

Vuorenmaa, Tommi A.

Working Paper

A wavelet analysis of scaling laws and long-memory in stock market volatility

Bank of Finland Research Discussion Papers, No. 27/2005

Provided in Cooperation with:

Bank of Finland, Helsinki

Suggested Citation: Vuorenmaa, Tommi A. (2005) : A wavelet analysis of scaling laws and long-memory in stock market volatility, Bank of Finland Research Discussion Papers, No. 27/2005, ISBN 952-462-254-8, Bank of Finland, Helsinki, <https://nbn-resolving.de/urn:nbn:fi:bof-20140807666>

This Version is available at:

<https://hdl.handle.net/10419/212025>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.



Tommi A. Vuorenmaa

A wavelet analysis of scaling laws and long-memory in stock market volatility



EUROJÄRJESTELMÄ
EUROSYSTEMET

Bank of Finland Research
Discussion Papers
27 • 2005

**Suomen Pankki
Bank of Finland
P.O.Box 160
FIN-00101 HELSINKI
Finland**

☎ + 358 10 8311

<http://www.bof.fi>

Bank of Finland Research
Discussion Papers
27 • 2005



Tommi A Vuorenmaa*

A wavelet analysis of scaling laws and long-memory in stock market volatility

The views expressed are those of the author and do not necessarily reflect the views of the Bank of Finland.

* Department of Economics, University of Helsinki.
University of California, San Diego.

<http://www.bof.fi>

ISBN 952-462-253-X
ISSN 0785-3572
(print)

ISBN 952-462-254-8
ISSN 1456-6184
(online)

Multiprint Oy
Helsinki 2005

A wavelet analysis of scaling laws and long-memory in stock market volatility

Bank of Finland Research
Discussion Papers 27/2005

Tommi A. Vuorenmaa
Monetary Policy and Research Department

Abstract

This paper investigates the dependence of average stock market volatility on the timescale or on the time interval used to measure price changes, which dependence is often referred to as the scaling law. Scaling factor, on the other hand, refers to the elasticity of the volatility measure with respect to the timescale. This paper studies, in particular, whether the scaling factor differs from the one in a simple random walk model and whether it has remained stable over time. It also explores possible underlying reasons for the observed behaviour of volatility in terms of heterogeneity of stock market players and periodicity of intraday volatility. The data consist of volatility series of Nokia Oyj at the Helsinki Stock Exchange at five minute frequency over the period from January 4, 1999 to December 30, 2002. The paper uses wavelet methods to decompose stock market volatility at different timescales. Wavelet methods are particularly well motivated in the present context due to their superior ability to describe local properties of times series. The results are, in general, consistent with multiscaling in Finnish stock markets. Furthermore, the scaling factor and the long-memory parameters of the volatility series are not constant over time, nor consistent with a random walk model. Interestingly, the evidence also suggests that, for a significant part, the behaviour of volatility is accounted for by an intraday volatility cycle referred to as the New York effect. Long-memory features emerge more clearly in the data over the period around the burst of the IT bubble and may, consequently, be an indication of irrational exuberance on the part of investors.

Key words: long-memory, scaling, stock market, volatility, wavelets

JEL classification numbers: C14, C22

Aalokeanalyysi osakemarkkinoiden volatiliteetin skaalauslaeista ja pitkämuistisuudesta

Suomen Pankin tutkimus
Keskustelualoitteita 27/2005

Tommi A. Vuorenmaa
Rahapolitiikka- ja tutkimusosasto

Tiivistelmä

Tässä työssä tutkitaan Helsingin pörssistä kerättyjen Nokian osakkeen hintatietojen avulla osakemarkkinoiden keskimääräisen hintavaihtelun, volatiliteetin, ja sen mittaamisessa käytetyn aikavälin, osakemarkkina-aineiston aikaskaalan, välistä riippuvuutta. Tätä riippuvuutta kutsutaan skaalauslaiksi ja siihen liittyvällä skaalauskertoimella tarkoitetaan käytetyn volatiliteettimittarin joustoa aikaskaalan suhteen. Tässä työssä tutkitaan erityisesti, onko osakemarkkina-aineiston skaalauskerroin sopusuunnassa yksinkertaisen satunnaispolkumallin kanssa ja onko se pysynyt vakaana ajan mittaan. Tutkimuksessa pohditaan myös investoijien heterogeenisuuden ja päivänsisäisten vaihtelujaksojen merkitystä volatiliteettiä selittävinä tekijöinä. Aaloke- eli väreanalyysia käyttäen osakemarkkinoiden volatiliteetti hajotetaan aikaskaaloittain. Aalokeanalyysin käyttö on tässä yhteydessä perusteltua, mikä johtuu tämän menetelmän erinomaisesta kyvystä kuvata aikasarjan lokaalisia ominaisuuksia. Suomen osakemarkkinat ovat tulosten mukaan moniskaalaiset. Skaalaustekijä ja volatiliteetin pysyvyyteen liittyvät parametrit eivät ole pysyneet vakioina ajan kuluessa, eivätkä niiden estimaatit tue satunnaispolkumallia. Mielenkiintoista kyllä, osakemarkkinoiden volatiliteettiin vaikuttaa tulosten mukaan erityinen päivänsisäinen vaihtelujakso, josta tutkimuksessa käytetään nimitystä New York -vaikutus. Volatiliteetin pitkän aikavälin ominaisuudet ilmenevät aineistossa selkeämmin IT-kuplan aikana kuin sen jälkeen, mikä saattaa olla todiste sijoittajien lyhytnäköisyydestä.

Avainsanat: pitkämuistisuus, skaalaus, osakemarkkinat, heilahtelut, aallokkeet (väreet)

JEL-luokittelu: C14, C22

Contents

Abstract.....	3
Tiivistelmä (abstract in finnish).....	4
1 Introduction	7
2 Wavelet decomposition of variance	8
2.1 Wavelet filters.....	8
2.2 Maximal overlap discrete wavelet transform	9
2.3 Wavelet variance	10
3 Long-memory volatility modeling	12
4 Empirical analysis	13
4.1 Data description.....	13
4.2 Preliminary data analysis.....	15
4.3 Multiresolution decomposition.....	18
4.4 Global scaling laws and long-memory	18
4.5 Local scaling laws and long-memory.....	23
4.6 Effects of volatility periodicity	24
5 Conclusions	31
References.....	34
Appendix Fractional differencing and long-memory	38

1 Introduction

Stock market volatility exhibits jumps and clustering (see eg Cont (2001)). For the last two decades the main emphasis has been put into the research of volatility clustering, also known as the ‘ARCH-effect’. The seminal articles of Engle (1982) and Bollerslev (1986) launched a huge interest in different kinds of (generalized) autoregressive conditional heteroskedastic ((G)ARCH) models (for a review, see eg Bollerslev et al (1992)).¹ The huge interest in the conditional variance stems from the fact that a correctly specified volatility model is useful for example in valuation of stocks and stock options and in designing optimal dynamic hedging strategies for options and futures.

The ARCH-models have still hard time explaining the stylized facts, however. This is partly because they typically model only one time-scale (usually a day) at time. But stock market data have no specific time-scale to analyze. A notable exception in this respect is the heterogenous ARCH model (Müller et al (1997)) which is based on the hypothesis of a heterogenous market (Müller et al (1993)). According to this hypothesis the stock market consists of multiple layers of investment horizons (time-scales) varying from extremely short (minutes) to long (years). The short horizons are thought to be related to speculation and the longer horizons to serious investing. Of course, the players in the stock market form a heterogenous group with respect to other reasons as well – such as perceptions of the market, risk profiles, institutional constraints, degree of information, prior beliefs, and other characteristics such as geographical locations – but as Müller et al (1993) argue, many of these differences translate to sensitivity to different time-scales. In fact, Müller et al (1997) show evidence that time-scale is one of the most important aspects in which trading behaviors differ. This is convenient because high-frequency data has made it possible to study scale dependent phenomena in the stock markets.

The incorporation of multiple time-scales into the analysis should improve the efficiency of risk management which requires scaling a risk measure (standard deviation, say) of one time-scale to another. The industry standard is to scale by the square-root of time, familiar from Brownian motion. But such scaling implicitly assumes that the *data generating process* (DGP) is made of *identically and independently distributed* (IID) random variables. This assumption is not reasonable for financial time series where the persistence in conditional second moments is universally found so strong and long-lasting that volatility is said to exhibit long-memory. Under such non-IID circumstances square-root scaling may be misleading (see Diebold et al (1997)).

Several authors have provided evidence of scaling laws in the FX markets (eg Müller et al (1990, 1997), Guillaume et al (1997), and Andersen et al (2000)). But a single scaling factor may not always be adequate. Gençay et al (2001a) have found that a different scaling region exists in the FX markets for intraday time-scales and for larger time-scales. It is thus interesting to see (i) if different scaling regions (known as multiscaling, see eg Fisher et al

¹In 2003 Engle was given a (half of) Nobel Prize in Economic Sciences ‘for methods of analyzing economic time series with time-varying volatility’.

