Available online at www.sciencedirect.com
SCIENCE DIRECT
Journal of Empirical Finance 11 (2004) 629–658
Journal of EMPIRICAL FINANCE
www.elsevier.com/locate/econbase

Structural change and long-range dependence in volatility of exchange rates: either, neither or both?

Claudio Morana[a], Andrea Beltratti[b,*]

[a] *University of Piemonte Orientale, Italy*
[b] *Istituto di Economia Politica, Università Commerciale L. Bocconi, Via Sarfatti, 25, 20100, Milan, Italy*

Accepted 20 March 2003
Available online 9 April 2004
This paper was presented at the International Conference on the Econometrics of Financial Markets, Delphi, 2001.

Abstract

In this paper, we test for the existence of long memory and structural breaks in the realized variance process for the DM/US$ and Yen/US$ exchange rates. While long memory is evident in the actual processes, a structural break analysis reveals that this feature is partially explained by unaccounted changes in regime. We then compare the forecasting performance of Markov switching models with that of an ARFIMA model. The results indicate that neglecting the break process is not important for very short term forecasting once it is allowed for a long memory component in the model, but that superior forecasts can be obtained at longer horizons by modelling both long memory and structural change.
© 2004 Elsevier B.V. All rights reserved.

JEL classification: C14; C22; C53; G10
Keywords: Long memory; Structural change; Forecasting

1. Introduction

Since Diebold (1986) and Lamoreaux and Lastrapes (1990), it is well known that misleading inference on the persistence of the volatility process may be caused by unaccounted structural breaks. Econometric models where structural change can be modelled endogenously have been proposed by Cai (1994) and Hamilton and Susmel

* Corresponding author. Tel.: +39-2-58-36-36-53-06; fax: +39-2-58-36-53-43.
E-mail address: andrea.beltratti@uni-bocconi.it (A. Beltratti).

0927-5398/$ - see front matter © 2004 Elsevier B.V. All rights reserved.
doi:10.1016/j.jempfin.2003.03.002

(1994) for the Markov switching ARCH model while So et al. (1998) and Duker (1997) have extended the Markov switching model to the stochastic volatility and GARCH frameworks, respectively.

Recent contributions have proposed the alternative explanation of long-range dependence to account for the persistence of the conditional variance process; see, for instance, Ding et al. (1993); Dacorogna et al. (1993); Baillie et al. (1996); Bollerslev and Mikkelsen (1996); Lobato and Savin (1998) and Andersen and Bollerslev (1997). In these papers, IGARCH behavior is interpreted as a spurious feature and as an approximation to the true long-memory behavior. In addition to empirical evidence, economic justifications have been provided for the presence of long-range dependence in the volatility of financial returns. For instance, Muller et al. (1997) have proposed an explanation based on the interaction in the market of agents with different time horizon. According to this hypothesis, long memory arises from the reaction of short-term dealers to the dynamics of a proxy for the expected volatility trend (coarse volatility) which causes persistence in the mean higher frequency volatility process (fine volatility). On the other hand, long-term dealers base their decisions on the fundamentals of the market, ignoring short-term movements. Some skeptical views about long memory have been offered by Granger and Hyung (1999) and Mikosch and Starica (1998).

The choice between structural breaks and long memory is important for a variety of financial applications, particularly in risk measurement, asset allocation and option pricing. The use of a structural break model would require that much attention is devoted to forecasting and quickly recognizing a change in the parameters of the econometric model, a task that is by no means simple lacking a theory describing why and how there may be a change in parameters. The use of a long-memory model seems a simpler practical alternative for financial decision-making, although the extent to which this is worth is an empirical question. Along these lines, Diebold and Inoue (1999) suggest that long memory may be a useful description, particularly for forecasting purposes, even if the data generating process shows structural breaks and weak dependence. However, the authors do not offer empirical evidence supporting their view. In a quite different context, various attempts have been proposed to compare the forecasting performance of linear, fixed coefficients models with that of time-varying parameters models. The results generally indicate that adding time-varying parameters may even worsen the forecasting performance of a simple linear regression.

In the paper, we analyze the properties of the realized variance processes of Andersen et al. (1998) for the Deutsche mark/US dollar and Japanese yen/US dollar exchange rates with the goal of studying whether explicitly considering structural breaks brings some benefit in terms of understanding and forecasting the economic forces driving volatility in the foreign exchange markets. The main results of the paper are the following. The realized volatility processes show long memory also when structural change is accounted for. In fact, when the hypothesis is tested using break-free series, we still find evidence of long memory, although the degree of persistence is weaker. We also compare the forecasting performance of switching regimes models and an ARFIMA model. We find that neglecting the break process is not important for very short term forecasting once it is allowed for a long-memory component in the model, but that superior forecasts can be obtained at longer horizons by modelling both long memory and structural change. This finding provides support to the conclusion that long memory in the variance series considered cannot be fully explained by unaccounted structural change and that forecasting may be safely based on a standard model with fixed coefficients.

After this introduction, the paper is organized as follows. In Section 2, we briefly summarize the findings in the literature which have provided empirical evidence of long memory and that have cast doubts on its nature. In Sections 3 and 4, we describe the data, investigate the long-memory hypothesis and test for structural breaks. In Section 5, we compare the forecasting performance of the models. Finally, in Section 6, we conclude.

2. Spurious or real long memory?

According to one popular definition, a stationary processes X_t is said to be long range dependent if the autocorrelation function is significantly different from zero at very long lags; that is

$$\rho_X(\tau) = c_\rho \tau^{2d-1}$$

where c_ρ is a positive constant, τ is the order of the autocorrelation and $d \in (0,0.5)$ is the coefficient of fractional integration, which is related to the Hurst exponent by the relation $H = 1/2 + d$. Differently from $I(0)$ processes, $I(d)$ stationary processes do not show an exponentially fast decay of the autocorrelation function but a slow hyperbolic decay.

The main implication of long memory is that shocks to the volatility process tend to have long-lasting effects. For instance, a fractional white noise process $(1-L)^d X_t = \varepsilon_t$ has a MA(∞) representation with coefficients that show a hyperbolic decay. Given this strong implication of long-range dependence, an accurate testing of the memory properties of the volatility process is required before such a model is chosen for volatility forecasting.

The empirical literature on long memory in volatility is vast and two opposite positions can be found, one in favor of the long-memory hypothesis, the other against it. To the first group belong the seminal papers of Dacorogna et al. (1993) and Ding et al. (1993) and the more recent works of Ding and Granger (1996a,b), Lobato and Savin (1998), Baillie et al. (1996), Bollerslev and Mikkelsen (1996), Breidt et al. (1998), Henry and Payne (1997), Andersen and Bollerslev (1997), Beltratti and Morana (1999), Andersen et al. (1998), Bollerslev and Wright (2000), Eibens (1999), Beltratti and Morana (2001). The evidence provided in the studies above is comprehensive. The memory properties of the volatility process have been analyzed with reference to both stock and exchange rate returns, considering daily and intradaily data over different time spans and subsamples, different econometric techniques, both parametric (FIGARCH, FIEGARCH, LMSV, ARFIMA) and semiparametric (stationarity tests, log periodogram regression, averaged periodogram and local Whittle estimators), and by modelling volatility also as an observed process (realized volatility process). The common conclusion that is drawn in these studies is that the volatility process of financial returns is well described by a long-memory process; that is, long memory appears to be a feature of the data generating process of financial returns. Coherently with the theoretical results of Chambers (1998), long memory has been found in the volatility of individual stocks as in that of aggregate indexes and the degree of persistence has been found to be invariant with respect to the frequency of the data.

On the other hand, to the second group of studies belong the papers of Ryden et al. (1998), Mikosch and Starica (1998), Granger and Hyung (1999), Liu (2000). The main conclusion that is drawn in these studies is that the strong persistence shown by the volatility process of financial returns is a spurious feature due to unaccounted structural change. Under this interpretation, long memory in volatility is due to unaccounted structural breaks. Cai (1994), Hamilton and Susmel (1994), So et al. (1998) and Duker (1997) have suggested a framework in which, by controlling for structural change, volatility modelling can be made more reliable. The main conclusion of these latter studies is that the volatility process can be successfully described by a model of switching heteroskedasticity. Economic interpretations of these shifts have also been provided, relating the dynamics of stock volatility to the business cycle, the volatility of fundamentals and the fad component of returns.

However, as shown by Ryden et al. (1998), a Markov switching model can replicate most of the properties of the raw and absolute returns indicated in Ding and Granger (1995) but predicts an exponentially fast decaying correlogram for the absolute returns, in contrast with the evidence largely provided in the literature. Mikosch and Starica (1998) have shown that the evidence of long memory in the conditional variance of returns for the S&P 500 index (see also Ding et al., 1993 and Bollerslev and Mikkelsen, (1996) may be an artifact due to unaccounted structural change. They have shown that the correlogram of the absolute returns is consistent with long memory (i.e., the correlogram follows a slow hyperbolic decay) only when data corresponding to different regimes for the unconditional variance are included in the sample used to compute the autocorrelation function. Moreover, Granger and Hyung (1999), reexamining the finding of Ding et al. (1993) and Ding and Granger (1996a,b), have analyzed the memory properties of the break-free absolute return process for the S&P 500 index. Their conclusion is that structural change is the likely cause of the observed long-memory behavior. Support to this conclusion has also been provided by Liu (2000) and explained on the basis of the fat tails of the duration distribution of the regimes. These latter findings are important for our study because the sample analyzed (1986–1996) is large enough to argue in favor of the presence of structural breaks in the exchange rates series analyzed. Whether indeed the low-frequency (persistent) volatility component is fully explained by the break process, as the latter author suggest, is the central question we attempt to answer in this paper.

3. Empirical results

In this paper, we focus on daily realized variance processes. The choice has both statistical and practical motivations. From a statistical point of view, computing daily realized variance processes yields a sample of 2445 observations, which is large enough to make the semiparametric analysis and the out-of-sample forecasting exercise meaningful. From a practical point of view, the daily frequency is relevant to the financial industry and to investors. For instance, risk management needs accurate forecast of daily and weakly volatility for implementation of value at risk models. Furthermore, in the case of quantitative asset allocation models, investors are interested in risk measurement at the daily or even lower frequencies.

The data employed in this study are 5 min returns for the mark–dollar and the yen–dollar exchange rates over the period December 1, 1986–December 1, 1996. The data have been provided by Olsen and Associates[1] and are the same as employed by Andersen et al. (1998).[2] The daily return is defined as the logarithmic difference between exchange rates recorded at 21 p.m. G.M.T. A number of thin trading/nontrading days have been excluded from the sample, namely, the weekend period (from Friday 21:05 GMT to Sunday 21:00 GMT), Christmas (December 24–26), New Year's (December 31–January 2), July Fourth (July 4 or July 3), Good Friday, Easter Monday, Memorial Day, Labor Day, Thanksgiving (and the day after).[3] After data cleaning, we are left with a sample size equal to 2445 complete days.

Daily realized variance processes have been obtained by summing 5 min squared returns over 280 successive observations, i.e., by computing

$$\sigma_t^2 = \sum_{n=1}^{280} r_{t,n}^2.$$

Theoretical results presented in Andersen et al. (1998) shows that the realized volatility process is a consistent (in the frequency of sampling) estimator of integrated variance. This implies that, arbitrarily, close approximation of integrated variance may be obtained at all the relevant horizons by summing high-frequency squared returns. This approach has clear advantages relative to the practice of squaring low-frequency returns because it yields virtually noise-free measures of volatility and allows therefore to model volatility as observed variables. On the other hand, squared low-frequency returns yield a noisy volatility proxy, which necessarily require the use of unobserved component models in order to avoid inconsistent estimation. In our sample, the correlation between the time series of realized variance and the time series of squared daily returns is 0.53 for the mark and 0.49 for the yen, suggesting that these series contain indeed different dynamics.

