A multi-horizon scale for volatility

Alexander Subbotin

A Multi-Horizon Scale for Volatility

Alexander SUBBOTIN

2008.20

A Multi-Horizon Scale for Volatility [*]

Alexander Subbotin [†]

March 3, 2008

Abstract

We decompose volatility of a stock market index both in time and scale using wavelet filters and design a probabilistic indicator for volatilities, analogous to the Richter scale in geophysics. The peak-over-threshold method is used to fit the generalized Pareto probability distribution for the extreme values in the realized variances of wavelet coefficients. The indicator is computed for the daily Dow Jones Industrial Average index data from 1896 to 2007 and for the intraday CAC40 data from 1995 to 2006. The results are used for comparison and structural multi-resolution analysis of extreme events on the stock market and for the detection of financial crises.

Keywords: Stock market, Volatility, Wavelets, Multi-resolution Analysis, Financial Crisis

J.E.L. Classification: G.10, G.14.

1   Introduction

Volatility of a stock market index is a natural candidate to characterize the state of the stock market and compare the impact of important events, such as financial crises. The need to compare events on the stock market arises in many applications: from structural analysis, studying the dynamics of the stock market in a general economic context, to asset allocation, for which measuring the severity and time bounds of extreme fluctuations is of interest. In principle, such analysis can be based on the values of some direct volatility estimates. But these estimates are hard to interpret and their magnitude varies for different assets. One way to overcome this inconvenience is to design a universal indicator, based on a probability transform of volatility measures.

Zumbach et al. (2000) proposed an intuitively appealing analogy with the Richter-Gutenberg scale in geophysics and constructed an indicator of market volatility, called a Scale of Market Shocks (SMS). The Richter-Gutenberg scale is a measure of the logarithm of the seismic wave amplitude, which is related to the total energy, liberated during an earthquake (Richter, 1958). The SMS maps the exchange rate volatility to a logarithmic scale, suggesting the analogy between volatility and mechanical work, i.e. the rate of change of energy in time. The indicator was originally applied to the currency market, but by construction it can be used for any traded asset, for which historical quotes are available.

The SMS accounts not only for the size but also for the scale of fluctuations. According to the heterogeneous market hypothesis, first proposed by Müller et al. (1997), the price dynamics is driven by the actions of market participants, operating at different frequencies. To make this idea clear, compare an institutional investor who operates over medium and long term targets and a small private investor exploiting short term market moves. Hardly do they share the same opinion on what is trend and fluctuation on the market. However, sometimes they would agree to characterize the situation on the market as a crisis or a crash. The underlying intuition is to describe volatility in the time domain and in the spectral domain simultaneously, thus attributing fluctuations to particular ranges of frequencies (scales). Economically, these scales can be associated to the decision taking and portfolio readjustment horizons of investors. The SMS takes the form of an average realized volatility across different scales.

Maillet and Michel (2003, 2005) adapted the multiscale approach to the stock market. The new indicator, called Index of Market Shocks (IMS), is designed for the detection and comparison of severity of different crises. The authors change the way of computing and aggregating volatilities at multiple frequencies. Besides, Instead of modeling their probability distributions, the principal components analysis is applied to estimate the hidden factors of volatility, affecting various scales. These factors are then subject to the

2

probability transform.

Following the approach suggested in Zumbach et al. (2000) and Maillet and Michel (2003, 2005), we study the multiscale indicators of volatility and propose a new Market Volatility Scale (MVS). Though the general idea of such indicators is very appealing, the existing implementations suffer from several important drawbacks, which motivates our work. In our view, the main problem with SMS is the scheme of multi-resolution analysis which is based on sequential smoothing rather than on the scale-by-scale decomposition, which does not allow to determine the relative importance of different horizons. This problem subsists for the IMS, which in addition makes the computation of the indicator more complicated by introducing the latent factors of volatility.

We argue that modeling latent factors instead of volatilities themselves leads to the loss of simplicity and interpretation errors, and propose to work with the scale components of volatility directly. The MVS, which we introduce in this paper, uses the realized wavelet variance to represent volatilities at multiple horizons. The log-normality assumption, which is used for the probability transform in the IMS, is clearly too restrictive. This is why we apply the peak-over-threshold approach to fit the generalized Pareto density for the tails of the realized variances' distribution at each scale. This allows more accurate estimation of the probability of extreme events, such as stock market crashes.

We test our approach on two types of data: high-frequency (15-minute intervals) observations of the French CAC40 index from 1995 to 2006 and daily observations of the Dow Jones Industrial Average (DJIA) index from 1896 to 2007. In these two examples the ranges of frequencies, to which we refer, are not the same, but the computation algorithm is quite similar. We detect periods of extreme volatility on the stock market, often provoked by financial crises, and study their structural characteristics using the scale-by-scale decomposition. This enables us to determine which component of volatility (short-term, mid-term or long-term) is more important in each detected period.

Potentially, the applications of multiscale decomposition of volatility can go far beyond event comparison and structural analysis. Most promising are volatility forecasting and asset allocation. These areas fall out of the scope of this paper, but the tools and ideas introduced here can be used in further studies.

The rest of the paper is organized as follows. In section 2 we briefly overview the heterogeneous market hypothesis, the existing multiscale indicators and discuss their main drawbacks. Section 3 introduces some wavelet formalism, used for the scale-by-scle decomposition of volatility. Section 4 defines the Market Volatility Scale and the associated computation algorithm. In section 4 we compute the MVS for the DJIA and CAC40 indexes and present the main results. Section 4 concludes.

2 Multiscale Indicators of Volatility

Multiscale indicators of volatility are based on the model of financial market with multiple horizons, known as the heterogeneous market hypothesis. First suggested in Müller et al. (1997), it attempts to explain well-known stylized facts about volatility of stock returns, such as long-range correlations in the time series of squared returns and absolute returns and volatility clustering (Ding et al., 1993; Lux, 1996; Andersen and Bollerslev, 1997; see Cont, 2001 for an overview on the subject). The central point is the scale-by-scale analysis of volatility which assumes that market data contains patterns specific to peculiar frequencies of observations and are thus of interest for different types of market agents. The latter may include intra-day speculators, daily traders, portfolio managers and institutional investors, each having their characteristic period of reaction to news and frequency of intervention in the market. A multiscale dynamic model for volatility was proposed in Müller et al. (1997) and Dacorogna et al. (1998), who represent current volatility as a function of squared returns over different horizons:

$$\sigma_t^2 = c_0 + \sum_{j=1}^{n} c_j \left(\sum_{i=1}^{j} r_{t-i} \right)^2 \tag{1}$$

where $c_k \geq 0$ for all $k = 0, ..., n$ (for $k = 0$ and $k = n$ the inequality is strict), and r_t is logarithmic return. Thus the sum $\sum_{i=1}^{j} r_{t-i}$ represents the log-return for the period of length j. By construction, the resulting HARCH (Heterogeneous ARCH) model captures the hierarchical structure of volatility correlations. More recently Corsi (2004) proposed a simple additive cascade model, inspired by the HARCH, which uses realized volatilities over some representative scales (daily, weekly, monthly) to model current daily realized volatility.

From a more theoretical perspective Ghashghaie et al. (1996) and Breymann et al. (2000), inspired by the model of turbulence in hydrodynamics, designed a multiplicative cascade model which decomposes the stochastic volatility process into a set of scales, so that the volatility at each scale depends on the volatility at lower frequency and some stochastic renewal process. The probability of renewal is chosen so that the mean interval between two renewals is equal to the time length of the associated scale. Arneodo et al. (1998) shows that under the assumptions of the model the covariance between logarithms of absolute returns at each scale should decrease with time as logarithmic function, so that returns simulated from the model exhibit dynamic patterns fairly close to empirical evidence. For the moment this model has mainly theoretical interest because the methods of its identification on real data are not elaborated.

Zumbach et al. (2000) retain the multiscale concept, but their aim is to characterize the vector of volatilities across various scales rather than impose

a dynamic model for predicting instantaneous volatility. The multiscale indicator, called the Scale of Market Shocks (SMS) is a probability transform of volatilities at different scales. It is based on an analogy with the Richter-Gutenberg scale in geophysics (Richter, 1958).

The Richter scale is a measure of the logarithm of the total energy E liberated during an earthquake, compared to a reference level E_0:

$$R \approx \ln\left(\frac{E_0}{E}\right) \qquad (2)$$

According to recent evidence, the probability of occurrence of large earthquakes grouped within temporal clusters of high seismic activity obeys the inverse power law (see Christensen et al., 2002; Mega et al., 2003). The probability to observe an earthquake of energy E reads:

$$p(E) \approx \left(\frac{E_0}{E}\right)^{\kappa}, \qquad (3)$$

with κ the scaling parameter. Using (3) the Richter scale can be rewritten:

$$R \approx \ln\left(\frac{E_0}{E}\right) = -\frac{1}{\kappa}\ln p(E) \qquad (4)$$

.

By analogy with (4), Zumbach et al. (2000) map the exchange rate volatility to a logarithmic scale. Volatility stands as a counterpart of mechanical work, i.e. the rate of change of energy in time. The corresponding formula is:

$$SMS_t = -\alpha \ln P(\sigma_t) \qquad (5)$$

with $P(\sigma_t)$ the distribution function associated with the volatility σ_t and α a scaling parameter. Note that the analogy is not direct because the probability distribution in (5) is not a pure power law and must be fitted.

The indicator is multiscale because the volatility used in (5) is computed separately for various "scales" of observations. An estimate of the realized price volatility for a given scale, used for that purpose, reads:

$$\widetilde{\sigma}_t^{(k)} = \left(\frac{\sum_{i=1}^{M_k-1} r_{t,\delta}^2}{M_k - 1}\right)^{1/2} \qquad (6)$$

with $\sigma_t^{(k)}$ the volatility of returns at scale k, $r_{t,\delta}$ the log-returns computed for time intervals δ of length $\tau_k(M_k-1)^{-1}$, τ_k the time length of scale k and M_k the number of observations available for that period. As the indicator was constructed for inhomogeneous tick-by-tick data, (6) could not be applied directly and the authors used smoothed volatilities, computed over moving windows, of the form:

$$\widetilde{\sigma}_t^{(k)} = \int_{t-2\tau_k}^{t} K(\phi)\omega_\phi^{(k)}\sigma_\phi^{(k)}d\phi \qquad (7)$$

with $K(\cdot)$ an appropriately chosen kernel function and $\omega_\phi^{(k)}$ a scale dependent weight. The final aggregated indicator is computed as the weighted average for all scales in the following way:

$$SMS_t = -\sum_{k=1}^{m} w_k \ln P\left(\sigma_t^{(k)}\right), \tag{8}$$

with w_k a convolution kernel, measuring the contribution of scale τ_k to the overall volatility, and $P(\cdot)$ is the complement of the cumulative density function.

A potential problem is that the scaling method suggested in Zumbach et al. (2000) gives no idea about the range of scales to be considered and their relative importance. Multi-resolution analysis consists in sequential smoothing rather than scale-by scale decomposition. The authors assert that the choice of the convolution kernel is not a crucial issue and it suffices to take a function which satisfies some regularity conditions and tends to zero at the edges of the range of frequencies, which is 15 minutes - 64 days in their case. To our knowledge, this choice is made *a priori* and is not based on any economic notion or on statistical properties of the time series. The assumption that the mass point, or the most important frequency, is located in the center of the range is also doubtful.

Besides that, the SMS does not account for the possible interdependence between scales. The aggregated indicator is a function of transformed univariate cumulative density functions and not of a multivariate *cdf*.

Maillet and Michel (2003, 2005) account for the dependence between scales and propose a new indicator called the Index of Market Shocks (IMS). Instead of scaled volatilities themselves, principal components are used in its definition:

$$IMS_t = -\sum_{i=1}^{q} \omega_i \log_2 P(c_i) \tag{9}$$

with $c_i, \ldots, c_q$ normalized principal components (or factors underlying multiscale volatility), and ω_i the weights of each component determined as the portion of variance of the data explained by the corresponding component. The IMS also uses a different method of multi-resolution analysis. Instead of computing volatilities for different scales by successively moving a window of fixed length, the IMS uses different frequencies of sampling within each interval to obtain the scale components of volatility.

Several shortcomings of the IMS can be mentioned. First, the method of estimating volatilities at different scales by varying sampling frequency has one important drawback which becomes very important if applied to low-frequency data. If the length of computation window is fixed and volatilities are calculated from the samples with different number of points, the results of estimation will be different in terms of statistical error and become incomparable. For example, if the length of computation window is fixed at 60

6

business days, the daily estimate is an average of 59 points and the monthly volatility is an average of only 3 points. Estimated monthly volatilities become counter-intuitively less smooth than the times series of daily volatility because of statistical errors. On the contrary, the daily average computed over a 3-months interval smooths out all interesting dynamics, corresponding in fact to the daily frequencies. Hence such multi-resolution analysis becomes inadequate.

