Long-Memory Analysis

Gilles Teyssiere

Long-memory in economics and finance is an important research topic as several economic variables exhibit the main characteristics of long-memory processes, i.e., a significant dependence between very distant observations and a pole in the neighborhood of the zero frequency of their spectrum. In particular, returns on financial assets are uncorrelated, while the series of absolute and squared returns display long-memory.

Long-memory in finance is still an empirical research topic. A structural microeconomic model based on interacting agents generating long-memory properties has been proposed by Kirman and Teyssière (1998). Statistical tools for measuring long-memory which only depend on weak assumptions on the data generating process are emerging in the research literature. This chapter focuses on new results on semiparametric tests and estimators.

All quantlets for long-memory analysis are contained in the quantlib `times` and become available after entering the instruction

```
library("times")
```

1 Introduction

`d = gph(y)`
: estimates the degree of long-memory of a time series by using a log-periodogram regression

`k = kpss(y)`
: calculates the KPSS statistics for $I(0)$ against long-memory alternatives

```
q = lo(y{, m})
```
calculates the Lo Statistics for long-range dependence

```
t = lobrob(y{, bdvec})
```
provides a semiparametric test for $I(0)$ of a time series

```
q = neweywest(y{, m})
```
calculates the Newey and West heteroskedastic and autocorrelation consistent estimator of the variance

```
d = roblm(x{, q, bdvec})
```
semiparametric average periodogram estimator of the degree of long-memory of a time series

```
d = robwhittle(x{, bdvec})
```
semiparametric Gaussian estimator of the degree of long-memory of a time series, based on the Whittle estimator

```
k = rvlm(x{, m})
```
calculates the rescaled variance test for $I(0)$ against long-memory alternatives

A stationary stochastic process $\{Y_t\}$ is called a **long-memory process** if there exist a real number H and a finite constant C such that the autocorrelation function $\rho(\tau)$ has the following rate of decay:

$$\rho(k) \sim C\tau^{2H-2} \quad \text{as} \quad \tau \to \infty. \tag{1}$$

The parameter H, called the **Hurst exponent**, represents the long-memory property of the time series. A long-memory time series is also said **fractionally integrated**, where the fractional degree of integration d is related to the parameter H by the equality $d = H - 1/2$. If $H \in (1/2, 1)$, i.e., $d \in (0, 1/2)$, the series is said to have **long-memory**. If $H > 1$, i.e., $d > 1/2$, the series is **nonstationary**. If $H \in (0, 1/2)$, i.e., $d \in (-1/2, 0)$, the series is called **antipersistent**.

Equivalently, a long-memory process can be characterized by the behaviour of its spectrum $f(\lambda_j)$, estimated at the harmonic frequencies $\lambda_j = 2\pi j/T$, with $j = 1, \ldots, [T/2]$, near the zero frequency:

$$\lim_{\lambda_j \to 0^+} f(\lambda_j) = C\lambda_j^{-2d} \tag{2}$$

where C is a strictly positive constant. Excellent and exhaustive surveys on long-memory are given in Beran (1994), Robinson (1994a) and Robinson and Henry (1998).

A long-memory process with degree of long-memory d is said to be integrated of order d and is denoted by $I(d)$. The class of long-memory processes generalises the class of integrated processes with integer degree of integration.

2 Model Indepependent Tests for $I(0)$ against $I(d)$

A stochastic process is $I(d)$ if it needs to be differentiated d times in order to become $I(0)$. We shall test for $I(0)$ against fractional alternatives by using more formal definitions.

In a first approach, we define a stochastic process $\{Y_t\}$ as $I(0)$ if the normalized partial sums follow a particular distribution. We only require is the existence of a consistent estimator of the variance for normalizing the partial sums. The tests presented here make use of the Newey and West (1987) heteroskedastic and autocorrelation consistent (HAC) estimator of the variance, defined as

$$\hat{\sigma}_T^2(q) = \hat{\gamma}_0 + 2\sum_{j=1}^{q} \left(1 - \frac{j}{1+q}\right) \hat{\gamma}_j, \quad q < T, \tag{3}$$

where $\hat{\gamma}_0$ is the variance of the process, and the sequence $\{\hat{\gamma}_j\}_{j=1}^{q}$ denotes the autocovariances of the process up to the order q. This spectral based HAC variance estimator depends on the user chosen truncation lag q. Andrews (1991) has proposed a selection rule for the order q.

