Trading Volume in Models of Financial Derivatives

Sam Howison and David Lamper

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

June 6, 2001

Abstract

This paper develops a subordinated stochastic process model for the asset price, 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. Both the asset price and the number of trades are used in parameter estimation. The underlying process is found to be fast mean reverting, and this is exploited to perform an asymptotic expansion. The implied volatility skew is then used to calibrate the model.

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. There are many variations on the lognormality assumption in the Black–Scholes model of option pricing, including statistical approaches [9] and a vast literature based essentially on the properties of Brownian motion.

It has been well documented in the empirical literature that although stock returns are normally distributed on a timescale of a month or greater, they exhibit significant departures from normality when shorter horizons are considered [48]. The assumption that daily returns are normally distributed suffers many shortcomings, with the empirical distribution exhibiting fat tails and skewness [19, 43]. A number of alternatives have been proposed that attempt to describe the probability density of the asset returns more accurately, of which popular approaches include: hyperbolic distribution [16]; Student's t-distribution [6]; Lévy stable non-Gaussian model [43, 51]; truncated Lévy flight [44]; multi-fractal processes [42].

The discrete-time modelling of volatility has also been particularly popular in the econometrics literature [4, 8]. Volatility is not found to be constant, but to vary over time and exhibit positive serial correlation, i.e. volatility clustering. The most successful models have been the family of autoregressive conditional heteroskedastic (ARCH) models [17], and its extension into GARCH [7] and more recently EGARCH [46]. In their most general form, ARCH models make the conditional variance at time t a function of exogenous and lagged endogenous variables, time, parameters, and past prediction errors.

An alternative to pure empirical study is to consider the process of price evolution itself. Indeed, one of the dominant themes in the academic literature since the 1960's has been the concept of an efficient market, which, loosely, states that security prices fully reflect all available information [20, 21]. Thus the price of a security can be thought of as adjusting rapidly to incorporate new information [33], and will depend on the behaviour of the information that influences financial markets. The concept of information flow is easy to grasp, but difficult to quantify. It is widely believed, however, that measures of market activity such as trading volume are related to the information flow, and may be suitable proxies for this unobservable process.

Trading volume can be decomposed into two components: the number of trades, and the size of trades, frequently referred to as just 'volume'. There is little question that the number of trades is intimately connected with the fundamental mechanism of trading which compounds the new information into prices, and is indicative of such an information flow. Much empirical research has focused on the simultaneous link between price and volume; for a summary of current literature see [37, 25]. This research has found a strong positive correlation between volume and the absolute price change, supporting the old Wall Street adage that "it takes volume to move prices". Asymmetric patterns have been found by some researchers, suggesting that volume is larger when prices move up than when they move down. The observation of volume is also a popular tool for the technical trader [28]. Furthermore, ARCH effects have been shown to vanish when volume is also included as an explanatory variable in the conditional variance equation [39]. This implies that ARCH effects reflect time-dependence in the process generating information flow to the market.

The relationship between the number of trades and volatility has also recently been investigated [35, 29]. In regressions of volatility on both volume and the number of transactions, the volatility-volume relation is rendered statistically insignificant. That is, it is the occurrence of transactions per se, and not their size, that generates volatility. Indeed, the lack of trading itself has been shown to reduce volatility during the lunch break of the Tokyo Stock Exchange [32]. This supports the contention that volatility is driven by the identical factors that generate trades, i.e. information. It has been suggested [18, 5] that volume itself may reflect the extent of market disagreement on the information received.

From a market microstructure perspective, price movements are caused primarily through the arrival of information. The dynamics by which this information is incorporated into the current price is addressed in the market microstructure literature, where many models of price formation have been proposed; for an overview of this topic see O'Hara [47]. Theoretical analyses suggest that trading volumes in financial markets may be determined by: liquidity effects; information flows; asymmetric information (or differences in opinion); and the quality of information. Trade size has no role in some market models, e.g. in the Kyle model [38] volume is not a factor in the price adjustment process, but has increasingly been used as a measure of the information content of financial and macroeconomic events [1, 36, 52].

We conclude that the concept of information is an important factor in determining changes in stock value, and both the number and size of trades, especially the former, are closely related to it.

In this paper we posit that the speed of evolution of the asset price process is determined by the information flow, with calendar time playing a secondary role. It is further assumed that the current information flow can be indirectly observed through the number of trades. This is compatible with the popular notion of the number of transactions being a guide to the pace of market activity. On less eventful days trading is slow and prices evolve slowly, whereas prices evolve faster with heavier trading when more information arrives [13]. In this framework, the stock price is a good candidate to be described by a subordinated stochastic process model.

Empirical evidence has shown that subordinated processes represent well the price changes of stocks and futures. As early as 1973, Clark [14] applied subordinated processes to cotton futures data. In Clark's model the daily price change is the sum of a random number of intra-day price changes. The events that are important to the pricing of a security occur at a random, not uniform, rate through time. The variance of the daily price change is thus a random variable with a mean proportional to the mean number of daily transactions. Clark argues that the trading volume is positively correlated with the number of intra-day transactions, and so the trading volume is positively correlated with the variability of the price change. More recently, Geman [26, 27] has found that the asset price follows a geometric Brownian motion with respect to a "timescale" (stochastic variable) defined by the number of transactions. This work supports the view that volatility is stochastic in real (calendar) time because of random information arrivals, but that it may be modelled as being stationary in information time.