Second and more important, replacing volatilities by their linear combinations (factors) changes the meaning of the indicator, which is not explicited in Maillet and Michel (2003). The authors claim that the IMS has a clear economic sense: since a logarithm in base 2 is used in (9), a one point increase of the IMS corresponds to the volatilities vector twice as unlikely. This is obviously true for the baseline Richter transformation in the form analogous to (5), but not for (9), because the multivariate function of a vector is not the same as the distribution function of a linear combination of its components. Hence a one point increase of the IMS corresponds to the factors vector half as likely. But the factors, obtained from the Principal Components Analysis (PCA) do not necessarily have positive or negative impact on the volatilities at all scales. They can have positive loadings in some scales and negative in others. Thus it is possible that an important change in some factor corresponds to a "redistribution" of volatility from one scale to another rather than increase in volatility. As the PCA is done on the variance-covariance matrix, the volatilities at different scales are treated as if they were all of equal importance to investors. All this makes the interpretation of the resulting indicator very problematic.

Modeling factors of volatility can be of interest if these factors are stable and meaningful. In our view, this is not the case. First, the definition of factors by the PCA is too restrictive: it relies on linear decorrelation and may not have the power to identify factors if they are not symmetrically distributed (Loretan, 1997) or if the underlying model is not a simultaneous linear mix. It is unrealistic to suppose that the underlying model for multi-scale volatility is a simultaneous mix, because it contradicts to the very idea of the heterogeneous market.

Economic intuition suggests that investors acting on lower scales adjust their market positions according to the change in information since the previous observation. Investors acting at high frequencies react to information more often. But as the information set in the first case encompasses all the information that has already been used by the short-term investors (possibly accompanied by some long-term factors), it is natural to suppose that the long-scale volatility is correlated with the lagged short-scale volatilities. Yet the PCA allows for the instantaneous decorrelation only and the resulting factors can hardly correspond to the true hidden factors of volatilities. Studying the properties of the volatilities vector empirically, we find that the factor loadings in the PCA are unstable in time and principal components

are dependent if lags are considered.[1] Unable to propose a simple and viable method of factor analysis, we prefer to model the volatility vector directly, but in a way different from that proposed in Zumbach et al. (2000).

Finally, the probability transform used to compute the indicator is an important issue. Maillet and Michel (2003) postulate the multi-normality for log-volatilities (and principal components) at all scales, though they report that the log-normality is violated for the high values of volatilities. This can affect the values of the IMS during extreme events. Meanwhile, these values are of great importance for analysis of financial crises. The SMS, in its turn, uses a more sophisticated modification of the log-normal *pdf*, used to account for the fat tails. In our view, this modification does not allow enough flexibility, because an *a priori* judgment about the scaling properties of financial time series has to be made. It can be accurate for some assets and not for others. We argue that to define a universal indicator we do not necessarily need an overall fit for volatility. Instead we focus on the extreme events only, using the peak-over-threshold approach to fit the generalized Pareto distribution for the tails (see section 4 for details).

As follows from the above discussion, the most important improvements to be made in multiscale indicators of volatility concern the multi-resolution analysis by itself and the fit of probability distribution for the volatilities vector. These improvements are discussed in the following section.

3   The Scale-by-Scale Decomposition of Volatility with Wavelets

In this section we propose a method of multiscale analysis of volatility, based on the wavelet filters. We introduce the realized variance of wavelet coefficients, which is used to define the Market Volatility Scale in the following section.

Wavelet analysis is widely used in signal processing in order to decompose a given time series called "signal" into a hierarchical set of approximations and details (multi-resolution analysis), and to decompose the energy of the signal on scale-by-scale basis. For a detailed description of the use of wavelets for time series analysis refer to Percival and Walden (2000) and Gençay et al. (2001a). Within the financial markets context, wavelets were used for multiscale analysis in several studies. An early application of wavelets can be found in Ramsey and Zhang (1995) who attempt to recover the dynamic structures of the foreign exchange rates. Capobianco (2004) identifies periodical patterns in intra-day stock returns. Fernandez and Lucey (2007) use wavelets to estimate value at risk for multiple investment horizons in a

---

[1]The discussion on factor analysis of multiscale volatility falls out of the scope of this paper, these preliminary results can be obtained from the author on request.

portfolio management perspective. Lee (2004) studies the effects of volatility spillovers from the US to the South Korean stock markets at multiple horizons.

The idea to use the wavelet decomposition for the analysis of volatility cascade can be found in Arneodo et al. (1998). Gençay et al. (2001a) used wavelets to investigate the scaling properties of foreign exchange volatility. They decompose the variance of a time series of returns and the covariance between two time series on a scale by scale basis through the application of wavelet transforms and show that no unique scaling parameter can be adapted for these time series. In particular, there are two different scaling parameters corresponding to intraday and longer scales. In a recent paper Fan et al. (2007) propose a wavelet-based method for removing the jump-component from high-frequency returns and estimating realized variance.

We use the Maximum Overlap Discrete Wavelet Transform (MODWT) to decompose variance both in time and scale. A brief overview of the wavelet formalism which we use can be found in the Appendix. Let $\widetilde{W}_{j,t}$ be the MODWT wavelet coefficients at level j and $\widetilde{V}_{J,t}$ the scaling coefficients corresponding to the maximum level of decomposition, obtained from filtering some process $x_t, t = 1, \ldots, T$. The scale-by-scale decomposition of variance then holds:

$$\|x\|^2 = \sum_{j=1}^{J} \|\widetilde{W}_j\|^2 + \|\widetilde{V}_J\|^2. \tag{10}$$

Now let $\overline{W}_{j,t}$ represent the random process resulting from filtering an infinite sequence x_t with the MODWT wavelet filter. The difference between $\overline{W}_{j,t}$ and $\widetilde{W}_{j,t}$ is that the latter is obtained by circularly filtering a finite sequence. The wavelet variance $\nu_{j,t}$ of the process x_t for the scale of length τ_j at time t is defined as the variance of $\overline{W}_{j,t}$:

$$\nu_{j,t} = \text{Var}\{\overline{W}_{j,t}\} = \text{E}\{(\overline{W}_{j,t} - \text{E}\{\overline{W}_{j,t}\})^2\} = \text{E}\{\overline{W}_{j,t}^2\} \tag{11}$$

If wavelet variance is constant in time, its natural estimator is:

$$\widetilde{\nu}_j = \frac{1}{N} \sum_{t=0}^{N-1} \widetilde{W}_{j,t}^2 \tag{12}$$

This estimator is biased because it includes wavelet coefficients depending on the boundary conditions (i.e. the circularity assumption, used to filter a finite-length time series). Its unbiased version can be computed by ignoring such coefficients:

$$\widehat{\nu}_j = \frac{1}{N - L_j + 1} \sum_{t=L_j-1}^{N-1} \widetilde{W}_{j,t}^2 \tag{13}$$

9

with L_j the number of coefficients, affected by the circularity assumption, which clearly depends on the length of the filter. The main disadvantage of this unbiased estimator is that it does not allow the exact decomposition of the sample variance, given by (10), so we prefer to keep all the wavelet coefficients and use (12). The importance of each scale can be interpreted as the portion of total variance corresponding to each scale. It can be estimated from the decomposition of energy, defined in equation (10), by replacing wavelet variances by their estimates.

In our context we need to characterize volatility dynamically, so the estimate (12), which suggests that the wavelet variance is constant, is not quite appropriate. Dynamic estimates of $\widetilde{\nu}_{j,t}$ can be obtained by applying (12) locally for some moving window in time. This is the approach we take. The length of the window, used to estimate the wavelet variances, is chosen depending on the frequency of available data. For the data sampled at 15-minute intervals we use daily windows (32 observations) without intersections, which allows computing one value per day (analogous to the way realized volatility is computed). We call this estimate "realized wavelet variance" by analogy to the usual realized variance estimator. For daily data, we use monthly moving windows and obtain monthly realized wavelet variance. The final formula reads:

$$\widetilde{\nu}_{j,t} = \frac{1}{K} \sum_{k=t-K+1}^{t} \widetilde{W}_{j,k}^2 \tag{14}$$

with K the computation window length.

This definition is close to the one proposed in Fan et al. (2007). One difference is that we do not remove jumps from the data. Such removal is desirable when the final goal is a non-parametric estimation of the continuous diffusion parameter and its forecasting. In the perspective of a multiscale structural analysis of volatility we prefer to include the variance related to jumps in its short-term component, because it may contain important information about extreme events on the market. Another difference with Fan et al. (2007) is that we consider wavelet variances at each scale separately instead of summing them over all horizons.

The practical implementation of the MODWT algorithm requires treating the following issues: (i) the choice of wavelet filter, (ii) handling boundary conditions, and finally (iii) the choice of the number of horizons, *i.e.* levels of decomposition (see Percival and Walden, 2000, pp.135 - 145).

As for the first problem, the family of orthogonal and compactly supported wavelets is of particular interest. This family includes Daubenchies extremal phase scaling filters (also called "daublet") and least asymmetric scaling filters (also called "symlets") that are used most often. For these wavelets, full reconstruction of the signal is possible and fast decomposition algorithms are available. The filters of each type can be of different width,

usually taken from 2 to 20.

Though there are no universal criteria for the choice of the type of wavelets and their width, this choice must be dictated by the objective to balance two considerations. On the one hand, wavelet filters of too short width can introduce undesirable artifacts (unrealistic blocks) into the multi-resolution analysis. As the width of the wavelet filter increases, it can better match to the characteristics of the time series, but on the other hand, the influence of boundary conditions becomes more severe, the localization of DWT coefficients decreases and the computation becomes more difficult. Percival and Walden (2000, p. 136) suggest that the best strategy is to choose the smallest width of wavelet that brings "reasonable" results. They also point out that this strategy often results in the choice of the least asymmetric filter with width equal to 8, denoted LA(8).

This wavelet was used in particular by Gençay et al. (2001a). They report however that the use of LA(4) did not significantly alter the results. In our case reducing the filter width had no significant influence on the results of the multiscale analysis. So we decided to apply the LA(4) wavelet for the multiscale decomposition and to use it as a matter of definition for the construction of the MVS. This choice was motivated by the parsimony arguments, but mostly by the need to minimize the impact of the boundary conditions (discussed below in details), so that the results for the latest observations of volatility could be adequate. The form of the wavelet and scaling functions for this wavelet is represented on Figure 1.

The problem (ii) of handling boundary conditions arises because real time series are of finite length and we need to apply wavelet filters at their beginning and end. One possible solution is to use the so-called "circular filtering" operation, which treats the beginning and the end of the time series as if it were a portion of some periodic sequence. For financial time series the circularity assumption is hardly appropriate. Another solution which we retain in the absence of better alternatives is the "reflection" assumption: the latest observations are reflected around the boundary to obtain a time series of the necessary length. This assumption affects the results of scale by-scale decomposition for several latest observations of volatility. If a period of extremely high volatility is observed at the final date T and it started at some date $T - n$, we apply the $MODWT$ for the period $[T - n; T]$ supposing that this high volatility period will last till $T + n$ in future, thus possibly under- or overestimating the persistence of fluctuations and potentially attributing them to a wrong scale. Anyway, introducing a rule of this kind is inevitable if we need to characterize currently observed fluctuations as short or long term, since we do not know *a priori* how long they will last.

The choice of the number of scales (iii) is a difficult issue. In some applications of multi-resolution analysis, it is dictated by some physical considerations of scale. For example, to decompose a cardiogram collected at a rate 180 samples per second, a natural level of decomposition can be 6,

because is associated with the physical scale 0.36 seconds, which permits to isolate large scale fluctuations ("baseline drift") that are of little interest (Percival and Walden, 2000, p. 145). Unfortunately, no such easy rule can be established for financial time series. Apparently for the data sampled at 15-minute intervals (32 observations per day) the level of decomposition must be higher than 5 in order to isolate intraday fluctuations. Our choice of the number of horizons is a compromise between the high enough portion of energy, explained by the details, and the accuracy of the approximation, which declines as the number of scales increases for a given number of observations. Concrete explications on the choice of the maximum decomposition level for particular time series are presented in the following section.

4   A New Multi-Horizon Market Volatility Scale

The wavelet decomposition allows computing estimates of variance for a range of scales. They can be used to construct synthetic indicators by applying the transformation of the form given by equation (5). But unlike the SMS and the IMS, described in section (2), we first construct the indicator for particular scales rather than aggregate them directly.

