A Parsimonious Continuous Time Model Of Equity Index Returns (Inferred From High Frequency Data)

Mascia Bedendo and Stewart D. Hodges
Financial Options Research Centre, University of Warwick
Coventry, CV4 7AL, United Kingdom

November 11, 2004

Abstract

In this paper we propose a continuous time model capable of describing the dynamics of futures equity index 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 it from the analysis of a data set of 5-minute returns on the S&P500 futures contract. We start with a very general specification. First we model the seasonal pattern in intraday volatility. Once we correct for this component, we aggregate intraday data into a daily volatility measure to reduce the amount of noise and its distorting impact on the results. We then employ this 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. The results from a Monte Carlo experiment indicate that a sample of returns simulated according to our model successfully replicates the main features observed in market returns.

Keywords: High frequency data; continuous time models; non-linear filtering.[1]

[1]International Journal of Theoretical and Applied Finance.

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*. 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 market returns in a consistent manner across different time horizons. Traditionally, continuous time models in finance have been estimated and tested on moderate frequency (normally daily) data. However, the asset returns generated from those models often manage to capture the dynamics of daily or weekly returns fairly accurately, but fail to mirror the behavior 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 area of research. 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 characterize the actual evolution of intraday returns and volatility. We will find that 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. 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 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 statistic and econometric techniques employed to estimate the model. Our main interest is to assess whether the data-driven, step-by-step criteria we propose for selecting the model (and subsequently refining it on the basis of intraday returns) enable 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, as that would introduce a lot of complexity to the analysis without contributing significantly to the main results. Therefore we use 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 relevant 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. The assessment of both the volatility and the density intraday estimates produced by our model is carried out in Section 5. A Monte Carlo simulation exercise of the complete model is performed in Section 6. Section 7 summarizes the main findings.

2 The Informative Content of Intraday 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 (not immediately related to the present work) has focused on deriving tailored models for intraday returns and volatility, capable of capturing their distinctive features. Models for the dynamics of transaction prices have been suggested, amongst the others, by [36].

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 realized volatility (see [38], [8]). The theoretical justification for this approximation ([11], [14]) is to be found in the theory of quadratic variation (see [29]). A complete asymptotic theory of the convergence of the realized volatility to the integrated volatility was derived by [15], [16], [17], under the assumptions that conditional returns are normally distributed and volatility follows either a diffusion specification or a Lévy process. They also considered extensions to account for the presence of a leverage effect. [12] and [31] discussed potential distortions and biases in the realized volatility measure.

An impressive number of papers have appeared in the last couple of years in this area, proposing various possible applications for the realized volatility measure. For a survey of this literature, the interested reader can consult [9] and [23].

The distribution and the time series properties of the realized volatility have also been studied. Examples in this context are given by [11] for exchange rates and [10] for both the index and its constituent stocks.

A third stream of literature 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. The stochastic time change is given by a measure of the intraday economic activity (e.g. trading volumes, proxy of integrated stochastic volatility). The underlying theoretical justification is that all arbitrage-free processes defining asset returns can be represented as time-changed Brownian motions,[2] where the time change (or business time) must account for information arrival and market activity. This stream of research originated from the pioneering paper by [22], 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 [6], [13] and, more recently, [28], and [19]. In relation to this area, we recall some relevant contributions of the "econophysics" literature, directed at analyzing the scaling behavior of the distribution of normalized returns over different time horizons (see [26], [3], [4]).

Our work directly relates to the second and the third streams of research. At an intermediate stage of our analysis, we derive a daily realized volatility quantity from intraday data, in order to obtain an almost noise-free measure which can provide reliable insights on the stochastic volatility dynamics and on the shape of the conditional distributions. Similarly to the literature on stochastic time changes, our purpose is to estimate a valid continuous time specification from high frequency data. However, contrary to the contributions listed above, we will find that, after correcting for the seasonality in volatility, no stochastic time change is necessary and the model can be set in calendar time.

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 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 analyze overnight log returns, calculated as the difference between the logarithm of the open price

[2]See [32] for the proof that any semimartingale is a time-changed Brownian motion.

and the logarithm of the closing price for the previous day. For the same reason, unlike most works on high frequency data, we retain the return on the first interval of the trading day, which mainly reflects the information accumulated overnight and shows a high volatility.

In 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 and a strong sample excess kurtosis. 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, by effect of the aggregation process which takes place over a longer time horizon.

Table 1: Summary sample statistics for 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

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 [35].[3] Although statistically significant for the first two and a half years, the serial correlation in the intraday futures returns always seems economically negligible. To ascertain that, we compute the bias in the variance induced by ignoring first and second order serial correlation, which is, respectively, -0.318% and -0.319% of the correct variance.[4] Therefore, it does not make any substantial difference if we remove the autocorrelation from our series or not. The percentage bid-ask spread is consistently small, around 0.06%. Our findings suggest that 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

Throughout the present section we develop a careful step by step analysis of the data, aimed at isolating its main distinctive features and their nature, and,

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

[4] Our measure for the bias in the variance is 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.

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.

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; 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;[5] ε_{it} symbolizes the conditional intraday log return, with zero mean, independent of both the seasonal and the stochastic volatility parts. Once both the seasonal and stochastic volatility components have been correctly modelled, ε_{it} should translate into a series which is independent across the intraday intervals.

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 these very general assumptions may easily be too strong. However, we do not take our assumptions for granted, but we attempt to test their validity as much as possible, as part of our data analysis. Also, it is worth emphasizing that 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.

A note on the terminology that we use and on the scaling of the model in (1). After ascertaining the deterministic nature of the seasonal volatility component, we choose to work with de-seasonalized returns, which involves scaling our model so that $E[\sigma_{it}] = 1$ and $E[|\varepsilon_{it}|] = 1$. Also, 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, [17]), absolute returns are less sensitive to large outliers and more reliable when the fourth moment of the distribution of returns is not finite. Finally, when we talk about 5-minute or intraday intervals, we also refer to the overnight interval, unless otherwise stated.

[5] 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.

In the following subsections we proceed to investigate the empirical features of the various components in (1).

4.1 The seasonal component

The plot of the average absolute returns (Fig. 1), computed across the time series of the single 5-minute intervals,[6] reveals an obvious U-pattern in intraday volatility, which was first documented by [39]. The average absolute returns keep on increasing for the first half an hour of the trading day, decline smoothly to their lowest level before noon and then increase again until the closure of the cash index market. The final spike in the last 15 minutes of the trading day is attributable to the post cash market trading.

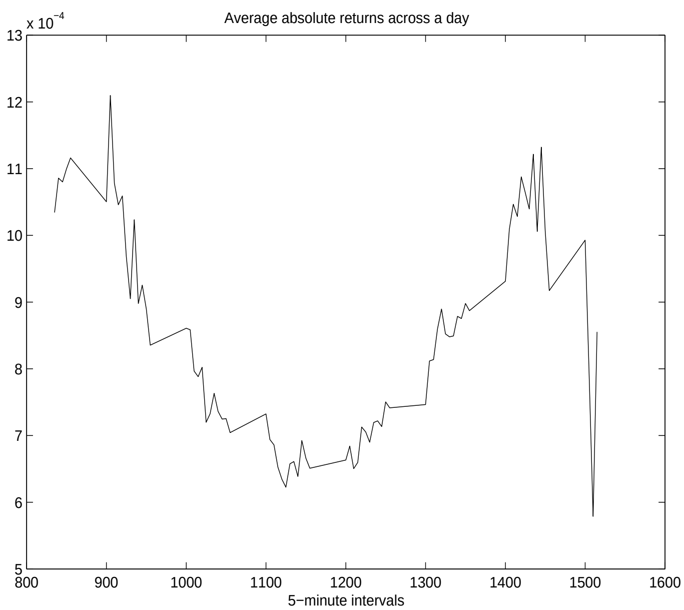

Figure 1: Seasonality pattern in intraday volatility.