Table 1 contains summary statistics for daily returns, log realized variance and realized variance for the whole sample period 1986–1996. As far as the returns are concerned, there is evidence of weak negative asymmetry and strong positive excess kurtosis. As a consequence, unconditional normality is strongly rejected. In addition, there is no sign of autocorrelation. Realized variance ranges in a much larger interval than returns. While the ratio between the absolute values of the maximum and the minimum for returns is about 1, in the case of the variance, the ratio is about 10 for the mark/dollar and 30 for the yen/dollar. The kurtosis of log variance is lower than that of returns and much lower than that of simple variance. However, normality is rejected for log variance as well. Finally, the Box–Ljung portmanteau test reveals a strong persistence in the variance and log variance processes, finding that may suggest the presence of long-range dependence in these series. In Fig. 1, we plot the autocorrelation function for the variance series up to 1000 lags, which reveals evidence of long memory. The autocorrelation function for both exchange

[1] See Dacorogna et al. (1993) and Muller et al. (1990) for a description of how these data are built starting from bid and ask prices collected in real time from various information sources.

[2] We are grateful to T. Andersen for kindly making us available the daily series employed in this study.

[3] See Andersen et al. (1998) for a full description of the construction of the data.

Table 1

Panel A: Summary statistics for daily returns and volatilities

	r_t (DM/US$)	σ_t^2 (DM/US$)	$\ln\sigma_t^2$(DM/US$)
Maximum	3.506	5.245	1.657
Minimum	− 3.244	0.052	− 2.961
Mean	− 0.007	0.529	− 0.897
S.D.	0.699	0.484	0.691
Skewness	− 0.007	3.708	0.345
Excess kurtosis	2.212	21.1	0.263
Normality test	0.000	0.000	0.000
$Q(5)$	0.130	0.000	0.000
$Q(20)$	0.629	0.000	0.000
$Q(200)$	0.656	0.000	0.000

Panel B: Summary statistics for daily returns and volatilities

	r_t (Yen/US$)	σ_t^2 (Yen/US$)	$\ln\sigma_t^2$(Yen/US$)
Maximum	4.050	10.097	2.312
Minimum	− 3.628	0.028	− 3.576
Mean	− 0.010	0.538	− 0.887
S.D.	0.692	0.522	0.702
Skewness	− 0.121	5.571	0.262
Excess kurtosis	3.566	63.7	0.523
Normality test	0.000	0.000	0.000
$Q(5)$	0.567	0.000	0.000
$Q(20)$	0.010	0.000	0.000
$Q(200)$	0.785	0.000	0.000

The table shows summary statistics for daily returns (r_t), realized variance (σ_t^2) and log realized variance ($\ln\sigma_t^2$) for the Deutsche mark–US dollar and Yen–US dollar exchange rates. The daily returns are constructed as the sum of the 5 min returns. The daily realized variance is constructed as the sum of the squares of the 5 min returns. The normality test is a $\chi^2(2)$ test for normality of the unconditional distribution. $Q(k)$ is the Box–Ljung portmanteau test for serial correlation up to the kth order.

rates decays slowly but is not statistically significant, according to the standard significance band, already after 60 lags. In addition, it becomes negative for some long lags.

However, there are many reasons why the autocorrelation function may lead to misleading inference on the persistence property of a time series. Firstly, if a very long span of time is used in the analysis, nonstationarity in the volatility process can easily be confused for long memory. Mikosch and Starica (1998) have shown that the correlogram for a series composed of two subsequent samples extracted from populations, each one strictly stationary but with different means, converges in probability, for large lags, to a constant given by the squared difference of the two sample means. Secondly, Hosking (1996) has shown that the limiting distribution of the autocorrelation function for a long-memory process is not standard normal. This result has implications for testing the significance of the autocorrelation function using standard bounds. As suggested by Mikosch and Starica (in press), the correlogram at long lags, although different from zero, may not be statistically significant. Using inappropriate critical values can, therefore, lead to invalid conclusions. Thirdly, a hyperbolic decay of the correlogram can be produced by a short-memory nonlinear model. Granger and Terasvirta (1999) and Diebold and Inoue

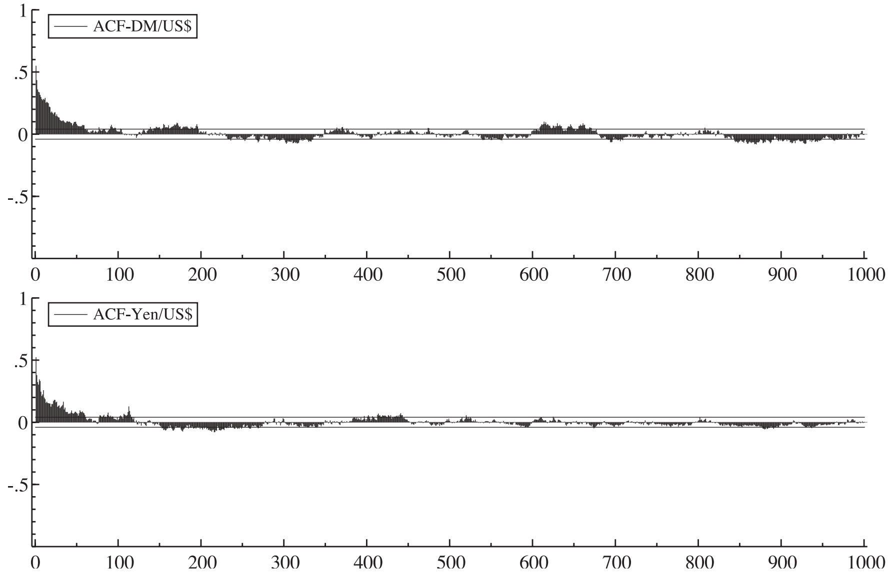

Fig. 1. Autocorrelation functions.

(1999) have provided some examples. Finally, Granger and Marmol (1997) have shown that if the series is noisy, that is, the series is composed of a long-memory process plus a simple noise component, the autocorrelation function is underestimated at all lags, with the degree of underestimation depending on the variance of the noise component.

All the results above are important for our study because the sample employed encompasses periods of market tranquillity as well as periods of prolonged market turbulence, so that the presence of structural breaks in the variance processes cannot be excluded a priori. In addition, the consequences of observational noise on the estimation of the autocorrelation function should not be neglected. Interestingly, evidence from Beltratti and Morana (2001) suggest that, while the volatility measure computed by aggregating intraday squared returns is much less noisy than the ex post volatility measure computed by squaring daily returns, still it contains a nonnegligible noise component. Estimation of a stochastic volatility model on the log variance process reveals that the standard deviation of the irregular and autoregressive components are equal to 0.41 and 0.19 for the mark/dollar and 0.38 and 0.26 for the yen/dollar, respectively.

Following the graphical analysis, therefore, we have investigated the long-memory hypothesis more rigorously by means of a battery of tests which are described in Appendix A. Several semiparametric estimators for the Hurst exponent have been proposed in the literature. In the paper, we employ the local Whittle estimator suggested by Kunsch (1987) and compare the estimates with those obtained using the averaged periodogram estimator (Robinson, 1994a), the log periodogram regression (Geweke and Porter-Hudak, 1983; Robinson, 1995b) and a new estimator proposed by Robinson (1998). Robinson (1995a,b) has shown that the local Whittle estimator dominates the averaged periodogram estimator and the log periodogram estimator under several respects. Firstly, it is asymptotically

Table 2

Panel A: Stationarity and semiparametric tests: DM/US$

	r_t (DM/US$)	σ_t^2 (DM/US$)	$\ln\sigma_t^2$ (DM/US$)
H_{LW}	0.492 (0.027)	0.999 (0.055)	0.999 (0.054)
H_{LP}	0.480 (0.003)	0.971 (0.010)	1.145 (0.010)
H_{AP}	0.528 (0.042)	0.917 (–)	0.942 (–)
H_{LM}	0.495 (–)	0.845 (–)	0.835 (–)
H_A	0.489 (0.016)	0.880 (0.016)	0.913 (0.015)
W	0.033 (0.857)	42.8 (0.000)	42.3 (0.000)
LM	0.032 (0.858)	443.7 (0.000)	390.2 (0.000)
LR	0.035 (0.852)	156.1 (0.000)	166.9 (0.000)

Panel B: Stationarity and semiparametric tests: Yen/US$

	r_t (Yen/US$)	σ_t^2 (Yen/US$)	$\ln\sigma_t^2$ (Yen/US$)
H_{LW}	0.543 (0.032)	0.999 (0.055)	0.999 (0.056)
H_{LP}	0.506 (0.004)	0.934 (0.010)	1.022 (0.011)
H_{AP}	0.598 (0.044)	0.925 (–)	0.937 (–)
H_{LM}	0.539 (–)	0.843 (–)	0.837 (–)
H_A	0.482 (0.016)	0.868 (0.016)	0.947 (0.016)
W	1.492 (0.222)	43.7 (0.000)	42.6 (0.000)
LM	1.758 (0.185)	441.6 (0.000)	400.3 (0.000)
LR	1.787 (0.181)	164.3 (0.000)	158.8 (0.000)

The table shows stationarity and semiparametric tests, with standard errors or p-values in brackets, for daily returns (r_t), realized variance (σ_t^2) and log realized variance ($\ln\sigma_t^2$) of the Deutsche mark–US dollar and Yen–US dollar exchange rates. H_i i=LW, LP, AP, LM, A is the estimate of the Hurst exponent using the Gaussian semiparametric estimator, the log periodogram estimator, the averaged periodogram estimator, the LM estimator, and the ARFIMA(0,d,0) model, respectively. LM is the LM test; W is the Wald test; LR is the Likelihood Ratio test. In the implementation, we used $q=0.5$ for the averaged periodogram estimator and $l=0$ for the semiparametric estimators.

efficient; secondly, it does not require the trimming of the lowest frequency estimates; thirdly, it does not assume Gaussianity; fourthly, it does not require the selection of a numerical value for any additional parameter.

In Table 2, we report the results of the semiparametric analysis.[4] As shown in the table, the results are similar for both exchange rates. As far as the return series are concerned, the null of weak dependence is never rejected. In fact, the estimated Hurst coefficients are close to the theoretical value of 0.5, which implies a zero value for the fractional differencing parameter. In addition, the null of $I(0)$ stationarity is (largely) not rejected by the LM, Wald and LR tests. On the other hand, the log-realized variance and the realized variance processes appear to be characterized by long memory. The Hurst

[4] The selected bandwidths tended to be higher for the log periodogram regression than for the other estimators. For the returns, the minimum bandwidth is 202 for the averaged periodogram (Yen/US$) and a maximum of 598 for the log periodogram estimator (DM/US$). For the variance processes, the minimum bandwidth is 69 for the averaged periodogram and a maximum of 163 for the log periodogram (both exchange rates). Finally, for the log variance processes, the minimum bandwidth is 68 (Yen/US$) for the averaged periodogram estimator and the maximum is 172 for the log periodogram estimator (DM/US$).

