A model for intraday volatility

Yongyang Cai, Baeho Kim, Matthew Leduc,
Kamil Szczegot, Yang Yixiao, Manuel Zamfir

June 15, 2007

Abstract

In this paper, we build an intraday model for volatility based on price change intensity. The quantity we model is thus named "volatensity". The model is a combination of an Autoregressive Conditional Duration (ACD) structure resembling that of Engle and Russel (1998) and an additional term, inspired by the literature on Hawkes processes. The ACD structure allows us to capture the long memory property of volatility using intraday information on price change events. The Hawkes part allows us to controls the speed of decay of volatility after jumps. Both the ACD term and the Hawkes term confer the model appropriate self-exciting (clustering) properties. The model is fitted to market data using MLE and simulations show a very accurate model fit. It also allows us to produce an intraday forecast of volatility and, by extension, a daily forecast.

Key words: intraday volatility model, volatensity, ACD, Hawkes processes, MLE, thinning method, intraday volatensity forecast.

Contents

1 Introduction

With the rapid development in computing power and storage capacity, data are being collected and analyzed at very high frequencies. In financial market, usually the quantity which the transaction events happened in a period of time is the key economic variable to be modeled or forecast, so it is natural to study the transaction timing. The financial market microstructure theories are typically tested on a transaction by transaction basis so again the timing of these transactions can be central to understanding the economics.

Transaction data inherently arrive in irregular time intervals, while standard econometric techniques are based on fixed time interval analysis. So a new technique has been researched as the alternative of the fixed interval analysis. The arrival times are treated as random variables which follow a point process. The dependence of the conditional intensity on past durations suggests that the model be called the autoregressive conditional duration (ACD) model. In this report, a new model that combines ACD model and Hawkes process for stock intraday or interday volatility is formulated. The conditional intensity function is parameterized in terms of past events in a way that seems particularly well suited for the transactions process. The fundamental application of the model is to measure and forecast the volatility of the stock transaction arrivals which is essentially the instantaneous quantity of transactions.

The following section introduces the basic concepts of the "volatensity" and Section 2 and Section 3 gives the detailed mathematical introduction for this model. Section 4 discusses the empirical data manipulation and analysis from the NYSE stock transaction database. Section 5 discusses the parameter calibration and volatility forecasting. Finally, Section 6 concludes and discuss future work.

2 The Volatensity

Let S_t be the price of a stock following the diffusion:

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma_t dW_t$$

where W_t is a standard Brownian Motion.

Let $y_t = \log(S_t)$ be the return on the stock. Then y_t satisfies:

$$dy_t = \left(\mu_t - \frac{\sigma_t^2}{2}\right) dt + \sigma_t dW_t$$

hence

$$\begin{aligned} \sigma_t^2 &= \lim_{h \to 0} \frac{1}{h} E\{(y_{t+h} - y_t)^2 | \mathcal{F}_t\} \\ &= \lim_{h \to 0} \frac{1}{h} \int_0^\infty P\{(y_{t+h} - y_t)^2 \geq \alpha | \mathcal{F}_t\} d\alpha \end{aligned} \tag{1}$$

Approximating the RHS of (1) by a Riemann sum on a compact interval, we get:

$$\begin{aligned}\sigma_t^2 &\approx \lim_{h\to 0}\frac{1}{h}\sum_{k=1}^{K} P\{|y_{t+h}-y_t|\geq\sqrt{k\epsilon}|\mathcal{F}_t\}\epsilon \\ &= \sum_{k=1}^{K}\lim_{h\to 0}\frac{1}{h}\,P\{|y_{t+h}-y_t|\geq\sqrt{k\epsilon}|\mathcal{F}_t\}\epsilon \end{aligned} \tag{2}$$

Let $N_t^{(\Delta y)}$ be the counting process that counts the number of jumps of y_t exceeding the fixed threshold Δy. Then its intensity $\lambda_t^{(\Delta y)}$ is given by:

$$\begin{aligned}\lambda_t^{(\Delta y)} &= \lim_{h\to 0}\frac{1}{h}\,P\{N_{t+h}^{(\Delta y)}-N_t^{(\Delta y)}>0|\mathcal{F}_t\} \\ &= \lim_{h\to 0}\frac{1}{h}\,P\{|y_{t+h}-y_t|>\Delta y|\mathcal{F}_t\}\end{aligned} \tag{3}$$

We will call the above intensity the **volatensity**. Also, for clarity of exposition, whenever the threshold Δy is understood from the context, it will be omitted from the notation.

