



ELSEVIER

Journal of Econometrics 98 (2000) 81–106

JOURNAL OF
Econometrics

www.elsevier.nl/locate/econbase

Semiparametric estimation of long-memory volatility dependencies: The role of high-frequency data

Tim Bollerslev^{a,b}, Jonathan H. Wright^{c,*}

^a*Department of Economics, Duke University, Durham, NC 27708-0097, USA*

^b*NBER, Cambridge MA 02138, USA*

^c*International Finance Division, Board of Governors of the Federal Reserve System, Washington, DC 20551, USA*

Received 1 April 1998; received in revised form 1 October 1999; accepted 1 October 1999

Abstract

Recent empirical studies have argued that the temporal dependencies in financial market volatility are best characterized by long memory, or fractionally integrated, time series models. Meanwhile, little is known about the properties of the semiparametric inference procedures underlying much of this empirical evidence. The simulations reported in the present paper demonstrate that, in contrast to log-periodogram regression estimates for the degree of fractional integration in the mean (where the span of the data is crucially important), the quality of the inference concerning long-memory dependencies in the conditional variance is intimately related to the sampling frequency of the data. Some new estimators that succinctly aggregate the information in higher frequency returns are also proposed. The theoretical findings are illustrated through the analysis of a ten-year time series consisting of more than *half-a-million* intradaily observations on the Japanese Yen–U.S. Dollar exchange rate. © 2000 Published by Elsevier Science S.A. All rights reserved.

JEL classification: C15; C22; F31

Keywords: Long memory; Log-periodogram regressions; Stochastic volatility; Temporal aggregation; High-frequency data; Exchange rates

* Corresponding author. Tel.: + 1-202-452-3696; fax: + 1-202-263-4843.

E-mail addresses: bollers@econ.duke.edu (T. Bollerslev), jonathan.h.wright@frb.gov (J.H. Wright).

1. Introduction

Following the work by Mandelbrot (1971), numerous studies have tested for evidence of slowly mean reverting dependencies in financial rates of returns. While none of these studies have found convincing evidence for long-memory dependencies in the level of returns (see e.g. Cheung and Lai, 1995; Hiemstra and Jones, 1997; Lo, 1991; Jacobsen, 1996), more recent findings suggest that the log-squared, squared or absolute returns are highly persistent processes. In particular, it appears that the long-run dependencies in the autocorrelation functions for these different measures of return volatility are best characterized by a slowly mean reverting hyperbolic rate of decay; see e.g. Andersen and Bollerslev (1997a), Dacorogna et al. (1993), Granger et al. (1997) and Lobato and Savin (1998). These slowly mean reverting dependencies are not well captured by the standard ARCH and stochastic volatility models hitherto applied in the literature, which generally imply that shocks to the second moments of the returns die out at a fast exponential rate of decay or else persist indefinitely in a forecasting sense. Building on the fractionally integrated ARMA, or ARFIMA, class of models (cf. Adenstedt, 1974; Granger and Joyeux, 1980; Hosking, 1981; McLeod and Hipel, 1978), a number of new parametric formulations have been proposed to better accommodate these systematic features in the variances of returns. These models include the long-memory ARCH models in Ding and Granger (1996) and Robinson (1991), the long-memory nonlinear moving average models in Robinson and Zaffaroni (1996) and Zaffaroni (1997), the fractionally integrated GARCH, or FIGARCH, model in Baillie et al. (1996) and the long-memory stochastic volatility models in Breidt et al. (1998) and Harvey (1998).¹

Meanwhile, the implementation of discrete-time econometric models invariably entails an implicit assumption regarding the length of the unit time interval. While this is typically dictated by the sampling frequency for the time series under study, many financial price series are now available on a virtually continuous, or tick-by-tick, basis. This raises important new questions related to the effects of temporal aggregation and the choice of sampling frequency when modeling the dynamics of speculative prices. In this regard, a parametric class of time series model is said to be closed under temporal aggregation if a parametric model from the same class, but with different parameter values, characterizes the data generating process across all observation frequencies. The standard ARMA class of models possesses this useful property; see e.g. Nijman and Palm (1990) and the references therein. The ARCH and stochastic volatility models typically

¹ The practical relevance of explicitly allowing for such long-range dependence when evaluating options contracts has also been explored and found to result in superior long-term prices in Bollerslev and Mikkelsen (1999).

employed in modeling volatility clustering are not, however, closed under temporal aggregation. In particular, models in the GARCH class are closed under temporal aggregation only under additional restrictive assumptions, as developed by Drost and Nijman (1993) and Drost and Werker (1996).

The present paper explores the estimation of long-memory volatility dependencies with data of different sampling frequencies through a detailed Monte Carlo study involving a fractionally integrated stochastic volatility model. We consider the estimation of the long-memory parameter using different volatility measures (the true latent fractionally integrated volatility process, log-squared returns, squared returns and absolute returns) and with data of different sampling frequencies. In all cases, the long-memory parameter is estimated by applying the semiparametric log-periodogram regression estimator of Geweke and Porter-Hudak (1983) and Robinson (1995a) to one of these volatility measures. We find that in the standard ARFIMA model, temporal aggregation makes relatively little difference to the accuracy of the estimates for d . The span of the data is of utmost importance in this situation. However, for empirically realistic parameter values, temporal aggregation induces large downward biases in the estimates for d in the squared, log-squared or absolute returns. Intuitively, temporal aggregation makes these volatility measures more noisy. We also propose a set of new volatility estimators that effectively extract the information in the high-frequency data about the longer run dependencies, while relying on a more manageable low-frequency time series for the actual estimation.² As such, our results demonstrate that when estimating volatility dynamics, it is not just the span of data that matters. The availability of high-frequency data allows for vastly superior and nearly unbiased estimation of the fractional differencing parameter that characterizes the long-run volatility dynamics.

The plan for the remainder of the paper is as follows. The model of fractionally integrated stochastic volatility is introduced in the next section and methods of inference are briefly reviewed. Section 3 presents estimation results for the degree of fractional integration in the volatility of the Japanese Yen–U.S. Dollar, ¥–\$, spot exchange rate based on a ten-year sample of more than *half-a-million* 5-min returns. Whereas the high-frequency estimates unambiguously point towards the existence of long-memory volatility dependencies, the results for the corresponding temporally aggregated daily returns are less clear cut. The main simulation results for the stochastic volatility model calibrated to the ten-year 5-min return sample are contained in Section 4. These findings clearly highlight the advantages afforded by high-frequency observations when estimating the degree of fractional integration in the volatility. Section 5 concludes.

² Along these same lines, the results in Andersen and Bollerslev (1998) illustrate that more frequently sampled observations significantly reduce the measurement error uncertainty when evaluating ex-post volatility forecasts.

2. Long-memory stochastic volatility

A long-memory time series may be defined as one which has a covariance function $\gamma(j)$ and a spectrum $f(\lambda)$ such that $\gamma(j)$ is of the same order as j^{2d-1} , as $j \rightarrow \infty$, and $f(\lambda)$ is of the same order as λ^{-2d} , as $\lambda \rightarrow 0$, for $0 < d < \frac{1}{2}$. Long-memory models were introduced to econometrics by authors such as Granger and Joyeux (1980) and Hosking (1981), while recent and comprehensive reviews of this topic are provided by Baillie (1996) and Robinson (1994a). The ARFIMA model is the leading parametric long-memory time series model. We say that $\{h_t\}_{t=1}^T$ is an ARFIMA (p, d, q) time series if

$$a(L)(1-L)^d h_t = b(L)\sigma_u u_t, \quad (1)$$

where u_t is i.i.d. with mean zero and variance one, L is the lag operator, the lag polynomials $a(L)$ and $b(L)$ (of orders p and q , respectively) are assumed to have all of their roots outside the unit circle, and the fractional differencing operator $(1-L)^d$ is defined by the binomial, or Maclaurin, series expansion. For d between -0.5 and 0.5 , $\{h_t\}$ is a stationary and invertible ARFIMA process. For $d > 0$, the ARFIMA time series model is a long-memory time series. The case where $d < 0$ is defined as antipersistence or negative long memory.

As motivated in the introduction, our focus in this paper is on models of asset returns with long memory in volatility. Let $\{y_t\}_{t=1}^T$ be a time series of asset returns such that

$$y_t = g(z_t)\sigma_\varepsilon \varepsilon_t,$$

where z_t is a long-memory time series and ε_t is an i.i.d. time series with mean zero and variance one. Since $\text{Var}(y_t|z_t) = g(z_t)^2\sigma_\varepsilon^2$, for certain functions $g(\cdot)$ this can be described as a long-memory stochastic volatility model (see Robinson, 1999). In this paper, we focus on a particular parametric model in this class, namely the fractionally integrated stochastic volatility model, introduced by Breidt et al. (1998). Other models in this class include the long-memory nonlinear moving average models in Robinson and Zaffaroni (1996) and Zaffaroni (1997).