The quantlet `neweywest` computes the Newey and West (1987) estimator of the variance of a unidimensional process. Its syntax is:

```
sigma = neweywest(y{, q})
```

where the input parameters are:

`y`
: the series of observations

`q`

optional parameter, which can be either a vector of truncation lags or a single scalar

The HAC estimator is calculated for all the orders included in the parameter `q`. If no optional parameter is provided, the HAC estimator is evaluated for the default orders `q` = 5, 10, 25, 50. The estimated HAC variances are stored in the vector `q`.

In the following example the HAC variance of the first 2000 observations of the 20 minutes spaced sample of Deutschmark-Dollar FX is computed.

```
library("times")
y = read("dmus58.dat")
y = y[1:2000]
q = 5|10|25|50
sigma = neweywest(y,q)
q~sigma
```

`longmem01.xpl`

As an output we get

```
Contents of _tmp

[1,]        5  0.0047841
[2,]       10  0.008743
[3,]       25  0.020468
[4,]       50  0.039466
```

2.1 Robust Rescaled Range Statistic

The first test for long-memory was devised by the hydrologist Hurst (1951) for the design of an optimal reservoir for the Nile river, of which flow regimes were persistent. Although Mandelbrot (1975) gave a formal justification for the use of this test, Lo (1991) demonstrated that this statistic was not robust to short range dependence, and proposed the following one:

$$Q_T = \frac{1}{\hat{\sigma}_T(q)} \left[\max_{1 \le k \le T} \sum_{j=1}^{k} (X_j - \overline{X}_T) - \min_{1 \le k \le T} \sum_{j=1}^{k} (X_j - \overline{X}_T) \right] \tag{4}$$

which consists of replacing the variance by the HAC variance estimator in the denominator of the statistic. If $q = 0$, Lo's statistic reduces to Hurst's R/S statistic. Unlike spectral analysis which detects periodic cycles in a series, the R/S analysis has been advocated by Mandelbrot for detecting nonperiodic cycles. Under the null hypothesis of no long-memory, the statistic $T^{-\frac{1}{2}}Q_n$ converges to a distribution equal to the range of a Brownian bridge on the unit interval:

$$\max_{0\leq t\leq 1} W^0(t) - \min_{0\leq t\leq 1} W^0(t),$$

where $W^0(t)$ is a Brownian bridge defined as $W^0(t) = W(t) - tW(1)$, $W(t)$ being the standard Brownian motion. The distribution function is given in Siddiqui (1976), and is tabulated in Lo (1991).

This statistic is extremely sensitive to the order of truncation q but there is no statistical criteria for choosing q in the framework of this statistic. Andrews (1991) rule gives mixed results. If q is too small, this estimator does not account for the autocorrelation of the process, while if q is too large, it accounts for any form of autocorrelation and the power of this test tends to its size. Given that the power of a useful test should be greater than its size, this statistic is not very helpful. For that reason, Teverovsky, Taqqu and Willinger (1999) suggest to use this statistic with other tests.

Since there is no data driven guidance for the choice of this parameter, we consider the default values for $q = 5$, 10, 25, 50. XploRe users have the option to provide their own vector of truncation lags.

Let's consider again the series of absolute returns on the 20 minutes spaced Deutschmark-Dollar FX rates.

```
library("times")
y = read("dmus58.dat")
ar = abs(tdiff(y[1:2000]))
lostat = lo(ar)
lostat
```

longmem02.xpl

Given that we do not provide a vector of truncation lags, Lo's statistic is computed for the default truncation lags. The results are displayed in the form of a table: the first column contains the truncation orders, the second columns contains the computed statistic. If the computed statistic is outside the 95% confidence interval for no long-memory, a star $*$ is displayed after that statistic.

```
Contents of lostat

[1,] " Order   Statistic"
[2,] "__________________ "
[3,] ""
[4,] "     5      2.0012 *"
[5,] "    10      1.8741 *"
[6,] "    25      1.7490 "
[7,] "    50      1.6839 "
```

This result illustrates the issue of the choice of the bandwidth parameter `q`. For `q` = 5 and 10, we reject the null hypothesis of no long-memory. However, when `q` = 25 or 50, this null hypothesis is accepted, as the power of this test is too low for these levels of truncation orders.