Under this notation, equation (2) rewrites:

$$\sigma_t^2 \approx \sum_{k=0}^{K}\lambda_t^{(\sqrt{k\epsilon})}\epsilon \tag{4}$$

so the volatility of the stock can be approximated by a finite linear combination of volatensities.

Due to the strong correlation between volatensities, we have observed heuristically that even one volatensity will make a good proxy for the actual volatility. Figure 1 depicts a comparison between the GARCH(1,1) model for volatility and the volatensity for the IBM stock, based on TAQ data from 01/05/2007. In the next sections we will present and analyze a model for volatensity.

3 ACD+Hawkes Model for Volatensity

3.1 Definition of the Model

Since the volatensity will be used a proxy for volatility, one would like to incorporate in the model the so-called "stylized facts of volatility" observed in practice:

1. Self-exciting process.
2. Power-law decay between jumps.
3. Slow autocorrelation decay.
4. Close to log-normal distribution.

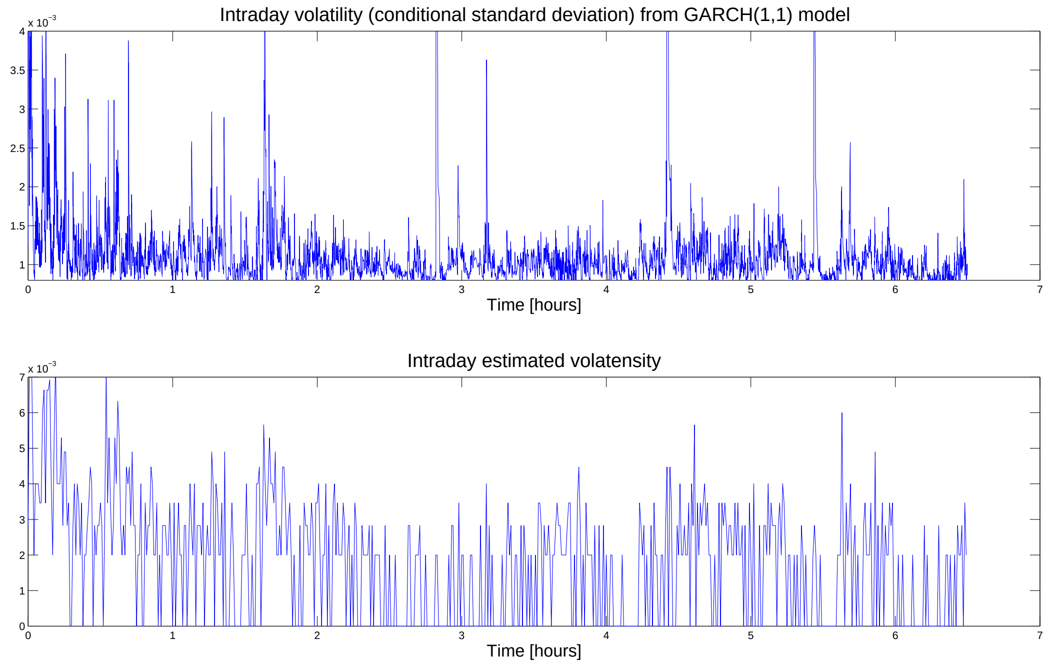

Figure 1: Comparison between GARCH(1,1) fitted into the return series and empirical volatensity.

Our proposed model is a hybrid ACD and Hawes process:

$$\lambda_t = \underbrace{\lambda_\infty(N_t)}_{\text{ACD part}} + \underbrace{\int_0^t g(t-s)dN_s}_{\text{Hawkes part}} \tag{5}$$

where

$$\frac{1}{\lambda_\infty(n)} = C + \sum_{i=1}^{m} \alpha_i (T_{n-i} - T_{n-i-1}) + \sum_{j=1}^{q} \beta_j \frac{1}{\lambda_\infty(n-j)} \tag{6}$$

and T_i is the time of the ith jump.

