A Parsimonious Continuous Time Model of Equity Returns (Inferred from High Frequency Data)

Mascia Bedendo
FORC
University of Warwick
CV4 7AL Coventry U.K.
Tel. +44 2476 522701
Fax. +44 2476 524167
forcmb@wbs.warwick.ac.uk

Stewart D. Hodges
FORC
University of Warwick
CV4 7AL Coventry U.K.
Tel. +44 2476 523606
Fax. +44 2476 524167
forcsh@wbs.warwick.ac.uk

May 11, 2003

Abstract

In this paper we propose a parsimonious continuous time model capable of describing the dynamics of futures equity returns at different time frequencies. Unlike several related works in the literature, we avoid specifying a model *a priori* and we attempt, instead, to infer our model from the analysis of a data set of 5-minute returns on the S&P500 futures contract. Throughout the entire paper we try to keep the modelling assumptions to a minimum (and to test them step by step), while retaining an adequate level of structure. We start with a very general specification for our model for futures equity returns. First we model the seasonal pattern in intraday volatility, which turns out to be deterministic and stationary through time. Once we correct for this component, we aggregate intraday data into a daily volatility measure to reduce the amount of noise in the data and its distorting impact on the results. We then employ this much less noisy daily measure to infer the structure of the stochastic volatility model and of the leverage component, as well as to obtain insights on the shape of the distribution of conditional returns. Our model is then refined at a high frequency level by means of a simple non-linear filtering technique which provides an intraday update of volatility and return density estimates on the basis of observed 5-minute returns. This method allows to capture all the information embedded in the intraday data which was lost in the daily aggregation. The results from a Monte Carlo experiment indicate that a sample of returns simulated according to our model well replicates the main features observed in market returns.

JEL Classification: C11, C51, C52, C53, G12.

1 Introduction

The increasing availability of high frequency data in finance has improved the empirical analysis of financial asset returns in several respects. In the first place, it has enabled the investigation of the dynamics of intraday volatility and returns *per se*, in consideration of the various market microstructure effects that characterise high frequency financial data. Secondly, and perhaps more importantly, it has enriched the information set available to develop and test continuous time models which are able to explain and replicate the dynamics of financial returns observed in the markets in a consistent manner across different time horizons. Traditionally, continuous time models in finance have been estimated and tested on moderate frequency (normally daily) financial data. In most cases, however, the asset returns generated from those models manage to capture the dynamics of daily or weekly returns fairly accurately, but fail to mirror the behaviour of high frequency financial returns. Therefore, intraday data can be usefully employed to derive a more consistent specification for a continuous time model.

The present work fits in this latter stream of literature, since its aim is to identify the simplest possible model which is both congruent with the specifications commonly adopted in this field and capable of replicating the essential features that characterise the evolution of intraday returns and volatility as observed in futures equity markets. A continuous time specification turns out to be the most convenient and appropriate one for such purpose.

A distinctive aspect of our study, which we consider a significant contribution to the related literature, is that we adopt a parsimonious approach and, throughout the different steps, we let the data suggest the model as much as possible, rather than imposing a model ourselves. The standard approach commonly followed by the literature consists of assuming from the beginning a particular specification for the model in all its components and using the data to estimate and test it. Instead we believe that a model for financial data should originate from the data itself, therefore here we avoid specifying a model *a priori*. We start with a very general model structure and we perform a careful step by step analysis of the data, recording the relevant features to be modelled, whose peculiar characteristics will actually drive the choice among different specifications. At each step we also look carefully for possible specification errors. Throughout the entire paper we try to keep the modelling assumptions to a minimum, while retaining an adequate level of structure. Our approach is also parsimonious in terms of the statistical and econometric techniques employed to estimate the resulting model. Our main interest here is in assessing whether the data-driven, step-by-step criteria we propose for selecting the model and subsequently refining it on the basis of intraday returns enables us to derive a valid specification that adequately explains the empirical features. Producing the most precise estimates for the parameters of our model is not our main concern, especially since this would require the implementation of very complicated econometric tools that would introduce a lot of complexity to the analysis without contributing significantly to the main results. We use therefore simple techniques that still produce reasonably accurate estimates.

At the conclusion of our analysis we propose a relatively simple specification, able to capture and model most of the aspects observed in equity index futures markets, namely: seasonality in intraday data, stochastic volatility and the presence of jumps, and a leverage effect. By means of a simple Bayesian filtering technique we also generate 5-minutes ahead volatility estimates and density estimates for the distribution of the intraday returns, whose accuracy is thoroughly assessed via both point and distributional forecast tests.

The paper is structured as follows. Section 2 introduces the related literature. Section 3 describes the data set. Section 4 details in its subsections the various steps of the data analysis and the modelling of each component, up to the identification and estimation of a simple, but accurate, model in continuous time. A Monte Carlo simulation exercise of the complete model is performed in Section 5. Section 6 summarises the main conclusions and a few suggestions for further research.

2 The Informative Content of High Frequency Data

During the last few years, the availability of high frequency data on financial assets has stimulated the production of a very rich literature. One stream of literature, which is not immediately related to the present work and, therefore, which is not explored in detail here, has focused on deriving tailored models for intraday returns and volatility. These models should be capable of capturing some distinctive features such as the significant serial correlation in returns induced by market microstructure effects, and the discreteness of transaction prices.[1]

A second stream of literature exploits the informative content of intraday data to obtain more accurate measures of the volatility of financial returns. Most of these studies approximate the volatility over a certain period, such as a day, with the sum of intraday squared, or absolute, returns, a measure called realised volatility.[2] The theoretical justification for this approximation[3] is to be found in the theory of quadratic variation (see

[1]Models for the dynamics of transaction prices have been suggested, amongst the others, by Rydberg and Shephard [2000, 2003], Rogers and Zane [1998] and Giot [2000].

[2]See Taylor and Xu [1997] and Andersen and Bollerslev [1998] among the authors who first proposed this measure.

[3]See Andersen, Bollerslev, Diebold and Labys [2001, 2003], Barndorff-Nielsen and Shephard [2001a].

Karatzas and Shreve [1988]). A complete asymptotic theory of the convergence of the realised volatility to the integrated volatility has been derived by Barndorff-Nielsen and Shephard [2001b, 2002, 2003], under the assumptions that conditional returns are normally distributed and volatility follows either a diffusion specification or a Lévy process. They have also considered extensions to account for the presence of a leverage effect. Bai, Russell and Tiao [2001], Andreou and Ghysels [2002] and Meddahi [2002] discuss potential distorsions and biases in the realised volatility measure.

An impressive number of papers have appeared in the last couple of years in this area, proposing various possible applications for the informative content of intraday data via the realised volatility measure.[4] Andersen and Bollerslev [1998] and Blair, Poon and Taylor [2001] employed realised volatility as a measure against which to compare daily volatility forecasts produced with a GARCH model. Some authors (Andersen, Bollerslev and Lange [1999], Blair, Poon and Taylor [2001], Martens [2001], Hol and Koopman [2002]) investigated whether out-of-sample volatility forecasts could be improved by using intraday data.

The distribution and the time series properties of the realised volatility have also been studied. Examples in this context are given by Andersen, Bollerslev, Diebold and Labys [2001, 2003] for exchange rates, Ebens [1999] for the Dow Jones Industrial Average index, Andersen, Bollerslev, Diebold and Ebens [2001] for both the index and its constituent stocks, Areal and Taylor [2002] for FTSE-100 index futures. Barndorff-Nielsen and Shephard [2002] fit a continuous time superposition of Ornstein-Uhlenbeck processes to the time series of realised volatility.

A related research area which has attracted academic interest in the last few years employs intraday data in order to estimate and test continuous time models in which financial returns are described by a time-changed Brownian motion or Lévy process where the stochastic time change is given by a measure of the intraday economic activity (e.g. trading volumes, proxy of integrated stochastic volatility). The theoretical justification for such an approach is that all arbitrage-free processes defining asset returns can be represented as time-changed Brownian motions,[5] where the time change (or business time) must account for information arrival and market activity. This stream of literature originates from the pioneering paper by Clark [1973], who showed how, once re-specified in the new business time (expressed in terms of the cumulative volume of activity), financial returns are virtually distributed according to a Gaussian law. Amongst the most relevant contributions in this field, we recall Andersen [1996] who investigates

[4] For a survey of this literature, the interested reader can consult Andersen, Bollerslev and Diebold [2001], Barndorff-Nielsen, Nicolato and Shephard [2002] and Dacorogna *et al.* [2001].

[5] See Monroe [1978] for the proof that any semimartingale is a time-changed Brownian motion.

the returns specification when the intraday information flow is modeled as a stochastic volatility process, and Ané and Geman [2000] who extend Clark's results by using the cumulative number of trades as a stochastic clock. Very recent studies include Geman, Madan and Yor [2001], Carr, Geman, Madan and Yor [2003], Carr and Wu [2003] where various kinds of Lévy processes are evaluated at a time change given by the integrated volatility, generally modelled as a square root process.

3 The Data Set

Our data set consists of 5-minute frequency intraday prices on the S&P500 stock index futures contract from September 15, 1997, to July 26, 2001. All prices are for the futures contract closest to maturity, except for the days within one week to expiration, when the next contract is considered, in order to always refer to the contract with the highest trading volume. Days that recorded transactions only for part of the entire trading day[6] have been excluded from the data set. We have also eliminated four days which exhibited very large returns on some intraday interval immediately followed by equally large returns of the opposite sign, which could be indicative of mistakes in recording the price. Some other days were originally missing from the data set. All in all, our final sample consists of intraday prices for 960 days.

The full trading day in the futures market at the Chicago Mercantile Exchange starts at 8:30 a.m. and ends at 3:15 p.m. Chicago time. Intraday log returns have been computed on the consecutive logarithmic closing prices for each of the 81 5-minute intervals that constitute a trading day. Since in modelling the intraday dynamics of returns and volatility it is important to take into account the close-to-open returns and their volatility, we also analyse overnight log returns, calculated as the difference between the logarithm of the open price and the logarithm of the closing price for the previous day.[7]

In the top half of Table 1 we report some sample statistics for the 5-minute and the overnight returns, which we consider separately, given the different nature and characteristics of the two series. The intraday returns display an almost zero sample mean, a sample standard deviation of 0.121%, positive sample skewness of 0.88 and a strong sample excess kurtosis of 35.48. As expected, the standard deviation of the overnight returns is considerably larger, as it refers to a longer temporal horizon. The higher moments are closer to those of a normal distribution, with a sample skewness of -0.38

[6]Most of those days are those immediately preceeding holidays, such as December 24 and December 31.

[7]For the same reason, unlike other works on high frequency data (see, for example, Andersen and Bollerslev [1997]), we retain the return on the first interval of the trading day, which mainly reflects the information accumulated overnight and shows a high volatility.

and an excess kurtosis of 3.04, by effect of the aggregation process which takes place over a longer time horizon.

Table 2 displays the values of first order autocorrelation coefficients in the series of high frequency returns for each year under analysis, together with the percentage bid-ask spread, estimated following Roll [1984].[8] Although statistically significant for the first two and a half years, the serial correlation in the intraday futures returns always seems to be economically negligible. To ascertain that, we have computed the bias in the variance induced by ignoring first and second order serial correlation, which turns out to be, respectively, -0.318% and -0.319% of the correct variance.[9] Our results suggest that it does not make any substantial difference if we remove the autocorrelation from our series or not. Therefore, here we do not need to worry about market microstructure issues such as the bid-ask bounce, which would bring a strong negative serial dependence and complicate the analysis further, by introducing a serious bias in the volatility measures.

4 Data Analysis and Derivation of the Model

In the present section we perform a careful step by step analysis of the data, aimed at isolating its main distinctive features and their nature, and, therefore, at providing directions for plausible model specifications.

We start by postulating a very general structure for our model of the dynamics of intraday returns, represented as follows:

$$r_{it} = s_{it}\sigma_{it}\varepsilon_{it} \tag{1}$$

for $i = 1, \ldots, 82$, $t = 1, \ldots, T$, where r_{it} represents the unconditional intraday (or overnight) log return for interval i at day t, de-meaned by the sample mean on the corresponding sub-interval; s_{it} identifies the volatility for sub-interval i at day t attributable to the seasonal pattern in intraday volatility; σ_{it} stands for the stochastic volatility component, independent of the seasonal component;[10] ε_{it} symbolises the conditional intraday log return, with zero mean, independent of both the seasonal and the stochastic volatility parts. Once all the seasonal and stochastic volatility components have been correctly modeled, this latter component should translate into a

[8] Roll simply defines a measure for the bid-ask spread in percentage of the geometric average of the average bid and ask prices as: $s_r \equiv -2\sqrt{|\operatorname{Cov}[R_{t-1}, R_t]|}$.

[9] Our measure for the bias in the variance has been obtained by comparing the variance of the returns with the variance of the residuals resulting from fitting, respectively, an AR(1) and an AR(2) process to the high frequency returns.

[10] The choice of such a specification for the volatility seemed natural since the empirical evidence indicates that both a periodic pattern and a stochastic volatility component exist in intraday volatility and need to be modelled.

series which is independent[11] across the intraday intervals.[12]

An important consideration needs to be made here. The structure in (1) is indeed a very rich specification which admits an infinite variety of models as special cases. Nevertheless, even this very general assumptions may easily be too strong, e.g. the distribution of ε_{it} could well depend on s_{it} or σ_{it}, or the stochastic volatility σ_{it} could depend on s_{it}. However, in this work we do not take our assumptions for granted, but we will test their validity as part of our data analysis. Another issue that could be raised about the model in (1) is its possible misspecification in the presence of a significant price clustering. The evidence of a small number of ticks per 5-minute volatility in our sample, ranging from 11 in the quietest moments of the day to 16 in the busiest time, suggests the absence of a strong price clustering.

A note on the terminology that will be used in the present work and on the scaling of the model in (1). Amongst the possible alternatives, after verifying the deterministic nature of the seasonal volatility component, we have chosen to work with de-seasonalised returns which translates into scaling our model so that $E[\sigma_{it}] = 1$ and $E[|\varepsilon_{it}|] = 1$.[13] The use we make here of the term volatility to denote σ_{it} is therefore quite unorthodox and, to be precise, we should refer to σ_{it} as "relative volatility".

In the following subsections we proceed to investigate the features of the data that characterise the nature of the components in Equation (1) and then, to propose and test accurate model specifications for each component.

4.1 The seasonal component

Fig. 1 displays the intraday patterns in average returns (plotted with 95% confidence intervals around zero) and average absolute returns, computed across the time series of the single 5-minute intervals.[14] No clear predictable pattern is discernible for the average returns, whereas an obvious U-pattern characterises the intraday volatility.[15]

[11] The assumption of independence is, in fact, justified by the absence of a significant autocorrelation structure, as ascertained in the previous section.

[12] Here we do not attempt to model the risk premium given that four years of data would not be a sufficiently long time span to obtain reliable estimates for such purpose.

[13] Throughout the paper, we prefer to use absolute, rather than squared, returns, to measure volatility. As largely documented in the existing literature (see, for an exhaustive discussion, Barndorff-Nielsen and Shephard [2003]), absolute returns are less sensitive to large outliers and more reliable when the fourth moment of the distribution of returns is not finite.