To reduce the number of horizons we aggregate some dyadic scales in order to define three characteristic horizons, called Short (S), Medium (M) and Long (L) scales. For $s \in \{S, M, B\}$ the synthetic variance reads:

$$v_t^s = \sum_{j \in H_s} \nu_{j,t} \tag{15}$$

with $\nu(\tau_j))$ the realized wavelet variance for the scale of length τ_j defined by equation (14) and H_s the set of horizons (scales) included in the range corresponding to s. The scale $s = L$ includes the highest level of decomposition and thus represents the realized variance of the scaling coefficients, corresponding to the approximation in terms of the MODWT. The scale-specific MVS is then defined by:

$$MVS_t^s = -\log_2 P(v_t^s) \tag{16}$$

with $P(v_t^s)$ the probability to observe the variance at the synthetic scale s above v_t^s.

We do not attempt to make an overall parameter fit for the volatilities' distribution (such as normal or log-normal). Instead we focus on the tails and propose to fit a Generalized Pareto distribution (GPD) for extremely high variances. For non-extreme variances we use a simple empirical probability transform. This makes the indicator more flexible and more accurate. The GPD density function reads:

$$f(x|\theta, \xi, \sigma) = \frac{1}{\sigma} \left(1 + \xi \frac{x - \theta}{\sigma}\right)^{-1 - \frac{1}{\xi}} \tag{17}$$

For $x > \theta$ with θ the threshold parameter, $\xi \neq 0$ the shape parameter and σ the scale parameter. For positive ξ, if σ and θ are equal, the GPD is equivalent to Pareto distribution. If $\xi < 0$ then x has zero probability to exceed $-\sigma/\xi$.

Subsequent realized variances are highly correlated in time, especially for the long scales. This leads to the clustering of extreme observations and the *iid* assumption, underlying the maximum likelihood estimation of parameters, is violated. A simple declustering procedure is envisaged (see Coles, 2001, Ch. 5). The details concerning the GPD fit can be found in the following section.

Once the MVS for the three scales has been computed, the rule for detecting periods of extreme volatility must be specified. Any definition of such periods is arbitrary, so we try to propose an intuitively plausible one. For that purpose we define two thresholds, noted Θ and Φ, which will be used to detect respectively the beginning and end of events. More precisely, the event starts when one of the three MVS components overcomes Θ and ends when all three components first fall below Φ. The thresholds can be chosen depending on the desired rareness of the events. Depending on applications, we use two values for Θ: $\log_2(100)$ and $\log_2(20)$, corresponding to the tail probabilities 0.01 and 0.05 respectively. For Φ we use the value $\log_2(10)$ for tail probability 0.1.

It is useful not only to characterize volatility at different scales, but to make an overall judgment about volatility on the market, which accounts for all scales. For that purpose we need to determine the importance of each scale. The additive decomposition of wavelet variance, described by (10), suggests that the magnitude of realized variances at different scales measures the impact of these scales onto the overall variance.

The problem is that a huge part of returns' variance can be attributed to the shortest scales, but the persistence of volatility at short scales is not necessarily the same as at long scales. The persistence, or the level of long-range dependence in some time series x_t is characterized by the power-law divergence at the origin of its spectrum:

$$\Theta_x(u) \sim c_\Theta |u|^{-\alpha} \tag{18}$$

with α the scaling parameter and c_Θ a constant. The scaling parameter is related to the Hurst exponent, denoted H, by:

$$H = \frac{\alpha + 1}{2} \tag{19}$$

If the values of H for some stationary process lie in the range $(0.5, 1)$, the process is said to have long memory.

The Hurst exponent can be estimated using wavelet variances. The wavelet variance is closely related to the spectral density function of the

time series. Namely, under some regularity conditions the wavelet variance for scale τ_j is approximately the mean value of the spectral density function within the interval of frequencies $[1/2^{j+1}, 1/2^j]$:

$$\nu_j \approx 2 \int_{1/2^{j+1}}^{1/2^j} \Theta_x(u) du \tag{20}$$

This gives us the possibility to study the scaling behavior separately for different ranges of frequencies. For a pure power law with spectral density $\Theta_X(u) \propto |u|^\alpha$, the wavelet variance is related to its scaling parameter α in the following way:

$$\nu_j \propto (\tau_j)^{-\alpha-1} \tag{21}$$

These properties underlie the estimation of the scaling parameter by regressing $\log \nu_j$ on $\log \tau_j$ (see Gençay et al., 2001a).

We use the so-called logscale diagram method, first proposed in Veitch and Abry (1999), to estimate the scaling parameter. The logscale diagram is essentially a log-log plot of variance estimates of the wavelet details at each scale, against scale, accompanied by the confidence intervals about these estimates at each scale. A weighted least estimate of the slope of the logscale diagram in some range of scales (alignment region) is transformed to the Hurst exponent using equation (19). The squared errors in the wavelet variance estimates are used as the weights. The estimator is semi-parametric, because prior to the estimation, an analysis phase is required to determine the lower and the upper cutoff scales, where linear alignment on the logscale diagram is adequate. It is unbiased even for samples of finite length and efficient. The method also provides an estimate of the constant c_Θ in equation (18). A detailed discussion of the statistical properties can be found in Veitch and Abry (1999).

Using a similar method, Gençay et al. (2001b) find two alignment regions on the logscale diagram for high frequency exchange rates volatilities. The scaling properties of the data change approximately at the daily horizon, so that intraday changes have significantly lower level of persistence. We reproduce their results with the 15-minute returns on the CAC40 index and find a similar phenomenon in the daily returns on the DJIA index, analyzed in a very long sample. In the latter case the break is observed at approximately 3-weeks horizon (see the following section for details).

We propose to define the weights of scales in a way that makes them dependent on the portion of total variance, attributed to each scale, but also on the level of persistence associated with it. The latter can be measured by the Hurst exponent. Our definition is inspired by the practice of volatility "annualization", common in the financial industry: a volatility estimate, computed over a short period of time, is multiplied by the square root of the length of a reference period (usually one year). The square root corresponds

14

to the Hurst exponent, equal to 0.5 for the Brownian motion process of log-returns. In our case the Hurst exponent is estimated for two ranges of scales. We use a 3-month reference scale instead of one year to reduce the impact of the errors in the Hurst exponent estimation on the weights of scales. The definition of the scale weights reads:

$$\omega^S = \frac{A_S R^{H_1}}{A_S R^{H_1} + A_M R^{H_2} + A_L R^{H_2}}$$
$$\omega^M = \frac{A_M R^{H_2}}{A_S R^{H_1} + A_M R^{H_2} + A_L R^{H_2}}$$
$$\omega^L = \frac{A_L R^{H_2}}{A_S R^{H_1} + A_M R^{H_2} + A_L R^{H_2}}$$
(22)

with $A_s, s \in \{S, M, L\}$ the mean variance for the corresponding scales, computed over a sufficiently long window (*i.e.* a direct measure of the range of scales in the decomposition of energy), R - the time length of the reference scale (we use 3 months), and $H_{1,2}$ the scaling factors for the short and medium-long scales respectively.

The final indicator, which aggregates the three scales, is then computed in the following way:

$$MVS_t = \omega^S MVS_t^S + \omega^M MVS_t^M + \omega^L MVS_t^L = - \sum_{s \in S, M, L} \omega^s \log_2 P(v_t^s) \quad (23)$$

The resulting formula (23) resembles the definition of the IMS in equation (9), but this similarity is spurious: all terms in (23) are defined in a different way. The computation algorithm for the MVS, described above, can be summarized as follows:

- choose the number of decomposition levels and compute the wavelet and scaling coefficients from absolute logarithmic returns;

- compute realized variances of wavelet and scaling coefficients;

- estimate the Hurst exponent for different ranges of frequencies, determine if a break point is present (*i.e.* the limit between two alignment regions on the logscale diagram with different scaling parameters) and compute the weights of scales;

- aggregate the realized wavelet variances to three representative horizons (short, medium and long scale);

- fit the GPD probability distribution for the tails of realized wavelet variances at each of the three synthetic scales;

- compute the MVS for each of the three synthetic scales and the weighted MVS, representative of all scales.

This algorithm is applied in the following section to compute the MVS for the stock index data.

5 Application to the Stock Market Data

In this section we compute the MVS for real stock market data and use the indicator to detect and characterize periods of extreme volatility. Two different data sets are used: values of the French CAC40 Index from 20/03/1995 to 29/12/2006 sampled at 15-minute intervals ($100, 881$ observations) and daily values of the Dow Jones Industrial Average Index (DJIA) from 26/05/1896 to 10/10/2007 ($28, 864$ observations). The first sample is the same as in Maillet and Michel (2003, 2005), but extended till 2006. Absolute values of the logarithmic returns are used to measure volatilities.

Obviously, in these two applications different ranges of scales are considered. The choice of the number of decomposition levels is a compromise between the will to represent most of the overall variance of returns by the details (and thus to attribute this variance to specific ranges of frequency) and the constraints, imposed by the finite length of the sample. For the CAC40 data we chose the 9-level decomposition with the LA(4) wavelet, so that the approximation represents the scale 8 days and more (256 observations) and accounts for about 10% of the variance. For the DJIA sample we use an 8-level decomposition with the approximation for the scale 128 days and more which explains about 16% of the variance. The results of the variance decomposition by the MODWT according to equation (10) is shown in Tables 1 and 2. In both cases the biggest portion of variance is attributed to the shortest scale, and for longer scales it decreases rapidly.

We then compute the realized variance of the wavelet and scaling coefficients at each scale. 32 observations are used to obtain daily realized variances of the CAC40 and 20 observations for the monthly realized variances of the DJIA[2]. For the CAC40 index a problem of treating overnight returns arises. We follow the procedure, suggested in Hansen (2005), to obtain an estimate of the realized variance of the whole day, which is a weighted average of the squared overnight return and the realized variance over the trading hours with the weights, minimizing the mean square error of the estimate. Optimal weights are determined for the conventional realized variance estimation, and then applied to the realized wavelet variances (i. e. the wavelet coefficient, corresponding to the overnight return, is taken with the same weight as would be attributed to the overnight return itself when computing the conventional realized variance). Overnight returns are generally larger in absolute value than 15-minute returns, but this exceedance is filtered out by the first level wavelet detail, so the Hansen correction for the overnight returns is important only for the short-term component of volatility.

The next step in the algorithm is the study of the scaling properties of volatility over different ranges of frequencies. Figures (2) and (6) represent

---

[2]The exact number of observations may differ for weeks with less than 5 trading days.

the logscale diagrams, used to estimate the Hurst exponent, as described in the previous section. Applying the logscale diagram method to different subsamples, we always find two alignment regions. For the CAC40 data: scales 2 to 5 and 6 to 9, and for the DJIA data: scales 2 to 4 and 5 to 8. For both datasets the scaling parameter is significantly higher for the longer scales, which means higher persistence. Similar results are reported by Gençay et al. (2001b) for the exchange rate volatility. On Figure 4 a logscale diagram for simulated log-normal returns is shown for comparison (the sample is of the same length as the actual DJIA data on Figure 6). Here the alignment is almost perfect for the scales 2-8 and no significant difference between the estimates of the scaling parameter for the short and long scale regions is present.

In order to account for the possible changes in the scaling behavior of the time series we recompute the Hurst exponent on a rolling basis: for the CAC40 it is reestimated each 512 observations (approximately each 16 days) using a window of 32,512 observations (approximately 4 years), for the DJIA it is reestimated each 256 observations (approximately each year) on a window of 13,000 observations (approximately 50 years). The point estimates of the Hurst exponent and the associated confidence intervals are represented on Figures 5 and 6. For the CAC40 high frequency data, the value of the parameter for the short scales is stable around 0.6, while the average for the long scales is about 0.8. The estimate of the Hurst exponent for the long scales falls after each period of extremely high volatility (2001, 2002-2003), and also during the recent period after the vanishing of the positive trend in returns, characteristic to years 2003-2005. For the construction of the aggregated MVS this implies that the relative importance of short scales increases for the periods of very high volatility and uncertainty.

The realized wavelet variances are then aggregated. For the CAC40 data, the short-term component of volatility, which corresponds to the "intraday", is the sum of the realized daily variances of wavelet coefficients at the scales of width inferior to 1 day (details 1-5), the medium term - at the scales of width between 1 and 16 days (details 6-9) and the long-term is the realized daily variance of the scaling coefficients (approximation at scale 9). For the DJIA data the short term aggregates the variances at the scales of less than 8 days width (details 1-4), the medium term is the range from 8 days to 128 days (details from 6 to 8) and the long term is the realized monthly variance of scaling coefficients (approximation at scale 8).