(1997)) appear in a stock market which has smaller turnover, lower liquidity, and higher transaction costs than the FX markets, (ii) if the scaling factor systematically differs from the Brownian, (iii) if the scaling factor is constant in time, and (iv) if the behavior can be explained by the heterogeneity of the players in the market or by other means such as intraday volatility periodicity. Moreover, because the scaling law is intimately related to the memory of the DGP, and because high-frequency data allows for superior estimation of the long-memory parameter (see Bollerslev and Wright (2000)), this paper sheds light on the behavior of long-memory in volatility in time. The data used are the 5-minute volatility series of Nokia Oyj at the Helsinki Stock Exchange around the burst of the IT-bubble. Period one represents the era of *'irrational exuberance'* (Shiller (2001)) and another its aftermath.

This paper attempts to answer all these issues using wavelet analysis. Wavelet analysis is a non-parametric method that allows for the study of time-scale dependent phenomena. It is akin to Fourier analysis but does not lose time dimension in the transformation which is useful for time-aligning. Wavelet analysis is locally adaptive. Jumps and clusters of volatility do not present a problem for its application and it is thus well-suited for the analysis of stock market data. This gives the wavelet method a distinct advantage over the standard frequency domain methods. In particular a wavelet-based OLS method allows for consistent estimation of long-memory.

The structure of this paper is as follows. Section 2 shortly describes the wavelet methodology and the concept of wavelet variance. In Section 3 a locally stationary long-memory stochastic volatility model is reviewed. Section 4 describes the empirical results. Section 5 concludes.

2 Wavelet decomposition of variance

There are nowadays number of good accounts on wavelet methodology. The reader could for example consult some of the following for more information than given in the next subsections: An introduction for economists is given by Schleicher (2002) and Crowley (2005). A more detailed description is Vuorenmaa (2004). A short review of wavelets in statistical time series is Nason and von Sachs (1999). A comprehensive treatment is Percival and Walden (2000). The close relationship to Fourier analysis is discussed in Priestley (1996). Mathematical proofs can be found in Härdle et al (1998).

2.1 Wavelet filters

A wavelet filter h_l acts as a high-pass (or more precisely, a band-pass) filter. This means that convolving a wavelet filter with data gives the (empirical) wavelet coefficients, ie the details with low-frequencies filtered out. To qualify

as a wavelet filter, h_l of length L must satisfy:

$$\sum_{l=0}^{L-1} h_l = 0,$$

$$\sum_{l=0}^{L-1} h_l^2 = 1, \text{ and } \sum_{l=0}^{L-1} h_l h_{l+2n} = 0,$$

for all nonzero integers n (see Percival and Walden (2000, p. 69)).

An example of a compactly supported Daubechies wavelet filter is the Haar wavelet of length $L = 2$,

$$h_0^{\text{Haar}} = 1/\sqrt{2} \text{ and } h_1^{\text{Haar}} = -1/\sqrt{2},$$

for which the above conditions are easily checked to hold. In general, however, the Daubechies filters have no explicit time-domain formulae. The values are tabulated instead (see eg Percival and Walden (2000)). Nevertheless, Daubechies filters are practical because they yield a *discrete wavelet transform* (DWT) that can be described in terms of generalized differences of weighted averages. This means that the Daubechies wavelet filters are capable of producing stationary wavelet coefficient vectors from ‘higher degree’ non-stationary stochastic processes.

The choice of a proper width can be somewhat tricky. A wider width wavelet filter prevents undesirable artifacts and results in a better match to the characteristic features of a time series. However, as the width gets wider, more coefficients are being unduly influenced by boundary conditions. This is associated with an increase in computational burden. Thus one should search for the smallest L that gives reasonable results. If one also wants to have the DWT coefficients be alignable in time, the optimal choice in empirical studies is often the *least asymmetric* wavelet filter of length 8 (LA(8)). This is the wavelet filter used in this paper, as well.

2.2 Maximal overlap discrete wavelet transform

The *maximal overlap discrete wavelet transform* (MODWT) is a non-orthogonal transform. With a proper algorithm the complexity of the MODWT is of the same order as of the fast Fourier transform. This is useful because in high-frequency finance the number of observations is high. Percival and Walden (2000, pp. 159–60) list some properties that distinguish the MODWT from the DWT. For the present needs, it is enough to mention that the MODWT can handle any sample size N and that the MODWT wavelet variance estimator (to be defined in the next subsection) is asymptotically more efficient than the estimator based on the DWT.

The MODWT is easily formulated using matrices (for details, see eg Gençay et al (2002, Ch. 4.5)). The length $(J + 1)N$ column vector of MODWT coefficients $\tilde{\mathbf{w}}$ is obtained by

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

where \mathbf{x} is a length N column vector of observations and $\widetilde{\mathcal{W}}$ is a $(J+1)N \times N$ non-orthogonal matrix defining the MODWT. The vector $\widetilde{\mathbf{w}}$ and the matrix $\widetilde{\mathcal{W}}$ consist of length N column subvectors $\widetilde{\mathbf{w}}_1, \dots, \widetilde{\mathbf{w}}_J, \widetilde{\mathbf{v}}_J$ and $N \times N$ submatrices $\widetilde{\mathcal{W}}_1, \dots, \widetilde{\mathcal{W}}_J, \widetilde{\mathcal{V}}_J$, respectively. The level j wavelet coefficients $\widetilde{\mathbf{w}}_j$ are associated with *changes* on a scale of length $\lambda_j \doteq 2^{j-1}$ and the scaling coefficients $\widetilde{\mathbf{v}}_J$ are associated with *averages* on a scale of length $2^J \doteq 2\lambda_J$. Only the wavelet coefficients are going to play role in the subsequent analysis.

In what follows, I will rely on the fact that the MODWT is capable of producing a scale-by-scale analysis of variance upon the energy decomposition:

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

This can be used to analyze phenomena consisting of variations over a range of different scales. In stock markets the usefulness of a decomposition with respect to time-scale was motivated in the introduction. The theoretical counterpart of variance decomposition is discussed next in the context of long-memory processes.

2.3 Wavelet variance

Consider an ARFIMA process $\{X_t\}$ (see Appendix) whose d th order backward difference Y_t is a stationary process with mean μ_Y (not necessarily zero). Then a Daubechies wavelet filter \widetilde{h}_l of width $L \geq d$ results in the j th wavelet coefficient process $\overline{w}_{j,t} \doteq \sum_{l=0}^{L_j-1} \widetilde{h}_{j,l} X_{t-l}$ being a stationary process. Now define the (time-independent or global) *wavelet variance* for $\{X_t\}$ at scale λ_j to be

$$\nu_X^2(\lambda_j) \doteq \mathbb{V}\{\overline{w}_{j,t}\},$$

which represents the contribution to the total variability in $\{X_t\}$ due to changes at scale λ_j . By summing up the time-scale specific wavelet variances, we get the variance of $\{X_t\}$:

$$\sum_{j=1}^{\infty} \nu_X^2(\lambda_j) = \mathbb{V}\{X_t\}.$$

The wavelet variance is well-defined for both stationary and non-stationary processes with stationary d th order backward differences as long as the width L of the wavelet filter is large enough. An advantage of the wavelet variance is that it handles both types of processes equally well. An unbiased² estimator of $\nu_X^2(\lambda_j)$ is

$$\widetilde{\nu}_X^2(\lambda_j) = \frac{1}{M_j} \sum_{t=L_j-1}^{N-1} \widetilde{w}_{j,t}^2,$$

²Unbiasedness tells us that the distribution of the estimate is centered around the unknown true value of the parameter.

where $M_j \doteq N - L_j + 1 > 0$ and $\tilde{w}_{j,t} \doteq \sum_{l=0}^{L_j-1} \tilde{h}_{j,l} X_{t-l \bmod N}$ are the (periodically extended) MODWT coefficients. Here only coefficients unaffected by the periodic boundary conditions are included in the sum; otherwise the estimator would be biased.

A Fourier spectrum decomposes the variance of a series across different frequencies. Because the scales that contribute the most to the variance of the series are associated with those coefficients with the largest variance, it is not surprising that the estimates of the wavelet variance can be turned into SDF estimates. The approximation improves as the width L of the wavelet filter increases because then $\tilde{h}_{j,l}$ becomes a better approximation to an ideal band-pass filter. Now, it is well-known that the periodogram is an inconsistent³ estimator of the Fourier spectrum. It follows that the popularly used GPH-estimator (Geweke and Porter-Hudak (1983)) based on the *ordinary least squares* (OLS) regression of the log-periodogram for frequencies close to zero is in general an inconsistent estimator of the long-memory parameter from a fractionally integrated process with $|d| < 1/2$. And although the GPH-estimator can be shown to be consistent under certain regularity conditions (Gaussianity, in particular), these are not realistic with financial data (volatility is not distributed normally).

Jensen (1999) has shown that a wavelet based OLS-estimator is consistent when the sample variance of the wavelet coefficients is used in the regression. Namely, using the wavelet variance of the DWT coefficients $w_{j,t}$,

$$\bar{\nu}_X^2(\lambda_j) = \frac{1}{2^j} \sum_{k=0}^{2^j-1} w_{j,k}^2, \quad (2.1)$$

we get that

$$\mathbb{V}\{w_{j,t}\} = \nu_X^2(\lambda_j) \rightarrow \sigma^2 2^{j(2d-1)},$$

as $j \rightarrow \infty$ (here σ^2 is a finite constant). By taking logarithms on both sides, we then obtain the (approximate) log-linear relationship

$$\log \nu_X^2(\lambda_j) = \log \sigma^2 + (2d - 1) \log 2^j, \quad (2.2)$$

from which the unknown d can be estimated consistently by OLS-regression by replacing ν_X^2 with its sample variance $\bar{\nu}_X^2$ of Equation (2.1). Jensen has also shown that in *mean square error* (MSE) sense the wavelet based OLS-estimator fares significantly better than the GPH-estimator. The asymptotic efficiency of this estimator can be further improved by using the MODWT coefficients instead of the DWT coefficients.