In §2 we introduce subordinated stochastic processes and define a model for the asset price, directed by an information process that is proxied by the number of transactions. The rate of information flow is proposed to follow a mean-reverting process in §2.4, and comparisons with a stochastic volatility model are subsequently made. Both the price and number of transactions are used to estimate the necessary model parameters, and the information flow is found to exhibit fast mean-reversion. An option pricing equation is derived in §3 and a numerical solution is then sought using both finite-difference and finite-element methods. Finally, in §3.4 an asymptotic expansion of the pricing equation is performed, which utilises the information contained within the implied volatility curve to calibrate the model.

2 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 days with many trades, there are usually larger than average price changes, indicating rapid price evolution. This can be noticed by observing market data, see e.g. Figure 1.

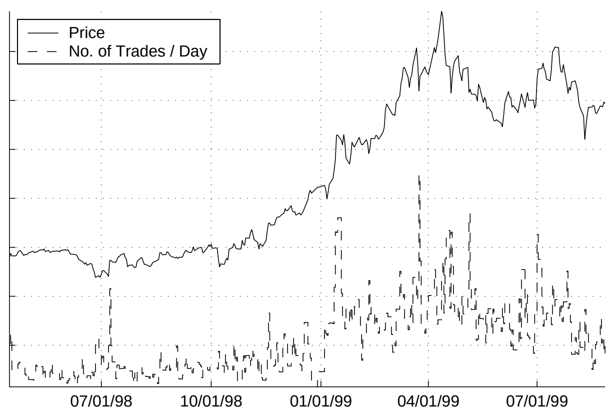

Figure 1: A price and transaction count chart for Dixons (Source: Primark DataStream). On days where a large number of transactions are realised, there is typically an increase in realised stock volatility.

2.1 Subordinated stochastic processes

Discrete stochastic processes are indexed by a discrete variable, usually time, in a straightforward manner: $X(0), X(1), \dots, X(t), \dots$; here $X(t)$ is the value that a particular realization of the stochastic process assumes at time t. Instead of indexing by the integers $0, 1, 2, \dots$ the process could be indexed by a set of numbers $\tau_1, \tau_2, \dots$ which 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 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* [22]. 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".

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 [27, 26] 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).

2.2 Information arrivals process: general considerations

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 work in continuous time and assume that there is a positive rate, or intensity, of information arrivals $I(t)$. Hence $n(t)$ is defined as

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

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

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

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 pricing of options in information time has also been considered by Chang *et al.* [11, 12], under the assumption that dn is a Poisson process.) The directing process $n(t)$ is related to the number of transactions and its estimation is discussed in §2.4.1.

2.3 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 timescale 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:

$$dS = \mu_n S\ dn + \sigma_n S\ dX_{n(t)}^{(1)} \tag{3}$$

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

$$R_{\Delta t}(t) \mid \Delta n \quad \sim \quad N(\mu_n \Delta n, \sigma_n^2 \Delta n), \tag{4}$$

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

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

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.

The unconditional centred moments of R are given by:

$$\begin{aligned}
\mathrm{E}[R_{\Delta t}] &= \mu_r = \mathrm{E}\Big[\mathrm{E}[R_{\Delta t} | \Delta n]\Big] = \mu_n \mathrm{E}[\Delta n], \\
\mathrm{Var}[R_{\Delta t}] &= \sigma_r^2 = \sigma_n^2 \mathrm{E}[\Delta n] + \mu_n^2 \mathrm{Var}[\Delta n], \\
\mathrm{m}_3[R_{\Delta t}] &= \mu_n^3 \mathrm{m}_3[\Delta n] + 3\mu_n \sigma_n^2 \mathrm{Var}[\Delta n], \\
\mathrm{m}_4[R_{\Delta t}] &= \mu_n^4 \mathrm{m}_4[\Delta n] + 6\mu_n^2 \sigma_n^2 \left(\mathrm{m}_3[\Delta n] + \mathrm{E}[\Delta n]\mathrm{Var}[\Delta n]\right) \\
&\quad + 3\sigma_n^4 \left(\mathrm{Var}[\Delta n] + (\mathrm{E}[\Delta n])^2\right),
\end{aligned}$$

where m_3 and m_4 represent the third and fourth centred moments respectively. The unconditional distribution of R is kurtotic and skewed compared to the normal distribution, because it is an average of diffuse (large Δn) and compact (small Δn) conditional densities. This can be seen by rewriting (4) as

$$R_{\Delta t} \quad = \quad \int_{\Delta n} N(\mu_n \Delta n, \sigma_n^2 \Delta n) p(\Delta n) d(\Delta n),$$

where $p(\Delta n)$ represents the probability distribution of Δn. This indicates that the unconditional distribution of the returns process is a mix of normal distributions, but is not itself normally distributed. The mixture of normal distributions hypothesis[1] (MDH), in which the asset return and trading volume are driven by the same underlying information flow or mixing variable, is a well-known representation of asset returns and has often been cited in the financial literature to

[1] In the mixture of distributions hypothesis a varying number of events occur each day that are relevant to the pricing of an asset. Let δ_{it} denote the ith intra-day equilibrium price increment on day t. This implies that the daily price increment, ϵ_t, is given by

$$\epsilon_t \quad = \quad \sum_{i=1}^{n_t} \delta_{it} \qquad \text{where} \quad \delta_{it} \sim_{i.i.d.} D(0, \sigma^2)$$

where D represents a symmetric distribution, and the variation in the mixing variable n_t, the number of events on day t, may be random, deterministic, and/or seasonal. In this setup it is clear that ϵ_t is drawn from a mixture of distributions, where the variance of each distribution depends on n_t. Furthermore, the volume of trades is also assumed to be related to the mixing variable n_t.

model the observed leptokurtosis in returns. Empirical tests are generally supportive of the model [30, 31], but a subsequent study was less encouraging [40].