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 conduct 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 perform, for each intraday interval, a two sample Student's t-test for mean equality (at 95% confidence level) on the average absolute returns of the two

[6] Overnight returns have not been included in the plot, since their average absolute values are not in line with the rest of the intraday data, and their inclusion would have distorted the analysis.

sub-samples we want to compare.[7] 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 derive the series of intraday ratios computed on the average normalized absolute returns of the two sub-samples of interest; average values and standard deviations for these series are reported in Table 3. Both the small percentages of rejections for the mean equality test (ranging from 0 to 12%) and the low 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. These findings are reinforced by the evidence of very

Table 3: Tests for equality of seasonal volatility patterns.

Sub-samples	% 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

similar shapes displayed by the seasonal patterns in intraday volatility for the various sub-samples (not reported here, but available upon request).

Therefore, over the time period spanned by our data we can safely assume a constant deterministic intraday seasonal pattern, represented by s_i.

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. The average absolute returns for the individual sub-intervals constitute simple estimates of the intraday seasonal component s_i, both for the 5-minute intervals and for the overnight returns. Different approaches have been proposed in the literature to obtain smoothed estimates of these seasonal coefficients (see, for example, the Flexible Fourier functions recommended by [7]). In Fig. 2 we show how good smoothed estimates can be easily obtained by fitting a set of cubic B-splines to the average absolute returns for the 5-minute intervals.[8]

Having obtained accurate estimates for the deterministic seasonal pattern in intraday volatility, we can now derive the time series of de-seasonalized un-

[7] To allow a comparison amongst subsets of data with different levels of volatility, we compute absolute returns normalized by the average of absolute returns across the day, taken as a volatility proxy for the day.

[8] 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.

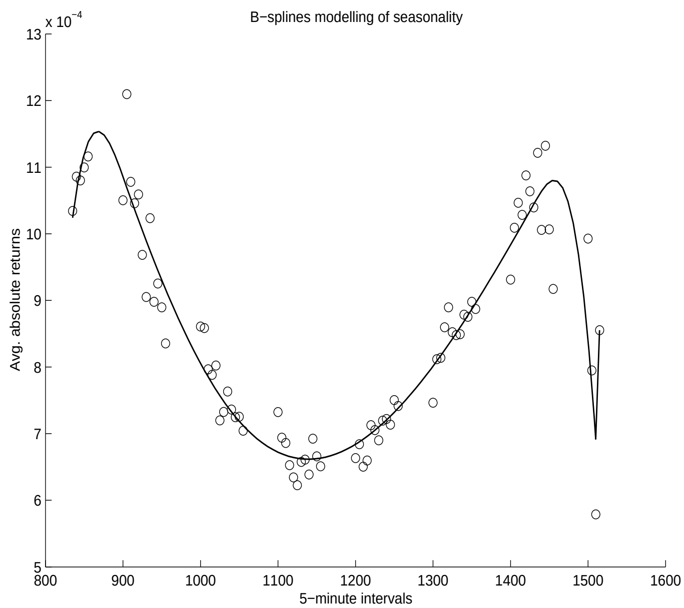

Figure 2: Smoothed B-spline estimation of seasonal coefficients.

conditional intraday returns, by dividing the 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-seasonalized unconditional overnight returns.

4.2 The stochastic volatility component: analysis of daily volatility estimates

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-seasonalized intraday return is a combination of independent realizations from a stochastic volatility process and from a conditional density. Therefore, the next step is to identify an appropriate stochastic volatility process capable of generating good intraday volatility estimates.

To investigate the presence and the nature of the stochastic volatility component, we start by plotting the autocorrelogram of the absolute de-seasonalized 5-minute returns for $4,100$ lags, corresponding to 50 days (Fig. 3). 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.

Our target is to model stochastic volatility at a high frequency level. 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. We

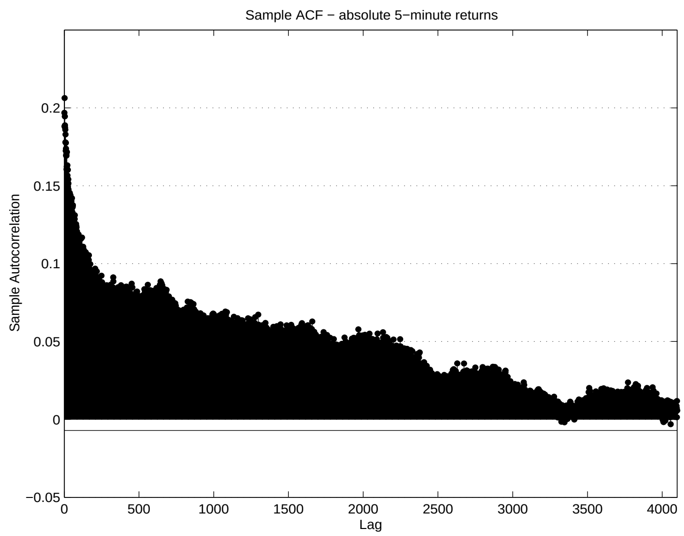

Figure 3: ACF of intraday absolute returns.

can eliminate most of this noise by working initially with daily averages of the intraday absolute returns, in order to understand the low frequency component of the volatility dynamics. The daily average of the 82 5-minute absolute unconditional returns is computed as 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 realized volatility measures mentioned earlier.

The model estimated on these volatility proxies will yield good estimates of daily volatility. Only if volatility were constant across the 5-minute subintervals of a same day, such estimates would also be accurate at an intraday level. We know that this is not the case, therefore at a high frequency level these estimates will show inaccuracies due to both some measurement error and intraday changes in the volatility. Instead, our daily volatility estimates will prove to be very useful at revealing most of the structure that our stochastic volatility model should possess in order to capture the essential features of the volatility dynamics, including the impact of the leverage component. Moreover, the daily estimates can be considered good enough for the purpose of normalizing the series of unconditional intraday returns and, consequently, providing more precise information on the shape of the conditional distributions.

The autocorrelogram for the daily average of absolute high frequency returns up to lag 50 is displayed in Fig. 4, 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

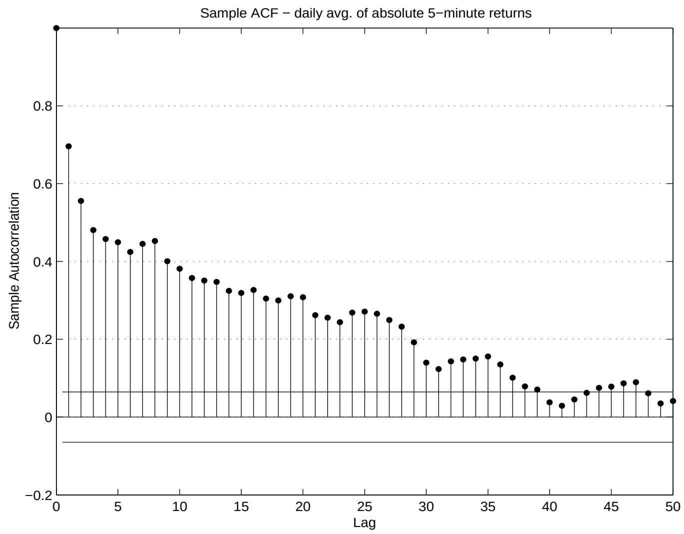

Figure 4: ACF of daily average absolute returns.

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 autocorrelograms at both daily and 5-minute level reveal 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.

As we will see later, the statistical techniques that we use to estimate a stochastic volatility model provide more accurate results the closer the series to be modelled is to a normal. Therefore, we choose 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.

4.3 The leverage effect

The leverage effect was first discussed by [18], who observed how 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. In order to investigate presence and magnitude of the leverage effect, we propose a specification which is very easy to estimate and well supported by our data.