Wavelet variance can be defined also locally. In this case only the wavelet coefficients "close" to the time point t are used. Namely, given $L > 2d(u)$, an unbiased estimator of local wavelet variance for $\{X_t\}$ at scale λ_j based upon the MODWT is

$$\tilde{\nu}_X^2(u, \lambda_j) = \frac{1}{K_j} \sum_{s=\tau_j}^{\tau_j+K_j} \tilde{w}_{j,t+s,T}^2, \quad (2.3)$$

³An estimator is said to be consistent if it gets closer and closer to the true value of the parameter as the number of observations grow.

where $u = t/T$ represents a time point in the rescaled time domain $[0, 1]$, K_j is a "cone of influence", and τ_j is an "offset" (Whitcher and Jensen (2000, p. 98)). The K_j includes only those wavelet coefficients where the corresponding observation made a significant contribution. An inconvenience of this approach is that K_j varies across scales and different filters. The tabulated values for the Daubechies family of wavelets are given in Whitcher and Jensen (2000) as well as the the values of 'offsets' τ_j for each wavelet filter $L > 2$.

Whitcher and Jensen (1999) have shown that when the MODWT is being applied to a *locally stationary* (in the sense of Dahlhaus (1997)) long-memory process $\{X_{t,T}\}$, then the level- j MODWT wavelet coefficients $\{\tilde{w}_{j,t,T}\}$ form a locally stationary process with mean zero and time-varying variance

$$\mathbb{V}\{\tilde{w}_{j,t,T}\} = \nu_X^2(u, \lambda_j) \rightarrow \sigma^2(u)2^{j[2d(u)-1]},$$

as $j \rightarrow \infty$ ($\sigma^2(u)$ is given in Whitcher and Jensen (2000)). Analogously, then,

$$\log \nu_X^2(u, \lambda_j) = \log \sigma^2(u) + [2d(u) - 1] \log 2^j, \quad (2.4)$$

from which the unknown $d(u)$'s can be estimated consistently by OLS by replacing ν_X^2 by its time-varying sample variance $\tilde{\nu}_X^2$ from Equation (2.3). Using simulations Whitcher and Jensen have shown that in the case of a globally stationary ARFIMA the median of $\hat{d}(u)$ accurately estimates the true value of d with a slight negative bias near the boundaries. They also found that when disturbed by a sudden shift in the long-memory parameter (to imitate local stationarity), the estimated d still performed well on both sides of the change although with a slight bias and increase in the MSE at the boundaries.

3 Long-memory volatility modeling

A discrete-time *stochastic volatility* (SV) model may be written as

$$y_t = \sigma \exp(h_t/2)\varepsilon_t,$$

where y_t denotes the demeaned return process $y_t = \log(P_t/P_{t-1}) - \mu$ (here P_t is the price of a stock and μ is the mean of returns), $\{\varepsilon_t\}$ is a series of IID random disturbances with mean 0 and variance 1, and the conditional variance $\{\sigma_t^2\}$ is modeled as a stochastic process $\{\log \sigma_t^2\} \doteq \{h_t\}$ and it is independent of $\{\varepsilon_t\}$.

There exist many specifications for the volatility scheme $\{h_t\}$, such as ARMA or random walk. Recently a *long-memory stochastic volatility* (LMSV) model proposed in Breidt et al (1998) has caught a lot of attention. In their model $\{h_t\}$ (log-volatility) is generated by fractionally integrated Gaussian noise,

$$(1 - B)^d h_t = \eta_t,$$

where $|d| < 1/2$, B is the lag operator, and $\eta_t \sim NID(0, \sigma_\eta^2)$. This model encompasses a 'short-memory' model when $d = 0$. But the LMSV model is a

stationary model and it thus ignores for example intraday volatility patterns, irregular occurrences of market crashes, mergers and political coups as noted by Jensen and Whitcher (2000). In particular, it may be that the long-memory parameter d is not constant over time. This motivated Jensen and Whitcher to introduce a non-stationary class of long-memory stochastic volatility models with time-varying parameters. In their model, the logarithmic transform of the squared returns is a locally stationary process that has a time-varying spectral representation. This means that the level of persistence associated with a shock to conditional variance – which itself is allowed to vary in time – is dependent on when the shock takes place (the shocks themselves, of course, still produce responses that persist hyperbolically).

More specifically, a locally stationary LMSV model is defined by

$$\begin{aligned} y_{t,T} &= \exp(H_{t,T}/2) \varepsilon_t, \\ \Phi(t/T, B)(1 - B)^{d(t/T)} H_{t,T} &= \Theta(t/T, B) \eta_t, \end{aligned}$$

where $|d(u)| < 1/2$, $\varepsilon_t \sim NID(0, 1)$ and $\eta_t \sim NID(0, \sigma_\eta^2)$ are independent of each other. The functions $\Phi(u, B)$ and $\Theta(u, B)$ are, respectively, order p and q polynomials whose roots lie outside the unit circle uniformly in u and whose coefficient functions, $\phi_j(u)$, for $j = 1, \dots, p$, and $\theta_k(u)$, for $k = 1, \dots, q$, are continuous on \mathbf{R} . The coefficient functions satisfy $\phi_j(u) = \phi_j(0)$, $\theta_k(u) = \theta_k(0)$ for $u < 0$, and $\phi_j(u) = \phi_j(1)$, $\theta_k(u) = \theta_k(1)$ for $u > 1$, and are differentiable with bounded derivatives for $u \in [0, 1]$. Notice that by setting $\Phi(u, B) = \Phi(B)$, $\Theta(u, B) = \Theta(B)$, and $d(u) = 0$ for all $u \in [0, 1]$, one gets the SV model of Harvey et al (1994). If, on the other hand, one sets $d(u) = d$ for all $u \in [0, 1]$, one gets the LMSV model.

4 Empirical analysis

4.1 Data description

The data consist of transactions of Nokia Oyj between January 4 (1999) and December 30 (2002). There are at least two good reasons for choosing Nokia. First, Nokia has been the market leader in the cellular phone industry for many years now and is thus representative of the bubble times of the late 1990s. Second, Nokia is a highly liquid stock at the *Helsinki Stock Exchange* (HEX) which in year 2003, for example, accounted for 62,1% of the total number of Nokia shares traded in the whole world. For comparison, the same percentage for the *New York Stock Exchange* (NYSE) was only 20,3% (the HEX (May 4, 2004)) where (and the NASDAQ) Nokia has the largest trading volume of cross-listed non-U.S. companies (Citigroup (June 23, 2004)).

The data were discretized by extracting the 5-minute prices P_t using the closest transaction price to the relevant time mark. Discretizing is necessary for the wavelet decomposition to be interpretable in terms of time-scales that capture a band of frequencies. Theoretically discretizing can be justified by assuming that the DGP does not vary significantly over short time intervals. For a liquid stock this holds true (there is no ‘non-synchronous trading’). The

Table 1: Periods I and III.

Time period	Trading (I)	AMT (I)	Trading (II)
1/4/99 – 8/31/00	10:30–17:30	17:30–18:00	–
4/17/01 – 3/27/02	10:00–18:00	18:03–18:30	18:03–21:00

5-minute returns were then formed as

$$r_{t,n} = 100 (\ln P_{t,n} - \ln P_{t,n-1}),$$

where $r_{t,n}$ denotes the return for intraday period n on trading day t , with $d \geq 1$ and $t = 1, \dots, T$. The 5-minute interval has been often used in the literature because it is usually the smallest interval that does not suffer badly ‘bid-ask bounce’ (see eg Campbell et al (1997, Ch. 3)). In the case of a missing observation for a specific time mark, the previous price was always used. Prices were adjusted for splits but not for dividends. Block trades were not controlled for either. These omissions are not critical for the subsequent analysis, however.

At the HEX an electronic trading system called *Helsinki Stock Exchange Automated Trading and Information System* (HETI) has been in use since 1990. This means that brokers trade electronically, the smallest ‘tick-size’ being 0.01 (in euros). Notice that the HEX did not have constant trading hours during 1999–2002. The changes in trading hours were mainly caused by an international pressure towards harmonization of exchange open hours in Europe. In particular, the long-run trend of longer trading days was suppressed by the weak market conditions after the burst of the IT-bubble. In what follows, I analyze subperiods I and III separately. This is because there is a reason to believe that a structural break took place in between these periods. For this reason I refer to the two periods as ‘I’ and ‘III’ instead of the more logical ‘I’ and ‘II’.

In Period I, from January 4 (1999) to August 31 (2000), continuous trading took place between 10:30 a.m. and 5:30 p.m. (total of 7 hours or 85 intraday 5-minute prices). Transactions between 5:30 and 6:00 p.m. were discarded because they belonged to *after market trading #1* (AMT (I)).⁴ Only one day, namely April 20 (2000), was an incomplete day and the missing observations were substituted by the last observed price. In total, there were 419 trading days resulting in 35,615 price observations.