Using equation (1), we can rewrite (3) as

$$dS = \mu_n S I(t)\; dt + \sigma_n S \sqrt{I(t)}\; dX_t^{(1)}, \tag{5}$$

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 $I(t)$. Empirical studies have found a strong link between the rate of information arrivals and observed short run volatility [24, 34, 50]. Ross [50] notes that in an arbitrage-free economy, the volatility of prices is directly related to the rate of flow of information to the market. Moreover, the I dependence of the drift in our model implies that an increase in volatility will result in an increase in the expected return. This is compatible with risk-averse agents who will demand compensation in the form of an increase in expected return for holding a risky asset, measured by the variance of such return.

2.4 The rate of information arrivals: a specific model

The process representing the rate of information arrivals $I(t)$ must take only positive values. Moreover, news arrivals are often positively autocorrelated. When an unanticipated news item occurs on a given day, more detailed disclosures tend to follow rapidly over the next few hours or days, and different interpretations of the circumstances leading to the event are formed. This tends to keep the story in the headlines for some time, suggesting that the information arrivals process should exhibit a positive autocorrelation, albeit over short time periods. Furthermore, it is reasonable to assume that the rate of information arrivals has a long-run equilibrium value, and the process is mean-reverting. On average, the amount of information released concerning an established company should not exhibit significant trending behaviour. Over a reasonable time period, a company (or an index or currency) may be either in or out of "the news", but the average frequency of such events does not tend to change significantly without a major change in the structure of the firm concerned. For a high growth company a trending information flow might be appropriate, but this will not be considered further in this paper.

With these considerations in mind, we model I by the mean-reverting random walk

$$dI = \alpha(\mu_I - I)\; dt + \beta I^{1/2}\; dX_t^{(2)}, \tag{6}$$

where α represents the rate of mean reversion and μ_I is the long-run mean-level of I. This is of the same form as the Cox, Ingersoll and Ross model of the interest rate [15]; this mean reverting process has also been used to model volatility directly [3]. Although this is just one possible distributional form it fulfils all the criteria stated above. It is not hard to generalise this approach (as in [23]) by defining I to be an explicit function of some stochastic process Y:

$$I = f(Y) \qquad \text{where} \qquad dY = P(Y,t)\ dt + Q(Y,t)\ dX_t^{(2)}.$$

This framework allows a specific underlying process Y to be used, and then I to be some function of this process.

A recent study of the distribution of stock return volatility itself [2] indicated that the unconditional distribution of the log standard deviation for a number of individual stocks in the Dow Jones Industrial Average all appeared approximately Gaussian. Since $\sqrt{I(t)}$ is related to the standard deviation of the asset in calendar time, see (5), this implies

$$d\Big[\log\left(\sqrt{I(t)}\right)\Big] = \alpha(\mu_I - \log\sqrt{I(t)})\ dt + \beta\ dX_t^{(2)}$$

which has the invariant distribution $N(\mu_I, \beta^2/2\alpha)$, giving

$$dI = 2\alpha I(\mu_I + \beta^2 - \log\sqrt{I})\ dt + 2\beta I\ dX_t^{(2)}$$

as a possible alternative for the evolution of I.

Returning to the model (6), the origin is non-attainable provided $2\alpha\mu_I/\beta^2 \geq 1$ [53]. The behaviour of I has the following properties: negative values are precluded; should I reach zero, it can subsequently become positive; the variance of I increases as I increases; there is a steady-state distribution for I. The probability density function for I at time t, conditional on its value at the current time s is given by [15]:

$$p(I(t),t;I(s),s) = c(s,t)e^{-u(s,t)-v(t)}\left(\frac{v(t)}{u(s,t)}\right)^{q/2} I_q[2(u(s,t)v(t))^{1/2}],$$

where

$$c(s,t) = \frac{2\alpha}{\beta^2(1-e^{-\alpha(t-s)})}, \qquad u(s,t) = cI(s)e^{-\alpha(t-s)},$$

$$v(t) = cI(t), \qquad q = \frac{2\alpha\mu_I}{\beta^2} - 1,$$

and I_q is a modified Bessel function of the first kind of order q.

The expected value and variance of I at time t, conditional on its value $I_0 = I(0)$, can be calculated from the differential form of I in (6):

$$\mathrm{E}[I(t)|I_0] = I_0 e^{-\alpha t} + \mu_I(1-e^{-\alpha t}).$$

The variance of I at time t, conditional on its initial value, is

$$\mathrm{Var}[I(t)|I_0] = \frac{\beta^2}{2\alpha}\left[\mu_I(1-e^{-\alpha t})^2 + 2I_0(e^{-\alpha t} - e^{-2\alpha t})\right],$$

from which the stationary variance is

$$\mathrm{Var}[I] \quad = \quad \sigma_I^2 \;=\; \frac{\mu_I \beta^2}{2\alpha}. \tag{7}$$

Finally, the conditional autocovariance can be derived. For $t > s > 0$,

$$\begin{aligned}
\mathrm{Cov}[I(t), I(s)|I_0] &= \mathrm{E}[I(t)I(s)|I_0] - \mathrm{E}[I(t)|I_0]\mathrm{E}[I(s)|I_0] \\
&= e^{-\alpha(t-s)}\mathrm{Var}[I(s)|I_0],
\end{aligned}$$

and unconditionally

$$\mathrm{Cov}[I(t), I(s)] \quad = \quad \sigma_I^2 e^{-\alpha|t-s|}. \tag{8}$$

The exponential rate of decorrelation of $I(t)$ is proportional to α, and so $1/\alpha$ can be thought of as a typical correlation time. Increasing α and keeping $\mathrm{Var}(I)$ fixed changes the degree of persistence of the mean reverting process I, without affecting the magnitude of the fluctuations. A large value of α will lead to burstiness, or clustering, in the driving process I. The common observation of volatility clustering in asset returns, that is, the tendency of large stock price changes to be followed by large stock price changes, but of unpredictable sign, can be modelled by fast mean-reversion of I and will be considered further in §3.4.