The distance between the current level of the index and its moving average quite naturally represents a new, stationary, variable, which is intuitively related to the volatility of the index. The average index level is computed as an exponentially weighted moving average M of closing log prices for the S&P500 stock index futures: $M_t = (1-\theta\Delta t)M_{t-\Delta t} + \theta \ln(S_{t-\Delta t})\Delta t$. The new stationary series is derived 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 for the following period $\ln(\hat{\sigma}_{t+\Delta t})$.

To quantify the leverage effect for variables measured at a daily level ($\Delta t = 1$) we proceed as follows. 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 maximizes (in absolute terms) the correlation between the series of daily log price movements and daily log volatility proxy. For this parametrization, we obtain a correlation of $\rho = -0.545$ between the two series, which confirms the existence of a strong leverage effect.

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:

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

The regression in (3), performed on daily measures, provides us with an estimate for κ of -4.34 (standard error 0.26) and with time series of daily residuals υ_t, whose evolution mirrors 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 υ_t could be adequately modelled by means of an AR(1) specification. Unfortunately, the ACF of the residuals from the AR(1) process (not reported here, but available upon request) indicates that this simple and appealing specification does not capture all the dynamics of the process. More complete specifications are then needed in order to achieve a satisfactory model for the stochastic volatility component.

4.4 A short memory model for the stochastic volatility

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 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, perhaps because we analyze only 4 years of data. [27] provided an alternative explanation 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) provides a sufficiently accurate description of the empirical results, while maintaining the nice properties of a short memory process.

We therefore explore the use of a model similar to [1], and represent log volatility in continuous time as the sum of two independent Ornstein-Uhlenbeck processes, each of them mean reverting towards the long run level of zero, given that $E[\sigma_t] = 1$.

Since we employ the residuals from Eq. (3) as our (ex-leverage) log volatility measure, the discrete time version of the model for the dynamics of the log volatility becomes:

$$
\begin{aligned}
\upsilon_t &= \ln(\sigma_t) + \xi_t \\
\ln(\sigma_t) &= \ln(\sigma_{s,t}) + \ln(\sigma_{l,t}) \\
\ln(\sigma_{s,t}) &= \rho_s \ln(\sigma_{s,t-\Delta t}) + \beta_s \sqrt{\Delta t}\omega_{s,t} \\
\ln(\sigma_{l,t}) &= \rho_l \ln(\sigma_{l,t-\Delta t}) + \beta_l \sqrt{\Delta t}\omega_{l,t}
\end{aligned}
\tag{4}
$$

The two log volatility components follow a Gaussian first-order autoregressive process with mean zero, autoregressive parameter $\rho_j = 1 - \alpha_j \Delta t$ (with mean reversion parameter α_j) and variance $\beta_j^2 \Delta t$. $\Delta t = 1$,[9] since we estimate the model on daily volatility proxies, whose measurement error ξ_t is small and can be easily bounded.[10]

The estimation is carried out by applying a Kalman filter algorithm to the state space system in (4). The estimated parameters, with standard errors in brackets, are displayed in Table 4.

We can clearly identify a transient volatility component, with $\alpha_s = 0.739$ corresponding to a half life of 0.94 days and a more persistent one, with $\alpha_l = 0.018$ and half life of approximately 37.5 days. Most of the short-run variance of the model can be attributed to the transient component, whereas 52% of the unconditional long-run variance is explained by the more persistent component.

[9] 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.

[10] 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 38.5% 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 5.3% of the total variance for the log volatility proxy on the unconditional returns.

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

The dynamics of the volatility has now been correctly captured.

In our case, the distributions of the residuals from both the state equations of the two components and the measurement equation exhibit little positive skewness (around 0.35) and excess kurtosis (around 0.4). 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.

4.5 Some insights on the conditional return densities

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 $\ln(\hat{\sigma})$. Although inaccurate as 5-minute volatility estimates, these constant intraday volatility measures 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 normalizing the unconditional de-seasonalized return, r_{it}/s_i, by the volatility estimate for day t made at the end of the previous day.

Once the volatility dynamics has been accurately modelled, if the conditional return distribution is Gaussian and independent from the volatility process, then the conditional intraday returns will 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) would be discernible. 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 $i = 1, \ldots, m$.

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. First we investigate how these sample statistics relate to the theoretical ones from a normal distribution, and then we discuss their stability across the 5-minute subintervals. We plot in Fig. 5 (top) the standard deviation of the time series of the conditional returns for the individual intervals, together with a straight horizontal line at $\sqrt{\pi/2}$, which represents the theoretical level of standard deviation under the assumption of normality for the distribution of conditional returns.[11] We can clearly detect a few spikes for some intraday intervals, that seem to suggest the fat-tailed nature of the conditional distribution. However, the spikes are mainly attributable to a very small number of outliers (around 15 for the whole data set, 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 fairly close to the theoretical $\sqrt{\pi/2}$ level.

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 individual intraday periods is reported in Fig. 5 (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. 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.[12]

Once we find evidence of near-Gaussian 5-minute conditional return densities, we focus on the stability of the shape of such distributions across the intraday intervals. The comparison amongst the values of standard deviation and excess kurtosis recorded for the individual subintervals reveals that the main source of variability in the shapes occurs as an effect of the few outliers already discussed above. Except for that, on average the values of the summary sample statistics computed for the intraday conditional distributions are quite flat across the 5-minute intervals of the day.

To summarize, the distribution of conditional returns computed by normalizing upon constant intraday volatility forecasts turns out to be surprisingly close to Gaussian and virtually the same across the different 5-minute intervals of the day. 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 subsection.

4.6 The estimation of the model at an intraday level

The model calibrated on a daily basis cannot accurately describe the actual dynamics of high frequency data. Therefore, the estimation of our continuous time specification must be refined by exploiting the information content of the

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

[12] The average skewness of the conditional returns for each of the intraday intervals has also been computed. However, its analysis is not particularly informative, given that, apart from very few exceptions, skewness coefficients do not depart significantly from zero.

5-minute return series. For this purpose we employ 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 seems to be supported by the analysis 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. 6. The fact that the two segments for $k \leq 80$ are different, with a higher serial correlation within the same day, indicates that the overnight period has a significant impact on the mean reversion of the model, an evidence in favor of a continuous time specification.

In deriving high frequency estimates, we start by dismissing the persistent component of the stochastic volatility process in (4), 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 modelled at a lower frequency level. To justify our choices, we have computed the proportions of the variance of log volatility 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.

First, we consider a standard diffusion model for the 5-minute volatility process. We will then introduce jumps in our volatility specification, which will significantly improve our return density estimates.

In order to describe the intraday volatility dynamics, we maintain the standard Gaussian Ornstein-Uhlenbeck specification employed in the daily model to characterize 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 avoid imposing strong structural assumptions, high frequency volatility estimates 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:[13]

$$P^*_{i,t} \approx \sum_{j=1}^{N} P^a_{(i,t),(j,t-\Delta t)} P_{j,t-\Delta t} \tag{5}$$

which will then be applied to the corresponding volatility values in the range in order to return the intraday variance estimate ${\sigma^*}^2 = \sum_{i=1}^{N} P^*_{i,t}\sigma^2_i$. Within

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

this framework, the discretization 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 the standard way, by equating the first two moments of the continuous process to those of its discretization.[14]

Under the assumption of normality for the conditional returns, justified on the basis of the results derived earlier, a density forecast for the unconditional de-seasonalized returns r_t^f is represented by a mixture of normal densities, where each component is a normal with zero mean 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:

$$r_t^f \sim \sum_{i=1}^{N} P_{i,t}^* N(0,\, \sigma_i \sqrt{\pi/2}) \tag{6}$$

Once the 5-minute unconditional de-seasonalized return r_t/s_t 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 $P_{i,t}^p$:

$$P_{i,t}^p \approx \frac{P_{i,t}^r P_{i,t}^*}{\sum_{j=1}^{N} P_{j,t}^r P_{j,t}^*} \tag{7}$$

which will replace the initial probabilities P_i in order to re-start the process. Volatility and return density forecasts are then 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, during which 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. Since the variance of the unconditional overnight returns is about 25 times the variance of the corresponding 5-minute returns, this choice is equivalent to expressing time in calendar terms during the trading day and in volatility terms overnight.[15]

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 realizations of our non-linear filtering model, given by:

[14] For the practical implementation of the model, we choose a log volatility range between −1.5 and 1.5, with step size equal to 0.1, roughly equal 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.

