

Semiparametric Density Forecasts of Daily Financial Returns from Intraday Data*

Mark Hallam[†] *City University London*
Jose Olmo, *University of Southampton*

March 9, 2013

*The authors would like to thank the Editor, the Associate Editor, two anonymous referees and seminar participants at both City University and the 2011 Econometrics Society European Meeting in Oslo for useful comments and suggestions.

[†]Corresponding author: Department of Economics, City University London, Whiskin Street, EC1R 0JD, United Kingdom. E-mail: mark.hallam.1@city.ac.uk

Abstract

In this paper we propose a new method for producing semiparametric density forecasts for daily financial returns from high-frequency intraday data. The daily return density is estimated directly from intraday observations that have been appropriately rescaled using results from the theory of unifractal processes. The method preserves information concerning both the magnitude and sign of the intraday returns and allows them to influence all properties of the daily return density via the use of nonparametric specifications for the daily return distribution. The out-of-sample density forecasting performance of the method is shown to be competitive with existing methods based on intraday data for exchange rate and equity index data.

JEL: C58, C22, G17

KEYWORDS: density forecasting, unifractal, high-frequency data, semiparametric

Over the past decade there has been a dramatic increase in the availability of intraday financial data, resulting in an extensive literature on the use of high-frequency data in financial econometrics. These data obviously allow for the study of financial market behaviour at intraday timescales, but they also contain potentially valuable information for longer timescales, which are arguably of more interest for most market participants. As a result, there have been efforts to incorporate intraday data into the modelling and forecasting of financial variables at daily or even lower frequencies.

The most notable example is provided by the large literature on realised volatility, a concept that was first properly formalised by Andersen et al. (2001). Realised volatility and related measures allow the unobservable daily volatility to be estimated from intraday returns and it has been found (see for example Andersen et al., 2003) that such measures can provide significant improvements in the modelling and forecasting of daily return volatility compared to models using only daily data.

Whilst return volatility is undoubtedly a variable of substantial interest, there are situations in finance in which information concerning just the first two moments of the distribution of returns is not sufficient. Perhaps most obviously, risk management problems, such as the calculation of Value-at-Risk, require knowledge of particular quantiles of the return distribution. In addition, it has been shown that higher moments, such as skewness and kurtosis, are time varying and relevant for problems of portfolio allocation and asset pricing (see for example Harvey and Siddique, 2000, or Dittmar, 2002).

However, as noted by Žikeš (2009), the use of intraday data to model and forecast characteristics of daily returns beyond the first two moments has not yet received much attention. Notable exceptions include Andersen et al. (2003), Giot and Laurent (2004), Clements, Galvão, and Kim (2008) and Maheu and McCurdy (2010), all of which extend the use of realised volatility measures to either the quantiles or the density function of daily returns. The methods used by these previous studies consist of two components. The first is a parametric time series model for volatility incorporating one or more realised volatility measures, which is used to model and produce point forecasts for daily volatility. The second component is a typically a parametric distributional assumption about daily returns, allowing density or quantile forecasts for daily returns to be produced from the point forecasts of daily realised volatility.

There are two potential weaknesses with this approach; firstly, the high-frequency data enter only through the realised volatility measures and so any information provided by the sign of the intraday returns is lost when they are squared. Furthermore, the intraday data can only directly influence the second moment of the daily return distribution. Secondly, with the exception of Clements, Galvão, and Kim (2008) who also consider an empirical distribution for returns, these previous studies require a specific parametric form to be chosen for the distribution of daily returns; choosing the most appropriate parametric distribution for financial returns is difficult, particularly in a dynamic context, and the density forecasts produced by misspecified parametric models will generally be misleading. The semiparametric quantile regression approach of Žikeš (2009) avoids the last of these problems, but only produces estimates for specific quantiles rather than the complete distribution or density.

In the current paper, a new approach is proposed for the estimation and forecasting of daily return densities from intraday data, based on the theory of unifractal processes. Under the assumption that the return process is unifractal, the distribution of returns at any pair of timescales is identical after rescaling by an appropriate factor; this factor can be estimated for a particular time series and used to rescale the intraday data for a given time period, such that they are equal in distribution to daily returns. The density of daily returns can then be estimated from these rescaled intraday observations.

The proposed method has two theoretical advantages compared to existing methods based on realised volatility. Firstly, the daily return density is estimated directly from these rescaled intraday observations (rather than squared or absolute values), thus preserving information contained in both the magnitude and sign of the intraday returns. Secondly, because a large sample of rescaled intraday observations are obtained for each trading day, it is possible to apply a range of estimators to these rescaled intraday returns to estimate the daily return density for a given trading day.

In particular, a nonparametric density estimation approach is proposed using a standard kernel density estimator, which allows the intraday data to influence all aspects of the daily return density without being complex to implement or computationally demanding. However, the use of nonparametric density estimators precludes the use of standard dynamic structures for forecasting and so a new method is proposed and developed that imposes a parametric dynamic structure directly on the time series of densities themselves, with the relevant parameters selected using concepts from the literature on density forecast combination.

The structure of the paper is as follows: Section 1 presents the relevant theory on unifractal processes and describes how these results can be employed to estimate the density of daily returns from intraday return data. Section 2 details the chosen estimation methods for both the scaling factor used to rescale the intraday data and the density function of the rescaled intraday data. Section 3 explores the issue of producing density forecasts for daily returns from a time series of estimated return densities. Section 4 presents an empirical application comparing the density forecasting performance to existing methods using intraday data on equity indexes and exchange rates and finally Section 5 concludes.

1 Unifractal & Multifractal Processes

In order to estimate the density of daily returns from intraday data a method is required for formally linking the characteristics of return distributions across different sampling frequencies. Instead of the realised volatility measures previously employed in the literature, the proposed method relies on results from the theory of self-affine or unifractal processes.

On an intuitive level, such stochastic processes exhibit some form of *scale invariance*, such that the behaviour of the process at one timescale is, after an appropriate transformation, identical in a statistical sense to that observed at another time scale. A large number of empirical studies have confirmed the existence of this type of distributional scaling behaviour in a wide range of financial time series¹ and this has led to the development of several asset pricing models that explicitly reproduce this distributional scaling behaviour².

The current section begins with a brief summary of the theoretical properties of these processes (with more detailed treatments found in Calvet and Fisher, 2002 or Kantelhardt, 2009), before exploring how these properties can be employed to estimate the density of daily returns from intraday return data.