In Period III, from April 17 (2001) to March 27 (2002), continuous trading included evening hours from 6:00 to 9:00 p.m (total of 11 hours or 133 intraday 5-minute prices). A technical break (no transactions), occurred every trading day between 6:00 and 6:03 p.m. After that continuous trading and AMT (I) took place simultaneously. This simultaneity required careful filtering. I decided to apply the following rule: prices that had a percentage price change of more than 3% relatively to the last genuine price recorded (before the technical break at 6:00 p.m.) were detected as artificial and replaced by the previous genuine price. There were no incomplete trading days. In total, there were 237 trading days resulting in 31,521 price observations.

⁴During AMT the trading price can fluctuate between the trading range established during continuous trading for round-lot trades ([Http://www.porssisaatio.fi](http://www.porssisaatio.fi)).

Table 2: Key figures of Periods I and III.

Period I						
Min.	1st Q.	Med.	Mean	3rd Q.	Max.	Std.
-11.61	-0.1118	0	$+3.747e - 03$	0.1161	10.97	0.3789
Period III						
Min.	1st Q.	Med.	Mean	3rd Q.	Max.	Std.
-11.16	-0.1523	0	$-6.798e - 04$	0.1534	14.05	0.3869

4.2 Preliminary data analysis

Statistical key figures of Periods I and III are summarized below (Table 2). Notice the following: First, Periods I and III are of approximately equal size which is convenient from statistical inference point of view. Second, Periods I and III represent turbulent and calm regimes, respectively: Period I is representative of the ‘IT-bubble’ and Period III of its aftermath.⁵ The volatilities of Periods I and III seem to differ from each other by simply glancing at the return series (see the bottom plots of Figs. 1 and 2).⁶ This observation is valuable because structural breaks can generate artificial long-memory (see eg Granger and Hyung (1999)). In particular, Mikosch and Stărică (2004) have argued that long-memory might be due to non-stationarity implying that stationary models are inappropriate over longer horizons. It is thus safer to analyze these two periods separately.

The sample *autocorrelation functions* (ACFs) of returns in Periods I and III differ from each other in a non-trivial way (see the top plots of Figs. 3 and 4). The opening of the HEX as well the U.S. markets (in New York) at 5:30 p.m. (Central European Time +1) has caused some statistically significant linear dependence in Period I.⁷ Of course, this does not necessarily imply arbitrage opportunities because transactions costs can be high. But considering the slightly different results of Period III, it seems that when the markets cooled down they in fact became ‘more efficient’ around the openings. A bit surprisingly, though, in Period I no significant negative autocorrelation of MA(1) type exists that is typically reported (Andersen and Bollerslev (1997b) found it to be -0.04 in the FX markets with 5-minute data) and attributed to bid-ask bounce. In Period III, a significant negative first-lag autocorrelation (-0.08) does appear, however. In the subsequent analysis the lag one dynamics have not been considered critical and therefore they have not been filtered out.

To proxy volatility, I use absolute returns instead of squared returns in order to prevent big jumps in volatility. The sample ACFs of absolute returns

⁵Polzehl et al (2004) find that the ‘2001 Recession’ in the U.S. might have started as early as October 2000 and ended as late as the Summer of 2003. This supports the division to two periods.

⁶Counterintuitively, the standard deviation of returns of Period I is actually smaller than that of Period III (0.3789 vs. 0.3869, respectively). Similarly, the means of absolute returns are 0.1874 and 0.2287, respectively. Wavelet variance decomposition will shed more light on this finding.

⁷The NYSE opens at 9:30 a.m. local time (Eastern Standard Time). When comparing the figures to each other, recall that the length of the trading day was different in Periods I and III.

Figure 1: Price and return series of Period I (IT-bubble period)





Figure 2: Price and return series of Period III (altermath period)



stay significantly positive for a long time in both periods, statistically as well as economically (see the bottom plots of Figs. 3 and 4). In Period III, for example, the first-lag autocorrelation (0.32) is well above the confidence interval (Andersen and Bollerslev (1997b) found 0.31). Clearly returns are not independent. Although the pattern is quite similar in both periods, there are some important differences here too. First, the ACF peaks higher in Period I than in Period III. This peak is caused by the large (on average) overnight return in Period I. The larger ‘overnight effect’ in Period I is most probably due the frequent news arrivals, the irrational exuberance (the hype) that took place during the bubble, and the shorter trading day at the HEX (so that information had more time to accumulate over night). Second, in Period I the first peak just prior to the highest peak is a reflection of the opening of the New York stock markets, here referred to as the ‘New York effect’.⁸ In Period III the New York effect in autocorrelations is smaller which is probably due the weaker link between the U.S. and European markets after the burst of the IT-bubble.

4.3 Multiresolution decomposition

In order to study volatility at different time-scales, the MODWT($J = 12$) is performed to absolute returns using LA(8) (with reflecting boundary). The first 12 wavelet levels with the corresponding time-scales and associated changes are listed below (see Table 3).⁹ Interpreting the time-scales in ‘calendar time’ requires care since the length of the trading day varied. So, for instance, in Period I the first 6 levels correspond to intraday (and daily) dynamics capturing frequencies $1/64 \leq f \leq 1/2$, ie oscillations with a period of $10 - 320$ ($2 * 5 - 64 * 5$) minutes. In Period III, on the other hand, the first 7 levels correspond to intraday (and daily) dynamics capturing frequencies $1/128 \leq f \leq 1/2$, ie oscillations with a period of $10 - 640$ minutes. In terms of changes, then, the 6th level in Period I corresponds to approximately a half of a trading day. In Period III this corresponds to the 7th level. These levels thus serve as natural watersheds between intraday and interday dynamics.

4.4 Global scaling laws and long-memory

Most of the total energy of volatility is located at the smallest time-scales (the highest frequencies). The relationship is approximately hyperbolic which is observed as an approximate linear relationship on a double-logarithmic scale. The Gaussian confidence interval for Period III suggests that there exist two

⁸Of course it is possible that other market places affect volatility at the HEX too but as will be demonstrated later (in Sec. 4.6), the average intraday volatility peaks consistently at the opening of the New York market. In the literature, volatility spillover effects – ‘*meteor showers*’ – have been reported for example by Engle et al (1990).

⁹An unfortunate consequence of the dyadic dilation is that time-scales become coarse rapidly so that not all of the potentially interesting scales are recovered. Thus the non-dyadic extension (Pollock and Lo Cascio (2004)) might be worthwhile to look at.

Figure 3: The sample ACFs of returns and absolute returns in Period I. The 95% confidence interval (dashed line) is for Gaussian white noise: $\pm 1.96/\sqrt{N}$.




Figure 4: The sample ACFs of returns and absolute returns in Period III.




Table 3: Wavelet levels and time-scales.

Level	Scale	Associated with changes of
1	1	5 min.
2	2	10 min.
3	4	20 min.
4	8	40 min.
5	16	80 min.
6	32	160 min. \approx 3 h.
7	64	320 min. \approx 5 h.
8	128	640 min. \approx 11 h.
9	256	1280 min. \approx 21 h.
10	512	2560 min. \approx 43 h.
11	1024	5120 min. \approx 85 h.
12	2048	10240 min. \approx 171 h.

different scaling regions in Period I – a finding that is similar to Gençay et al (2001a) in the FX markets. A break is located at the 7th level associated with 320-minute changes (see Fig. 5).

One might expect a similar break at the 8th level in Period III because of its longer trading day (11 hours versus 7 hours), but this does not happen. The difference between the scaling laws of Periods I and III is most evident at level 6. Period I experienced more middle-sized jumps in this particular time-scale than Period III did. This observation is however not enough to explain the extremely jumpy look of Period I. Because jumps are high-frequency events, they should be well captured by the 1st level. This intuition is confirmed by the 1st level wavelet variance of Period I which lies outside the 95% confidence interval of Period III, as well. So the ‘more volatile’ outlook of Period I is caused by the different dynamics at levels 1 and 6 corresponding to 5-minute and approximately 3-hour changes, respectively.

As explained in the introduction, the difference in the overall level of volatility can be attributed to specific time-scales that correspond to certain type of players in the market. More precisely, the jumps at the 1st level measure the flow of new information and the general level of nervousness. This is the type of information short-run speculators find valuable. Because most of the big jumps at this level are caused by overnight returns, short-run speculators have probably rebalanced their positions at the market opening(s). The difference at the 6th level is not so easily attributable to any specific group of investors (such as speculators operating at a daily interval), however. This is because the volatility seasonality that is particularly strong in Period I may have affected the scaling law. This possibility will be studied more carefully later (in Sec. 4.6).

As shown earlier (in Sec. 2.3), scaling laws are intimately related to the memory of the DGP. The observed initial rapid decay of the sample autocorrelation followed by a very slow rate of dissipation is characteristic of slowly mean-reverting fractionally integrated processes that exhibit hyperbolic

Figure 5: The wavelet variances of Periods I (cts line) and III (dashed) on a double-logarithmic scale using LA(8) with reflecting boundary. The Gaussian 95% confidence interval (dotted) of Period III has been drawn to address the significance.




Table 4: OLS-regression results.

