Trading Volume and Stochastic Volatility

Sam Howison and David Lamper

Oxford Centre for Industrial and Applied Mathematics,
Mathematical Institute, 24–29 St. Giles', Oxford OX1 3LB, UK.

Abstract. This paper develops a subordinated stochastic process model for asset prices, where the directing process is identified as information. Motivated by recent empirical and theoretical work, we make use of the under-used market statistic of transaction count as a suitable proxy for the information flow. An option pricing formula is derived, and comparisons with stochastic volatility models are drawn.

1 Introduction

Derivative pricing depends crucially on the assumptions made concerning the distributional properties of the asset price. Without some model of the underlying price process, it is impossible to price a derivative. The standard model of asset prices for derivative pricing is

$$\frac{dS}{S} = \mu\ dt + \sigma\ dX$$

in which μ and σ are constants. The model works well, but the predictions do not agree completely with market prices. There are many variations on the lognormality assumption, but it is likely that there is no universal remedy. Popular approaches include: other random walks, e.g. constant elasticity of variance [1]; local volatility surfaces [13]; stochastic volatility models [6,10]; other stochastic processes, e.g. the hyperbolic distribution [4] or truncated Lévy [11]; and more general statistical approaches [2].

There is also a vast literature on market microstructure; for an overview see [12]. However, it is not easy to make the link with more global (longer time-scale) problems such as derivative pricing. Many models of price formation have been proposed, with the key element that price movements are primarily due to the arrival of information in the form of buy/sell orders.

In this paper we formulate a model based on the old Wall Street adage that "it takes volume to move prices". The idea is to use some measure of trading volume as a proxy for the information events, and relate price changes to such information events rather than simply the passage of calendar time. This is supported by recent empirical evidence by Geman [7,8], who found that the asset price follows a geometric Brownian motion with respect to a "timescale" (stochastic variable) defined by the number of transactions. In this framework the stock price is a good candidate to be described by a subordinated stochastic process model, which we will introduce in the next section.

2 Subordinated stochastic processes

Discrete stochastic processes are indexed by a discrete variable, usually time, in a straightforward manner: $X(0), X(1), \ldots, X(t), \ldots$; here $X(t)$ is the value that a particular realization of the stochastic process assumes at time t. Instead of indexing by the integers $0, 1, 2, \ldots$ the process could be indexed by a set of numbers $\tau_1, \tau_2, \ldots$ where these numbers are themselves a realisation of a stochastic process with positive increments. That is, if $\tau(t)$ is a positive and increasing stochastic process, a new process $X(\tau(t))$ may be formed. The resulting nonstationary process $X(\tau(t))$ is said to be *subordinated* to $X(t)$, called the *parent process*, and is *directed* by $\tau(t)$, called the *directing process* or the *subordinator* [5]. If the increments of the directing process $\tau(t)$ are not independent, this technique is known as a general stochastic time change. The process $\tau(t)$ is often referred to as a "stochastic clock".

Many problems can be formulated in terms of subordinated processes. Take, for example, the expected remaining lifetime of a hard drive, denoted by X, which may be considered a stochastic function of time, i.e. $X(t)$, with a negative drift subject to the condition it is always positive. However, the expected lifetime of a hard drive in a server environment is clearly different from a desktop, and the lifetime process may alternatively be considered a stochastic function of the total amount of data transferred by the drive, i.e. $X(n)$, where n itself is a stochastic function of time. In situations where the drive is accessed continuously, the process describing the remaining lifetime of the drive will evolve quickly. Conversely, when the drive is little used, X evolves slowly with time. Formulating this example in the subordinated process framework, the expected remaining lifetime of the drive is the parent process; the amount of data transferred is the directing process.

Formulating our model in a subordinated process framework, the total number of information arrivals, denoted by $n(t)$, is assumed to drive the market, i.e. $n(t)$ represents the directing process of the market. The timescale regulated by $n(t)$ is hereafter referred to as information time, and is distinct from calendar time t. Both the asset price S and the cumulative number of trades N are dependent on the number of information arrivals, and are regarded as the observable parent processes. We utilise the observation by Geman [8,7] that a stationary, lognormal distribution for S can be achieved through a stochastic time change, where the directing process is found to be well approximated by the number of trades (up to a constant). This is compatible with the popular notion of the number of transactions being a guide to the pace of market activity.

3 Market model

The total price change of an asset in any fixed calendar time interval, such as a day, reflects the accumulation of a random number of arrivals of many small bits of information, proxied by the number of trades. On less eventful days trading is slow and prices evolve slowly, whereas prices evolve faster with heavier trading when more information arrives. Empirical evidence has shown that subordinated processes represent well the price changes of stocks and futures. As early as 1973, Clark [3] applied subordinated processes to cotton futures data.