1.1 A Review of Unifractal and Multifractal Processes

The distributional scaling behaviour of a unifractal or self-affine process can be defined by a simple expression that links the distribution of the process at different sampling intervals. Formally, unifractal or self-affine processes can be defined in the following way:

Definition 1.1. A process is said to be *self-affine* or *unifractal* if for some $H > 0$, all $c \geq 0$ and all $t_1, t_2, \dots, t_k \geq 0$ it obeys the distributional scaling relationship

$$\{X(ct_1), X(ct_2), \dots, X(ct_k)\} \stackrel{d}{=} \{c^H X(t_1), c^H X(t_2), \dots, c^H X(t_k)\} \quad (1.1)$$

which can be expressed more compactly as:

$$X(ct) \stackrel{d}{=} c^H [X(t)] \quad (1.2)$$

If the increments of the process are stationary, then the distributional scaling law of (1.2) also holds for the increments of the process at any given sampling interval:

$$X(t + c\Delta t) - X(t) \stackrel{d}{=} c^H [X(t + \Delta t) - X(t)] \quad (1.3)$$

The parameter H is known as the self-affinity index and describes the relationship between the distributions of the return process at different timescales. Common examples of unifractal processes in finance include the standard Brownian motion, for which $H = 1/2$, and also the more general fractional Brownian motion, for which $H \neq 1/2$. In the current context, under the assumption of unifractality these scaling laws imply that the distribution of returns at different timescales or sampling intervals is identical after rescaling by a factor that depends on the characteristics of

the particular return process (via H) and the difference between the two sampling intervals (via c).

One can also consider the more general class of multifractal processes, which allow for a more flexible relationship between distributions across different sampling frequencies. In the case of a multifractal process, equation (1.2) is generalised to:

$$X(ct) \stackrel{d}{=} c^{H(c)}[X(t)] \quad (1.4)$$

where the scaling factor c^H has been replaced by the more general function of c , $c^{H(c)}$. An alternative characterisation of scaling behaviour is often used in the multifractal case, where it can be shown (see for example Mandelbrot et al., 1997) that a stochastic process $X(t)$ with increments $X(t + \Delta t) - X(t)$ is multifractal if these increments are stationary and satisfy:

$$E[|X(t + \Delta t) - X(t)|^q] = c(q)(\Delta t)^{\tau(q)+1} \quad (1.5)$$

where $c(q)$ and $\tau(q)$ are deterministic functions of q . The function $\tau(q)$ in (1.5) is referred to as the *scaling function* and describes how different moments of the absolute increments of the process $X(t)$ scale with the sampling interval, Δt . The scaling function is strictly concave for a multifractal process but linear for a unifractal process and as with the self-affinity index above, can be estimated for a particular time series using various methods (see the survey by Kantelhardt, 2009).

1.2 Estimating Daily Return Densities from Intraday Data

To proceed we require the following assumption:

A.1. The stochastic logarithmic price process, $X(t)$, is unifractal (or self-affine) and has stationary increments $Y^\Delta(t)$, where $Y^\Delta(t) = X(t + \Delta) - X(t)$ is the return process over the interval Δ ³.

Whilst there are numerous empirical studies confirming the existence of distributional scaling behaviour in a wide range of financial assets (see footnote 2 for examples), these studies are typically concerned with the more general multifractal case, which includes the unifractal scaling of assumption A.1 as a special case. Whether a given sample of data is consistent with assumption A.1 is primarily an empirical issue and will be discussed further in the context of the current dataset during the empirical exercise.

For ease of exposition, the current subsection focuses on the specific example of estimating daily return densities from 5-minute intraday data. However, under assumption A.1 the distributional scaling laws of (1.2) and (1.3) hold for any pair of timescales and so the method could be used to estimate the density of returns for any given sampling interval from those at some higher frequency. However, if attention is restricted to a specific pair of sampling intervals then a weaker condition than A.1 would be sufficient; it would then only be strictly necessary for (1.2) to hold for the sampling intervals of interest and not for all possible sampling intervals as is the case for a true unifractal process⁴.

Assume that a series of 5-minute returns are observed for a financial asset over a given period and denote this set of intraday returns and the corresponding probability density function by $Y_{I,t}$ and $f(y_I)$ respectively, where the I subscript is used to indicate returns at the intraday frequency. Under assumption A.1 the corresponding return process must satisfy the distributional scaling laws of equation (1.3). If we denote the density of daily returns over the same period by $f(y_D)$, then in the context of the current example these scaling laws imply that:

$$f(y_D) = f(c^H y_I) \tag{1.6}$$

From (1.6), the density of daily returns is equal to the density of the 5-minute intraday returns, when these intraday returns have been appropriately rescaled by the factor c^H . The self-affinity index, H , can be estimated from the intraday data using various methods and from (1.1) the value of c is determined solely by the relative lengths of the two sampling intervals; for a market with 24-hour trading, as is typical for exchange rates, the appropriate value of c in the current example would be 288, since there are 288 5-minute returns observed over a 24-hour period. For a market with shorter trading hours c will take a smaller value.

Therefore, under the assumption that the return process is unifractal, an estimate of the relevant rescaling factor, c^H , can be obtained from the intraday data and the density of daily returns for a given period can be estimated from the intraday returns observed over the same period via the distributional scaling laws of the previous section.

2 Estimating the Density of Daily Returns from Intraday Data

Estimation of the daily return density from the intraday data occurs in two stages: in the first stage, the self-affinity index is estimated from the intraday returns observed over a given time period and the resulting estimate is then used to rescale the intraday returns as discussed in Section 1. In the second stage, the probability density function of these rescaled returns is estimated using the chosen density function estimator, providing an estimate of the daily return density for the same time period. Sections 2.1 and 2.2 will discuss in turn the methods employed for the first and second stages of estimation, respectively.

2.1 Estimation of the Self-Affinity Index

Numerous estimators are available for the self-affinity index of a unifractal process and several studies⁵ have demonstrated that the relative performance of these estimators can vary substantially in small samples. Due to their typically strong performance relative to other estimators and their suitability for the current dynamic estimation environment, attention was restricted to the popular Detrended Fluctuation Analysis (DFA) and Detrended Moving Average (DMA) estimators. The DMA estimator is detailed below, since it was generally found to produce superior density forecasting performance to DFA for the current dataset. Further details of these and other estimators can be found in the survey by Kantelhardt (2009).

For a given time series of observations, $y(t)$, observed at discrete times $t = 1, \dots, T$, the DMA estimator is implemented in the following way. Initially, multiple filtered versions of the original series are produced using a standard moving average (MA) filter with a range of different MA window sizes, n , from $n = n_{min}$ to $n = n_{max}$ ⁶. The filtered series produced from each window size, n , is given by:

$$y_n^{MA}(t) = \frac{1}{n} \sum_{k=0}^{n-1} y(t-k)$$

For each moving average filtered series, $\{y_n^{MA}(t) : n_{min} \leq n \leq n_{max}\}$, the value of σ_n^2 (sometimes referred to as the generalised variance) is calculated, where:

$$\sigma_n^2 = \frac{1}{T-n} \sum_{i=n}^T [y(i) - y_n^{MA}(i)]^2 \quad (2.1)$$

For a unifractal process, $\sigma_n \propto n^H$ and so an estimate of the self-affinity index, H , can be obtained from the estimated slope coefficient in a linear regression of the logarithm of σ_n obtained from equation (2.1) on the logarithm of n , for a range of values of n .