The invariant distribution $p_\infty(I)$ satisfies the steady-state forward Kolmogorov equation

$$\frac{1}{2}\frac{d^2}{dI^2}[\beta^2 I p_\infty(I)] - \frac{d}{dI}[\alpha(\mu_I - I)p_\infty(I)] \quad = \quad 0, \tag{9}$$

from which we find that $p_\infty(I)$ is the gamma distribution $\Gamma(\omega, \nu)$, with $\omega = 2\alpha/\beta^2$ and $\nu = 2\mu_I\alpha/\beta^2$. A positive skew and excess kurtosis are predicted:

$$\mathrm{Skew}[I] \quad = \quad \frac{2}{\sqrt{\omega\mu_I}} \qquad \text{and} \qquad \mathrm{Kurt}[I] \;=\; 3 + \frac{6\mu_I}{\omega}. \tag{10}$$

2.4.1 Parameter estimation

The information intensity $I(t)$ is a hidden process and is not directly observable. However the directing process $n(t)$ can be well approximated by the cumulative number of transactions $N(t)$. By treating S&P500 returns as a subordinated process, Geman [26] calculated the moments of a directing process necessary to achieve normality of these returns in the informational timescale through a numerical optimisation. Remarkably, the values of this directing process greater than one were perfectly matched by the moments of the number of transactions. Thus the number of transactions in a given time interval, ΔN, up to a constant, is assumed to be equal to the change in value of the directing process, Δn. In this manner, the current information flow $I(t)$ can be estimated from the number of transactions,

$$I(t) \quad \approx \quad \frac{\Delta n}{\Delta t} \;=\; \frac{\Delta N}{\Delta t} - const. \tag{11}$$

where *const* is a recentring parameter which can be thought to represent a background rate of trades independent of new information concerning the price, i.e. driven by liquidity rather than informational considerations.

Moments of the directing process greater than one are thus identical to the moments of the number of transactions. Hence the variance of the number of trades ΔN is identical to the variance of Δn,

$$\mathrm{Var}[\Delta N] \;=\; \mathrm{Var}[\Delta n] \;=\; \frac{2\sigma_I^2}{\alpha^2}\left(\alpha\Delta t + e^{-\alpha\Delta t} - 1\right).$$

where Δt represents the frequency of the available data. By studying how the variance of ΔN changes with Δt, a value of α can be obtained. Subsequently, a value of σ_I can be calculated from the variance at a fixed Δt. The error in estimating the variance is proportional to the reciprocal of the number of data points, and using the highest frequency data available, in this case daily, will give the best estimate of σ_I.

An estimate of the value of I during the sampling interval Δt can be obtained from the transaction data using (11). In order to define all the parameters detailing the process for $I(t)$, it is necessary to consider the next highest moment and calculate the skew of the transaction distribution. From (10) and (11),

$$\mathrm{Skew}[\Delta N] \;=\; \mathrm{Skew}[\Delta n] \;\approx\; \mathrm{Skew}[I] \;=\; 2\sigma_I/\mu_I,$$

enabling μ_I and subsequently the value of I applicable for each discrete time-step to be determined.

An additional estimate of the rate of mean reversion α can be obtained by considering the autocovariance of I, or equivalently ΔN. From (8) and (11),

$$\begin{aligned}\mathrm{Cov}[\Delta N(t+j\Delta t), \Delta N(t)] &\approx \Delta t^2 \mathrm{Cov}[I(t+j\Delta t), I(t)] \\ &= \Delta t^2 \sigma_I^2 e^{-\alpha j \Delta t}\end{aligned}$$

where j is an integer. Thus α can be determined by regressing the log of the lagged covariances of the number of transactions against the time lag $j\Delta t$, the gradient of which gives an additional estimate of α. A comparison of the two separate estimates of α can give an approximate error indication involved in this estimation process.

The remaining parameters μ_n and σ_n can be estimated from the asset price data and expressed in terms of μ_r and σ_r, the annualised drift and volatility of the asset return. The moments of the returns process were stated in §2.3, and can be written as

$$\begin{aligned}\mu_n &= \mu_r/\mu_I, \\ \sigma_n^2 &= \frac{\sigma_r^2}{\mu_I} - \frac{2\mu_r^2\sigma_I^2}{\mu_I\alpha^2}(\alpha - 1 + e^{-\alpha}). \qquad (12)\end{aligned}$$

It was found that since α was large, see Table 1, the second term on the right-hand side of (12) was negligible. An excellent approximation for the volatility per information event is $\sigma_n \approx \sigma_r/\sqrt{\mu_I}$ which will subsequently be used.

The parameter estimation technique was undertaken on a number of FTSE-250 stocks using daily price and transaction count data. 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. 2. With these stocks, there was good agreement between the values of α calculated from the scaling of the variance with Δt and the autocovariance, differing by typically 10% or less.