Levels	Period I			Period III		
	Coeff.	SE	\hat{d}	Coeff.	SE	\hat{d}
2 – 10	−0.6204	0.0502	0.1898	−0.6408	0.0206	0.1796
2 – 6	−0.4443	0.0911	0.2778	−0.5643	0.0440	0.2178
7 – 10	−0.3773	0.0227	0.3114	−0.5777	0.0400	0.2112

rate of decay (ie long-memory).¹⁰ Using Equation (2.2), the estimation of the fractional differencing parameter d is done for Periods I and III by the OLS. The same type of approach has been used by Jensen (2000) and Tkacz (2000), for example. Following Ray and Tsay (2000) the standard errors obtained from regression theory are used to judge the significance. The estimates of d support the conjectured long-memory since they fall in the interval $(0, 1/2)$ (see Table 4). Period I has systematically a slightly larger value than Period III. The coefficients using levels $j_1 = 2, \dots, 6$ and $j_2 = 7, \dots, 10$ in Period III do not differ significantly but the coefficients in Period I may (and are to be discussed later).

The relatively short time-span of the data (approx. 1.5 years) may be criticized. But the recent evidence (eg Andersen and Bollerslev (1997a,b)) suggest that the performance of the long-memory estimate from the volatility series may be greatly enhanced by increasing the observation frequency instead of time-span (in contrast to estimating long-memory dependencies in the mean). Bollerslev and Wright (2000) have for example argued that high-frequency data allows for vastly superior and nearly unbiased estimation of d .

4.5 Local scaling laws and long-memory

The assumption of a constant long-memory structure may not always be reasonable. Bayraktar et al (2003) tackled the problem of time-varying long-memory by segmenting the data before its estimation. But this scheme might not always be sufficient. Whitcher and Jensen (2000) have argued that the ability to estimate local behavior by applying a partitioning scheme to a global estimating procedure is inadequate when compared with an estimator designed to capture time-varying features.

I therefore apply the MODWT-based methodology laid out in Whitcher and Jensen (2000) for estimating the local long-memory parameter $d(t)$. But in contrast to Jensen and Whitcher (2000) who used log-squared returns to proxy volatility, I use absolute returns to prevent an ‘inlier problem’ (taking a logarithm of a number close to zero would generate an outlier). In order to let Equation (2.4) hold, I then implicitly assume that absolute results are generated by a locally stationary process. Considering the jumps and the clustering of volatility, this assumption seems more reasonable than covariance

¹⁰Most ARCH-models exhibit exponential rate of decay and thus fail in this respect (see eg Breidt et al (1998)).

Table 5: Key figures of time-varying long-memory.

Period I						Jarque-Bera		
Min.	1st Q.	Med.	Mean	3rd Q.	Max.	X^2	df	p
-0.3426	0.2084	0.3117	0.3020	0.4056	1.0313	5.5039	2	0.0638
Period III						Jarque-Bera		
Min.	1st Q.	Med.	Mean	3rd Q.	Max.	X^2	df	p
-0.3287	0.1663	0.2505	0.2412	0.3262	0.6624	3.6219	2	0.1635

stationarity. Only the results using the levels 2 – 10 in the OLS-regression are reported here because they gave the most stable results.

The local long-memory parameter estimates of Periods I and III behave similarly (see Figs. 6 and 7): the estimate of $d(t)$ tends to stay in the long-memory interval $(0, 1/2)$ although big jumps in the original series pull the estimate downwards and ‘out of bounds’. Fortunately however, the visits outside the stationary interval of $(-1/2, 1/2)$ are short-lived and the process is mean reverting. Moreover, the estimate stabilizes during less volatile times. For example, a steady increase in the price increases the estimate of long-memory consistently with the definition of long-memory. The median of the local long-memory parameter estimate of Period I is again larger than the median of Period III (see Table 5). The estimates are unconditionally Gaussian which would help the modeling of the behavior of long-memory in time. Although this idea is not yet explored further here, one may try to find a stochastic structure for $d(t)$. One would then have to consider the effect of structural breaks more seriously.¹¹ Here, however, no sign of a structural break is visible in either period. In fact, one of the main reasons for the division of the data to Periods I and III was to avoid this problem altogether.

4.6 Effects of volatility periodicity

On average, the shape of intraday volatility is similar in Periods I and III (see Figs. 8 and 9). After the highly volatile first 5 minutes (when the overnight returns are excluded) the average volatility calms down smoothly and stabilizes. At afternoon hours, however, the behavior of volatility becomes abrupt again. The first peak occurs at 3:35 p.m. and the next one at 4:35. The former peak is most probably due to regular U.S. macro news announcements¹² and the latter is the New York effect. There is also a small but distinct 5-minute peak half an hour later at 5:05 which is probably caused by macro news, too. In Period III the highest peak is at 6:05 p.m. (and right after it) when AMT (I) begins. This is just an artifact of which the 3%-filter (see Sec. 4.1) was unable to remove totally. The last 5 minutes of trading also experience a sudden but small increase in volatility in both periods. In general, then, the

¹¹This is because structural breaks might affect the estimate of $d(t)$ upwards (see Granger and Hyung (1999)). The timing and size of the breaks would then become an equally important issue.

¹²The most important U.S. macro news announcements are released at 8:30 and 10:00 a.m. Eastern Standard Time (see Andersen and Bollerslev (1998) and Bollerslev et al (2000)).

Figure 6: Local long-memory parameter estimates of Period I. Return and price series are plotted below to align the features in time.




Figure 7: Local long-memory parameter estimates of Period III.




Table 6: OLS-regression results of periodicity cleaned volatility.

Levels	Period I			Period III		
	Coeff.	SE	\hat{d}	Coeff.	SE	\hat{d}
2 – 10	−0.4483	0.0334	0.2758	−0.5468	0.0155	0.2266
2 – 6	−0.3326	0.0955	0.3337	−0.6218	0.0069	0.1891
7 – 10	−0.5429	0.0199	0.2286	−0.4692	0.0320	0.2654

average volatility pattern is an ‘inverse-J’.¹³

The wavelet method could be used to annihilate the intraday dependencies. Unfortunately, by considering the interdaily and longer dynamics (ie the wavelet smooth of level $J \geq 6$) as proposed by Gençay et al (2001b), I was not able to reproduce the hyperbolic decay in the sample autocorrelation function of the filtered series. The intraday seasonalities were therefore removed by the *Fourier Flexible Form* (FFF) (see Andersen and Bollerslev (1997c, 1998)), instead. The FFF has been successfully applied in the stock markets previously (see eg Martens et al (2002)). I chose to settle for the minimum number of sinusoids that gave a reasonable fit (3 and 4 for Periods I and III, respectively). The regression results (not shown here) tell us that volatility for the market opening ($n = 1$) and U.S. macro news ($n = 61$) in Period I increased by 3.14 and 1.85 percent, respectively. In Period III the effects were a bit weaker, accounting for 2.25 and 1.45 percent, respectively. So markets reacted more strongly in Period I than in Period III which is not surprising. By comparing the average volatility patterns of Periods I and III, the New York effect is relatively (although not absolutely) a bit larger in Period I, as well.

In order to compare the scaling laws of the periodicity filtered returns to the original series, the overnight returns must first be omitted. This reduces the total energy of the series but the form of the scaling laws remains similar (see the upper subplots of Fig. 10). The removal of the intraday periodicity has smoothed out the kink at the 6th level, however. In Period I, in particular, the wavelet variances at larger than the 6th level have increased considerably. The law is still not totally linear. The first region includes time-scales smaller than an hour and the second one the rest up till 2560 minutes (approx. 43 hours). Almost all of the periodicity-filtered wavelet variances are significantly different from the original (with the overnight returns excluded). In Period III, the change in the distribution of energy across the scales is not that dramatic but the slight kink at the 6th level has disappeared, as well.

The estimates of d increased significantly after the removal of the seasonality in both periods (see Table 6). This is in contrast what Bayraktar et al (2003) have found. They argued the OLS-based wavelet variance estimation to be robust to seasonalities. Here the removal of the relatively stronger intraday periodicity in Period I actually caused a larger change in the estimate of d than in Period III; in Period I the change is from 0.1898 to 0.2758 while in Period III the change is from 0.1796 to 0.2266.

¹³Similar patterns are found in the New York stock markets (see eg Wood et al (1985)) and the FX markets (see eg Andersen and Bollerslev (1997a,c)), although in the latter case the periodicity is associated with the opening and closing of various financial centers around the world.

Figure 8: Average intraday volatility of Period I (cts line) and its FFF fit (dashed).




Figure 9: Average intraday volatility of Period III (cts line) and its FFF fit (dashed).




Figure 10: The scaling laws of the overnight return excluded series (cts line in the top subplots) are below the original (dark), especially in Period I. The removal of the volatility periodicity smooths out the level 6 kink (dark line in the bottom subplots). The Gaussian 95% confidence intervals (dotted thin) are also drawn.