A brief note is required at this point regarding the issue of intraday seasonality. It is well known that intraday financial data can exhibit strong deterministic seasonal patterns that pose a problem for certain estimation methods. Because of this, the DMA estimator was applied both to the raw 5-minute returns and to intraday data seasonally adjusted using the method of Andersen et al. (2003). The resulting estimates of the self-affinity index from both cases were similar, as was the predictive ability of the resulting density forecasts with neither approach having a consistent advantage over the other. This is perhaps due to the fact that the chosen DMA estimator automatically performs local detrending across various window sizes in the process of estimating the self-affinity index, thus eliminating some or all of the seasonal patterns. Following this finding, it was decided to use the simple unadjusted intraday data for estimation of the self-affinity index, rather than the seasonally adjusted data.

2.2 Estimation of Daily Return Densities

Given that a large number of rescaled intraday returns can be obtained even for a single trading day, it is possible to apply a variety of estimators for the probability density function, including both parametric and nonparametric methods. Whilst the possibility of using nonparametric estimation methods is perhaps most interesting, one example of each class of estimator will be employed for the empirical analysis in order to investigate the potential gains from such a nonparametric approach. The current density function estimators are both intentionally simple, but could be replaced with more complex estimators without substantial modifications to the method.

For the parametric case, a standard 3-parameter location-scale t -distribution is fitted to the rescaled intraday returns using maximum likelihood. The mathematical details of this approach will not be discussed here, since the techniques involved are standard and have been discussed in detail elsewhere.

In the nonparametric case, the standard kernel estimator for a univariate density function is employed (originally due to Parzen, 1962). Consistent with previous notation, the observed series of T intraday returns are denoted by $\{Y_{I,t}\}_{t=1}^T$ and the rescaled intraday returns are denoted by $\{Y_{R,t}\}_{t=1}^T$, where $Y_{R,t} = c^{\hat{H}} Y_{I,t}$ and \hat{H} is an estimate of the self-affinity index. The kernel density estimator for the density of the rescaled intraday returns $f(y_R)$ is then given by:

$$\hat{f}(y_R) = \frac{1}{hT} \sum_{t=1}^T k\left(\frac{Y_{R,t} - y_R}{h}\right) \quad (2.2)$$

where $k(\cdot)$ is a nonnegative and bounded kernel function and the parameter h is a bandwidth or smoothing parameter.

Standard regularity conditions on the kernel function and bandwidth parameter h (see for example Li and Racine, 2006) guarantee uniform consistency of the kernel estimator in (2.2) in the case of independent data. However, it is well known that financial returns at the daily and intraday frequencies considered here typically exhibit some form of serial dependence. However, the same uniform almost sure rate of convergence for the standard kernel estimator is preserved in the case of weakly dependent data, provided that the structure and strength of serial dependence satisfies certain conditions⁷. In the current context, we require the following assumption:

A.2 The intraday return process $Y_{I,t}$ observed over each individual trading day is strictly stationary and either ρ -mixing with ρ -mixing coefficients satisfying $\rho(j) = O(j^{-(1+\epsilon)})$, or α -mixing with α -mixing coefficients satisfying $\alpha(j) = O(j^{-(1+\epsilon)})$, for some $\epsilon > 0$.

It should be noted that A.2 only requires the conditions on serial dependence to hold for each individual trading day of intraday data in isolation and not for the whole intraday return process over the complete sample period. This is because, as discussed during the following section, the daily return density for each trading day is estimated just from the rescaled intraday returns observed over that trading day and not from multiple days. Intuitively it seems plausible that dependence between the intraday returns in each period and those at the start of trading becomes increasingly small as we move towards the end of the trading day. This argument could fail if multiple days of intraday data were used for estimation, since the patterns of intraday seasonality present in intraday returns may introduce long-range dependence in the intraday return process, which is not permitted by assumption A.2. In addition, estimates of the long-range dependence parameter from the complete sample of intraday data and the autocorrelation functions for longer lag lengths provide no strong evidence for the presence of long-range dependence in the level of the intraday return series at longer horizons. Assumptions A.1 and A.2 lead to the following proposition, the proof of which follows from standard consistency arguments for nonparametric methods (see for example Li and Racine, 2006):

Proposition 1 Under assumptions A.1 and A.2, the kernel estimator for the probability density function of rescaled intraday returns given in equation (2.2) is a consistent estimator of the probability density function of daily returns.

3 Forecasting the Density of Daily Returns

In order to produce one-step-ahead out-of-sample density forecasts for daily returns from the estimated daily return densities obtained using the previously presented methods, we must impose a dynamic structure to describe the evolution of the daily return density over time. For both the parametric and nonparametric approaches of the previous section a simple autoregressive structure is employed, however the implementation necessarily differs depending on which class of estimator is used for the density function.

Section 3.1 provides a formal description of the forecasting environment and Sections 3.2 and 3.3 present the dynamic structures used for density forecasting in the parametric and nonparametric cases. Finally, Section 3.4 discusses how the form of these dynamic structures can be estimated in practice in order to produce forecasts for the daily return density.

3.1 The Forecasting Environment

It is assumed that a series of intraday returns are observed over a period of T days, together with a corresponding series of daily returns. A standard rolling window scheme is employed to produce one-step-ahead density forecasts for daily returns from this intraday data; specifically, at day m , an estimate of the self-affinity index, H , is produced using the intraday data from day 1 up to day m . This estimate is denoted by \hat{H}_m and is then used to rescale the intraday data for the same m -day period.

The density of the rescaled intraday data for each of the m days is then estimated using the methods presented in Section 2.2 to produce a time series of m estimated daily return densities, denoted by $\{\hat{f}_t(y) : 1 \leq t \leq m\}$, with the estimated density for each trading day produced using only the rescaled returns observed during that day. These m estimated daily return densities are used to produce in-sample estimates of the relevant parameters of the chosen dynamic structure and these estimated values are then used to produce an out-of-sample one-step-ahead forecast for the density of daily returns at time $m + 1$. This density forecast is denoted by $\tilde{f}_m(y)$, with the subscript indicating that the forecast is conditional on the information available at time m . The estimation window is then rolled forward by one day and the procedure is repeated to produce a one-step-ahead density forecast for use on day $m + 2$, denoted $\tilde{f}_{m+1}(y)$, using the intraday data from day 2 up to day $m + 1$.

This procedure can be repeated over the rest of the sample to produce a sequence of N out-of-sample one-step-ahead density forecasts for daily returns, where $N = T - m$. Finally it should be noted that each time the m -day estimation window is rolled forward, an updated estimate of the self-affinity index is produced. This permits the scaling properties of the return process to change over time, allowing for more flexible scaling behaviour compared to assuming a constant global value of the self-affinity index.

3.2 Dynamic Structure for Density Forecasting: Parametric Case

As previously stated, a simple autoregressive dynamic structure is employed in order to produce one-step-ahead density forecasts from the time series of estimated

densities obtained using the methods outlined in the previous section. In the case of parametric specifications for the daily return density this is relatively simple, since the dynamic structure describing the evolution of the return density over time can be imposed via the parameters of the chosen distribution; this allows point forecasts for the distributional parameters to be produced that in turn provide a density forecast for daily returns.