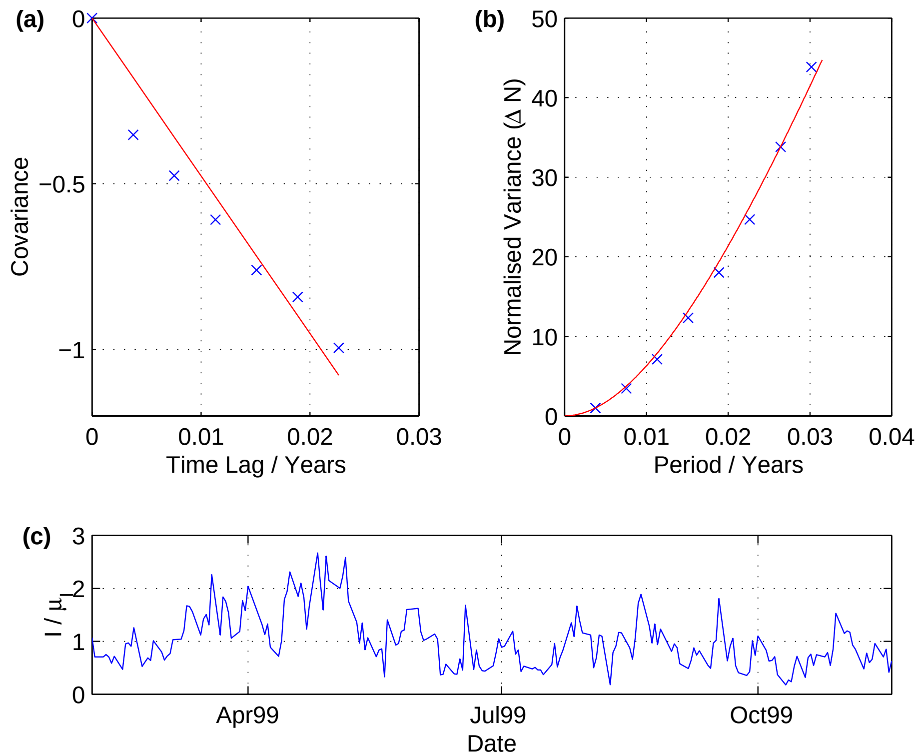

Figure 2: Estimation of parameters of $I(t)$ from transaction count data for Thames Water (Source: Primark Datastream). (**a**) Estimation of α using lagged covariance data. (**b**) Scaling of variance with Δt. The solid lines represent the fitted function. (**c**) Extracted $I(t)$ over estimation window.

On recent market entrants, e.g. high-tech, the process of parameter estimation was less successful. These stocks exhibit a strong increasing trend in the number of transactions, which is not accounted for by the mean-reverting choice of I. These stocks typically have days with an exceptional number of trades, sometimes over ten times the normal average. This has a significant impact on the value of the skew measured, from which σ_I is calculated. Correspond-

Stock	α	σ_I/μ_I	σ_r
Thames Water	55	0.5	0.19
British Airways	60	0.7	0.3
Bass	100	0.8	0.22
Pilkington	80	0.6	0.42
Baltimore Technologies	140	1.4	2.3
Oxford GlycoSciences	200	1.0	2.9
Morse Holdings	100	1.0	1.3

Table 1: Sample values (annualised units).

ingly, parameter stability is then reduced, with a strong dependence on which of these abnormal days are included in the sample period. There is generally poor agreement between the two estimates of α and a worse than expected fit of the variance of the number of transactions versus the time interval. For a summary of calculated parameters for a number of stocks, see Table 1.

3 Derivative pricing

3.1 The pricing equation

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. The option price can then be expressed as a solution of the parabolic partial differential equation

$$\begin{aligned}\frac{\partial V}{\partial t} + \frac{1}{2}\sigma_n^2 S^2 I \frac{\partial^2 V}{\partial S^2} + \frac{1}{2}q^2\frac{\partial^2 V}{\partial I^2} + \rho\sigma_n S q\sqrt{I}\frac{\partial^2 V}{\partial S \partial I}& \\ + rS\frac{\partial V}{\partial S} + (p-\lambda q)\frac{\partial V}{\partial I} - rV &= 0.\end{aligned} \tag{13}$$

Here $\lambda(S,I,t)$ is the market price of (information arrival intensity) risk, which is determined by the agents in the market and depends on their aggregate risk aversion, as well as liquidity considerations and other factors.

The pricing equation (13) can be re-expressed in terms of the parameters calculated in §2.4, namely:

$$p = \alpha(\mu_I - I), \qquad q = \beta I^{1/2} = \sigma_I\sqrt{\frac{2\alpha I}{\mu_I}}, \qquad \sigma_n \approx \frac{\sigma_r}{\sqrt{\mu_I}},$$

giving

$$\frac{\partial V}{\partial t} + \frac{\sigma_r^2 S^2 I}{2\mu_I}\frac{\partial^2 V}{\partial S^2} + \frac{\alpha\sigma_I^2 I}{\mu_I}\frac{\partial^2 V}{\partial I^2} + \rho\sigma_r\sigma_I\sqrt{2\alpha}S\frac{I}{\mu_I}\frac{\partial^2 V}{\partial S\partial I}$$
$$+ rS\frac{\partial V}{\partial S} + \left(\alpha(\mu_I - I) - \lambda\sigma_I\sqrt{\frac{2\alpha I}{\mu_I}}\right)\frac{\partial V}{\partial I} - rV = 0. \tag{14}$$

This equation must be solved with appropriate payoff and boundary conditions.

3.2 Relation to stochastic volatility models

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 (SVM). The pricing PDE (13) can be compared with the result for a general SVM, where it is assumed that the volatility of the asset S can be defined to be an explicit function of some stochastic process Y, i.e. $\sigma = f(Y)$. Equating the driving process Y with the mean reverting process I, then (13) can be interpreted as a SVM with $f(Y) = \sigma_n\sqrt{Y}$.