2.1. Fractionally integrated stochastic volatility

The fractionally integrated stochastic volatility (FISV) model specifies that

$$y_t = \exp(h_t/2)\sigma_\varepsilon \varepsilon_t, \quad (2)$$

where h_t is an ARFIMA time series, as defined in Eq. (1), and it is further assumed that ε_t and u_t are i.i.d. standard normal and mutually independent. This is a long-memory stochastic volatility model with $z_t = h_t$, and $\text{Var}(y_t|h_t) = \exp(h_t)\sigma_\varepsilon^2$.

Although the time series h_t is assumed to be unobservable, its persistence properties are propagated to the observable series, $\log(y_t^2)$. In particular, on

squaring both sides of Eq. (2) and taking logarithms,

$$\log(y_t^2) = \mu_\xi + h_t + \xi_t,$$

where $\mu_\xi = \log(\sigma_\varepsilon^2) - E \log(\varepsilon_t^2)$ and $\xi_t = \log(\varepsilon_t^2) - E \log(\varepsilon_t^2)$. Hence, $\log(y_t^2)$ is the sum of a Gaussian ARFIMA process and independent non-Gaussian noise. Consequently,

$$\text{Cov}(\log(y_t^2), \log(y_{t-j}^2)) \sim j^{2d-1}, \tag{3}$$

as $j \rightarrow \infty$, while the spectrum of $\log(y_t^2)$ is of the same order as λ^{-2d} , as $\lambda \rightarrow 0$.³ This observation forms the basis for the application of the traditional log-periodogram estimation procedure.

2.2. Log-periodogram regression

Fully efficient estimation of an ARFIMA model necessarily requires that the orders of the autoregressive and moving average polynomials be specified. However, the long-memory dependency is characterized solely by the fractional differencing parameter, or d . Estimates of this parameter may be obtained by relatively simply semiparametric methods that remain agnostic about the short-run dynamics in $a(L)$ and $b(L)$; indeed semiparametric methods can be used to estimate the parameter d of a general long-memory time series. In practice, the log-periodogram regression estimate is by far the most widely used such semiparametric procedure.⁴ This estimator was first proposed by Geweke and Porter-Hudak (1983), but a proof of its consistency and limiting distribution were elusive, until Robinson (1995a) provided these, in the Gaussian case.

The basic idea is simple. If the time series under study exhibits fractional integration, the spectrum should be log-linear for frequencies close to zero with a slope equal to $-2d$. Specifically, let $I(\lambda)$ denote the sample periodogram evaluated at frequency λ . The log-periodogram estimate of d , or \hat{d} , is then given by minus the estimate of β_1 in the linear regression equation

$$\log(I(\lambda_j)) = \beta_0 + \beta_1 \log[4 \sin^2(\lambda_j/2)] + \zeta_j, \quad j = l + 1, \dots, m, \tag{4}$$

where $0 < l < m \ll T$, and $\lambda_j = 2\pi j/T$ denotes the j th Fourier frequency based on a sample of T observations. The heuristic motivation for the estimator is clear since $4 \sin^2(\lambda/2) \simeq \lambda^2$ for λ close to zero, and $f(\lambda) \sim \lambda^{-2d}$ as $\lambda \rightarrow 0$. The estimator originally proposed by Geweke and Porter-Hudak (1983) sets the *trimming parameter* $l = 0$, and so uses all of the m lowest frequencies. Meanwhile, the

³ Throughout this paper, the \sim notation means that the limit of the ratio of the quantities on the left- and right-hand sides of the symbol is a finite positive constant.

⁴ There are a number of alternative semiparametric estimators, including the ‘Gaussian’ estimator proposed by Robinson (1995b).

proof provided by Robinson (1995a) leaves out some of the very lowest frequencies.⁵ In particular, both l and the *bandwidth parameter*, m , must converge to infinity while $l/m \rightarrow 0$ and $m/T \rightarrow 0$, in which case, under additional mild conditions,

$$\sqrt{m}(\hat{d} - d) \rightarrow N(0, \frac{\pi^2}{24}).$$

The log-periodogram estimate is therefore consistent and asymptotically normal, but converges at a rate slower than the usual \sqrt{T} . Alternatively, the variance of \hat{d} may also be consistently estimated by the conventional OLS variance estimate from the regression in (4). In the present paper, our interest is not in estimating an ARFIMA process, but rather in estimating the FISV model. So, in the FISV model, we seek to estimate d semiparametrically by applying the log-periodogram regression to $\log(y_t^2)$. Of course, in the FISV model, $\log(y_t^2)$ is given by the sum of an ARFIMA process and non-Gaussian noise, and so the asymptotic theory provided by Robinson (1995a) does not formally apply to this case.

In our empirical analysis of the high-frequency ¥–\$ exchange rate below, we supplement the asymptotic standard error from the regression in (4) with a frequency domain bootstrap confidence interval, based on the ideas in Franke and Härdle (1992) (see also Berkowitz and Diebold, 1998; Berkowitz and Kilian, 1999). Specifically, let $\hat{\beta}_0$ and $\hat{\beta}_1$ denote the OLS estimates of β_0 and β_1 in Eq. (4), respectively, and let $\hat{f}(\lambda_j) = \exp(\hat{\beta}_0)[4 \sin^2(\lambda_j/2)]^{\hat{\beta}_1}$ denote the corresponding estimated spectrum at low frequencies (Fourier frequencies λ_{l+1} through λ_m). We then take bootstrap draws of the periodogram at these Fourier frequencies, by multiplying the estimated spectrum by i.i.d. exponential random variables, and hence build up a bootstrap distribution of the log-periodogram regression estimate. A confidence interval can be read off from the percentiles of this bootstrap distribution.

2.3. Alternative volatility measures

Although semiparametric estimation of the fractional integration parameter for the FISV model is most naturally based on the log-squared time series, $\log(y_t^2)$, many researchers have relied on the time series of squared returns, y_t^2 , or absolute returns, $|y_t|$, as measures of the ex-post volatility. Interestingly, these alternative measures, in particular the absolute returns, tend to result in equally strong empirical evidence for fractional integration. This is to be expected.

To understand why, note that from the definition of the FISV model in Eq. (2), $y_t^2 = \exp(h_t + \log(\sigma_\varepsilon^2))\varepsilon_t^2$ and $|y_t| = \exp(h_t + \log(\sigma_\varepsilon^2))^{1/2}|\varepsilon_t|$. Furthermore, by the

⁵ Hurvich et al. (1998) provide a version of this proof which sets $l = 0$; it entails some additional regularity conditions.

properties of power transforms of normals (see e.g: Andersen, 1994), it follows that

$$\begin{aligned} & \text{Cov}[\exp(h_t + \log(\sigma_\varepsilon^2)), \exp(h_{t-j} + \log(\sigma_\varepsilon^2))] \\ &= \{E[\exp(h_t + \log(\sigma_\varepsilon^2))]\}^2 \{\exp[\text{Cov}(h_t, h_{t-j})] - 1\} \\ &= \{E[\exp(h_t + \log(\sigma_\varepsilon^2))]\}^2 [j^{2d-1} + o(j^{2d-1})] \sim j^{2d-1}, \end{aligned}$$

as $j \rightarrow \infty$, and, similarly,

$$\text{Cov}[\exp(h_t + \log(\sigma_\varepsilon^2))^{1/2}, \exp(h_{t-j} + \log(\sigma_\varepsilon^2))^{1/2}] \sim j^{2d-1},$$

as $j \rightarrow \infty$. The autocovariances of the squared and absolute returns therefore decay at the same rate as the autocovariances of h_t , i.e.

$$\text{Cov}(y_t^2, y_{t-j}^2) \sim j^{2d-1} \tag{5}$$

and

$$\text{Cov}(|y_t|, |y_{t-j}|) \sim j^{2d-1}, \tag{6}$$

as $j \rightarrow \infty$. These results were shown by Andersen and Bollerslev (1997a) and also follow from more general results in Robinson (1999). The log-periodogram estimates calculated from the log-squared, squared, or absolute returns (i.e. based on Eqs. (4)–(6)) can therefore be expected to be consistent⁶ for the same value of d .

2.4. Temporal aggregation

Many financial time series are now available on a virtually continuous, or tick-by-tick, basis. As such, the choice of a unit discrete time interval and the corresponding estimate for d may appear rather arbitrary, especially since the ARFIMA model is not closed under temporal aggregation. Nevertheless, Chambers (1998) shows that the spectrum for the temporally aggregated process, $h_t^{(k)} \equiv h_{tk+1-k} + h_{tk+2-k} \cdots + h_{tk}$, for $t = 1, 2, \dots, T/k$, remains log-linear for frequencies close to zero, diverging at the same rate of λ^{-2d} as $\lambda \rightarrow 0$, for all k .