For the chosen location-scale t -distribution each of the three distributional parameters (location, scale and degrees of freedom) can be modelled separately using a standard univariate $AR(p)$ model; in-sample estimates of the relevant autoregressive parameters can be obtained from the m estimated daily return densities over the in-sample period and these estimates can then be used to produce one-step-ahead out-of-sample forecasts for the distributional parameters. The only complication is that some additional constraints must be imposed to ensure the forecasted parameter values from an autoregressive structure are within the permitted parameter space for the distribution. An alternative approach was also tested in which the autoregressive dynamic structure was imposed on the moments of the rescaled intraday returns instead of the distributional parameters; from one-step-ahead forecasts for the mean, variance and kurtosis of the rescaled returns, forecasts for the parameters of the location-scale t -distribution can be uniquely determined. Performance of this alternative was similar overall, but typically slightly inferior to that obtained by directly modelling the distributional parameters.

3.3 Dynamic Structure for Density Forecasting: Nonparametric Case

When the daily return densities are estimated using nonparametric methods such as the kernel estimator of Section 2.2, the approach outlined above is clearly inapplicable. Instead, an alternative method is now developed based on imposing a parametric dynamic structure on the evolution of the complete probability density. It is assumed that the entire density at time $t + 1$ depends on several past densities, with this temporal dependence again assumed to follow an autoregressive structure. Whilst this choice is perhaps slightly ad-hoc, autoregressive structures have been shown to work well in the context of quantiles (see for example Engle and Manganelli, 2004).

The simplest case would be to assume that each of the lagged densities have a constant coefficient, making the density of daily returns at time $t + 1$ a simple weighted sum of the return densities from t to $t - p + 1$, which can be expressed as:

$$f_{t+1}(y) = \beta_1 f_t(y) + \beta_2 f_{t-1}(y) + \dots + \beta_p f_{t-p+1}(y) + u_{t+1}(y) \quad (3.1)$$

The D subscripts indicating daily return densities have been suppressed for notational simplicity and the error terms $u_{t+1}(y)$ are a martingale difference sequence for all values of y in the domain of $f_{t+1}(y)$. The true daily return density in each period, $f_s(y)$, is unknown, but replacing with the corresponding density estimate obtained using the method of Section 3, $\hat{f}_s(y)$, gives:

$$\hat{f}_{t+1}(y) = \gamma_1 \hat{f}_t(y) + \gamma_2 \hat{f}_{t-1}(y) + \dots + \gamma_p \hat{f}_{t-p+1}(y) + v_{t+1}(y) \quad (3.2)$$

where again, the error terms $v_{t+1}(y)$ are martingale difference sequences for all values of y . In-sample estimates of the autoregressive parameters $\{\gamma_i : 1 \leq i \leq p\}$ can be produced from the time series of estimated daily return densities, with the resulting estimates denoted by $\{\hat{\gamma}_i : 1 \leq i \leq p\}$. When combined with the estimated daily return densities for periods t to $t - p - 1$, these in-sample estimates of the autoregressive parameters can then be used to produce a one-step-ahead out-of-sample density forecast for the density of daily returns for day $t + 1$ from:

$$\tilde{f}_t(y) = \hat{\gamma}_1 \hat{f}_t(y) + \hat{\gamma}_2 \hat{f}_{t-1}(y) + \dots + \hat{\gamma}_p \hat{f}_{t-p+1}(y) \quad (3.3)$$

Note that $\tilde{f}_t(y)$ is used to denote the forecast of the daily return density made conditional on the information available at time t (but for use at time $t + 1$) and the tilde is used to distinguish it from the in-sample density estimate at time t . Clearly constraints need to be imposed on the values of the estimated parameters to ensure that the density forecast produced by (3.3) is always a valid probability density, but for the autoregressive structure above this is guaranteed simply by constraining the estimated parameters sum to unity and are all non-negative.

Although the dynamic structure of equation (3.1) is simple to interpret and implement, one potential limitation is that it does not allow the dependence between the density at time t and that at time $t - s$ to vary across different regions of the density⁸. It could be generalised by replacing the constant coefficients, β_1, \dots, β_p , with functions of y , but ensuring that the resulting density forecast is a valid probability density function is no longer straightforward and so this is left as a possible area for future research.

3.4 Estimating Autoregressive Parameter Values

For the case of parametrically estimated densities the parameter values for the univariate autoregressive models required to produce forecasts can be estimated straightforwardly using standard techniques. For the case of nonparametrically estimated densities the situation is again more complex and alternative techniques must be developed.

The simplest option is to impose some fixed vector of values for the autoregressive parameters $\{\gamma_q : 1 \leq q \leq p\}$ in (3.2) over all time periods. Although this approach may seem overly simplistic, it has been shown in the literature on forecast combination that a simple average of forecasts can perform better than a combination chosen to minimise some statistical loss function (see Timmermann, 2006). Although the current problem is not identical to that of forecast combination, a simple specification with $\gamma_q = 1/p$ for $q = 1 \dots p$ is included in the empirical exercise to test whether the same ad-hoc dynamic structure can also perform well in the current context.

Assuming instead that for a given period we wish to identify the autoregressive parameter vector that produces the most accurate one-step-ahead out-of-sample density forecast, then an appropriate loss function needs to be selected to formally define what constitutes the ‘best’ forecast. Given that a probability density is defined across a range of values, conventional point-based measures of accuracy cannot be applied and an alternative loss function appropriate for probability densities must be employed.

As a loss function the current method employs the well-known logarithmic score, which is closely related to the Kullback-Liebler information criterion (KLIC). The logarithmic score has been employed in the literature on both density forecast comparison and the optimal combination of density forecasts; the second of these applications (see Hall & Mitchell, 2007) is particularly relevant, since the problem of choosing the parameter values for the simple autoregressive structure of equation (3.2) is mathematically equivalent to that of identifying the optimal weights for the commonly studied linear combination of density forecasts.

Denoting the one-step-ahead density forecast for daily returns produced at time $t - 1$ for use in period t as before by $\tilde{f}_{t-1}(y_t)$ and the actual daily return observed at time t by y_t^* , the average logarithmic score over the periods $t = 1, \dots, S$ is given by:

$$\frac{1}{S} \sum_{t=0}^{S-1} \ln \tilde{f}_{t-1}(y_t^*) \quad (3.4)$$

Given that better forecasting models should on average assign higher probabilities to the outcome that actually occurred, higher values of the average logarithmic score provide evidence of superior predictive ability. The autoregressive coefficients in (3.2) can therefore be chosen in order to maximise the average logarithmic score. More formally, the vector of estimated autoregressive parameters $\hat{\gamma}$, is obtained as the solution to:

$$\hat{\gamma} = \arg \max_{\gamma} \frac{1}{S} \sum_{t=0}^{S-1} \ln \tilde{f}_t(y_{t+1}; \gamma) \quad \text{s.t.} \quad \sum_{i=1}^p \gamma_i = 1 \text{ and } \gamma_i \geq 0 \text{ for } i = 1, \dots, p$$

The constraints ensure that the resulting density forecast is a valid probability density function and the maximisation problem is solved using a numerical optimisation procedure.

