Short time-scale in S&P 500 volatility

Jean-Pierre Fouque* George Papanicolaou† Ronnie Sircar‡
Knut Solna§

May 2, 2002

Abstract

We use model independent statistical tools to demonstrate the presence of a short time-scale, on the order of days, in S&P 500 volatility. The data is high-frequency and we discuss the inherent problems due to the "day effect". We show that the presence of a well-separated longer time-scale does not affect our analysis. On the other hand, we explain why this short scale has to be dealt with in option pricing. We assess the sensitivity of this estimate using simulations from a stochastic volatility model calibrated roughly to the data.

Contents

*Department of Mathematics, NC State University, Raleigh NC 27695-8205, *fouque@math.ncsu.edu*.
Work partially supported by NSF grant DMS-0071744.
†Department of Mathematics, Stanford University, Stanford CA 94305, *papanico@math.stanford.edu*.
‡Department of Operations Research & Financial Engineering, Princeton University, E-Quad, Princeton, NJ 08544, *sircar@princeton.edu*. Work partially supported by NSF grant DMS-0090067.
§Department of Mathematics, University of California, Irvine CA 92697, *ksolna@math.uci.edu*.

1 Introduction & Outline

1.1 Estimation of time-scales of volatility mean reversion

Estimating rates of mean reversion of the S&P 500 volatility means that we want to identify characteristic time scales in the data. Long time-scales (on the order of several months) have been studied extensively (see [10] in the context of the Dow Jones index, and other references cited therein). Here, we demonstrate the presence of a well-identified *short* time-scale (on the order of a few days). In order to identify this scale, we analyze high-frequency data using the empirical structure function, or variogram, of the log absolute returns. We can alternatively use the spectrum of the log absolute returns. This is presented in detail, along with the results, in Section 2.

One way of assessing the sensitivity of this estimate, and the one we follow, is to simulate numerically the returns from a stochastic volatility model with both a long and a short scale, and then to use the variogram or the spectrum. This sensitivity analysis requires a rough calibration of the chosen model so that the simulated data reproduce well the observed variogram and spectrum.

1.2 Fast mean reversion and scales

We are interested in the short-run behaviour of the S&P 500. Thus, we look at high frequency, intraday data so that we can resolve scales on the order of days. At this frequency, a sample corresponding to one year of data is sufficient. We illustrate with data from 1994 and 1999-2000.

Many papers (for example Stein [16], Diz and Finucane [7], Engle and Mustafa [9] and Bollerslev and Mikkelsen [5]) consider daily data over much longer time periods. By looking at data over several years they estimate a longer scale on the order of months. Our analysis is complementary to this, as we find an additional short range time-scale. This has also been identified in various financial time series and by various other methods [1, 6, 2, 3]. We discuss this further in Section 2.3.

The estimation of the order of the fast time-scale is done in a model independent manner. However, the "day effect", which is the systematic intraday variation of volatility (see Figure 5), must be accounted for properly (Section 2.4). Andersen and Bollerslev [2] discuss high frequency data in which they also observe the day effect. They note that direct parameter estimation with the generalized method of moments (GMM) performs poorly. They get around this issue by low pass filtering the high frequency data in order to concentrate on long term mean reversion effects. Since we are interested in short term dynamics we have to deal with intraday data differently. We observe the day effect in the periodic oscillations in the variogram (Figure 3, top) as well as in the spectrum (Figure 6, top). We then account for it in quantifying the fast scale. This procedure is validated by replicating numerically the data (Figure 3 bottom, Figure 4, Figure 7, Figures 9-12).

1.3 Consequences for Option Pricing

The presence of a short scale has important consequences for option pricing. In Section 3, we show how this fast scale can be dealt with (through an asymptotic analysis). We argue that in fact this scale has to be dealt with and cannot be ignored, as we see from S&P 500 options data.

We summarize and conclude in Section 4.

2 Volatility Statistics of the S&P 500

Our source for the data that is analyzed in this section is the Berkeley Options Database described in [14]. This gives us the quoted S&P 500 index corresponding to index option quotes for the year 1994. In Section 2.7, we present analogous findings with more recent (1999-2000) data.

2.1 Notation

The S&P 500 index is recorded at irregular time intervals every few seconds. Let S_n denote the S&P 500 index recorded at time t_n. We average the data over five minute intervals so that we have 72 data points per day. We collapse the time by eliminating overnights, weekends and holidays so that we have 251 trading days with $72 \times 251 = 18072$ data points per year.

We introduce the *normalized fluctuations* of the data

$$\bar{D}_n = \frac{2(\bar{S}_n - \bar{S}_{n-1})}{\sqrt{\Delta t}(\bar{S}_n + \bar{S}_{n-1})}. \tag{1}$$

Here $\{\bar{S}_n\}$ are the five minute averaged S&P 500 data and Δt is the five minute time increment in annualized units. It makes very little difference in the results if the five-minute averaging of the data is done on the S_n or on the normalized fluctuations. This and other sensitivity issues regarding estimation are considered further in the Appendix. All of our analysis is based on the 18072 data points for the normalized fluctuations $\{\bar{D}_n\}$ for 1994. In the framework of the general class of stochastic volatility models

$$dS_t = \mu S_t dt + \sigma_t S_t dW_t\,, \tag{2}$$

the continuous analog of the normalized fluctuations can be written formally as

$$\frac{1}{\sqrt{\Delta t}}(\frac{\Delta S_t}{S_t} - \mu \Delta t) = \sigma_t \frac{\Delta W_t}{\sqrt{\Delta t}} \tag{3}$$

where $\sigma_t,\ t \geq 0$ is the (positive) volatility process, μ is a constant, and ΔW_t is the increment of a Brownian motion. The subtraction of the incremental mean return $\mu \Delta t$ is omitted in the discrete normalized fluctuations $\{\bar{D}_n\}$ because it is negligibly small. Based on (3) we model the normalized fluctuation process by

$$\bar{D}_n = \sigma_n \epsilon_n \tag{4}$$

where $\{\epsilon_n\}$ is a sequence of independent identically distributed Gaussian random variables with mean zero and variance one and $\{\sigma_n\}$ is the volatility at time t_n.

We will analyze the log absolute value of the normalized fluctuations

$$X_n = \log |\bar{D}_n| = \log \sigma_n + \log |\epsilon_n| \tag{5}$$

which is the sum of the discrete log-volatility process and a discrete white noise process.

2.2 Variogram analysis

We begin with a study of the *empirical structure function* or *variogram* of X_n:

$$V_j^N = \frac{1}{N} \sum_n (X_{n+j} - X_n)^2 \tag{6}$$

where j is the lag and N is the total number of points.

2.2.1 Variogram from S&P 500 Data

The **top plot** in Figure 1 shows the variogram (dotted line) of the log fluctuations $X_n = \log(|\bar{D}_n|)$ after applying a ten point median filter to compensate for the singular noise $\log |\epsilon_n|$. The first observation is the clear periodic component. This due to the systematic intraday effect which we discuss and model in Section 2.4. If, after removal of this component, the variogram were flat, we would conclude there is no short time-scale on the order of days. To gauge the extent of curvature, we fit the variogram to

$$V_j^N \approx 2\gamma^2 + 2\nu^2(1 - e^{-\alpha j \Delta t}), \tag{7}$$

a convenient three parameter family, where the vertical intercept $2\gamma^2$ gives twice the variance γ^2 of the noise $\log |\epsilon_n|$ and $1/\alpha$ quantifies the characteristic time-scale by analogy with exponential decorrelation in Markov processes.

Remark The expression (7) would be exact (in the limit $N \to \infty$), if σ_t was the exponential of an Ornstein-Uhlenbeck (OU) process (Y_t), with α being the rate of mean-reversion and ν^2 the variance of its invariant distribution:

$$\sigma_t = e^{Y_t}; \qquad dY_t = \alpha(m - Y_t)\, dt + \nu\sqrt{2\alpha}\, dZ_t. \tag{8}$$

This follows from

$$\text{covariance}(Y_0, Y_t) = \nu^2 e^{-\alpha t}$$

under the invariant distribution $\mathcal{N}(m, \nu^2)$. With this parameterization,

1. ν represents the typical size of the fluctuation of Y, and consequently of volatility. Empirically, this is neither extremely small nor extremely large.