In the same way, while the FISV model is not closed under temporal aggregation, the rate of decay of the autocovariance function of squared returns is invariant to the length of the return interval. Specifically, let $y_t^{(k)} = y_{tk+1-k} + y_{tk+2-k} \cdots + y_{tk}$ denote the continuously compounded k -period return, so that $y_t^{(k)2} = y_{tk+1-k}^2 + y_{tk+2-k}^2 \cdots + y_{tk}^2 + \text{cross-product terms}$. Because the FISV model implies that the one-period return is serially uncorrelated, the

⁶ This is, of course, not a formal consistency proof. The fact that the covariances decay at the rate j^{2d-1} does not necessarily imply that the spectrum diverges at the rate λ^{-2d} (this will however be ensured if the autocovariances are quasi-monotonically convergent to zero). In any case, we do not have the Gaussianity required in the consistency proof of Robinson (1995a).

cross-product terms all have zero mean. Hence, it follows that

$$\text{Cov}(y_t^{(k)2}, y_{t-j}^{(k)2}) = \sum_{h=-k-1}^{k-1} (k - |h|)\text{Cov}(y_t^2, y_{t-jk-h}^2) \sim (jk)^{2d-1} \sim j^{2d-1}, \quad (7)$$

as $j \rightarrow \infty$, see also Andersen and Bollerslev (1997a). Although the analysis in Robinson (1999) of instantaneous nonlinear transforms of long-memory measures does not formally apply to the squared temporally aggregated data, we conjecture that, for a wide class of distributions, the autocovariances of log-squared or absolute temporally aggregated returns should eventually decay at the same long-run hyperbolic rate, i.e.

$$\text{Cov}[\log(y_t^{(k)2}), \log(y_{t-j}^{(k)2})] \sim j^{2d-1} \quad (8)$$

and

$$\text{Cov}[|y_t^{(k)}|, |y_{t-j}^{(k)}|] \sim j^{2d-1}, \quad (9)$$

as $j \rightarrow \infty$.

In practice, with prices available over a fixed time period, the number of return observations is inversely related to the length of the return interval, so that with a sample of T one-period high-frequency returns, temporal aggregation results in T/k k -period returns. In the simulations reported below, we find that with a fixed sampling interval T , the properties of the log-periodogram estimates for d based on the relationships in Eqs. (7)–(9) deteriorate significantly as the return horizon, or k , lengthens. This is not true for the log-periodogram estimate for d based on the latent $h_t^{(k)}$ process. This motivates our final set of volatility measures and estimators, explicitly designed to overcome the computational burden involved in analyzing large high-frequency data sets, while maintaining the additional information about the longer run volatility dependencies inherent in the high-frequency returns.

In particular, suppose that the autocovariances of some high-frequency time series v_t decay at the rate j^{2d-1} . Let the temporally aggregated series be denoted by $[v_t]^{(k)} = v_{tk+1-k} + v_{tk+2-k} \cdots + v_{tk}$ for $t = 1, 2, \dots, T/k$. Collecting terms of the same lag length, it follows that the autocovariance between $[v_t]^{(k)}$ and $[v_{t-j}]^{(k)}$ is

$$\text{Cov}([v_t]^{(k)}, [v_{t-j}]^{(k)}) = \sum_{h=-k-1}^{k-1} (k - |h|)\text{Cov}(v_t, v_{t-jk-h}) \sim (jk)^{2d-1} \sim j^{2d-1}, \quad (10)$$

as $j \rightarrow \infty$. Thus, instead of estimating the degree of fractional integration in the volatility directly from long time-series of high-frequency returns, d may be estimated from shorter samples of temporally aggregated log-squared, squared or absolute returns, where the aggregation takes place after the nonlinear

transformation;⁷ that is $[\log(y_t^2)]^{(k)}$, $[y_t^2]^{(k)}$ and $[|y_t|]^{(k)}$. Note that, in contrast to the log-squared, squared or absolute temporally aggregated low-frequency returns in Eqs. (7)–(9), the volatility measures based on Eq. (10) explicitly rely on high-frequency data. Of course, the quality of the approximation in Eq. (10), and the corresponding estimates, will depend on the functional form being used.

In summary, there are many potential volatility measures (log-squared, squared and absolute returns, with any choice of the temporal aggregation parameter k , and with aggregation either before or after the nonlinear transformation). The main focus of this paper is on the comparison of the estimators of d , using these alternative volatility measures.

3. Foreign exchange rate volatility

A number of studies have argued for the existence of long memory in the volatility of equity and foreign exchange returns. While the majority of these studies have relied on relatively long time spans of daily or lower frequency returns, Andersen and Bollerslev (1997a) have recently uncovered strong evidence for long-memory dependencies with a one-year time series of high-frequency Deutschemark–U.S. Dollar spot exchange rate returns. The estimation results for the ten-year ¥–\$ intraday return series detailed below complement these findings, and set the stage for the subsequent Monte-Carlo study.

3.1. Data and preliminary summary statistics

The spot exchange rate data were collected and provided by Olsen and Associates in Zürich, Switzerland. The full sample spans the period from December 1, 1986 through December 1, 1996. The returns are calculated as the logarithmic difference between the linearly interpolated average of the mid-point of the bid and the ask for the two nearest quotes, resulting in a total of 288 5-min return observations per day.⁸ Although the foreign exchange market is officially open 24 h a day and 365 days a year, the trading activity slows decidedly during the weekend period. In order to avoid confounding the evidence by such weekend patterns, we simply excluded all returns from Friday 21:00 Greenwich Mean Time (GMT) through Sunday 21:00 GMT; a similar

⁷ Note that with daily data, $[y_t^2]^{(2)}$ corresponds directly to the ex-post monthly variance measure commonly employed in the empirical finance literature; see e.g. Schwert (1989). Similar ex-post daily variance measures based on high-frequency intradaily returns have also been analyzed by Andersen and Bollerslev (1998) and Hsieh (1991), among others.

⁸ For a more detailed description of the activity patterns in the ¥–\$ foreign exchange market and the method of data capturing and filtering that underlie the return calculations, we refer to Dacorogna et al. (1993) and Müller et al. (1990).

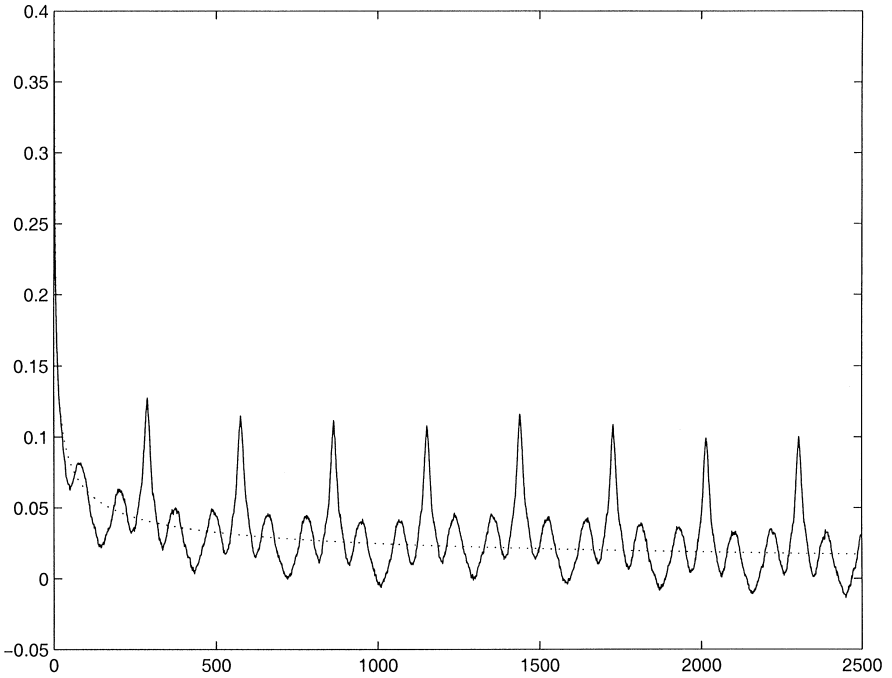


Fig. 1. Sample and benchmark model population autocorrelograms. The solid line graphs the sample autocorrelogram for the 5-min log-squared ¥–\$ series. The sample period extends from December 1, 1986 through December 1, 1996, for a total of $T = 751,392$ observations. The dotted line refers to the theoretical autocorrelogram for the FISV model defined by Eqs. (2) and (11) with parameters $d = 0.3$, $\sigma_\varepsilon^2 = 4.10^{-4}$, $\phi = 0.6$ and $\sigma_u^2 = 0.25$.

weekend no-trade convention was adopted by Andersen and Bollerslev (1997a,b). This leaves us with a sample 2609 days, for a total of $T = 2609 \times 288 = 751,392$ 5-min y_t return observations. The corresponding sample sizes for the hourly, $y_t^{(12)}$ and daily $y_t^{(288)}$ returns are 62,616 and 2609, respectively.

Consistent with the notion of efficient markets, the 5-min returns are approximately mean zero and serially uncorrelated.⁹ At the same time, the evidence for volatility clustering is overwhelming. For instance, the lag-1 sample autocorrelation coefficient for the log-squared 5-min returns is 0.373, which is overwhelmingly significant at any level. Meanwhile, the longer run dependence in the autocorrelogram, depicted in Fig. 1, is masked by a distinct repetitive pattern.