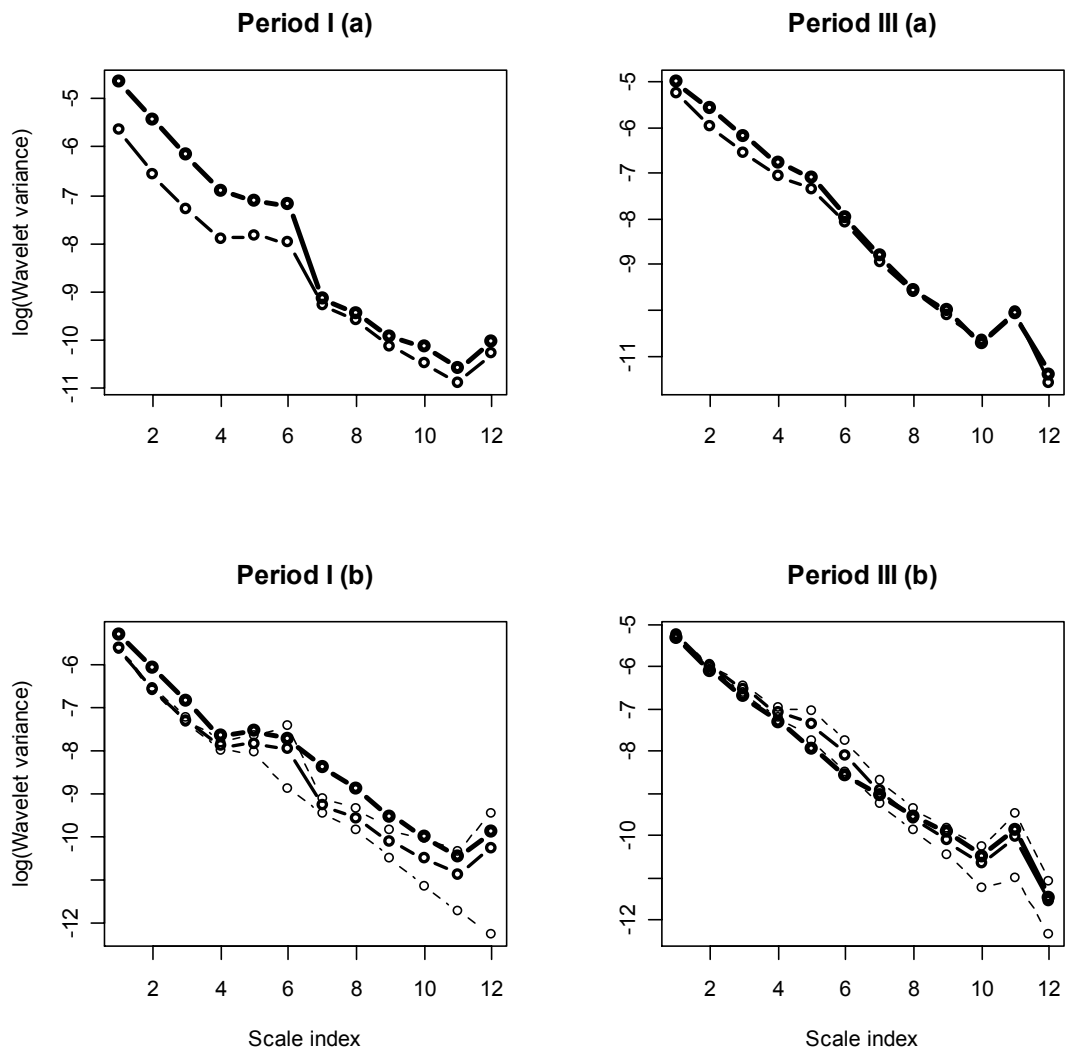


Table 7: Key figures of time-varying long-memory in periodicity cleaned volatility.

Period I						Jarque-Bera		
Min.	1st Q.	Med.	Mean	3rd Q.	Max.	X^2	df	p
-0.3737	0.2591	0.3464	0.3415	0.4300	1.1873	12.322	2	0.0021
Period III						Jarque-Bera		
Min.	1st Q.	Med.	Mean	3rd Q.	Max.	X^2	df	p
-0.3446	0.1992	0.2803	0.2682	0.3481	0.7073	5.3449	2	0.0691

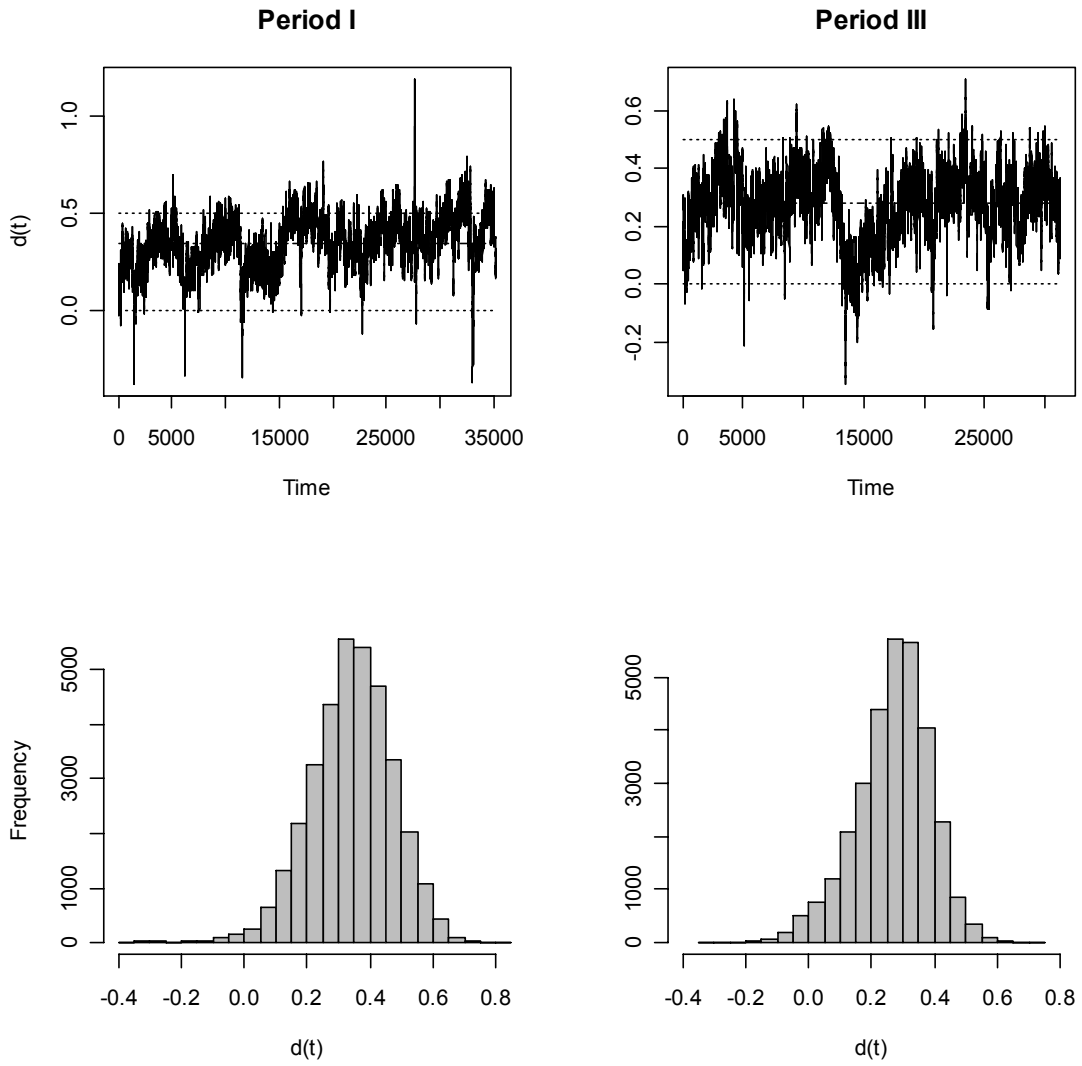
The removal of the intraday periodicity affected the local long-memory estimates, as well. The median of $d(t)$ increased by approximately 0.03 in both periods (see Table 7): in Period I the increase is from 0.3117 to 0.3464 and in Period III from 0.2505 to 0.2803 – a more modest increase than in the global analysis, though. Notice that in Period I the FFF method mistakingly enlarged the amplitude of certain intraday jumps, thus making the range of the estimates of $d(t)$ wider and the unconditional distribution only nearly Gaussian. But in between the jumps the path of $\hat{d}(t)$ became steadier (see the left-hand part of Fig. 11). The 1st and 3rd quantile confirm that the unconditional distribution is more concentrated around the mean value. In Period III the path of $\hat{d}(t)$ has stabilized too (see the right-hand part of Fig. 11). The range of periodicity cleaned returns has there become only a bit wider and the unconditional distribution remains Gaussian which is convenient from the modeling point of view.

5 Conclusions

This paper has studied the time-varying behavior of scaling laws and long-memory in volatility using wavelet analysis. In short, the results show that (i) different scaling regions (multiscaling) may appear in stock markets and not only in the FX markets, (ii) the scaling factor is systematically different from the Brownian square-root, (iii) the scaling factor is not constant in time, and that (iv) the behavior can be explained for a significant part by an intraday volatility periodicity called the New York effect. And because the scaling law is intimately related to the memory of the DGP, stock market volatility exhibits long-memory (as claimed by many before), long-memory varies in time, and that a significant part of this time-variation can be explained by the New York effect. Long-memory appears to have been significantly stronger in the bubble period than its aftermath.