We will consider two forms for the decay function g:

$$g_{\text{general}}(x) = \frac{1}{\sum_{l=0}^{L} c_l x^l} \tag{7}$$

and

$$g_{\text{special}}(x) = \frac{a}{(b+x)^c} \tag{8}$$

	ACD part	Hawkes part
self-exciting	√	√
slow autocorrelation decay	√	
power-law decay of intensity between jumps		√
suitability for intraday modelling	√	√
ease of simulation	√	√
feasibility of calibration	√	√

Table 1: Comparison between ACD part and Hawkes part.

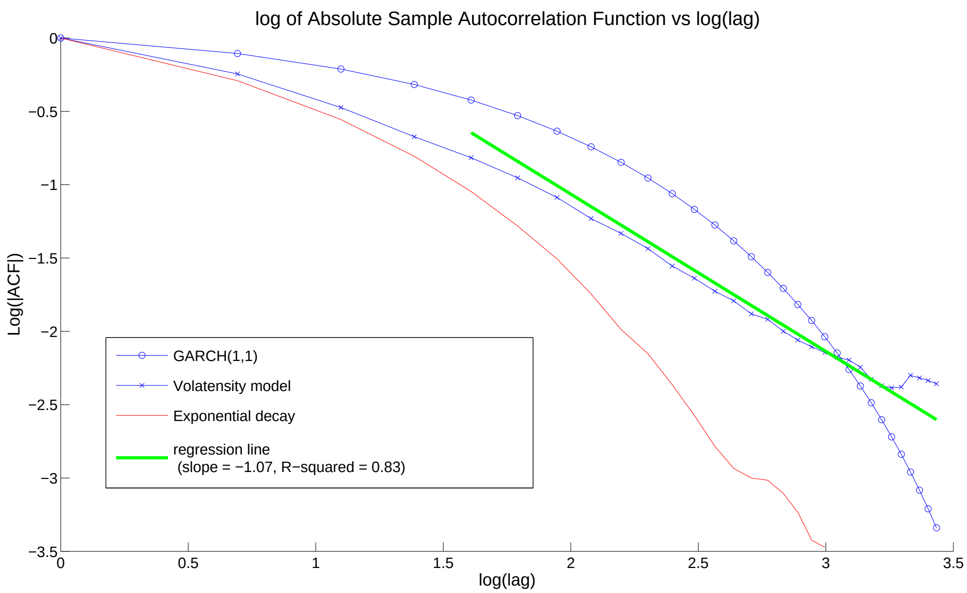

Figure 2: Log of absolute sample autocorrelation function vs log(lag).

3.2 Justification of the Model

We elected to create our model by combining the ACD and Hawkes models, in order to be able to incorporate the properties (mentioned above) of the volatility observed in practice. As the table below and also Figure 2 shows, just the Hawkes part would not be able to produce slow decay of autocorrelation, while just the ACD part would not be able to generate power-law decay between jumps.

A second reason to use the combined model is that while the Hawkes term generates short-term volatility bursts, the ACD part generates medium and long-term one. The Hawkes part is responsible for the narrow peaks, while the ACD part is responsible for the wave-shaped base level of the process. This is shown in the following figure that depicts a sample path of the process.