Optimal bandwidth estimation was carried out using a GAUSS routine written by Marc Henry, to whom the authors are grateful.

coefficient is below but close to 1 for both exchange rates, and, in general, it is higher for the log-variance process than for the variance process. Finally, the local Whittle and the log periodogram estimator yield the highest estimates (the estimate ranges from a minimum of 0.93 for the yen/dollar variance process and a maximum of 1.15 for the mark/dollar log variance process), while the LM estimator (Robinson, 1998) produces the lowest ones (the estimates are for both process close to 0.84).[5]

As a comparison, we have also estimated an ARFIMA model. After some experimentation, an ARFIMA(0,d,0) was chosen as the best specification for all of the processes.[6] The comparison with the parametric estimator is important because parametric analysis does not require the selection of a bandwidth, although it has the drawback that a misspecified short-run structure may lead to unreliable estimates of the degree of long memory. As shown in Table 2, the point estimates of the coefficient of fractional integration are close to those obtained by Andersen et al. (1998) using the log periodogram estimator and setting the bandwidth to $T^{4/5}$ (514), and have an intermediate magnitude relatively to the estimates obtained using the semiparametric methods and setting the bandwidths according to optimal bandwidth theory. The implied Hurst exponents for the variance processes are in fact 0.88 and 0.87, while for the log variance process are 0.91 and 0.95 for the mark and the yen, respectively. The LM, Wald and LR tests always reject the null of weak dependence. The parametric analysis, therefore, confirms the results of the semiparametric analysis: the processes are covariance stationary and long memory.

Given the absence of a statistically significant short-run structure, it may be argued that the most reliable estimates of the Hurst exponent are those provided by the parametric estimator because the bandwidths selected according to optimal bandwidth theory are very conservative. A drawback of the semiparametric analysis is that the estimates may be sensitive to the bandwidth m, i.e., the number of the periodogram ordinates used in the computation of the statistics. The theory of optimal bandwidth selection is based on the minimization of the asymptotic MSE, and the rationale underlying this practice is that, to study the memory properties of a long-range dependent series, one should look at the periodogram for frequencies close to zero. Looking at the periodogram for frequencies very far from zero may lead to misleading conclusions due to the bias induced by the short-term correlation structure on the estimate of the Hurst exponent when the true process is an ARFIMA model and not a simple fractional white noise. Minimizing the MSE is then an effective way of trading-off dispersion and bias. However, on the basis of

[5] Kunsch (1986) has suggested that a weakly dependent process with certain trends can be confused for a long-memory process. To avoid confusion, the lowest frequency ordinates of the periodogram should be trimmed out. By analyzing the dependence of the estimates on the bandwidth, we have found that trimming the lowest ordinates of the periodogram my lead to unreliable results, particularly when m is small. The original log periodogram estimator should therefore be preferred (see Hurvich and Beltrao, 1994 for a similar finding).

[6] No residual serial correlation is left in the data after fractional differencing, apart from the Yen/US$ variance process. Although the AIC criterion would select a less parsimonious ARFIMA(5,d,0) specification, apart from the DM/US$ variance process for which the most parsimonious model should also be preferred, the estimated autoregressive parameters are very close to zero and do not have any relevant impact on the fitted signal: the correlation between the fitted process from the ARFIMA(0,d,0) and the ARFIMA(5,d,0) is in fact about 0.99. Normality as well as lack of ARCH effects in the residuals are rejected in all of the cases. Details about ML estimation of ARFIMA models can be found in Sowell (1992).

our results, the realized variance and log variance series appear to follow a fractional white noise process, so that the minimization of the MSE would require the selection of a higher number of ordinates than the one selected by the procedure implemented. For instance, following Hurvich et al. (1998), for the log-periodogram estimator, the asymptotic MSE can be written as $\mathrm{MSE}(\hat{d}) = \frac{4\pi}{81}(2\tau^*)^2 \frac{m^4}{T^4} + \frac{\pi^2}{24m}$, where m is the number of periodogram ordinates, T is the sample size, the smoothness parameter is $\tau^* = (f^{*\prime\prime}(0)/2f^*(0))$, and the function $f^*(\cdot)$ describes a short-term correlation structure in the spectral density of the process $f(\lambda) = |1 - \exp(i\lambda)|^{1-2\hat{H}} f^*(\lambda)$. Then, for a fractional white noise process, we have $\mathrm{MSE}(\hat{d}) = \frac{\pi^2}{24m}$ and $\lim_{m\to\infty} \mathrm{MSE}(\hat{d}) = 0$. The bandwidth selected by Andersen et al. (1998), using the graphical selection procedure of Taqqu and Teverovsky (1998), supports this conclusion.

4. Testing for structural breaks

Because long memory may be a spurious feature when structural change is not accounted for, we investigated whether structural breaks may characterize the DGP of the two exchange rates considered. The hypothesis of structural change is reasonable because our sample cover 10 years of data, from December 1, 1986 through December 1, 1996. We implemented the tests of Inclan and Tiao (1994) and Bai (1997). We also estimated the break processes using a recurrent state Markov switching model for the log variance processes, which allows the intercept to switch across regimes (see Appendix C for a description of the methodology).[7] It is important to notice that, as suggested by Granger and Hyung (1999) and Hsu (2001), applying a break detection approach can lead to spurious conclusions concerning the number of breaks if the DGP truly follows an $I(d)$ process without breaks. In fact, as shown by Hsu (2001), the Bai (1997) test will reject the null of no structural change with probability one when the process has long-memory innovations and no structural breaks. In addition, it may identify a spurious break in the middle of the sample. On the other hand, Granger and Hyung (1999) have shown that the number of breaks detected increases with the value of d and it is zero only if the process is truly $I(0)$. Estimating d from a series from which spurious breaks have been removed should yield a value of $d<0$; finding that can be employed in practice to decide whether the $I(d)$ property of a series is an artifact of unaccounted occasional breaks or a true feature of its DGP.

4.1. Results

The Markov switching model allows joint estimation of the break process and the autoregressive dynamics. However, we have found that the estimated break processes resulting from the joint estimation show too much variability to be considered an accurate

[7] In a recent paper Timmerman (2001) has suggested to employ a nonrecurrent state Markov switching model for break process estimation. However, this latter model allows only for an ex post evaluation of the presence of a regime change. We have therefore employed a standard recurrent state model, which has the advantage of allowing to forecast the break process.

representation of the low-frequency volatility component. We have analyzed MS–AR models with up to three regimes, changing intercepts and variances and/or autoregressive parameters. The selection of the best model was not straightforward. For instance, allowing for regime dependence in all of the parameters of the model suggested that a single regime model may capture the characteristics of the data. This is in contrast with the outcome of LR tests for regime selection, which suggest that a three-regime specification should be selected for both exchange rates when allowing only for a switching intercept and variance. Restricting regime dependence to the variance and intercept components only, or reducing the number of regimes from three to two, did not improve the results of the specification analysis or the interpretability of the estimated break processes. In fact, even the inclusion of five lags was not sufficient to eliminate serial correlation from the estimated residuals, and the estimated own transition probabilities indicated low persistence of all the regimes. Some experimentation suggested to include only two lags in the model and to allow both the intercept and the variance to be regime dependent.

While the estimated standardized residuals reveal the presence of autocorrelation and nonnormality, the estimated regimes are more interpretable, although the duration of the regimes is still very short to be regarded as plausible. For instance, for the mark/dollar, the estimated own transition probabilities (and the average duration of the regimes) are 0.88 (8.5 days) for the low-volatility regime, 0.95 (19 days) for the average volatility regime, and 0.51 (2 days) for the high-volatility regime. On the other hand, for the yen/dollar the estimates are 0.90 (11 days) for the low-volatility regime, 0.84 (6.4 days) for the average volatility regime, and 0.99 (72 days) for the high-volatility regime. Neglecting the autoregressive dynamics did not help the estimation of the low-frequency volatility component, yielding a break process still characterized by very frequent changes (Fig. 2).

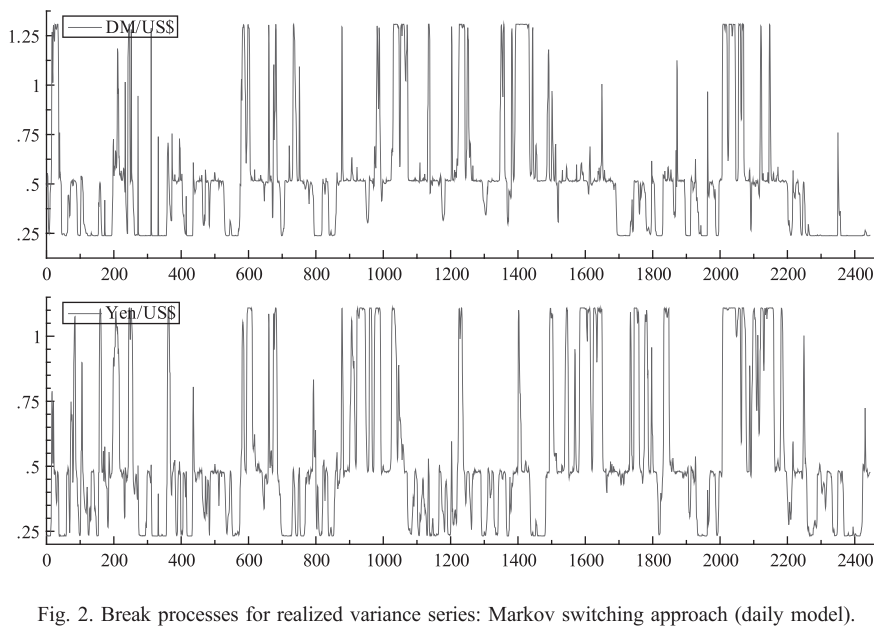

Fig. 2. Break processes for realized variance series: Markov switching approach (daily model).

Table 3

Panel A: Markov switching daily break processes, daily data

	$\ln\sigma_t^2$ (DM/US$)	$\ln\sigma_t^2$ (Yen/US$)
μ_1	− 1.558 (0.022)	− 1.584 (0.029)
μ_2	− 0.785 (0.024)	− 0.857 (0.028)
μ_3	0.146 (0.059)	− 0.019 (0.033)
p_{11}	0.943	0.910
p_{22}	0.941	0.915
p_{33}	0.849	0.905
d_1	16.25	10.38
d_2	17.99	9.95
d_3	11.80	8.88

Panel B: Markov switching daily break processes, monthly data

	$\ln\sigma_t^2$ (DM/US$)	$\ln\sigma_t^2$ (Yen/US$)
μ_1	− 1.265	− 1.142
μ_2	− 0.562	− 0.433
p_{11}	0.995	0.996
p_{22}	0.995	0.994
d_1	198.9	230.7
d_2	202.8	119.5

The table shows the switching unconditional means (μ_i) with standard errors in brackets for the log realized variance processes ($\ln\sigma_{tj}^2$) for the Deutsche mark–US dollar and Yen–US dollar exchange rates. p_{ii} are the own transition probabilities, while d_i indicates the duration of each regime.

Because the break process should reflect a low-frequency volatility component, one may argue that neglecting high-frequency dynamics, i.e., neglecting daily volatility extreme realizations, may lead to a more reliable estimate using a Markov switching model. This can be accomplished using low-frequency data for the log variance processes. We found that using monthly observations provided a reliable description of the break process for the two exchange rates.[8] After some experimentation, a two-regime specification was selected for the monthly models, contrary to the three-state specification selected for the daily models.