2. The quantity $1/\alpha$ represents the intrinsic time-scale of the factor Y.

For simplicity, we have assumed the Brownian motions (Z_t) and (W_t) are independent.

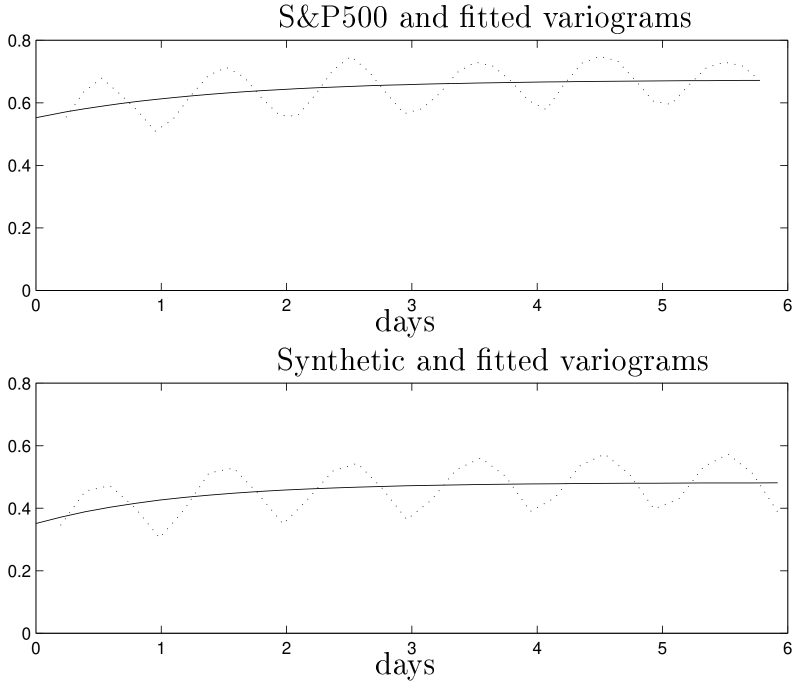

Figure 1: *The empirical variogram for the S&P* 500 *, top plot dotted line and for simulated data based on our stochastic volatility model, bottom plot, dotted line. The solid lines are exponential fits from which the rate of mean-reversion can be obtained.*

2.2.2 Implications

We point out the following:

1. Intraday Variations Note in Figure 1 the periodic component corresponding to day effects in the data that is not in the model (4) but will be incorporated below. The solid line is a fitted exponential obtained by nonlinear least squares regression. The estimate was obtained by first computing the average of the empirical variogram over the first and last days and then doing a one dimensional least squares fit for α.

2. Result of the Fit The estimated mean-reversion time, $1/\alpha$ in equation (7), is 1.5 days which identifies the presence of a short time-scale in S&P 500 volatility. We analyze the sensitivity in the estimate of $1/\alpha$ in Appendix A and find that it is approximately ± 0.4 days.

2.2.3 Variogram from Synthetic Data

The bottom plot in Figure 1 shows the estimate of the variogram for simulated log normalized fluctuations, *where the simulations incorporate a model of the day effect detailed below.* They are also median filtered as the S&P 500 data. The solid line is the estimated exponential fit obtained by nonlinear least squares.

To generate simulated normalized fluctuations, we used a discretization of the exponential OU model (8). We needed rough estimates of the parameters m, α, ν. The variogram estimation gives $\nu \approx .26$ and $1/\alpha \approx 1.5$ days which corresponds to $\alpha \approx 167$ in annualized units.

Then the computation of the average square volatility $\overline{\sigma}^2 = E[\bar{D}_n^2] = e^{2m+2\nu^2}$ leads to an estimate of m from the data. We use $\overline{\sigma} = 0.07$. Throughout, we have assumed the independence of the two Brownian motions. It is simple to introduce a correlation

$$d\langle W_t, Z_t\rangle = \rho\, dt$$

in the model used to generate the synthetic data. Extensive simulations show that the estimates of the parameters $\nu, \overline{\sigma}, \alpha$ are insensitive to the correlation ρ in the range $[-0.8, 0]$ that we expect it to be. This is shown in Figure 2.

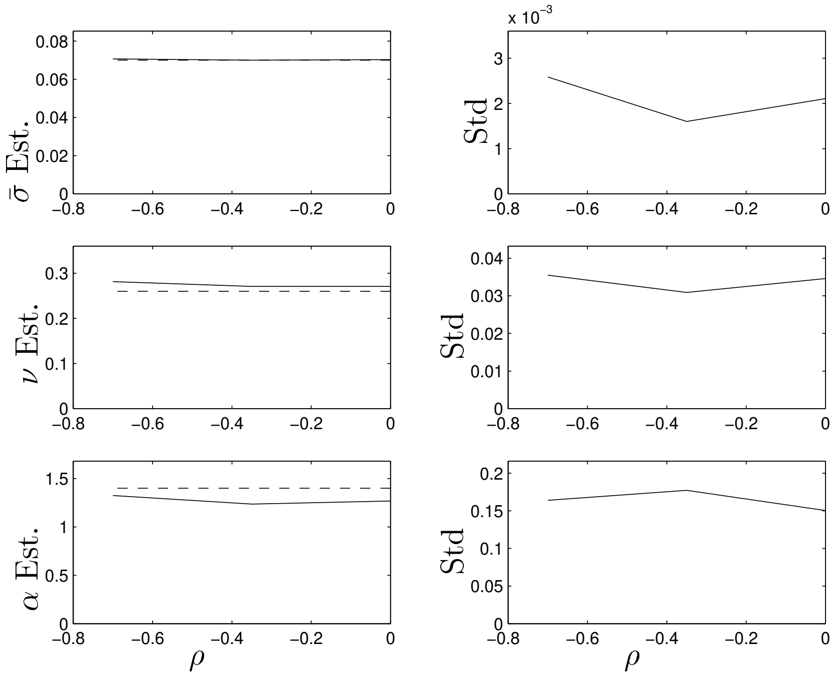

Figure 2: *Sensitivity of the estimated parameters $\bar{\sigma}$, ν and α on the correlation ρ in the range [-0.8,0].*

Observe how the variogram analysis separates the periodic daily component from the stochastic behavior in the volatility without any prior knowledge on these two components. This will also be apparent in the spectrum analysis presented below in Section 3.5. In contrast, the performance of GARCH estimates are strongly affected by this intraday effect described in the following section.

2.3 Effect of Longer Scales

A longer sample of data with the variogram extended to larger lags would reveal the presence of a longer characteristic time-scale in the S&P 500 , as documented in numerous studies. This is on the scale of months or longer, well-separated from the short scale we have identified here. It is natural to ask how the presence of this scale would affect our empirical analysis. In fact it does not, as we now explain. Let us consider a model with two components driving the volatility, one fast and one slow. This is a special case of the three-component model analyzed by LeBaron [15], who shows that a three-factor multiscale stochastic volatility model can lead to an apparent power-law variogram at medium lag lengths. Chernov *et al.* [6] also suggest a two-factor stochastic volatility model with one factor mean-reverting on a short scale and the other on a longer scale.

We consider a stochastic volatility model of the form

$$\sigma_t = e^{Y_t} e^{L_t},$$

where (Y_t) and (L_t) are independent OU processes. The first (Y_t) has characteristic mean-reversion time $1/\alpha$ on the order of days, while the second (L_t) mean-reverts on the order of months. We denote the rate of mean-reversion and the variance of the equilibrium distribution of (L_t) by α_L and ν_L^2 respectively. The variogram formula (7) becomes

$$V_j^N \approx 2\gamma^2 + 2\nu^2(1 - e^{-\alpha j \Delta t}) + 2\nu_L^2(1 - e^{-\alpha_L j \Delta t}).$$

For the range of lags we a looking at, that is $j\Delta t$ up to a week, $\alpha_L j \Delta t$ is small and the last term is negligible. Hence the long scale plays no role in our previous variogram estimation. At the opposite extreme, observe that if we had looked at a longer data sample and included larger lags, then the first term would contribute only a constant $(2\nu^2)$ except at the very short lags.