3.1 Information arrivals process

The total number of information arrivals $n(t)$ is assumed to be large, yet display a significant daily variation. It must be a positive, increasing function of time, and hence $dn \geq 0$. We assume that there is a positive rate, or intensity, of information arrivals $I(t)$. Hence $n(t)$ is defined as the solution of

$$dn = I(t)\ dt. \tag{1}$$

We propose to model $I(t)$ by the stochastic process

$$dI = p(I,t)\ dt + q(I,t)\ dX_t^{(2)},$$

where p and q are as yet unspecified functions of I and t, which must, however, be such that $I(t) \geq 0$, and $dX_t^{(2)}$ is the increment of Brownian motion in calendar time, i.e. $[dX_t^{(2)}]^2 = dt$.

The directing process $n(t)$ is well approximated by the cumulative number of transactions $N(t)$. The number of transactions in a given time interval, ΔN, subject to the addition of a constant, is assumed to be equal to the change in value of the directing process, Δn. Thus moments of the directing process greater than 1 are directly matched by the moments of the number of transactions.

3.2 Asset price models

A stochastic time change is made from calendar time to information time to achieve a stationary, lognormal model of the asset price S in the informational timescale. Alternatively, this may be regarded as a change in the frame of reference or as time deformation, since the relevant time scale promoting normality of returns is no longer calendar time but information time. Hence the subordinated process S can be described by the usual lognormal random walk in this timescale:

$$dS = \mu_n S\ dn + \sigma_n S\ dX_{n(t)}^{(1)}, \tag{2}$$

where μ_n and σ_n are constants representing the drift and volatility of the asset return per information event respectively. The increment of Brownian motion $dX_{n(t)}^{(1)}$ evolves in the informational timescale, i.e.

$$\left[dX_{n(t)}^{(1)}\right]^2 = dn(t) \ = \ I(t)\ dt$$

from (1), and is distinct from, but may be correlated with, $dX^{(2)}$.

The return in a time interval Δt at time t, $R_{\Delta t}(t)$, can be expressed as

$$R_{\Delta t}(t) \mid \Delta n \sim N(\mu_n \Delta n, \sigma_n^2 \Delta n),$$

where $\Delta n = n(t + \Delta t) - n(t)$, indicating the conditional normality of this process. The variance, conditional on the value of Δn, is

$$\mathrm{Var}[R_{\Delta t} \mid \Delta n] = \sigma_n^2 \Delta n.$$

However, the number of information arrivals Δn in a time period Δt is not constant, but is a stochastic variable. Thus $R_{\Delta t}$ exhibits *conditional heteroskedasticity*, that is the conditional variance of $R_{\Delta t}$ is not constant. Furthermore, if Δn were assumed to be serially correlated in a discrete setting, the variance of $R_{\Delta t}$ would be an ARCH process.

Using equation (1), we can rewrite (2) as

$$dS = \mu_n S I(t)\ dt + \sigma_n S \sqrt{I(t)}\ dX_t^{(1)},$$

where $dX_t^{(1)}$ is the realisation of $dX^{(1)}$ in calendar time. It can be seen that the rate of information arrivals $I(t)$ drives the volatility of stock returns in calendar time. Thus price variability in our model depends on the flow of information into the market, both in the drift and volatility terms. Volatility is often associated with the amount of information arriving into the market, and this model proposes that stochastic volatility is directly linked to the rate of information flow $I(t)$.

We now propose a specific model for the information flow, modelling $I(t)$ by the mean reverting random walk

$$dI = \alpha(\mu_I - I)\ dt + \beta I^{1/2}\ dX_t^{(2)},$$

where α represents the rate of mean reversion and μ_I is the long-run mean-level of I. The information intensity $I(t)$ is a hidden process and is not directly observable. However, in this model the moments greater than one of the directing process $n(t)$ are identical to the moments of the number of transactions. This observation can be utilised to obtain information about the underlying process $I(t)$. The rate of mean reversion α can be obtained by considering how the variance of the number of transactions ΔN scales with time, or looking at the autocovariance. A comparison of the two separate estimates of α gives an approximate error indication. Fits of the variance and covariance were successful on many of the 'old economy' stocks, where there was typically no strong trending element in the number of transactions and relatively few days with exceptional trading behaviour, see e.g. Fig. 1. The remaining parameters can be estimated from the price series.