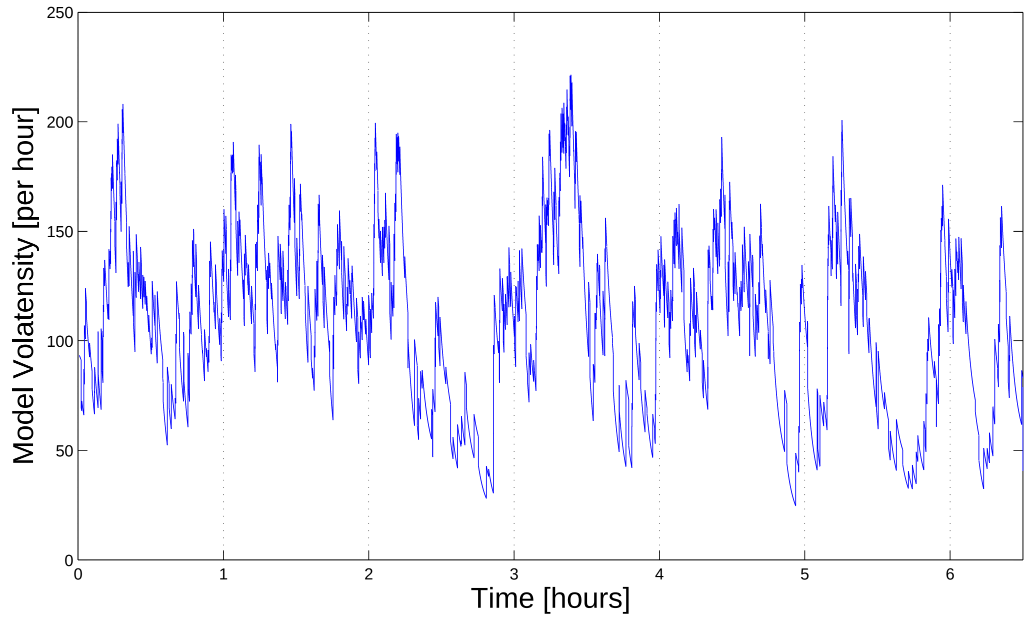

Figure 3: Reproducing stylized features of volatensity.

4 Empirical Volatensity

Transaction data inherently arrive in irregular time intervals, while standard econometric techniques are based on fixed time interval analysis. Frequently traded stocks will have transactions every few seconds. If a short time interval is chosen, there will be many intervals with no new information and heteroskedasticity of a particular form will be introduced into the data. On the other hand, if a long interval is chosen, the micro structure features of the data will be lost. In particular, multiple transactions will be averaged and the characteristics and timing relations of individual transactions will be lost, mitigating the advantages of moving to transaction data in the first place.

The problem becomes more complicated when one realizes that the rate of arrival of transaction type data may vary over the course of the day, week, or year making the choice of an "optimal" interval more difficult. For stocks, Studies have found a daily pattern over the course of the trading day (see Engle & Russell 1998): activity is higher near the open and the close then in the middle of the day. For currency markets, there are clear periods of high and low activity as markets around the world open and close.

This paper applies the volatensity model to IBM transactions data. The data were abstracted from the NYSE Trade and Quote (TAQ) database. The data set contains intraday transactions data (trades and quotes) for all securities listed on the New York Stock Exchange (NYSE) and American Stock Exchange (AMEX), as well as Nasdaq National Market System (NMS) and SmallCap issues. The original data is given with one second accuracy. For transactions falling at the same second, we take the average price as our QUOTE for that second. And our program takes out national holidays, weekends and off-hours from multiday data, as well as any data before/after trading hours 9:30-16:00. The volatensity models defines an event arrival using transaction arrival times for which the price has changed and absolute log return between two consecutive transactions is larger than the threshold δy. An event arrival time is now characterized not just by a transaction occurring at a new price, but a transaction occurring at a new price such that the absolute log return between this transaction and its last transaction is higher than δy. Thus, the empirical volatensity process can be obtained by the average number of events basesd on fixed time interval analysis.

We download intraday data (IBM, 01/05/2007, 9:30am-4:00pm) from TAQ[1].

