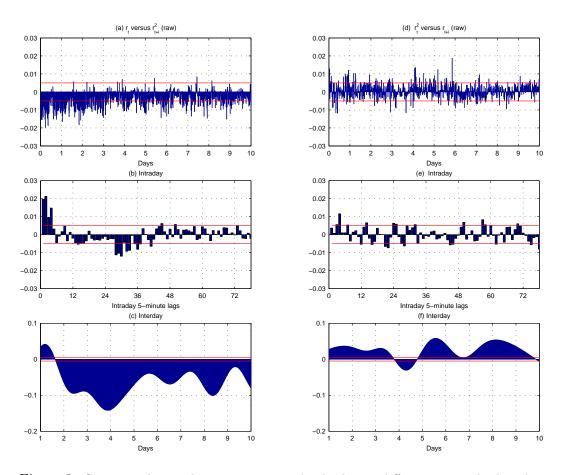We can compare our findings to those obtained with traditional methodology

13

Figure 3: Cross correlations between return and volatility at different timescales based on the wavelet analysis: (a) raw data at 5-minutes: return r_t versus squared return r^2_{t+i} (b) intraday scale: intraday returns versus intraday volatility (squared rough wavelet coefficients) (c) interday scale: interday return versus interday volatility (squared smooth wavelet coefficients). (d) raw data at 5-minutes: squared return r^2_t versus future return r_{t+i} (e) intraday scale: intraday volatility (squared rough wavelet coefficient) versus future intraday returns (f) interday scale: interday volatility (squared smooth wavelet coefficient) versus future interday returns. 95 percent confidence intervals are plotted as solid horizontal lines. The raw data sample period is September 19, 1994 - October 16, 2002 (149,525 5-minutes, 1,917 days). Data source: Olsen Data Inc. (www.olsendata.com).

of multihorizon analysis. One approach to obtain a lower frequency volatility from the high-frequency data is aggregation. In this approach, 5-minute returns are aggregated for a certain time period and the lower frequency returns obtained after
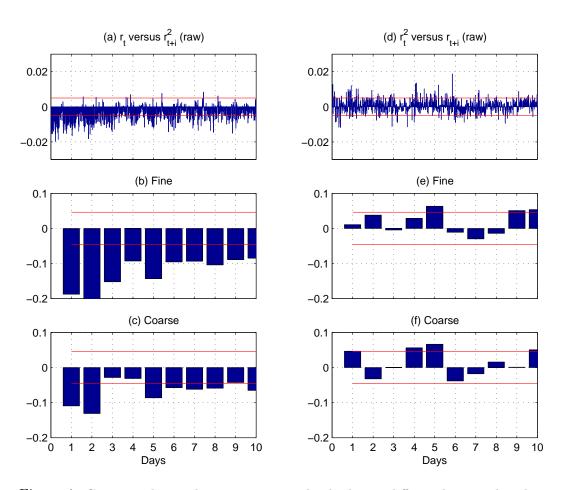
Figure 4: Cross correlations between return and volatility at different horizons based on the fine (realized) and coarse volatilities: (a) raw data at 5-minutes: return r_t versus future volatility (squared return r_{t+i}^2) (b) daily return versus future daily realized volatility (c) daily return versus future daily coarse volatility (d) raw data at 5-minutes: squared return r_t^2 versus future return r_{t+i} (e) daily realized volatility versus future daily returns (f) daily coarse volatility versus future daily returns. 95 percent confidence intervals are plotted as solid horizontal lines. The raw data sample period is September 19, 1994 - October 16, 2002 (149,525 5-minutes, 1,917 days). Data source: Olsen Data Inc. (www.olsendata.com).

aggregation are squared. This volatility measure is labeled as *coarse* volatility (Dacorogna *et al.*, 2001). For example, daily coarse volatility is obtained by aggregating

5-minute returns and squaring the resulting aggregated return for each day:

$$r_{ct}^2 = \left(\sum_{i=1}^{78} r_{it} \right)^2 \qquad t = 1, 2, \ldots, 1,917. \tag{12}$$

Another measure is the *fine* (or realized) volatility. In this approach, squared returns at high frequency are aggregated to obtain a lower frequency measure of volatility. For example, daily fine volatility is obtained by aggregating 5-minute squared returns

$$r_{ft}^2 = \sum_{i=1}^{78} r_{it}^2 \qquad t = 1, 2, \ldots, 1,917. \tag{13}$$

Figure 4 plots sample cross-correlations between the current return and fine (realized) and coarse volatilities at daily horizon. Regarding the leverage effect, the magnitude of the cross-correlation between the current daily return and future daily fine volatility is similar to what we obtained with the wavelet analysis at interday scale. An important difference is that the wavelet based cross-correlation is weak at the beginning and becomes stronger after two days. Similarly, sample cross-correlation between the current return and future coarse volatility is negative. However, the magnitude of the cross-correlation is smaller as compared to fine volatility. Figure 4(e,f) provides no evidence for volatility feedback hypothesis: sample cross-correlation between the current daily volatility (fine or coarse) and the future daily return is statistically not significant whereas the wavelet based cross-correlations at interday scales are significant for several days. We conclude that the wavelet based cross-correlation analysis captures important information at interday scales where realized or coarse volatility measures may fail.