[14] Overnight returns have not been included in the plots. Their average and average absolute values of, respectively, 0.02% and 0.41% are clearly not in line with the rest of the intraday data and their inclusion would have distorted the analysis.

[15] The presence of a U-pattern in intraday volatility of stock returns was first documented by Wood, McInish and Ord [1985] and Harris [1986], whereas the impact of this seasonal component on the volatility dynamics was first investigated by Andersen and Bollerslev [1997].

The average absolute returns start out at about 0.1% at the market opening, keep on increasing for the first half an hour up to 0.12%, decline smoothly to the lowest level of 0.062% before noon and then increase again to 0.113% at the closure of the cash index market. The alternance of drop and rise in the last 15 minutes of the trading day is attributable to the post cash market trading.

In principle, the model in (1) allows for a seasonal volatility component that changes through time t. To test for the stability of the seasonal pattern across different moments, we have analysed the shape of the intraday periodicity in volatility on subsamples computed on single days of the week (Fig. 2), on the 50% highest and the 50% lowest volatility days and on the first and the second half of the entire sample period (Fig. 3).[16] The visual inspection of the plots reveals the presence of very similar and almost indistinguishable patterns in intraday volatility. We have also conducted more formal tests of equality between intraday volatility patterns for, respectively, high and low volatility days, first and second half of the sample, each trading day of the week and the overall sample. We first performed, for each intraday interval, a two sample Student's t-test for mean equality (at 95% confidence level) on the average normalised absolute returns of the two subsamples we wanted to compare. The percentage of sub-intervals on which the null hypothesis of equal means is rejected (on the total of 82 intervals) is displayed in the second column of Table 3. We then derived the series of intraday ratios computed on the average normalised absolute returns of the two subsamples of interest; average values and standard deviations for these series have also been reported in Table 3. Both the small percentages of rejections for the mean equality test (ranging from 0 to 12%) and the little dispersion of the ratios of intraday volatility coefficients around the average level of one (with values for the standard deviation between 0.055 and 0.075) seem to support the stability of the seasonal pattern. Therefore, our results suggests that over the time period spanned by our data we can safely assume constant deterministic intraday seasonal pattern, which can then be appropriately represented by s_i.[17]

Since the intraday periodicity in the return volatility has a strong impact on the dynamic properties of intraday returns, it is essential to correct for this component in order to reveal and model the stochastic volatility dynamics present in our series. The average absolute returns for the individual sub-intervals constitute simple and accurate estimates of the in-

[16] The comparison amongst subsets of data with different levels of volatility has been made possible by using absolute returns normalised by the average of absolute returns across the day, taken as a volatility proxy for the day.

[17] There is a chance that different seasonal patterns arise during daylight saving periods, which we have not investigated. However, we expect the impact of a correction in this direction to be much less important for equity data than it would be for exchange rate data.

traday seasonal component in volatility s_i, both for the 5-minute intervals and for the overnight returns.[18] Different approaches have been proposed in the literature to obtain smoothed estimates of these seasonal coefficients. In a recent comparative paper, Martens, Chang and Taylor [2002] recommend the Flexible Fourier Functions approach, first suggested by Andersen and Bollerslev [1997]. Following this technique, smoothed estimates of the seasonal components (only for the 5-minute subperiods) are produced by fitting Flexible Fourier Functions (which consist of combinations of linear, quadratic, trigonometric functions and dummy variables) to multiplicative coefficients based on the average absolute returns. In Fig. 4 we show how good smoothed estimates can be obtained much more easily by fitting a set of cubic B-splines to the average absolute returns for the 5-minute intervals.[19] The cubic B-splines present several advantages over the FFF approach: they are easier to fit, more flexible, more general and they do not rely on *ad hoc* specifications for the inclusion of dummy variables.

Having obtained fairly accurate smoothed estimates for the deterministic intraday seasonal pattern in volatility,[20] we proceed to derive the time series of de-seasonalised unconditional intraday returns, obtained by dividing the de-meaned unconditional returns r_{it} by the corresponding estimate of periodicity in volatility s_i. In the same way we compute the time series of de-seasonalised unconditional overnight returns. Sample higher moments for the series of de-seasonalised unconditional returns (both 5-minute and overnight) are displayed in the bottom half of Table 1.

4.2 The stochastic volatility component

4.2.1 Methodology

Once we have adjusted for the intraday periodicity in volatility, the model in (1) translates into a mixture process, such that each de-meaned and de-seasonalised intraday return is a combination of independent realisations from a stochastic volatility process and from a conditional density. Therefore, the next step to take in order to produce sensible forecasts for high frequency returns is to identify an appropriate stochastic volatility process capable of generating good intraday volatility estimates.

[18] Taylor and Xu [1997] proposed a similar adjustment for seasonality, based on averages of squared returns.

[19] Once again, the overnight period has been excluded from the analysis. As estimate of the overnight seasonal volatility component we use the average absolute overnight returns. Also note the presence of a spike in our cubic B-splines curve, due to a knot placed to capture the drop-and-rise movement typical of the futures contract.

[20] To ensure that no relevant information has been lost as an effect of the smoothing procedure, we have also computed de-seasonalised high frequency returns on the average intraday absolute returns, and re-estimated our model on this series. The changes in the results turned out to be negligible.

Before providing the details of the analysis, let us briefly discuss the intuition and the motivation that drove us through the following steps. Our target is to model stochastic volatility at a high frequency level, but the data available clearly does not enable us to investigate accurately how volatility changes at that frequency, since a 5-minute return gives an extremely noisy measure of intraday volatility. However, if the stochastic volatility does not move too much within a single day, we can easily derive a very accurate daily proxy for the volatility from the average of the 5-minute absolute returns and employ that to analyse the evolution of volatility on a daily basis.

Obviously the estimated model will not be appropriate for producing good intraday volatility estimates (for which it will have to be re-estimated at a higher frequency level); instead, it will provide us with good initial estimates of daily volatility, which reveal the structure that our stochastic volatility model should have in order to capture the essential characteristics of the volatility changes that need to be modelled, including the component attributable to the leverage effect. If volatility were constant during the day, these estimates would also be accurate at an intraday level. We know that this is not the case, therefore at a high frequency level such estimates will show inaccuracies due to both some measurement error and intraday changes in the volatility. However, our daily estimates can still be considered good enough for the purpose of conditioning the series of unconditional intraday returns upon them and consequently, obtaining more precise information on the shape of the conditional distribution, which will turn out to be essential in order to re-estimate the stochastic volatility model at an intraday level and to produce better 5-minute volatility estimates.

4.2.2 The analysis of daily volatility estimates

To investigate the presence and the nature of the stochastic volatility component, we start by plotting the autocorrelogram of the absolute de-seasonalised 5-minute[21] returns for $4,100$ lags, corresponding to 50 days (Fig. 5, top half). The highly significant serial correlation in absolute intraday returns over many lags reveals an important stochastic component in volatility. The slow decay of the autocorrelation coefficients through time indicates the persistence of such component.

However, any single 5-minute absolute return obviously gives a very poor estimate of volatility, as confirmed by the strongly irregular pattern of the ACF. In presence of a large, slow-decaying, noisy component at a 5-minute level, we can best examine the volatility dynamics by taking the daily average of absolute intraday unconditional returns as our volatility proxy, as

[21] In our notation, 5-minute should not be interpreted literally since, in general, it also refers to the overnight interval.

follows:

$$\hat{\sigma_t} = \frac{1}{m}\sum_{i=1}^{m}\left|\frac{r_{it}}{s_i}\right| \tag{2}$$

The measure in (2) directly relates to the realised volatility measures mentioned earlier, which have been recently recommended by several authors as very accurate estimates of the latent daily volatility.

The autocorrelogram for the daily average of absolute high frequency returns up to lag 50 is displayed at the bottom of Fig. 5 to provide a comparison with the one at the top and therefore to ensure that our measure for the volatility at a daily level reproduces the basic characteristics displayed by the volatility estimates at an intraday level (hopefully highlighted by the reduction in the noise). As expected, the elimination of most of the noise produces an overall increase in the level of serial correlation for the daily volatility, which is around four times as much as the intraday level. Also the averaging process has the obvious effect of drastically reducing the very high autocorrelation recorded in intraday volatility for the first 150-200 lags. The visual inspection of the serial correlation in absolute returns at both daily and 5-minute level reveals that the stochastic volatility factor seems to be the result of two components: a) a fast mean reverting component; b) a more persistent component, which appears to decline almost linearly in time.[22]

As we will explain more in detail later, the statistical techniques that we use to estimate a stochastic volatility model provide more reliable and unbiased results the closer the series to be modelled is to a normal. Therefore, we have chosen to work with the time series of the logarithm of the daily volatility proxy $\ln(\hat{\sigma_t})$, whose skewness of 0.40 and excess kurtosis of 0.31 are much closer to the corresponding moments of a Gaussian than those of the volatility proxy itself (equal to, respectively, 1.94 and 6.63).

In the following subsections, we first investigate the impact of the leverage effect and suggest a model to account for the changes in volatility induced by this component and then we explore a way of modelling the (daily) dynamics of the volatility in order to capture the features described above.

4.2.3 The leverage effect

The leverage effect was first discussed by Black [1976] who observed that the amplitude of the volatility of a stock tends to increase when its price drops. However, a direct comparison between volatility and stock prices is not possible, since the first series is stationary and the second one is not. Therefore, in order to investigate presence and magnitude of the leverage

[22] It is worth noticing that the ACF is informative when the underlying model is linear and the variables are Gaussian, which is not the case in our context. When we deal with non-linearities in the model and heavy tailed distributions, the ACF might suggest the presence of spurious long memory effects (see Davis and Mikosch [2000]).

effect, we construct a new, stationary, variable which measures how distant the current price is from its average level.

A very simple specification, which actually seems to be well supported by our data, consists in computing an exponentially weighted moving average M of daily closing log prices for the S&P500 stock index futures as $M_t = (1-\theta)M_{t-1} + \theta\ln(S_{t-1})$, for $t = 1, \ldots, T$ and in deriving the new stationary series as $\ln(S_t) - M_t$. A measure of the leverage effect is then given by the correlation between $\ln(S_t) - M_t$ and the log volatility proxy for the following day $\ln(\hat{\sigma}_{t+1})$. The initial value M_0 is set equal to the initial log price and we choose $\theta = 0.03$ (corresponding to a half life of 23 days), which is the value that maximises (in absolute terms) the correlation between log price movements and log volatility series. For this parametrisation, we obtain a correlation of $\rho = -0.545$ between the two series, whose scatter and time series plots are shown in Fig. 6. Our findings indicate a strong leverage effect whose impact on the volatility dynamics needs to be adequately modelled.

In order to separate the changes in volatility induced by the leverage effect from those arising from the dynamics of the stochastic volatility component we propose the following specification:

$$\ln(\hat{\sigma_t}) = \kappa(\ln(S_{t-1}) - M_{t-1}) + \upsilon_t \tag{3}$$

The regression in (3) provides us with an estimate for κ of -4.34 (standard error 0.26) and with time series of the residuals υ_t, whose evolution should thus mirror the dynamics of the (ex-leverage) stochastic volatility. Given that the ACF inspection carried out in the previous section suggests the presence of both a transient and a more permanent component in the volatility process and that the leverage effect turns out to be quite persistent, we start by assessing whether the volatility expressed by the residuals could be adequately modelled by means of an AR(1) specification.[23] Unfortunately, this simple and appealing specification is immediately ruled out, as indicated by the ACF of the residuals from the AR(1) process (Fig. 7, top), which clearly highlights the existence of a more persistent dynamics ignored by our model.

More complete specifications capable of taking into account this feature are then needed in order to achieve a satisfactory model for the stochastic volatility component.

4.2.4 A short memory model

Given the slow, almost hyperbolic, decay in the sample autocorrelogram for the stochastic volatility, which seems to suggests the presence of long memory effects, we have attempted to model the volatility component by

[23]The continuous time equivalent of an AR(1) model is a standard Ornstein-Uhlenbeck process.

means of long memory ARFIMA(p,d,q) processes of different kinds. Quite surprisingly, none of the specifications chosen is supported by our data set. Gallant, Hsu and Tauchen [1999] provided an alternative explaination for such a slowly decaying dynamics by showing how, for an appropriate choice of parameters, the sum of two AR(1) processes also exhibits long memory features. Modelling the stochastic volatility as a sum of two AR(1) or equivalently, in continuous time framework, with a superposition of Ornstein-Uhlenbeck processes would consent to accurately describe the empirical results, while maintaining the nice properties of a short memory process.

We therefore explore the use of a model similar to Alizadeh, Brandt and Diebold [2002] and represent the log volatility in continuous time as:

$$\ln(\sigma_t) = \ln(\sigma_{s,t}) + \ln(\sigma_{l,t}) \tag{4}$$

with:

$$\begin{aligned} d\ln(\sigma_{s,t}) &= -\alpha_s dt \ln(\sigma_{s,t}) + \beta_s \sqrt{dt} dW_{s,t} \\ d\ln(\sigma_{l,t}) &= -\alpha_l dt \ln(\sigma_{l,t}) + \beta_l \sqrt{dt} dW_{l,t} \end{aligned}$$

In this parametrization, the log volatility $\ln(\sigma_t)$, which is our latent state variable, evolves like a sum of two independent Ornstein-Uhlenbeck processes, each of them mean reverting towards the long run level of zero,[24] with mean reversion parameters α_j. Since we estimate the model on daily variables, $dt = 1$.[25] Once we discretize the model in (4), the two components of the log volatility follow a Gaussian first-order autoregressive process with mean zero, autoregressive parameter $\rho_j = 1 - \alpha_j$ and variance β_j^2.

Since we choose as volatility proxy the residuals from Equation (3), the model we estimate in discrete time to describe the dynamics of the log volatility is:

$$\upsilon = \ln(\sigma_t) + \xi_t \tag{5}$$

$$\begin{aligned} \ln(\sigma_t) &= \ln(\sigma_{s,t}) + \ln(\sigma_{l,t}) \\ \ln(\sigma_{s,t}) &= \rho_s \ln(\sigma_{s,t-1}) + \beta_s \omega_{s,t} \\ \ln(\sigma_{l,t}) &= \rho_l \ln(\sigma_{l,t-1}) + \beta_l \omega_{l,t} \end{aligned}$$

The estimation has been carried out by applying a Kalman filter algorithm to the state space system in (5). If the measurement equation errors ξ_t were normally distributed, we could obtain exact maximum likelihood estimates

[24] In our multiplicative model, the expected value for the stochastic volatility is one, therefore the level at which the log volatility must revert is zero. In fact, statistical estimates of the long run mean turned out to be not significantly different from zero.

[25] The empirical issue of the choice of dt at intraday level, in view of the overnight market closure, will be discussed later on in the paper.

of the model and the linear projections produced by the Kalman filter procedure would represent conditional expectation. As we will discuss later, in our case the measurement errors turn out not to be Gaussian. However, quasi-maximum likelihood procedures can still produce consistent estimates of the parameters.[26]

The estimated parameters, with standard errors in brackets, are displayed in Table 4. We can clearly distinguish between a transient volatility component, with $\alpha = 0.734$ corresponding to a half life of 0.94 days and a permanent one, with $\alpha = 0.018$ and half life of approximately 37.5 days. Most of the short-run variance of the model can be attributed to the transient component, as the values of the β coefficients suggest, whereas 52% of the unconditional long-run variance is explained by the more persistent component. The ACF of the residuals from the two factor AR(1) specification (Fig. 7, bottom) reveals how all the dynamics of the volatility has now been correctly captured.