The aggregated realized variances for the three synthetic scales, defined above, are present on Figures 7 and 8 along with the traditional realized variance estimates. For both indexes, the profile of the short scale wavelet variance resembles almost perfectly the profile of the overall realized variance, though the magnitude of peaks is different. The peaks in the medium and long scale components generally coincide with the peaks in the realized variance, but many short-term fluctuations are filtered out completely. For

17

the CAC40 data we observe extreme long-term volatility in 1998, 2000, 2001 and 2002. Note that during the boom of 2003-2006, when the overall realized variance was relatively low, it was largely attributed to the long scale. The positive trend in returns translates into a gradual increase of the long-term variance, unlike the sharp and abrupt splashes of volatility observed during the stock market crises. For the DJIA data the most important period of extreme long-term variance is recorded during the Great Depression of 1929-1934. It largely surpasses in amplitude and duration the latest crashes in 1987 and 2001.

We report the results of the GPD fit for the CAC40 index in Figure 9 and Table 3 and for the DJIA index in Figure 10 and Table 4. The parameters are estimated for the overall realized variance and for the three scales, based on the whole sample and on two subsamples of equal length. The threshold parameter θ in the probability density function (17) equals to the 0.9 empirical quantile for the CAC40 and 0.8 empirical quantile for the DJIA. This choice allows both good quality of alignment and sufficient number of observations (after declustering procedure, described below) for the estimation. The study of the mean excess plot, commonly used to determine the threshold, results in roughly the same values for the threshold. The estimates of parameters σ and ξ are obtained by the maximum likelihood method. The associated confidence intervals are computed using the profile log-likelihood functions, as suggested in Gilli and Këllezi (2006). This method, based on reparameterizing the GPD distribution as a function of the unknown quantile, allows easy interval estimation of the quantile as well as of the parameters. Conventional bootstrap interval estimates for σ and ξ, not reported here, yield similar results.

Maximum likelihood estimation is based on the *iid* assumption for the data, which is obviously violated for the realized wavelet variances, especially at the medium and long scales. To overcome this problem we use a declustering procedure (see Coles, 2001). The cluster starts when the variance exceeds the threshold θ and ends when r consecutive observations fall below this threshold. We fix the value of r to be 20 daily observations of the realized wavelet variance for the high frequency CAC40 data and one monthly observation for the low frequency DJIA data. We then pick the maxima within each cluster and use them instead of raw data for the parameter estimation.

Note that the quality of the GPD fit is generally better for the variances at different scales than for the overall realized variances, especially for the quantiles above 0.97 (see the Q-Q probability plots on Figures 9 and 10). The probability law is relatively close to Pareto distribution for all scales and both subsamples (for the Pareto law we must have $\xi > 0$ and $\theta = \sigma$). For the CAC40 data, the Pareto conditions on parameters are accurately verified for the second subsample, while for the first subsample the point estimate of the scale parameter σ is considerably above the value of θ, though its

95% confidence interval estimate still contains θ. The same conslucions are globally true for the DJIA, except for the long term component, for which the estimate of σ is too low. But in the latter case the heterogeneity between two samples (prior and after year 1952) is much more important, which is obviously due to the evolution of the size and liquidity of the market since the beginning of the XXth century. Notably, the 0.8 empirical quantile for the realized monthly variance fell almost twice from 4.2×10^{-7} to 2.2×10^{-8}, which is mainly attributed to short and medium term. To account for evolution, we further estimate the parameters of the GPD distribution for the DJIA data using 50-year rolling windows, thus the recent values of the MVS are unaffected by the events before the mid 1950s.

Our study of the probability distribution evidences in favor of the analogy between volatility and earthquakes: probability of extreme volatility clusters and high seismic activity clusters can both be described by the Pareto law. Thus a tool similar to the Richter-Gutenberg scale in geophysics, can be adequate to measure volatility in finance.

Finally we proceed with the computation of the MVS for each scale and its aggregation. Figure 6 shows the values of the CAC40 index, the realized variance and the aggregated MVS, computed as a weighted average over the three scale components, described above. The areas colored in gray correspond to the periods of extremely high volatility (probability less than 0.01), as detected by the MVS. The quantitative information about these periods is reported in Table 5. The beginning of each period is the date when the MVS for one of the scales overcomes the threshold $\log_2(100) = 6.6439$, which corresponds to 0.01 tail probability. The end of such period is the date when the MVS for all scales falls below the threshold $\log_2(10) = 3.3219$, which corresponds to 0.1 tail probability.

We report the duration of the periods, determined as the number of trading days between the first and the last date, the maximum value of the MVS over the period (for the three scale and for the weighted average[3]) and the sum of the MVS values for the days, when one of the components was above the critical threshold of $\log_2(100)$, taken over the period of extreme volatility (for the three scales and for the weighted average). Each of the measures can be chosen to asses the gravity of the situation on the stock market. Another important characteristics is the scale of the MVS component, which first broke the critical threshold.

We detect six periods, corresponding to our criteria of extreme volatility. The period associated to the Asian crisis (from 27/8/98 to 14/10/98) was the longest in duration (34 open days) and in terms of the sum of the weighted MVS values above the threshold. The crisis is found to be the most "long-term" one, as evidenced by the values of the third component

---

[3]Note that the column reporting the maximum of the weighted MVS is not the mean of maxima at each scale, because the latter occur at different times.

of the MVS. The beginning of the crisis was signaled by the medium-term component of the indicator. The period 19/7/02-9/8/02 is also characterized by persistently high values of the long-term volatility component, which was the first to break the critical threshold. The weighted average MVS is close to, but never breaks the threshold during these crises, which means that the peaks of volatility at different scales did not match in time.

The crisis following the September 11, 2001 events (11/9/01 - 1/10/01) ranks third in terms of the sum of the weighted MVS values, but it is structurally different from the two previously mentioned crises. The volatility is higher at the short and medium scales, clearly because the beginning of the crisis was sudden and driven by the factors, exogenous to the financial market. The shock was strong enough to affect the long-term component, though it does not overcome the critical value. The crisis of 2001 has the highest value of the weighted MVS over the whole sample (7.8, or tail probability about 0.005). Two periods of high volatility in year 2000 are characterized by high volatility on the short and medium scale. The maximum aggregated MVS for the period 14/4/00 - 3/5/00, which marks the beginning of a long downturn in the market lasting till 2003 and is associated to the burst of the so-called "new economy" bubble, is 7.4 and ranks second after the 2001 crisis. Finally, fluctuations in May 2006 are labeled as extreme at the medium scale only. The weighted MVS stays below the critical threshold.

The MVS presented so far is suitable for the *a posteriori* structural analysis of market events, but it cannot be used for the real time applications, because the MVS values, as defined above, are not adapted. At each date we use all available data to perform the scale-by-scale wavelet transform and estimate the probabilities. Alternatively, we construct an adapted indicator for the CAC40 index. The part of the sample before 2001 (6 years) is used for initial estimates, which are updated dynamically as the new data arrives. So only the information prior to each computation date is used. At each iteration the last computed value of the MVS is retained. As we noted before, this value depends on the assumptions used for the boundary conditions. Here we use the reflection assumption, which implies that if at the date of analysis we observe extreme volatility over some period of time in the past, we expect that it will last exactly the same period of time in future. This assumption affects the scale-by-scale decomposition and is the cost to pay for the non-anticipating property of the MVS.

The results, obtained for the CAC40 index with the adapted MVS, are reported in Figure 12 and Table 6. Clearly, using a smaller sample to estimate probabilities produces more periods of extreme volatility. We detect three periods in 2002 instead of one and a new 3-week period in 2003. But structurally the results are the same. The 2001 crash is still the biggest in terms of the average MVS and concentrates mainly on the medium and short scales (but it also surpasses the 2002 crisis by the sum of the MVS values). As before, the crisis in 2002 is characterized by the high long-term

volatility: in July and October 2002 the long-term components of the MVS mark the beginning of the extreme volatility periods. Generally, the conclusions made with the adapted indicator are consistent with those presented before.

The results of similar analysis for the DJIA index are reported in Figure 6 and Table 7. In this case the definitions of the short, medium and long term are not the same - all the range of scales is shifted toward lower frequency and the realized variance is computed over monthly intervals. It is less variable than the daily realized variance, and critical levels of the MVS are chosen lower than for the high frequency data: $\log_2(20)$ for the beginning and $\log_2(10)$ for the end of the extreme volatility periods, which corresponds to 0.05 and 0.10 tail probabilities respectively.

As explained above, we estimate the distribution of variances over 50-year rolling windows (so the estimates till year 1944 are not adapted and further estimates are adapted). This convention allows to account for the changes in the structure of the financial market (and consequently, in the distribution of volatilities), but as a side-effect it produces an artificial shift in the 1990s - the period when the Great Depression falls out of the sample. Thus, the market volatility in late 1973 - early 1974, related to the oil crisis, Vietnam war and the Watergate scandal in the USA, followed by the resignation of president Nixon, is clearly higher than normal, as can be seen on Figure 6. But it does not reach the critical threshold, defined relatively to the extreme events of the 1930-s. Compared to the more recent events (say, from 1950 to 2000), this period would probably be detected as a period of extreme volatility.

This example demonstrates that the MVS is a relative indicator, since the underlying probability distribution is estimated from the historical data. Thus particular attention should be paid to the choice of the sample length and threshold values, especially when the sample is heterogeneous as in the DJIA case. However, this does not undermine the value of the indicator: it allows for the scale-by-scale structural analysis and comparison of events with the past, even though its absolute values and thresholds depend on arbitrary choices.

The results in Table 7 confirm that the Great Depression had an unprecedentedly strong affect on the stock market. The period of extreme volatility lasts for 54 months from the crash in 1929 till early 1934, when the recovery starts. The volatility shock associated to the Black Thursday (October 24, 1929) was so strong that all the three components of the MVS simultaneously break the critical threshold. The market downturn during the quick recovery (1937-1938), following the Great Depression, was first detected in the medium-term component of the MVS. The subsequent events, including those related to the beginning of the World War II, had a smaller impact on the volatility of the American stock market. Surprisingly, the panic provoked by the bankruptcy of the Penn Central Road in May 1970

had a stronger impact on the mid-term component of the MVS, than the following recession of 1973-1974, though, as it can be seen on Figure 6, the latter period is characterized by a persistently higher-than-average level of volatility, measured by the aggregated MVS.

The end of year 1982 was marked by the change in the trend of the index returns (from the bear market in 1967 - 1982 to the bull market in 1982-2000), accompanied by a period of extremely high short-term volatility. The following sharp rise in stock prices is interrupted by the crash of October 1987 (known by the Black Monday of October 19, 1987, though the first significant fall of 3.81% occured 5 days before). The weighted MVS in October 1987 has even higher value than the maximum recorded during the Great Depression, though the overall impact during the subsequent yearly period, measured by the sum of MVS extremes, is considerably lower.

The beginning of the Asian crisis in 1997 and its bottom in 1998, coupled with the Russian crisis, did not drive the long-term MVS component above the critical threshold. The medium and short-term volatility, however, was very high, resulting in extreme values of the weighted MVS. A rather long period from 2000 to 2003 is characterized by the burst of the "new economy" bubble. The exogenous shock of September 11, 2001 aggravates the situation, already featured by very high volatility. At the bottom of this crisis (late 2002 - early 2003) we detect the second period of extreme volatility since 1987, that affects the long-term component of the MVS in a significant way.

6   Conclusion

We design a new indicator of the stock market volatility, called the Market Volatility Scale, which characterizes the fluctuations both in time and frequency domains. Unlike the other multiscale indicators, proposed in the existing literature, it is based on the wavelet variances of returns, which are closely related to the spectrum of the time series of returns and allow an accurate scale-by-scale decomposition of the variance. We introduce a time-varying wavelet variance estimate, analogous to the conventional realized variance, but measuring variation in a particular range of frequencies.

The realized variances of the wavelet and scaling coefficients are aggregated to three representative scales (short term, medium term and long term). The MVS is computed at each date as the logarithm of the probability to observe the realized wavelet variance higher than its value at this date. This approach is inspired by the Richter-Gutenberg scale, used to measure the magnitude of earthquakes. Our scale is thus universal and easy to interpret, which is helpful for the comparison of events. When a logarithm to the base 2 is used for the definition, a one-point increase in the MVS corresponds to a volatility observation twice as unlikely.