[15] We have also attempted to estimate our model in calendar time only and in volatility time only, by applying the same measure to all subintervals, but such alternative specifications have been rejected by the data.

$$L(r_T) \approx \prod_{t=1}^{T} \left(\sum_{j=1}^{N} P^r_{j,t} P^*_{j,t} \right) \tag{8}$$

is maximized 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.[16] 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 between complexity and accuracy) we have found that estimated values of $\alpha_s = 0.6$ and $\beta_s = 0.28$ meet these requirements.

We have employed 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 recently developed in the literature: auxiliary particle filtering techniques for volatility filtering (see [34], [21]), Markov Chain Monte Carlo (MCMC) methods ([20], [25]) and GMM procedures ([37]) for parameter estimation of a variety of diffusion and jump-diffusion processes. 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 the best possible estimation accuracy is not our main concern.

Once the model is fully parametrized, 5-minute volatility estimates and return density estimates can be extracted. The time series of conditional returns is now obtained by normalizing 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 (Fig. 7, top) and the excess kurtosis (bottom) across the individual subintervals. A few spikes due to the presence of very large outliers, rather than to the effect of some external source of information not captured by our model, can still be easily detected. If we ignore these outliers, we observe values for both the standard deviation and the excess kurtosis 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. In accordance with the hypothesis of a Gaussian specification, the summary sample statistics for the conditional return densities also exhibit very similar values across all the intervals of the trading day.

These findings are entirely in line with our expectations: volatility estimates updated at a high frequency level can account for most of the fat-tailness left in the conditional return density after normalizing upon the volatility estimates which remain constant across the day. The assumption of a Gaussian conditional

[16] The likelihood function for mixture models is known to be unbounded at some points on the edge of the parameter space (see [30]). In our case, however, we do not attempt to maximize the likelihood *per se*, and we only use it to discriminate between various set of parameters that satisfy the volatility constraint.

return distribution turns out to work surprisingly well, actually better than expected.

At this stage, 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.

4.7 A process with jumps in intraday volatility

The introduction of jumps can take place in the return process, in the volatility process, or in both. To avoid arbitrary assumptions, we analyze the nature of the outliers (i.e. 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.

To investigate if the increased volatility consequent to a jump is persistent, we run the regression $(\mid r_{t+\Delta t}/s_{t+\Delta t} \mid - \sigma^*_{t+\Delta t}) = a + b(\mid r_t/s_t \mid - \mid \overline{r/s} \mid)$ at a 5-minute level on both the entire sample and the sub-sample 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. Also, the temporal distribution of the outliers highlights a significant clustering in the incidence of jumps, which again contradicts the i.i.d. assumption of the jumps in returns. Our empirical results indicate that the outliers exhibits more the features of jumps in volatility than those of jumps in returns. 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 [33], [24]). 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.

We then need to ascertain how persistent the impact of jumps on the intraday volatility is, in order to decide to which volatility process (transient or permanent) the jumps should be added. In Fig. 8 we plot the average (computed across all outliers) difference between post-jump volatility levels and the average 5-minute volatility level across the 82 subintervals preceding the jump.[17] We do that for several post-jump intervals, ranging from five minutes to three day after the outlier has occurred. The rapid decay in the volatility difference suggests that the inclusion of the jumps can be safely restricted to the transient volatility component.

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

$$d\ln(\sigma_{s,t}) = -\alpha_s dt \ln(\sigma_{s,t}) + \beta_s \sqrt{dt} dW_{s,t} + \sum_{i=1}^{N(t)} Y_i - \lambda dt E[Y] \qquad (9)$$

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

[17] The intraday volatility is approximated by 5-minute absolute returns.

jump magnitudes. A compensated jump process has been chosen to maintain the mean of the volatility process unchanged.

We choose to model the jumps in volatility such that the corresponding outliers in returns Y_i^r follow a power law distribution: $P(Y_i^r > x) \sim x^{-\zeta}$, for x above 3 standard deviations. This specification is in line with the findings of a rich literature (see, for instance, [3], [4]), which show that conditional financial returns over various time scales have a power law distribution. [2] have also modelled the fast-decaying component of the volatility with a power law specification. The estimate for ζ is obtained via calibration of the power law specification to the frequency at which empirical outliers occur, which yields a value of $\zeta = 4$. In order to fit the discrete version of the stochastic volatility model in (9) into our trinomial grid structure, we need to work with jumps of discrete size, expressed as a multiple of our step size Δ_r. The jump intensities are derived from the power law probabilities of the corresponding outliers in returns, originated from the jumps in volatility.

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 difference that the log volatility process evolving on the grid is now the mean reverting model augmented by the jumps component. The transition probabilities must be recomputed, following the procedure for discretized jump diffusion processes suggested by [5]. 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 are equal to $\alpha_s = 0.7$ and $\beta_s = 0.25$.

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

5 The appraisal of intraday volatility and density estimates

The assessment of our high frequency volatility and return density estimates is based on statistical techniques borrowed from both point and density forecast evaluation practice.

Point forecast evaluation techniques are used to assess the 5-minute volatility estimates, through a comparison with the absolute value of de-seasonalized high frequency returns, taken as a proxy of the actual intraday volatility level. In line with the existing literature, we first regress the absolute return on the volatility prediction, $|r_t/s_t| = \alpha + \beta\sigma_t^* + \epsilon_t$. 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 R^2 is large. However, the presence of a very noisy component in our volatility proxy induces very small R^2 coefficients *per se*: the regression performed on returns simulated from exact volatility forecasts yields an $R^2 = 0.2244$, which is indicative of the best we could expect to achieve.

We also report a standard Mean Absolute Deviation measure, determined as simple average of the absolute deviations of the volatility forecasts from the volatility proxies.

The findings from the point forecast evaluation of constant and changing 5-minute volatility estimates are displayed in Table 5.

Table 5: Intraday volatility estimates evaluation.

	Daily update	5-minute update	5-minute with jumps
Regression			
α	-0.0164*	-0.0487	-0.0382
	(0.0139)	(0.0083)	(0.0093)
β	1.0505	1.0416	1.0216
	(0.0158)	(0.0173)	(0.0195)
R^2	0.0895	0.1596	0.1698
MAD	0.6433	0.6364	0.6328

* not significantly different from zero at 5% level.

As expected, the results indicate a very poor forecasting performance in all cases, given the distortion induced 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 higher R^2 (0.16 against 0.09) of the regression, as well as a slightly smaller MAD (0.635 against 0.64).

Density forecast evaluation techniques are employed to assess the intraday density forecasts for the returns.[18] Following a standard procedure, from the sequence of 5-minutes ahead density forecasts $f_t(r)$, we derive the series of the probability integral transforms of the realized intraday returns as follows:

$$z_t = \int_{-\infty}^{r_t/s_t} f_t(r)\,\mathrm{d}r \tag{10}$$

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, follows an i.i.d. $N(0,1)$.

To guarantee more robust results against possible misspecifications of different type, several goodness-of-fit techniques have been implemented. The popular Kolmogorov-Smirnov and Watson statistics have been chosen to test

[18]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.

for uniformity. The normality is assessed via the Jarque-Bera test, 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). The results from the density forecast tests are reported in Table 6.

Table 6: Distributional forecast evaluation.