In general, SVMs [23, 41] have many benefits over the standard Black–Scholes model. Many SVMs give more realistic probability density functions for the asset, e.g. fat tails. The skew of the distribution can be incorporated by correlating the two Brownian motions. It can be proved that for any uncorrelated SVM, the implied volatility is convex with a minimum at the forward price of the stock [49]. Thus uncorrelated SVMs imply a volatility smile.

3.3 Numerical solution

In general, equation (14) has no analytical solution, and hence a numerical solution must be sought. Both finite-difference and finite element methods are considered.

Throughout this section we consider the solution appropriate for a standard vanilla call option, with strike K. However, the techniques can be easily extended to any vanilla derivative. Suitable boundary conditions in the S dimension are: $S \to 0$, $V \to 0$ and as $S \to \infty$, $\partial^2 V/\partial S^2 \to 0$. Alternatively, knowledge of the asymptotic value for the call as $S \to \infty$ can be used, i.e. as $S \to \infty$, $V = S - Ke^{-r(T-t)}$ (plus exponentially small terms). Since the payoff is a linear function of S, an appropriate boundary condition for large I is to let $\partial^2 V/\partial I^2 = 0$. A boundary condition for small I is less clear, but since there is a positive lower bound to I, it is assumed that $\partial V/\partial I = 0$. It was found that the final solution is insensitive to the exact boundary condition in this direction.

The final term left to specify is the market price of risk $\lambda(S, I, t)$, which allows the model to be calibrated against observed market prices. Information concerning the market price of risk can be obtained by fitting the model to existing prices on the market, but this is a computationally intensive task. To demonstrate the solution of this model, an arbitrary choice will be made in which the market price of risk is assumed to be of the form $\lambda(S, I, t) = \text{constant} \cdot \sqrt{I}$. This is convenient because it is non-trivial, but has been specified purely to demonstrate the solution technique.

3.3.1 Finite-difference methods

The Black-Scholes equation and its generalisations are ideal candidates for a finite-difference solution since they are typically linear and contain dominant diffusive terms that lead to smooth solutions. For a practical guide to the application of finite-difference schemes in option pricing, see [53, 54].

A sample graph of the implied volatility surface for the case of $\rho = 0$ is included as Figure 3, obtained from an ADI solution. However, in the most general case a correlation is present and hence the two-factor equation must be solved with a mixed derivative term present. This can be accommodated in an explicit scheme or an ADI framework [45]; the effect of changing the correlation is demonstrated in Figure 4.

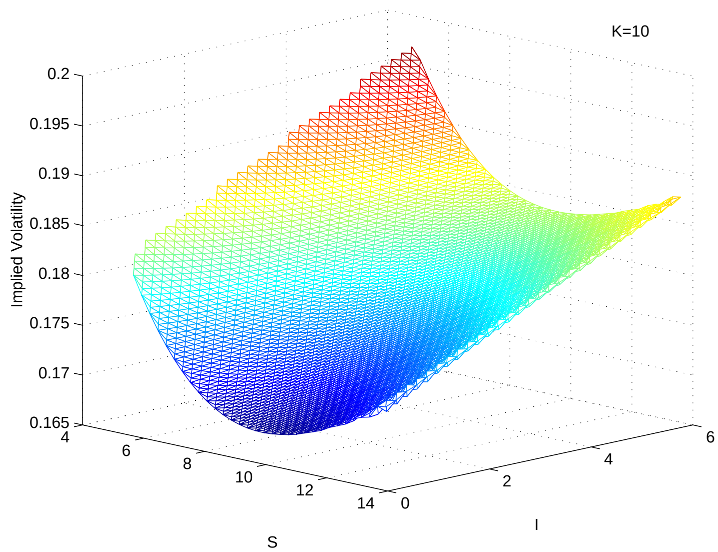

Figure 3: Implied volatility surface: $K = 10$, $T = 0.5$, $\alpha = 60$, $\sigma_I/\mu_I = 0.7$ and $\rho = 0$.

3.3.2 Finite element method

Prices were also obtained via a Galerkin finite-element approach implemented through the finite-element generic PDE package *Fastflo*. The mesh is concentrated near the strike K and stretched out near the boundaries, which promotes

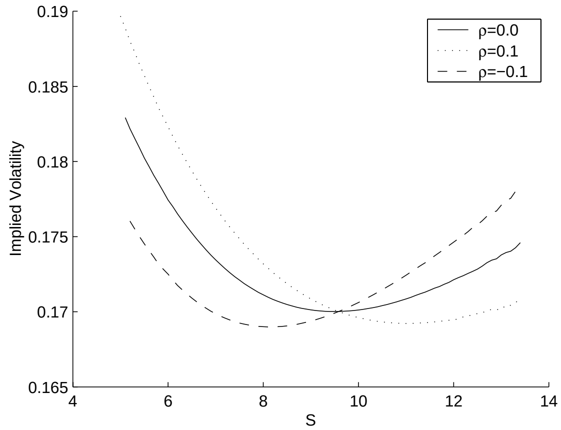

Figure 4: The effect of changing the correlation: $K = 10$, $T = 0.5$, $\alpha = 60$ and $\sigma_I/\mu_I = 0.7$.

high accuracy in the region of the strike. Natural boundary conditions are imposed in the I direction, but the final solution was not found to be sensitive to the exact form of these due to the coarse meshing in the region of the boundaries. The results obtained can be compared with the finite-difference methods described in the previous section, see Table 2.

	$S = 9.5$	$S = 10.0$	$S = 10.5$
ADI scheme	0.4490	0.7510	1.1269
Fastflo	0.4491	0.7509	1.1265

Table 2: Comparison of numerical results for a European call with $K = 10$, $T = 0.5$, $\alpha = 60$, $\sigma_I/\mu_I = 0.8$ and $\rho = 0$.