We model the probability distribution of volatilities in a flexible way, using a generalized Pareto distribution fit for the clusters of extreme volatility. So we focus on the right tail, applying an empirical probability integral transform for the rest of the probability distribution. This choice is motivated by the need to be as accurate as possible in characterizing the rareness of extreme events.

The MVS is first computed separately for each scale, and then the results are aggregated to provide a general judgment on the market volatility at all scales. The aggregated MVS takes the form of the weighted average over the three scales. The weights are dependent on the scaling properties of the time series of returns, which are studied using the logscale diagrams. We show that the level of the long range dependence in volatility is not the same for different ranges of frequencies, supporting earlier findings reported for the high-frequency exchange rate volatility. Using the weights, depending on the Hurst exponent, we measure the importance of scales not only by the portion of the variance, attributed to them (clearly, most variance is due to the short scale), but also by the level of persistence, which is higher for the medium and long scales.

The MVS is computed for the daily Dow Jones Industrial Averages index data from 1896 to 2007 and for the intraday CAC40 data from 1995 to 2006. It is used for the detection and scale-by-scale analysis of the periods of extreme volatility. In particular, we report which scales of fluctuations (short, medium or long-term) were most important during each period, and at which scale the critical threshold, used to determine the beginning of the period, was first overcome. The events can be ranked by the maximum value of the MVS at each scale and by the sum of the MVS values over the whole period of extreme volatility.

Clearly, the potential scope of application of the indicator and of the underlying variance decomposition goes far beyond the comparison of events. We mention several areas of further research which are of particular interest. From the theoretical viewpoint, it is important to understand the mechanics of the multiscale price dynamics. Wavelet variances can be an appropriate tool in testing market microstructure models of this type, both on simulation and real-data basis.

From a more practical perspective, the multiscale view of volatility can be suggestive in asset allocation and portfolio management. A multiscale indicator can be applied for conditioning the asset allocation process in a regime-switching framework. Another important area of application is volatility forecasting. In this paper we focused on structural rather than dynamic properties of volatility. But the interdependence of scales and the additivity property of wavelet variances make appealing the idea of the scale-by-scale autoregression, which can be used to forecast the components of volatility and then aggregate them.

References

T. Andersen and T. Bollerslev. Heterogeneous information arrivals and return volatility dynamics: Uncovering the long run in high frequency data. *Journal of Finance*, 52:975–1005, 1997.

A. Arneodo, J. Muzy, and D. Sornette. Casual cascade in stock market from the 'infrared' to the 'ultraviolet'. *European Physical Journal B*, 2: 277–282, 1998.

W. Breymann, S. Ghashghaie, and P. Talkner. A stochastic cascade model for fx dynamics. *International Journal of Theoretical and Applied Finance*, 3:357–360, 2000.

E. Capobianco. Multiscale analysis of stock index return volatility. *Computational Economics*, 23(3):219–237, 2004.

K. Christensen, L. Danon, T. Scanlon, and P. Bak. Unified scaling law for earthquakes. *Physical Review Letters*, 88(17):178501–504, 2002.

S. Coles. *An Introduction to Statistical Modeling of Extreme Values*. Springer-Verlag, London, 2001.

R. Cont. Empirical properties of asset returns: Stylized facts and statistical issues. *Quantitative Finance*, 1(2):223–236, 2001.

F. Corsi. A simple long memory model of realized volatility. *Working Paper, University of Southern Switzerland*, 2004.

M. Dacorogna, U. Muller, R. Dave, R. Olsen, and O. Pictet. Modelling short-term volatility with garch and harch models. In Dunis-Zhou, editor, *Nonlinear Modelling of High Frequency Financial Time Series*, pages 161–176. John Wiley, 1998.

Z. Ding, C. Granger, and R. Engle. A long memory property of stock market returns and a new model. *Journal of Empirical Finance*, 1(1): 83–106, 1993.

J. Fan, Y. Fan, and J. Jiang. Multi-scale jump and volatility analysis for high-frequency financial data. *Journal of American Statistical Association*, 102:618–631, 2007.

V. Fernandez and M. Lucey. Portfolio management under sudden changes in volatilityand heterogeneous investment horizons. *Physica A: Statistical Mechanics and its Applications*, 375(2):612–624, 2007.

R. Gençay, F. Selçuk, and B. Whitcher. An introduction to wavelets and other filtering methods in finance and economics. 2001a.

R. Gençay, F. Selçuk, and B. Whitcher. Scaling properties of foreign exchange volatility. *Physica A*, pages 249–266, 2001b.

S. Ghashghaie, W. Breymann, J. Peinke, P. Talkner, and Y. Dodge. Turbulent cascades in foreign exchange markets. *Nature*, 381:767–770, 1996.

M. Gilli and E. Këllezi. An application of extreme value theory for measuring financial risk. *Computational Economics*, 27(2-3):207–228, 2006.

P. Hansen. A realized variance for the whole day based on intermittent high-frequency data. *Journal of Financial Econometrics*, 3(4):525–554, 2005.

H. Lee. International transmission of stock market movements: A wavelet analysis. *Applied Economics Letters*, 11:197–201, 2004.

M. Loretan. Generating market risk scenarios using principal component analysis. In *The Measurement of Aggregate Market Risk*, number 7. Publications of the Committee on the Global Financial System (CGFS), Bank for International Settlements, 1997.

T. Lux. Long-term stochastic dependence in financial prices: Evidence from the german stock market. *Applied Economics Letters*, 3:701–706, 1996.

B. Maillet and T. Michel. An index of market shocks based on multiscale analysis. *Quantitative Finance*, 3(2):88–97, 2003.

B. Maillet and T. Michel. The impact of the 9/11 events on the american and french stock markets. *Review of International Economics*, 13(3):597–611, 2005.

M. Mega, P. Allegrini, P. Grigolini, V. Latora, L. Palatella, A. Rapisarda, and S. Vinciguerra. Power law time distributions of large earthquakes. *Physical Review Letters*, 90:188501, 2003.

U. Müller, M. Dacorogna, R. Dave, R. Olsen, O. Pictet, and J. Von Weizsacker. Volatilities of different time resolutions - analyzing the dynamics of market components. *Journal of Empirical Finance*, 4:213–239, 1997.

D. Percival and A. Walden. Wavelet methods for time series analysis. 2000.

J. Ramsey and Z. Zhang. The analysis of foreign exchange data using waveform dictionaries. *Working paper, C.V. Starr Center for Applied Economics, Department of Economics, New York University, NYC, USA*, 1995.

C. Richter. *Elementary Seismology*. Freeman, San Francisco, 1958.

D. Veitch and P. Abry. A wavelet based joint estimator of the parameters of long-range dependence. *IEEE Transactions on Information Theory*, 45 (3):878–897, 1999.

G. Zumbach, M. Dacorogna, J. Olsen, and R. Olsen. Measuring shocks in financial markets. *International Journal of Theoretical and Applied Finance*, 3(3):347–355, 2000.

Appendix: Wavelet Transforms

Any wavelet decomposition is based on the so-called wavelet filter and the associated scaling filter. A real-valued Discrete Wavelet Transform (DWT) filter $h_l : l = 0, \ldots, L-1$ of width L is an infinite sequence such as $h_l = 0$ for $l < 0$ and $l > L$, which sums to zero, has unit energy and is orthogonal to its even shifts:

$$\sum_{l=0}^{L-1} h_l = 0, \sum_{l=0}^{L-1} h_l^2 = 1, \sum_{l=0}^{L-1} h_l h_{l+2n} = 0 \qquad \text{(A-1)}$$

The "quadrature mirror" filter corresponding to h_l is called the scaling filter and is given by:

$$g_l = (-1)^{l+1} h_{L-1-l}. \qquad \text{(A-2)}$$

A Maximum Overlap Discrete Wavelet Transform (MODWT) wavelet filter and scaling filter are rescaled versions of the corresponding DWT filters defined by $\widetilde{h}_l = h_l/\sqrt{2}$ and $\widetilde{g}_l = g_l/\sqrt{2}$ respectively. Note that equation A-2 also holds for the MODWT filters. The usage of the MODWT filters instead of the DWT aims mainly to define a transform which preserves the ability to carry out multi-resolution analysis and variance decomposition, but does not depend ont the choice of a starting point for a time series, which is the case when the outputs from the recursive application of the DWT filters are downsampled in the so-called "pyramid" algorithm. In the MODWT algorithm wavelet coefficients $\widetilde{W}_{1,t}$ and scaling coefficients $\widetilde{V}_{1,t}$ are obtained by filtering the circular shift of the vector X_t rather than X_t itself:

$$\widetilde{W}_{1,t} = \sum_{l=0}^{L-1} \widetilde{h}_l x_{t-l \bmod N}, \quad \widetilde{V}_{1,t} = \sum_{l=0}^{L-1} \widetilde{g}_l x_{t-l \bmod N} \qquad \text{(A-3)}$$

The first stage of the MODWT pyramid algorithm consists in finding solution to the equation:

$$\begin{bmatrix} \widetilde{W}_1 \\ \widetilde{V}_1 \end{bmatrix} = \begin{bmatrix} \widetilde{B}_1 \\ \widetilde{A}_1 \end{bmatrix} x, \qquad \text{(A-4)}$$

with $\widetilde{W}_1$ and $\widetilde{B}_1$ defined according to (A-3). The first level maximum overlap detail is given by $\widetilde{D}_1 = \widetilde{B}_1^T \widetilde{W}_1$ and the first level approximation is $\widetilde{S}_1 = \widetilde{A}_1^T \widetilde{V}_1$, so the signal admits the following additive (first-level) decomposition:

$$x = \widetilde{B}_1^T \widetilde{W}_1 + \widetilde{A}_1^T \widetilde{V}_1 = \widetilde{D}_1 + \widetilde{S}_1 \qquad \text{(A-5)}$$

The energy of the signal, in its turn, can also be decomposed:

$$\|x\|^2 = \|\widetilde{W}_1\|^2 + \|\widetilde{V}_1\|^2. \qquad \text{(A-6)}$$

27

The wavelet and scaling MODWT coefficients for higher levels of decomposition $j = 2, \ldots, J$ are obtained by applying recursive formulas of the pyramid algorithm without downsampling of the outputs:

$$\widetilde{W}_{j,t} = \sum_{l=0}^{L-1} \tilde{h}_l \widetilde{V}_{j-1,t-2^{j-1}l \bmod N}, t = 0, \ldots, N-1; \qquad \text{(A-7)}$$

$$\widetilde{V}_{j,t} = \sum_{l=0}^{L-1} \tilde{g}_l \widetilde{V}_{j-1,t-2^{j-1}l \bmod N}, t = 0, \ldots, N-1. \qquad \text{(A-8)}$$

The j-th level of the pyramid algorithm is defined by:

$$\begin{bmatrix} \widetilde{W}_j \\ \widetilde{V}_J \end{bmatrix} = \begin{bmatrix} \widetilde{B}_j \\ \widetilde{A}_j \end{bmatrix} \widetilde{V}_{j-1}, \qquad \text{(A-9)}$$

Finally, the j-th level detail $\widetilde{D}_j$ and the J-th level smooth $\widetilde{S}_j$ are computed:

$$\widetilde{D}_j = \widetilde{A}_1^T \cdots \widetilde{A}_{j-1}^T \widetilde{B}_j^T \widetilde{W}_j, \quad \widetilde{S}_j = \widetilde{A}_1^T \cdots \widetilde{A}_{J-1}^T \widetilde{A}_J^T \widetilde{V}_J. \qquad \text{(A-10)}$$

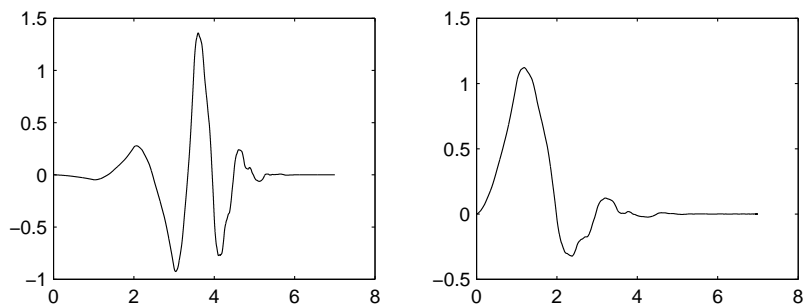The multi-resolution analysis of signal X is thus its scale-by-scale additive decomposition:

$$x = \sum_{j=1}^{J} \widetilde{D}_j + \widetilde{S}_J \qquad \text{(A-11)}$$

The scale-by-scale decomposition of energy reads:

$$\|x\|^2 = \sum_{j=1}^{J} \|\widetilde{W}_j\|^2 + \|\widetilde{V}_J\|^2. \qquad \text{(A-12)}$$

Figures and Tables

Figure 1: The form of the LA(4) Wavelet

The figure represents the form of the wavelet function ϕ (left) and scaling function ψ (right) for the wavelet filter LA(4)

Table 1: The MODWT variance decomposition for the CAC40 returns

| Level	D1	D2	D3	D4	D5	D6	D7	D8	D9	A
τ_j	1	2	4	8	16	32	64	128	256	> 256
$\widetilde{\nu}(\tau_j) \times 10^5$	0.17	0.12	0.07	0.04	0.03	0.01	0.01	0.00	0.00	0.05
$\widetilde{\nu}(\tau_j / \sum_j \widetilde{\nu}(\tau_j)$	0.39	0.26	0.15	0.09	0.06	0.02	0.01	0.01	0.01	0.10

Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. The MODWT decomposition of variance with 9 levels is used (9 details and an approximation). τ_j is the time length of each scale (one unit of time corresponds to 15 minutes). The variance of the wavelet coefficients at the first level (detail D1) corresponds to the variance at physical scale 15 minutes. The variance at level $j > 1$ represents the range from $15 \times 2^{j-2}$ to $15 \times 2^{j-1}$ minutes. The variance of the scaling coefficients (A) at level 9 represents the range of scales superior to $15 * 2^8 = 3840$ minutes, which is roughly 8 trading days. $\widetilde{\nu}(\tau_j) \times 10^5$ is the biased estimate of wavelet variance according to formula (12). The last line of the table contains the portion of the total variance, attributed to each scale.