In more detail, global analysis revealed differences between the turbulent ‘IT-bubble’ period and its calmer aftermath period. A break in the scaling law occurs at a daily frequency in the bubble period. This means that traditional time-scale invariant scaling (eg by square-root of time) across relevant time-scales would then have been improper; intraday speculators and longer-term investors faced different type – not just different magnitude – of risks. Because such a break is not found in the aftermath period, the form of the scaling law is able to evolve in time. The interpretation is as follows: The

Figure 11: Unconditional distributions of the estimate of $d(t)$.



bubble period was characterized by stronger overnight and New York effects as a consequence of ‘irrational exuberance’. The frenzy behavior of intraday speculators can for a significant part explain the change in the form of the scaling law. This is because after adjusting for the strong New York effect in the bubble period, the scaling law tends to linearize (not totally, though).

Local analysis revealed that the magnitude (ie not only the shape) of the scaling law is able to evolve in time, as well. The intimate relationship between the scaling law and long-memory then implies that long-memory is not constant over time. Because the local adaptiveness of wavelets allows for the detection of the change points, a locally stationary LMSV model could be applied in modeling. The usefulness of such a model is exemplified by the fact that long-memory was found to be significantly stronger in the bubble period than its aftermath. Before estimating such a model, however, the volatility periodicity should first be taken care of by some suitable method. This is because the removal of periodicity has the tendency to increase and stabilize the local estimate of long-memory. In this paper the FFF was considered an adequate method.

The reason for long-memory in volatility and its time-varying behavior were not addressed in this paper in detail. It is however unlikely that in this particular case the observed long-memory would be caused by structural breaks because of the data division to two periods. It is therefore probable that the bubble period experienced truly stronger long-memory caused by irrational exuberance which, on the other hand, was driven by factors such as loose monetary policy in the U.S. (for a good account on why the markets went crazy, see Lee (2004)). The exact identification of structural breaks should however prove to be useful in order to show this. This would also help the modeling of long-memory in time. Structural breaks can be identified by wavelet methods and it is left for future research.

References

- Andersen, T.G. – Bollerslev, T. (1997a) **Answering the Critics: Yes, ARCH Models Do Provide Good Volatility Forecasts.** NBER Working Paper 6023.
- Andersen, T.G. – Bollerslev, T. (1997b) **Heterogenous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-run in High Frequency Returns.** Journal of Finance 52, 975–1005.
- Andersen, T.G. – Bollerslev, T. (1997c) **Intraday Periodicity and Volatility Persistence in Financial Markets.** Journal of Empirical Finance 4, 115–158.
- Andersen, T.G. – Bollerslev, T. (1998) **Deutsche Mark-Dollar Volatility: Intraday Activity Patterns, Macroeconomic Announcements, and Longer Run Dependencies.** Journal of Finance 53, 219–265.
- Andersen, T.G. – Bollerslev, T. – Diebold, F.X. – Labys, P. (2000) **The Distribution of Realized Exchange Rate Volatility.** NBER Working Paper 6961.
- Baillie, R.T. (1996) **Long Memory Processes and Fractional Integration in Econometrics.** Journal of Econometrics 73, 5–59.
- Bayraktar, E. – Poor, H.V. – Sircar, K.R. (2000) **Estimating the Fractal Dimension of the S&P 500 Index Using Wavelet Analysis.** Submitted manuscript, Princeton University.
- Beran, J. (1994) **Statistics for Long Memory Processes.** Volume 61 of Monographs on Statistics and Applied Probability. Chapman and Hall, New York.
- Bollerslev, T. (1986) **Generalized Autoregressive Conditional Heteroskedasticity.** Journal of Econometrics 31, 307–327
- Bollerslev, T. – Chou, R.Y. – Kroner, K.F. (1992) **ARCH Modeling in Finance; A Review of the Theory and Empirical Evidence.** Journal of Econometrics 52, 5–59.
- Bollerslev, T. – Cai, J. – Song, F.M. (2000) **Intraday Periodicity, Long Memory Volatility, and Macroeconomic Announcement Effects in the US Treasury Bond Market.** Journal of Empirical Finance 7, 37–55.
- Bollerslev, T. – Wright, J.H. (2000) **Semiparametric Estimation of Long-Memory Volatility Dependencies: The Role of High-Frequency Data.** Journal of Econometrics 98, 81–106.
- Breidt, F.J. – Crato, N. – De Lima, P. (1998) **The Detection and Estimation of Long Memory in Stochastic Volatility.** Journal of Econometrics 83, 325–348.

- Campbell, J.Y. – Lo, A.W. – MacKinlay, A.C. (1997) **The Econometrics of Financial Markets**. Second corrected printing, Princeton University Press.
- Crowley, P.M. (2005) **An Intuitive Guide to Wavelets for Economists**. Bank of Finland Research. Discussion Papers No. 1.
- Cont, R. (2001) **Empirical Properties of Asset Returns: Stylized Facts and Statistical Issues**. *Quantitative Finance* 1, 223–236.
- Dahlhaus, R. (1997) **Fitting Time Series Models to Nonstationary Processes**. *The Annals of Statistics* 25, 1–37.
- Diebold, F.X. – Hickman, A. – Inoue, A. – Schuermann, T. (1997) **Converting 1-Day Volatility to h-Day Volatility: Scaling by \sqrt{h} is Worse than You Think**. Wharton Financial Institutions Center Working Paper.
- Engle, R.F. (1982) **Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation**. *Econometrica* 50, 987–1008.
- Engle, R.F. – Ito, T. – Lin, W-L. (1990) **Meteor Showers or Heat Waves? Heteroskedastic Intra Daily Volatility in the Foreign Exchange Market**. *Econometrica* 58, 525–542.
- Fisher, A.J. – Calvet, L.E. – Mandelbrot, B.B. (1997) **Multifractality of Deutschemark / US Dollar Exchange Rates**. Cowles Foundation Discussion Paper 1166.
- Gençay, R. – Seluk, F. – Whitcher, B. (2001a) **Scaling Properties of Foreign Exchange Volatility**. *Physica A* 289, 249–266.
- Gençay, R. – Selçuk, F. – Whitcher, B. (2001b) **Differentiating Intraday Seasonalities Through Wavelet Multi-Scaling**. *Physica A* 289, 543–556.
- Gençay, R. – Seluk, F. – Whitcher, B. (2002) **An Introduction to Wavelets and Other Filtering Methods in Finance and Economics**. Academic Press.
- Geweke, J. – Porter-Hudak, S. (1983) **The Estimation and Application of Long Memory Time Series Models**. *Journal of Time Series Analysis* 4, 221–238.
- Granger, C.W.J. – Joyeux, R. (1980) **An Introduction to Long Memory Time Series Models and Fractional Differencing**. *Journal of Time Series Analysis* 1, 15–29.
- Granger, C.W.J. – Hyung, N. (1999) **Occasional Structural Breaks and Long Memory**. University of California San Diego Discussion Paper 99-14.
- Guillaume, D.M. – Dacorogna, M.M. – Dave, R.D. – Müller, U.A. – Olsen, R.B. – Pictet, O.V. (1997) **From the Bird’s Eye to the Microscope: a Survey of New Stylized Facts of the Intra-Daily Foreign Exchange Markets**. *Finance and Stochastics* 1, 95–129.

- Harvey, A. – Ruiz, E. – Shephard, N. (1994) **Multivariate Stochastic Variance Models**. *Review of Economic Studies* 61, 247–264.
- Hosking, J.R.M. (1981) **Fractional Differencing**. *Biometrika* 68, 165–176.
- Härdle, W. – Kerkycharian, G. – Picard, D. – Tsybakov, A. (1998) **Wavelets, Approximation, and Statistical Applications**. Springer, New York.
- Jensen, M.J. (1999) **Using Wavelets to Obtain a Consistent Ordinary Least Squares Estimator of the Long-Memory Parameter**. *Journal of Forecasting* 18, 17–32.
- Jensen, M.J. (2000) **An Alternative Maximum Likelihood Estimator of Long-Memory Processes Using Compactly Supported Wavelets**. *Journal of Economic Dynamics and Control* 24, 361–387.
- Jensen, M.J. – Whitcher, B. (2000) **Time-Varying Long-Memory in Volatility: Detection and Estimation with Wavelets**. Technical report, University of Missouri and EURANDOM.
- Lee, T. (2004) **Why the Markets Went Crazy: And What it Means for Investors**. Palgrave Macmillan.
- Martens, M. – Chang, Y.-C. – Taylor, S.J. (2002) **A Comparison of Seasonal Adjustment Methods when Forecasting Intraday Volatility**. *Journal of Financial Research* XXV, 283–299.
- Mikosch, T. – Stărică, C. (2004) **Non-Stationarities in Financial Time Series, the Long Range Dependence and the IGARCH Effects**. Chalmers University of Technology, Submitted.
- Müller, U.A. – Dacorogna, M.M. – Olsen, R.B. – Pictet, O.V. – Schwarz, M. – Morgenegg, C. (1990) **Statistical Study of Foreign Exchange Rates, Empirical Evidence of a Price Change Scaling Law, and Intraday Analysis**. *Journal of Banking and Finance* 14, 1189–1208.
- Müller, U.A. – Dacorogna, M.M. – Dave, R.D. – Pictet, O.V. – Ward, R.B. (1993) **Fractals and Intrinsic Time, a Challenge to Econometricians**. Olsen Ltd. Working Paper.
- Müller, U.A. – Dacorogna, M.M. – Davé, R.D. – Olsen, R.B. – Pictet, O.V. – von Weizsäcker, J.E. (1997) **Volatilities of Different Time Resolutions – Analyzing the Dynamics of Market Components**. *Journal of Empirical Finance* 4, 213–239.
- Nason, G.P. – von Sachs, R. (1999) **Wavelets in Time Series Analysis**. *Philosophical Transactions of the Royal Society of London, Series A* 357, 2511–2526.
- Percival, D.B. – Walden, A.T. (2000) **Wavelet Methods for Time Series Analysis**. Cambridge University Press.