3.4 Asymptotics

The asymptotic analysis follows the approach adopted by Fouque, Papanicolaou and Sircar [23]. The main advantages are a reduction in the number of parameters required, and the ability to utilise the information contained within the implied volatility curve for calibration purposes.

The rate of mean-reversion of the process $I(t)$ was estimated from historical data for a number of different stocks in §2.4.1. From Table 1 it can be concluded that $I(t)$ *does* exhibit fast mean reversion, and we now define the dimensionless small quantity $\epsilon = 1/\alpha$, where α is in annualised units and the number 1 has inferred dimensions of years. The stochastic representation of $I(t)$, defined in

(6), can now be rewritten as

$$dI \quad = \quad \frac{1}{\epsilon}(\mu_I - I)\ dt + \sigma_I\sqrt{\frac{2I}{\epsilon\mu_I}}\ dX_t,$$

where σ_I is the long-run volatility defined in (7). We now express (14) as

$$\left(\frac{1}{\epsilon}\mathcal{L}_0 + \frac{1}{\sqrt{\epsilon}}\mathcal{L}_1 + \mathcal{L}_2\right) V \quad = \quad 0,$$

where

$$\begin{aligned}
\mathcal{L}_0 &= \frac{\sigma_I^2 I}{\mu_I}\frac{\partial^2}{\partial I^2} + (\mu_I - I)\frac{\partial}{\partial I}, \\
\mathcal{L}_1 &= \rho\sigma_r\sigma_I\sqrt{\frac{2}{\mu_I}}SI\frac{\partial^2}{\partial S\partial I} - \lambda\sigma_I\sqrt{\frac{2I}{\mu_I}}\frac{\partial}{\partial I}, \\
\mathcal{L}_2 &= \frac{\partial}{\partial t} + \frac{\sigma_r^2 S^2 I}{2\mu_I}\frac{\partial^2}{\partial S^2} + rS\frac{\partial}{\partial S} - r,
\end{aligned}$$

and the solution for V is expanded in powers of $\sqrt{\epsilon}$:

$$V(S,I,t) \quad \sim \quad V_0(S,I,t) + \sqrt{\epsilon}V_1(S,I,t) + \epsilon V_2(S,I,t) + \cdots$$

We then perform the standard procedure in which at each order the solution contains an undetermined function which is found by a solvability condition at the next order in the expansion. The zero order term $V_0(S,t)$ is the solution of the Black-Scholes equation with a constant volatility of σ_r, and the corrected price can be expressed as

$$V \quad = \quad V_0 - \sqrt{\epsilon}(T-t)\left(A_2 S^2\frac{\partial^2 V_0}{\partial S^2} + A_3 S^3\frac{\partial^3 V_0}{\partial S^3}\right) + \mathcal{O}(\epsilon), \tag{15}$$

where

$$\begin{aligned}
A_2 &= \sqrt{2}\sigma_I\sigma_r^2\left(\rho\sigma_r\Big\langle g'(I)I/\mu_I\Big\rangle - \frac{1}{2}\Big\langle\lambda g'(I)\sqrt{I/\mu_I}\Big\rangle\right), \\
A_3 &= \frac{1}{\sqrt{2}}\rho\sigma_r^3\sigma_I\Big\langle g'(I)I/\mu_I\Big\rangle,
\end{aligned}$$

and where $g(I)$ is the solution of $\mathcal{L}_0 g(I) = I/\mu_I - 1$; here angled brackets represent an expectation with respect to $p_\infty(I)$, as defined in (9). An explicit dependence on I enters only in the $\mathcal{O}(\epsilon)$ term.

A key point is the universality of this formula. Any fast mean-reverting stochastic volatility model will lead to a first-order correction of this form. The A_2 term is a volatility level correction, and depends on the market price of risk, whereas the A_3 term depends entirely on the correlation coefficient ρ and the third derivative of the Black-Scholes option price with respect to S.

3.4.1 Implied volatilities

The implied volatility θ can be approximated by a linear function of the logged forward moneyness:

$$\theta \quad \sim \quad \sigma_r + \sqrt{\epsilon}\left(a + b\frac{\log M}{(T-t)}\right) + \mathcal{O}(\epsilon),$$

where

$$a \quad = \quad \frac{3A_3 - 2A_2}{2\sigma_r} \qquad \text{and} \qquad b \; = \; \frac{A_3}{\sigma_r^3}.$$

The coefficients a and b can be estimated by OLS regression using the implied volatility skew. The values of the parameters A_2 and A_3 can subsequently be calculated:

$$A_2 \quad = \quad \sigma_r\left(\frac{3}{2}b\sigma_r^2 - a\right) \qquad \text{and} \qquad A_3 \; = \; b\sigma_r^3.$$

For example, using market data for Glaxo Wellcome call options, expiry Oct 99, values of $\sqrt{\epsilon}A_2 = 1.9 \times 10^{-3}$ and $\sqrt{\epsilon}A_3 = -1.5 \times 10^{-2}$ were obtained, confirming these quantities are small.

These coefficients can be used to directly price an option using (15), with no need to estimate the market price of risk. This is superior to regarding implied volatilities as the market's rational expectations of future volatility, since statistical evidence shows little or no correlation between implied volatility and subsequent realised volatility [10].