The distributions of the residuals from both the state equations of the two components and the measurement equation have been analysed (Fig. 8). All the residual series exhibit positive skewness and excess kurtosis which lead to a rejection of their normality. However, for all the distributions, skewness and fat tailness are not too pronounced and this "approximate" Gaussianity should ensure a reasonable efficiency of both the Gaussian quasi-maximum likelihood estimates and the consequent inferences about the latent volatility process.[27]

4.2.5 Some insights on the conditional return densities

As mentioned earlier, the estimation of the stochastic volatility model on a daily basis provides us with both a structure for the dynamics of the stochastic volatility component, and estimates of the (log) volatility level, adjusted daily according to the new value for the log volatility proxy $\ln(\hat{\sigma})$. Such estimates will be fairly accurate at a daily level, but unsatisfactory at a 5-minute level, given that a daily update is equivalent to assuming that

[26] The model has also been estimated by maximising the spectral log-likelihood function for the sum of two AR(1) processes. The two estimation methods produce essentially the same results.

[27] The variance of the measurement error associated with our log volatility proxy should not be very large. In fact the distribution of the residuals from the measurement equation includes both the noise component and the sampling variation from the conditional distribution of the log volatility proxy, which in practice are very difficult to separate. However, the variance of the error term can be used as an upper bound to the percentage of the total variance attributable to measurement error. In our example, it amounts to 0.052, which is the 38.51% of the total variance of the log volatility measure. A lower bound on the variance explained by measurement error is obtained by calculating what the variance of the log volatility proxy from conditional returns would be if the conditional distribution of the returns was normal. In our case it is equal to 0.0072, which corresponds to the 5.30% of the total variance for the log volatility proxy on the unconditional returns.

the intraday volatility estimates for the 82 subintervals of a same day are all identical.[28]

Since it is realistic not to expect large changes in the volatility to occur within a single day, these constant intraday volatility estimates can be usefully employed to extract information on the distribution of conditional returns, as a necessary preliminary step to perform in view of refining the estimates of our model at a 5-minute level. The time series of conditional returns is obtained by normalising the unconditional de-meaned, de-seasonalised return, r_{it}/s_i, by the volatility estimate for day t made at the end of the previous day.[29]

If the volatility dynamics was accurately modelled and the conditional return distribution was independent from the volatility process, then conditional intraday returns should be identically distributed across all intervals of the day and no changes in the shape of their density (i.e. more fat-tailed in intervals of higher activity and less fat-tailed when there are less transactions on the market) should be discernible.[30]

In order to empirically assess such hypotheses, we start by computing summary sample statistics of the time series of conditional returns for each of the 82 intraday intervals. We plot in Fig. 9 (top) the standard deviation of the time series of the conditional returns for the individual intervals. We can clearly detect a few spikes for some intraday intervals that seem to suggest the fat-tailed nature of the conditional distribution and its variability across subintervals. However, looking more carefully at our data, we can see that the spikes are mainly attributable to a very small number of outliers (around 15 for the whole dataset, i.e. less than 0.020% of the total observations) that distort the tails of the distributions over some intervals (not necessarily the busiest ones). Disregarding the spikes, the standard deviation of the conditional distribution of intraday returns turns out to be surprisingly flat across the different subintervals of the day. It is also quite close to the value of $\sqrt{\pi/2}$ (the straight horizontal line in the plot), which represents the theoretical level of standard deviation under the assumption of normality for the distribution of conditional returns.[31]

[28] A rigorous evaluation of both the volatility forecasts and the resulting density forecasts for the returns is postponed to one of the following sections, where a comparison will be drawn with the improved forecasts from high frequency updating.

[29] A remark is appropriate here: in order to avoid tautological results, when deriving the series of conditional returns we always condition on the information available before the actual return was observed. In other words, the return r_{it} observed on interval i of day t will be conditioned on the volatility estimate produced the day before (for daily-updated volatility forecasts) or the 5-minutes before (for 5-minute updated volatility forecasts).

[30] Similarly, given our choice of scaling, if we refer to the distribution of absolute conditional intraday returns, the following properties should hold: $E[|\varepsilon_{it}|] = 1$ and $Var[|\varepsilon_{it}|]$ constant for $i = 1, \ldots, m$.

[31] Under the assumption of normally distributed conditional returns, their variance must be equal to $\pi/2$ in order to satisfy the condition $E[|\varepsilon_{it}|] = 1$.

However, the fact that the volatility of the empirical distribution is persistently higher than the theoretical normal one seems to suggest that the conditional distribution is more fat-tailed than a Gaussian. The plot showing the average excess kurtosis of the conditional returns for the various intraday periods is reported in Fig. 9 (bottom). Again, we observe a high level of excess kurtosis over some intervals of the day, which is mainly due to the presence of few sporadic outliers, as discussed above, rather than to the effect of some external source of information not captured by our model. If we removed these outliers, the excess kurtosis for the overall conditional distributions would be around 2, pushing the distributions much closer to a Gaussian.

To summarise, the distribution of conditional returns computed by normalising upon constant intraday volatility forecasts turns out to be virtually the same across the different 5-minute intervals of the day and surprisingly close to Gaussian. However, it exhibits a small degree of fat-tailness that could be explained by the changes in volatility across the day that our simplified estimates do not capture. The investigation of this aspect will be the object of the following section.

4.3 The estimation of the model at an intraday level

As previously stated, volatility and density estimates produced by the model calibrated on a daily basis are not satisfactory at an intraday level. In order to obtain estimates that accurately describe the actual dynamics of high frequency data, the estimation of our continuous time specification must be refined by exploiting the information content of the 5-minute return series. The aim of the present section is to derive improved estimates for our continuous time model by means of a simple non-linear filtering technique in which the update occurs every 5 minutes, based on observed intraday market returns.

The assumption of a continuous time specification for our model is supported by the inspection of the serial correlation $\rho_{t,t+k}$ for 5-minute absolute returns within the same day (t and $t+k$ belong to the same day, $k \leq 80$) and between one day and the following (t and $t+k$ belong to adjacent days, $k \leq 161$), reported in Fig. 10. The two segments clearly seem to belong to the same curve, with no sudden break in the series, which is in favour of a continuous time specification.

In deriving high frequency estimates, we start by dismissing the persistent component of the stochastic volatility process in (5), since we expect the contribution of the fast mean-reverting part to be predominant for such purpose. We also ignore the impact of the leverage effect at intraday level, given that this component is quite persistent and, therefore, its effect should be better investigated and modeled at a lower frequency level.[32]

[32]To justify our choices, we have computed the proportions of the variance of log volatil-

In what follows, we first consider a standard diffusion model for the 5-minute volatility process, of the kind described in (5) for the transient component. We then introduce jumps in our volatility specification which will significantly improve our return density forecasts.

4.3.1 A simple diffusion process for intraday volatility

In order to describe the intraday volatility dynamics, we maintain the standard Gaussian Ornstein-Uhlenbeck specification employed in the daily model to characterise the evolution of the transient component, and we make use of the information available on high frequency returns to obtain improved estimates for the parameters of the process.

To impose the least possible structural assumptions on the derivation of high frequency volatility estimates, the latter are obtained and updated through a simple non-linear filtering technique based on observed intraday market returns. A range of possible discrete values for the log volatility $\ln(\sigma_j)$ for $j = 1, \ldots, N$ is specified, together with the corresponding set of initial probabilities P_j assigned to each value. These initial probabilities are then combined with the transition probabilities $P^a_{i,j}$ between log volatility values j and i to produce a discrete set of prior probabilities P^*_i for $i = 1, \ldots, N$ as follows:[33]

$$P^*_{i,t} \approx \sum_{j=1}^{N} P^a_{(i,t),(j,t-1)} P_{j,t-1} \tag{6}$$

which will then be applied to the corresponding volatility values in the range in order to return the intraday volatility estimate $\sigma^* = \sum_{i=1}^{N} P^*_{i,t}\sigma_i$. Within this framework, the discretisation of our continuous time volatility process is achieved by evolving the analogous discrete mean reverting process on a trinomial grid structure. The resulting transition probabilities, which we assume constant, are derived in Appendix A.1.[34]

Under the assumption of normality for the conditional returns, justified on the basis of the results derived in the previous section, a density forecast for the unconditional de-seasonalised returns r^f_t is represented by a mixture of normal densities, where each component is a normal with zero mean

---

ity innovations at 5-minute frequency attributable to each component: the leverage effect and the persistent volatility component explain, respectively, less than 5% and 4% of the total variance and, therefore, both components can be safely disregarded for the purpose of improving the high frequency volatility process.

[33]For simplicity of exposition, here t denotes the intraday moment previously indicated as it, hence $t = 1, \ldots, 82T$.

[34]For the practical implementation of the model, we chose a log volatility range between -1.5 and 1.5, with step size equal to 0.1, roughly corresponding to three times the estimated volatility of the mean reverting process. Alternative choices for the volatility range and the step size have been investigated, and the results do not seem to differ too significantly.

and standard deviation equal to one of the volatility values in the range multiplied by $\sqrt{\pi/2}$ and the mixing probabilities are given by the prior probabilities for the individual values in the volatility range:

$$r_t^f \sim \sum_{i=1}^{N} P_{i,t}^* N(0,\, \sigma_i \sqrt{\pi/2}) \tag{7}$$

Once the 5-minute unconditional de-seasonalised return r_t/s is observed, the probability $P_{i,t}^r$ that such return represents an observation from each of the Gaussian components of the mixture is computed and a Bayesian probability update is applied to the set of prior probabilities, producing a corresponding set of posterior probabilities $P_{i,t}^p$:

$$P_{i,t}^p \approx \frac{P_{i,t}^r P_{i,t}^*}{\sum_{j=1}^{N} P_{j,t}^r P_{j,t}^*} \tag{8}$$

which will replace the initial probabilities P_i in order to re-start the process. Unlike in the previous case, here volatility and return density forecasts are updated every 5 minutes on the basis of the actual evolution of returns observed in the financial market.

An important empirical issue concerning the implementation of our continuous time specification is the choice of the time step Δt. Our data seems to support a time step equal to 1/106 for 5-minutes intervals and 25/106 for the overnight period,[35] given that on average de-seasonalised unconditional overnight returns are about 25 times the corresponding 5-minute returns.[36]

We then need estimates of both the mean reverting coefficient α_s and the volatility parameter β_s such that the likelihood that the observed returns are realisations of our non-linear filtering model, given by:

$$L(r_T) \approx \prod_{t=1}^{T} \left(\sum_{j=1}^{N} P_{j,t}^r P_{j,t}^* \right) \tag{9}$$

is maximised and that, on average, the volatility of the intraday changes in the log volatility estimates is equal to the volatility parameter β_s of the process.[37] Working on a grid of possible values for α_s and β_s^2 (spaced at a step of, respectively, 0.05 and 0.01, which turns out to be a good compromise

[35] It is worth noticing that during the 25 steps of the overnight period, the process evolves only on the grid and the Bayesian update of the probability does not occur until the opening price for the day is known.

[36] Alternative choices for time steps of $\Delta t = 1/82$ for all intervals as well as $\Delta t = 1/288$ for 5-minute and $\Delta t = 207/288$ for overnight periods have been considered. However, the results turned out to be less satisfactory and robust.

[37] The likelihood function for mixture models is known to be unbounded at some points on the edge of the parameter space (see Kiefer [1978]). In our case, however, we do not attempt to maximise the likelihood *per se*, and we only use it to discriminate between various set of parameters that satisfy the volatility constraint.

between complexity and accuracy) we have found that estimated values of $\alpha_s = 0.6$ and $\beta_s = 0.28$ meet these requirements.

In the present work, we employ very simple filtering and estimation techniques to produce step-by-step volatility and return density forecasts and to obtain estimates of the relevant parameters. Much more sophisticated econometric methods have been developed recently in the literature: auxiliary particle filtering techniques have been suggested for volatility filtering in continuous and discrete time models by Pitt and Shephard [1999], Durham and Gallant [2002], Chib, Nardari and Shephard [2002], Johannes, Polson and Stroud [2002]. Markov Chain Monte Carlo (MCMC) methods provide very precise parameter estimates for a variety of diffusion and jump-diffusion models (see Elerian, Shephard and Chib [2001], Eraker [2001], Eraker, Johannes and Polson [2003]). Estimation strategies based on GMM procedures have been proposed, amongst the others, by Singleton [2001] and Bollerslev and Zhou [2002]. The implementation of such techniques would certainly improve the accuracy of our results but at the cost of an increased complexity which would not be justified in our context given that, as already stated, the best possible estimation accuracy is not our main concern. Therefore we choose to use simple techniques that still produce reasonably accurate estimates, as will be shown later.

Once the model is fully parametrised, 5-minute volatility estimates and return density estimates can be extracted. The time series of conditional returns is obtained by normalising unconditional returns r_{it} upon the intraday volatility forecasts σ^*_{it} computed 5-minutes earlier, and the analysis aimed at investigating shape and constancy of the conditional return distribution across the different subintervals of the day is replicated.

Again we compute summary sample statistics for the time series of conditional returns for each of the 82 intraday intervals and we report plots of the standard deviation (top Fig. 11) and the excess kurtosis (bottom Fig. 11) across the individual subintervals. A few spikes due to the presence of very large outliers can still be easily detected. If we ignore these outliers, we observe values for both the standard deviation and the excess kurtosis very similar across all the intervals of the trading day and very close to the values we would have for normally distributed conditional returns, with a standard deviation closely oscillating around the value of $\sqrt{\pi/2}$ and an average excess kurtosis of 0.9.

These findings are entirely in line with our expectations: volatility estimates which are updated at a high frequncy level can account for most of the fat-tailness left in the conditional return density after normalising with respect to volatility estimates which remain constant across the day. The assumption of a Gaussian distribution to describe returns computed by conditioning on accurate 5-minute volatility estimates turns out to work surprisingly well, actually better than expected.

At this stage, in order to correct for the presence of the outliers and for the residual fat-tailness while maintaining a specification valid in continuous time, we propose to introduce jumps in the model.

4.3.2 A process with jumps in intraday volatility

The introduction of jumps can take place in the returns process, in the volatility process, or in both. To avoid arbitrary assumptions on the most appropriate specification (and, therefore, to maintain it as parsimonious as possible), we analyse the nature of the outliers (identified with all conditional returns larger, in absolute value, than $3\sqrt{\pi/2}$) to decide whether they are more likely to represent jumps in returns or in volatility.