Table 2: The MODWT variance decomposition for the DJIA returns

Level	D1	D2	D3	D4	D5	D6	D7	D8	A
τ_j	1	2	4	8	16	32	64	128	> 128
$\widetilde{\nu}(\tau_j) \times 10^5$	2.52	1.27	0.66	0.44	0.32	0.24	0.21	0.14	1.12
$\widetilde{\nu}(\tau_j / \sum_j \widetilde{\nu}(\tau_j)$	0.36	0.18	0.10	0.06	0.05	0.03	0.03	0.02	0.16

Source: Dow Jones Indexes, daily values of the DJIA index from 26/05/1896 to 10/10/2007. The MODWT decomposition of variance with 8 levels is used (8 details and approximation). τ_j is the time length of each scale (one unit of time corresponds to one trading day). The variance of the wavelet coefficients at the first level (detail D1) corresponds to the variance at physical scale one day. The variance at level $j > 1$ represents the range from $\times 2^{j-2}$ to 2^{j-1} trading days. The variance of the scaling coefficients (A) at level 9 represents the range of scales superior to $2^7 = 128$ days. $\widetilde{\nu}(\tau_j) \times 10^5$ is the biased estimate of wavelet variance according to formula (12). The last line of the table contains the portion of the total variance, attributed to each scale.

Figure 2: The Hurst exponent for the CAC40 index volatility for two sub-samples
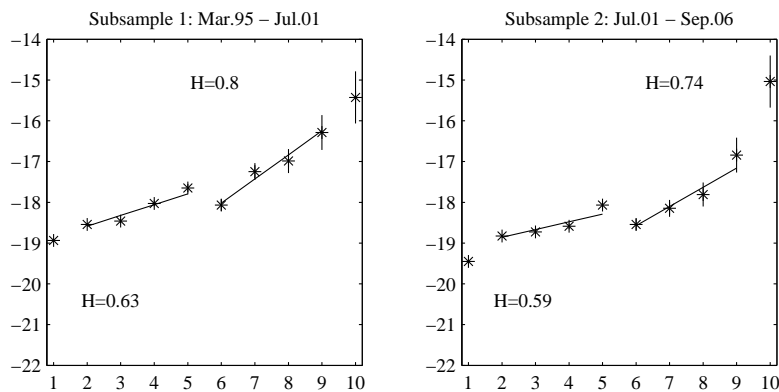
Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. The scaling factor is estimated by the Veitch-Abry logscale diagram method. The logscale diagram is generated for two subsamples of equal length (50,440 observations each) and examined to find a lower cutoff and upper cutoff scales, where alignment (a straight line) is observed. Two alignment regions are found: scales from 2 to 5 and scales from 6 to 9. The estimate of the scaling factor is obtained from the estimate of the slope of the of the WLS regression, as described in the text. Computation by the author.

Figure 3: The Hurst exponent for the DJIA index volatility for two subsamples
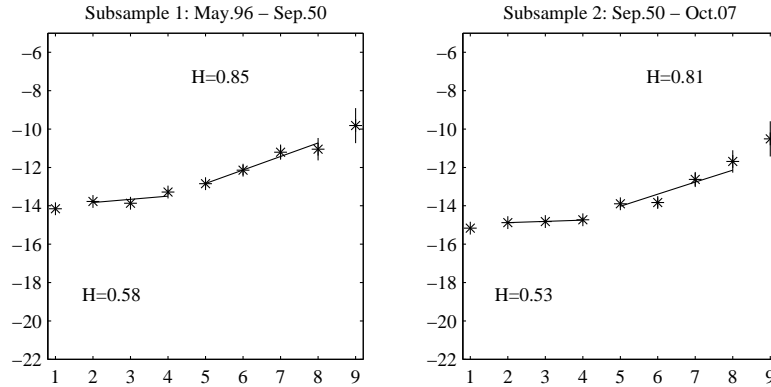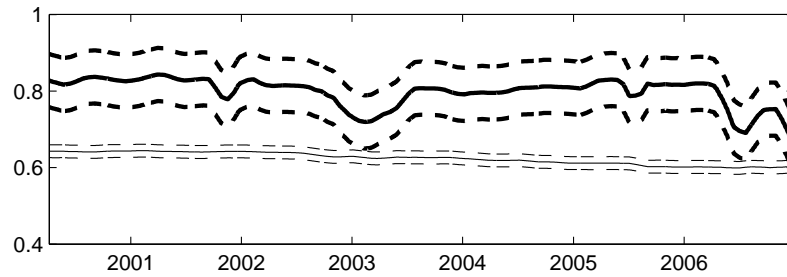

Source: Dow Jones Indexes, daily values of the DJIA index from 26/05/1896 to 10/10/2007. The scaling factor is estimated by the Veitch-Abry logscale diagram method. The logscale diagram is generated for two subsamples of equal length (14,432 observations each) and examined to find a lower cutoff and upper cutoff scales, where alignment (a straight line) is observed. Two alignment regions are found: scales from 2 to 4 and scales from 5 to 8. The estimate of the scaling factor is obtained from the estimate of the slope of the of the WLS regression, as described in the text. Computation by the author.

Figure 4: Estimation of the scaling parameter for the simulated Gaussian data
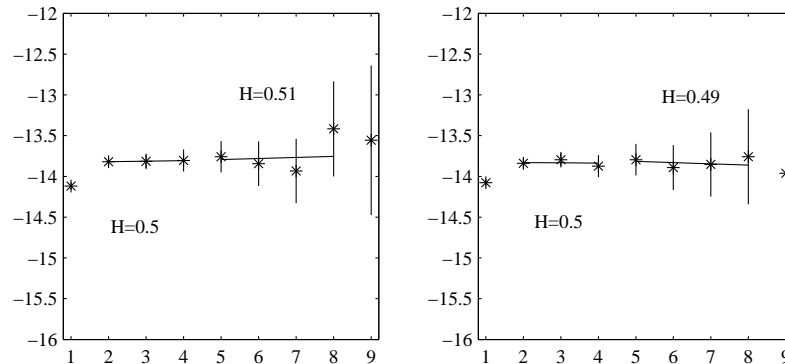
Logscale diagram for two simulated samples of Gaussian data of the sames size (14,432 observations) and of the same mean and variance as the data used for the DJIA index on figure (6). The scaling factor is then estimated by the Veitch-Abry logscale diagram method. The same two alignment regions as on figure (6) are explored. No significant change in slope of the WLS regression line is observed.

Figure 5: Dynamic estimation of the Hurst exponent for the CAC40 index volatility
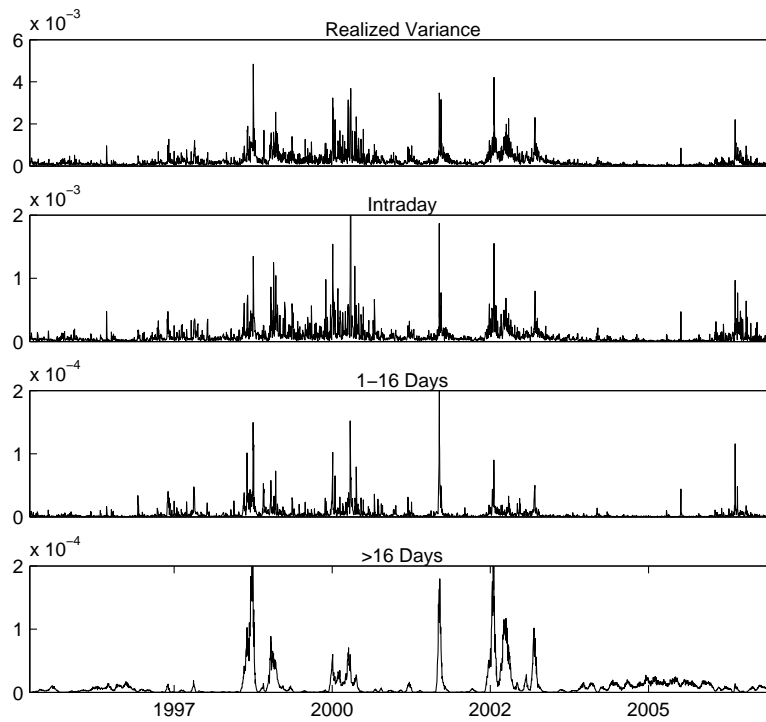

Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. The scaling factor is reestimated each 512 observations (approximately each 16 days) on a window of 32,512 observations (approximately 4 years) length. The estimation is made by the Veitch-Abry logscale diagram method, described in the text. Computation by the author.

Figure 6: Dynamic estimation of the Hurst exponent for the DJIA index volatility
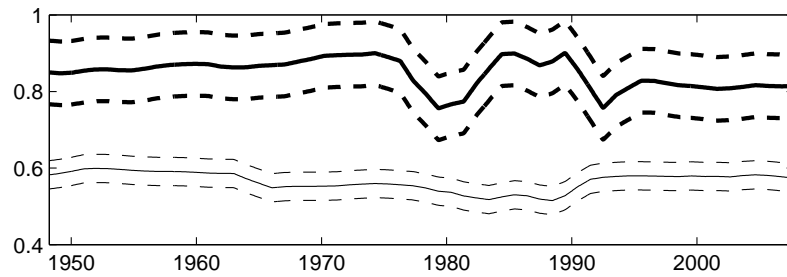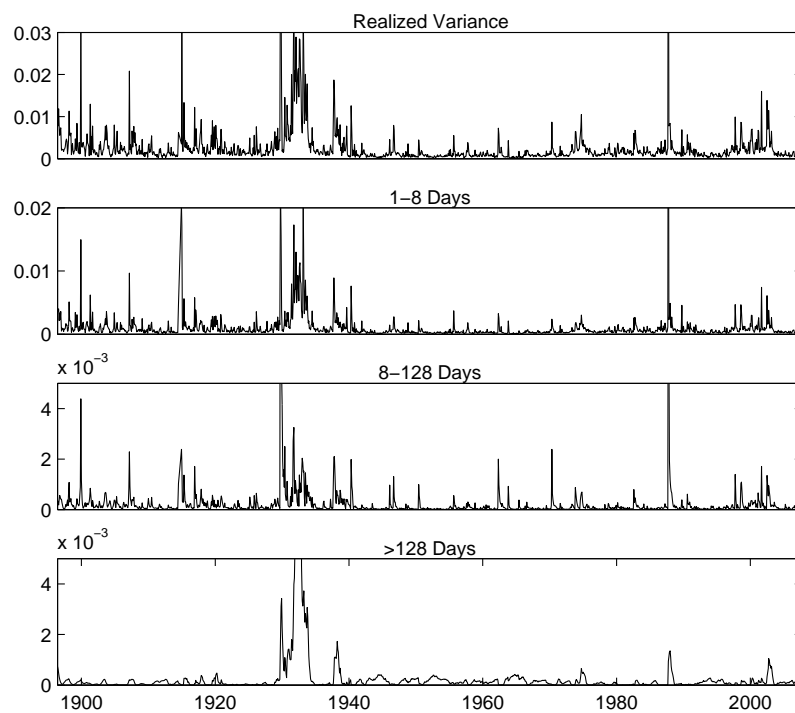

Source: Dow Jones Indexes, daily values of the DJIA index from 26/05/1896 to 10/10/2007. The scaling factor is reestimated each 256 observations (approximately each year) on a window of 13000 observations (approximately 50 years) length. The estimation is made by the Veitch-Abry logscale diagram method, described in the text. Computation by the author.