	Uniformity tests		**Normality tests**	**LR tests**	
	K-S (1.36)	Watson (0.19)	J-B (5.99)	LR1 (3.84)	LR2 (7.81)
Entire sample (78720 obs)					
Daily	3.76*	3.32*	13,995*	249.18*	1,418.0*
5-minute	3.67*	4.66*	2,271*	2.36	4.36
5-m. jumps	3.82*	5.69*	461*	0.98	6.54
4 Sub-samples					
Daily					
Low vol.	2.08*	1.09*	3,452*	51.78*	348.61*
Medium low	2.19*	1.32*	2,724*	62.80*	277.21*
Medium high	2.48*	0.75*	2,322*	68.76*	359.43*
High vol.	2.62*	0.67*	5,844*	67.04*	445.56*
5-minute					
Low vol.	2.79*	2.69*	975*	2.37	14.58*
Medium low	2.40*	1.96*	686*	1.06	5.16
Medium high	2.08*	1.04*	312*	2.11	5.54
High vol.	2.88*	0.50*	425*	1.41	11.91*
5-m. jumps					
Low vol.	1.62*	0.31*	75*	1.25	30.09*
Medium low	2.21*	0.50*	81*	1.89	9.33*
Medium high	2.33*	1.25*	142*	0.98	8.40*
High vol.	3.61*	1.81*	186*	1.51	8.80*

* rejected at 5% level.

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 sub-samples according to the level of the volatility forecast.

The null hypothesis that the return density forecasts represent accurate predictions of the actual distribution of the returns is generally rejected by all our goodness-of-fit statistics, for both constant and changing intraday volatility estimates. However, a substantial improvement in the forecasting performance is recorded 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 are now much closer to their critical values. The normal Q-Q plots for the case of changing volatility estimates (Fig.

9) display empirical quantiles fairly close to the normal ones, especially for the model with jumps. The fat-tailness induced by the jumps in volatility seems to correct for most of the misspecification in the tails recorded for both the daily updating method and the intraday method without jumps.

In the light of our findings, we can conclude that in order to correctly model the dynamics at an intraday frequency, our estimates must be updated every 5 minutes on the basis of the current value of 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 (4) which works well at 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 with a non-linear intraday filtering model).

6 A Monte Carlo Simulation Exercise

Throughout the previous sections 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 continuous time model for the intraday returns r_t as follows:[19]

$$\begin{aligned}
r_t &= s_i \sigma_t \varepsilon_t \\
dM_t &= \theta(\ln(S_t) - M_t)dt \\
\ln(\sigma_t) &= \kappa(\ln(S_{t-dt}) - M_{t-dt}) + \ln(\sigma_{s,t}) + \ln(\sigma_{l,t}) \\
d\ln(\sigma_{s,t}) &= -\alpha_s dt \ln(\sigma_{s,t}) + \beta_s \sqrt{dt} dW_{s,t} + \sum_{i=1}^{N(t)} Y_i - \lambda dt E[Y] \\
d\ln(\sigma_{l,t}) &= -\alpha_l dt \ln(\sigma_{l,t}) + \beta_l \sqrt{dt} dW_{l,t}
\end{aligned} \tag{11}$$

where:

1. s_i denotes the deterministic seasonal component of the intraday volatility, estimated by fitting smoothed cubic B-splines to the average absolute returns for the individual subintervals.

2. σ_t stands for the stochastic intraday volatility component, independent of s_i, whose dynamics is driven by:

    - a leverage effect component, as a proportion κ of the lagged distance between the (log) level of the index and its exponentially weighted moving average M. The values for the parameters, $\kappa = -4.34$ and $\theta = 0.03$, are derived at a daily level, as explained in Section 4.3;

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

- a slowly-decaying volatility component, $\sigma_{l,t}$, whose log is modelled as a Gaussian Ornstein-Uhlenbeck process. The estimates for the parameters, $\alpha_l = 0.018$ and $\beta_l = 0.044$, are those obtained on a daily basis in Section 4.4;
- a fast-decaying volatility component, $\sigma_{s,t}$, whose log is modelled as a compensated jump diffusion process, with Gaussian innovations independent of the innovations of the persistent volatility component. The jumps Y_i follow a power law distribution, with parameter $\zeta = 4$, obtained via calibration to the empirical frequency of the actual outliers. The parameter values $\alpha_s = 0.7$ and $\beta_s = 0.25$ are obtained as a result of the intraday non-linear filtering procedure described in Sections 4.6 and 4.7.

3. ε_t denotes the intraday conditional return, independent of both s_i and σ_t. Given our findings from the analysis of the shape of the conditional return distribution, realizations for ε_t are obtained by sampling from the Gaussian density $N(0, \sqrt{\pi/2})$.

In order to: 1) test whether the dynamics of the high frequency returns generated from our model in (11) 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 our model, parametrized as specified above.

For simplicity, we only simulate one full sample, whose properties will be compared to the empirical ones, with the purpose of verifying whether our data set could actually represent a random sample generated from the model. Four more samples are simulated to assess the estimation technique.

We start by looking at the plots of 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. 10). We then aggregate the simulated high frequency values to derive daily log volatility proxies as averages of absolute de-seasonalized returns, and daily measures of leverage. The time series of these daily simulated variables are contrasted with their daily empirical counterparts (Fig. 11, bottom). The dependence between log volatility and leverage component from simulated data has also been investigated, via scatter plot (Fig. 11, top) and computation of the correlation coefficient, equal to $\rho = -0.527$. 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 both high frequency and daily level, suggest that the model in (11) 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

Table 7: Estimates from simulated samples.

Samples	1	2	3	4	5	Avg.	Std.Dev.	Data
Daily model								
κ	-4.52	-4.20	-4.07	-4.66	-4.24	-4.338	0.244	-4.34
α_s	0.67	0.68	0.64	0.66	0.62	0.660	0.032	0.73
α_l	0.01	0.02	0.02	0.02	0.01	0.016	0.005	0.02
β_s	0.21	0.23	0.18	0.22	0.23	0.214	0.021	0.21
β_l	0.05	0.03	0.03	0.05	0.03	0.038	0.011	0.04
ρ	-0.53	-0.48	-0.43	-0.44	-0.51	-0.478	0.043	-0.54
Intraday model								
α_s	0.75	0.70	0.65	0.70	0.60	0.680	0.057	0.65
β_s	0.22	0.23	0.21	0.20	0.24	0.220	0.016	0.24
ζ	4.0	4.0	5.0	4.0	5.0	4.4	0.548	4.0

generating process. Following the steps of our data analysis, we start by investigating the seasonal component, whose pattern, for all five simulated samples, is indistinguishable from the one shown by the market data (results not reported here, but available on request). 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 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. The actual frequencies of the empirical outliers are employed to re-estimate the power-law parameter ζ. 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, as they are quite close to the actual parameters of the data generating process. However, we can detect an underestimate of the incidence of jumps, as well as of the mean reversion and volatility parameters of the diffusion component. On the whole, our findings suggest that the estimates produced by applying our simple techniques are quite reliable and adequate for our purposes.

7 Conclusions and Further Work

In the present work we have attempted to build a simple and accurate continuous time model capable of describing and replicating the dynamics of both high and moderate frequency index returns. In our approach, we have performed 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.