2.4 The day effect

The day effect is put into a one-factor stochastic volatility model by replacing (4) with

$$\bar{D}_n = f(Y_n) g_n \epsilon_n, \tag{9}$$

where we allow for more general functions f than the exponential. Here Y_n is the discrete OU process, ϵ_n a unit white noise process and g_n a deterministic periodic function that models the systematic intraday variations in the volatility. The estimate of g_n from the S&P 500 data is shown in Figure 3.

When the day effect is included in the simulations then the synthetic variogram in Figure 1, bottom, is very close to the one for the S&P 500 shown in Figure 1, top, except for the horizontal intercept. This intercept depends on the variance of the process $\log(|\epsilon_n|)$ which is sensitive to the fine properties in the tail of the distribution of ϵ_n. These are properties which, like the correlation ρ, we do not need to model and estimate precisely from data. From the point of view of derivatives pricing or hedging it is more natural to relate ρ to the skew of implied volatility and calibrate it *indirectly* from S&P 500 options data as explained in Section 4.

2.5 Spectrum

Instead of using the variogram or structure function we can go into the frequency domain and analyze the empirical spectrum of the log fluctuations X_n given by (5). The analog of (7) for the spectrum is the Lorentzian model

$$\Gamma_\omega^N \approx 2\gamma^2 + \frac{2\nu^2}{\pi} \frac{\alpha}{\alpha^2 + \omega^2}, \tag{10}$$

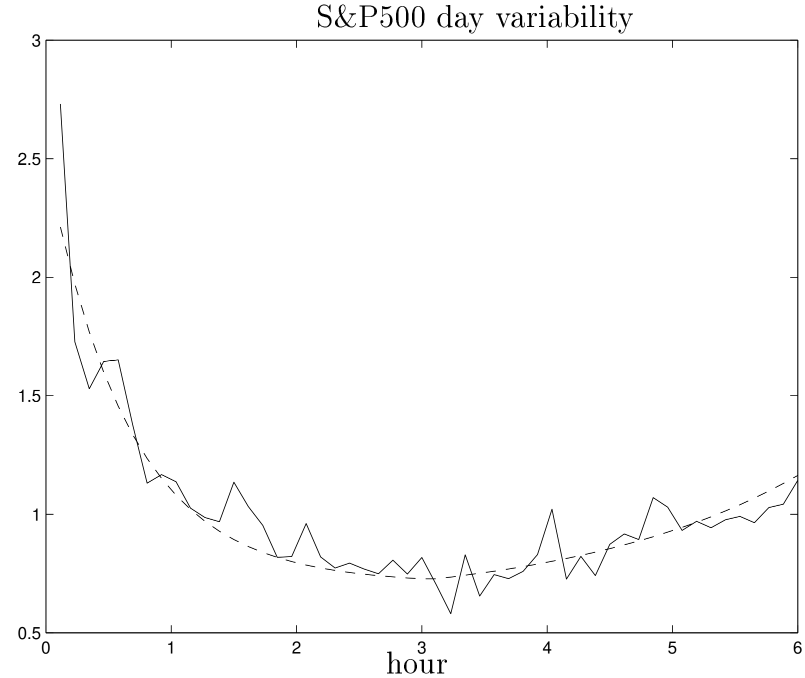

Figure 3: *Estimated intraday variability envelope (dashed line). The solid line is the root mean square of the observed normalized fluctuations $\bar{D}_n$ as a function of trading hour, averaged over 250 days to give the variability envelope g_i. We model g_i by a sum of two exponentials $g(x) = a_1 + a_2 \exp(-x/l_1)$ for $x \in (0, 1/2)$ and $g(x) = a_3 + a_4 \exp((x-1)/l_2)$ for $x \in (1/2, 1)$ with $(a_1, a_2, l_1) = (.7, 1.9, .1)$ and $(a_3, a_4, l_2) = (.6, .5, .3)$. The parameters were obtained by a least squares fit. Note that the volatility is strongest at the beginning of the trading day.*

ignoring again the correlation between the noise $\log(|\epsilon_n|)$ and the OU process Y_n as well as the day effect that is just an additive impulse fuction on the right at the frequency of one day.

In Figure 4, dotted line top, we show the spectrum for the S&P 500 data. The Lorentzian model is shown by the solid line. Note the good match. The day effect is the impulse at the frequency of one day.

The spectrum is computed using the MATLAB routine 'spectrum', which computes the averaged periodogram with a Hanning window. The 1994 year data is divided into segments of 14 trading days and averaging is done over 17 segments that contain nearly all the 251 trading day data. The bottom plot is the spectrum of two realizations from the estimated model for $\bar{D}_j$. The same Lorentzian fit for both is the solid line. It is qualitatively very similar to that of the real data at the top. The daily variations are not perfectly sinusoidal and we see also some higher order harmonics at frequencies corresponding to submultiples of days. Without the median filtering these higher order harmonics become visible also in the real data as can be seen in Figure 12. As in Figure 1 we see that the white noise part for the simulated data is somewhat smaller than that of the real data, as was explained in connection with Figure 1.

It is important to note that even though the Lorentzian spectrum is clearly visible for small frequencies, the variability in the spectrum in this frequency range is fairly high. For this reason we use the variogram rather than the spectrum to estimate the rate of mean-

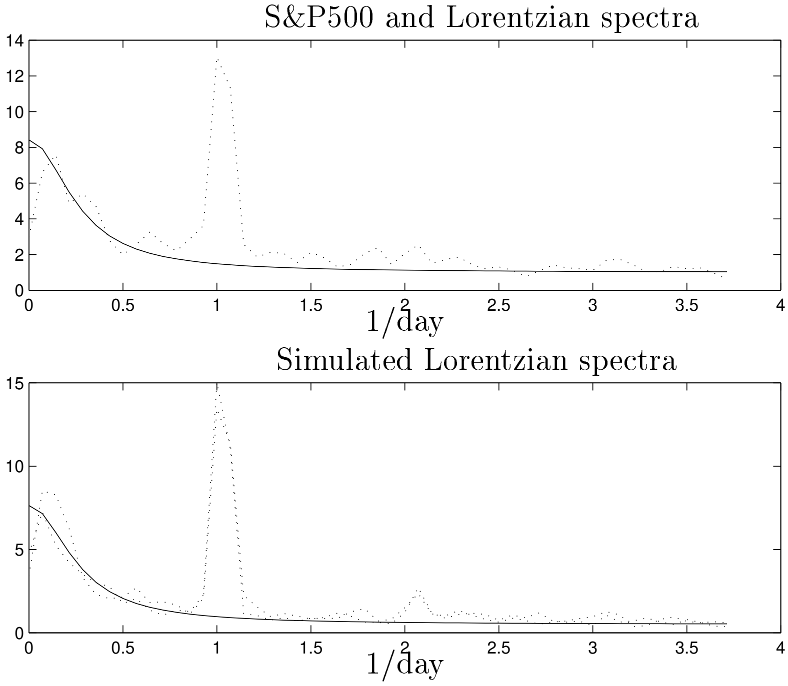

Figure 4: *The top plot is the spectrum (dotted line) of* $\log|\bar{D}_n|$ *for the real S&P* 500 *data. The Lorentzian model is shown by the solid line. The bottom plot is the spectrum of two simulated realizations from the estimated model for* $\bar{D}_j$. *The same Lorentzian model as in the previous figure is shown by the solid line. Note the good match, including the day effect.*

reversion α of the fast component.

2.6 Histograms

In Figure 5 we make use of scale separation to estimate separately the components in the normalized fluctuations $D_n = f(Y_n)g_n\epsilon_n$. The estimate of $f(Y_n)g_n$ is the square root of the local mean of $\bar{D}_n^2$ obtained with a uniform filter of width about two hours, which is small relative to the variation of $f(Y_n)g_n$ but large relative to the variation of ϵ_n. The histogram of the resulting estimates is shown in the top left plot for the S&P 500 data. The estimate of ϵ_n is obtained from the ratio of $\bar{D}_n$ and the estimated value for $f(Y_n)g_n$. The histogram of these estimates is given in the top right plot. It has a Gaussian form, as expected, but with somewhat fatter tails. In the bottom plots we show the same quantities but starting with simulated data for D_n. Note that the resulting estimates for the marginal distributions of the components of the normalized fluctuations conform with those obtained from the S&P 500 data and shown in the top plots.