In order to investigate how persistent the increased volatility consequent to a jump turns out to be, the regression $(\mid r_{t+1} \mid - \sigma^*_{t+1}) = a + b(\mid r_t \mid - \mid \bar{r} \mid)$ has been run at a 5-minute level on both the entire sample and the subsample where r_t are all outliers. The estimated coefficients of $a = -0.014$ (s.e. 0.010) and $b = 0.027$ (s.e. 0.014) for the entire sample and $a = 0.03$ (s.e. 0.058) and $b = 0.109$ (s.e. 0.017) for the outliers suggest that the impact of the jumps seems to persist and not to die out immediately as the nature of jumps in returns would predict. The inspection of the temporal distribution of the outliers highlights a significant clustering in the incidence of jumps, which contradicts the i.i.d. assumption made for jumps in returns. Our empirical results indicate that the outliers exhibits more the features of jumps in volatility than those of jumps in returns.[38]

The continuous time process for the dynamics of intraday volatility then becomes:

$$d\ln(\sigma_{s,t}) = -\alpha_s dt \ln(\sigma_{s,t}) + \beta_s \sqrt{dt} dW_{s,t} + \sum_{i=1}^{N(t)} Y_i - \lambda dt E[Y] \tag{10}$$

where $N(t)$ denotes the total number of jumps in dt (arrivals of a Poisson process with intensity λ) and Y_i are i.i.d. random variables corresponding to the Poisson jump magnitudes.[39] In order to fit the discrete version of this stochastic volatility model into our grid structure, we need to work with jumps of discrete size (i.e. Y_i will have a discrete distribution), expressed as a multiple of our step size Δ_r. To simplify the analysis, jump sizes and intensities are assumed to be constant.

The discrete values for the jump magnitudes and the respective intensities are obtained via calibration with the empirical features of the outliers.

[38]This is in line with some recent findings which point out how models with diffusive stochastic volatility and jumps in returns are incapable of capturing the empirical features of equity returns (see Pan [2002], Bates [2000], Duffie, Singleton and Pan [2000], Eraker, Johannes and Polson [2003]). A more rigorous specification would also allow for jumps in returns. For simplicity, here we restrict our attention to jumps in volatility, which still yields good results.

[39]A compensated jump process has been chosen to maintain the mean of the volatility process unchanged.

Our outliers can be roughly grouped into three categories: small (up to 6 standard deviations), medium (between 6 and 12 standard deviations) and large (above 12 standard deviations). This suggests three possible kinds of jumps in volatility: a small jump, of size equal to $3\Delta_r$ which, given our log volatility model, is roughly equivalent to an outlier in conditional returns of 4 standard deviations; a medium jump of size $10\Delta_r$, corresponding to 8 standard deviations; a large jump of size $20\Delta_r$ roughly equivalent to 22 standard deviations. A simple investigation of the frequency of outliers reveals that, on average, a jump occurs every other week, therefore we choose $\lambda = 1/10$ as our overall jump intensity expressed on a daily basis. A more detailed analysis of the frequency of jump incidence for each group (given by the number of jumps occurred within that group divided by the total number of days in the sample) returns frequencies of 0.00729 (= 7/960) for large jumps, 0.027 (= 26/960) for medium jumps and 0.0708 (= 68/960) for small jumps. On the basis of these empirical frequencies, properly rescaled by the overall jump intensity $\lambda = 1/10$, we derive probabilities of jump incidence for each group equal to $p^l = 8\%$ for the biggest jumps, $p^m = 27\%$ for intermediate jumps and $p^s = 65\%$ for small jumps.[40]

Once the jump sizes and intensities have been specified, the Bayesian filtering procedure illustrated in the previous section can be entirely replicated here, with the only difference that the log volatility process evolving on the grid is now the mean reverting model augmented by the jumps component. Therefore, the transition probabilities must be recomputed, following Amin [1993] (details in Appendix A.2). The values for log volatility range, initial probabilities, step size and time step are the same as before. The estimates for the remaining parameters of the volatility process produced by our methods are equal to values of $\alpha_s = 0.7$ and $\beta_s = 0.24$.

As before, we obtain 5-minutes ahead volatility and return density forecasts, whose accuracy in both absolute and relative[41] terms needs to be adequately assessed.

4.4 The appraisal of intraday volatility and density estimates

The present section focusses on the assessment of our high frequency volatility and return density estimates through the implementation of statistical techniques borrowed from both point and density forecast evaluation practice.

Point forecast evaluation techniques are used to assess the 5-minute

[40] To ensure that our choices of jumps sizes and intensities are not too arbitrary, we have replicated the analysis using alternative values obtained from different groupings of the outliers. As long as we allow for at least three kinds of jumps in volatility, the results do not differ too much.

[41] In comparison with constant intraday estimates and changing intraday estimates without jumps.

volatility estimates, through a comparison with the absolute value of demeaned, de-seasonalised high frequency unconditional returns, taken as a proxy of the actual intraday volatility level. In line with the existing literature on volatility forecasts evaluation,[42] we first regress the absolute return on the volatility prediction, $|r_t/s| = \alpha + \beta\sigma_t^* + \epsilon_t$. The forecast is unbiased only if $\alpha = 0$ and $\beta = 1$ and, what is most important for a good prediction, has got small forecast errors if the standard error of the regression is small and R^2 is large. However, the presence of a strong noisy component in our volatility proxy induces very high standard errors and very small R^2 coefficients: the regression performed on returns simulated from exact volatility forecasts returns a standard error of 0.6687 and $R^2 = 0.2244$.

We also report a measure P of the proportion of volatility explained by the forecasts, first introduced by Blair, Poon and Taylor [2001], which compares the amount of variations in the forecast errors with that in actual volatility, according to:

$$P = 1 - \frac{\sum (|r_t/s| - \sigma_t^*)^2}{\sum(|r_t/s| - E[|r_t/s|])^2}$$

Values closer to one indicate better forecasts, with a small variation in the forecast errors. Finally, a standard Mean Absolute Deviation measure, determined as simple average of the absolute deviations of the volatility forecasts from the volatility proxies, is computed.

The findings from the point forecast evaluation of constant and changing 5-minute volatility estimates are displayed in Table 5. As expected, the results indicate a very poor forecasting performance in all cases, but we should keep in mind that they are heavily distorted by the noise in the high frequency absolute returns. In a relative comparison, the forecasts updated on an intraday basis (with and without jumps) perform significantly better than the ones updated on a daily basis, as suggested by a smaller standard error (0.77 against 0.96) and a higher R^2 (0.16 against 0.09) of the regression, as well as by a higher value of the P statistic (0.14 vs. 0.09) and a slightly smaller MAD (0.635 against 0.64).

Density forecasts evaluation techniques are employed to evaluate the intraday density forecasts for the returns.[43] Following a standard procedure, first applied to density forecast evaluation problems by Diebold, Gunther

[42] See, for a comprehensive review, Granger and Poon [2003].

[43] We briefly recall what our intraday density forecasts for the returns look like, under the assumption of normally distributed conditional returns. When the volatility forecasts stay the same across the day, the 5-minutes ahead density forecast for the unconditional returns on each of the intraday intervals is given by a Gaussian with zero mean and standard deviation equal to the forecasted volatility for that day multiplied by $\sqrt{\pi/2}$. For changing intraday volatility forecasts, the density forecast for the returns is represented by the mixture of normal densities derived in the previous section.

and Tay [1998], from the sequence of 5-minutes ahead density forecasts $f_t(r)$, we derive the series of the probability integral transforms of the realised intraday unconditional returns taken with respect to the corresponding density forecasts as follows:

$$z_t = \int_{-\infty}^{r_t/s} f_t(r)\,\mathrm{d}r \tag{11}$$

If the forecasts and the true densities coincide, then the sequence of PITs is distributed as i.i.d. $U(0,1)$. Equivalently, the sequence of transformed PITs, where a transformation to normality is applied to the PITs series (see Berkowitz [2001]), follows an i.i.d. $N(0,1)$.

To guarantee more robust results against possible misspecifications of different type in the forecasted distributions, several goodness-of-fit techniques have been implemented. The popular Kolmogorov-Smirnov, Anderson-Darling and Watson statistics have been chosen to test for uniformity. The normality is assessed via Jarque-Bera and Doornik-Hansen tests, as well as via normal Q-Q plots. Two likelihood ratio tests are performed to test for independence (LR1) and for the joint hypothesis of independent observations with zero mean and unit variance (LR2). Since the size of our sample is huge (82 observations for 960 days), virtually any distributional forecast, even a very good one, can be easily rejected. To overcome, at least partially, this problem, we have sorted our sample in four subsamples according to the level of the volatility forecast.

The results from the density forecasts tests on both the entire sample and the four groups are reported in Table 6. The null hypothesis that the return density forecasts represent accurate predictions of the actual distribution of unconditional returns is generally rejected by all our goodness-of-fit statistics on the entire sample as well as on the four subsamples, for both constant and changing intraday volatility estimates. However, a substantial improvement is recorded in the production of density forecasts when volatility estimates are updated every 5 minutes, which becomes even more striking when jumps are introduced in the volatility process. The values of the goodness-of-fit statistics (especially of the normality tests) are now much closer to their critical values and the normal Q-Q plots for the case of changing volatility estimates (Fig. 12 and 13) display empirical quantiles fairly close to the normal ones, especially for the model with jumps. More specifically, it seems that the fat-tailness induced by the jumps in volatility corrects for most of the misspecification in the tails recorded for both the daily updating method and the intraday method without jumps.

To summarise, in the light of our findings, we can conclude that in order to produce accurate volatility and return density estimates at an intraday frequency, we need to rely on a non-linear filtering technique where the forecasts are updated every 5 minutes on the basis of the current value of unconditional returns observed in the market. Also, a simple diffusion process for the intraday volatility is not appropriate and a specification which allows for jumps is to be preferred.

A stochastic volatility model of the kind in (5) which works well at both high and lower frequency level can be obtained by combining the permanent component whose parameters are estimated on a daily basis with the transient component (with jumps) whose parameters are estimated on a non-linear intraday filtering model.

5 A Monte Carlo Simulation Exercise

Throughout the previous section the relevant features in the evolution of the observed returns have been carefully isolated, studied and modelled: all the individual components have then been assembled together to produce a complete 5-minute continuous time model as follows:[44]

$$\begin{aligned}
r_t &= s_i \sigma_t \varepsilon_t \\
dM_t &= \theta(\ln(S_{t-1}) - M_{t-1})dt \\
\ln \sigma_t &= \kappa(\ln(S_{t-1}) - M_{t-1}) + \ln \sigma_{s,t} + \ln \sigma_{l,t} \\
d\ln(\sigma_{s,t}) &= -\alpha_s dt \ln(\sigma_{s,t}) + \beta_s \sqrt{dt} dW_{s,t} + \sum_{i=1}^{N(t)} Y_i - \lambda dt E[Y] \\
d\ln(\sigma_{l,t}) &= -\alpha_l dt \ln(\sigma_{l,t}) + \beta_l \sqrt{dt} dW_{l,t}
\end{aligned} \tag{12}$$

In order to: 1) test whether the dynamics of unconditional high frequency returns generated from our model in (12) does actually mirror the empirical one; 2) assess whether our simple estimation procedure produces reliable estimates, we perform a simple Monte Carlo simulation experiment.

A number of 82 intraday unconditional returns is generated each day for a total of 960 days according to the model in Equation (12) where:

- The deterministic seasonal coefficients are given by their smoothed estimates.
- The parameters for the leverage component are those estimated at a daily level in Section 4.2.
- Realisations of the stochastic volatility are generated from the two-factor model where both components have Gaussian innovations and the parameters are estimated on a daily basis for the permanent component and on an intraday basis with jumps for the transient component.

[44] Again, t denotes time on a 5-minute, and not daily, basis, and dt indicates the intraday interval.

- Realisations of the conditional returns are obtained by sampling from the Gaussian density $N(0, \sqrt{\pi/2})$.

For simplicity, we only simulate five full samples. In fact, when the focus in on assessing whether our data set could actually represent a random sample generated from the model, we restrict the comparison between the features of the simulated returns and those of the empirical data to one sample only. Instead, all five samples are employed to assess the estimation technique.

We start by looking at the plots of higher moments, skewness and excess kurtosis, computed across the time series of high frequency returns for each of the intraday intervals, which indicate very similar values for both simulated and observed returns (Fig. 14). We then aggregate the simulated high frequency values to derive daily log volatility proxies as averages of absolute de-meaned, de-seasonalised returns, and daily measures of leverage. The time series of these daily simulated variables (Fig. 15, bottom) are contrasted with their daily empirical counterparts (Fig. 6, bottom) to check for possible significant differences in the evolution of simulated and observed volatility proxy and leverage measure. The dependence between log volatility and leverage component from simulated data has also been investigated, via scatter plot (Fig. 15, top) and computation of the correlation coefficient, equal to $\rho = -0.533$. The results are very encouraging, since both the temporal evolution of simulated volatility proxy and leverage measure and their correlation structure closely resemble the empirical ones. These findings at high frequency, as well as daily, level, suggest that the model in (12) seems capable of capturing and replicating the most significant features observed in futures equity returns.

To evaluate the adequacy of the estimation techniques employed so far, we have derived estimates of our model from each of the five simulated samples and compared the resulting parameters with the actual parameters of the data generating process. Following the steps of our data analysis, we start by investigating the seasonal component, whose pattern (not reported here) is indistinguishable from the one shown by the market data for all five simulated samples. Daily measures of log volatility and leverage computed on simulated data are then used to obtain estimates for the leverage model through the regression in 3 and for the two-factor stochastic volatility model via Kalman filter on the residuals from the previous regression. The estimates, displayed in Table 7, are in all cases very close to the original parameters of the process from which the samples have been simulated, and only the mean reversion parameter of the transient volatility component is slightly underestimated in all samples. In relative terms, the larger (but still quite small in absolute terms) dispersion can be observed for the estimates of the parameters of the permanent volatility component.

As before, the non-linear filtering technique with intraday updating of volatility and return density estimates is implemented in order to refine the high frequency volatility process. First we produce estimates of the volatility specification without jumps and we employ the resulting volatility forecasts to obtain a series of conditional returns. Again the inspection of the outliers provides us with information on the characteristics of the jumps. For simplicity, we maintain the discrete magnitudes of the three kinds of jumps unchanged (equal to $3\Delta_r$, $10\Delta_r$ and $20\Delta_r$, with $\Delta_r = 0.1$) and we only re-estimate the overall and the individual jump intensities on the basis of their frequency of incidence. Finally, we re-estimate the parameters of mean reversion and volatility of volatility on the grid. The estimates for the log volatility process with jumps, shown in the bottom part of Table 7, are fairly satisfactory since they turn out to be quite close to the actual parameters of the data generating process. However, we can detect an underestimate of both the overall incidence of jumps, with jumps occurring every 3 or 4 weeks for four out of five simulated samples, and the intensity of the biggest jumps, since no evidence of the presence of large jumps can be found for two samples. The mean reversion and volatility parameters of the diffusion component seem also to be a little underestimated. On the whole, our findings suggest that the estimates produced by applying our simple techniques are quite reliable and adequate for our purposes.

6 Conclusions and Further Work

In the present work we have attempted to build a simple, although accurate, continuous time model capable of describing and replicating the dynamics of both high and moderate frequency index returns, by performing a careful analysis of a set of intraday data, aimed at: 1. identifying the relevant features that need to be modelled; 2. investigating the best possible model specification, without imposing too much structure *a priori* and by testing step by step the assumptions made.