Let us briefly recap the stages that led us to the identification of the complete model. We have started by specifying a very general multiplicative structure for the model of 5-minute unconditional returns, as a function of a seasonal and a stochastic volatility component, and of intraday conditional returns. We have then examined the nature of the seasonal component in intraday volatility, which has proved to be deterministic. As for the analysis of the stochastic volatility component, given the large amount of noise present in high frequency data, we have derived much less noisy daily average measures of volatility, which represent a considerably more useful starting point for studying the volatility dynamics. On the basis of these daily volatility proxies we have first investigated the presence of a leverage effect in our data and devised a simple specification for its modelling. Following the evidence of the existence of both a transient and a persistent feature in the volatility, we have then explored how to model the ex-leverage volatility dynamics. A two-factor short memory volatility model has been successfully estimated on the daily volatility measures. The volatility estimates obtained from the daily model have also been employed to provide insights on the distribution of the conditional returns, which has turned out to be very close to a Gaussian and fairly stable across the various subintervals of the day. At this stage, the estimation of our model has been refined at an intraday level by exploiting the information content of the 5-minute return series, so as to obtain a specification capable of describing the dynamics of high frequency data. The fine-tuning of the model has been performed by means of a simple non-linear filtering technique, in which the estimates are updated every five minutes, following the observed intraday market returns. Finally, in order to account for the presence of some outliers and for the residual fat-tailness in the model, we have introduced jumps in our volatility specification.

An attractive feature of our work that would deserve further investigation is the possibility to obtain simplified versions of our general model specification, which will possess the correct properties for various specific time horizons of interest.

Acknowledgments

We are grateful to Roel Oomen, and Gianluca Fusai, as well as an anonymous referee and participants at EFMA 2003 and Financial Options Research Centre conferences for helpful comments and suggestions. We also thank FORC

corporate members for financial support.

References

[1] S. Alizadeh, M.W. Brandt and F.X. Diebold, Range-based estimation of stochastic volatility models, *Journal of Finance* **57** (2002) 1047–1091.

[2] L. Amaral, X. Gabaix, P. Gopikrishnan, V. Plerou and H.E. Stanley, Economic Fluctuations and Anomalous Diffusion, *Physical Review E* **62** (2000) 3023–3026.

[3] L. Amaral, P. Gopikrishnan, M. Meyer, V. Plerou and H.E. Stanley, Scaling of the distribution of fluctuations of financial market indices, *Physical Review E* **60** (1999) 5305–5316.

[4] L. Amaral, P. Gopikrishnan, M. Meyer, V. Plerou and H.E. Stanley, Scaling of the distribution of price fluctuations of individual companies, *Physical Review E* **60** (1999) 6519–6529.

[5] K.I. Amin, Jump diffusion option valuation in discrete time, *Journal of Finance* **48** (1993) 1833–1863.

[6] T.G. Andersen, Return volatility and trading volume: an information flow interpretation of stochastic volatility, *Journal of Finance* **51** (1996) 169–204.

[7] T.G. Andersen and T. Bollerslev, Intraday periodicity and volatility persistence in financial markets, *Journal of Empirical Finance* **4** (1997) 115–158.

[8] T.G. Andersen and T. Bollerslev, Answering the skeptics: yes, standard volatility models do provide accurate forecasts, *International Economic Review* **39** (1998) 885–905.

[9] T.G. Andersen, T. Bollerslev and F.X. Diebold, Parametric and nonparametric measurements of volatility, in *Handbook of Financial Econometrics*, Ed. Y. Ait-Sahalia and L.P. Hansen (2004), forthcoming.

[10] T.G. Andersen, T. Bollerslev, F.X. Diebold and H. Ebens, The distribution of realized stock return volatility, *Journal of Financial Economics* **61** (2001) 43–76.

[11] T.G. Andersen, T. Bollerslev, F.X. Diebold and P. Labys, Modeling and forecasting realized volatility, *Econometrica* **71** (2003) 579–626.

[12] E. Andreou and E. Ghysels, Rolling-sampling volatility estimators: some new theoretical, simulation and empirical results, *Journal of Business and Economic Statistics* **20** (2002) 363–376.

[13] T. Ané and H. Geman, Order flow, transaction clock and normality of asset returns, *Journal of Finance* **55** (2000) 2259–2284.

[14] 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* **63** (2001a) 167–241.

[15] O.E. Barndorff-Nielsen and N. Shephard, Estimating quadratic variation using realised variance *Journal of Applied Econometrics* **17** (2001b) 457–477.

[16] 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* **64** (2002) 253–286.

[17] O.E. Barndorff-Nielsen and N. Shephard, Realised power variation and stochastic volatility *Bernoulli* **9** (2003) 243–265.

[18] F. Black, Studies of stock price volatility change, in *Proceedings of the 1976 American Statistical Association*, Business and Economical Statistics Section, American Statistical Association (1976).

[19] P. Carr and L. Wu, Time-changed Lévy processes and Option Pricing, *Journal of Financial Economics* **17** (2004) 113–141.

[20] S. Chib, O. Elerian and N. Shephard, Likelihood inference for discretely observed non-linear diffusions, *Econometrica* **69** (2001) 959–993.

[21] S. Chib, F. Nardari and N. Shephard, Markov Chain Monte Carlo methods for stochastic volatility models, *Journal of Econometrics* **108** (2002) 281–316.

[22] P.K. Clark, A subordinated stochastic process with finite variance for speculative prices, *Econometrica* **41** (1973) 135–156.

[23] M.M. Dacorogna, R. Gençay, U. Muller, R.B. Olsen and O.V. Pictet, *An Introduction to High-Frequency Finance*, Academic Press: New York (2001).

[24] D. Duffie, K. Singleton and J. Pan, Transform analysis and asset pricing for affine jump-diffusions, *Econometrica* **68** (2000) 1343–1376.

[25] B. Eraker, M. Johannes and N. Polson, The impact of jumps in volatility and returns, *Journal of Finance* **58** (2003) 1269–1300.

[26] X. Gabaix, P. Gopikrishnan, V. Plerou and H.E. Stanley, A theory of power-law distributions in financial market fluctuations, *Nature* **423** (2003) 267–270.

[27] 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* **81** (1999) 617–631.

[28] H. Geman, D. Madan and M. Yor, Time Changes for Lévy processes, *Mathematical Finance* **11** (2001) 79–96.

[29] I. Karatzas and S.E. Shreve, *Brownian motion and stochastic calculus*, Springer Verlag (1988) New York.

[30] N.M. Kiefer, Discrete Parameter Variation: efficient estimation of a switching regression model, *Econometrica* **46** (1978) 427–434.

[31] N. Meddahi, A theoretical comparison between integrated and realized volatility, *Journal of Applied Econometrics* **17** (2002) 479–508.

[32] I. Monroe, Processes that can be embedded in Brownian motion, *Annals of Probability* **6** (1978) 42–56.

[33] J. Pan, The jump-risk premia implicit in options: evidence from an integrated time series study, *Journal of Financial Economics* **63** (2002) 3–50.

[34] M.K. Pitt and N. Shephard, Filtering via simulation: auxiliary particle filter, *Journal of the American Statistical Association* **94** (1999) 590–599.

[35] R. Roll, A simple implicit measure of the effective bid-ask spread in an efficient market, *Journal of Finance* **39** (1984) 1127–1140.

[36] T.H. Rydberg and N. Shephard, Dynamics of trade-by-trade price movements: decomposition and models, *Journal of Financial Econometrics* **1** (2003) 2–25.

[37] K. Singleton, Estimation of affine asset pricing models using the empirical characteristic function, *Journal of Econometrics* **102** (2001) 111–141.

[38] S.J. Taylor, X. and Xu, The incremental volatility information in one million foreign exchange quotations, *Journal of Empirical Finance* **4** (1997) 317–340.

[39] R.A. Wood, T.H. McInish and J.K. Ord, An investigation of transaction data for NYSE stocks, *Journal of Finance* **25** (1985) 723–739.