2.2 The KPSS Statistic

Equivalently, we can test for $I(0)$ against fractional alternatives by using the KPSS test Kwiatkowski, Phillips, Schmidt and Shin (1992), as Lee and Schmidt (1996) have shown that this test has a power equivalent to Lo's statistic against long-memory processes. The two KPSS statistics, denoted by η_t and η_μ, are respectively based on the residuals of two regression models: on an intercept and a trend t, and on a constant μ. If we denote by S_t the partial sums $S_t = \sum_{i=1}^{t} \hat{e}_i$, where $\hat{e}_t$ are the residuals of these regressions, the KPSS statistic is defined by:

$$\eta = T^{-2} \sum_{t=1}^{T} S_t^2 / \hat{\sigma}_T^2(q) \tag{5}$$

where $\hat{\sigma}_T^2(q)$ is the HAC estimator of the variance of the residuals defined in equation (3). The statistic η_μ tests for stationarity against a long-memory alternative, while the statistic η_t tests for trend-stationarity against a long-memory alternative.

The quantlet `kpss` computes both statistics. The default bandwidths, denoted by L_0, L_4 and L_{12} are the one given in Kwiatkowski, Phillips, Schmidt and Shin (1992). We evaluate both tests on the series of absolute returns `ar` as follows:

```
library("times")
y = read("dmus58.dat")
ar = abs(tdiff(y[1:2000]))
kpsstest = kpss(ar)
kpsstest
```

longmem03.xpl

The quantlet `kpss` returns the results in the form of a table. The first column contains the truncation order, the second column contains the type of the test: `const` means the test for stationary sequence, while `trend` means the test for trend stationarity. The third column contains the computed statistic. If this statistic exceeds the 95% critical value, a $*$ symbol is displayed. The last column contains this critical value.

Thus, XploRe returns the following table:

```
 Contents of kpsstest

[1,] "   Order   Test   Statistic  Crit. Value "
[2,] "_________________________________________ "
[3,] ""
[4,] " L0 =  0   const    1.8259 *   0.4630"
[5,] " L4 =  8   const    1.2637 *   0.4630"
[6,] " L12= 25   const    1.0483 *   0.4630"
[7,] " L0 =  0   trend    0.0882    0.1460"
[8,] " L4 =  8   trend    0.0641    0.1460"
[9,] " L12= 25   trend    0.0577    0.1460"
```

2.3 The Rescaled Variance V/S Statistic

Giraitis, Kokoszka and Leipus (1998) have proposed a centering of the KPSS statistic based on the partial sum of the deviations from the mean. They called it a rescaled variance test V/S as its expression given by

$$V/S = \frac{1}{T^2 \hat{\sigma}_T^2(q)} \left[\sum_{k=1}^{T} \left(\sum_{j=1}^{k} (Y_j - \overline{Y}_T) \right)^2 - \frac{1}{T} \left(\sum_{k=1}^{T} \sum_{j=1}^{k} (Y_j - \overline{Y}_T) \right)^2 \right] \quad (6)$$

can equivalently be rewritten as

$$V/S = T^{-1}\frac{\hat{V}(S_1, \dots, S_T)}{\hat{\sigma}_T^2(q)}, \tag{7}$$

where $S_k = \sum_{j=1}^{k}(Y_j - \overline{Y}_n)$ are the partial sums of the observations. The V/S statistic is the sample variance of the series of partial sums $\{S_t\}_{t=1}^{T}$. The limiting distribution of this statistic is a Brownian bridge of which the distribution is linked to the Kolmogorov statistic. This statistic has uniformly higher power than the KPSS, and is less sensitive than the Lo statistic to the choice of the order q. For $2 \leq q \leq 10$, the V/S statistic can appropriately detect the presence of long-memory in the levels series, although, like most tests and estimators, this test may wrongly detect the presence of long-memory in series with shifts in the levels. Giraitis, Kokoszka and Leipus (1998) have shown that this statistic can be used for the detection of long-memory in the volatility for the class of ARCH(∞) processes.

We evaluate the V/S statistic with the quantlet `rvlm` which has the following syntax:

```
vstest = rvlm(ary{, q})
```

where

`ary`
: is the series

`q`
: is a vector of truncation lags. If this optional argument is not provided, then the default vector of truncation lags is used, with `q` = 0, 8, 25.