For each of the rolling m -day in-sample estimation windows, an estimate of the autoregressive parameter vector, γ , is obtained as the solution to the above optimisation problem for the $m - p$ in-sample forecasts in the current m -day window. The resulting in-sample estimate of the parameter vector, $\hat{\gamma}$, is then used to produce an out-of-sample one-step-ahead density forecast for daily returns, for use in the following period.

The method for density forecast comparison employed in the empirical section uses an alternative scoring rule known as the continuous ranked probability score (CRPS). An equivalent numerical optimisation procedure could be employed to select the values of the autoregressive parameter values based on the CRPS instead of the logarithmic score, however this alternative is much more computationally demanding and does not appear to provide substantial gains in predictive ability for the current dataset.

4 Empirical Application

The current section applies the new semiparametric density forecasting framework to both foreign exchange and equity data in order to compare the performance of the

resulting density forecasts with those of existing methods. Section 4.1 describes the dataset employed for the empirical analysis and Section 4.2 details the benchmark density forecasting model used for comparison. Section 4.3 discusses the statistical test employed to compare the relative density forecasting performance of the models and finally Section 4.4 presents the empirical results.

4.1 Data

The data used were obtained from Olsen Associates and consist of intraday 5-minute observations from 3rd January 2007 until 31st December 2010 on the Euro (EUR) and Japanese Yen (JPY) exchange rates against the US Dollar (USD) and the levels of the S&P500 and NASDAQ-100 equity indexes. The choice of 5-minute data was guided by the desire to exploit as much of the potentially valuable intraday information as possible, whilst avoiding the distortions caused by market microstructure effects typically encountered at very short sampling intervals⁹.

All weekends and other non-trading days were removed from the raw 5-minute price or level data. For the equity indexes this was straightforward, since these markets have well-defined opening hours. The list of non-weekend closures was constructed from the historical list of holidays on the NYSE website, however there were also 9 days for which the market was open, but for reduced hours; the analysis was performed with these partial trading days both removed and included, but the choice had no significant effect on the results.

For the exchange rate series trading is typically continuous, but slows substantially over weekends and certain holidays. Following Andersen et al. (2001), the end of each 24-hour trading day was taken to be 21:00 GMT and the 48-hour weekend periods between Friday 21:05 and Sunday 21:00 GMT were removed from the 5-minute price series. Apart from weekends, only for Christmas Day and New Year's Day was trading noticeably slower for the EUR/USD and JPY/USD markets and so only these holidays were omitted from the exchange rate series. The analysis was also performed with a larger list of holidays removed from the EUR/USD and JPY/USD series, but again this did not significantly influence the results.

This process leaves a sample size of 1008 and 1037 trading days for the equity index and exchange rate series respectively. Continuous 5-minute returns were then constructed from the first difference of the log-price series for each asset, with the first 5-minute return for each day calculated between the closing price in the previous trading day and the opening price in the current day (thus including any overnight or weekend effects). Daily return series were also constructed from the last 5-minute price observed in each trading day for use in forecast comparison.

As previously discussed in Section 1, the proposed density forecasting method is only strictly valid when the distributional scaling behaviour of the return process is consistent with that of a unifractal process, rather than the more general class of multifractal processes. In practice the method should still be applicable even if this assumption is not satisfied exactly, provided that the distributional scaling behaviour of a unifractal process still provides a good approximation of the true scaling behaviour of the process. Nonetheless, larger deviations from unifractal scaling are still likely to invalidate the approach and so is important to check the validity of this assumption for the particular return series of interest.

Although there is no formal statistical method for testing whether the scaling behaviour of a given sample is consistent with that of a unifractal or a multifractal

process, an informal graphical assessment can be based on the scaling function, $\tau(q)$ from equation (1.5). As previously discussed in Section 1.1, this function is strictly concave for a multifractal process and linear for a unifractal process with equation $\tau(q) = Hq - 1$, where H is the self-affinity index.

The solid lines in Figure 1 are the estimated scaling functions obtained from the 5-minute data for each of the series under the assumption of multifractality, using the standard partition or structure function approach (see Kantelhardt, 2009 for details). The dashed lines in each sub-plot are the linear scaling functions obtained under the assumption of unifractality, with the self-affinity index (and therefore the slope of $\tau(q)$) estimated using the DMA estimator of Section 2.1. These functions are plotted over the domain $0 \leq q \leq 5$, which is a common choice in empirical studies of scaling behaviour in asset returns.

From Figure 1 it can be seen that although the estimated scaling functions are strictly concave for all series, suggesting multifractal rather than unifractal distributional scaling, the degree of nonlinearity varies for the different series; it appears to be lowest for the EUR/USD data and highest for the JPY/USD data, with the two equity index series lying somewhere in between. The differences observed between the estimated scaling functions for the pair of exchange rates and the pair of equity indexes suggest that the type of financial asset under consideration may not in isolation provide any strong a priori information regarding the likely distributional scaling properties of the return process. Whether these observed deviations from unifractal scaling behaviour are sufficiently small for the proposed method to perform well is an empirical issue, which will be considered later in the current section during the density forecasting exercise.

4.2 Benchmark Density Forecasting Model

For the empirical exercise it is necessary to have one or more existing density forecasting methods to compare the performance of the new unifractal method against. Whilst there are many possibilities, attention has been restricted to a single established benchmark model using intraday data. This takes the form of the autoregressive realised volatility (AR-RV) model of Andersen et al. (2003), which fits a univariate autoregressive model to the time series of (logarithmic, demeaned) daily realised volatility measures. Density forecasts for daily returns can then be produced by combining these point forecasts of volatility with the empirical observation that daily returns are approximately normally distributed if standardised by their corresponding (time-varying) realised volatilities for each day and their constant sample mean. Following Andersen et al. (2003), a 5th order AR-RV(5) model was used initially and this choice was also found to produce the best average density forecasting performance for the dataset employed here. Finally, the same m -day rolling window estimation and forecasting scheme was used for both benchmark density forecasting methods as described in Section 3 for the unifractal method.

4.3 Method for Density Forecast Comparison

The method used for out-of-sample density forecast comparison is the test of equal predictive ability proposed by Gneiting and Ranjan (2011). The test assumes that two competing forecasting models are used to produce one-step-ahead out-of-sample density forecasts for the variable of interest, y . Consistent with the previous nota-




Figure 1: Estimated unifractal and multifractal scaling functions. Solid lines correspond to estimated scaling functions for the multifractal case (obtained using the partition function estimator) and dashed lines correspond to the unifractal estimates (obtained using the DMA estimator).

tion, it is assumed that N density forecasts are produced by each forecasting method and the forecasts produced by the two models at time t (for use at time $t + 1$) are denoted by $\tilde{f}_t(y)$ and $\tilde{g}_t(y)$, respectively.