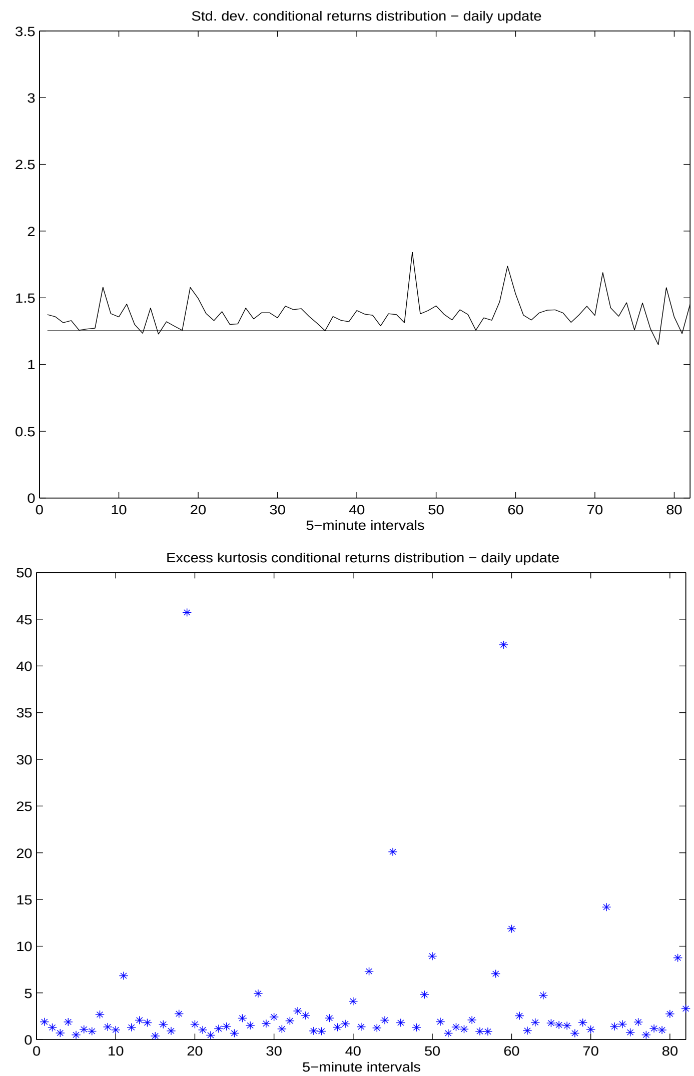

Figure 5: Standard deviation and excess kurtosis of conditional returns by intraday intervals using daily updated volatility estimates.

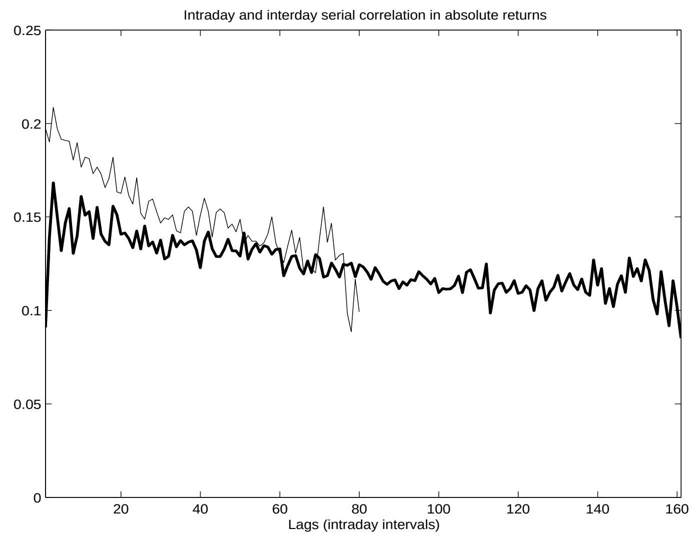

Figure 6: Intraday and interday serial correlation of absolute returns.

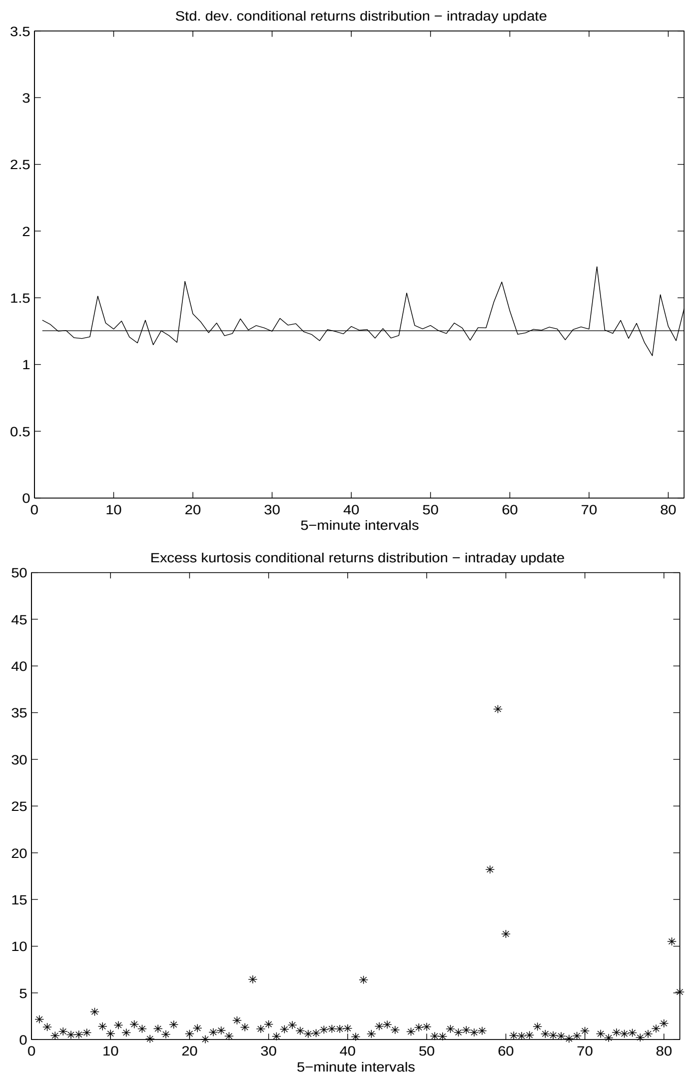

Figure 7: Standard deviation and excess kurtosis of conditional returns by intraday intervals using 5-minute updated volatility estimates.

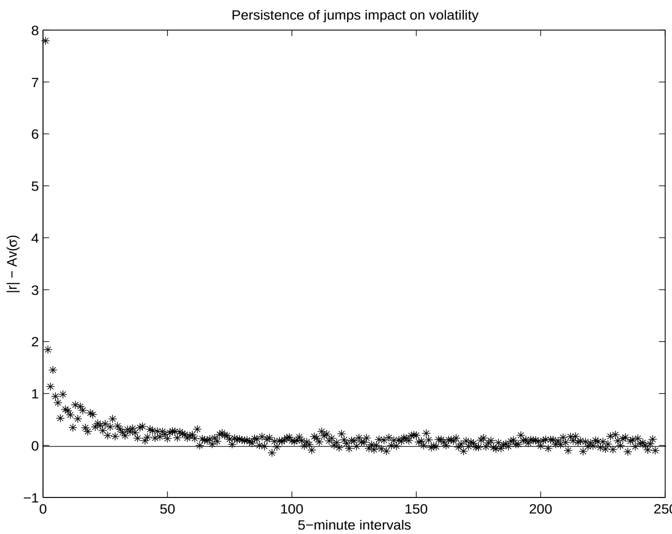

Figure 8: Persistence of jumps impact on volatility measured in 5-minute intraday intervals.