32

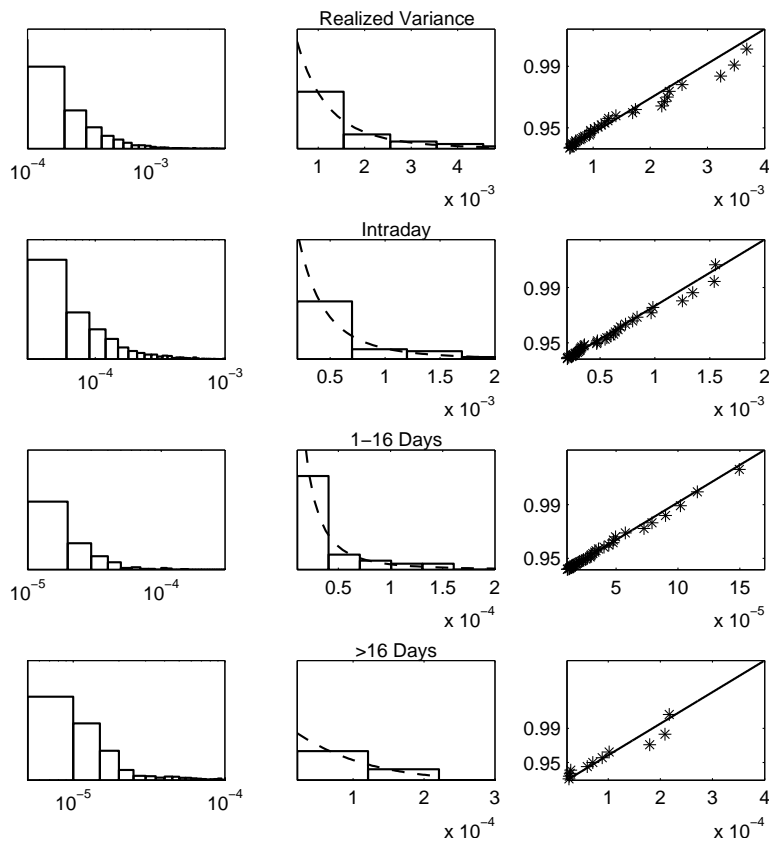Figure 7: Realized wavelet variance of the CAC40 returns

Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. Realized daily variance includes overnight returns according to Hansen (2005) method. "Intraday" realized daily variance is the sum of the realized wavelet variances for the scales of width inferior to 1 day (details 1-5), "1-16 Days" realized daily variance is the sum of the realized wavelet variances for the scales of width between 1 and 16 days (details 6-9), "> 16 Days" is the realized daily variance of the scaling coefficients (approximation at scale 9). The LA(4) filter is used for the wavelet decomposition. Computation by the author.

Figure 8: Realized wavelet variance of the DJIA returns

Figure 9: GPD fit for the realized wavelet variance of the CAC40 returns

Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. The figure represents the results of the GPD fit for the realized daily variance of the index and realized daily variances for different scale ranges, obtained by aggregation of realized wavelet variances as described in the text. Each line of figures shows (from left to right): the histogram of all data in logarithmic scale (2,954 observations), the fit of the GPD density for the upper tail of the distribution and the probability QQ-plot. Maximum likelihood estimation. The threshold parameter is chosen as the 0.9 empirical quantile of each distribution. Computation by the author.

Table 3: Parameters of the GPD fit for the realized wavelet variance of the CAC40 returns

	$\Theta \times 10^5$	ξ	$\sigma \times 10^5$	$Q_{99} \times 10^5$
All Sample: 03/1995 - 12/2006				
IV	33.8372	0.5236	47.7984	380.6894
		[0.2292; 1.0365]	[17.6859; 81.2570]	[380.6890; 381.0697]
Intraday	11.7858	0.501	20.8789	157.0575
		[0.2371; 0.9427]	[7.7256; 35.4937]	[157.0572; 157.2143]
$1 - 32$ days	0.5389	0.6782	1.2298	12.5549
		[0.3862; 1.1530]	[0.4554; 2.0903]	[12.5545; 12.5671]
> 32 days	1.3778	1.2792	0.8117	30.0373
		[0.6486; 2.5726]	[0.3008; 2.0290]	[30.0368; 30.0669]
Subample: 03/1995 - 06/2001				
IV	37.7046	0.2088	76.643	356.7373
		[-0.0447; 0.7812]	[28.3583; 130.2927]	[356.7370; 357.0937]
Intraday	13.4835	0.4493	24.2599	166.9273
		[0.1445; 1.0630]	[8.9766; 41.2415]	[166.9270; 167.0939]
$1 - 32$ days	0.7395	0.5163	1.4482	11.1065
		[0.1929; 1.1439]	[0.5363; 2.4614]	[11.1061; 11.1173]
> 32 days	0.9093	1.5068	0.4341	26.9212
		[0.6417; 3.7712]	[0.1169; 1.7797]	[26.9207; 26.9477]
Subample: 06/2001 - 12/2006				
IV	30.3373	0.9677	19.7181	379.8839
		[0.4517; 2.0316]	[7.2961; 49.2947]	[379.8835; 380.2634]
Intraday	10.1567	0.7716	11.9121	150.4925
		[0.3270; 1.6935]	[4.4079; 29.7800]	[150.4922; 150.6427]
$1 - 32$ days	0.4029	0.9508	0.8557	15.0357
		[0.4785; 1.8779]	[0.3169; 2.1388]	[15.0353; 15.0503]
> 32 days	1.5412	1.3049	0.492	19.9626
		[0.6061; 2.9390]	[0.1823; 2.0168]	[19.9623; 19.9823]

Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. The fit is obtained by the maximum likelihood method. The threshold parameter is chosen as the 0.9 empirical quantile of each distribution. Computation by the author.

Figure 10: GPD fit for the realized wavelet variance of the DJIA returns
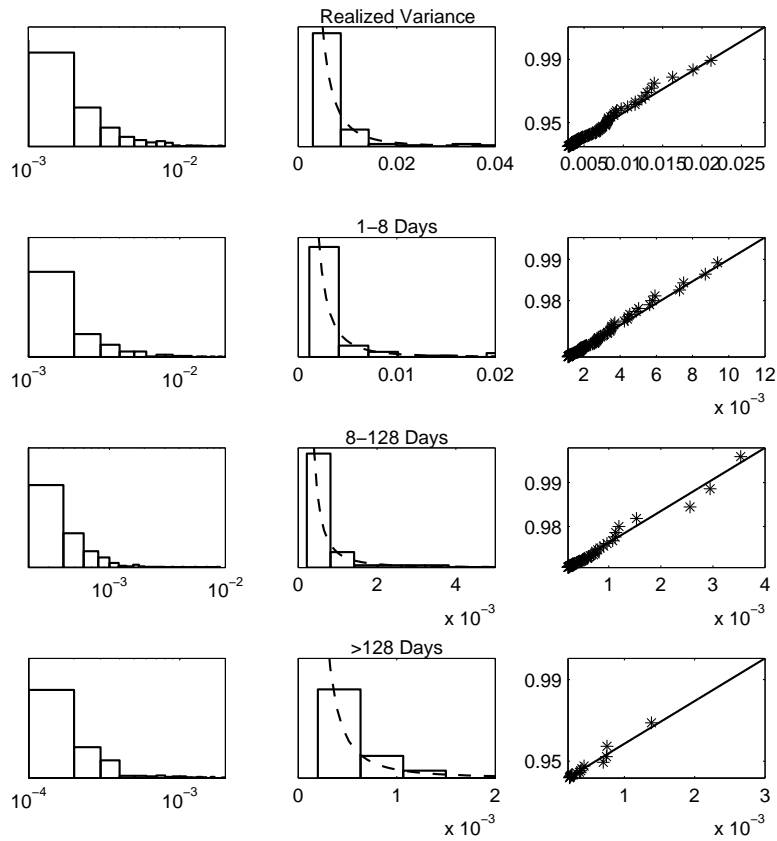
Source: Dow Jones Indexes, daily values of the DJIA index from 26/05/1896 to 10/10/2007. The figure represents the results of the GPD fit for the realized monthly variance of the index and realized monthly variances for different scale ranges, obtained by aggregation of the realized wavelet variances, as described in the text. Each line of figures shows (from left to right): the histogram of all data in logarithmic scale (2,954 observations), the fit of the GPD density for the upper tail of the distribution and the probability QQ-plot. Maximum likelihood estimation. The threshold parameter is chosen as the 0.8 empirical quantile of each distribution. Computation by the author.

Table 4: Parameters of the GPD fit for the realized wavelet variance of DJIA returns

	$\Theta \times 10^5$	ξ	$\sigma \times 10^5$	$Q_{99} \times 10^5$
		All Sample: 06/1896 - 10/2007		
IV	298.5244	0.5681	232.7840	789.4352
		[0.3649; 0.8599]	[160.6212; 395.7325]	[789.4348; 790.2242]
1-8 D	113.7658	0.6246	102.9759	340.7957
		[0.4107; 0.9305]	[71.0539; 175.0588]	[340.7954; 341.1362]
8-128 D	23.9888	0.8194	25.2972	89.2564
		[0.5448; 1.2205]	[17.4555; 43.0049]	[89.2560; 89.3453]
>128 D	21.3159	1.0740	5.9134	40.2136
		[0.6234; 1.8614]	[2.1884; 14.7833]	[40.2132; 40.2535]
		Subample: 06/1896 - 04/1952		
IV	410.5140	0.4377	375.9861	1127.3479
		[0.1977; 0.8500]	[259.4308; 639.1760]	[1127.3476; 1128.4750]
1-8 D	155.3517	0.6163	111.7318	400.0798
		[0.3554; 1.0292]	[77.0952; 189.9437]	[400.0794; 400.4795]
8-128 D	34.7769	0.8072	27.1548	104.1403
		[0.4696; 1.3586]	[10.0475; 46.1629]	[104.1399; 104.2440]
>128 D	22.9118	1.3701	3.9641	39.3502
		[0.7086; 2.7493]	[1.4670; 9.9099]	[39.3499; 39.3892]
		Subample: 05/1952 - 10/2007		
IV	216.6324	0.7106	121.6385	503.8929
		[0.4073; 1.2040]	[45.0065; 206.7851]	[503.8924; 504.3963]
1-8 D	81.5683	0.6619	60.7362	219.5061
		[0.3779; 1.1242]	[22.4727; 103.2512]	[219.5056; 219.7251]
8-128 D	14.0994	1.0602	10.3458	46.7705
		[0.6698; 1.6811]	[3.8284; 17.5875]	[46.7702; 46.8169]
>128 D	19.5670	0.9022	5.0805	33.6049
		[0.4055; 1.9608]	[1.8802; 12.7008]	[33.6046; 33.6382]

Source: Dow Jones Indexes, daily values of the DJIA index from 26/05/1896 to 10/10/2007. Maximum likelihood estimation. The threshold parameter is chosen as the 0.8 empirical quantile of each distribution. Computation by the author.
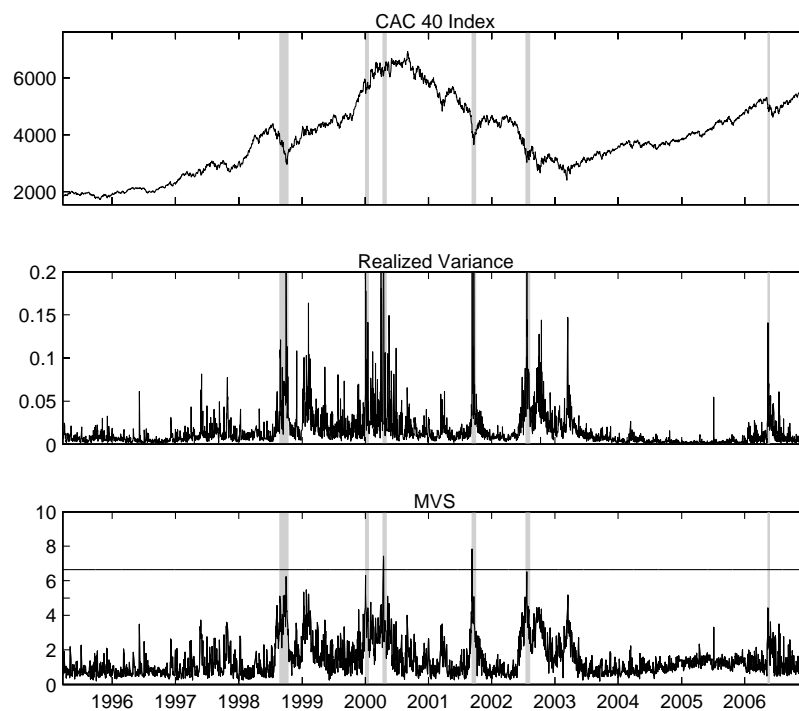
Figure 11: Market Volatility Scale for the CAC40 index

Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. The figure shows (from top to bottom) the values of the index, the realized variance and the aggregated Market Volatility Scale (MVS), which is a weighted average of the three scale components, as described in the text. Regions colored in gray correspond to the periods of extremely high volatility (tail probability less than 0.01) as detected by the MVS (see table 5). Computation by the author.