4 Conclusions

This paper proposes a simple yet powerful methodology to investigate volatility return dynamics at different time scales. The proposed method is based on a wavelet multi-scaling approach which decomposes the data into its low- and high-frequency components. Using this methodology, we showed that the leverage effect is weak at short-horizons, and becomes prominent at lower frequencies (long horizons) at the interday scale. On the other hand, the volatility feedback is a very short lived phenomenon. The positive correlation between the current volatility and future returns becomes dominant after one day at the interday scale, providing evidence that the risk and return are positively correlated. We also compared our results to traditional measures of volatility at different horizons and showed that the wavelet based cross-correlation analysis captures important information at interday scales where realized or coarse volatility measures may fail.

16

References

Ahn, D., Boudoukh, J., Richardson, M., and Whitelaw, R. F. (2002). Partial adjustment or stale prices? Implications from stock index and futures return autocorrelations. *Review of Financial Studies*, **15**, 655–689.

Amihud, Y. and Mendelson, H. (1986). Trading mechanisms and stock returns: An empirical investigation. *Journal of Finance*, **42**, 533–553.

Barclay, M. J. and Hendershott, T. (2003). Price discovery and trading after hours. *Review of Financial Studies*, **16**, 1041–1073.

Bekaert, G. and Harvey, C. R. (1997). Emerging equity market volatility. *Journal of Financial Economics*, **43**, 29–77.

Bekaert, G. and Harvey, C. R. (2003). Emerging markets finance. *Journal of Empirical Finance*, **10**, 3–56.

Bekaert, G. and Wu, G. (2000). Asymmetric volatility and risk in equity markets. *The Review of Financial Studies*, **13**, 1–42.

Black, F. (1976). Studies of stock price volatility changes. In *Proceedings of the 1976 Meetings of the Business and Economic Statistics Section, American Statistical Association,* 177-181.

Bollerslev, T. and Zhou, H. (2003). Volatility puzzles: A unified framework for gauging return-volatility regressions. *Finance and Economics Discussion Series 2003-40,* Washington: Board of Governors of the Federal Reserve System.

Bouchaud, J. P. and Potters, M. (2000). *Theory of Financial Risks: From Statistical Physics to Risk Management.* Cambridge University Press, Cambridge, UK.

Campbell, J. Y. and Hentschel, L. (1992). No news is good news: An asymmetric model of changing volatility in stock returns. *Journal of Financial Economics*, **31**, 281–318.

Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1997). *The Econometrics of Financial Markets.* Princeton University Press, Princeton, NJ.

Christie, A. A. (1982). The stochastic behaviour of common stock variances–value, leverage and interest rate effects. *Journal of Financial Economics*, **10**, 407–432.

Dacorogna, M., Gençay, R., Muller, U., Olsen, R., and Pictet, O. (2001). *An Introduction to High-Frequency Finance.* Academic Press, San Diego.

French, K. R., Schwert, G. W., and Stambaugh, R. F. (1987). Expected stock returns and volatility. *Journal of Financial Economics*, **19**, 3–30.

Gençay, R. and Selçuk, F. (2004). Intraday dynamics of stock market returns and volatility. Manuscript, Department of Economics, Simon Fraser University.

Gençay, R., Selçuk, F., and Whitcher, B. (2001a). Differentiating intraday seasonalities through wavelet multi-scaling. *Physica A*, **289**, 543–556.

Gençay, R., Selçuk, F., and Whitcher, B. (2001b). *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. Academic Press, San Diego.

Gençay, R., Selçuk, F., and Whitcher, B. (2001c). Scaling properties of foreign exchange volatility. *Physica A*, **289**, 249–266.

Gençay, R., Selçuk, F., and Whitcher, B. (2003). Systematic risk and time scales. *Quantitative Finance*, **3**, 108–116.

Gençay, R., Selçuk, F., and Whitcher, B. (2004a). Asymmetry of information flow between volatilities across time scales. Manuscript, http://www.sfu.ca/∼rgencay/whmm.pdf.

Gençay, R., Selçuk, F., and Whitcher, B. (2004b). Multiscale systematic risk. *Journal of International Money and Finance*, (forthcoming).

Glosten, L., Jagannathan, R., and D.Runkle (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, **48**, 1779–1801.

Harris, L. (2003). *Trading&Exchanges: Market Microstructure for Practitioners*. Oxford University Press, Oxford.

Harrison, P. and Zhang, H. H. (1999). An investigation of the risk and return relation at long horizons. *The Review of Economics and Statistics*, **81**, 399–408.

Nelson, B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, **59**, 347–370.

O'Hara, M. (1995). *Market Microstructure Theory*. Blackwell Publishers Ltd, Oxford.

Percival, D. B. and Mofjeld, H. O. (1997). Analysis of subtidal coastal sea level fluctuations using wavelets. *Journal of the American Statistical Association*, **92**, 868–880.

Percival, D. B. and T., W. A. (2000). *Wavelet Methods for Time Series Analysis.* Cambridge Press, Cambridge.

Ramsey, J. B. (1999). The contribution of wavelets to the analysis of economic and financial data. *Philosophical Transactions of the Royal Society of London A*, **357**, 2593–2606.

Ramsey, J. B. (2002). Wavelets in economics and finance: Past and future. *Studies in Nonlinear Dynamics & Econometrics*, **6**, Article 1.

Roll, R. (1984). A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, **39**, 1127–1139.

Selçuk, F. (2004). Asymmetric stochastic volatility: An application to emerging stock markets. *Mimeo,* Bilkent University, Department of Economics, Ankara, Turkey.

Wu, G. (2001). The determinants of asymmetric volatility. *The Review of Financial Studies*, **14**, 837–859.