⁹ The first-order sample autocorrelation coefficient equals -0.062 , which is highly statistically significant when judged by the conventional asymptotic standard error of $\sqrt{1/T}$. However, the coefficient is numerically small and is readily explained by the strategic positioning of discrete bid and ask quotes; see e.g. the discussion in Bollerslev and Domowitz (1993).

Equally pronounced periodic autocorrelograms for other log-squared speculative returns and time periods have previously been reported in the literature by Andersen and Bollerslev (1997b) and Dacorogna et al. (1993), among others. However, this periodicity is directly attributable to the existence of strong intradaily volatility patterns associated with the opening and closing of the various financial centers around the world. Abstracting from this daily pattern, the overall slow rate of decay in the 5-min autocorrelations is striking and highly suggestive of long-memory dependencies. As such, semiparametric frequency domain procedures that explicitly ignore the intradaily periodicities are ideally suited to estimating the rate of this apparent hyperbolic decay.

3.2. Log-periodogram estimates

Implementation of the log-periodogram regression in Eq. (4) requires a choice of the trimming and bandwidth parameters, l and m . For the results reported below we took $l = 10$. Informal experimentation revealed the estimates to be robust with respect to this particular choice. Several recent studies have been concerned with finding ‘good’ bandwidths for long-memory models; see e.g. Delgado and Robinson (1996), Geweke (1998), Hurvich et al. (1998), Lobato and Robinson (1998), and Robinson (1994b). General optimal conditions that would be applicable in the present context have so far proven elusive, but the ‘rule-of-thumb’ employed in most of the applied literature dictates a bandwidth equal to $[T^{1/2}]$, or $m = 867$ for the 5-min return series analyzed here.

In the empirical estimation, we therefore explored a range of different bandwidth choices.¹⁰ The estimates of d , using 5-min log-squared returns for a range of values of m are graphed in Fig. 2. Fig. 2 also shows the associated 95% confidence intervals based on the asymptotic standard errors from the regression in Eq. (4), i.e. $\sqrt{\pi^2/24m}$, as well as the bootstrap confidence intervals discussed in Section 2.2. The two sets of confidence intervals are virtually indistinguishable. Also, for values of m in excess of 750, the estimates for d all lie within a fairly narrow band close to 0.3. This value of d , along with the various summary statistics discussed above, forms the basis for the simulation design in the next section.

3.3. Alternative volatility measures and temporal aggregation

Before discussing the simulation results, it is informative to briefly review and compare the estimates for long-memory volatility dependencies constructed

¹⁰ A related issue concerns the choice of m for different sampling frequencies. No matter what the sampling frequency, the same value of m corresponds to cycles of the same periodicity in calendar time. This is an argument for using the same bandwidth for different levels of temporal aggregation. However, following this rule literally could arguably result in unreasonably large values of the bandwidth parameter when the sampling frequency is low.




Fig. 2. Log-periodogram estimates for 5-min log-squared ¥-\\$ returns. The solid line graphs the log-periodogram regression estimates for the 5-min log-squared series as a function of the bandwidth, m . The sample period extends from December 1, 1986 through December 1, 1996, for a total of $T = 751,392$ observations. The dotted lines give the $+/- 1.96$ standard error bands based on the asymptotic standard error from the regression in Eq. (4), i.e. $\sqrt{\pi^2/24m}$. The dashed and dotted lines give the 95% bootstrap confidence intervals, discussed in text.

from the alternative volatility measures outlined in Section 2. Fig. 3 plots the log-periodogram estimates of d against the bandwidth m , using these alternative ex-post volatility measures, i.e. log-squared, squared and absolute returns, for different levels of temporal aggregation (5-min, hourly and daily), and their counterparts aggregating after transformation. The results for $m = 1000$ (400 for the daily data) along with the associated standard errors are also reported in Table 1.¹¹ The estimates of d based on 5-min log-squared and squared returns are similar (around 0.3), with those based on 5-min absolute returns slightly higher. Meanwhile, the estimates for d based on the log-squared, squared or

¹¹ We also used the alternative Gaussian semiparametric estimator proposed by Robinson (1995b). These estimates were very close to the conventional log-periodogram estimates, and are consequently not reported in the table (e.g. for $m = 1000$, the Gaussian estimates were 0.355, 0.322 and 0.449 for log-squared, squared and absolute 5-min returns, respectively).

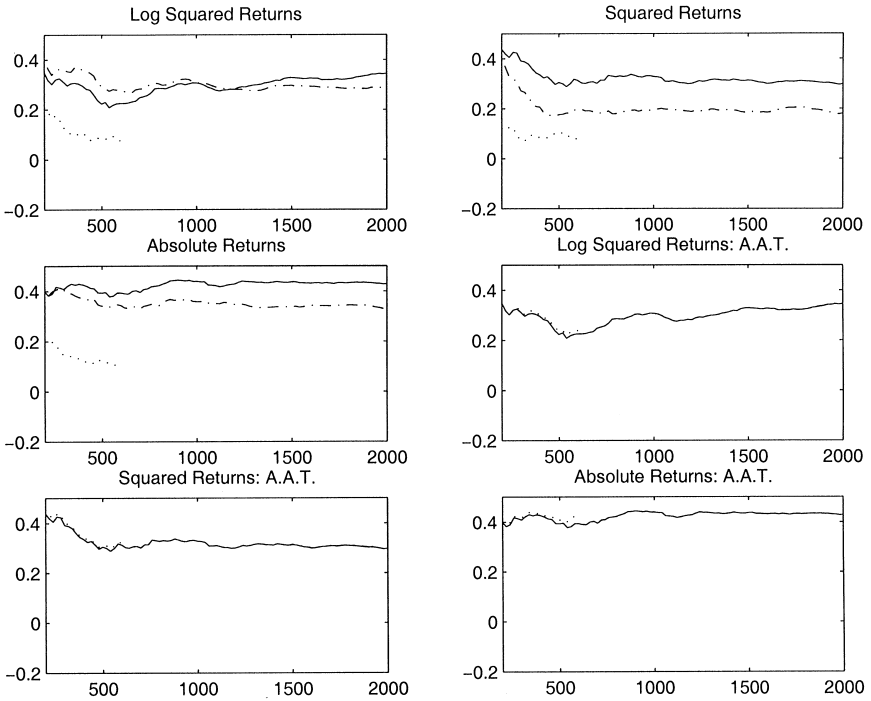


Fig. 3. Log-periodogram estimates for ¥–\$ returns at different aggregation levels. The solid line graphs the log-periodogram regression estimate for d based on the 5-min ¥–\$ series as a function of the bandwidth, m . The sample period extends from December 1, 1986 through December 1, 1996, for a total of $T = 751,392$ observations. The dashed and dotted lines refer to the estimates with hourly data, while the daily estimates are represented by the dotted lines. The different volatility estimators are defined in the text.

absolute *hourly* returns are generally lower. In turn, the estimates based on the log-squared, squared or absolute *daily* returns are much lower, and typically around 0.1.

In contrast to the low persistence implied by the three one-day return volatility measures, the point estimates for the temporally aggregated daily volatility series (aggregated after the nonlinear transformation) all correspond closely with the point estimates obtained directly from the high-frequency returns. For instance, the estimate for d calculated from the daily $[\log(y_t^2)]^{(288)}$ time series equals 0.302, whereas the estimate from $\log(y_t^2)$ is 0.307. However, whereas the estimates for $k = 1$ are based on $T = 751,392$ observations, the temporal aggregation with $k = 288$ reduces the sample size to a much more manageable 2609 daily ex-post volatility observations.

In summary, the empirical results for the ¥–\$ returns in Table 1 clearly suggest that alternative volatility measures can yield very different conclusions, and that

Table 1
Long-memory parameter estimates for ¥–\$ volatility^a

k	1	12	288
$\log(y_t^{(k)2})$	0.307 (0.020) [0.260, 0.350]	0.308 (0.020) [0.264, 0.354]	0.107 (0.032) [0.020, 0.188]
$y_t^{(k)2}$	0.328 (0.020) [0.279, 0.371]	0.195 (0.020) [0.151, 0.240]	0.083 (0.032) [0.005, 0.164]
$ y_t^{(k)} $	0.439 (0.020) [0.392, 0.485]	0.358 (0.020) [0.315, 0.404]	0.129 (0.032) [0.042, 0.212]
$[\log(y_t^2)]^{(k)}$	0.307 (0.020) [0.260, 0.350]	0.307 (0.020) [0.259, 0.351]	0.302 (0.032) [0.220, 0.383]
$[y_t^2]^{(k)}$	0.328 (0.020) [0.279, 0.371]	0.328 (0.020) [0.282, 0.377]	0.344 (0.032) [0.261, 0.383]
$[y_t]^{(k)}$	0.439 (0.020) [0.392, 0.485]	0.439 (0.020) [0.395, 0.482]	0.434 (0.032) [0.349, 0.508]

^aNotes: The table reports the log-periodogram regression estimates for the degree of fractional integration in ¥–\$ return volatility. The numbers in parentheses give the asymptotic standard errors from the regression in Eq. (4). The numbers in square brackets refer to the 95% bootstrap confidence intervals discussed in text. The sample period extends from December 1, 1986 through December 1, 1996. The sample size for the 5-min returns is 751,392. The columns labelled $k = 1$, $k = 12$ and $k = 288$ refer to the 5-min data, aggregated hourly data and aggregated daily data, respectively. The different ex-post volatility measures underlying the log-periodogram regression estimates are defined in the first column: $[\log(y_t^2)]^{(k)}$, $[y_t^2]^{(k)}$ and $[|y_t|]^{(k)}$ refer to the data aggregated after transformation, as discussed in text. The trimming parameter is fixed at $l = 10$. The estimates for $k = 1, 12$ are based on a bandwidth of $m = 1000$, whereas the estimates for $k = 288$ set $m = 400$.

high-frequency data may be critically important for reliable inference concerning longer run volatility dependencies. The simulations in the next section support this conjecture.