At the conclusion of our analysis we propose a model where the seasonal intraday volatility component is deterministic and constant through time, the stochastic volatility component follows a two-factor mean reverting process with jumps in the transient factor, where volatility forecasts are updated every 5 minutes and the conditional return distribution is Gaussian and fairly constant across the subintervals.

An additional attractive feature of our work is that from the general model specification, which is consistent for both high and moderate frequency data, we can easily obtain simplified versions which have the correct properties for the specific time horizon of interest. This aspect would deserve further investigation.

A Transition probabilities for the high frequency volatility process

In what follows we derive the transition probabilities between volatility values, associated with the discretised versions of the two continuous time processes (without and with jumps) chosen in 4.3 to model the dynamics of intraday volatility.

A.1 Diffusion model

The discrete analogous of the continuous time model proposed to describe the evolution of 5-minute volatility, obtained via Euler discretisation,[45] is given by:

$$\Delta \ln(\sigma_{s,t}) = -\alpha_s \Delta t \ln(\sigma_{s,t}) + \beta_s \sqrt{\Delta t}\omega_s$$

with $\omega_s \sim N(0,1)$. In order to fit this structure into our non-linear filtering technique, we let each log volatility value in our range $\ln(\sigma_j)$ for $j = 2, \ldots, N-1$, evolve according to a trinomial tree. Between t and $t+\Delta t$ the log volatility (equal to $\ln(\sigma_i)$ in t) can go up by the amount Δ_r (step size) to level $\ln(\sigma_{i-1})$ with probability p_u, down by $-\Delta_r$ to level $\ln(\sigma_{i+1})$ with probability p_d or stay at level $\ln(\sigma_i)$ with probability p_e. Here we assume that both the transition probabilities p_u, p_d, p_e and the step size are constant. Following a standard procedure, the transition probabilities are obtained by equating the first two moments of the discrete time process to the corresponding moments of the continuous time model:

$$\begin{aligned}
p_u\Delta_r + p_d(-\Delta_r) &= E[\Delta \ln(\sigma_i)] = -\alpha_s \Delta t \ln(\sigma_i) \\
p_u(\Delta_r)^2 + p_d(-\Delta_r)^2 &= var[\Delta \ln(\sigma_i)] + E^2[\Delta \ln(\sigma_i)] = \beta_s^2 \Delta t + [-\alpha_s \Delta t \ln(\sigma_i)]^2 \\
p_u + p_d + p_e &= 1
\end{aligned}$$

for $i = 2, \ldots, N-1$. For $i = 1$, $p_e^{(1)} = p_e$ and $p_d^{(1)} = 1 - p_e^{(1)}$ and, similarly, for $i = N$, $p_e^{(N)} = p_e$ and $p_u^{(N)} = 1 - p_e^{(N)}$. Δ_r should be chosen close to three times the standard deviation of the continuous model ($\beta_s\sqrt{dt}$) to ensure an adequate representation of the process in a discrete framework.

A.2 Jump diffusion model

The discrete time version of our process for intraday stochastic volatility with jumps is given by:

$$\Delta \ln(\sigma_{s,t}) = -\alpha_s \Delta t \ln(\sigma_{s,t}) + \beta_s \sqrt{\Delta t}\omega_s + \sum_{i=1}^{N(\Delta t)} Y_i - \lambda \Delta t E[Y]$$

[45] In this context, the Euler discretisation should not introduce a significant bias, since we work with high frequency data, which are frequently spaced.

The probability of a generic jump occurring in a time step Δt is equal to $\lambda\Delta t$ and we assume that multiple jumps cannot happen in a single time step. As before, the step size is equal to Δ_r. Let p^l denote the probability of a large jump of size $20\Delta_r$, p^m the probability of an intermediate jump of $10\Delta_r$ and $p^s = 1 - p^l - p^m$ the probability of a small jump of size $3\Delta_r$.

Between t and $t+\Delta t$ the log volatility (equal to $\ln(\sigma_i)$ in t) can evolve according to the trinomial structure analysed before, with probability $1-\lambda\Delta t$ (no jumps occurring), or it can jump up by $3\Delta_r$, $10\Delta_r$ or $20\Delta_r$ with probabilities $\lambda\Delta t p^s$, $\lambda\Delta t p^m$, $\lambda\Delta t p^l$, respectively. Again, the transition probabilities are obtained by equating the first two moments of the discrete and the continuous time process:

$$\begin{aligned}
(1-\lambda\Delta t)(p_u\Delta_r + p_d(-\Delta_r)) + \lambda\Delta t(p^s 3\Delta_r + p^m 10\Delta_r + p^l 20\Delta_r) &= \\
-\alpha_s \Delta t \ln(\sigma_i) & \\
(1-\lambda\Delta t)(p_u(\Delta_r)^2 + p_d(-\Delta_r)^2) + \lambda\Delta t(p^s(3\Delta_r)^2 + p^m(10\Delta_r)^2 + p^l(20\Delta_r)^2) &= \\
\beta_s^2\Delta t + \lambda\Delta t(p^s(3\Delta_r)^2 + p^m(10\Delta_r)^2 + p^l(20\Delta_r)^2) + [-\alpha_s\Delta t\ln(\sigma_i)]^2 & \\
p_u + p_d + p_e = 1 &
\end{aligned}$$

with the following restrictions dictated by the grid structure:

- $\lambda = 0$, $p_e^{(1)} = p_e$, $p_u = 0$ and $p_d^{(1)} = 1 - p_e^{(1)}$ for $i = 1$ (highest volatility value, no possibility of upward movements).
- $\lambda = 0$ for $i = 2, 3$ (no possibility of jumps).
- $p^s = 1$, $p^m = p^l = 0$ for $i = 4, \ldots, 10$ (only small jumps possible).
- $p^m = 1 - p^s$, $p^l = 0$ for $i = 11, \ldots, 20$ (only small and intermediate jumps possible).

Table 1: Summary sample statistics for intraday returns.

Unconditional Intraday Returns		
	5-minute	Overnight
Mean	-9.95E-09	0.0002
Std. Dev.	0.121%	0.575%
Skewness	0.883	-0.378
Excess Kurtosis	35.476	3.044
Unconditional De-seasonalised Intraday Returns		
	5-minute	Overnight
Mean	-1.15E-05	1.9E-17
Std. Dev.	1.409%	1.408%
Skewness	0.908	-0.378
Excess Kurtosis	35.622	3.044

Table 2: First order serial correlation of intraday returns.

	1997-98	1999	2000	2001
Serial correlation	-0.0793*	-0.0901*	-0.0491	-0.0364
Critical value	-0.0372	-0.0423	-0.0493	-0.0570
Estimated BA spread (%)	-0.064%	-0.064%	-0.059%	-0.053%

* statistically significant at 5% confidence level.

Table 3: Tests for equality of seasonal volatility patterns.

Subsamples	% of rejection mean equality test	Average value ratio intraday coeff.	Std. dev. ratio intraday coeff.
Monday	0.0%	1.002	0.057
Tuesday	12.20%	0.998	0.075
Wednesday	2.44%	0.998	0.056
Thursday	0.0%	1.002	0.055
Friday	4.88%	1.001	0.064
First-second half	12.20%	1.004	0.074
High-low volatility	7.32%	1.000	0.064

Table 4: Coefficient estimates for two-factor AR(1) model.

	ρ	$\alpha(=1-\rho)$	β^2	Half life in days	Unconditional var. $\beta^2/(1-\rho^2)$
Transient	0.261 (0.095)	0.739	0.044 (0.570)	0.94	0.046
Permanent	0.982 (0.008)	0.018	0.0019 (0.321)	37.50	0.051

Table 5: Comparison intraday volatility estimates.

	Daily	5-minute	5-minute with jumps
Regression			
alpha	-0.0164*	-0.0421	-0.0487
	(0.0139)	(0.0097)	(0.0083)
beta	1.0505	1.0252	1.0416
	(0.0158)	(0.0207)	(0.0173)
S.E. regression	0.9620	0.7653	0.7786
R^2	0.0895	0.1677	0.1596
P-statistic	0.0882	0.1429	0.1387
MAD	0.6433	0.6343	0.6364

* not significantly different from zero at 5% level.

Table 6: Distributional forecast evaluation.

	Uniformity tests			**Normality tests**		**LR tests**	
	K-S (1.36)	A-D (2.49)	Watson (0.19)	J-B (5.99)	D-H (5.99)	LR1 (3.84)	LR2 (7.81)
Entire sample (78720 obs)							
Daily	3.76*	65.66*	3.32*	13,995*	8,255*	249.18*	1,418.0*
5-minute	3.67*	31.66*	4.66*	2,271*	1,733*	2.36	4.36
5-m. jumps	3.90*	41.17*	6.23*	895*	751*	1.23	5.29
4 Subsamples							
Daily							
Low vol.	2.08*	19.99*	1.09*	3,452*	1,948*	51.78*	348.61*
Medium low	2.19*	17.87*	1.32*	2,724*	1,698*	62.80*	277.21*
Medium high	2.48*	15.83*	0.75*	2,322*	1,503*	68.76*	359.43*
High vol.	2.62*	15.83*	0.67*	5,844*	3,127*	67.04*	445.56*
5-minute							
Low vol.	2.79*	18.59*	2.69*	975*	693*	2.37	14.58*
Medium low	2.40*	13.18*	1.96*	686*	497*	1.06	5.16
Medium high	2.08*	7.64*	1.04*	312*	252*	2.11	5.54
High vol.	2.88*	3.52*	0.50*	425*	345*	1.41	11.91*
5-m. jumps							
Low vol.	3.61*	24.34*	3.99*	480*	373*	2.58	33.04*
Medium low	2.42*	13.89*	2.07*	244*	196*	0.55	6.68
Medium high	2.07*	5.84*	0.80*	152*	133*	0.20	2.60
High vol.	2.92*	3.65*	0.59*	120*	108*	0.71	8.51*

* rejected at 5% level.

Table 7: Estimates from simulated samples.

Samples	1	2	3	4	5	Avg.	Std.Dev.	Data
Daily model								
κ	-4.74	-4.02	-4.94	-3.89	-4.15	-4.35	0.464	-4.34
α_s	0.632	0.608	0.683	0.647	0.666	0.647	0.029	0.734
α_l	0.024	0.023	0.035	0.010	0.013	0.021	0.010	0.018
β_s	0.179	0.220	0.203	0.235	0.231	0.214	0.023	0.210
β_l	0.042	0.032	0.059	0.052	0.047	0.046	0.010	0.043
ρ	-0.55	-0.45	-0.57	-0.45	-0.48	-0.50	0.057	-0.545
Intraday model								
α_s	0.7	0.6	0.6	0.6	0.7	0.64	0.055	0.7
β_s	0.21	0.22	0.20	0.24	0.21	0.216	0.015	0.24
p^s	0.7	0.6	0.7	0.65	0.76	0.682	0.060	0.65
p^m	0.3	0.32	0.27	0.35	0.20	0.288	0.057	0.27
p^l	0.0	0.08	0.03	0.0	0.04	0.03	0.033	0.08
λ	0.066	0.05	0.10	0.066	0.05	0.066	0.020	0.10

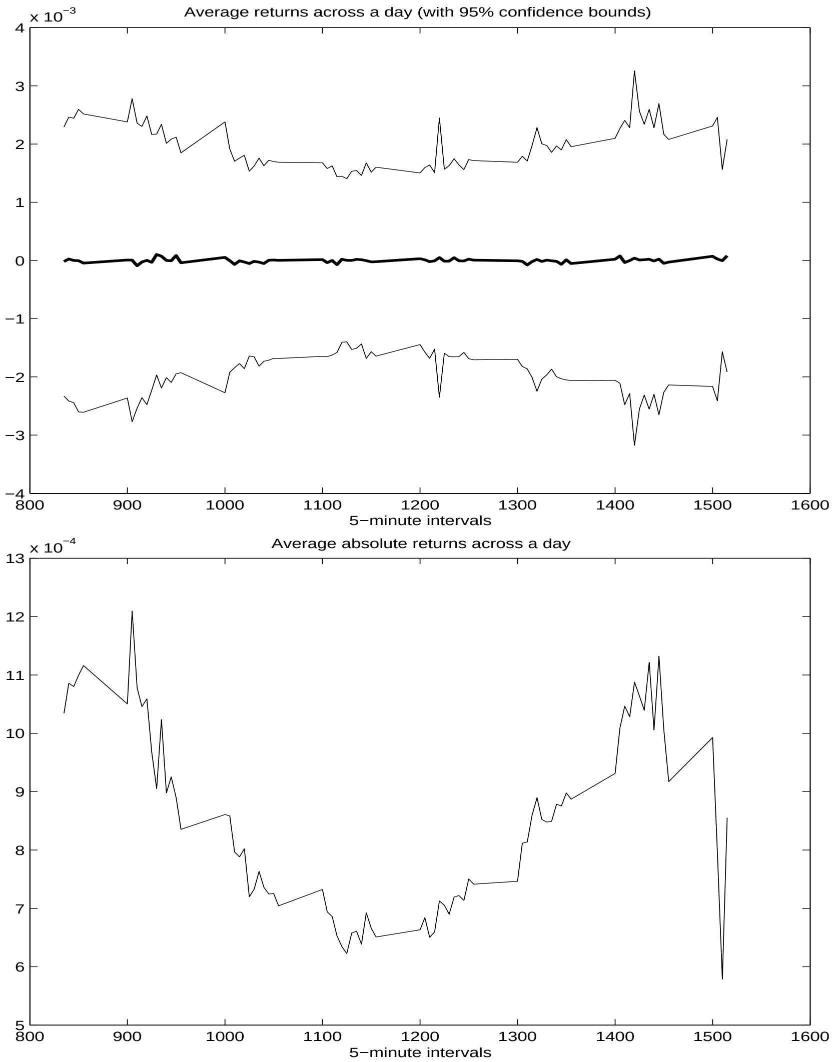

Figure 1: Intraday patterns of average returns and volatility.