- Pollock, D.S.G. – Lo Cascio, I. (2004) **Adapting Discrete Wavelet Analysis to the Circumstances of Economics**. Manuscript, Queen Mary, University of London.
- Polzehl, J. – Spokoiny, V. – Stărică, C. (2004) **When Did the 2001 Recession Really Start?** Chalmers University of Technology, Submitted.
- Priestley, M.B. (1996) **Wavelets and Time-Dependent Spectral Analysis**. Journal of Time Series Analysis 17, 85–103.
- Ray, B.K. – Tsay, R.S. (2000) **Long-Range Dependence in Daily Stock Volatilities**. Journal of Business and Economic Statistics 18, 254–262.
- Shiller, R.J. (2001) **Irrational Exuberance**. Broadway Books, New York.
- Schleicher, C. (2002) **An Introduction to Wavelets for Economists**. Monetary and Financial Analysis Department, Bank of Canada, Working paper 2002-3.
- Tkacz, G. (2000) **Estimating the Fractional Order of Integration of Interest Rates Using a Wavelet OLS Estimator**. Bank of Canada Working Paper 2000-5.
- Vuorenmaa, T. (2004) **A Multiresolution Analysis of Stock Market Volatility Using Wavelet Methodology**. Licentiate Thesis, University of Helsinki.
- Whitcher, B. – Jensen, M.J. (2000) **Wavelet Estimation of a Local Long Memory Parameter**. Exploration Geophysics 31, 94–103.
- Wood, R.A. – McInish, T.H. – Ord, J.K. (1985) **An Investigation of Transaction Data for NYSE Stocks**. Journal of Finance 40, 723–739.

Appendix

Fractional differencing and long-memory

A real-valued discrete parameter *fractional* ARIMA (ARFIMA) process $\{X_t\}$ is often defined with a binomial series expansion,

$$(1 - B)^d X_t \doteq \sum_{k=0}^{\infty} \binom{d}{k} (-1)^k X_{t-k},$$

where B is the lag operator and $\binom{a}{b} \doteq \frac{a!}{b!(a-b)!} = \frac{\Gamma(a+1)}{\Gamma(b+1)\Gamma(a-b+1)}$. These models were originally introduced by Granger and Joyeux (1980) and Hosking (1981).

In ARFIMA models the ‘long-memory’ dependency is characterized solely by the *fractional differencing parameter* d . A time series is said to exhibit long-memory when it has a covariance function $\gamma(j)$ and a spectrum $f(\lambda)$ such that they are of the same order as j^{2d-1} and λ^{-2d} , as $j \rightarrow \infty$ and $\lambda \rightarrow 0$, respectively. For $0 < d < 1/2$, an ARFIMA model exhibits long-memory, and for $-1/2 < d < 0$ it exhibits antipersistence. In practice, the range $|d| < 1/2$ is of particular interest because then an ARFIMA model is stationary and invertible (Hosking (1981)).

Details of long-memory can be found in Beran (1994), for example. Concerning fractionally integrated processes in econometrics, see Baillie (1996).

**BANK OF FINLAND RESEARCH
DISCUSSION PAPERS**

ISSN 0785-3572, print; ISSN 1456-6184, online

- 1/2005 Patrick M. Crowley **An intuitive guide to wavelets for economists.** 2005. 68 p. ISBN 952-462-188-6, print; ISBN 952-462-189-4, online.
- 2/2005 Tuomas Saarenheimo **Ageing, interest rates, and financial flows.** 2005. 30 p. ISBN 952-462-190-8, print; ISBN 952-462-191-6, online.
- 3/2005 Timo Vesala **Relationship lending and competition: Higher switching cost does not necessarily imply greater relationship benefits.** 2005. 27 p. ISBN 952-462-192-4, print; ISBN 952-462-193-2, online.
- 4/2005 Johanna Sinkkonen **Labour productivity growth and industry structure. The impact of industry structure on productivity growth, export prices and labour compensation.** 2005. 88 p. ISBN 952-462-196-7, print; ISBN 952-462-197-5, online.
- 5/2005 Matti Virén **Why do capital intensive companies pay higher wages?** 2005. 23 p. ISBN 952-462-198-3, print; ISBN 952-462-199-1, online.
- 6/2005 Maritta Paloviita – Matti Virén **The role of expectations in the inflation process in the euro area.** 2005. 33 p. ISBN 952-462-200-9, print; ISBN 952-462-201-7, online.
- 7/2005 Marketta Henriksson **Productivity differentials and external balance in ERM II.** 2005. 32 p. ISBN 952-462-202-5, print; ISBN 952-462-203-3, online.
- 8/2005 Matti Virén **Government size and output volatility: is there a relationship?** 2005. 24 p. ISBN 952-462-204-1, print; ISBN 952-462-205-X, online.
- 9/2005 Karlo Kauko **Bank interest rates in a small European economy: Some exploratory macro level analyses using Finnish data.** 2005. 44 p. ISBN 952-462-206-8, print; ISBN 952-462-207-6, online.
- 10/2005 Essi Eerola – Niku Määttänen **The optimal tax treatment of housing capital in the neoclassical growth model.** 2005. 30 p. ISBN 952-462-210-6, print; ISBN 952-462-211-4, online.
- 11/2005 Karlo Kauko **The mixed oligopoly of cross-border payment systems.** 2005. 31 p. ISBN 952-462-212-2, print; ISBN 952-462-213-0, online.

- 12/2005 Patrick M. Crowley – Jim Lee **Decomposing the co-movement of the business cycle: a time-frequency analysis of growth cycles in the euro area.** 2005. 67 p. ISBN 952-462-214-9, print; ISBN 952-462-215-7, online.
- 13/2005 Jarmo Pesola **Banking fragility and distress: An econometric study of macroeconomic determinants.** 2005. 95 p. ISBN 952-462-216-5, print; ISBN 952-462-217-3, online.
- 14/2005 Karlo Kauko **The demand for money market mutual funds.** 2005. 29 s. ISBN 952-462-218-1, print; ISBN 952-462-219-X, online.
- 15/2005 Matti Virén **Fiscal policy in the 1920s and 1930s. How much different it is from the post war period's policies?** 2005. 27 p. ISBN 952-462-220-3, print; ISBN 952-462-221-1, online.
- 16/2005 Alistair Milne **What's in it for us? Network effects and bank payment innovation.** 2005. 27 p. ISBN 952-462-223-8, print; ISBN 952-462-224-6, online.
- 17/2005 Hilde C Bjørnland – Kai Leitemo **Identifying the interdependence between US monetary policy and the stock market.** 2005. 43 p. ISBN 952-462-225-4, print; ISBN 952-462-226-2, online.
- 18/2005 Juha Tarkka – Mika Kortelainen **International economic spillovers and the liquidity trap.** 2005. 86 p. ISBN 952-462-227-0, print; ISBN 952-462-228-9, online.
- 19/2005 Kari Kemppainen **Assessing effects of price regulation in retail payment systems.** 2005. 32 p. ISBN 952-462-231-9, print; ISBN 952-462-232-7, online.
- 20/2005 Kai Leitemo – Ulf Söderström **Robust monetary policy in a small open economy.** 2005. 53 p. ISBN 952-462-233-5, print; ISBN 952-462-234-3, online.
- 21/2005 Hanna-Leena Männistö **Forecasting with a forward-looking DGE model – combining long-run views of financial markets with macro forecasting.** 2005. 52 p. ISBN 952-462-235-1, print; ISBN 952-462-236-X, online.
- 22/2005 Maritta Paloviita **Comparing alternative Phillips curve specifications: European results with survey-based expectations.** 2005. 26 p. ISBN 952-462-237-8, print; ISBN 952-462-238-6, online.
- 23/2005 Alistair Milne **Standard setting and competition in securities settlement.** 2005. 29 p. ISBN 952-462-239-4, print; ISBN 952-462-240-8, online.

- 24/2005 Yehning Chen – Iftekhar Hasan **The transparency of the banking industry and the efficiency of information-based bank runs.** 2005. 38 p. ISBN 952-462-241-6, print; ISBN 952-462-242-4, online.
- 25/2005 Matti Virén **Inflation expectations and regime shifts in the euro area.** 2005. 23 p. ISBN 952-462-243-2, print; ISBN 952-462-244-0, online.
- 26/2005 Mikko Puhakka **The effects of aging population on the sustainability of fiscal policy.** 2005. 29 p. ISBN 952-462-251-3, print; ISBN 952-462-252-1, online.
- 27/2005 Tommi A. Vuorenmaa **A wavelet analysis of scaling laws and long-memory in stock market volatility.** 2005. 39 p. ISBN 952-462-253-X, print; ISBN 952-462-254-8, online.

Suomen Pankki
Bank of Finland
P.O.Box 160
FI-00101 HELSINKI
Finland



.2343