4. Monte-Carlo evidence

A number of simulation studies have addressed the performance of the log-periodogram regression procedure for estimating long-memory dependencies in the mean; see e.g. Agiakloglou et al. (1992), Cheung (1993), Hurvich and Beltrao (1994), Hurvich et al. (1998) and Janacek (1982). A common finding to emerge from these studies concerns the large finite sample biases in empirically realistic situations with macroeconomic data. Furthermore, it has also become

apparent that the magnitude of this small sample bias depends crucially upon the time span of the data, whereas the frequency of the observations is much less important.¹² In contrast, the present findings demonstrate that when estimating the degree of fractional integration in the volatility, the performance of the estimates may be greatly enhanced by increasing the observation frequency.

4.1. Simulation design

The simulation experiment is modelled after the empirical results for the log-squared 5-min ¥–\$ returns discussed in the previous section. However, in order to keep the design tractable, we do not incorporate the strong intraday periodicities in the volatility evident in Fig. 1. Since the log-periodogram regression is calculated solely from the longer run interdaily ordinates, this should not alter any of the qualitative findings.

Specifically, data is generated by the FISV model, where the latent volatility factor, h_t , is assumed to follow the ARFIMA (1, d , 0) model

$$(1 - \phi L)(1 - L)^d h_t = \sigma_u u_t. \quad (11)$$

The data generated in the simulation is treated as ‘5-min’ data: aggregating by factors of 12 and 288 gives artificial series of ‘hourly’ and ‘daily’ returns. The parameter configuration for the benchmark FISV model defined by Eqs. (2) and (11) specifies that $d = 0.3$, $\phi = 0.6$, $\sigma_u^2 = 0.25$, $\sigma_\varepsilon^2 = 0.0004$ and $T = 2^{19} = 524,288$. This large sample size is directly in line with the number of 5-min return observations available for the ¥–\$ series analyzed above.¹³ The mean volatility level, as measured by $E \log(y_t^2) = E \log(\varepsilon_t^2) + \log(\sigma_\varepsilon^2) = -9.09$, closely matches the sample mean for the log-squared 5-min ¥–\$ returns of -8.75 . Furthermore, the population autocorrelation for $\log(y_t^2)$, given by the dotted line in Fig. 1, is also in close accordance with the overall slow decay in the corresponding sample autocorrelations.¹⁴

Beran (1994) provides a review of the standard approaches for simulating ARFIMA models, such as the Cholesky decomposition of the covariance matrix of the time series or the algorithm of Davies and Harte (1987). In the present

¹² Similarly, when testing for unit roots in the mean, the power of such tests is intimately related to the calendar time span of the data; see e.g. Shiller and Perron (1985) for some of the earliest evidence along these lines.

¹³ In order to utilize the fast Fourier transform, the sample size must be an integer power of 2.

¹⁴ Following Breidt et al. (1998) and Hosking (1981), the lag j autocorrelation for $\log(y_t^2)$ for the FISV model takes the form

$$\begin{aligned} & (-1)^j \frac{\sigma_u^2}{1 - \phi^2} \frac{\Gamma(1 - 2d)}{\Gamma(1 - j - d)\Gamma(j + 1 - d)} [F(1, d + j, 1 - d + j; \phi) + F(1, d - j, 1 - d - j; \phi)] \\ & + \frac{\pi^2}{2} 1(j = 0), \end{aligned}$$

where $1(\cdot)$ denotes the indicator function and $F(\cdot, \cdot, \cdot; \cdot)$ is the hypergeometric function.

context, with a sample size in excess of half-a-million, none of these methods would be computationally feasible. While it might be possible to simulate an approximate long-memory time series by applying a truncated version of the $(1 - L)^d$ filter, truncation would necessarily have to take place at a very high power of L and this algorithm would be very slow. Instead, we rely on the recent developments in Parke (1996) for generating a fractional white noise process as the weighted sum of i.i.d. shocks; a detailed description of this procedure is provided in the appendix. Passing this process through an AR(1) filter yields the desired ARFIMA (1, d , 0) model for the latent volatility, or h_t . All of the resulting simulation results are based on a total of 1000 replications.

4.2. Results

Fig. 4 plots the simulated bias of the log-periodogram estimates of d against the bandwidth m , for all the volatility measures and different levels of temporal aggregation analyzed in the previous section. For comparison, Fig. 4 also shows the bias of the log-periodogram estimate applied to the latent fractionally integrated volatility process, h_t . This latter estimator is, of course, not a feasible alternative in practice. Consistent with the empirical results in the previous section, the biases in the estimates are not very sensitive to the choice of m , once it reaches a ‘reasonable’ value. In the following, we shall therefore concentrate on the results for $m = 1000$ ($m = 400$ for the ‘daily’ estimates).¹⁵

The kernel estimates of the associated probability densities are shown in Fig. 5. Table 2 also gives the median and 2.5 and 97.5 percentiles of the simulated distribution of the log-periodogram estimates of d , for these values of m . The results in Table 2 and Figs. 4 and 5 show a slight downward bias in the estimate of d using ‘5-min’ log-squared returns. This bias gets worse using $\log(y_t^{(12)^2})$ and is much worse with $\log(y_t^{(288)^2})$. This same pattern emerges for the estimates based on the squared and absolute returns. Overall, the log-periodogram estimates based on the absolute and log-squared returns appear to have quite similar properties, whereas the estimates based on the squared returns have more downward bias and larger variance. The first three panels in Fig. 5 also vividly illustrate these biases from standard linear temporal aggregation in the FISV model.

This severe aggregation bias does not appear in the simple ARFIMA model, where aggregating to a lower frequency, while keeping the span of the data constant, makes little difference in terms of bias and only slightly increases the dispersion. For instance, the log-periodogram estimates, in Table 2, based on $h_t^{(k)}$ have median values 0.297, 0.299 and 0.356, for $k = 1, 12$ and 288, respectively.

¹⁵ The truncation parameter is fixed at $l = 10$ in all cases.

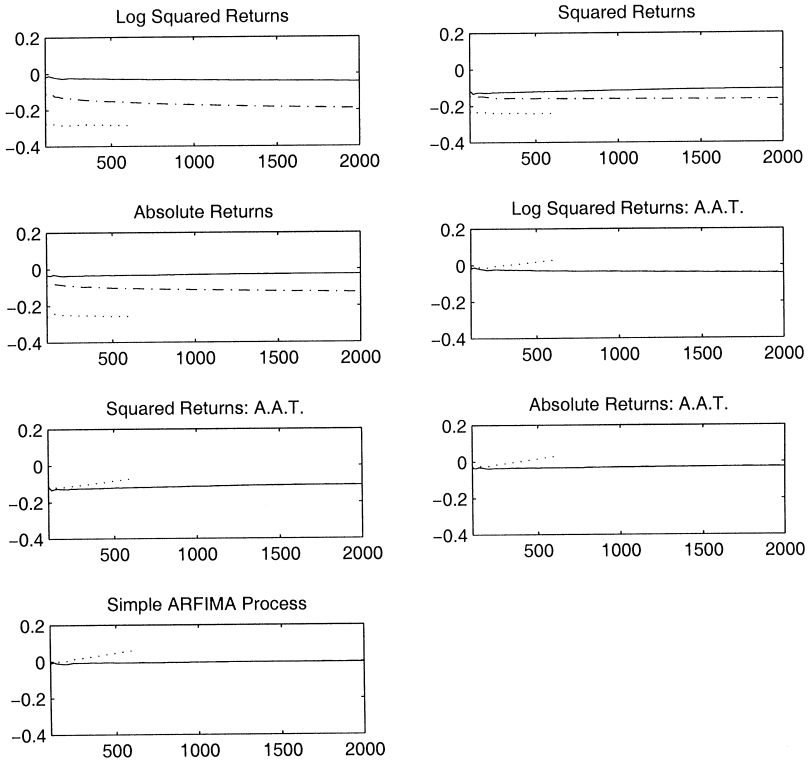


Fig. 4. Simulated bias of estimators. The solid lines graph the bias in the log-periodogram estimate for d based on the simulated ‘5-min’ return series as a function of the bandwidth, m . The dashed and dotted lines give the bias for the ‘hourly’ data ($k = 12$), while the biases for the simulated ‘daily’ ($k = 288$) estimates are represented by the dotted lines. The different volatility estimators are defined in the text. The data is generated by the FISV model in Eqs. (2) and (11) with parameters $T = 2^{19}$, $d = 0.3$, $\sigma_{\varepsilon}^2 = 4.10^{-4}$, $\phi = 0.6$ and $\sigma_u^2 = 0.25$.