The loss function employed by the test is the continuous ranked probability score¹⁰ (CRPS), generalised to allow more importance to be placed on forecast accuracy in particular regions of the density via the use of a weighting function. The value of the weighted CRPS for the forecast produced by the first model for use in period $t + 1$, denoted by $S(\tilde{f}_t, y_{t+1})$, is given by:

$$S(\tilde{f}_t, y_{t+1}) = 2 \int_0^1 \left(\mathbf{I}\{y_{t+1} \leq \tilde{F}_t^{-1}(\alpha)\} - \alpha \right) \left(\tilde{F}_t^{-1}(\alpha) - y_{t+1} \right) w(\alpha) d\alpha \quad (4.1)$$

where $\tilde{F}_t(y)$ is the CDF forecast produced at time t obtained from the PDF forecast for the same period $\tilde{f}_t(y)$, $\mathbf{I}\{\cdot\}$ is an indicator function and $w(\alpha)$ is a weighting function; the authors suggest several possible forms for $w(\alpha)$, which allow more weight to be placed on forecast accuracy different regions of the density, such as the centre or tails.

Whenever a closed form expression for (4.1) is unavailable, it can be approximated easily to any degree of accuracy using the method outlined by Gneiting and Ranjan (2011). The average value of the weighted CRPS in (4.1) can be calculated for each of the two density forecasting models over the N out-of-sample periods (for period $m + 1$ until period T) as:

$$\bar{S}^f = \frac{1}{N} \sum_{t=m}^{T-1} S(\tilde{f}_t, y_{t+1}) \quad \text{and} \quad \bar{S}^g = \frac{1}{N} \sum_{t=m}^{T-1} S(\tilde{g}_t, y_{t+1}) \quad (4.2)$$

A formal test can then be based on the following test statistic:

$$t = \frac{\bar{S}^f - \bar{S}^g}{\hat{\sigma}_n / \sqrt{N}} \quad (4.3)$$

where $\hat{\sigma}_n^2$ is a standard heteroskedasticity and autocorrelation consistent estimator for the asymptotic variance of $\sqrt{N}(\bar{S}^f - \bar{S}^g)$.

Under the null hypothesis that the two density forecasting models have equal predictive ability, the test statistic in (4.3) is asymptotically normally distributed, with the null rejected at the $\alpha\%$ significance level if $|t| > z_{\alpha/2}$, where $z_{\alpha/2}$ is the $(1 - \alpha/2)$ quantile of the standard normal distribution. Given that lower values of the CRPS correspond to better forecasts, in the case of rejection, the forecasting model f should be chosen when the sample value of the test statistic is positive and model g when it is negative.

4.4 Empirical Results

For the empirical density forecasting exercise two different lengths of rolling in-sample window (values of m in the previous notation) were used for parameter estimation: the first is a relatively typical choice of 250 working days and the second is a much shorter period of 50 working days. In principle, models based on intraday data may be able to produce better forecasts for shorter estimation windows than

methods using daily data if they can effectively exploit the additional information contained in the large number of intraday returns available. The shorter 50-day in-sample window length has been included to explore whether this is the case for the proposed unifractal method and the AR-RV benchmark model.

Throughout the following the new unifractal density forecasting approach is referred to as the autoregressive unifractal (AR-UF) model, with the parametric and nonparametric density variants denoted by AR-UF $P(p)$ and AR-UFNP (p) respectively. As with the benchmark AR-RV model, the optimal order of the autoregressive structures for the AR-UF methods was also typically found to be 5. For the AR-UFNP variant employing the nonparametric kernel density estimator, the standard normal kernel was chosen for the kernel function $k(\cdot)$ in equation (2.2) and the simple normal reference rule-of-thumb was used to select the bandwidth parameter h^{11} .

Finally, the values for the minimum and maximum window sizes for the DMA estimator of the self-affinity index that were found to be approximately optimal in this context were $n_{min} = 5$ for all series and $n_{max} = 100$ and 300 for the equity and exchange rate series respectively. These values of n_{max} coincide approximately with the number of 5-minute intraday returns observed during a trading day for each series, suggesting that it is optimal to estimate the scaling behaviour of the process just over the range of sampling intervals that are of direct interest for the current application.

Table 1 contains sample values of the CRPS-based test statistic for equal predictive ability between the new unifractal semiparametric method and the AR-RV benchmark density forecasting model. In addition to the simple unweighted version of the test, several of the weighting functions suggested by Gneiting and Ranjan (2011) have been employed here to place more weight on accuracy in the centre, left and right tails respectively.

From Table 1 it can be seen that in the majority of cases the null of equal predictive ability cannot be rejected, implying that the new AR-UF approach matches the performance of the existing AR-RV method. Comparing the results across the columns of Table 1, there do not appear to be any universal patterns in relative forecasting performance across the regions of the density, although arguably the performance of the AR-UF models is somewhat stronger in the left tail of the density. This suggests that the method should perform well in risk management applications, such as the calculation of Value at Risk or expected shortfall.

The relative performance of the AR-UF models appears to be stronger for the EUR/USD and S&P500 data, as indicated by the larger proportion of negative sample values for the test statistic and larger number of significant values for these two series. The situation is however reversed for the JPY/USD data where the majority of the sample values are positive, though almost always too small to be statistically significant.

Given that the deviation from true unifractal distributional scaling previously noted during Section 4.1 appears to be smallest for the EUR/USD data and largest for the JPY/USD data, this observed variation in forecasting performance across the assets could be attributable to differences in their scaling properties. The large difference in the number of observations available may also be an additional factor, with the intraday exchange rate series containing over three times more observations than the equity index series. Even if the distributional scaling properties of two return series were identical, a greater number of observations should allow more accurate estimation and ultimately result in better density forecasts. Given the

Table 1: Density Forecast Comparison Against AR-RV(5) Benchmark

	None	Centre	Left Tail	Right Tail
EUR/USD: 250-day in-sample period				
AR-NPUF(5) - estimated AR parameters	-1.291	-1.531	-1.221	-1.020
AR-NPUF(5) - fixed AR parameters	-1.753*	-1.816*	-1.485	-1.474
AR-UFP(5)	-1.337	-1.107	-1.051	-1.487
EUR/USD: 50-day in-sample period				
AR-UFNP(5) - estimated AR parameters	-0.452	-0.650	-0.897	0.061
AR-UFNP(5) - fixed AR parameters	-1.058	-0.987	-1.473	-0.498
AR-UFP(5)	-0.738	-0.587	-1.185	-0.165
JPY/USD: 250-day in-sample period				
AR-UFNP(5) - estimated AR parameters	1.690*	1.673*	1.009	1.987**
AR-UFNP(5) - fixed AR parameters	1.208	1.021	0.167	1.891*
AR-UFP(5)	0.175	0.311	0.156	0.155
JPY/USD: 50-day in-sample period				
AR-UFNP(5) - estimated AR parameters	0.971	1.016	0.522	1.181
AR-UFNP(5) - fixed AR parameters	0.540	0.530	0.506	0.422
AR-UFP(5)	-0.068	0.051	-0.080	-0.037
NASDAQ100: 250-day in-sample period				
AR-UFNP(5) - estimated AR parameters	-0.833	0.446	-0.523	-0.817
AR-UFNP(5) - fixed AR parameters	-1.996**	-0.693	-2.111**	-1.166
AR-UFP(5)	1.580	1.557	1.473	1.356
NASDAQ100: 50-day in-sample period				
AR-UFNP(5) - estimated AR parameters	0.012	0.100	0.191	-0.171
AR-UFNP(5) - fixed AR parameters	-1.562	-1.486	-1.381	-1.190
AR-UFP(5)	-0.111	0.091	0.527	-0.740
S&P500: 250-day in-sample period				
AR-UFNP(5) - estimated AR parameters	-0.752	0.295	-1.327	0.206
AR-UFNP(5) - fixed AR parameters	-1.937*	-1.056	-2.295**	-0.576
AR-UFP(5)	1.338	1.521	1.452	0.808
S&P500: 50-day in-sample period				
AR-UFNP(5) - estimated AR parameters	-1.000	-0.948	-1.358	-0.472
AR-UFNP(5) - fixed AR parameters	-1.646*	-1.718*	-2.295**	-0.413
AR-UFP(5)	-0.318	0.099	0.483	-1.197