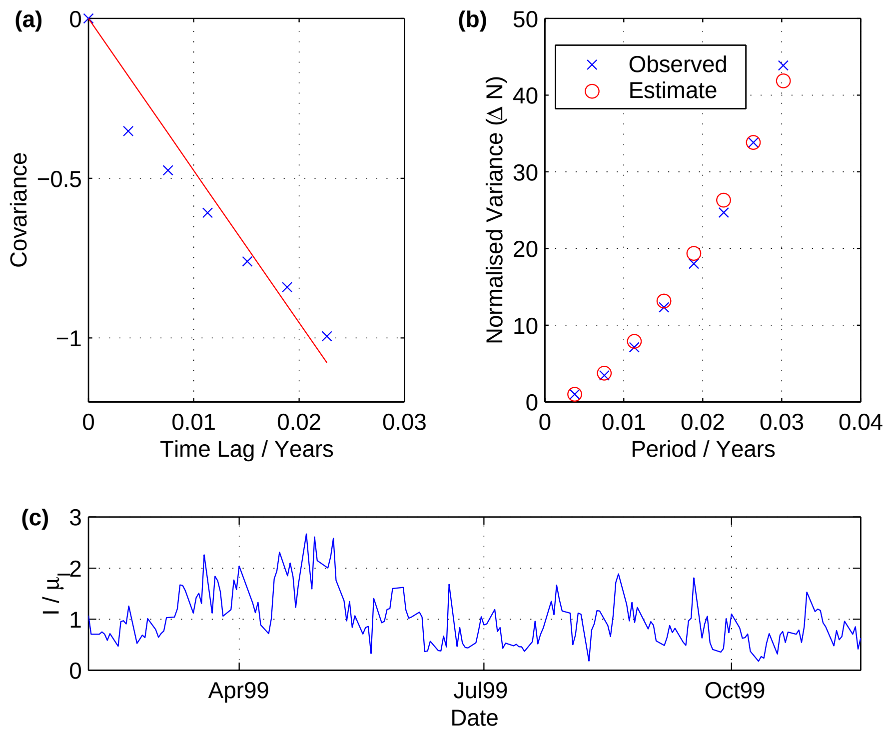

Fig. 1. Estimation of parameters of $I(t)$ from transaction count data for Thames Water (Source: Primark Datastream). (**a**) Estimation of α using lagged covariance data. (**b**) Scaling of variance with Δt. (**c**) Extracted $I(t)$ over estimation window

4 Concluding remarks and applications to derivative pricing

The rate of information arrivals is not a traded asset. Unlike the Black–Scholes case it is no longer sufficient to hedge solely with the underlying asset, but nevertheless arbitrage assumptions force the prices of different derivative products to be mathematically consistent. Because we have two sources of randomness, we set up a portfolio containing one option, with value denoted by $V(S, I, t)$, a quantity $-\Delta$ of the asset and a quantity $-\Delta_1$ of a separate liquid option with value $V_1(S, I, t)$ in a manner exactly analogous to stochastic volatility models [6]. The option price can then be expressed as a solution of the parabolic partial differential equation

$$\begin{aligned}\frac{\partial V}{\partial t} + \frac{1}{2}\sigma_n^2 S^2 I \frac{\partial^2 V}{\partial S^2} + \frac{1}{2} q^2 \frac{\partial^2 V}{\partial I^2} + \rho \sigma_n S q \sqrt{I} \frac{\partial^2 V}{\partial S \partial I} \\ + rS\frac{\partial V}{\partial S} + (p - \lambda q)\frac{\partial V}{\partial I} - rV \ = \ 0.\end{aligned}$$

Here $\lambda(S, I, t)$ is the market price of (information arrival intensity) risk. Clearly, the observed volatility depends on the rate of information arrivals I. When a large amount of information is arriving in the market place, I is above average, our stochastic clock runs faster and the observed asset volatility increases. Hence it is natural to interpret our model as a stochastic volatility model. There are a number of possible approaches to solving this model, both numerical and asymptotic and space does not permit a fuller discussion here; for further details see [9].

References

1. Stan Beckers, *The constant elasticity of variance model and its implications for option pricing*, Journal of Finance **35** (1980), no. 3, 661–673.
2. Jean-Philippe Bouchaud and Marc Potters, *Theory of financial risks*, Cambridge University Press, 2000.
3. Peter K. Clark, *A subordinated stochastic process model with finite variance for speculative prices*, Econometica **41** (1973), no. 1, 135–155.
4. Ernst Eberlein and Ulrich Keller, *Hyperbolic distributions in finance*, Bernoulli **1** (1995), 281–299.
5. William Feller, *An introduction to probability theory and its applications*, vol. II, John Wiley and Sons, 1966.
6. Jean-Pierre Fouque, George Papanicolaou, and K. Ronnie Sircar, *Derivatives in financial markets with stochastic volatility*, Cambridge University Press, 2000.
7. Hélyette Geman and Thierry Ané, *Stochastic subordination*, Risk Magazine **9** (1996), 145–149.
8. ———, *Order flow, transaction clock, and normality of asset returns*, The Journal of Finance **55** (2000), no. 5, 2259–2284.
9. Sam Howison and David Lamper, *Trading volume in models of financial derivatives*, OCIAM / MFG Working Paper (2000).
10. Alan L. Lewis, *Option valuation under stochastic volatility: with Mathematica code*, Finance Press, 2000.
11. Andrew Matacz, *Financial modelling and option theory with the truncated Levy distribution*, International Journal of Theoretical and Applied Finance **3** (2000), no. 1, 143–160.
12. Maureen O'Hara, *Market microstructure theory*, Blackwell, 1995.
13. Riccardo Rebonato, *Volatility and correlation*, John Wiley, 1999.