2.7 Analysis of 1999-2000 S&P 500 Data

We present, for comparison purposes, the analogs of the variogram (Figure 1), the spectrum (Figure 4) and the intraday variability (Figure 3), this time using S&P 500 data from the period July 1999 through June 2000. These are shown in Figures 6, 7 and 8 respectively.

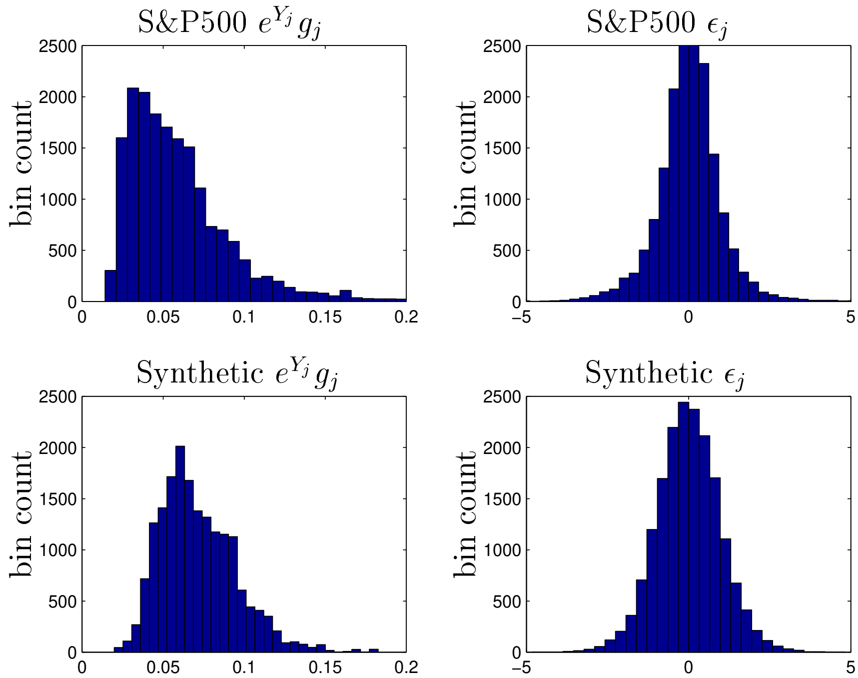

Figure 5: *Histogram of* $f(Y_n)g_n$ *and* ϵ_n *for the S&P* 500 *is shown on top. The bottom plots are the same quantities for simulated normalized fluctuations* $\bar{D}_n$.

The data was prepared by TickData and was available to us at a frequency of one record per minute. The exponential fit in Figure 6 recovers a mean-reversion time of 1.7 days.

3 Fast mean-reversion and option pricing

In this section, we discuss the presence of a fast and a slow scale on the pricing of options within a stochastic volatility model.

3.1 Scales under the Pricing Measure

Within the context of a simple two-component stochastic volatility model, similar to that discussed in Section 2.3, we demonstrate that if one of the two scales could be ignored under the pricing measure, there is an asymmetry in the choice of which scale to ignore. Since correlation between equity returns and volatility shocks is important in derivative pricing, we include them here explicitly.

We look at the model

$$dS_t = \mu S_t\, dt + \sigma_t S_t\, dW_t, \tag{11}$$

$$\sigma_t = e^{Y_t} e^{L_t}, \tag{12}$$

$$dY_t = \alpha(m - Y_t)\, dt + \nu\sqrt{2\alpha}(\rho\, dW_t + \rho'\, dZ_t), \tag{13}$$

$$dL_t = \alpha_L(m_L - L_t)\, dt + \nu_L\sqrt{2\alpha_L}(\rho_L\, dW_t + \rho'_L\, dZ_{L,t}), \tag{14}$$

written under the physical measure, where W, Z and Z_L are independent standard Brownian motions, and $\rho' = \sqrt{1-\rho^2}$, ρ'_L defined similarly. For simplicity of exposition, we ignore the

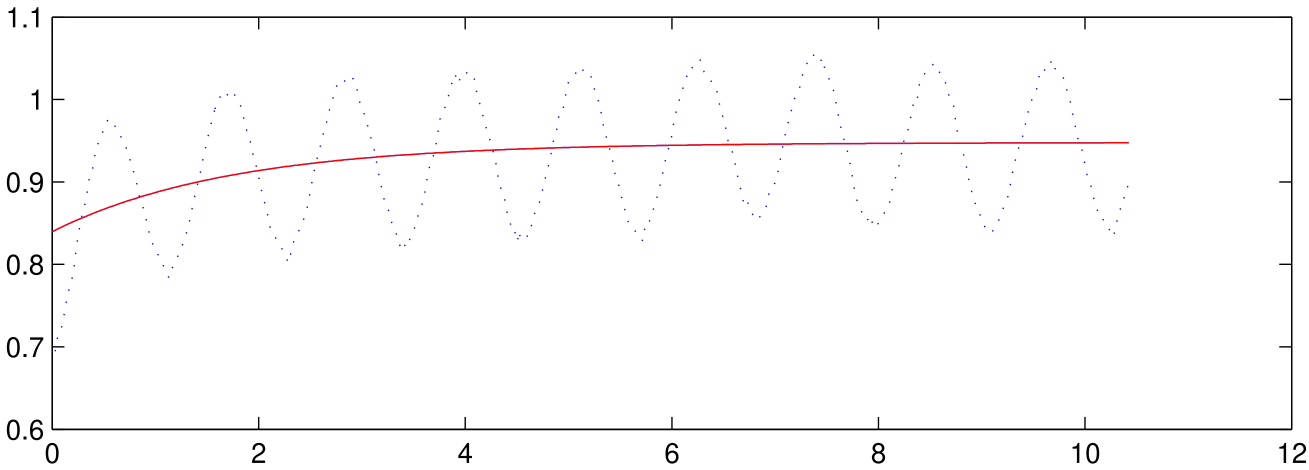

Figure 6: *The empirical variogram for the S&P* 500 *, dotted line, for the 1999-2000 data. The solid line is the exponential fit from which the rate of mean-reversion can be obtained.*

cutoff of the exponential function in (12) that is needed for (σ_t) to satisfy the condition of Girsanov's theorem.

In (13), α is large, modeling the fast scale we identified in the data analysis above, and in (14), α_L is of order one, modeling the longer scale that is described in numerous empirical studies.

Under the pricing measure, the evolution becomes

$$\begin{aligned}
dS_t &= rS_t\,dt + e^{Y_t}e^{L_t}S_t\,dW_t^\star, \\
dY_t &= \left[\alpha(m-Y_t) - \nu\sqrt{2\alpha}\left(\rho\frac{(\mu-r)}{\sigma_t} + \lambda_t\rho'\right)\right]dt + \nu\sqrt{2\alpha}\,(\rho dW_t^\star + \rho'\,dZ_t^\star) \\
dL_t &= \left[\alpha_L(m_L-L_t) - \nu_L\sqrt{2\alpha_L}\left(\rho_L\frac{(\mu-r)}{\sigma_t} + \lambda_{L,t}\rho_L'\right)\right]dt + \nu_L\sqrt{2\alpha_L}\left(\rho_L dW_t^\star + \rho_L'\,dZ_{L,t}^\star\right)
\end{aligned}$$

where λ and λ_L are the volatility risk premia for each of the two factors. We introduce the reduced notation

$$\begin{aligned}
\Lambda &= \rho\frac{(\mu-r)}{\sigma_t} + \lambda_t\rho' \\
\Lambda_L &= \rho_L\frac{(\mu-r)}{\sigma_t} + \lambda_{L,t}\rho_L'.
\end{aligned}$$

Choice of λ and λ_L corresponds to a choice of Λ and Λ_L

By choice of these premia, this model can be reduced to a one-component stochastic volatility model (with one time-scale) as is usually studied in the literature and used to calibrate the implied volatility surface. We have the two following extreme scenarios.