The last three volatility measures in Table 2 and Figs. 4 and 5 are constructed by summing the high-frequency log-squared, squared or absolute returns; i.e. $[\log(y_t^2)]^{(k)}$, $[y_t^2]^{(k)}$ and $|y_t|^{(k)}$, respectively. Applying the nonlinear transforms before aggregation effectively circumvents the aggregation bias. The ‘hourly’ estimates ($k = 12$) have a distribution that is very close to that of their high-frequency counterparts ($k = 1$). The ‘daily’ estimates ($k = 288$) have approximately the same mean and median as their high-frequency counterparts, though with slightly higher dispersion, and the ‘daily’ estimates based on the log-squared and absolute returns are nearly unbiased (and median unbiased). These results are again consistent with the empirical estimates reported in the previous

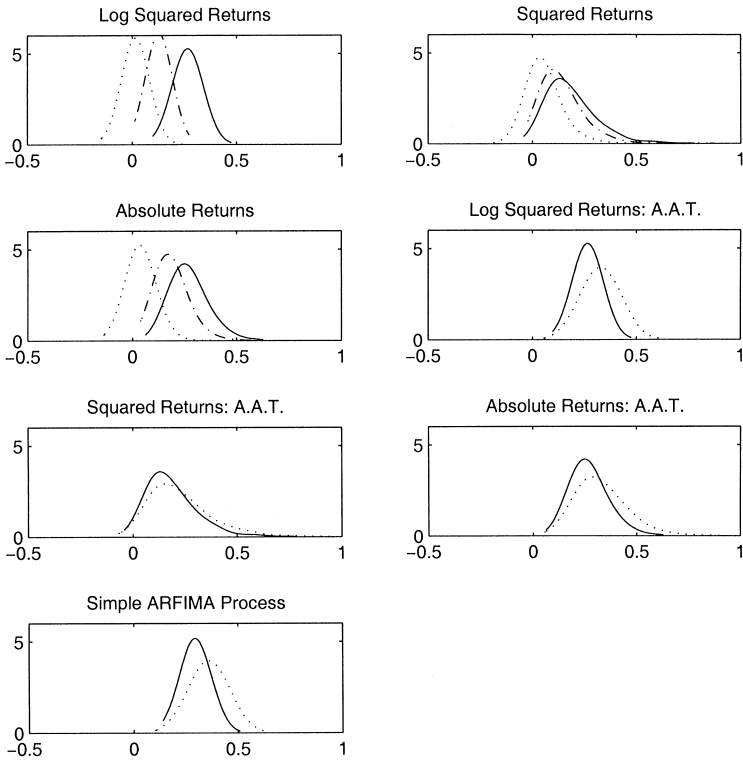


Fig. 5. Simulated densities of log-periodogram estimators. The figure graphs the Gaussian kernel estimates of the simulated probability density functions for the log-periodogram regression estimates of d . The solid lines give the densities for the ‘5-min’ returns and $m = 1000$. The dashed and dotted lines give the densities for the ‘hourly’ data ($k = 12$) and $m = 1000$, while the densities for the ‘daily’ ($k = 288$) estimates and $m = 400$ are represented by the dotted lines. The different volatility estimators are defined in the text. The data is generated by the FISV model in Eqs. (2) and (11) with parameters $T = 2^{19}$, $d = 0.3$, $\sigma_\varepsilon^2 = 4.10^{-4}$, $\phi = 0.6$ and $\sigma_u^2 = 0.25$.

section. However, it is important to note that, even though these estimates are based on much fewer time series observations, the corresponding low-frequency volatility measures cannot be constructed without the underlying high-frequency data. Accordingly, these results underscore the benefits provided by intradaily returns when estimating interdaily volatility dynamics.

The Monte-Carlo simulations clearly illustrate that temporal aggregation greatly increases the bias in the log-periodogram estimate of d for the log-squared, squared or absolute returns, but not for $h_t^{(k)}$, $[\log(y_t^2)]^{(k)}$, $[y_t^2]^{(k)}$ or $|y_t|^{(k)}$. Our explanation for this finding is easiest to convey in the context of the squared returns. By the standard intuition about the effects of temporal aggregation, h_t and $h_t^{(k)}$ should both lead to similar conclusions about the long-memory

Table 2
 Monte-Carlo properties of long-memory parameter estimates for FISV model^a

k	1	12	288
$\log(y_t^{(k)2})$	0.256 (0.185, 0.330)	0.143 (0.021, 0.286)	0.017 (-0.105, 0.139)
$y_t^{(k)2}$	0.179 (0.030, 0.443)	0.120 (0.009, 0.359)	0.045 (-0.068, 0.278)
$ y_t^{(k)} $	0.267 (0.170, 0.418)	0.177 (0.075, 0.345)	0.039 (-0.077, 0.176)
$h_t^{(k)}$	0.297 (0.221, 0.370)	0.299 (0.222, 0.372)	0.356 (0.177, 0.531)
$[\log(y_t^2)]^{(k)}$	0.256 (0.185, 0.330)	0.257 (0.187, 0.330)	0.327 (0.158, 0.489)
$[y_t^2]^{(k)}$	0.179 (0.030, 0.443)	0.179 (0.030, 0.443)	0.198 (0.000, 0.612)
$[y_t]^{(k)}$	0.267 (0.170, 0.418)	0.269 (0.171, 0.420)	0.313 (0.131, 0.615)

^aNotes: The table reports the median value of the simulated distribution of the log-periodogram regression estimates for d with the 2.5 and 97.5 percentiles of this simulated distribution in parentheses. The different ex-post volatility measures underlying the log-periodogram regression estimates are defined in the first column: $[\log(y_t^2)]^{(k)}$, $[y_t^2]^{(k)}$ and $[|y_t|]^{(k)}$ refer to the data aggregated after transformation, as discussed in text. The FISV model parameters are $T = 2^{19}$, $d = 0.3$, $\sigma_\varepsilon^2 = 4.10^{-4}$, $\phi = 0.6$ and $\sigma_u^2 = 0.25$. The trimming parameter is fixed at $l = 10$. The estimates for $k = 1, 12$ are based on a bandwidth of $m = 1000$, whereas the estimates for $k = 288$ set $m = 400$.

parameter. The same is true for y_t^2 and $[y_t^2]^{(k)}$. However, $[y_t^{(k)}]^2$ is equal to $[y_t^2]^{(k)}$ plus $k(k-1)$ zero-mean cross-product terms. Thus, $[y_t^{(k)}]^2$ provides a more noisy volatility measure than $[y_t^2]^{(k)}$ for $k > 1$. Using a more noisy volatility measure is likely to bias the estimate of d downwards; the higher is k , the more noisy is $[y_t^{(k)}]^2$, and so the larger the bias should be.

Of course, in practical applications any inference concerning the true value of d will have to rely on appropriately defined test statistics. Fig. 6 shows the densities of the t -statistic for testing the hypothesis that $d = 0.3$ (the true value), using the asymptotic standard errors from the regression in (4), with ‘5-min’ log-squared, squared and absolute returns as well as the latent h_t ARFIMA process. Four different bandwidth parameters are reported; $m = 500, 1000, 1500$ and 2000 , while $l = 10$. Table 3 gives the simulated probabilities that each of these t -statistics is less than 1.96 in absolute value (i.e. the coverage of a nominal 95% confidence interval).

The sampling distributions of the t -statistics always have higher variance than a standard normal. This is even true for the simple ARFIMA process. This tendency gets more pronounced as m increases. The excessively small standard

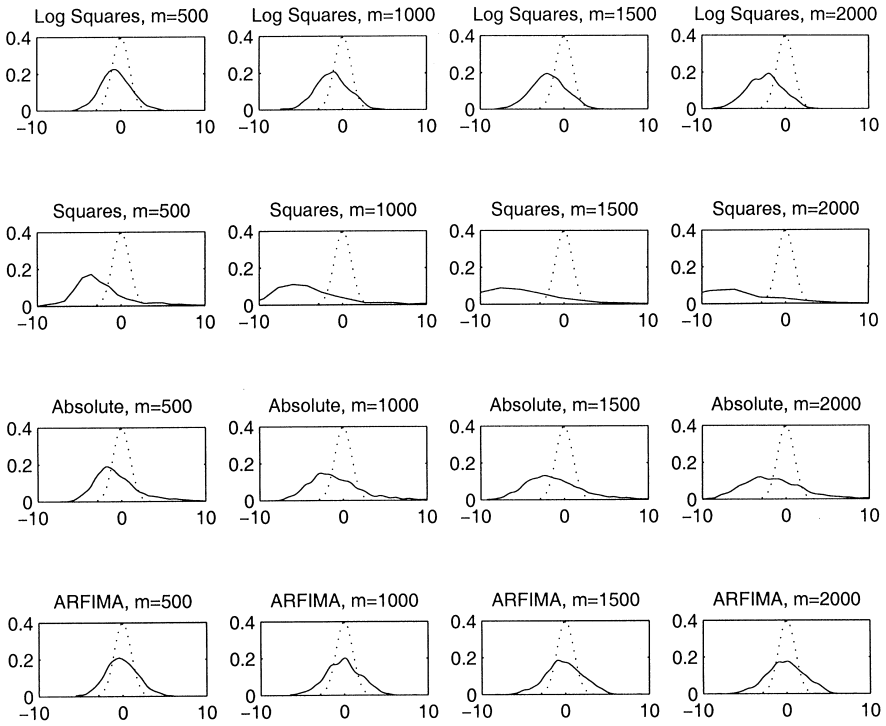


Fig. 6. Simulated densities of t -statistics. The solid lines graph the Gaussian kernel estimates of the simulated probability density functions for the t -statistics testing $d = 0.3$ based on the ‘5-min’ return estimates for d and the asymptotic standard error from the regression in Eq. (4). The dotted lines give the standard normal density, for comparison.

