International Journal of Forecasting 30 (2014) 863–881

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

Forecasting daily return densities from intraday data: A multifractal approach

Mark Hallam [a,*], Jose Olmo [b]

[a] Koç University, Turkey
[b] University of Southampton, United Kingdom

ARTICLE INFO

Keywords:
Density forecasts
Volatility forecasting
Multifractal
Unifractal
Intraday
Finance

ABSTRACT

This paper proposes a new approach for estimating and forecasting the moments and probability density function of daily financial returns from intraday data. This is achieved through a new application of the distributional scaling laws for the class of multifractal processes. Density forecasts from the new multifractal approach are typically found to provide substantial improvements in predictive ability over existing forecasting methods for the EUR/USD exchange rate, and are also competitive with existing methods when forecasting the daily return density of the S&P500 and NASDAQ-100 equity index.

1. Introduction

Over the past decade, there has been a dramatic increase in the availability of intraday financial data, resulting in an extensive body of literature on the use of high-frequency data in financial econometrics. These data obviously allow for the study of financial market behaviours at intraday timescales, but they also contain information that may potentially be of value 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 body of literature on realised volatility, a concept that was introduced by Andersen and Bollerslev (1998) and subsequently formalised by Andersen, Bollerslev, Diebold, and Labys (2001). The daily realised volatility is obtained by summing the squared intraday returns observed during that day, and can then be used as an estimate of the unobservable daily volatility. It has been found (see for example Andersen, Bollerslev, Diebold, & Labys, 2003; Andersen, Bollerslev, & Lange, 1999) that the use of high-frequency data in the form of these realised volatility measures can provide significant improvements in the modelling and forecasting of the daily return volatility compared to models using only daily data.

Whilst the return volatility is undoubtedly a variable which is of substantial academic and practical interest, many situations are encountered 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 and expected shortfall, require a knowledge of particular quantiles of the return distribution. In addition, numerous studies, including those of Brooks, Burke, Heravi, and Persand (2005) and Harvey and Siddique (2000), have shown that higher moments such as skewness and kurtosis are time varying, and there is empirical support for these higher moments being relevant in problems of portfolio allocation and asset pricing (see for example Dittmar, 2002; Harvey & Siddique, 2000, amongst others).

However, as was noted by Žikěs (2009), the use of intraday data for modelling and forecasting the characteristics

* Corresponding author.
E-mail address: mhallam@ku.edu.tr (M. Hallam).

http://dx.doi.org/10.1016/j.ijforecast.2014.01.007
0169-2070/

of daily return distributions beyond the first two moments is not a subject that has yet received much attention. Notable exceptions include Andersen et al. (2003), Clements, Galvão, and Kim (2008), Giot and Laurent (2004) and Maheu and McCurdy (2010), all of which extend the use of realised volatility measures to either quantiles of daily returns or the entire daily return density. The methods used by these previous studies to link the realised volatility measures produced from intraday data to the density (or quantiles) of daily returns 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 the daily volatility. The specific parametric volatility model has varied, with Giot and Laurent (2004) using an ARCH-type model, Clements et al. (2008) considering mixed data sampling (MIDAS) and heterogenous autoregressive (HAR) models, and Andersen et al. (2003) using a bivariate VAR. The second component is a parametric distributional assumption about daily returns, allowing density or quantile forecasts for daily returns to be produced from the point forecasts of the daily realised volatility. This is typically based on the finding of Andersen et al. (2001) that daily returns are approximately normally distributed when standardised by their corresponding daily realised volatilities, although Clements et al. (2008) also explored the use of an empirical distribution estimated from the data.

In response to these theoretical limitations, Hallam and Olmo (2014) propose a method for estimating and forecasting the probability density of daily returns from intraday data, based on a new application of distributional scaling laws for the class of unifractal processes. These processes possess a form of scale invariance, such that the distribution of the process at a given timescale is related to that at any other timescale through a distributional scaling law.

Under the assumption of unifractality, the form of this distributional scaling can be estimated for a given sample of data and it has been demonstrated that these estimates can be used to rescale the intraday returns appropriately such that they are equal in distribution to daily returns; the density of daily returns can then be estimated from these rescaled intraday observations directly. However, it should be noted that the applicability of the proposed approach only requires unifractal distributional scaling to be present locally over each estimation window and for the range of timescales that are of direct interest, rather than globally over all sampling intervals and sub-periods, as would be the case for a true unifractal process in the traditional sense.

In contrast to existing methods, information concerning both the magnitude and the sign of intraday returns can be incorporated into the estimates of the daily return density. Furthermore, this approach also allows the use of nonparametric density estimation methods, thus removing the need to impose a specific parametric form for the density of daily returns. The empirical application of Hallam and Olmo (2014) suggests that the proposed unifractal density forecasting method produces density forecasts that perform well when the true scaling behaviour of the return processes is sufficiently close to that of a unifractal process, even if it is not exactly unifractal. However, it also appears that the predictive ability of the unifractal approach can be affected adversely by larger deviations from the unifractal distributional scaling behaviour that is required in order for the method to be theoretically valid.

The current paper therefore proposes an alternative approach to the production of density forecasts for daily returns from intraday data, based on distributional scaling laws for the more general class of multifractal processes. Compared to unifractal processes, multifractal processes allow for a more flexible scaling relationship between return distributions at different sampling frequencies, overcoming a key theoretical limitation of the previous method. However, whilst the multifractal approach of the current paper permits more flexible distributional scaling behaviours than the earlier unifractal approach, the implementation of the method is more restrictive in some respects, most notably in its requirement of a parametric form being selected for the daily return distribution. Nonetheless, the proposed method still allows the intraday data to influence the properties of the daily return density directly, beyond the second moment. In particular, the approach allows the kurtosis of daily returns to be estimated from the intraday data directly, and incorporated into the forecasts of the density of daily returns.

The aim of the current work is to formalise this alternative multifractal approach and explore whether the additional flexibility which it permits, in terms of distributional scaling behaviours, allows it to produce accurate density forecasts, in spite of the more restrictive implementation it requires compared to the competing unifractal approach. The density forecasting performance of the proposed multifractal approach is compared to that of benchmark models from the GARCH and realised volatility literature, in addition to the unifractal approach of Hallam and Olmo (2014), in an empirical application using a dataset of 5-min intraday equity and exchange rate data.

The structure of the paper is as follows: Section 2 presents the relevant theory on unifractal and multifractal processes and describes how these results can be used to link the properties of the return process at different sampling frequencies. Section 3 then discusses how these concepts can be applied in practice for estimating and forecasting the moments of daily returns in the multifractal case, and ultimately for forecasting the daily return density. Section 4 presents the empirical application of the new multifractal approach, and finally, Section 5 concludes.

2. Unifractal and multifractal processes

In order to estimate the density of daily returns from intraday data, a method for formally linking the characteristics of return distributions across different sampling frequencies is required. In contrast to the existing work in the literature, which has been based on realised volatility measures, the proposed method relies instead on results from the theory of unifractal and multifractal processes.

On an intuitive level, such stochastic processes exhibit some form of *scale invariance*, such that the behaviour of the process observed at one timescale is, after an appropriate transformation, identical, in a statistical sense, to that