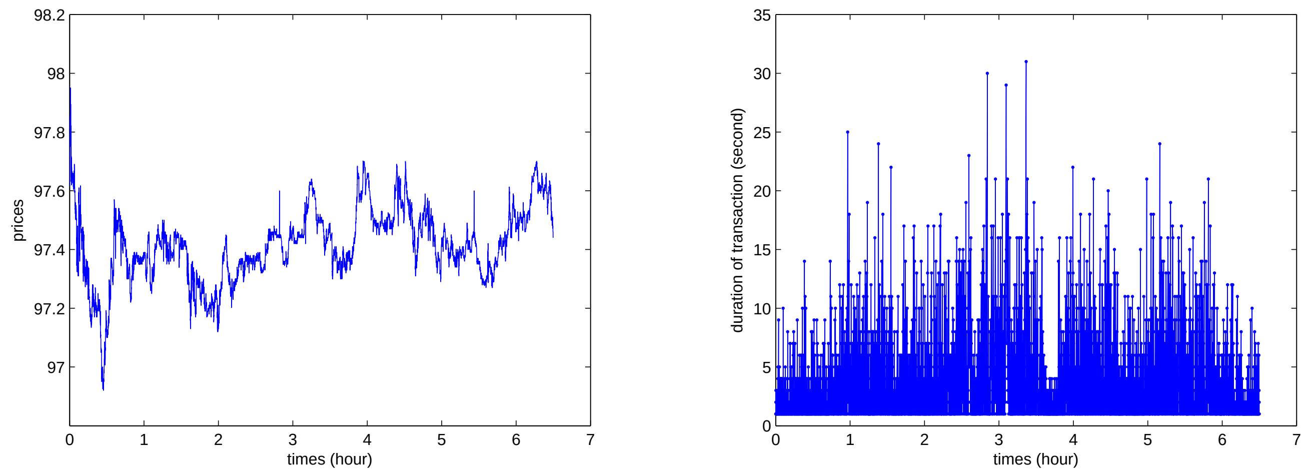

(a) Dynamics of IBM price.

(b) duration on transactions.

Figure 4: Empirical analysis on IBM, 01/05/2007, 9:30am-4:00pm.

Figure 4 (a) shows the dynamics of actual intraday price of IBM. Once we decide appropriate level of threshold to define price event, we can compute event durations from the historical data as described in Figure 4 (b). Then we calculate and draw the volatensity from the input transaction or price jump times. To obtain the average intraday volatensity graph for IBM in Figure 5, we assume that window size is 3 minutes and the price jump threshold is 0.00003 which means that

$$\left| \log \left(\frac{\text{price}^{IBM}(i)}{\text{price}^{IBM}(i-1)} \right) \right| > 0.00003. \tag{9}$$

5 Parameter Calibration and Volatensity Forecast

5.1 Maximum Likelihood Estimation of Parameters

Having observed historical price event times $\{t_i\}_{i\in\{1,\cdots n\}}$, we need to calibrate parameters by MLE using the following log-likelihood function:

$$\log \mathcal{L}\left(\theta; \{t_i\}_{i\in\{1,\cdots n\}}\right) = \sum_{i=1}^{n} \left\{ -\int_{t_{i-1}}^{t_i} \lambda_\theta(t)dt + \log \lambda_\theta(t_i) \right\}. \tag{10}$$

[1]https://wrds.wharton.upenn.edu/wrdsauth/members.cgi

First step to do is cleaning and converting the market data. We deleted trades outside the standard trading hours from 9:30am to 4:00pm. Then we took average the trade price in single time points, and determined the counting process based on a fixed threshold.

As a next step, we estimate model parameters using two-step procedure. Before the actual MLE, we performed least-squares fitting of the volatensity curve to the empirical volatensity to obtain a starting point. Then, we performed MLE estimation of (10) using the starting point numerically. Figure 5 illustrates the estimated historical volatensity, Model fitted in-data volatensity, and averaged out-of-data volatensity. After applying our two-step estimation procedure using `fmincon` function in MATLAB, we get

1. sample mean = 1.00
2. sample variance = 1.15
3. Ljung-Box test (with lags 10-80) accepts the null hypothesis that the model fit is adequate
4. Nice Q-Q plot visualization as Figure 6 (a)
5. Nice KS-test visualization as Figure 6 (b)

For statistical goodness-of-fit tests, we just applied Meyer's (1971) theorem that time-changed interarrival times should be independent exponentials.

5.2 Thinning Methods to Simulate Point Processes

Lewis and Shedler (1978) suggests a simple and relatively efficient method by thinning for simulating nonhomogeneous Poisson process. Moreover, Ogata (1981) extended the applications of Lewis' thinning simulation algorithm to any point processes that is absolutely continuous with respect to the standard Poisson process. The thinning method to simulate point processes is outlined in Algorithm 1. Note that the thinning methods require only evaluations of the conditional intensity function at certain points. This enables the simulation to be a lot faster without loss of its theoretical robustness.