This quantlet returns the results in the form of a table: the first column contains the order of truncation q, the second column contains the estimated V/S statistic. If this statistic is outside the 95% confidence interval for no long-memory, a star $*$ symbol is displayed. The fourth column displays the 95% critical value. Thus the instruction

```
library("times")
y = read("dmus58.dat")
ar = abs(tdiff(y[1:2000]))
vstest = rvlm(ar)
vstest
```

longmem04.xpl

returns

```
Contents of vstest

[1,] "   Order  Statistic  Crit. Value "
[2,] "_________________________________"
[3,] ""
[4,] "      0    0.3305 *   0.1869"
[5,] "      8    0.2287 *   0.1869"
[6,] "     25    0.1897 *   0.1869"
```

2.4 Nonparametric Test for $I(0)$

Lobato and Robinson (1998) nonparametric test for $I(0)$ against $I(d)$ is also based on the approximation (2) of the spectrum of a long-memory process. In the univariate case, the t statistic is equal to:

$$t = m^{1/2}\hat{C}_1/\hat{C}_0 \quad \text{with} \quad \hat{C}_k = m^{-1}\sum_{j=1}^{m}\nu_j^k I(\lambda_j) \quad \text{and} \quad \nu_j = \ln(j) - \frac{1}{m}\sum_{i=1}^{m}\ln(i), \tag{8}$$

where $I(\lambda) = (2\pi T)^{-1}|\sum_{t=1}^{T} y_t e^{it\lambda}|^2$ is the periodogram estimated for a degenerate band of Fourier frequencies $\lambda_j = 2\pi j/T, j = 1, \ldots, m \ll [T/2]$, where m is a bandwidth parameter. Under the null hypothesis of a $I(0)$ time series, the t statistic is asymptotically normally distributed. This two sided test is of interest as it allows to discriminate between $d > 0$ and $d < 0$: if the t statistic is in the lower fractile of the standardized normal distribution, the series exhibits long-memory whilst if the series is in the upper fractile of that distribution, the series is antipersistent.

The quantlet `lobrob` evaluates the Lobato-Robinson test. Its syntax is as follows:

```
l = lobrob(ary{, m})
```

where

`ary`
: is the series,

`m`
: is the vector of bandwidth parameters. If this optional argument is missing, the default bandwidth suggested by Lobato and Robinson is used.

The results are displayed in the form of a table: the first column contains the value of the bandwidth parameter while the second column displays the corresponding statistic. In the following example, the Lobato-Robinson statistic is evaluated by using this default bandwidth:

```
library("times")
y = read("dmus58.dat")
ar = abs(tdiff(y[1:2000]))
l = lobrob(ar)
l
```

`longmem05.xpl`

which yields

```
Contents of l

[1,] "Bandwidth   Statistic "
[2,] "_____________________ "
[3,] ""
[4,] "   334      -4.4571"
```

In the next case, we provide a vector of bandwidths `m`, and evalutate this statistic for all the elements of `m`. The sequence of instructions:

```
library("times")
y = read("dmus58.dat")
```

```
ar = abs(tdiff(y[1:2000]))
m = #(100,150,200)
l = lobrob(ar,m)
l
```

longmem06.xpl

returns the following table:

```
Contents of l

[1,] "Bandwidth   Statistic "
[2,] "______________________ "
[3,] ""
[4,] "   100      -1.7989"
[5,] "   150      -2.9072"
[6,] "   200      -3.3308"
```

3 Semiparametric Estimators in the Spectral Domain

These estimators are based on the behaviour of the spectrum of a long-memory time series near the zero frequency, and are estimated in the frequency band $(0, m]$, where m is a bandwidth parameter less than or equal to $[n/2]$, where [.] denotes the integer part operator. The idea is that the range of frequencies between zero and m captures the long term component, whilst the remainder of the frequencies capture the local variations which could be linear or nonlinear. These estimators are denoted semiparametric in the sense that they depend on a bandwidth parameter m.

3.1 Log-periodogram Regression

Under the assumption of normality, Geweke, and Porter-Hudak (1983) assumed that the spectrum $f(\lambda)$ near the zero frequency can be approximated by

$$f(\lambda) = C\{(4\sin^2(\lambda_j/2)\}^{-d} \tag{9}$$

and then propose to estimate the long-memory parameter d with the following spectral regression:

$$\log\{I(\lambda_j)\} = c - d\log\{4\sin^2(\lambda_j/2)\} + \varepsilon_j \tag{10}$$

where n is the sample size. We consider for this estimator only harmonic frequencies λ_j, with $j \in (l, m]$, where l is a trimming parameter discarding the lowest frequencies and m is a bandwidth parameter.

```
library("times")
y = read("dmus58.dat")
ar = abs(tdiff(y[1:2000]))
d = gph(ar)
d
```

longmem07.xpl

We obtain the following output:

```
Contents of d

[1,]  0.088369
```

3.1.1 Average periodogram estimator

The Robinson (1994b) averaged periodogram estimator is defined by:

$$\hat{d} = \frac{1}{2} - \frac{\ln\left(\hat{F}(q\lambda_m)/\hat{F}(\lambda_m)\right)}{2\ln(q)}, \tag{11}$$

where $\hat{F}(\lambda)$ is the average periodogram

$$\hat{F}(\lambda) = \frac{2\pi}{n}\sum_{j=1}^{[n\lambda/2\pi]} I(\lambda_j). \tag{12}$$