The results of the Markov switching analysis are reported in Table 3. A number of features can be noticed. Firstly, the estimated regimes are very persistent; that is, the own transition probabilities are close to one for both the variance processes. Secondly, the two processes show a similar duration for the low-volatility state (231 days for the yen/dollar exchange and 199 days for the mark/dollar), while the mark/dollar show a higher duration for the high-volatility state (203 days for the mark/dollar and 120 days for the yen/dollar). The estimated break processes indicate that the low-frequency volatility component is different for the two exchange rates, as is also shown by the smoothed probabilities plotted in Fig. 3. Although there is some evidence of comovement of the two variance processes, the null of a single-volatility factor can be rejected on the basis of the different behavior of

[8] Monthly realized variance series have been computed by summing the 5 min returns over 5600 successive observations.

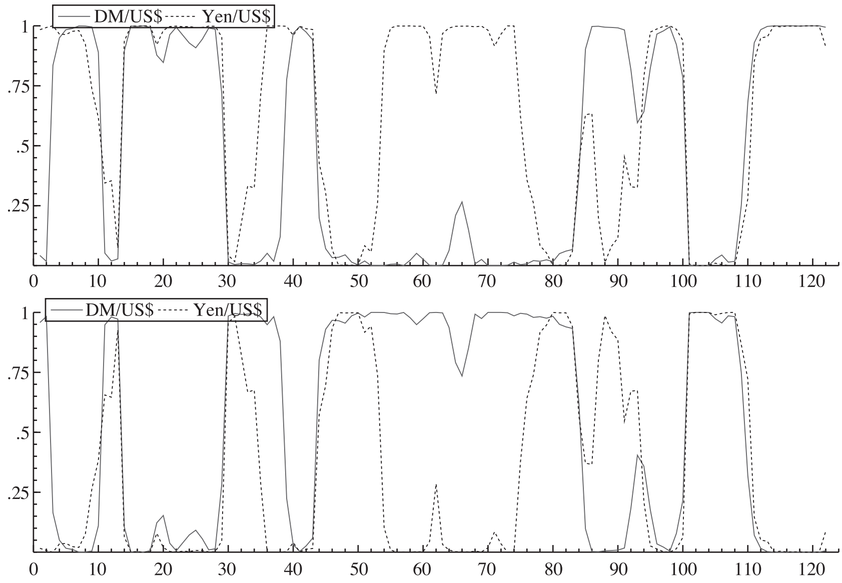

Fig. 3. Estimated smoothed probabilities (monthly model for log realized variances), low-volatiltity regime (top plot) and high-volatility regime (bottom plot).

the break processes in the early 1990s (observations 52 through 78).[9] The number of breaks detected for the mark is 9, while for the yen, the number of breaks is 12. Finally, the yen shows higher volatility levels in both states than the mark. From the monthly model, a daily break process can be obtained by means of a step function; that is, the daily smoothed probabilities are obtained from the monthly probabilities by keeping constant their value over each sequence of 20 observations. The transition matrix for the daily model is then obtained from the monthly one by means of the relation between the own transition probabilities and the duration of the regime. For the two-state model, we have

$$m_i = \frac{p_{ii,m}}{1 - p_{ii,m}}.$$

where m_i is the duration in months for state $i = 1, 2$, and $p_{ii,m}$ is the own transition probability for state i. By rewriting

$$p_{ii,m} = \frac{m_i}{1 + m_i}$$

we get

$$p_{ii,d} = \frac{20m_i}{1 + 20m_i}$$

[9] A Probit model indicated that the probability that the two exchange rates are in the same regime is about 0.7, and that this probability is statistically different from one.

and $p_{ji,d} = 1 - p_{ii,d}$. Finally, the daily volatility levels for the two states can be recovered from the monthly values by computing

$$\hat{\mu}_{d,t} = \frac{\exp(\hat{\mu}_{m,t} + 0.5V(\ln\sigma^2_{m,t}))}{20}$$

so that the daily break-free log variance process can be obtained as

$$(\ln\sigma^2_{t,j})_{d,bf} = \ln\sigma^2_{t,j} - \ln\hat{\mu}_{d,t,j}$$

The Inclan and Tiao (1994) and Bai (1997) approaches yield different results. The number of breaks estimated is 15 for the mark and 14 for the yen when the Inclan and Tiao (1994) method is used, and 8 and 7, respectively, when the Bai (1997) method is used. Figs. 4 and 5 show that the break processes obtained from the Inclan and Tiao and the Markov switching approaches are similar (apart from the longer high-volatility period detected by the Markov switching model in the middle of the sample for the mark). On the other hand, in the case of the yen, the two break processes are always very close. From Figs. 4 and 5, it can also be noticed that little is lost by assuming recurrent breaks, i.e., by constraining the volatility switches between two levels.

Finally, in Table 4, we report the results of the stationarity tests carried out on the break-free log variances. There is no evidence of spurious breaks for both exchange rates. In general, as for the raw variance series, the estimates of the Hurst exponent provided by the log periodogram regression tended to be higher than those obtained from the other methods employed, and the LM, Wald and LR tests reject the null of weak dependence.

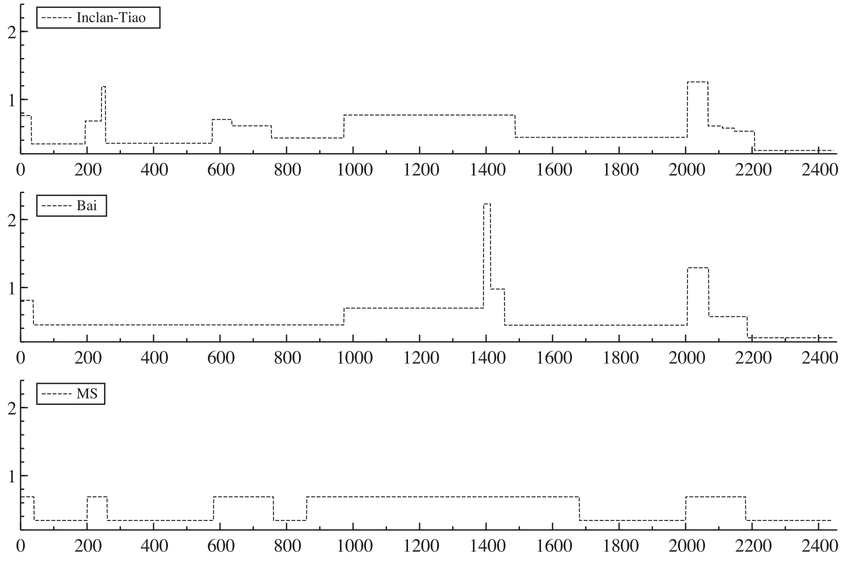

Fig. 4. Break processes for DM/US$ realized variance.

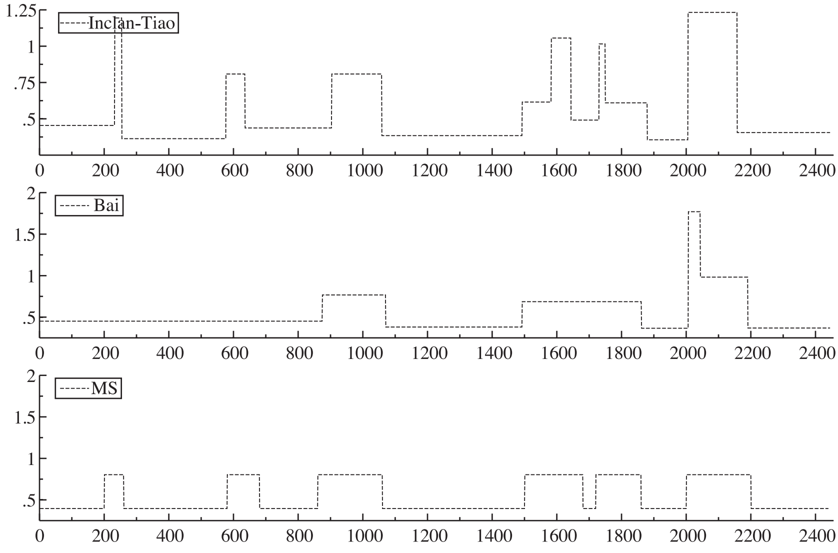

Fig. 5. Break processes for Yen/US$ realized variance.

On the other hand, the estimates obtained using the ARFIMA(0,d,0) model[10] tended to be higher than those obtained using the semiparametric estimators. Again, on the basis of the results of the specification analysis of the ARFIMA models, it can be argued that the number of ordinates selected by the optimal bandwidth procedure is very low and that conclusions on the degree of long memory of the break-free log variance series should be drawn preferably from the parametric estimator, rather than from the semiparametric estimators. For the Markov switching break-free processes, the ARFIMA estimator suggests that little persistence is explained by the structural breaks. The estimated fractional differencing parameters are in fact smaller than those obtained from the raw series but still close to 0.40 (0.36 for the mark and 0.41 for the yen).

Overall, the results suggest that accounting for structural breaks has an impact on the persistence properties of the variance series and that structural change is an important feature of the data which should be modelled. This result is also consistent with the findings of Beine and Laurent (2000), using a Markov switching FIGARCH model.

To gauge further evidence on the joint presence of both long memory and structural change, we have carried out a robustness exercise aiming to assess the impact of modelling the break process as a step function rather than as a smoothly evolving function. One possible advantage of this latter specification is that the conclusions on the long-memory

[10] As for the raw data, a parsimonious ARFIMA(0,d,0) was found sufficient to model the break-free processes according to the Ljung–Box test. Only for the break-free Yen/US$ variance processes there is some evidence of serial correlation in the residuals. However, including up to five autoregressive terms did not improve the specification of the model. Moreover, the inclusion of the autoregressive terms caused the estimated fractional differencing parameter to assume negative values, results that are in contrast with the semiparametric analysis.

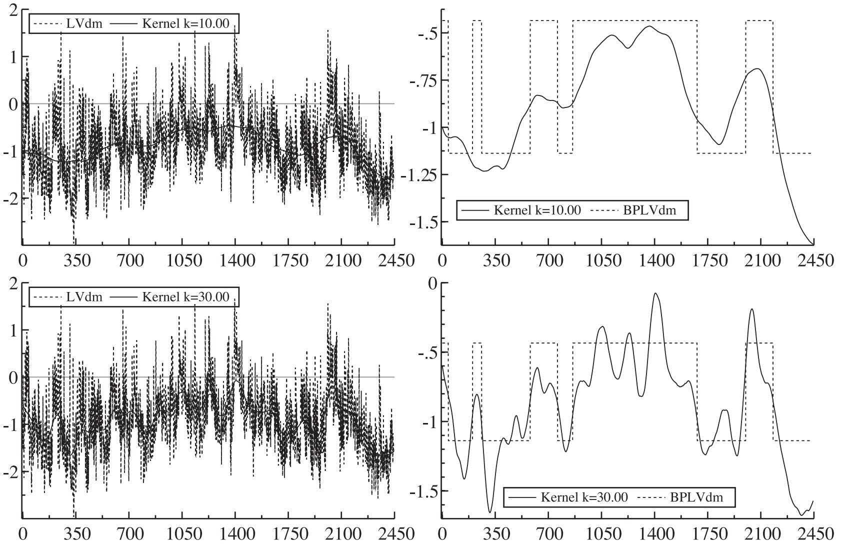

Fig. 6. DM/US$ log variance process (LVdm), Kernel interpolator (Kernel) and Markov switching break process (BPLVdm).

properties of the break-free processes should be less affected by the problem of accurately selecting the exact timing of the switches between regimes.[11] To estimate a smoothly evolving break process, we have followed Beran (1999) and employed a Kernel approach. In Figs. 6 and 7, we plot the Kernel smoother obtained allowing for a bandwidth equal to 10 and 30. The lower the bandwidth, the smoother is the estimated break process. As shown in the plot, setting the bandwidth equal to 30 is sufficient to capture the break points estimated by the Markov switching model. In Fig. 8, we plot the estimated Hurst exponent for the break-free processes against the bandwidth. As shown in the plots, in correspondence of the stable regions, there is clear evidence of long memory according to all the estimators (estimates range between a minimum of 0.66 and a maximum 0.91). Points estimates obtained by the ARFIMA model are 0.3450 (0.0169) and 0.3955 (0.0156), for the mark (30 and 10 parameters), 0.40 (0.0175) and 0.4329 (0.0163) for the yen. The estimates are in line with the results obtained using the Markov switching procedure, allowing to conclude against the evidence of spurious long memory and validating the use of step functions to model break processes.