Table 3

Coverage of confidence intervals in log-periodogram regressions with ‘5-min’ data^a

m	500	1000	1500	2000
$\log(y_t^2)$	0.697	0.574	0.480	0.405
y_t^2	0.240	0.157	0.127	0.116
$ y_t $	0.531	0.418	0.377	0.363
h_t	0.710	0.635	0.620	0.610

^aNotes: The table reports the effective coverage of the 95% confidence intervals for d based on the ‘5-min’ return estimates for d and the asymptotic standard error from the regression in Eq. (4). The Monte-Carlo design is as specified in Table 2.

errors, combined with modest downward biases in the log-periodogram estimates of d cause the confidence intervals to have low coverage, especially for large m . Thus, while high-frequency data allow the construction of approximately unbiased point estimates for the long-memory volatility

Table 4
Sensitivity analysis: Monte-Carlo properties for model M1^a

k	1	12	288
$\log(y_t^{(k)2})$	0.230 (0.139, 0.323)	0.100 (-0.031, 0.254)	0.009 (-0.275, 0.298)
$y_t^{(k)2}$	0.186 (0.034, 0.457)	0.103 (0.009, 0.304)	0.028 (-0.241, 0.384)
$ y_t^{(k)} $	0.253 (0.152, 0.436)	0.134 (0.027, 0.319)	0.031 (-0.256, 0.304)
$h_t^{(k)}$	0.301 (0.205, 0.398)	0.319 (0.219, 0.420)	0.370 (-0.024, 0.836)
$[\log(y_t^2)]^{(k)}$	0.223 (0.139, 0.323)	0.239 (0.153, 0.330)	0.330 (-0.038, 0.762)
$[y_t^2]^{(k)}$	0.186 (0.034, 0.457)	0.200 (0.045, 0.473)	0.161 (-0.120, 0.861)
$[y_t]^{(k)}$	0.253 (0.152, 0.436)	0.266 (0.165, 0.449)	0.295 (-0.085, 0.881)

^aNotes: This is the counterpart to Table 2, but where the model parameters are altered from their benchmark values to specify that the sample size is reduced to $T = 2^{16}$.

parameter, conventional inference procedures should be interpreted with some caution.¹⁶

4.3. Sensitivity analysis

In this section we report some sensitivity analysis, analyzing the consequences of modifications to the benchmark model. Specifically we consider the following three parameter configurations, each of which alters one parameter relative to the benchmark model:

- (i) Model M1: $T = 2^{16}$, $d = 0.3$, $\phi = 0.6$, $\sigma_\varepsilon^2 = 4 \times 10^{-4}$, $\sigma_u^2 = 0.25$.
- (ii) Model M2: $T = 2^{19}$, $d = 0.3$, $\phi = 0.8$, $\sigma_\varepsilon^2 = 4 \times 10^{-4}$, $\sigma_u^2 = 0.25$.
- (iii) Model M3: $T = 2^{19}$, $d = 0.4$, $\phi = 0.6$, $\sigma_\varepsilon^2 = 4 \times 10^{-4}$, $\sigma_u^2 = 0.25$.

The results for model M1 are reported in Table 4. The format of the table is identical to Table 2, and gives the median, 2.5 and 97.5 percentiles of the sampling distributions for the log-periodogram estimates. Tables 5 and 6 give the counterparts for models M2 and M3, respectively. The qualitative effects of temporal aggregation are the same in each of these three models as in the benchmark model. In model M1, the sample size is reduced to $2^{16} = 65,536$,

¹⁶ However, the log-periodogram estimates of d , for the 5-min ¥–\$ returns, reported in Table 1, are all more than *ten* standard errors away from zero, and the simulated bias in these estimates is a *downward* bias.

Table 5
Sensitivity analysis: Monte-Carlo properties for model M2^a

<i>k</i>	1	12	288
$\log(y_t^{(k)2})$	0.283 (0.201, 0.369)	0.215 (0.122, 0.321)	0.061 (− 0.061, 0.210)
$y_t^{(k)2}$	0.037 (− 0.003, 0.335)	0.031 (− 0.001, 0.288)	0.017 (− 0.022, 0.270)
$ y_t^{(k)} $	0.208 (0.065, 0.466)	0.175 (0.041, 0.421)	0.072 (− 0.018, 0.312)
$h_t^{(k)}$	0.293 (0.210, 0.381)	0.294 (0.211, 0.381)	0.324 (0.154, 0.519)
$[\log(y_t^2)]^{(k)}$	0.283 (0.201, 0.369)	0.284 (0.202, 0.371)	0.319 (0.147, 0.508)
$[y_t^2]^{(k)}$	0.037 (− 0.003, 0.335)	0.037 (− 0.001, 0.336)	0.028 (− 0.020, 0.339)
$[y_t]^{(k)}$	0.208 (0.065, 0.466)	0.209 (0.065, 0.469)	0.204 (0.028, 0.598)

^aNotes: This is the counterpart to Table 2, but where the model parameters are altered from their benchmark values to specify that $\phi = 0.8$.

Table 6
Sensitivity analysis: Monte-Carlo properties for model M3^a

<i>k</i>	1	12	288
$\log(y_t^{(k)2})$	0.377 (0.324, 0.428)	0.329 (0.202, 0.450)	0.120 (− 0.068, 0.308)
$y_t^{(k)2}$	0.195 (0.027, 0.441)	0.158 (0.006, 0.399)	0.093 (− 0.032, 0.350)
$ y_t^{(k)} $	0.349 (0.275, 0.471)	0.302 (0.201, 0.451)	0.160 (− 0.033, 0.439)
$h_t^{(k)}$	0.397 (0.344, 0.446)	0.399 (0.345, 0.447)	0.434 (0.313, 0.554)
$[\log(y_t^2)]^{(k)}$	0.377 (0.324, 0.428)	0.379 (0.326, 0.429)	0.428 (0.308, 0.537)
$[y_t^2]^{(k)}$	0.195 (0.027, 0.441)	0.195 (0.028, 0.438)	0.223 (0.018, 0.513)
$[y_t]^{(k)}$	0.349 (0.275, 0.471)	0.351 (0.276, 0.472)	0.394 (0.257, 0.594)

^aNotes: This is the counterpart to Table 2, but where the model parameters are altered from their benchmark values to specify that $d = 0.4$.

corresponding to an approximate one-year time span of 5-min returns. While the ‘one year daily’ and even ‘one year hourly’ returns are totally uninformative about the long memory in the volatility, high-frequency data over the same time span quite reliably identify a value of d greater than zero. In model M2, ϕ is

raised to 0.8, while in model M3, d is raised from 0.3 to 0.4. The magnitudes of all the biases in these models are close to those in the benchmark model.

5. Concluding remarks

Semiparametric methods are ideally suited to the empirical analysis of long-memory volatility dependencies in exchange rate data, both because of the complexity of these volatility dependencies and because of the sample sizes available. Semiparametric methods, applied to a sample of more than half-a-million high-frequency spot exchange rate returns clearly suggest the presence of long-memory volatility dependencies. Meanwhile, much lower estimates for the long-memory parameter are obtained with the temporally aggregated longer horizon returns. Our simulations of a fractionally integrated stochastic volatility model, calibrated to the high-frequency returns, demonstrate that downward bias in the semiparametric estimates of the degree of fractional integration is to be expected when low-frequency data is used. In contrast, when estimating long-run dependencies in the mean, the time span of the data is of utmost importance and simply increasing the frequency of the observations over a fixed time span does little to enhance the quality of the estimates. We also propose some new low-frequency volatility measures, constructed from nonlinearly aggregated high-frequency returns that allow for nearly unbiased estimation of the long-memory volatility parameter over relatively short calendar time spans. As such, our results hold the promise of more generally being able to distinguish between genuine long-memory effects in the volatility and occasional structural breaks, through the judicious empirical analysis of high-frequency returns.