1. Reduction to Fast Scale By choosing, for instance,

$$\begin{aligned}
\Lambda_L &= 0 \\
\Lambda &= \tilde{\Lambda}(Y+L) + \frac{\sqrt{\alpha}L}{\nu\sqrt{2}} + \frac{\alpha_L(m_L-L)}{\nu\sqrt{2\alpha}},
\end{aligned}$$

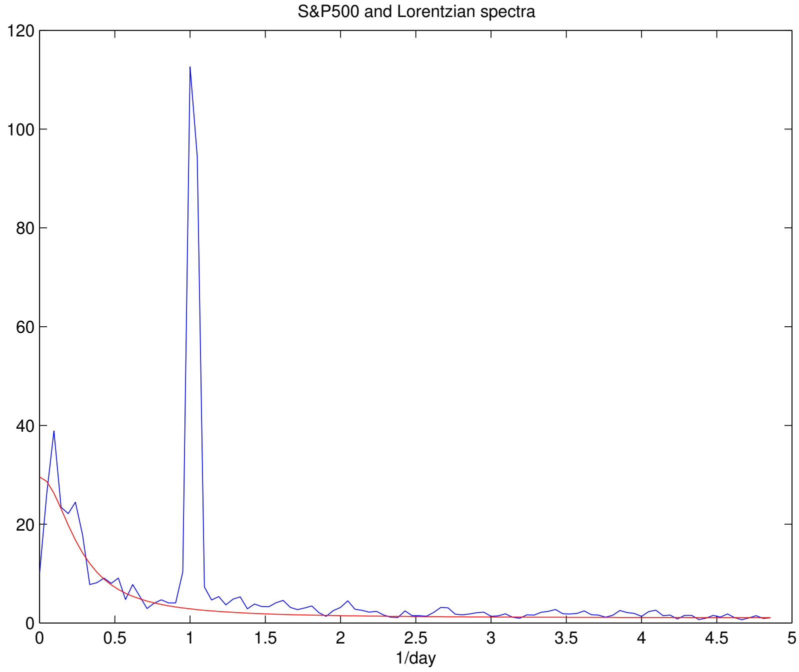

Figure 7: *Spectrum (solid line) of* $\log|\bar{D}_n|$ *for the 1999-2000 S&P* 500 *data. The Lorentzian model (omitting the day effect) is shown by the smooth curve.*

where $\tilde{\Lambda}$ is some function of the combined factor $\tilde{Y} = Y + L$, the model reduces to

$$\begin{aligned} dS_t &= rS_t dt + e^{\tilde{Y}_t} S_t dW_t^\star \\ d\tilde{Y}_t &= (\alpha(m - \tilde{Y}_t) - \nu\sqrt{2\alpha}\,\tilde{\Lambda}(\tilde{Y}_t))dt + \nu\sqrt{2\alpha}\,d\hat{Z}_t^\star + \nu_L\sqrt{2\alpha_L}\,d\hat{Z}_{L,t}^\star, \end{aligned}$$

where $\hat{Z}^\star$ is the Brownian motion $\rho W^\star + \rho' Z^\star$ and similarly for $\hat{Z}_L^\star$. This is a stochastic volatility model driven by one autonomous factor which is on the fast time-scale of order $1/\alpha$. Observe that the variance of the noise is of order α (and that the variance of the invariant distribution of the reduced factor $\tilde{Y}$ is of order one), simply because α_L is negligible when α is large, so that to this leading order the model is equivalent to just ignoring the long time scale.

2. Reduction to the Longer Time-Scale Similarly one can choose the Λ's so that the rate of mean-reversion of $Y + L$ becomes α_L corresponding to the longer scale. For instance, we can choose

$$\begin{aligned} \Lambda &= 0 \\ \Lambda_L &= \tilde{\Lambda}(Y + L) + \frac{\sqrt{\alpha_L}\,Y}{\nu_L\sqrt{2}} + \frac{\alpha(m - Y)}{\nu_L\sqrt{2\alpha_L}}, \end{aligned}$$

for some function $\tilde{\Lambda}$ of the combined factor $\tilde{Y} = Y + L$. However, the variance of the noise will always be of order α which is large. This is not the same as ignoring the short scale component from the beginning. It has been observed that when fitting one stochastic volatility component to option prices, this large variance shows up [4]. One

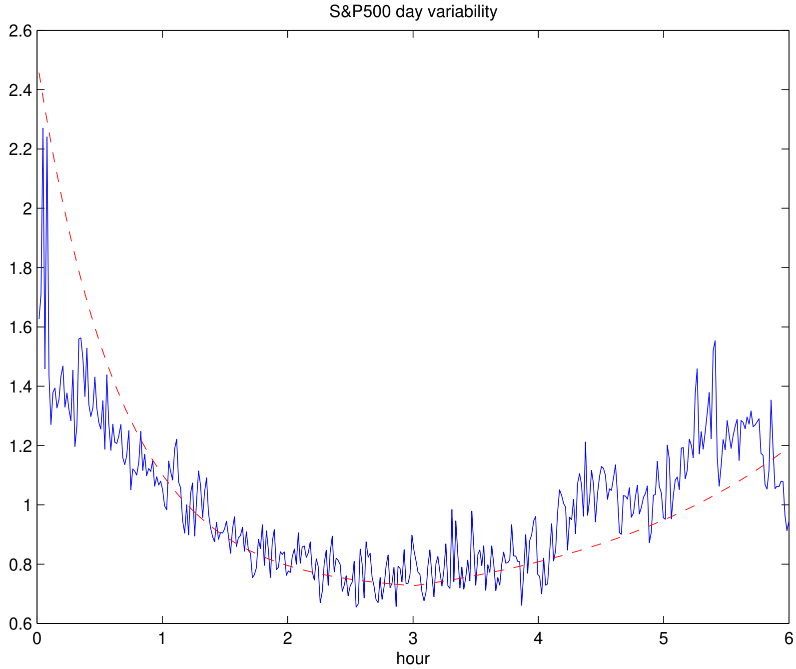

Figure 8: *Intraday variability in the 1999-2000 data. The dashed line is the envelope that we estimated for the 1993 data in Figure 3.*

> proposed remedy is to add jumps [8], but this also typically leads to an unrealistically high (risk-neutral) probability of large jumps. In response to this, Chernov *et al.* [6] suggest the incorporation of a short-run fast volatility factor to explain the large skews at short maturities.
>
> Furthermore, in this case, the variance of the reduced factor $\tilde{Y}$ becomes large since the variance of the invariant distribution of the driving noise term is of order α as above, but now the mean reversion, or the pull back term, is of order α_L, which is relatively small. Note also that the correlation between the white noise driving the stock and the white noise driving the volatility in *both* reductions is the correlation associated with the *fast* factor. It is thus the noise term that drives the fast factor that is the primary one and it cannot be ignored.
>
> In this example, it is apparent that the presence of the short scale could explain the discrepancy between the size of the noises under the physical and risk neutral measures.

It turns out that the first reduction to a one-component fast scale volatility driving process under the risk neutral measure leads to an asymptotic theory which simplifies the option pricing problem and yields explicit approximations that are very tractable for calibration. Notice that if the market priced options in this way, it would imply a large $\mathcal{O}(\sqrt{\alpha})$ risk premium for the slow factor to explain it within these models. (It would need an even larger $\mathcal{O}(\alpha)$ risk premium for the fast factor in the case of the other reduction). We do not claim to justify either of these extremely specific choices, but merely give these as examples of the implicit assumption of ignoring one or other factor if the real world dynamics are well-represented by the two-factor stochastic volatility model (11-14).