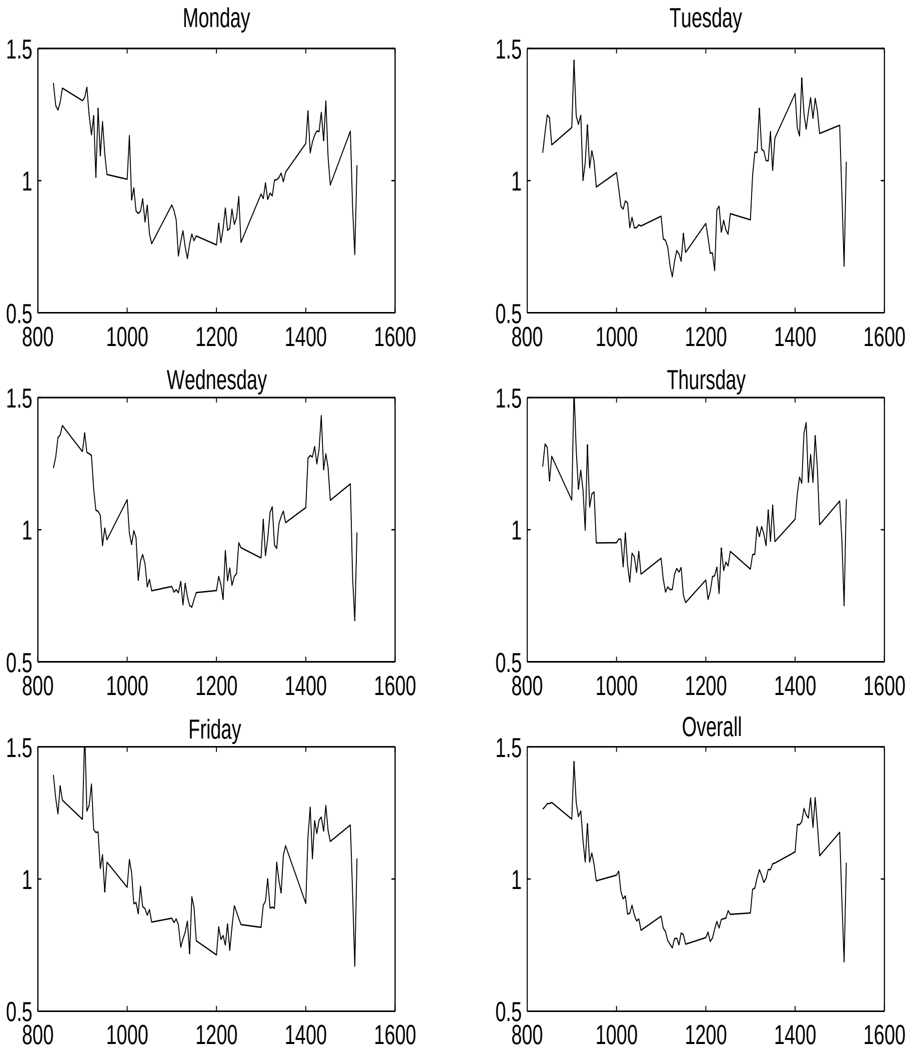

Figure 2: Seasonality patterns for days of the week.

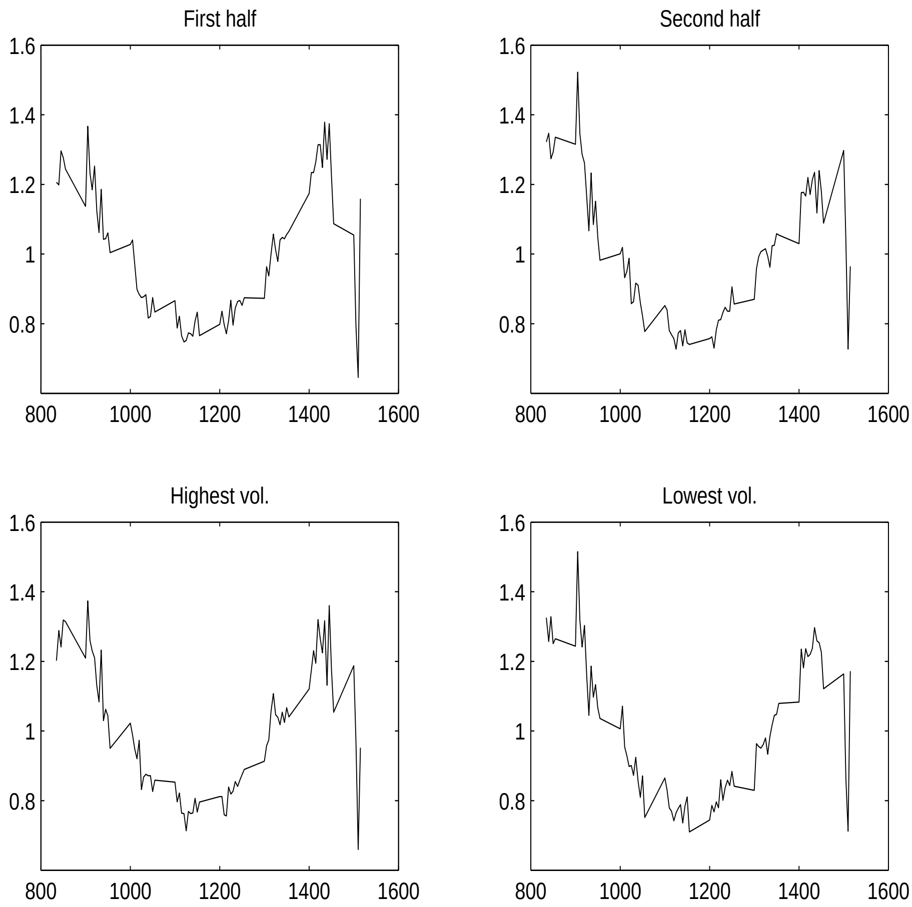

Figure 3: Seasonality patterns for different subsamples.

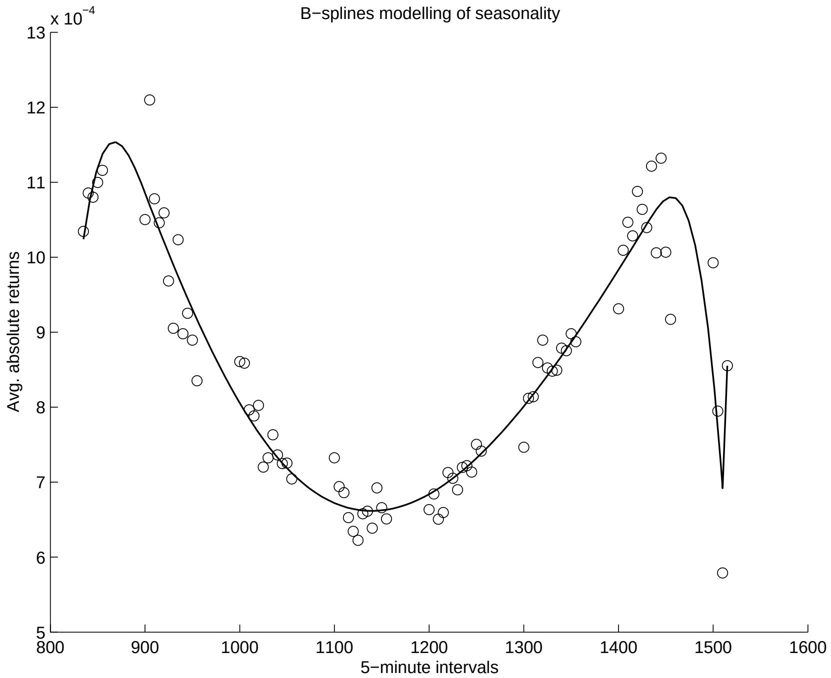

Figure 4: Smoothed B-spline estimation of seasonal coefficients.

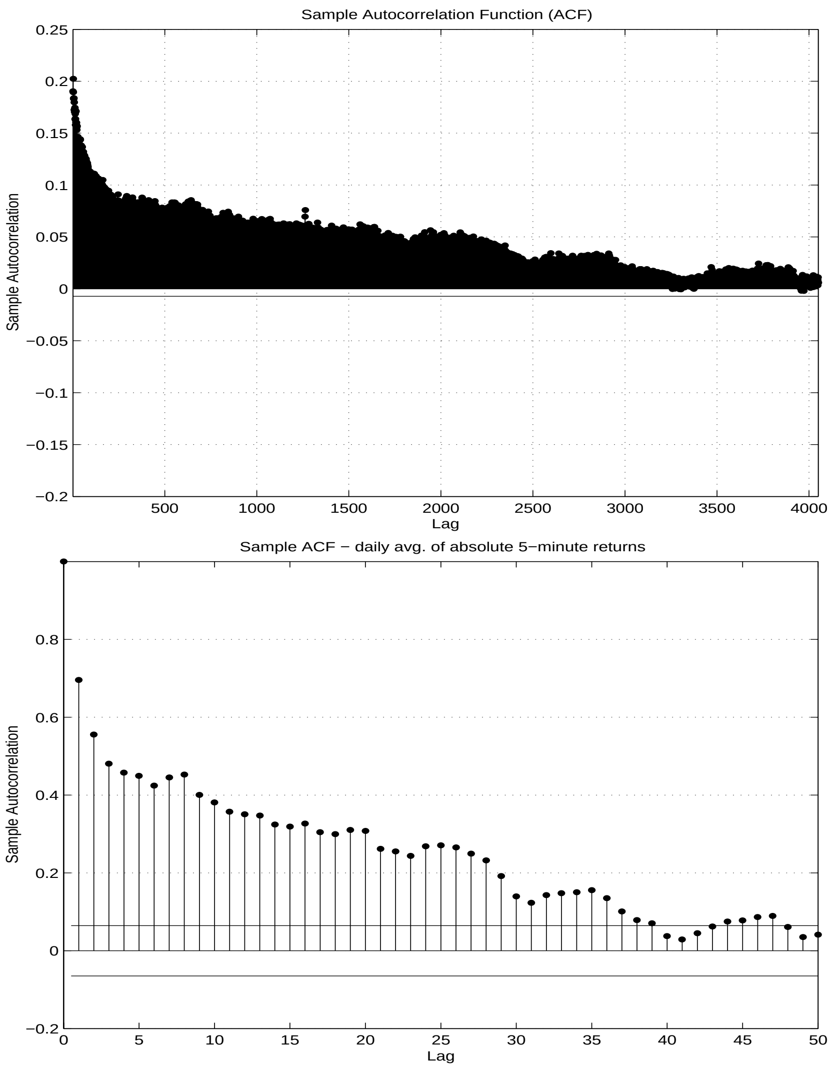

Figure 5: ACF of intraday absolute returns and daily average absolute returns.

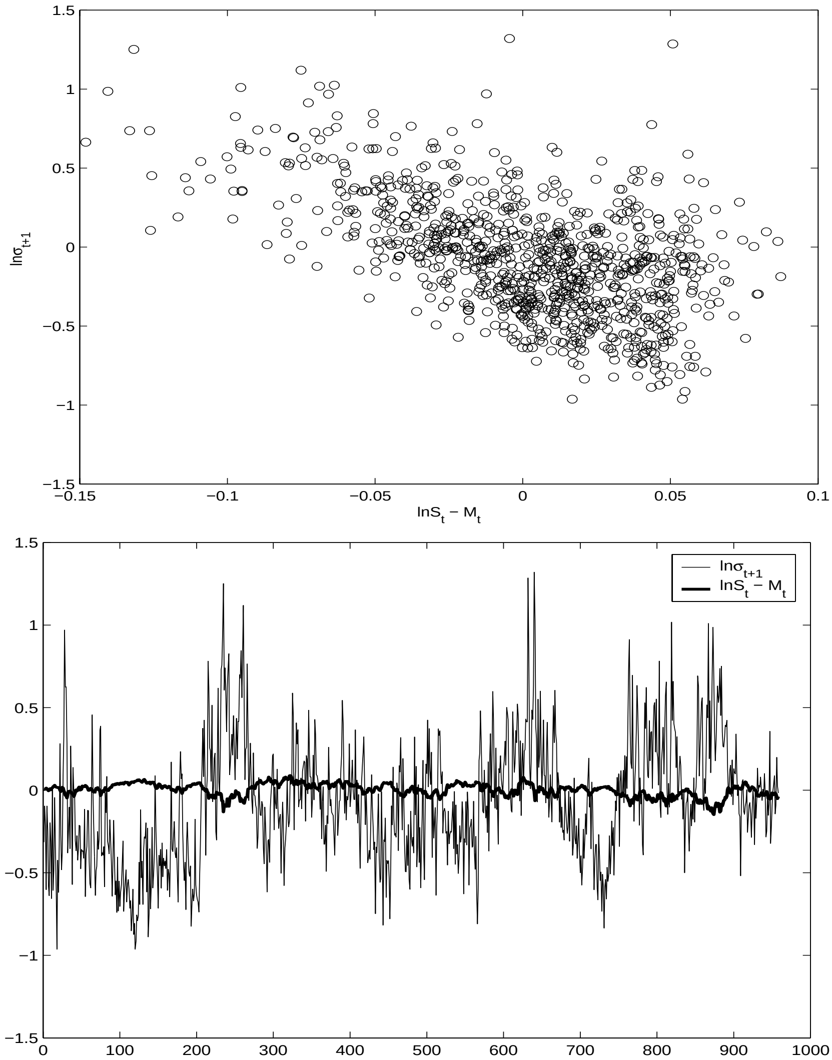

Figure 6: Scatter plot and time series of leverage measure against log volatility proxy - Market data.

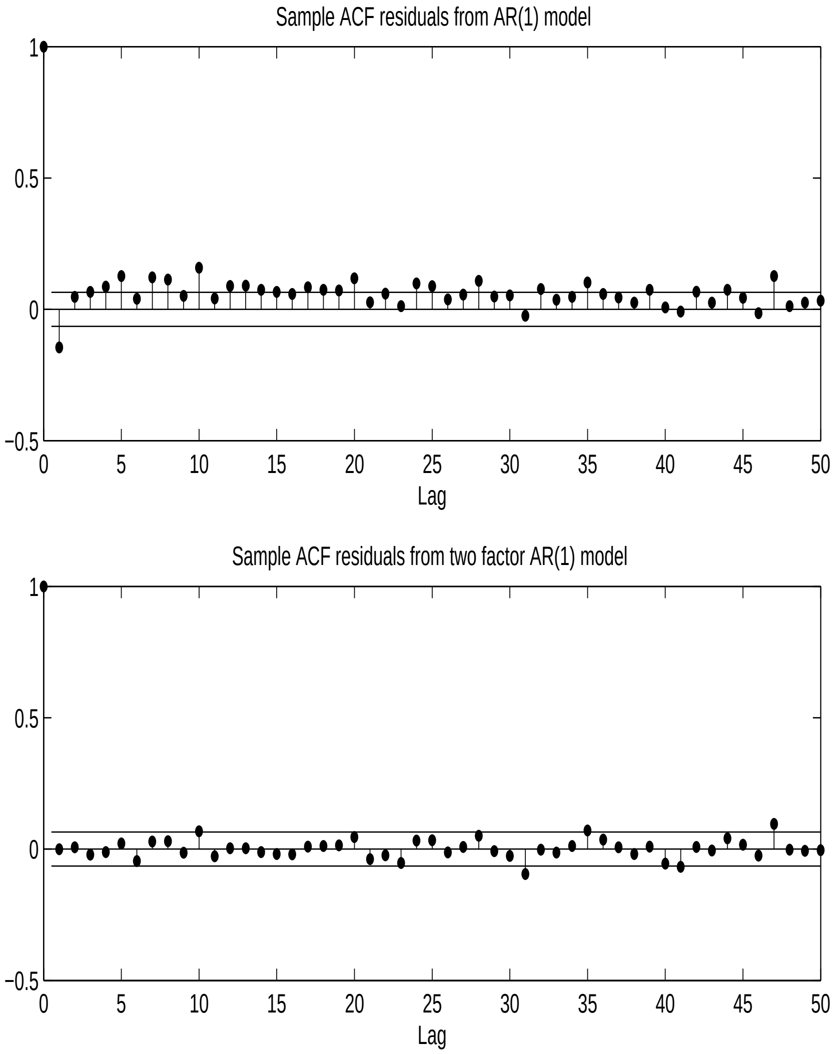

Figure 7: ACF of residuals from AR(1) and two-factor AR(1) specifications.