By construction, the estimated parameter $\hat{d}$ is $< 1/2$, i.e., is in the stationarity range. This estimator has the following asymptotic distribution if $d < 1/4$

$$\sqrt{m}(\hat{d} - d) \sim \mathrm{N}\left(0, \frac{\pi^2}{24}\right) \tag{13}$$

We evaluate the degree of long-memory with this estimator as follows:

```
library("times")
y = read("dmus58.dat")
ar = abs(tdiff(y[1:2000]))
d = roblm(ar)
d
```

longmem08.xpl

We obtain the following output:

```
Contents of d

[1,] "     d     Bandwidth     q   "
[2,] "______________________________"
[3,] ""
[4,] "  0.0927      500       0.5"
[5,] "  0.1019      250       0.5"
[6,] "  0.1199      125       0.5"
```

3.2 Semiparametric Gaussian Estimator

The Robinson (1995a) semiparametric estimator, suggested by Künsch (1987), is based on the approximation (2) of the spectrum of a long-memory process in the Whittle approximate maximum likelihood estimator. An estimator of the fractional degree of integration d is obtained by solving the minimization problem:

$$\{\hat{C}, \hat{d}\} = \arg\min_{C,d} L(C,d) = \frac{1}{m}\sum_{j=1}^{m}\left\{\ln\left(C\lambda_j^{-2d}\right) + \frac{I(\lambda_j)}{C\lambda_j^{-2d}}\right\}, \tag{14}$$

where $I(\lambda_j)$ is evaluated for a range of harmonic frequencies $\lambda_j = 2\pi j/n$, $j = 1, \ldots, m \ll [n/2]$ bounded by the bandwidth parameter m, which increases with the sample size n but more slowly: the bandwidth m must satisfy

$$\frac{1}{m} + \frac{m}{n} \to 0 \quad \text{as} \quad n \to \infty. \tag{15}$$

If $m = n/2$, this estimator is a Gaussian estimator for the parametric model $f(\lambda) = C\lambda^{-2d}$. After eliminating C, the estimator $\hat{d}$ is equal to:

$$\hat{d} = \arg\min_{d}\left\{\ln\left(\frac{1}{m}\sum_{j=1}^{m}\frac{I(\lambda_j)}{\lambda_j^{-2d}}\right) - \frac{2d}{m}\sum_{j=1}\ln(\lambda_j)\right\}. \tag{16}$$

Although this Gaussian estimator has no closed form, it is more efficient than the averaged periodogram estimator since

$$\sqrt{m}(\hat{d} - d) \sim \mathrm{N}\left(0, \frac{1}{4}\right). \tag{17}$$

Furthermore, Velasco (1998) has considered the nonstationary case, i.e., where $d \geq 0.5$, and has shown that, with tapered data, this estimator is consistent for $d \in (-1/2, 1)$ and asymptotically normal for $d \in (-1/2, 3/4)$, i.e., the statistical properties are robust to nonstationary but nonexplosive alternatives.

The quantlet `robwhittle` computes this local Whittle estimator. Its syntax is:

```
d = robwhittle(ary{, m})
```

where

`ary`
: is the series

`m`
: is a vector of bandwidth parameters. If this optional argument is not provided, the default bandwith vector `m` $= [T/4], [T/8], [T/16]$, where T denotes the sample size.

The results are displayed in the form of a table, the first column contains the value of the bandwidth parameter, while the second column contains the estimated value of d.

The instructions

```
library("times")
y = read("dmus58.dat")
ar = abs(tdiff(y[1:2000]))
d = robwhittle(ar)
d
```

longmem09.xpl

yield the following table

```
Contents of d

[1,] "     d     Bandwidth"
[2,] "____________________"
[3,] ""
[4,] "  0.0948      500"
[5,] "  0.1078      250"
[6,] "  0.1188      125"
```

References

Andrews, D. W. K. (1991). Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation, *Econometrica* **59**: 817–858.