It is interesting to note that our finding of long memory in the break-free series (Table 4) is in contrast with the finding of Granger and Hyung (1999), although a possible explanation for this discrepancy can be found in the use of the log periodogram estimator and in the selection of a relatively low number of ordinates of the periodogram ($m = \sqrt{T}$ in their study). It can be concluded that, possibly, both long memory and structural change are characteristics of the data generating process of the variance series considered, and,

[11] We are grateful to a referee for pointing this to us.

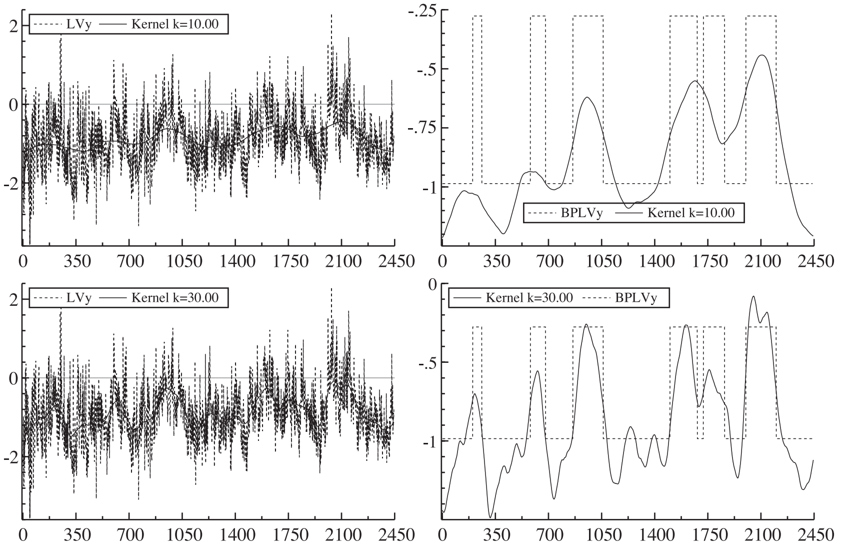

Fig. 7. Yen/US$ log variance process (LVdm), Kernel interpolator (Kernel) and Markov switching break process (BPLVdm).

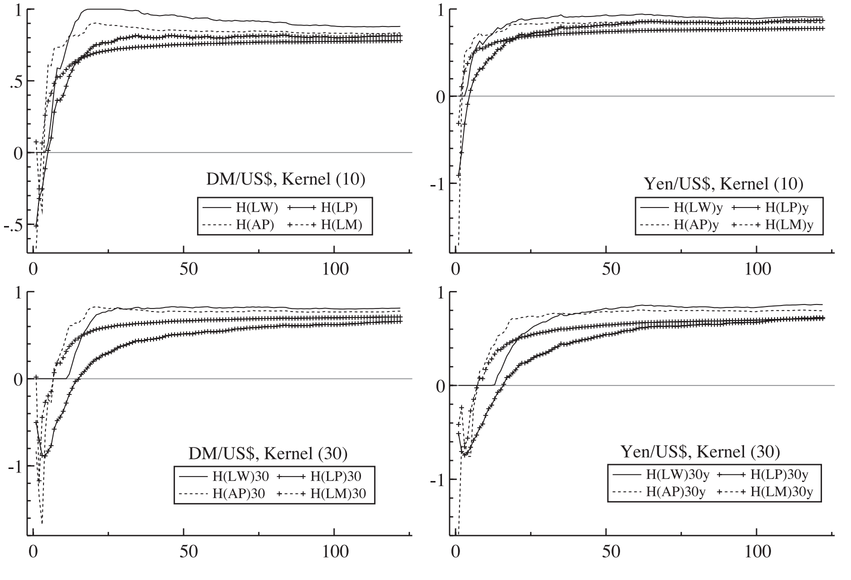

Fig. 8. Break-free processes: Hurst exponent (LW: local Whittle, LP: log periodogram, AP: averaged periodogram, LM: Lagrange multiplier). The bandwidth ranges between 10 and 1220 (each point on the *x*-axis corresponds to 10 periodogram ordinates).

Table 4

Panel A: Stationarity and semiparametric tests: break-free series (DM/US$)

	$\ln\sigma^2_{t\mathrm{IT}}$(DM/US$)	$\ln\sigma^2_{t\mathrm{B}}$(DM/US$)	$\ln\sigma^2_{t\mathrm{MS}}$(DM/US$)
H_{LW}	0.731 (0.057)	0.781 (0.059)	0.835 (0.047)
H_{LP}	0.918 (0.010)	0.935 (0.012)	0.940 (0.007)
H_{AP}	0.796 (–)	0.832 (–)	0.879 (–)
H_{LM}	0.771 (–)	0.773 (–)	0.758 (–)
H_{A}	0.863 (0.016)	0.869 (0.016)	0.866 (0.016)
W	23.3 (0.000)	96.3 (0.000)	48.1 (0.000)
LM	54.7 (0.000)	465.3 (0.000)	204.6 (0.000)
LR	40.6 (0.000)	256.4 (0.000)	118.7 (0.000)

Panel B: Stationarity and semiparametric tests: break-free series (Yen/US$)

	$\ln\sigma^2_{t\mathrm{IT}}$(Yen/US$)	$\ln\sigma^2_{t\mathrm{B}}$(Yen/US$)	$\ln\sigma^2_{t\mathrm{MS}}$(Yen/US$)
H_{LW}	0.637 (0.058)	0.679 (0.060)	0.737 (0.040)
H_{LP}	0.685 (0.010)	0.727 (0.011)	0.829 (0.004)
H_{AP}	0.612 (0.065)	0.659 (0.059)	0.818 (0.009)
H_{LM}	0.674 (–)	0.715 (–)	0.731 (–)
H_{A}	0.898 (0.017)	0.911 (0.017)	0.911 (0.017)
W	23.3 (0.000)	43.9 (0.000)	42.5 (0.000)
LM	54.7 (0.000)	135.6 (0.000)	146.2 (0.000)
LR	40.6 (0.000)	88.9 (0.000)	74.3 (0.000)

The table shows stationarity and semiparametric tests, with standard errors or p-values in brackets, for the daily break-free realized variance (σ^2_{tj}) and log realized variance ($\ln\sigma^2_{tj}$) of the Deutsche mark–US dollar and Yen–US dollar exchange rates, using the Inclan–Tiao (j=IT) and Bai (j=B) approaches. H_i i=LW, LP, AP, LM, A, is the estimate of the Hurst exponent using the Gaussian semiparametric estimator, the log periodogram estimator, the averaged periodogram estimator, the LM estimator, and the ARFIMA(0,d,0) model. LM is the LM test; W is the Wald test; LR is the Likelihood Ratio test. In the implementation, we used q=0.5 for the averaged periodogram estimator and l=0 for the semiparametric estimators.

therefore, structural change cannot be taken as the only explanation of volatility persistence.

5. A forecast-based comparison

A conclusion that can be drawn from the stationarity analysis is that a source of long memory in the realized variance series is unaccounted structural change. The evidence of long memory is weaker once structural change is accounted for. As suggested by Diebold and Inoue (1999), if the true data generating process is not characterized by long memory, an ARFIMA model may be useful for forecasting purposes. It may be of some interest, therefore, to compare the performance of models which allows for long memory and no structural change with models that explicitly account for structural breaks and assume that break-free volatility is a weak dependent process or a long-memory process. Among the approaches employed in this paper, only the Markov switching model allows for ex ante forecasting of the break process. In what follows, therefore, we have restricted the analysis to this latter approach.

Two different Markov switching models have been estimated. The first model allows the break-free log realized variance processes to follow a weak dependent process, while in the second model, the break-free log variance follows a fractional white noise process. By comparing the forecasting performance of these two models, it is possible to asses whether modelling long memory may be avoided; that is, whether a standard autoregressive model may be useful for volatility forecasting once structural change is accounted for. The forecasting performance of the Markov switching models has then been compared with that of a standard ARFIMA(0,d,0) model. This latter comparison is important to assess whether it is indeed important to model the break process or whether a long-memory model may be appropriate also when long-range dependence is spurious.

The forecasts from the Markov switching models have been obtained in two steps. Firstly, the Markov switching daily break process estimated from the monthly model is forecasted; secondly, the break-free volatility process is forecasted using an ARFIMA(0,d,0) model and an AR(8) model for the mark, and an AR(10) for the yen. As shown in Tables 5 and 6, the AR specification selected are appropriate for the data at hand according to standard specification tests. From the tables, it can be noticed that the switching models provide a better representation of the volatility DGP according to the specification tests, particularly the MS–AR model. The models were re-estimated using the first 1200 observations in the sample, using the last 1245 observations for model

Table 5
Models for daily log variance, DM/US$

	MS–AR	ARFIMA	MS–ARFIMA
μ_1	− 1.265	− 0.974 (0.368)	− 1.265
μ_2	− 0.562		− 0.562
D		0.413 (0.015)	0.366 (0.016)
π_1	0.373 (0.024)		
π_2	0.100 (0.023)		
π_3	0.034 (0.023)		
π_4	0.059 (0.022)		
π_5	0.094 (0.022)		
π_6	− 0.023 (0.022)		
π_7	0.044 (0.020)		
π_8	0.062 (0.020)		
Normality	0.000	0.000	0.000
$Q(20)$	0.003	0.000	0.003
$Q(40)$	0.211	0.012	0.054
$Q^2(20)$	0.000	0.000	0.000
$Q^2(40)$	0.000	0.000	0.000

The table shows Maximum Likelihood Estimates of Markov switching and ARFIMA models for the daily log variance of the Deutsche mark–US dollar exchange rate from December 1, 1986, to December 1, 1996, for a total of 2445 observations. Robust standard errors are reported in parenthesis. The Ljung–Box portmanteau tests for up to kth-order serial correlation in the standardized residuals (squared standardized residuals) is denoted by $Q(k)$ [$Q^2(k)$]. Normality is the Bera–Jarque test for normality. The table also shows the p-values of the various tests. MS–AR denotes the two-step Markov switching autoregressive model, ARFIMA denotes the standard ARFIMA model and MS–ARFIMA denotes the two-step Markov switching ARFIMA model.