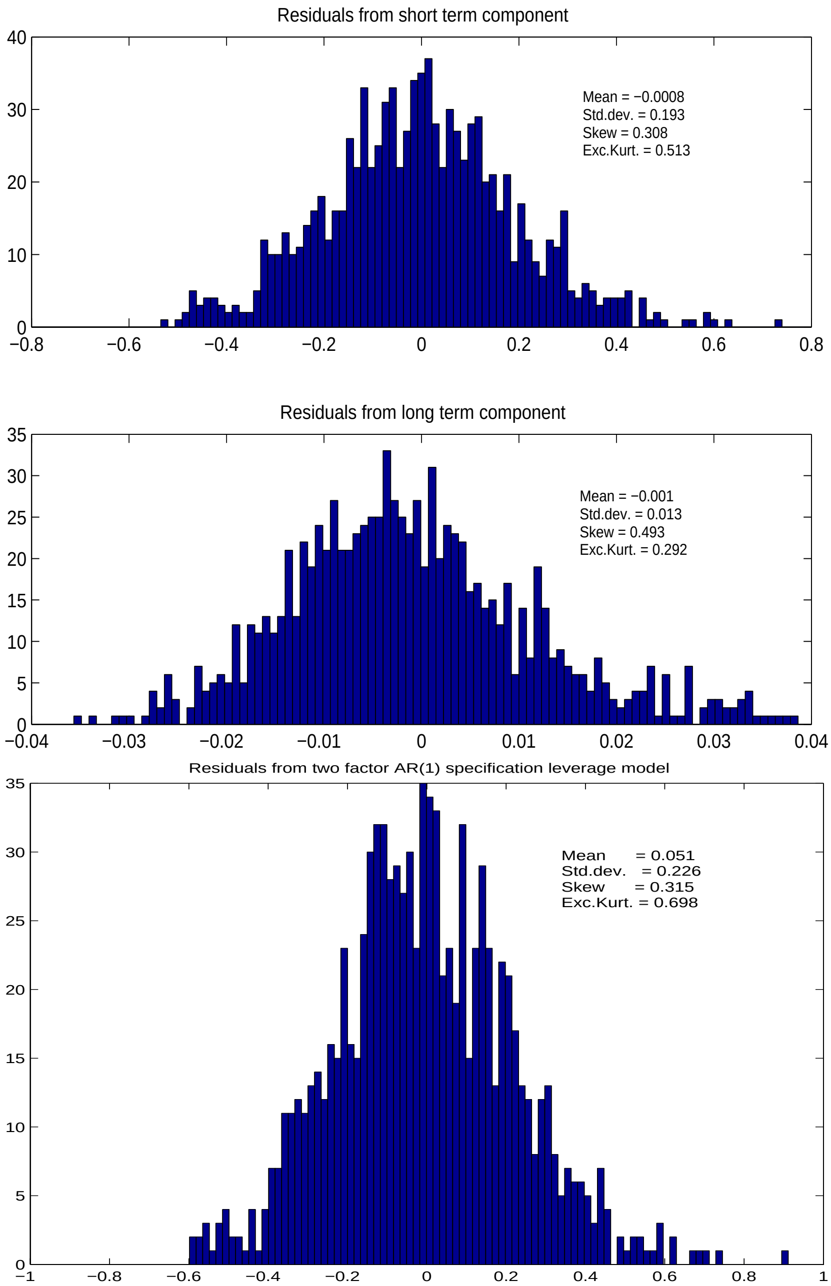

Figure 8: Histograms of residuals from state and measurement equations.

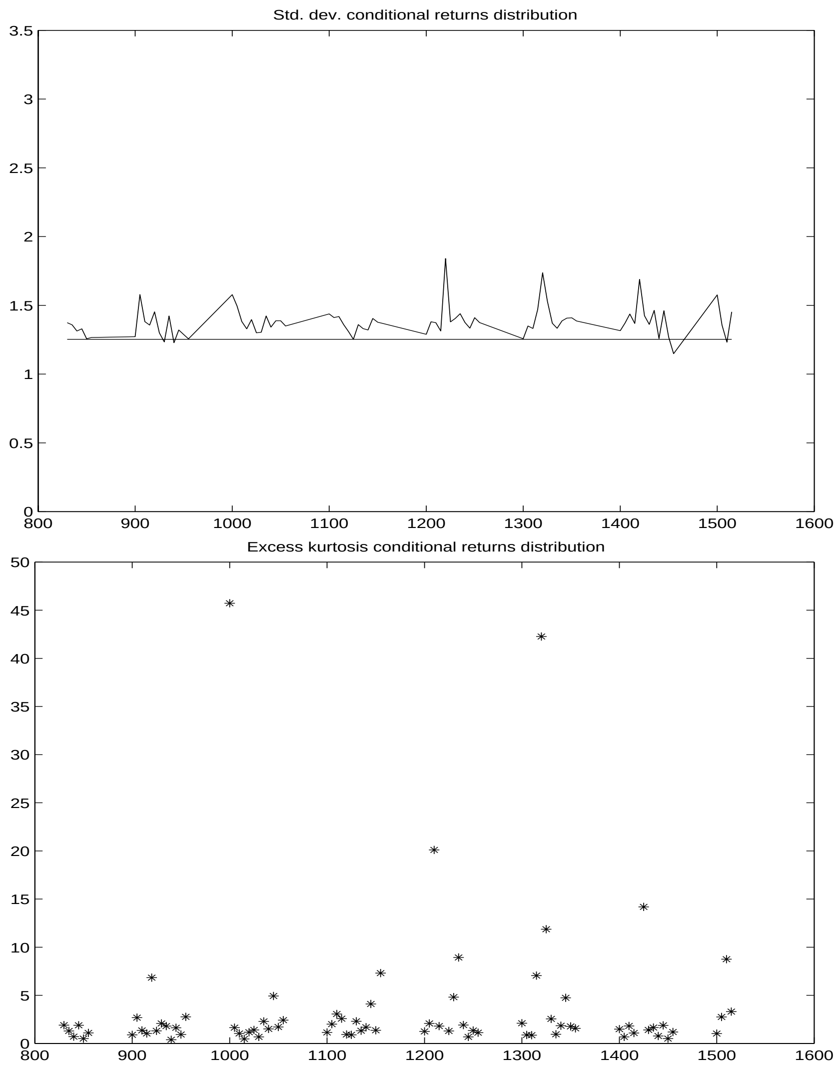

Figure 9: Standard deviation and excess kurtosis of conditional returns - constant intraday volatility estimates.

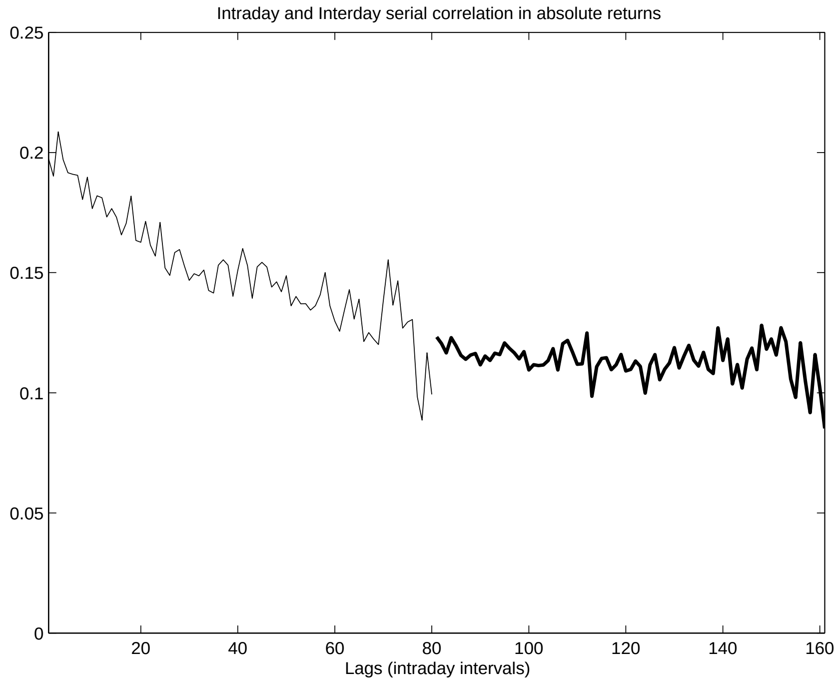

Figure 10: Intraday and interday serial correlation of absolute returns.

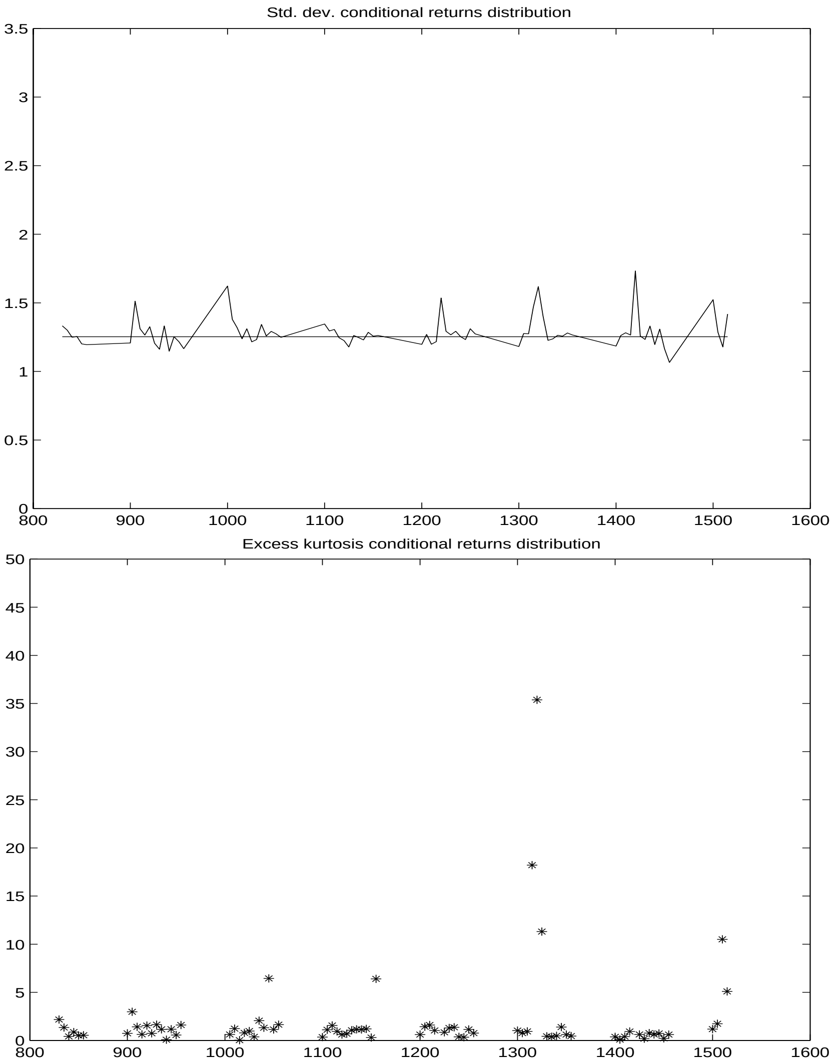

Figure 11: Standard deviation and excess kurtosis of conditional returns - changing intraday volatility estimates.

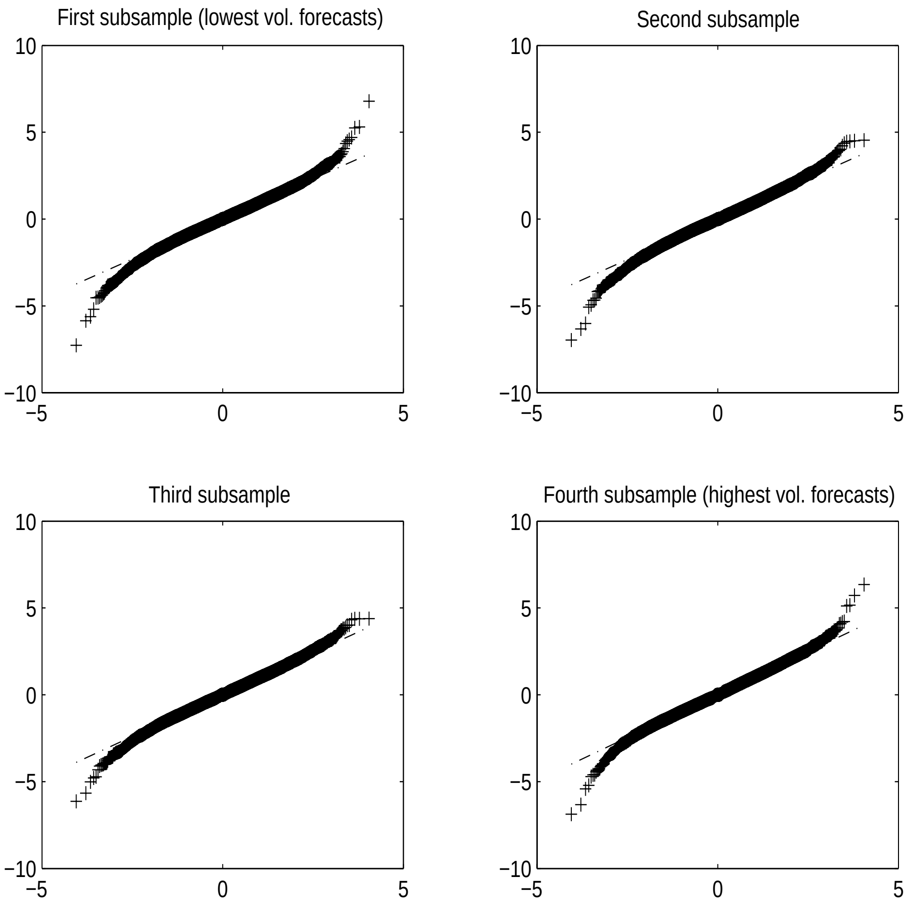

Figure 12: Normal QQ plots - return density forecasts with changing intraday volatility without jumps.