3.2 Option Pricing under Fast-Scale Volatility

We consider European derivatives with terminal payoff $h(S_T)$ when we have one fast mean-reverting stochastic volatility factor. The *no arbitrage* price P_t depends on the present stock price or index S_t *and* the present level of the volatility-driving process Y_t. It is given by the risk-neutral expected discounted payoff

$$P_t = \mathbb{E}^{Q(\Lambda)}\{e^{-r(T-t)}h(S_T)|S_t, Y_t\},$$

where $Q(\Lambda)$ represents probabilities in the risk-neutral world, Λ is the volatility risk premium. Under this measure,

$$\begin{aligned} dS_t &= rS_t dt + \sigma_t S_t dW_t^\star, \\ \sigma_t &= f(Y_t), \\ dY_t &= \left[\alpha(m - Y_t) - \nu\sqrt{2\alpha}\Lambda(Y_t)\right] dt + \nu\sqrt{2\alpha}\left(\rho dW_t^\star + \rho'\, dZ_t^\star\right), \end{aligned} \tag{15}$$

where $(W^\star, Z^\star)$ are independent Brownian motions under $Q(\Lambda)$. We shall assume that $\Lambda(y)$ is a function of y only. This can be validated *a posteriori* when the formula presented below is fitted to implied volatility data. We keep the functions $f(y)$ and $\Lambda(y)$ in (15) unspecified and we will see that when we have fast mean-reversion their exact forms are not needed.

The main result of an asymptotic analysis of this problem [11] says that when volatility is fast mean-reverting, we can approximate the derivative price P_t in the stochastic volatility environment by a more complicated Black-Scholes type formula in a constant volatility environment.

The procedure is as follows:

1. Let $(\overline{S}_t)_{t\geq 0}$ be the Black-Scholes lognormal model:

   $$d\overline{S}_t = r\overline{S}_t dt + \overline{\sigma}\,\overline{S}_t dW^\star,$$

   with $\overline{\sigma}$ the estimated historical volatility. Price the derivative with this model. That is, find

   $$P_{BS}(t, S) = \mathbb{E}^\star\{e^{-r(T-t)}h(\overline{S}_T)|\overline{S}_t = S\}.$$

2. The stochastic volatility price P_t is well-approximated by the price *in the Black-Scholes model* $\tilde{P}_t$ with the same terminal payoff $h(\overline{S}_T)$ and a payout rate $H(t, \overline{S}_t)$ between times t and T, given by

   $$H(t, S) = V_3 S^3 \frac{\partial^3 P_{BS}}{\partial S^3}(t, S) + V_2 S^2 \frac{\partial^2 P_{BS}}{\partial S^2}(t, S),$$

   where V_2 and V_3 are **small constants** related in a complex way to the original parameters (α, ν, m, ρ) and the unspecified functions f and Λ. That is, $\tilde{P}$ is given by the **modified Black-Scholes formula**

   $$\tilde{P}_t = P_{BS}(t, S_t) - (T - t)H(t, S_t)\,. \tag{16}$$

   The correction to the Black-Scholes price, which may be positive or negative, accounts for volatility randomness in a robust model-independent way.

Whenever there is an explicit formula for the Black-Scholes price of the derivative, for example calls, puts, binaries, barriers, there is an explicit formula for $\tilde{P}_t$. The correction to the usual Black-Scholes price can be more than 10% for S&P 500 options.

The error in the asymptotic approximation (16) is of order less than $1/\sqrt{\alpha}$ which is small when mean-reversion is fast, as in first scenario of Section 3.1. Notice that this approximation does *not* depend on the present level of volatility, which is not directly observable and usually difficult to estimate.

3.3 Calibration

Where simplification because of fast mean-reversion is extremely useful is in the parameter estimation problem. Having obtained $\overline{\sigma}$, the historical volatility as in Section 3, we can estimate V_2 and V_3 from the implied volatility skew. This is defined by

$$P^{\text{obs}} = P_{BS}(t, S_t; T, K, I) \tag{17}$$

where P^{obs} is the observed price of the option and the implied volatility $I = I(t, S, K, T, \overline{\sigma}, r)$ depends on all the variables involved in the determination of the price. If we take $h(S_T) = (S_T - K)^+$, a call option, compute the approximation described above to the stochastic volatility price and then work out the implied volatility I, we get a very simple formula

$$I = a\frac{\log(K/S)}{(T-t)} + b, \tag{18}$$

where

$$V_3 = -\overline{\sigma}^3 a \quad \text{and} \quad V_2 = \overline{\sigma}\left((\overline{\sigma} - b) - a(r + \frac{3}{2}\overline{\sigma}^2)\right). \tag{19}$$

It turns out that a has the same sign as ρ, so that a downward sloping skew indicates a negative correlation.

The calibration procedure is as follows:

1. Fit at-the-money implied volatilities to a straight line in the composite variable called the *log-moneyness-to-maturity-ratio* (LMMR)

$$\text{LMMR} := \frac{\log\left(\frac{\text{Strike Price}}{\text{Stock Price}}\right)}{\text{Time to Maturity}}.$$

   Estimate the slope a and the intercept b.

2. Estimate $\overline{\sigma}$, the historical volatility from stock price returns, and compute V_2 and V_3 using the formulas above.

3. Price any other European options by (16)

We stress again that this is *not* model specific. It does not depend on the choice of a function $f(y)$ in (15) or a particular ergodic driving diffusion Y. There is no need to estimate $(\alpha, \nu, m, \rho, \Lambda)$, as well as μ, separately. Only the V_2, V_3 (equivalently the a, b), which contain these, are needed. Nor is the present value Y_t required.

Most importantly, a similar procedure holds for pricing barrier options and American options [13].

3.4 Stability of Parameter Estimates

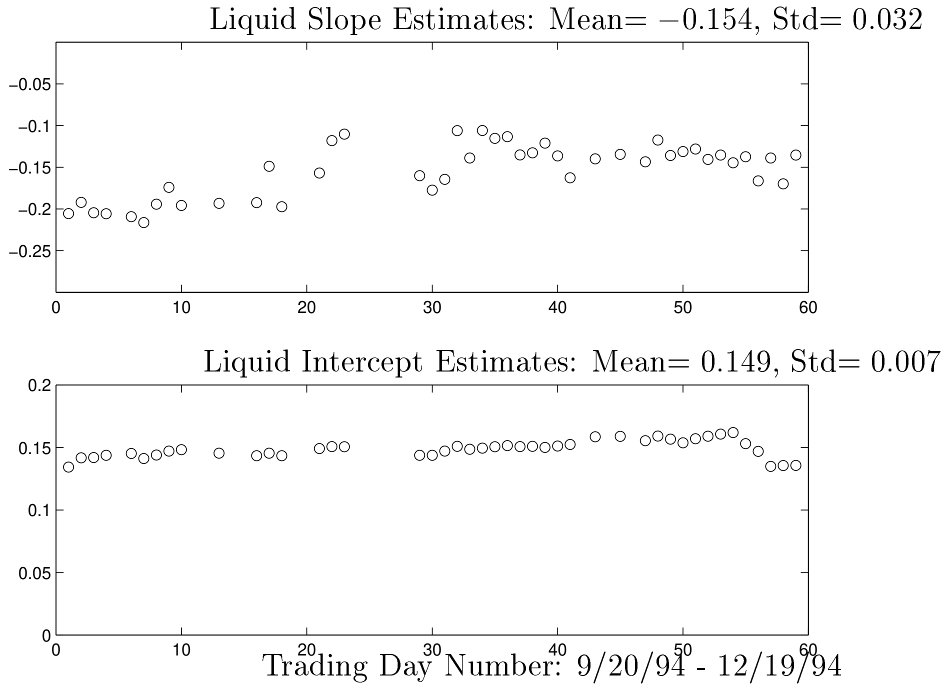

Figure 9: *Daily fits of S&P 500 European call option implied volatilties to a straight line in LMMR, excluding days when there is insufficient liquidity (16 days out of 60).*

We tested *a posteriori* the feasibility of the theory-predicted LMMR linefit for actual implied volatility data. We show in Figure 9 daily estimates of the slope and intercept coefficients a and b from fitting Black-Scholes implied volatilities from observed S&P 500 European call option prices

$$I^{obs}(t,S;K,T) = a\left(\frac{\log(K/S)}{T-t}\right) + b.$$

We note from the results that the slope coefficients a are small. This supports independently the presence of a fast time-scale in the S&P 500 and validates the choice of the first scenario and thus the use of the asymptotic formula. We also find that the estimates of a and b over a sixty-day period are stable as attested to by the standard deviations of the daily estimates shown in Figure 9. Improvement of the stability of these estimate is important and we address it in [12].