Table 6
Models for daily log variance, Yen/US$

	MS–AR	ARFIMA	MS–ARFIMA
μ_1	− 1.142	− 0.956 (0.610)	− 1.142
μ_2	− 0.433		− 0.433
D		0.447 (0.016)	0.411 (0.017)
π_1	0.432 (0.0233)		
π_2	0.106 (0.025)		
π_3	0.020 (0.024)		
π_4	0.045 (0.025)		
π_5	0.072 (0.024)		
π_6	− 0.009 (0.023)		
π_7	− 0.006 (0.023)		
π_8	− 0.006 (0.021)		
π_9	− 0.004 (0.023)		
π_{10}	0.0471 (0.021)		
Normality	0.000	0.000	0.000
$Q(20)$	0.314	0.133	0.087
$Q(40)$	0.518	0.342	0.238
$Q^2(20)$	0.000	0.000	0.000
$Q^2(40)$	0.000	0.000	0.000

The table shows Maximum Likelihood Estimates of Markov switching and ARFIMA models for the daily log variance of the Yen–US dollar exchange rate from December 1, 1986, to December 1, 1996, for a total of 2445 observations. Robust standard errors are reported in parenthesis. The Ljung–Box portmanteau tests for up to kth-order serial correlation in the standardized residuals (squared standardized residuals) is denoted by Q(k) [$Q^2(k)$]. Normality is the Bera–Jarque test for normality. The table also shows the p-values of the various tests. MS–AR denotes the two-step Markov switching autoregressive model, ARFIMA denotes the standard ARFIMA model and MS–ARFIMA denotes the two-step Markov switching ARFIMA model.

evaluation.[12] The out-of-sample forecasting exercise therefore considers 1245 one-step-ahead forecasts, 1241 five-step-ahead forecast and 1235 ten-step-ahead forecasts. The forecasting performance of the models has been evaluated by means of regressions of the actual values on the forecasts of each model and a forecast encompassing regression of the actual values on the forecast of all the models, constraining the slope parameters to sum to one. According to the first criterion, an accurate forecasting model should show a zero intercept term and a slope parameter equal to one. The R^2 of this regression measures the degree of linear association between the actual and forecasted values. On the other hand, the forecasting encompassing criterion evaluates whether it is possible to improve the forecasting performance by generating forecasts from a combined model.

As shown in Table 7, all the models show a similar forecasting performance at the one-step horizon in terms of intercept and slope parameters and R^2. In none of the cases, the intercept and the slope parameters are statistically different from the expected values, and the R^2 are very close. According to the forecast encompassing regression, however,

[12] See Hamilton (1994) for a description of an alghorithm to compute filtered probabilities, employed in the paper to generate the out-of-sample forecasts from the Markov switching model.

Table 7
Out-of-sample forecasting analysis

	MS–A	ARFIMA	MS–AR	MS–A	ARFIMA	MS–AR
1 Step						
α	−0.073	−0.071	−0.102	−0.012	−0.018	−0.060
	(0.037)	(0.034)	(0.035)	(0.022)	(0.021)	(0.027)
β	0.969	1.027	1.026	1.004	0.933	1.108
	(0.073)	(0.072)	(0.070)	(0.050)	(0.045)	(0.062)
R^2	0.443	0.451	0.451	0.338	0.332	0.350
δ_I	−0.216	0.276	0.940	−0.067	0.226	0.843
	(0.310)	(0.253)	(0.249)	(0.256)	(0.187)	(0.146)
5 Steps						
α	−0.311	−0.393	−0.558	0.070	−0.038	0.0436
	(0.202)	(0.204)	(0.206)	(0.173)	(0.184)	(0.210)
β	0.976	1.071	1.072	0.980	0.948	1.197
	(0.088)	(0.092)	(0.088)	(0.078)	(0.075)	(0.099)
R^2	0.498	0.482	0.508	0.430	0.403	0.439
δ_I	0.412	−0.200	0.788	0.651	−0.102	0.451
	(0.191)	(0.155)	(0.157)	(0.148)	(0.114)	(0.076)
10 Steps						
α	−0.660	−0.926	−1.418	0.297	0.065	−0.969
	(0.454)	(0.507)	(0.510)	(0.412)	(0.474)	(0.497)
β	0.993	1.111	1.139	0.961	0.938	1.228
	(0.099)	(0.110)	(0.110)	(0.093)	(0.095)	(0.117)
R^2	0.512	0.478	0.532	0.432	0.337	0.448
δ_I	0.625	−0.471	0.846	0.863	−0.338	0.477
	(0.151)	(0.126)	(0.122)	(0.117)	(0.091)	(0.060)

The table shows the results of the forecasting analysis. Left-hand side columns are for the DM/US$, right-hand side columns are for the Yen/US$. α and β are the constant and the slope parameters in the regression of the actual values on the forecasts (with Newey–West standard errors), and R^2 is the coefficient of determination of the regression. δ_i are the parameters of the forecast encompassing regression $y_t = \delta_1 \hat{y}_{\mathrm{MS},t} + \delta_2 \hat{y}_{\mathrm{A},t} + \delta_3 \hat{y}_{\mathrm{MSA},t}$ with $\Sigma_i \delta_i = 1$. MS–AR denotes the two-step Markov switching autoregressive model, ARFIMA denotes the standard ARFIMA model and MS–A denotes the two-step Markov switching ARFIMA model.

the MS–AR model would seem to be the preferred model. At longer forecasting horizon (5–10 steps ahead), the Markov switching models seems to perform better than the standard ARFIMA model in terms of R^2 and forecast encompassing criterion. Interestingly, the forecast encompassing criterion does not allow to select a clear winner between the MS–AR and the MS–ARFIMA models at the 5–10-step-ahead horizon, although the MS–ARFIMA model shows a better performance than the MS–AR model in terms of bias.

Overall, it can be concluded that modelling the break process is not important for daily volatility forecasting once it is allowed for a long-memory component in the model, and that neglecting long memory leads to an inferior forecasting model at longer forecasting horizons. The conclusion of Diebold and Inoue (1999) does seem to have empirical support for a short daily horizon, although we find that long memory seems to be a feature of the volatility data generating process.

6. Conclusions

We have investigated the memory properties of the realized volatility process introduced by Andersen et al. (1998) for the mark/dollar and yen/dollar exchange rates over the period 1986–1996. Our first aim was to evaluate whether the evidence for long memory uncovered by various papers may be considered to some extent spurious. An analysis conducted with various semiparametric models has produced evidence for structural breaks. Tests show that the evidence of long memory is weaker once structural change is accounted for, but long memory is clearly a feature of the data generating volatility process. As far as explaining volatility is concerned, therefore, both structural breaks and long memory are important features of the data. Forecasting is however a partly different story. The existence of structural breaks implies the necessity to forecast the next structural break, a task by no means simple given the relatively infrequent number of breaks present in the volatility data. We have therefore considered the Markov switching model proposed by Hamilton (1999) as a tool for predicting the timing of the breaks as well as taking into account the change in the coefficients. We have considered both a simple autoregressive version and a more sophisticated long-memory version in order to obtain more evidence about the relevance of long memory.

The results of the forecasting analysis provide some empirical support to the conclusion of Diebold and Inoue (1999) that a pure long-memory model may be useful for forecasting purposes, even when the true data generating process shows weak dependence. In fact, we have found that modelling the break process is not important for very short term forecasting (one-step ahead) once the model considers a long-memory component. However, at longer forecasting horizons (5–10 steps ahead), accounting for both long memory and structural change leads to a superior forecasting model.

Our conclusion is that structural changes in currency volatility are important for statistical analysis. In our opinion, our results raise interesting implications for theoretical analysis as well. In many cases, the goal of the researcher is not forecasting a time series but rather trying to understand the causes of the past evolution, e.g., for assessing what is the most relevant economic theory. In the theory of rational expectations, agents are assumed to know the mapping between exogenous and endogenous variables. In the presence of structural breaks, it may be hard to believe that agents have the possibility to use the true mapping in order to form expectations, as Lucas (1986) himself pointed out. The existence of structural breaks in past data and the relevance of structural breaks for forecasting provides strong evidence in favor of the Theory of Rational Beliefs proposed by Kurz (1997) in a series of papers. The Theory of Rational Beliefs generalizes rational expectations in the direction of allowing for the perception of nonstationarity in the data generation processes. Under these conditions, agents may rationally select time-varying forecasting models which are all consistent with the long-run evidence. Based on our results, this theory is likely to be essential in understanding the dynamics of the foreign currency market. We suspect that there are many other markets characterized by similar findings, although unfortunately, few other markets may be studied with data sets as rich as the one we have used here.

Acknowledgements

The authors are grateful to conference participants, the editor, R.T. Bailie, an anonymous referee for constructive comments and to T. Andersen for providing us with their series of daily volatility obtained from the Olsen and Associates high-frequency data set. The usual disclaimer applies. Financial support from MURST, Fondi Ex-60% (Morana), Centro di Economia Monetaria e Finanziaria Paolo Baffi and Bocconi University (Beltratti) are gratefully acknowledged.

Appendix A. Stationarity test and semiparametric methodologies

Given the $I(d)$ process x_t, the Hurst exponent is estimated using the approach of Kunsch (1987) and Robinson (1995b). This requires to minimize the following objective function:

$$Q(C,H_{\mathrm{LW}}) = \frac{1}{m}\sum_{j=1}^{m}\left(\log C\lambda_j^{1-2H} + \frac{\lambda_j^{2H-1}}{C}I(\lambda_j)\right)$$

where $I(\lambda_j)$ is the periodogram at frequency $\lambda_j = 2\pi j/T$, $j = 1,\ldots,m$, m is the bandwidth parameter, C is a positive constant. For $H < 0.5$, the process is anti-persistent; for $H > 0.5$, it is long memory; and for $H = 0.5$, it is weakly dependent. It is shown that

$$\sqrt{m}(\hat{H}_{\mathrm{LW}} - H) \xrightarrow{d} N\left(0,\frac{1}{4}\right).$$

An alternative consistent estimator for H, proposed by Robinson (1998), is

$$\hat{H}_{\mathrm{LM}} = \frac{\sum_{j=1}^{m}(1-2v_j)I(\lambda_j)}{\sum_{j=1}^{m}(2-2v_j)I(\lambda_j)}$$

where $v_j = \log j - \frac{1}{m}\sum_{j=1}^{m}\log j$. Once H is estimated, an estimate of d can be obtained from the relation $d = H - 1/2$. The corresponding standard error is $\hat{\sigma}_{\hat{d}} = \hat{\sigma}_{\hat{H}}$. We denote this estimator as H_{LM} because it can be derived from the LM test of Lobato and Robinson (1997).

A consistent but less efficient estimate of the fractional differencing parameter can be obtained by the log periodogram regression of Geweke and Porter-Hudak (1983), as modified by Robinson (1995a,b)

$$\ln I(\lambda_j) = c + d(-2\log\lambda_j) + \mu_j \qquad j = l,\ldots,m$$

where l is a trimming parameter. It is shown that

$$\sqrt{m}(\hat{d}_{\mathrm{LP}} - d) \xrightarrow{d} N\left(0, \frac{\pi^2}{24}\right).$$

Another estimator widely used in empirical work is obtained from the averaged periodogram (Robinson, 1994a; Lobato and Robinson, 1996)

$$\hat{H}_{\mathrm{AP},q} = 1 - \frac{1}{2\ln q}\ln\left\{\frac{\hat{F}(qm)}{\hat{F}(\lambda m)}\right\}$$

where $\hat{F}(\lambda) = \frac{2\pi}{n}\sum_{j=1}^{[\lambda n/2\pi]} I(\lambda_j)$. The limiting distribution of $H_{\mathrm{AP},q}$ is

$$m^{1/2}(\hat{H}_{\mathrm{AP},q} - H) \xrightarrow{d} N\left(0, \frac{(1+q^{-1}-2q^{1-2H})}{(\ln q)^2}\,\frac{(1-H)^2}{(3-4H)}\right)$$

for $1/2 \leq H \leq 3/4$, and

$$m^{2-2H}(\hat{H}_{\mathrm{AP},q} - H) \xrightarrow{d} N\left(0, \frac{(1-q^{2H-2})}{(\ln q)}\,\frac{(1-H)\Gamma(2(1-H))\cos((1-H)\pi)}{(2\pi)^{2-2H}}P\right)$$

as $T \to \infty$, where P is a random variable with unknown distribution, for $3/4 < H < 1$.