The test statistic is normally distributed under the null of equal predictive ability and the test statistic is constructed such that significant negative (positive) values imply the new unifractal (AR-RV benchmark) method provides superior density forecasting performance. Significance at the 10%, 5% and 1% levels is indicated by one, two or three asterisks, respectively.

weaker performance of the method for the JPY/USD data, the issue of sample size does appear to be of lesser importance than the scaling properties of the return process, however the density forecasting performance for the S&P500 and NASDAQ-100 series might be closer to that for the EUR/USD data if a larger number of intraday observations were available for estimation.

Finally, comparing the performance of the different variants of the AR-UF method it is clear that the simpler fixed parameter variant of the AR-UFNP model performs better than that in which the parameters are estimated via the KLIC; although the differences are sometimes small, this result is consistent across the results in Table 1. It may seem surprising that the more restrictive variant of the AR-UFNP model can consistently outperform the more flexible specification, but several points should be noted.

Firstly, as previously noted, empirical studies in the literature on density forecast combination often find that a simple average of forecasts performs better than a combination chosen to minimise some statistical loss function. A possible explanation put forward for this finding (see Timmermann, 2006) is the problem of estimation error; in principle a time-varying parameter vector may be able to provide better performance than a simple fixed parameter vector, but if in practice it is not possible to accurately estimate the optimal parameter values then the simple equally weighted version may perform better on average. The problem of density forecast combination is closely related to that of autoregressive parameter estimation in the current context and so similar results may also hold here.

The strong performance of a fixed parameter vector placing equal weight on each lagged density is however harder to justify in the current time series context, where it might be expected that the more recent densities would have larger weights, as is typically the case when fitting standard autoregressive models in a time series context. However, in the current application the dynamic structure is being imposed on complete densities and not single observations, so it is not necessarily true that the same patterns should hold here.

Finally, the loss function used to estimate the autoregressive parameter values is based on the logarithmic score rather than the CRPS employed for forecast comparison and there is no a priori reason to expect the optimal forecasts in a log-score sense to coincide with those in a CRPS sense. It is possible in principle to estimate the autoregressive parameters using the same CRPS loss function, however the substantially greater computational requirements of this approach make it impractical in the current rolling forecasting environment. Faster and less precise discrete approximations to the true CRPS were tested, but the resulting forecasts performed worse than the current method, whilst still being substantially slower to compute.

The forecasting performance of the final AR-UFNP specification that employs a parametric specification for the daily return density is more variable. For the shorter 50-day in-sample period the differences in predictive ability between the AR-UFNP and AR-UFNP specifications are typically small, but for the longer 250-day in-sample period the differences are often more significant; for the JPY/USD data the AR-UFNP model provides a large increase in predictive ability over the two AR-UFNP specifications, but for the two equity index series the situation is reversed.

Compared to the AR-UFNP specification, the AR-UFNP model imposes a more restrictive form for the daily return density for each trading day, but at the same time permits a more flexible dynamic structure by allowing each of the distributional parameters to evolve independently over time, rather than imposing a single

autoregressive structure on the complete density function. The differences in relative forecasting performance between these two specifications across different assets suggest that for some return series allowing for flexibility in the distributional form is more important than that in dynamic structure and for others the converse is true.

The results in Table 1 provide a comparison of average predictive ability over the complete out-of-sample period. However, it is also possible to calculate the CRPS differential for each of the days in the out-of-sample period individually to examine whether the relative forecasting performance of the two methods varies systematically over time. This period-by-period CRPS differential for the AR-RV benchmark is plotted in Figure 2 over the complete 750-day out-of-sample period for the EUR/USD and S&P500 data, together with the daily realised volatility as a proxy for the latent daily return variance. The period-by-period CRPS differentials of Figure 2 are constructed using the unweighted version of the CRPS and the longer 250-day rolling estimation window. As with the values of the test statistics reported in Table 1 above, a negative value of the period-by-period CRPS differential for a given day implies that the unifractal method had superior predictive ability for that day, with larger negative values (in absolute terms) implying a larger improvement over the benchmark AR-RV method.

From Figures 2a and 2c, it is clear that the predictive ability of the new unifractal forecasting method relative to the AR-RV benchmark varies over the length of the out-of-sample period. Perhaps most notably, for the S&P500 data the relative performance of the unifractal method appears to be stronger during the more volatile period in late 2008 and early 2009, corresponding to the most severe part of the recent financial crisis. This graphical observation is also supported numerically, with the average value of the CRPS differential for the 12 month period from the start of Q3 2008 until the start of Q3 2009 being -0.0137, compared to an average of -0.0033 for the 750 day out-of-sample period as a whole.

5 Conclusion

This paper has presented a new method for producing semiparametric density forecasts for daily financial returns using high-frequency intraday data. Through a new application of results from the theory of unifractal processes the intraday returns are appropriately rescaled and the density of daily returns for each trading day is estimated directly from these rescaled high-frequency observations, allowing for the use of both parametric and non-parametric estimators for the daily return density.

In contrast to previous methods using realised volatility measures to estimate and forecast daily return densities, the approach presented here utilises information about both the sign and magnitude of the intraday returns and allows this intraday information to influence aspects of the daily return density beyond the second moment. In addition, the ability to use nonparametric density estimation techniques avoids the potential difficulties encountered in selecting a suitable parametric model for asset returns.

The density forecasting performance of the unifractal method was compared against an autoregressive realised volatility model in an empirical application with 5-minute intraday data on exchange rate and equity index data. The empirical results show that the unifractal method is competitive with the AR-RV benchmark




Figure 2: Period-by-period CRPS differential and daily realised volatility.

and is able to provide statistically significant improvements in density forecasting performance in several situations. In addition, the gains in predictive ability provided by the new unifractal method seem to increase during periods of high return volatility, such as the most severe part of the financial crisis in late 2008 and early 2009.