4 Summary and Conclusions

In this paper, we have demonstrated the presence of a short time-scale in S&P 500 volatility, on the order of 1.5 days. This is done within a large class of stochastic volatility models, within which the volatility factors decorrelate exponentially. Through a simple illustrative example with two time-scales, one short, the other long, we have shown that our finding is not affected by the presence of the longer scale. Our model-independent tools are the variogram and spectrum of log absolute returns. We have identified the short scale in S&P 500 high-frequency data from 1994. Our analysis of more recent data, from 1999-2000, gives the same result.

In the context of simple two-factor continuous diffusion stochastic volatility models, which are the subject of much discussion in the current literature, we have argued that this short time scale cannot be ignored, and we propose to deal with it through an asymptotic approximation which has been presented in [13] and summarized in Section 3.

In conclusion, the short time-scale is in the data, it has to be dealt with in option pricing and this can be done in a very tractable way through the asymptotic method.

A Sensitivity of estimation

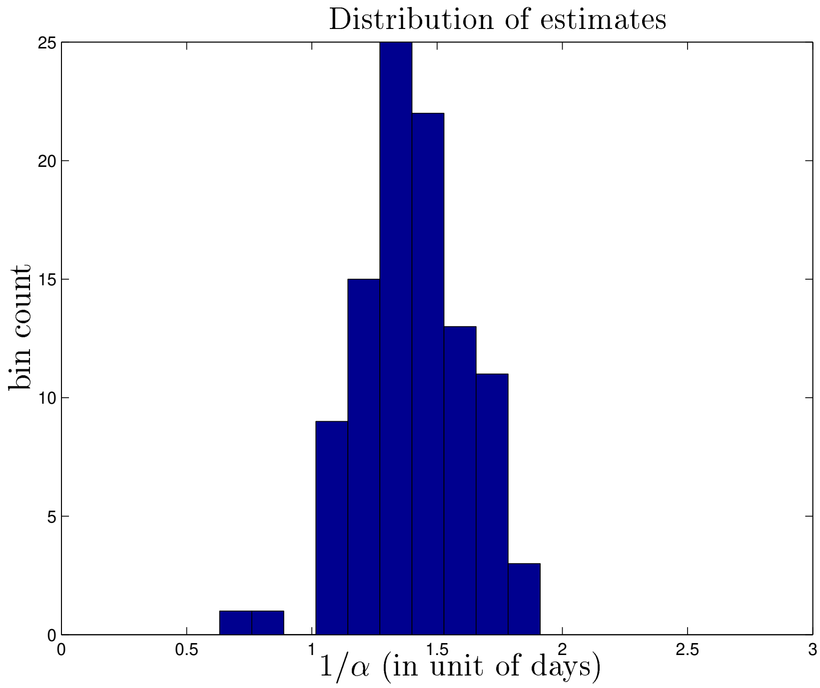

Figure 10: *Empirical distribution of the estimator for the rate of mean reversion α.*

In this Appendix we assess with numerical simulations the sensitivity of the estimate of the rate of mean-reversion α based on the variogram or on the spectrum of the log of the normalized fluctuations $\bar{D}_n$.

In Figure 10 we estimate the standard deviation of the estimate for α. We simulate several realizations of the normalized fluctuations and estimate the rate of mean-reversion from the variogram by nonlinear least squares fitting. The figure shows the distribution of the estimates when the rate of mean-reversion in the model is one day. Because of the day effect the estimator for α based on (7) is slightly biased by about $-.2$ days. In the figure we have compensated for the systematic bias. The standard deviation for the estimator of α is .4 days.

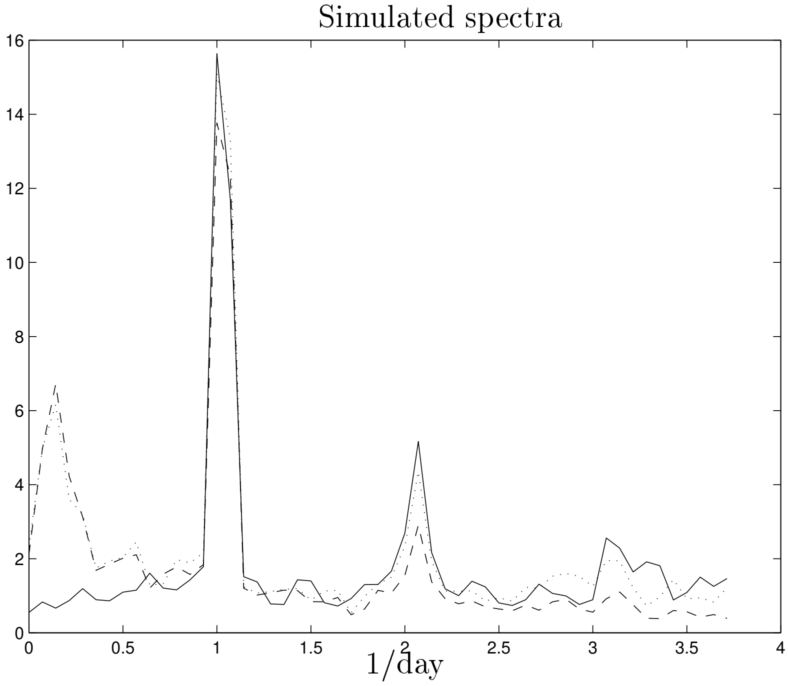

Figure 11: *Effect of median filtering on the spectrum for simulated data.*

Figure 11 shows that the median filter that we use above has only a small influence on the low frequencies. The dotted and dashed lines are spectra for a simulated realization of the normalized fluctuations. The dotted line is the spectrum without median filtering and the dashed line is the spectrum with median filtering. We see that even though the Lorentzian part of the spectrum is not affected much by the median filtering the day effect component in the spectrum is slightly blurred. The figure also illustrates that the rate of mean reversion α, shows up clearly in the power spectrum. The solid line is the spectrum of a realization when the rate of mean-reversion is one year rather than the estimated 1.5 days. We see that in that case the spectrum has a qualitative different form at low frequencies.

In the bottom plot of Figure 12 we show the spectrum of the real S&P 500 log normalized fluctuations when we do not include the median filtering. The data are then strongly intermittent. However, for low frequencies, the spectrum has the same qualitative structure as that in Figure 4 where we incorporated the median filtering step. Note that now the higher order harmonics of the day-effect are more visible. The top plot shows the variogram of the S&P 500 without median filtering. It has a somewhat more erratic shape than that in Figure 1. Note the change in scale on the vertical axis. The white noise component in the signal is now very strong relative to the coherent part. The estimated rate of mean-reversion is still 1.5 days, however. The fit with (7) is shown by the solid line.

Recall how we constructed the $\bar{D}_n$'s from the high frequency S&P 500 data. In Section 2, we *first* averaged S_n over 5 minute intervals and then constructed the $\bar{D}_n$. Alternatively, we could have formed the D_n's first and then averaged over 5 minute intervals to reduce the size of the data set. We then get the variogram shown in the top plot of Figure 13. Comparing with Figure 1 we find that, as expected, this change in strategy affects only the white noise component of the normalized fluctuations, not the crucial structural part that