The LM test for the null H_0:$H = 0.5$ against a two-sided alternative is

$$\mathrm{LM} = m\left(\frac{\sum_{j=1}^{m} v_j I(\lambda_{jT})}{\sum_{j=1}^{m} I(\lambda_{jT})}\right)^2 \xrightarrow{d} \varkappa^2_{(1)}$$

The Wald test is

$$W = 4m(\hat{H}_{\mathrm{LM}} - H)^2 \xrightarrow{d} \varkappa^2_{(1)}$$

and the Likelihood Ratio test is

$$\mathrm{LR} = 2m\left\{\log\frac{\sum_{j=1}^{m} I(\lambda_{jT})}{\sum_{j=1}^{m} \lambda_j^{2H-1} I(\lambda_{jT})} + (2\hat{H}_{\mathrm{LM}} - 1)\frac{1}{m}\sum_{j=1}^{m}\log\lambda_j\right\} \xrightarrow{d} \varkappa^2_{(1)}$$

Appendix B. Bandwidth selection

A theory of optimal bandwidth selection, based on the minimization of the asymptotic MSE, has been proposed for various semiparametric estimators (Robinson, 1994a,b;

Henry and Robinson, 1996; Hurvich et al., 1998), and efforts have been made to make it implementable (Delgado and Robinson, 1996; Henry, in press). Following Henry (in press), the optimal bandwidths can be stated as follows:

$$m^{*}_{\mathrm{LW}} = \left(\frac{3}{4\pi}\right)^{4/5}\left|\tau^{*} + \frac{d_x}{12}\right|^{-2/5} T^{4/5}$$

for the Gaussian semiparametric estimator and $-1/2 < d < 1/2$;

$$m_{\mathrm{LP}} = \left(\frac{27}{512\pi^2}\right)^{1/5} |\tau^{*}|^{-2/5} T^{4/5}$$

for the log periodogram estimator, $-1/2 < d < 1/2$ and $\tau^{*} \neq 0$;

$$m_{\mathrm{AP}_1} = \left(\frac{(3-2d_x)^2}{4(2\pi)^4(1-4d_x)}\right)^{1/5}\left|\tau^{*} + \frac{d_x}{12}\right|^{-2/5} T^{4/5}$$

for the averaged periodogram estimator with $0 < d < 1/4$, while for $1/4 < d < 1/2$

$$\begin{aligned} m_{\mathrm{AP}_2} = \frac{T^{\frac{2}{3-2d}}}{2\pi}\Bigg\{ & \frac{2\Gamma(1-2d_x)\cos\left(\left(\frac{1}{2}-d_x\right)\pi\right)(3-2d_x)}{8}\Bigg(\frac{2d_x-3}{2d_x(\tau^{*}+\frac{d}{12})} + \frac{1}{\left|\tau^{*}+\frac{d_x}{12}\right|} \\ & \times \Bigg[\frac{(3/2d_x)^2}{(2d_x)^2(d_x+\frac{1}{2})^2} + 32\left(\frac{1}{2}-d_x\right)\Bigg\{\frac{1}{2d_x(4d-1)} - \frac{1}{(4d_x+1)(d_x+\frac{1}{2})^2} \\ & - \frac{4\Gamma(2d_x)^2}{\Gamma(4d_x+2)}\Bigg\}\Bigg]^{\frac{1}{2}}\Bigg)\Bigg\}^{\frac{1}{3-2d_x}}; \end{aligned}$$

finally, for the LM estimator proposed by Robinson (1998), we used

$$m_{\mathrm{LM}} = \left(\frac{3}{4\pi}\right)^{4/5} |\tau^{*}|^{-2/5} T^{4/5}.$$

Optimal bandwidths and the corresponding estimates of d are then derived together using the recursion

$$\begin{aligned} \hat{d}_x^{(k)} &= d(\hat{m}^{(k)}), \\ \hat{m}^{(k+1)} &= m_i(\hat{d}_x^{(k)}, \tau^{*}) \end{aligned}$$

$i =$ LW,AW, for the Gaussian semiparametric estimator and the averaged periodogram estimator, while for the log periodogram estimator and the LM estimator, the formula simply is $\hat{d}_x = d(\hat{m}_j)$ $j =$ LP,LM. The recursions were started setting $\hat{m}^{(0)} = [T^{4/5}]$ and an estimate of the smoothness parameter $\tau^{*}=(f^{*\prime\prime}(0)/2f^{*}(0))$, where the function $f^{*}(\cdot)$ describes a short-term correlation structure in the spectral density of the process $f(\lambda)=|1-\exp(i\lambda)|^{1-2\hat{H}}f^{*}(\lambda)$ was obtained using the least square regression suggested by Delgado and Robinson (1996)

$$I(\lambda_j) = \sum_{k=0}^{2} Z_{jk}(\hat{H})\hat{\beta}_k + \hat{\varepsilon}_j \qquad j = 1,\ldots,\hat{m}^{(0)}$$

where $Z_{jk}(\hat{\mathrm{H}}) = |1-\exp(i\lambda_j)|^{1-2\hat{\mathrm{H}}}\lambda_j^k/k!$ τ^{*} is then estimated by $\hat{\beta}_2/2\hat{\beta}_0$.

Appendix C. Break process estimation

Inclan and Tiao (1994) have shown that, under the null of homogeneous unconditional variance, the centered and normalized cumulative sum of squares $C_k/(C_T)-k/T$ has the following limiting distribution

$$D_k^{*} = \sqrt{T/2}\left(\frac{C_k}{C_T}-\frac{k}{T}\right) \xrightarrow{d} W^0 \qquad k = 1,\ldots,T$$

where W^0 is a Brownian bridge, $C_k=\sum_{t=1}^{k} r_t^2$, $C_T=\sum_{t=1}^{T} r_t^2$, $r_t \sim \mathrm{NID}(0,\sigma^2)$ is the return, k/T is the break fraction, and T is the sample size. Hence, from the limiting distribution of D_k^{*}, the theoretical quantiles of the test statistic of interest, $\max_k|D_k^{*}|$, can be derived. Inclan and Tiao (1994) find that the asymptotic 5% critical value (1.358) is a good approximation for sample size larger than 200 observations. We have therefore used this critical value in our study. The procedure suggests to determine the dating of the potential structural breaks in an iterative way. In the first run, the entire sample is explored and the statistic D_k^{*} is computed for $k=1,\ldots,T$. Let us denote the maximum value of the statistic over the explored range $\max_k|D_k^{*}|_i$, assuming that the maximum is attained in correspondence of observation i. Then, a first break point is determined in correspondence of observation i if the statistic $\max_k|D_k^{*}|_i$ exceeds the critical value for the selected significance level. Then, in successive runs, the subsamples determined by the break point $1,\ldots,i$ and $i+1,\ldots,T$ are investigated and new break points are determined following the same procedure. For each of the break points found, the approach is repeated until the null of no structural change cannot be rejected. In the implementation of the procedure, we only explored subsamples of size larger than 50 observations (see Inclan and Tiao, 1994 for further details).

Bai (1997) has suggested a similar iterative procedure useful to detect structural breaks in the mean of a process. Therefore, to implement the test, we used the realized variance processes (y_t). With reference to the model

$$\begin{aligned} y_t &= \mu_i + x_t \\ x_t &= \sum_{j=0}^{\infty} a_j \varepsilon_{t-j}, \end{aligned}$$

the suggested test statistic is

$$\sup F_T = \sup_{T\eta \le \tau \le T(1-\eta)} \frac{\bar{S}_T - S_T(k)}{\hat{\sigma}^2_\varepsilon} \xrightarrow{d} \sup_{T\eta \le \tau \le T(1-\eta)} \frac{|B(\tau) - \tau B(1)|}{\tau(1-\tau)},$$

where $B(\cdot)$ is standard Brownian motion on [0,1],$\eta \in (0,1/2)$, $\tau = k/(T)$ is the break fraction, $\bar{S}_T = \sum_{t=1}^{T}(y_t - \bar{y})^2$ is the residual sum of squares computed over the entire sample, $\bar{y}_k$ is the sample mean for the first k observations, $\bar{y}^*_k$ is the sample mean for the last $T-k$ observations, $S_T(k) = \sum_{t=1}^{k}(y_t - \bar{y}_k)^2 + \sum_{t=k+1}^{T}(y_t - \bar{y}^*_k)^2$ is the residual sum of squares for the entire sample, allowing for the break in the mean, and $\hat{\sigma}^2_\varepsilon$ is a consistent estimator of $a(1)^2\sigma^2_\varepsilon$. The algorithm works as follows.

In the first run, the entire sample is explored and the statistic $(\bar{S}_T - S_T(k)/\hat{\sigma}^2_\varepsilon)$ is computed for $k=1,\ldots,T$. Let us denote the maximum value of the statistic over the explored range, neglecting a proportion η of observations at the beginning and at the end of the sample, $F_{Ti} = \sup_{T\eta \le \tau \le (1-\eta)} (\bar{S}_T - S_T(k)/\hat{\sigma}^2_\varepsilon)$, assuming that the maximum is attained in correspondence of observation i. Then, a first break point is determined in correspondence of observation i if the statistic sup F_{Ti} exceeds the critical value for the selected significance level. Then, in successive runs, the subsamples determined by the break point $1,\ldots,i$ and $i+1,\ldots,T$ are investigated as described above.[13] For each of the break points found, the procedure is repeated until the null of no structural change cannot be rejected. As for the Inclan and Tiao (1994) method, we only explored subsamples of size larger than 50 observations. Following Bai and Perron (1998), we set $\eta = 0.05$. The 5% critical value is 9.63, and is taken from Table II of Bai and Perron (1998). A consistent estimate of the error variance was obtained by allowing as many breaks as detected by the Inclan and Tiao (1994) approach.

Once the number and dating of the breaks has been determined, a break-free realized variance series can be computed as follows. Firstly, compute dummy variables (d_k) according to the dating of the breaks. Then, run the OLS regression

$$\sigma^2_{t,j} = \alpha + \sum_{k=1}^{p} \theta_k d_k + \varepsilon_{t,\sigma^2} \quad j = \text{DM, Yen}$$

using the realized variance process as dependent variable. The estimated break process is then simply the fitted value in the regression above, i.e., $\hat{\mu}_{t,j} = \hat{\sigma}^2_{t,j}$. Following Granger and Hyung (1999), the break-free log variance series have been obtained as $(\ln\sigma^2_{t,j})_{bf} = \ln\sigma^2_{t,j} - \ln\hat{\sigma}^2_{t,j}$.