Arteche, J. (1998). Semiparametric LM Tests on Seasonal/Cyclical Long-Memory, *Mimeo University of the Basque Country.*

Arteche, J. and Robinson P. M. (1998). Semiparametric Inference in Seasonal and Cyclical Long-Memory Processes, *LSE Sticerd DP* **EM/98/359**.

Beran, J. (1994). *Statistics for Long-Memory Processes*, Chapman and Hall.

Brockwell, P. J. and Davis, R. A. (1991). *Time Series: Theory and Methods*, Springer-Verlag, New York.

Geweke, J. and Porter-Hudak, S. (1983). The Estimation and Application of Long-Memory Time Series Models *Journal of Time Series Analysis.* **4**: 221–237.

Giraitis, L., Kokoszka, P. S. and Leipus, R. (1998). Detection of Long-memory in ARCH Models, *Mimeo LSE and University of Liverpool, Department of Mathematics.*

Giraitis, L., Robinson, P. M. and Surgailis, D. (1998). A Model for Long-Memory Conditional Heteroskedasticity, *Mimeo LSE.*

Giraitis, L., Kokoszka, P. S., Leipus, R. and Teyssière, G. (1998). Semiparametric Estimation of the Intensity of Long-Memory in Conditional Heteroskedasticity, *Statistical Inference for Stochastic Processes*, forthcoming.

Hurst, H. E. (1951). Long-Term Storage Capacity of Reservoirs, *Transactions of the American Society of Civil Engineers* **116**: 770–799.

Kirman, A. P. and Teyssière, G. (1998). Microeconomic Models for Long-Memory in the Volatility of Financial Time Series, *Mimeo Humboldt University, SFB 373, and GREQAM Marseilles.*

Kwiatkowski, D., Phillips, P. C. B., Schmidt, P. and Shin, Y. (1992). Testing the Null Hypothesis of Stationarity Against the Alternative of a Unit Root: How Sure Are We that Economic Series Have a Unit Root, *Journal of Econometrics* **54**: 159–178.

Künsch, (1987). Statistical Aspects of Self-Similar Processes, *Proceedings of the First World Congress of the Bernouilli Society,* Yu Prohorov and V. V. Sazanov editors. VNU Science Press, Utrecht, **1**, 67–74.

Lee, D. and Schmidt, P. (1996). On the Power of the KPSS Test of Stationarity Against Fractionally-Integrated Alternatives, *Journal of Econometrics* **73**: 285–302.

Lo, A. W. (1991). Long-Term Memory in Stock Market Prices, *Econometrica* **59**: 1279–1313.

Lobato, I. and Robinson, P. M. (1998). A Nonparametric Test for I(0), *Review of Economic Studies*, forthcoming.

Mandelbrot, B. B. (1975). Limit Theorems of the Self-normalized Range for Weakly and Strongly Dependent Processes, *Z. Wahr. Verw. Geb.* **31**: 271–285.

Newey, W. K. and West, K. D. (1987). A Simple Positive Definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix, *Econometrica* **55**: 703–705.

Robinson, P. M. (1994). Time Series with Strong Dependence, *Advances in Econometrics, Sixth World Congress*, CUP: 47–95.

Robinson, P. M. (1994). Semiparametric Analysis of Long-Memory Time Series, *Annals of Statistics* **22**: 515–539.

Robinson, P. M. (1995). Log-Periodogram Regressions of Times Series with Long Range Dependence, *Annals of Statistics* **23**: 1048–1072.

Robinson, P. M. (1995). Gaussian Semiparametric Estimation of Long Range Dependence, *Annals of Statistics* **23**: 1630–1661.

Robinson, P. M. and Henry, M. (1998). Long and Short Memory Conditional Heteroskedasticity in Estimating the Memory Parameter in Levels, *LSE Sticerd DP* **EM/98/357**

Siddiqui, M. (1976). The Asymptotic Distribution of the Range and Other Functions of Partial Sums of Stationary Processes, *Water Resources Research* **12**: 1271–1276.

Teverovsky, V., Taqqu, M. S. and Willinger, W. (1999). A Critical Look at Lo's Modified R/S Statistic, *Journal of Statistical Planning and Inference*, forthcoming.

Velasco, C. (1998). Gaussian Semiparametric Estimation of Non-Stationary Time Series, *Journal of Time Series Analysis*, forthcoming.