Directions for future research are numerous, but could include the use of additional criteria for density forecast comparison or evaluation; in particular, the performance of the new semiparametric method in real-world applications such as risk management and asset pricing could be investigated. Alternatively, it could be explored whether the density of daily returns for a given period, conditional on the return observed in the previous period, could be estimated directly from the intraday data in an analogous way by exploiting similar distributional scaling results. Density forecasts for daily returns could then be produced simply by updating the relevant conditioning information in the estimated conditional density function for each day, thus avoiding the need to impose a parametric dynamic specification in order to produce forecasts.

References

- G. Amisano and R. Giacomini. Comparing density forecasts via weighted likelihood ratio tests. *Journal of Business & Economic Statistics*, 25(2):177–190, 2007.
- T. G. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys. The distribution of realized exchange rate volatility. *Journal of the American Statistical Association*, 96(453):42–55, 2001.
- T. G. Andersen, T. Bollerslev, F. X. Diebold, and P. Labys. Modeling and forecasting realized volatility. *Econometrica*, 71(2):579–625, 2003.
- A. Bashan, R. Bartsch, J. Kantelhardt, and S. Havlin. Comparison of detrending methods for fluctuation analysis. *Physica-Section A*, 387(21):5080–5090, 2008.
- C. Bender, T. Sottinen, and E. Valkeila. Arbitrage with fractional brownian motion? *Theory of Stochastic Processes*, 13(29):23–34, 2007.
- L. Calvet and A. Fisher. Multifractality in asset returns: theory and evidence. *Review of Economics and Statistics*, 84(3):381–406, 2002.
- L. Calvet and A. Fisher. How to forecast long-run volatility: regime switching and the estimation of multifractal processes. *Journal of Financial Econometrics*, 2(1): 49–83, 2004.
- M. P. Clements, A. Galvão, and J. Kim. Quantile forecasts of daily exchange rate returns from forecasts of realized volatility. *Journal of Empirical Finance*, 15(4): 729–750, 2008.
- D. Delignieres, S. Ramdani, L. Lemoine, K. Torre, M. Fortes, and G. Ninot. Fractal analyses for 'short' time series: A re-assessment of classical methods. *Journal of Mathematical Psychology*, 50(6):525–544, 2006.
- T. Di Matteo. Multi-scaling in finance. *Quantitative finance*, 7(1):21–36, 2007.
- T. Di Matteo, T. Aste, and M. Dacorogna. Long-term memories of developed and emerging markets: using the scaling analysis to characterize their stage of development. *Journal of Banking & Finance*, 29(4):827–851, 2005.
- R. Dittmar. Nonlinear pricing kernels, kurtosis preference, and evidence from the cross section of equity returns. *Journal of Finance*, 57(1):369–403, 2002.
- R. Engle and S. Manganelli. Caviar: Conditional autoregressive value at risk by regression quantiles. *Journal of Business & Economic Statistics*, 22(4):367–381, 2004.
- P. Giot and S. Laurent. Modelling daily value-at-risk using realized volatility and arch type models. *Journal of Empirical Finance*, 11:379–398, 2004.
- T. Gneiting and R. Ranjan. Comparing density forecasts using threshold- and quantile-weighted scoring rules. *Journal of Business & Economic Statistics*, 29 (3):411–244, 2011.

- C. Harvey and A. Siddique. Conditional skewness in asset pricing tests. *Journal of Finance*, 55(3):1263–1295, 2000.
- J. Kantelhardt. Fractal and multifractal time series. In *Encyclopedia of Complexity and Systems Science*. Springer, 2009.
- Q. Li and J. Racine. *Nonparametric Econometrics: Theory and Practice*. Princeton University Press, 2006.
- J. Maheu and T. McCurdy. Do high-frequency measures of volatility improve forecasts of return distributions? *Journal of Econometrics*, 160(1):69–76, 2010.
- B. Mandelbrot, A. Fisher, and L. Calvet. A multifractal model of asset returns. *SSRN*, 1997.
- K. Matia, Y. Ashkenazy, and H. Stanley. Multifractal properties of price fluctuations of stocks and commodities. *Europhysics Letters*, 61(3):422–428, 2003.
- J. Mielniczuk and P. Wojdyllo. Estimation of hurst exponent revisited. *Computational Statistics & Data Analysis*, 51(9):4510–4525, 2007.
- J. Muzy, J. Delour, and E. Bacry. Modelling fluctuations of financial time series: from cascade process to stochastic volatility model. *The European Physical Journal B-Condensed Matter*, 17:537–548, 2000.
- E. Onali and J. Goddard. Unifractality and multifractality in the italian stock market. *International review of financial analysis*, 18(4):154–163, 2009.
- E. Parzen. On estimation of a probability density function and mode. *The Annals of Mathematical Statistics*, 33(3):1065–1076, 1962.
- F. Selçuk and R. Gençay. Intraday dynamics of stock market returns and volatility. *Physica-Section A*, 367:375–387, 2006.
- A. Timmermann. Forecast combinations. In *Handbook of Economic Forecasting, Volume I*. Elsevier, 2006.
- F. Žikeš. Semiparametric conditional quantile models for financial returns and realized volatility. *Working Paper*, 2009.

Notes

¹Examples include Calvet and Fisher (2002), Matia, Ashkenazy, and Stanley (2003), Di Matteo, Aste, and Dacorogna (2005), Selçuk and Gençay (2006), Di Matteo (2007) and Onali and Goddard (2009).

²These include the multifractal model of asset returns of Mandelbrot, Fisher, and Calvet (1997), the Markov switching multifractal model of Calvet and Fisher (2004) and the multifractal random walk of Muzy et al. (2000) .

³Whilst not directly relevant in the current context, as noted by an anonymous referee this assumption may introduce the possibility of arbitrage opportunities due to the properties of some unifractal processes (see for example Bender et al., 2007).

⁴This weaker condition would allow for the distributional scaling relationship to change or break down at very short or long sampling intervals that are outside the range of interest.

⁵See for example Delignieres et al. (2006), Mielniczuk and Wojdylo (2007) and Bashan et al. (2008).

⁶The issue of optimally selecting the values of n_{min} and n_{max} will be discussed further during the empirical exercise of Section 5.

⁷Further details of kernel estimation for dependent data, together with a brief summary of relevant mixing conditions, can be found in Li and Racine (2006).

⁸This limitation was not encountered when a similar autoregressive structure was applied by Engle and Manganelli (2004) in the literature on conditional quantiles, because the autoregressive coefficients were estimated separately for each of the individual quantiles, allowing them to vary across the distribution.

⁹In the realised volatility literature a 5-minute sampling interval has typically been found to be a good compromise between these two factors (see for example Andersen et al., 2001).

¹⁰The earlier weighted likelihood ratio (WLR) test of Amisano and Giacomini (2007) is similar in spirit, but uses the logarithmic score of Section 3.3. However, it has subsequently been demonstrated that the WLR test is not guaranteed to produce valid inference when a weighting function is used.

¹¹More complex plug-in and cross-validation approaches were also tested, but the former did not produce significant improvements in density forecasting performance and the high computational requirements of the latter made it impractical in the current rolling estimation context.