Finally, we have estimated the break process using a Markov switching model. Consider a two-regime model ($k=2$) for the unconditional mean and let $\boldsymbol{\eta}$ be a vector consisting of the mean elements μ_1, μ_2 and the variance of the error process in the model

[13] For instance, assuming that the first break point occurs in correspondence of observation i, the test is repeated considering the following residual sum of squares: $\bar{S}_i = \sum_{t=1}^{i}(y_t - \bar{y}_i)^2$ and $S_i(k) = \sum_{t=1}^{k}(y_t - \bar{y}_{i,k})^2 + \sum_{t=k+1}^{i}(y_t - \bar{y}^*_i)^2$ for the first subsample, and $\bar{S}_{i+1} = \sum_{t=i+1}^{T}(y_t - \bar{y}_{i+1})^2$ and $S_{i+1}(k) = \sum_{t=i+1}^{k}(y_t - \bar{y}_{i+1,k})^2 + \sum_{t=k+1}^{T}(y_t - \bar{y}^*_{i+1})^2$ for the second subsample, where $\bar{y}_j$, $\bar{y}_{j,k}$ and $\bar{y}^*_k$ are the means computed over the corresponding subsamples.

$y_t = \mu_{s_t} + \varepsilon_t$, with $\varepsilon_t \sim N(0,\sigma^2)$. The transition between states is governed by a Markov chain whose realizations take on values in $\{1,\ldots,k\}$, $p(s_t = j | s_{t-1} = i) = p_{ij}$, with $\sum_{j=1}^{k} p_{ij} = 1$. Let $\boldsymbol{p}=(p_{11},p_{12},\ldots,p_{kk})'$ the $(k^2 \times 1)$ vector of transition probabilities. The econometrician is supposed to observe only the realizations of the variable y_t but not of the state s_t. The unknown parameters can be collected in the vector $\boldsymbol{\lambda}=(\boldsymbol{p}',\boldsymbol{\eta}')'$ and maximum likelihood estimates of the parameters of the model can be obtained via the Expectation–Maximization algorithm. See Hamilton (1990) for further details.

The break process is then computed as

$$\hat{\sigma}^2_{t,j} = \hat{\mu}_{t,j} = \sum_{s=1}^{k} \hat{p}_{t,s,j}\hat{\mu}_{s,j}$$

where $\hat{p}_{t,s,j}$ is the estimated probability that the observation t of process j belongs to state s and $\hat{\mu}_{s,j}$ is the estimated value of the mean in the sth state. In the empirical implementation, we have used at most three regimes. The multiregime model should allow to identify different equilibrium states for the log realized volatility process, namely, a low-, average- and high-volatility state.

References

Andersen, T.G., Bollerslev, T., 1997. Heterogeneous 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., Diebold, F.X., Labys, P., 1998. The distribution of exchange rate volatility, NBER Working paper, no. 6961.

Bai, J., 1997. Estimating multiple breaks one at a time. Econometric Theory 14, 315–352.

Bai, J., Perron, P., 1998. Estimating and testing linear model with multiple structural changes. Econometrica 66, 47–78.

Baillie, R.T., Bollerslev, T., Mikkelsen, O.H., 1996. Fractionally integrated generalised autoregressive conditional heteroskedasticity. Journal of Econometrics 74, 3–30.

Beine, M., Laurent, S., 2000. Structural Change and Long Memory in Volatility: New Evidence from Daily Exchange Rates. University of Liège manuscript, Liège.

Beltratti, A., Morana, C., 1999. Computing value at risk with high frequency data. Journal of Empirical Finance 6, 431–455.

Beltratti, A., Morana, C., 2001. Comparing Realized Variance and GARCH Models, Quaderni di Ricerca, Centro di Economia Monetaria e Finanziaria Paolo Baffi. Bocconi University, Milan.

Beran, J., 1999. SEMIFAR Models—A Semiparametric Framework for Modelling Trends, Long Range Dependence and Nonstationarity. University of Kostanz manuscript, Kostanz.

Bollerslev, T., Mikkelsen, H.O., 1996. Modeling and pricing long memory in stock market volatility. Journal of Econometrics 73, 151–184.

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.

Cai, J., 1994. A Markov model of switching regime ARCH. Journal of Business and Economic Statistics 12 (3), 309–316.

Chambers, M., 1998. Long memory and aggregation in macroeconomic time series. International Economic Review 39, 1053–1072.

Dacorogna, M.M., Muller, U.A., Nagler, R.J., Olsen, R.B., Pictet, O.V., 1993. A geographical model for the daily and weekly seasonal volatility in the foreign exchange market. Journal of International Money and Finance 12, 413–438.

Delgado, M.A., Robinson, P.M., 1996. Optimal spectral bandwidth for long memory. Statistica Sinica 6, 97–112.
Diebold, F., 1986. Comment on "Modeling the persistence of conditional variance", by R. Engle and T. Bollerslev. Econometric Reviews 5, 51–56.
Diebold, F.X., Inoue, A., 1999. Long Memory and Structural Change, manuscript. New York University, Stern School of Business, Department of Finance.
Ding, Z., Granger, C.W.J., 1995. Some properties of absolute return. An alternative measure of risk. Annales d'Economie et de Statistique 40, 67–91.
Ding, Z., Granger, C.W.J., 1996a. Varieties of long memory models. Journal of Econometrics 73, 61–77.
Ding, Z., Granger, C.W.J., 1996b. Modeling volatility persistence of speculative returns: a new approach. Journal of Econometrics 73, 185–215.
Ding, Z., Granger, C.W.J., Engle, R.F., 1993. A long memory property of stock market returns and a new model. Journal of Empirical Finance 1, 83–106.
Duker, M.J., 1997. Markov switching in GARCH processes and mean-reverting stock market volatility. Journal of Business and Economic Statistics 15 (1), 26–34.
Eibens, H., 1999. Realized stock volatility, mimeo. The Johns Hopkins University, Department of Economics.
Geweke, J., Porter-Hudak, S., 1983. The estimation and application of long memory time series. Journal of Time Series Analysis 4, 221–238.
Granger, C.W.J., Hyung, N., 1999. Occasional structural breaks and long memory, mimeo. University of California, San Diego, Department of Economics.
Granger, C.W.J., Marmol, F., 1997. The correlogram of a long memory process plus a simple noise, mimeo. University of California, San Diego, Department of Economics.
Granger, C.W.J., Terasvirta, T., 1999. A simple nonlinear time series model with misleading linear properties. Economics Letters 62, 161–165.
Hamilton, J.D., 1990. Analysis of time series subject to changes in regime. Journal of Econometrics 45, 39–70.
Hamilton, J.D., 1994. Time Series Analysis. Princeton Univ. Press, Princeton, USA.
Hamilton, J.D., 1999. A new approach to the economic analysis of nonstationary time series and the business cycle. Econometrica 57 (2), 357–384.
Hamilton, J.D., Susmel, R., 1994. Autoregressive conditional heteroskedasticity and changes in regime. Journal of Econometrics 64, 307–333.
Henry, M., 2001. Robust automatic bandwidth for long memory. Journal of Time Series Analysis 22 (3), 293–316.
Henry, M., Payne, R., 1997. An investigation of long range dependence in intra-day foreign exchange rate volatility, LSE Financial Market Group, Discussion Paper n. 264.
Henry, M., Robinson, P.M., 1996. Bandwidth choice in Gaussian semiparametric estimation of long range dependence. In: Robinson, P.M., Rosemblatt, M. (Eds.), Athens Conference in Applied Probability and Time Series Analysis: Volume II: Time Series Analysis. In memory of E.J. Hannan. Springer-Verlag, New York, pp. 220–232.
Hosking, J.R.M., 1996. Asymptotic distributions of the sample mean, autocovariance and autocorrelations of long memory processes. Journal of Econometrics 73, 261–284.
Hsu, C.C., 2001. Change point estimation in regression with $I(d)$ variables. Economics Letters 70, 147–155.
Hurvich, C.M., Beltrao, K.I., 1994. Automatic semiparametric estimation of the memory parameter of a long memory time series. Journal of Time Series Analysis 15, 285–302.
Hurvich, C.M., Deo, R., Brodsky, J., 1998. The mean squared error of Geweke and Porter–Hudak's estimator of the memory parameter of a long memory time series. Journal of Time Series Analysis 19, 19–46.
Inclan, C., Tiao, G.C., 1994. Use of cumulative sum of squares for retrospective detection of change in variance. Journal of the American Statistical Association 89, 913–923.
Kunsch, H.R., 1986. Discrimination between monotonic trends and long range dependence. Journal of Applied Probability 23, 1025–1030.
Kunsch, H.R., 1987. Statistical aspects of self similar processes. In: Prohorov, Y., Sazanov, V.V. (Eds.), Proceedings of the First World Congress of the Bernoulli Society, vol. 1. VNU Science Press, Utrecht, pp. 67–74.
Kurz, M., 1997. On Rational Belief Equilibria. In Endogenous Economic Fluctuations. Springer, Berlin.

Lamoreaux, C.G., Lastrapes, W.D., 1990. Persistence in variance, structural change and the GARCH model. Journal of Business and Economic Statistics 8, 225–234.

Liu, M., 2000. Modeling long memory in stock market volatility. Journal of Econometrics 99, 139–171.

Lobato, I.N., Robinson, P.M., 1996. Averaged periodogram estimation of long memory. Journal of Econometrics 73, 303–324.

Lobato, I.N., Robinson, P.M., 1997. A nonparametric test for $I(0)$. Review of Economic Studies 65, 475–496.

Lobato, I.N., Savin, N.E., 1998. Real and spurious long memory properties of stock market data. Journal of Business and Economic Statistics 16 (3), 261–268.

Lucas, R., 1986. Adaptive behavior and economic policy. Journal of Business 59, S401–S426.

Mikosch, T., Starica, C., 1998. Change of structure in financial time series, long range dependence and the GARCH model, mimeo. University of Groningen, Department of Mathematics.

Mikosch, T., Starica, C., 2000. Limit theory for the sample autocorrelations and extremes of a GARCH (1,1) process. Annals of Statistics 28, 1427–1451.

Muller, U.A., Dacorogna, R.B., Olsen, 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.

Muller, U.A., Dacorogna, M.M., Davé, R.D., Olsen, R.B., Pictet, O.V., von Weizsacker, J.E., 1997. Volatilities of different time resolutions—analyzing the dynamics of market components. Journal of Empirical Finance 4, 213–239.

Robinson, P.M., 1994a. Semiparametric analysis of long memory time series. Annals of Statistics 22, 515–539.

Robinson, P.M., 1994b. Rates of convergence and optimal spectral bandwidth for long range dependence. Probability Theory and Related Fields 99, 443–473.

Robinson, P.M., 1995a. Log periodogram regression of time series with long range dependence. The Annals of Statistics 23, 1048–1072.

Robinson, P.M., 1995b. Gaussian semiparametric estimation of long range dependence. The Annals of Statistics 23, 1630–1661.

Robinson, P.M., 1998. (Comment). Journal of Business and Economic Statistics 16 (3), 276–279.

Ryden, T., Terasvirta, T., Asbrink, S., 1998. Stylized facts of daily return series and the hidden Markov model. Journal of Applied Econometrics 13, 217–244.

So, M.K.P., Lam, K., Li, N.K., 1998. A stochastic volatility model with Markov switching. Journal of Business and Economic Statistics 16 (2), 244–253.

Sowell, F., 1992. Maximum likelihood estimation of stationary univariate fractionally integrated time series models. Journal of Econometrics 53, 165–188.

Taqqu, M.S., Teverovsky, V., 1998. Semi-parametric graphical estimation techniques for long memory data. In: Robinson, P.M., Rosemblatt, M. (Eds.), Time Series Analysis in Memory of E.J. Hannan. Sringer-Verlag, New York, pp. 420–432.

Timmerman, A., 2001. Structural breaks, incomplete information and stock prices. University of California, San Diego. Discussion Paper n. 2.