Table 5: Periods of extreme volatility detected by the MVS for the CAC40 index

| Start	End	Length	Scale	Maximum MVS				Aggregated MVS			
				W	S	M	L	W	S	M	L
27/8/98	14/10/98	34	M	6.2	6.8	7.5	7.5	143.4	4.8	6.7	125.9
4/1/00	21/1/00	13	SM	6.3	7.2	6.7	4.0	53.6	7.2	6.7	55.3
14/4/00	3/5/00	10	M	7.4	9.5	7.6	4.0	42.2	5.9	7.6	44.4
11/9/01	1/10/01	14	SM	7.8	7.8	10.5	6.5	58.1	7.8	10.5	50.7
19/7/02	9/8/02	15	L	6.5	7.2	6.5	7.2	63.4	4.0	3.8	64.3
17/5/06	24/5/06	5	M	4.4	5.9	7.0	2.3	15.6	5.9	7.0	17.4

Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. The beginning of the period of extreme volatility is the date when the MVS for one of the scales overcomes the threshold $\log_2(100) = 6.6439$, which corresponds to 0.01 probability. The end of such period is the date when the MVS for all scales falls below the threshold $\log_2(10) = 3.3219$, which corresponds to 0.1 tail probability. Duration of the period is the number of trading days between the first and the last date. The scale refers to the component of the MVS, which first broke the critical threshold: "S" stands for the short scale (intraday), "M" stands for the medium scale $(1 - 16$ days) and "L" stands for the long scale $(> 16$ days). The following four columns report the maximum value of the MVS during the period of extreme volatility: "W" stands for the weighted average (aggregated) MVS over three scales, "S", "M" and "L" for the short, medium and long scales respectively. The next four columns contain the sum of the MVS values for the dates, when one of the scale components was above $\log_2(100)$, over the period of extreme volatility. The same notations for the scales are used. Computation by the author.

Table 6: Periods of extreme volatility detected by the adapted version of the MVS for the CAC40 index

| Start	End	Length	Scale	Maximum MVS				Aggregated MVS			
				W	S	M	L	W	S	M	L
11/9/01	22/10/01	29	M	9.2	9.3	12.4	7.0	104.6	1.8	7.9	88.6
27/6/02	8/7/02	7	S	5.9	6.8	3.8	4.1	24.3	6.8	3.8	20.9
24/7/02	28/8/02	25	ML	8.4	9.0	7.8	8.0	96.2	5.5	7.8	83.8
4/10/02	15/11/02	30	L	5.5	5.7	4.6	7.6	95.6	5.7	2.4	85.9
18/3/03	7/4/03	14	SM	6.3	7.1	6.7	4.9	49.0	7.1	6.7	41.1
18/5/06	25/5/06	5	SM	5.1	7.7	8.9	0.5	15.8	7.7	8.9	16.0
1/6/06	2/6/06	1	S	5.2	7.8	3.6	1.0	5.2	7.8	3.6	1.0

Source: Euronext, values of the CAC40 index from 20/03/1995 to 29/12/2006 at 15-minute intervals. The same as in Table (5), but the adapted version of the MVS is used.

Figure 12: Adapted version of the MVS for the CAC40 index
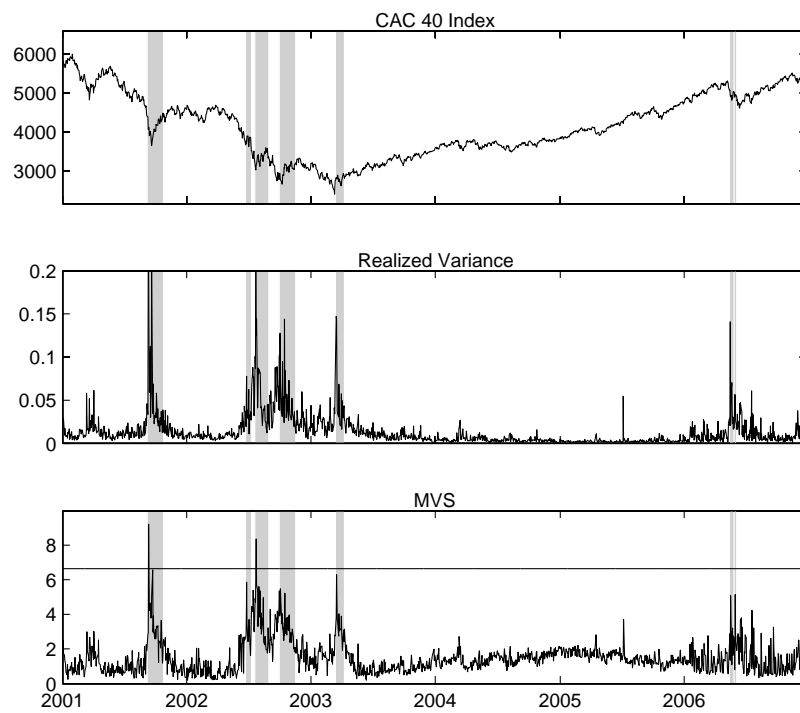
Figure 13: Market Volatility Scale for the DJIA index
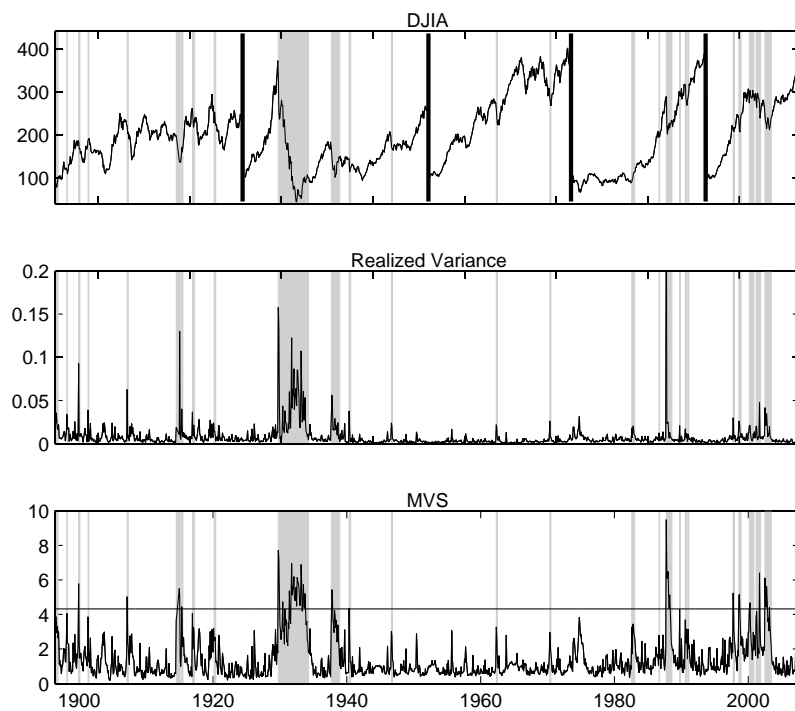
Source: Dow Jones Indexes, daily values of the DJIA index from 26/05/1896 to 10/10/2007. The figure shows (from top to bottom) the values of the index (bold vertical lines correspond to the dates when the index is rebased to 100 for the visualisation purpose), the realized variance and the aggregated MVS, which is a weighted average of the three scale components, as described in the text. Regions colored in gray correspond to the periods of extremely high volatility (tail probability less than 0.05) as detected by the MVS (see table 7). Computation by the author.

Table 7: Periods of extreme volatility detected by the MVS for the daily DJIA index

Start	End	Length	Scale	Maximum MVS				Aggregated MVS			
				W	S	M	L	W	S	M	L
06/1896	11/1896	5	L	3.9	4.2	3.2	5.0	17.3	3.9	1.4	19.9
03/1898	04/1898	1	SM	4.1	4.7	4.4	2.3	4.1	4.7	4.4	2.3
12/1899	02/1900	2	SM	5.8	7.0	6.8	2.1	8.2	7.0	6.8	4.1
05/1901	06/1901	1	S	3.9	5.1	3.9	1.0	3.9	5.1	3.9	1.0
03/1907	05/1907	2	SM	5.0	6.1	5.7	2.1	7.5	6.1	5.7	4.1
07/1914	02/1915	3	S	5.5	7.7	5.8	1.2	14.3	4.5	4.3	16.1
05/1915	07/1915	2	SM	4.4	4.9	4.8	3.3	7.5	4.9	4.8	5.7
12/1916	04/1917	4	SM	4.1	5.0	5.2	1.9	12.3	5.0	5.2	10.2
03/1920	06/1920	3	L	2.7	2.9	1.3	4.4	7.7	2.5	0.5	9.3
10/1929	04/1934	54	SML	7.7	8.0	9.0	7.2	240.9	8.0	9.0	214.3
09/1937	12/1938	15	M	5.4	5.9	5.5	6.0	52.3	4.2	4.4	41.3
05/1940	08/1940	3	SM	4.4	5.5	5.4	0.6	8.2	5.5	5.4	4.4
09/1946	11/1946	2	M	3.0	3.4	4.7	0.9	5.8	3.4	4.7	4.1
05/1962	07/1962	2	M	3.3	3.9	5.4	0.6	5.9	3.9	5.4	4.0
05/1970	07/1970	2	M	3.0	3.7	5.6	0.3	5.1	3.7	5.6	3.2
08/1982	09/1982	1	SM	3.3	4.5	4.6	0.1	3.3	4.5	4.6	0.1
10/1982	03/1983	3	S	3.4	4.8	4.1	1.7	9.7	4.8	4.1	8.2
09/1986	10/1986	1	S	2.8	4.7	2.1	0.1	2.8	4.7	2.1	0.1
10/1987	08/1988	10	SML	9.5	10.8	9.2	8.0	52.8	10.8	9.2	43.6
10/1989	11/1989	1	S	4.3	6.4	4.2	0.3	4.3	6.4	4.2	0.3
08/1990	11/1990	3	SM	3.7	4.7	4.9	0.5	9.2	4.7	4.9	7.2
02/1991	02/1991	1	S	3.1	4.4	2.8	0.4	3.1	4.4	2.8	0.4
10/1997	12/1997	2	SM	5.2	6.4	6.3	1.0	8.4	6.4	6.3	4.4
08/1998	12/1998	4	SM	5.2	6.4	5.9	2.2	17.3	6.4	5.8	14.0
03/2000	11/2000	8	S	4.7	5.3	4.2	3.9	23.0	5.3	4.2	18.8
03/2001	07/2001	4	M	4.2	4.8	5.0	2.3	12.3	4.1	5.0	9.8
09/2001	11/2001	2	SM	6.4	7.2	6.4	3.8	9.9	7.2	6.4	7.1
07/2002	06/2003	11	SML	6.1	6.6	5.9	6.8	46.1	6.6	5.9	39.7

Source: Dow Jones Indexes, daily values of the DJIA index from 26/05/1896 to 10/10/2007. The beginning of the period of extreme volatility is the month when the MVS for one of the scales overcomes the threshold $\log_2(20) = 4.3219$, which corresponds to 0.05 probability. The end of such period is the month when the MVS for all scales falls below the threshold $\log_2(10) = 3.3219$, which corresponds to 0.1 probability. Duration of the period is the number of months between the first and the last date. The scale refers to the component of the MVS, which first broke the critical threshold: "S" stands for short scale (1-8 days), "M" stands for the medium scale (from 8 to 128 days) and "L" stands for the long scale (more than 128 days). The following four columns report the maximum value of the MVS over the period of extreme volatility: "W" stands for the weighted average (aggregated) MVS over the three scales, "S", "M" and "L" for the short, medium and long scales respectively. The next four columns contain the sum of the MVS values, when one of the scale components was above $\log_2(20)$, over the period of extreme volatility. The same notations for the scales are used. Computation by the author.