Acknowledgements

The first author would like to acknowledge the financial support by a grant from the NSF to the NBER. We are also grateful to Olsen and Associates for making the high-frequency exchange rates available, as well as Torben G. Andersen, Bo Honoré, Mark Watson, Peter Robinson and two anonymous referees for their extremely helpful comments on earlier versions of this paper. Needless to say, we are solely to blame for any errors. The views in this paper should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System or of any person associated with the Federal Reserve System.

Appendix A. Parke's method for simulating an ARFIMA series

Parke (1996) showed that an ARFIMA $(0, d, 0)$ time series can alternatively be expressed as an infinite sum of shocks each having a random duration with a certain distribution. This idea forms the basis for our method of simulating a long fractionally integrated time series. Specifically, let

$$x_t = \sum_{s=-\infty}^t g_{s,t} \sigma_\varepsilon \varepsilon_s, \quad (\text{A.1})$$

where ε_s is i.i.d. $N(0, 1)$, $g_{s,t} = 1(t - s \leq \lambda_s)$, and λ_s is a discrete i.i.d. random variable with support $\{0, 1, 2, 3, \dots\}$ such that

$$P(\lambda_s \geq h) = \frac{\Gamma(h + d)\Gamma(2 - d)}{\Gamma(h + 2 - d)\Gamma(d)}, \quad h = 0, 1, 2, 3, \dots$$

and the sequences λ_s and ε_s are mutually independent. The resulting x_t series is then a Gaussian ARFIMA $(0, d, 0)$ series. In simulating the fractionally integrated series in Section 4, the summation in (A.1) was started at $s = 0$, not $s = -\infty$, but the algorithm was allowed a long 'burn in' period. For example, to simulate a series of length 524,288, a series of length 600,000 was first generated in this way and the first 75,712 observations then discarded. The sample covariance function of the series generated in this way closely matches its population counterpart.

References

- Adenstedt, R., 1974. On large-sample estimation for the mean of a stationary random sequence. *Annals of Statistics* 2, 1095–1107.
- Agiakloglou, C., Newbold, P., Wohar, M., 1992. Bias in an estimator of the fractional differencing parameter. *Journal of Time Series Analysis* 14, 235–246.
- Andersen, T.G., 1994. Stochastic autoregressive volatility: a framework for volatility modelling. *Mathematical Finance* 4, 75–102.
- Andersen, T.G., Bollerslev, T., 1997a. 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., 1997b. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115–158.
- Andersen, T.G., Bollerslev, T., 1998. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review* 39, 885–905.
- Baillie, R.T., 1996. Long-memory processes and fractional integration in econometrics. *Journal of Econometrics* 73, 5–59.
- Baillie, R.T., Bollerslev, T., Mikkelsen, H.O., 1996. Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 74, 3–30.
- Beran, J., 1994. *Statistics for Long-Memory Processes*. Chapman and Hall, New York.
- Berkowitz, J., Diebold, F.X., 1998. Bootstrapping multivariate spectra. *Review of Economics and Statistics* 80, 664–667.

- Berkowitz, J., Kilian, L., 1999. Recent developments in bootstrapping time series, *Econometric Reviews*, forthcoming.
- Bollerslev, T., Domowitz, I., 1993. Trading patterns and prices in the interbank foreign exchange market. *Journal of Finance* 48, 1421–1443.
- Bollerslev, T., Mikkelsen, H.O., 1999. Long-term equity anticipation securities and stock market volatility dynamics. *Journal of Econometrics* 92, 75–99.
- 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.
- Chambers, M.J., 1998. Long-memory and aggregation in macroeconomic time series. *International Economic Review* 39, 1053–1072.
- Cheung, Y.W., 1993. Tests for fractional integration: a Monte Carlo investigation. *Journal of Time Series Analysis* 14, 331–345.
- Cheung, Y.W., Lai, K., 1995. A search for long-memory in international stock market returns. *Journal of International Money and Finance* 14, 597–615.
- Dacorogna, M.M., Müller, U.A., Nagler, N.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.
- Davies, R.B., Harte, D.S., 1987. Tests for Hurst effect. *Biometrika* 74, 95–102.
- Delgado, M.A., Robinson, P.M., 1996. Optimal spectral bandwidth for long-memory. *Statistica Sinica* 6, 97–112.
- Ding, Z., Granger, C.W.J., 1996. Modeling volatility persistence of speculative returns: a new approach. *Journal of Econometrics* 73, 185–215.
- Drost, F.C., Nijman, T.F., 1993. Temporal aggregation of GARCH processes. *Econometrica* 61, 909–927.
- Drost, F.C., Werker, B., 1996. Closing the GARCH gap: continuous time GARCH modelling. *Journal of Econometrics* 74, 31–57.
- Franke, J., Härdle, W., 1992. On bootstrapping kernel spectral estimates. *Annals of Statistics* 20, 121–145.
- Geweke, J., 1998. Comment on real and spurious long-memory properties of stock-market data (by I. Lobato and N.E. Savin). *Journal of Business and Economic Statistics* 16, 269–271.
- Geweke, J., Porter-Hudak, S., 1983. The estimation and application of long-memory time series models. *Journal of Time Series Analysis* 4, 221–237.
- Granger, C.W.J., Ding, Z., Spear, S., 1997. Stylized Facts of the Temporal and Distributional Properties of Daily Speculative Markets. Unpublished manuscript, UC San Diego.
- Granger, C.W.J., Joyeux, R., 1980. An introduction to long-memory time series models and fractional differencing. *Journal of Time Series Analysis* 1, 15–29.
- Harvey, A.C., 1998. Long-memory in stochastic volatility. In: Kinght, J., Satchell, S. (Eds.), *Forecasting Volatility in Financial Markets*. Butterworth-Heinemann, London.
- Hiemstra, C., Jones, J.D., 1997. Another look at long-memory in common stock returns. *Journal of Empirical Finance* 4, 373–410.
- Hosking, J.R.M., 1981. Fractional differencing. *Biometrika* 68, 165–176.
- Hsieh, D.A., 1991. Implications of observed properties of daily exchange rate movements. *Journal of International Financial Markets* 1, 61–71.
- Hurvich, C.M., Beltrao, B.K., 1994. Automatic semiparametric estimation of the long-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 square error of Geweke and Porter-Hudak's estimates of the memory parameter of a long-memory time series. *Journal of Time Series Analysis* 19, 19–46.
- Jacobsen, B., 1996. Long-term dependence in stock returns. *Journal of Empirical Finance* 3, 393–417.

- Janacek, G.J., 1982. Determining the degree of differencing for time series via the log spectrum. *Journal of Time Series Analysis* 3, 177–183.
- Lo, A., 1991. Long-term memory in stock market prices. *Econometrica* 59, 1279–1313.
- Lobato, I., Robinson, P.M., 1998. A nonparametric test for $I(0)$. *Review of Economic Studies* 65, 475–495.
- Lobato, I., Savin, N.E., 1998. Real and spurious long-memory properties of stock-market data. *Journal of Business and Economic Statistics* 16, 261–268.
- Mandelbrot, B., 1971. When can price be arbitrated efficiently? A limit to the validity of the random walk and Martingale models. *Review of Economics and Statistics* 53, 225–236.
- McLeod, A.I., Hipel, K.W., 1978. Preservation of the rescaled adjusted range. *Water Resources Research* 14, 491–518.
- Müller, U.A., Dacorogna, M.M., Olsen, R.B., Pictet, O.V., Schwarz, M., Morgeneegg, 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.
- Nijman, T.E., Palm, F.C., 1990. Parameter identification in ARMA processes in the presence of regular but incomplete sampling. *Journal of Time Series Analysis* 11, 239–248.
- Parke, W.R., 1996. What is a fractional unit root? Unpublished manuscript, University of North Carolina.
- Robinson, P.M., 1991. Testing for strong serial correlation and dynamic conditional heteroskedasticity in multiple regression. *Journal of Econometrics* 47, 67–84.
- Robinson, P.M., 1994a. Time series with strong dependence. In: Sims, C.A. (Ed.), *Proceedings of the Sixth World Congress on Advances in Econometrics*. Cambridge University Press, Cambridge.
- 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 for time series with long range dependence. *Annals of Statistics* 23, 1048–1072.
- Robinson, P.M., 1995b. Gaussian semiparametric estimation of long range dependence. *Annals of Statistics* 23, 1630–1661.
- Robinson, P.M., 1999. The memory of stochastic volatility models. Unpublished manuscript, London School of Economics.
- Robinson, P.M., Zaffaroni, P., 1996. Modeling nonlinearity and long-memory in time series. In: Cutler, C., Kaplan, D. (Eds.), *Nonlinear Dynamics and Time Series: Building a Bridge between the Natural and Statistical Sciences*. American Mathematical Society, Providence, RI.
- Schwert, G.W., 1989. Why does stock market volatility change over time? *Journal of Finance* 44, 1115–1153.
- Shiller, R.J., Perron, P., 1985. Testing the random walk hypothesis: power versus frequency of observations. *Economics Letters* 18, 381–386.
- Zaffaroni, P., 1997. Gaussian estimation of long-range dependent volatility in asset prices. Unpublished manuscript, London School of Economics.