Algorithm 1 The algorithm to simulate default times $\{t_j^*\}_{j\geq 1}$ by thinning method.

(Step 1) Fix time horizon $T > 0$ and initial value $\lambda(0)$.

(Step 2) Set $j = 1$ and $\tau_0 = 0$.

(Step 3) Construct $\overline{\lambda} > 0$ such that $\lambda(t) \leq \overline{\lambda}$ for all $t \in [\tau_0, T]$ almost surely.

(Step 4) Generate $\tau^* \sim \exp(\overline{\lambda})$ as the first homogeneous Poisson arrival time with intensity $\overline{\lambda}$.

(Step 5) If $\tau_0 + \tau^* \leq T$, then set $\bar{t} = \tau_0 + \tau^*$ as a candidate default time. Otherwise, stop.

(Step 6) Generate $u \sim unif(0, 1)$ independently.

(Step 7) If $u \leq \frac{\lambda(\bar{t})}{\overline{\lambda}}$, set $t_j^* = \bar{t}$ with updating $\lambda(t_j^*)$ by adding self-affecting term and set $j := j + 1$.

(Step 8) Set $\tau_0 = \bar{t}$ and go to **(Step 3)**.

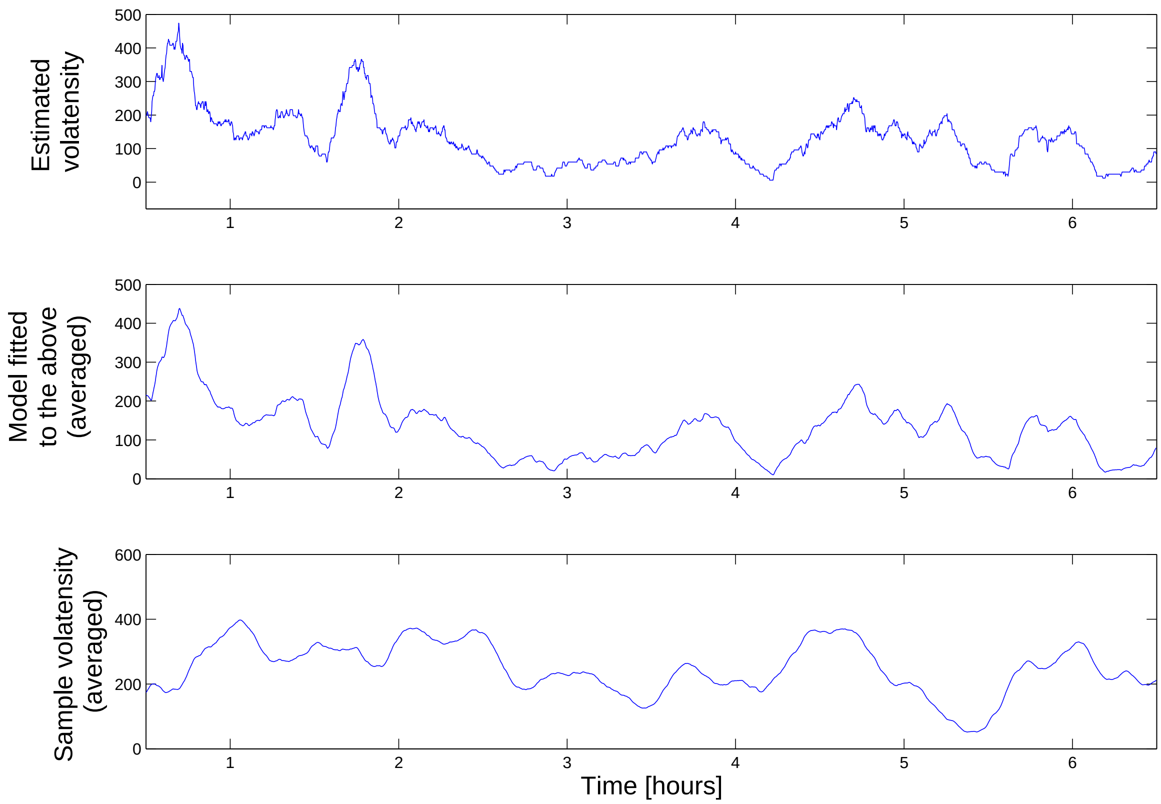

Figure 5: Historical vs. fitted volatensity

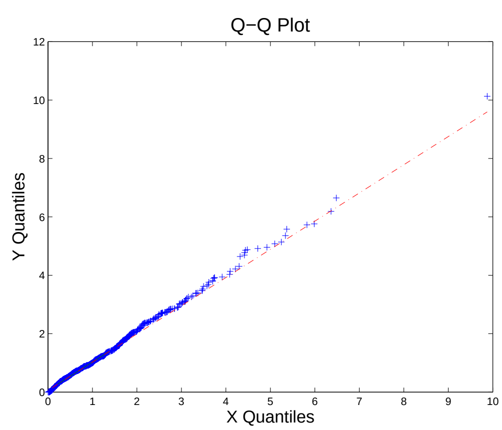

(a) Q-Q Plot

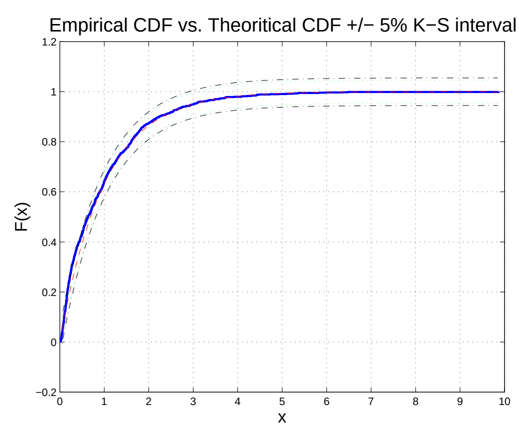

(b) KS Test

Figure 6: Statistical goodness-of-fit tests of parameter estimation.

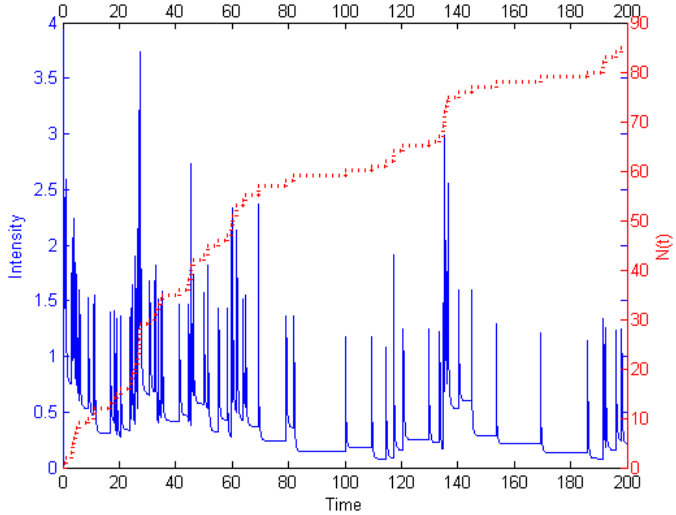

Figure 7: A sample path of intensity and the corresponding counting process via thinning method.

In its simple and efficient implementation, the method obviates the need for numerical integration and reverse-engineering of the intensity process. In fact, the key to this research is to apply the thinning methods to simulate self-affecting point processes effectively.

Hence, once we have calibrated parameters to the market data, we can forecast the volatensity in the future by simulation. Figure 8 illustrates the result of calibration and forecasting volatensity with 90% confidence interval.