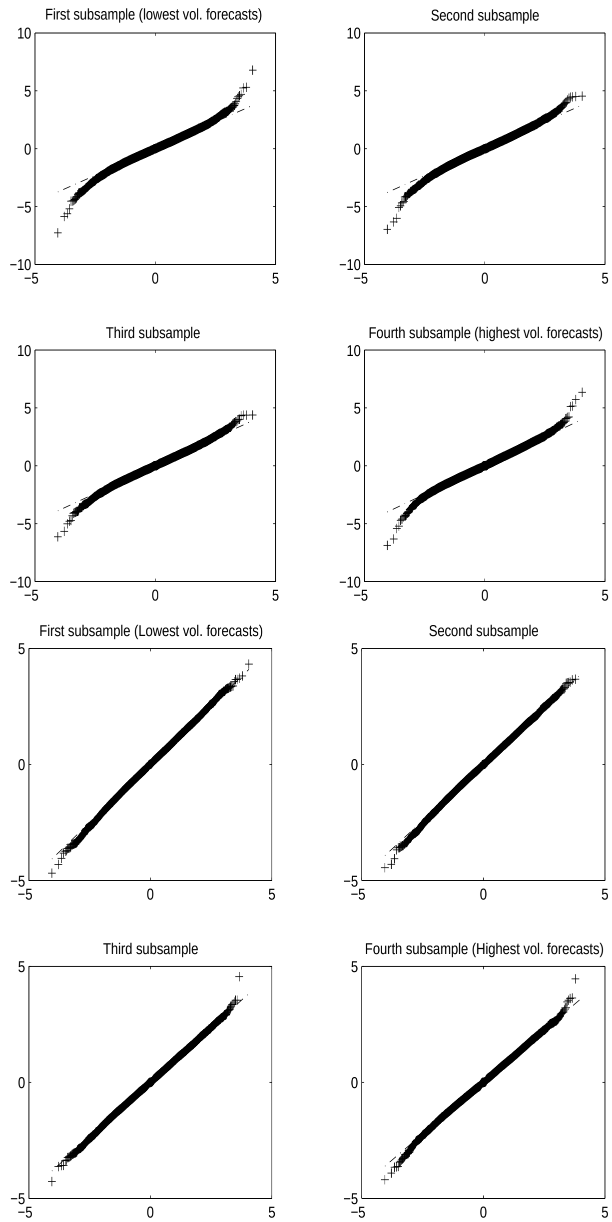

Figure 9: Normal QQ plots - return density forecasts using changing intraday volatility without (top) and with (bottom) jumps in volatility.

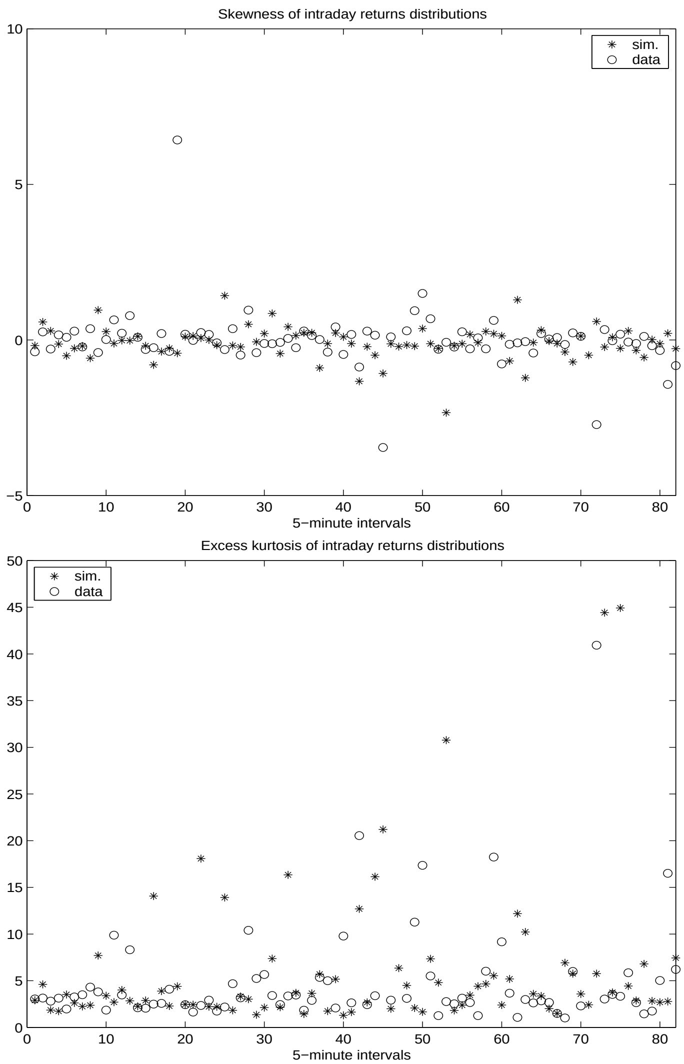

Figure 10: Skewness and excess kurtosis for intraday simulated and observed returns.

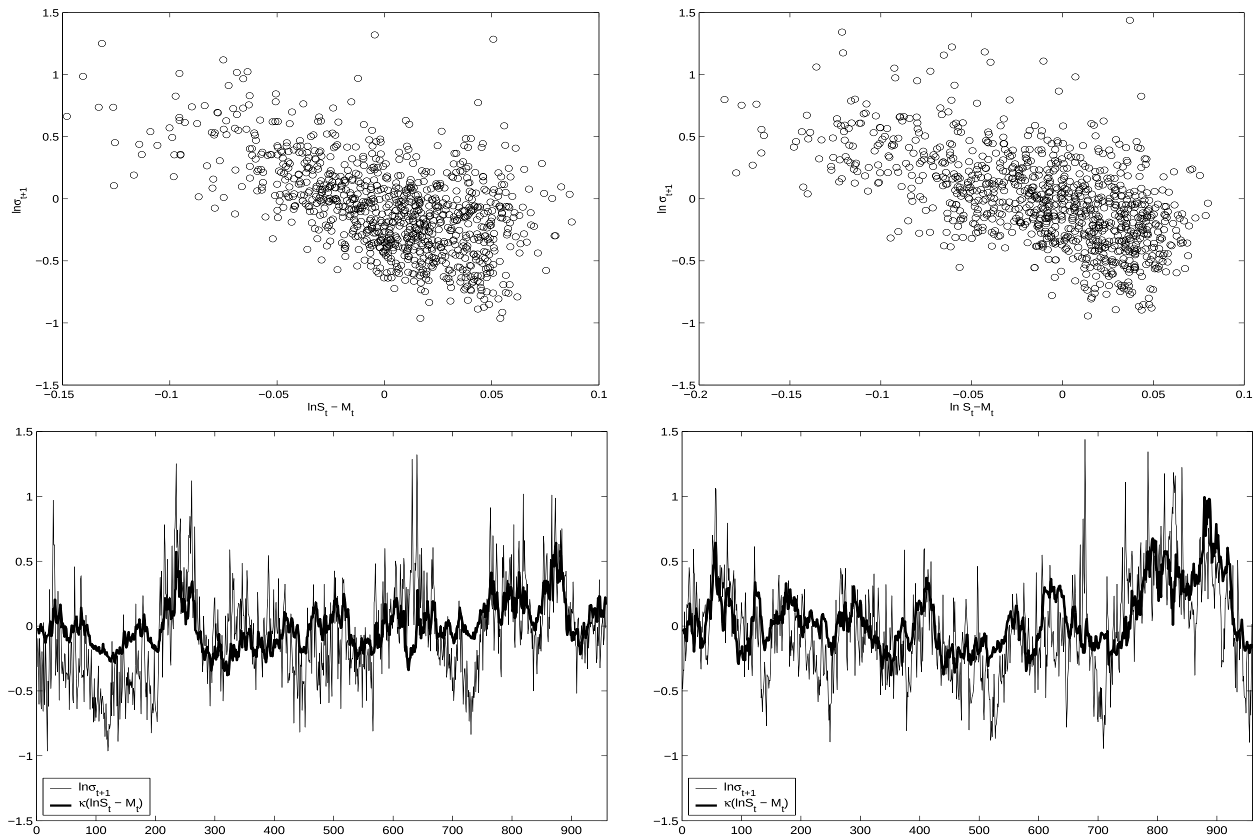

Figure 11: Scatter plot and time series of leverage measure against log volatility proxy - Market data (left) and simulated sample (right).