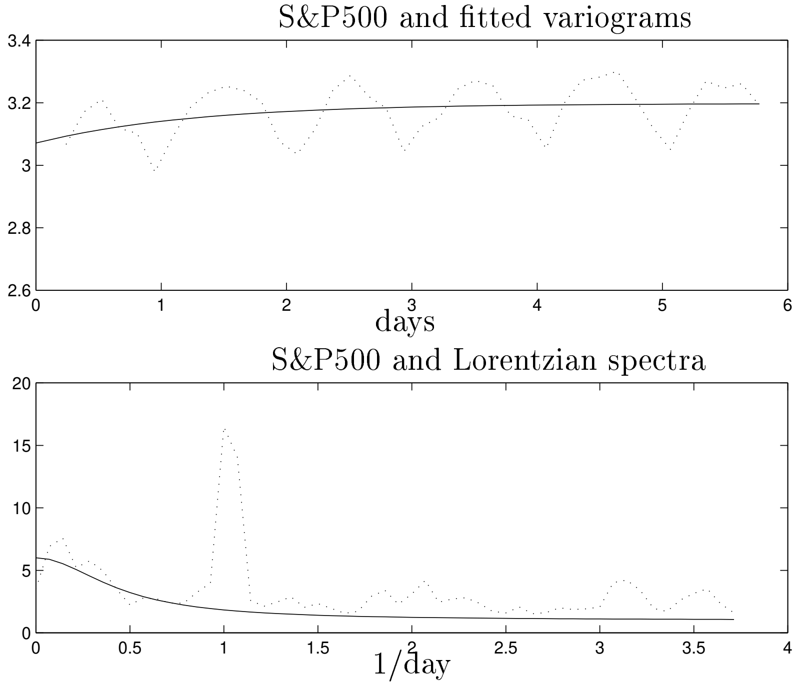

Figure 12: *Effect of median filtering on S&P 500 data.*

gives the estimate of the rate of the mean-reversion, which is 1.4 days according to the solid curve. In the bottom plot we consider the same data, that is, we compute the D_n's before we do the smoothing, and then compute the spectrum. Its shape (dotted line) is similar to that of Figure 4. The Lorentzian spectrum is shown by the solid line. We see that only the white noise part of the normalized fluctuations is affected by changing the smoothing procedure.

In Figure 14 we examine the sensitivity of the estimate of the rate of mean-reversion to the length of the averaging interval. In the top plot the averaging interval is 10 minutes, rather than 5, and the estimated rate of mean-reversion is 1.5 days. In the bottom plot the length of the averaging intervals is 2.5 minutes. The estimate of the mean-reversion rate is now 1.4 days. The estimates are obtained by non linear least squares regression as in the case with 5 minute averaging.

References

[1] S. Alizadeh, M. Brandt, and F. Diebold. Range-based estimation of stochastic volatility models. *Journal of Finance*, 2001. To appear.

[2] T. Andersen and T. Bollerslev. Intraday periodicity and volatility persistence in financial markets. *J. Empirical Finance*, 4:115–158, 1997.

[3] T. Andersen and T. Bollerslev. DM-dollar volatility: Intraday activity patterns, macroeconomic announcements, and longer run dependencies. *J. Finance*, 53:219–65, February 1998.

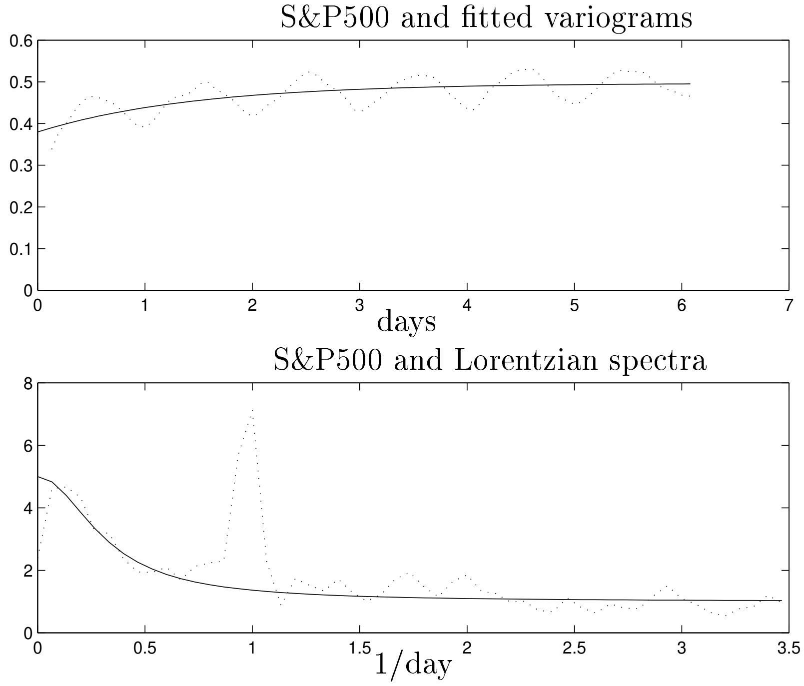

Figure 13: *Effect of the five minute averaging.*

[4] G. Bakshi, C. Cao, and Z. Chen. Empirical performance of alternative option pricing models. *J. Fin.*, 52(5), December 1997.

[5] T. Bollerslev and H. Mikkelsen. Long-term equity anticipation securities and stock market volatility dynamics. *J. Econometrics*, 92:75–99, 1999.

[6] M. Chernov, R. Gallant, E. Ghysels, and G. Tauchen. Alternative models for stock price dynamics. Preprint, Columbia University, January 2001.

[7] F. Diz and T. Finucane. Do the options markets really overreact? *J. Futures Markets*, 13(3):299–312, 1993.

[8] D. Duffie, J. Pan, and K. Singleton. Transform analysis and option pricing for affine jump-diffusions. *Econometrica*, 68:1343–76, 2000.

[9] R. Engle and C. Mustafa. Implied ARCH models from options prices. *J. Econometrics*, 52:289–311, 1992.

[10] R. Engle and A. Patton. What good is a volatility model? *Quantitative Finance*, 1:237–245, March 2001.

[11] J.-P. Fouque, G. Papanicolaou, and K. R. Sircar. Mean-Reverting Stochastic Volatility. *International Journal of Theoretical and Applied Finance*, 3(1):101–142, January 2000.

[12] J.-P. Fouque, G. Papanicolaou, K.R. Sircar, and K. Solna. Maturity cycles in implied volatility. *In preparation*, 2002.

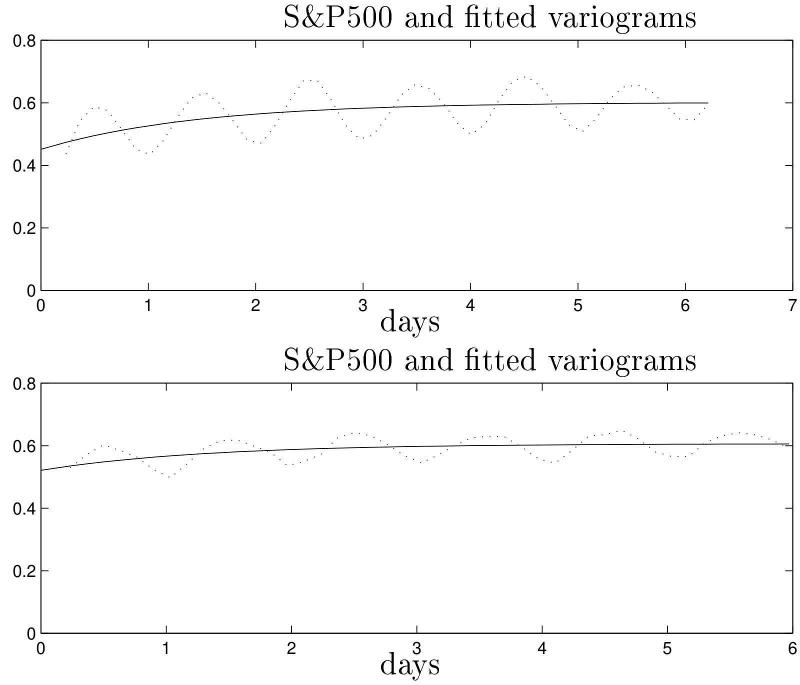

Figure 14: *Sensitivity to the length of the averaging interval.*

[13] J.-P. Fouque, G. Papanicolaou, and R. Sircar. *Derivatives in Financial Markets with Stochastic Volatility.* Cambridge Univerisy Press, 2000.

[14] Institute of Business and Economic Research, University of California at Berkeley. *The Berkeley Options Data Base User's Guide*, June 1995. Report number 1922.

[15] B. LeBaron. Stochastic volatility as a simple generator of apparent financial power laws and long memory. *Quantitative Finance*, 1(6):621–31, November 2001.

[16] J. Stein. Overreactions in the options market. *J. Finance*, XLIV(4):1011–1023, 1989.