6 Conclusion and Future work

We have built a model that combines the properties of the Autoregressive Conditional Duration (ACD) by Engle and Russel (1998) and Hawkes processes. The model was fitted using MLE and it was shown that the model presented the desired features of volatility. In fact, the model exhibits clustering (self-exciting behavior), long memory (power-law decay in autocorrelation) and power law decay in volatility after a jump (this decay is driven by the Hawkes term). A next step would be to test the model using a volatility trading strategy. Such a strategy can be constructed with straddles. It suffices to form a straddle portfolio daily with a put and a call of maturity between 15 and 60 days and strike price as close as possible to the underlying so as to have a portfolio as close to being delta neutral as possible. The daily return on such a portfolio is

$$R_t = \frac{(C_t + P_t) - (C_{t-1} + P_{t-1})}{C_{t-1} + P_{t-1}}$$

A way to test the model is to construct a daily forecast for volatility (scaled appropriately) and use it in the Black-Scholes formula to compute the next day's value of the straddle. If it is higher than the current day's value, we may take a long position in the straddle and vice versa. This would allow to test the model's ability to produce a daily forecast.

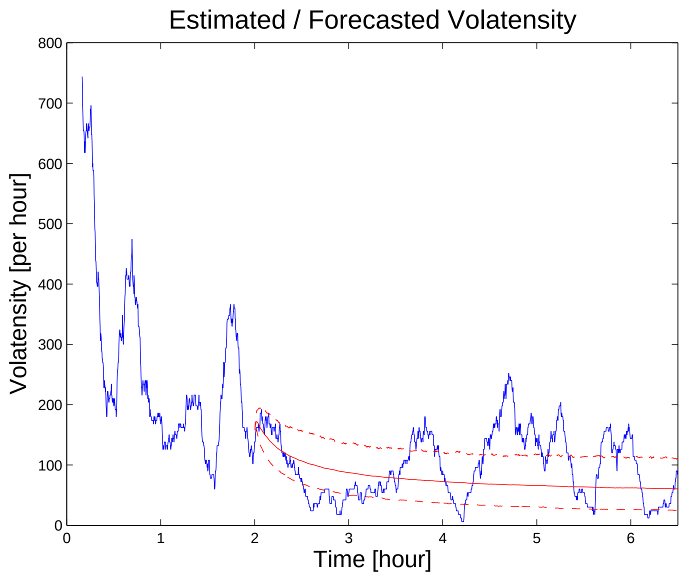

Figure 8: Forecasted volatensity after parameter estimation with 90% confidence interval.

It may be the case that an extra term to take the leverage effect into account would improve the model's performance. Such a model would then have the form

$$\lambda_\infty(n)^{-1} = C + \sum_{i=1}^{m} \alpha_i (T_{n-i} - T_{n-i-1}) + \sum_{j=1}^{q} \beta_j \lambda_\infty (n-j)^{-1}$$
$$+ L l_{\{(P_{T_{n-1}} - P_{T_{n-2}})<0\}} (P_{T_{n-1}} - P_{T_{n-2}})$$

where $l_{\{(P_{T_{n-1}} - P_{T_{n-2}})<0\}}$ is an indicator function that is 1 when the condition is fulfilled and 0 otherwise. This is left for future work.

References

[1] Lisa Borland and Jean Philippe Bouchaud, *On a multi-timescale statistical feedback model for volatility fluctuations*, 2005, arXiv.org:physics/0507073

[2] Engle, R.F. and J.R. Russell, *Autoregressive Conditional Duration: a New Model for Irregularly Spaced Transaction Data*, 1998, Econometrica, 66, 1127-1162.

[3] Eymen Erris, Kay Giesecke and Lisa R. Goldberg, *Pricing Credit from the Top Down with Affine Point Processes*, 2006, Working Paper, Stanford University.

[4] Kay Giesecke and Pascal Tomecek, *Dependent Events and Changes of Time*, 2005, Working Paper, Cornell University.

[5] Alan G. Hawkes, *Spectra of Some Self-exciting and Mutually Exciting Point Processes*, Biometrika, 58(1):83-90, 1971.

[6] P.A.W. Lewis and G.S. Shedler, *Simulation of Nonhomogeneous Poisson Processes by Thinning*, Naval Logistics Quarterly, 26(3):403-413, 1979.

[7] Yosihiko Ogata, *On Lewis' Simulation Method for Point Processes*, IEEE Information Theory, IT-27, pp. 23-31, 1986.