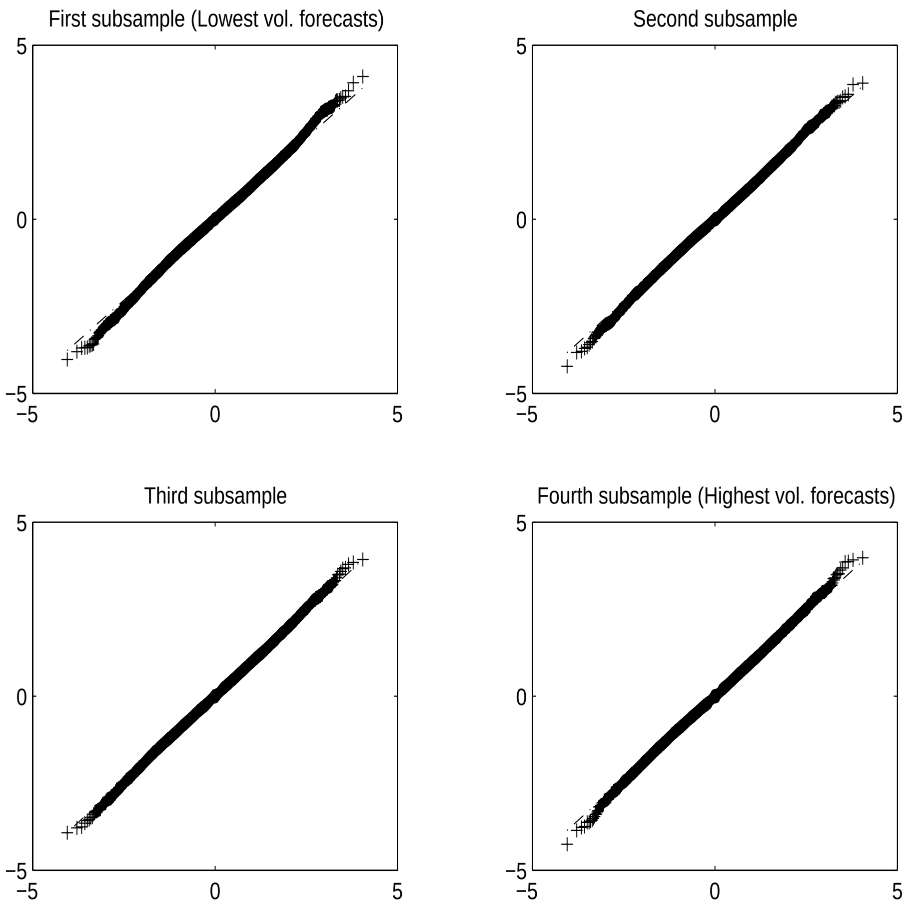

Figure 13: Normal QQ plots - return density forecasts with changing intraday volatility with jumps.

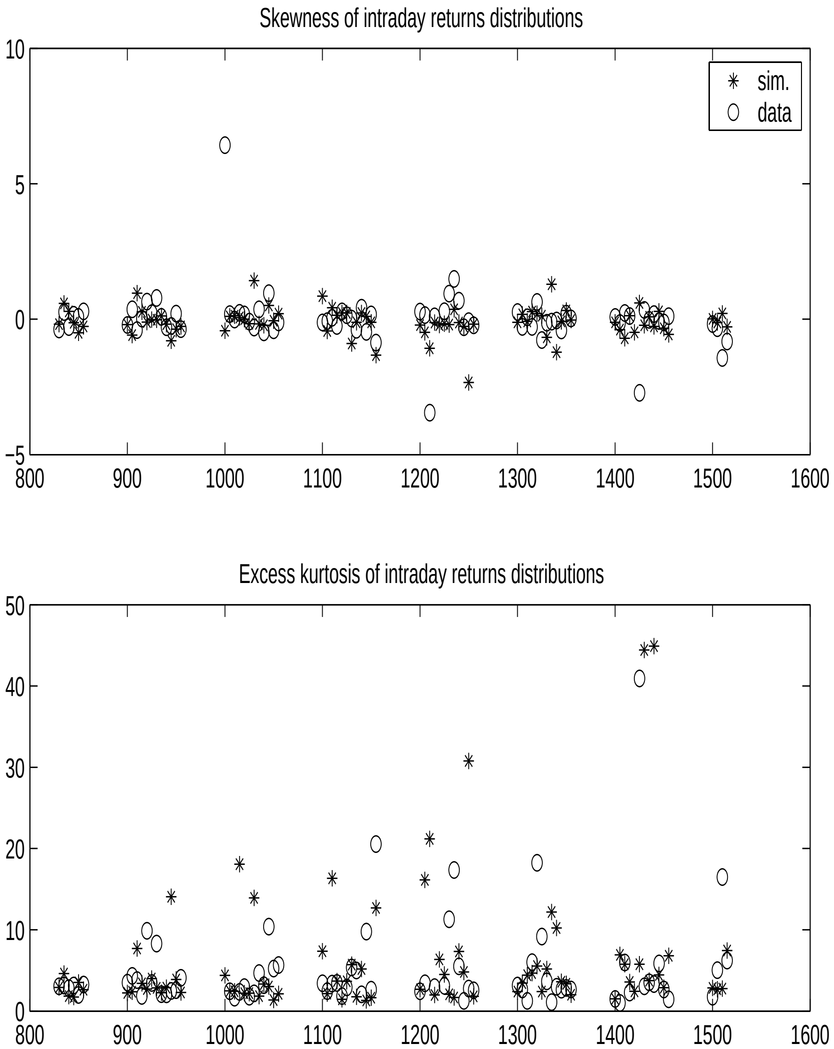

Figure 14: Skewness and excess kurtosis for intraday simulated returns.

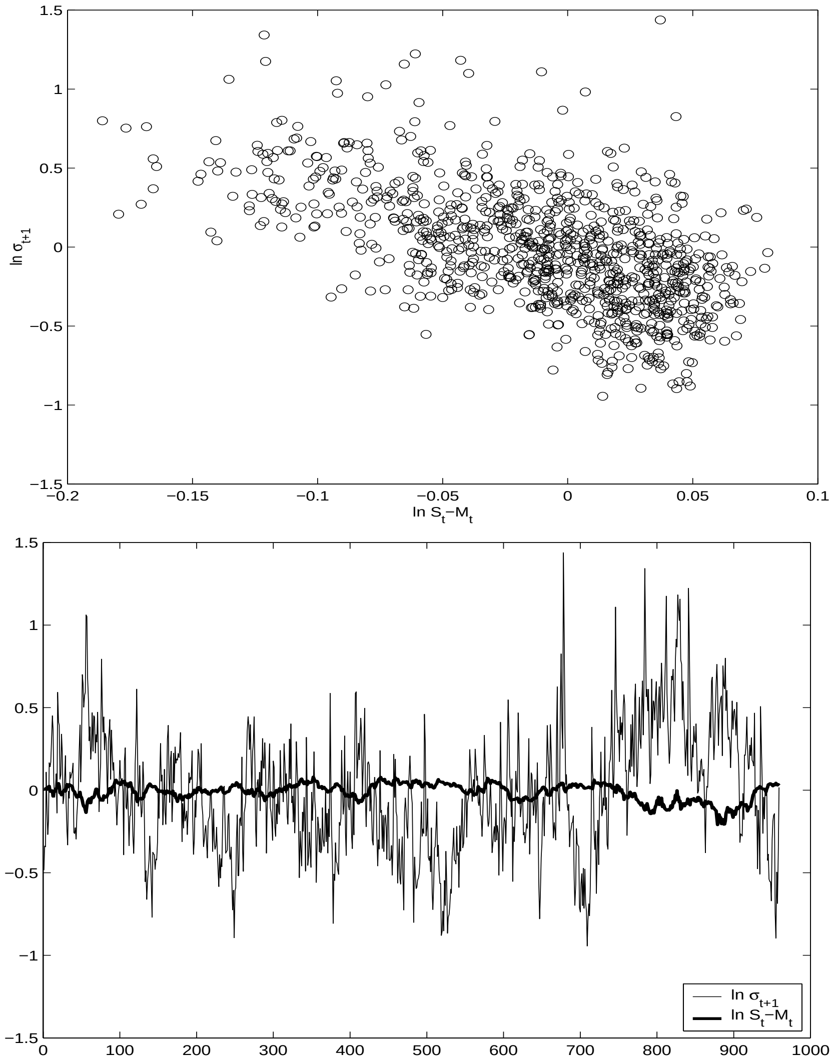

Figure 15: Scatter plot and time series of leverage measure against log volatility proxy - Simulated sample.

References

[1] *Brownian motion and stochastic calculus.* Springer Verlag, 1988.

[2] *An Introduction to High-Frequency Finance.* Academic Press: New York, 2001.

[3] S. Alizadeh, M.W. Brandt, and F.X. Diebold. Range-based estimation of stochastic volatility models. *Journal of Finance*, 2002.

[4] K.I. Amin. Jump diffusion option valuation in discrete time. *Journal of Finance*, 1993.

[5] T.G. Andersen. Return volatility and trading volume: an information flow interpretation of stochastic volatility. *Journal of Finance*, 1996.

[6] T.G. Andersen and T. Bollerslev. Answering the skeptics: yes, standard volatility models do provide accurate forecasts. *International Economic Review*, 1998.

[7] T.G. Andersen, T. Bollerslev, and F.X. Diebold. Parametric and nonparametric measurements of volatility. In *Handbook of Financial Econometrics*.

[8] T.G. Andersen, T. Bollerslev, F.X. Diebold, and H. Ebens. The distribution of realized stock return volatility. *Journal of Financial Economics*, 2001.

[9] T.G. Andersen, T. Bollerslev, F.X. Diebold, and P. Labys. The distribution of exchange rate volatility. *Journal of the American Statistical Association*, 2001.

[10] T.G. Andersen, T. Bollerslev, F.X. Diebold, and P. Labys. Modeling and forecasting realized volatility. *Econometrica*, 2003.

[11] T.G. Andersen, T. Bollerslev, and S. Lange. Forecasting financial market volatility: sample frequency vis-à-vis forecast horizon. *Journal of Empirical Finance*, 1999.

[12] E. Andreou and E. Ghysels. Rolling-sampling volatility estimators: some new theoretical, simulation and empirical results. *Journal of Business and Economic Statistics*, 2002.

[13] T. Ané and H. Geman. Order flow, transaction clock and normality of asset returns. *Journal of Finance*, 2000.

[14] N. Areal and S.J. Taylor. The realised volatility of FTSE-100 futures prices. *Journal of Futures Markets*, 2002.

[15] X. Bai, J. Russell, and G.C. Tiao. Beyond Merton's utopia: effects of non-normality and dependence on the precision of variance estimates using high frequency financial data. Working Paper, University of Chicago.

[16] O.E. Barndorff-Nielsen, E. Nicolato, and N. Shephard. Some recent developments in stochastic volatility modelling. *Quantitative Finance*, 2002.

[17] O.E. Barndorff-Nielsen and N. Shephard. Non-Gaussian OU based models and some of their uses in financial economics. *Journal of the Royal Statistical Society B*, 2001a.

[18] O.E. Barndorff-Nielsen and N. Shephard. Estimating quadratic variation using realised variance. *Journal of Applied Econometrics*, 2001b.

[19] O.E. Barndorff-Nielsen and N. Shephard. Econometric analysis of realised volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society B*, 2002.

[20] O.E. Barndorff-Nielsen and N. Shephard. Realised power variation and stochastic volatility. *Bernoulli*, 2003.

[21] D. Bates. Post-87 crash fears in the S&P500 futures options. *Journal of Econometrics*, 2000.

[22] J. Berkowitz. Testing density forecasts with applications to risk management. *Journal of Business and Economic Statistics*, 2001.

[23] F. Black. Studies of stock price volatility changes.

[24] B. Blair, S. Poon, and S.J. Taylor. Forecasting S&P100 volatility: the incremental information content of implied volatilities and high frequency returns. *Journal of Econometrics*, 2001.

[25] T. Bollerslev and H. Zhou. Estimating stochastic volatility diffusion using conditional moments of integrated volatility. *Journal of Econometrics*, 2002.

[26] P. Carr, H. Geman, D. Madan, and M. Yor. Stochastic volatility for Lévy processes. *Mathematical Finance*, 2003.

[27] P. Carr and L. Wu. Time-changed Lévy processes and option pricing. *Journal of Financial Economics*, 2003.

[28] S. Chib, F. Nardari, and N. Shephard. Markov Chain Monte Carlo methods for stochastic volatility models. *Journal of Econometrics*, 2002.

[29] P.K. Clark. A subordinated stochastic process with finite variance for speculative prices. *Econometrica*, 1973.

[30] R.A. Davis and T. Mikosch. The sample autocorrelations of financial time series models. In *Nonlinear and Nonstationary Signal Processing.*

[31] D. Duffie, K. Singleton, and J. Pan. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, 2000.

[32] G. Durham and R. Gallant. Numerical techniques for maximum likelihood estimation of continuous-time diffusion processes. *Journal of Business and Economic Statistics*, 2002.

[33] H. Ebens. Realized stock volatility. Working Paper, John Hopkins University.

[34] O. Elerian, S. Chib, and N. Shephard. Likelihood inference for discretely observed non-linear diffusions. *Econometrica*, 2001.

[35] B. Eraker. MCMC analysis of diffusion models with applications to finance. *Journal of Business and Economic Statistics*, 2001.

[36] B. Eraker, M. Johannes, and N. Polson. The impact of jumps in volatility and returns. *Journal of Finance*, 2003.

[37] R. Gallant, C.T. Hsu, and G. Tauchen. Using daily range data to calibrate volatility diffusions and extract the forward integrated variance. *The Review of Economics and Statistics*, 1999.

[38] H. Geman, D. Madan, and M. Yor. Time changes for Lévy processes. *Mathematical Finance*, 2001.

[39] P. Giot. Time transformations, intraday data and volatility models. *Journal of Computational Finance*, 2000.

[40] C.W.J. Granger and S. Poon. Forecasting volatility in financial markets. *Journal of Economic Literature*, 2003.

[41] L. Harris. A transaction data study of weekly and intradaily patterns in stock returns. *Journal of Financial Economics*, 1986.

[42] E. Hol and S.J. Koopman. Stock index volatility forecasting with high frequency data. Tinbergen Institute Discussion Papers.

[43] M. Johannes, N. Polson, and J. Stroud. Nonlinear filtering of stochastic differential equations with jumps. Working paper, Columbia University.

[44] N.M. Kiefer. Discrete parameter variation: efficient estimation of a switching regression model. *Econometrica*, 1978.

[45] M. Martens. Forecasting daily exchange rate volatility using intraday returns. *Journal of International Money and Finance*, 2001.

[46] M. Martens, Y.C. Chang, and S.J. Taylor. A comparison of seasonal adjustment methods when forecasting intraday volatility. *The Journal of Financial Research*, 2002.

[47] N. Meddahi. A theoretical comparison between integrated and realized volatility. *Journal of Applied Econometrics*, 2002.

[48] I. Monroe. Processes that can be embedded in Brownian motion. *Annals of Probability*, 1978.

[49] J. Pan. The jump-risk premia implicit in options: evidence from an integrated time series study. *Journal of Financial Economics*, 2002.

[50] M.K. Pitt and N. Shephard. Filtering via simulation: auxiliary particle filter. *Journal of the American Statistical Association*, 1999.

[51] L. Rogers and O. Zane. Designing and estimating models of high-frequency data. Working paper, University of Bath.

[52] R. Roll. A simple implicit measure of the effective bid-ask spread in an efficient market. *Journal of Finance*, 1984.

[53] T.H. Rydberg and N. Shephard. A modelling framework for the prices and times of trades made on the NYSE. In *Nonlinear and Nonstationary Signal Processing*.

[54] T.H. Rydberg and N. Shephard. Dynamics of trade-by-trade price movements: decomposition and models. *Journal of Financial Econometrics*, 2003.

[55] K. Singleton. Estimation of affine asset pricing models using the empirical characteristic function. *Journal of Econometrics*, 2001.

[56] S.J. Taylor and X. Xu. The incremental volatility information in one million foreign exchange quotations. *Journal of Empirical Finance*, 1997.

[57] R.A. Wood, T.H. McInish, and J.K. Ord. An investigation of transaction data for NYSE stocks. *Journal of Finance*, 1985.