4 Concluding remarks

In this paper we have described a model for the asset price where time is subordinated to a underlying stochastic process representing the number of trades. The model is consistent with normality of returns in this new timescale, with the leptokurtosis of the asset price, observed in calendar time, being due to variations in the directing process. We show that stock volatility is directly related to this underlying process, thereby predicting the observed positive correlation between volatility and the number of transactions. An option pricing formula was subsequently derived, and interpreted within the framework of a stochastic volatility model. The underlying process was found to be fast mean–reverting, and this was exploited to perform an asymptotic expansion of the pricing formula. Using this technique there is no need to specify a market price of risk, and the implied volatility skew can be used to calibrate the model. The proposed model for the information flow was not found to be applicable to all stocks, and this is a suitable area of further research.

Acknowledgements

David Lamper is grateful to EPSRC for financial support.

References

[1] Torben G. Andersen, *Return volatility and trading volume: An information flow interpretation of stochastic volatility*, The Journal of Finance **51** (1996), no. 1, 169–204.

[2] Torben G. Andersen, Tim Bollerslev, Francis X. Diebold, and Heiko Ebens, *The distribution of stock return volatility*, NBER Working Paper No. W7933, October 2000.

[3] Clifford A. Ball and Antonio Roma, *Stochastic volatility option pricing*, Journal of Financial and Quantitative Analysis **29** (1994), no. 4, 589–607.

[4] Anil K. Bera and Matthew L. Higgins, *ARCH models: Properties, estimation and testing*, Journal of Economic Surveys **7** (1993), no. 4, 305–366.

[5] Hendrik Bessembinder, Kalok Chan, and Paul J. Seguin, *An empirical examination of information, differences of opinion, and trading activity*, Journal of Financial Economics **40** (1996), 105–134.

[6] Robert C. Blattberg and Nicholas J. Gonedes, *A comparison of the stable and student distributions as statistical models for stock prices*, Journal of Business **47** (1974), no. 2, 244–280.

[7] Tim Bollerslev, *Generalized autoregressive conditional heteroskedasticity*, Journal of Econometrics **31** (1986), 307–327.

[8] Tim Bollerslev, Ray Y. Chou, and Kenneth F. Kroner, *ARCH modelling in finance. A review of the theory and empirical evidence*, Journal of Econometrics **52** (1992), 5–59.

[9] Jean-Philippe Bouchaud and Marc Potters, *Theory of financial risks*, Cambridge University Press, 2000.

[10] Linda Canina and Stephan Figlewski, *The information content of implied volatility*, The Review of Financial Studies **6** (1993), no. 3, 659–681.

[11] Carolyn W. Chang and Jack S.K. Chang, *Option pricing with stochastic volatility: Information-time vs. calendar-time*, Management Science **42** (1996), no. 7, 974–991.

[12] Carolyn W. Chang, Jack S.K. Chang, and Kian-Guan Lim, *Information-time option pricing: Theory and empirical evidence*, Journal of Financial Economics **48** (1998), no. 2, 211–242.

[13] Tarun Chordia and Bhaskaran Swaminathan, *Trading volume and cross-autocorrelations in stock returns*, Journal of Finance **55** (2000), no. 2, 913–935.

[14] Peter K. Clark, *A subordinated stochastic process model with finite variance for speculative prices*, Econometica **41** (1973), no. 1, 135–155.

[15] John C. Cox, Jonathan E. Ingersoll, and Stephen A. Ross, *A theory of the term structure of interest rates*, Econometrica **53** (1985), no. 1, 385–407.

[16] Ernst Eberlein and Ulrich Keller, *Hyperbolic distributions in finance*, Bernoulli **1** (1995), 281–299.

[17] Robert F. Engle, *Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation*, Econometrica **50** (1982), no. 4, 987–1008.

[18] Thomas W. Epps and Mary Lee Epps, *The stochastic dependence of security price changes and transaction volumes: Implications for the mixture-of-distributions hypothesis*, Econometrica **44** (1976), no. 2, 305–321.

[19] Eugene F. Fama, *The behaviour of stock market prices*, Journal of Business **38** (1965), 34–105.

[20] ———, *Efficient capital markets: A review of theory and empirical work*, Journal of Finance **25** (1970), no. 2, 383–417.

[21] ———, *Efficient capital markets II*, Journal of Finance **25** (1991), no. 5, 1575–1617.

[22] William Feller, *An introduction to probability theory and its applications*, vol. II, John Wiley and Sons, 1966.

[23] Jean-Pierre Fouque, George Papanicolaou, and K. Ronnie Sircar, *Derivatives in financial markets with stochastic volatility*, Cambridge University Press, 2000.

[24] Kenneth R. French and Richard Roll, *Stock return variances: The arrival of information and the reactions of traders*, The Journal of Financial Economics **17** (1986), 5–26.

[25] A. Ronald Gallant, Peter E. Rossi, and George E. Tauchen, *Stock prices and volume*, The Review of Financial Studies **5** (1992), no. 2, 199–242.

[26] Hélyette Geman and Thierry Ané, *Stochastic subordination*, Risk Magazine **9** (1996), 145–149.

[27] ———, *Order flow, transaction clock, and normality of asset returns*, The Journal of Finance **55** (2000), no. 5, 2259–2284.

[28] Ramazan Gencay and Thanasis Stengos, *Moving average rules, volume and the predictability of security returns with feedforward networks*, Journal of Forecasting **17** (1998), 401–414.

[29] Parameswaran Gopikrishnan, Vasiliki Plerou, Xavier Gabaix, and H. Eugene Stanley, *Statistical properties of share volume traded in financial markets*, cond-mat/0008113 **60** (1999), no. 2, 1390–1400.

[30] Lawrence Harris, *Cross-security tests of the mixture of distributions hypothesis*, Journal of Financial and Quantitative Analysis **21** (1986), no. 1, 39–46.

[31] ———, *Transaction data tests of the mixture of distributions hypothesis*, Journal of Financial and Quantitative Analysis **22** (1987), no. 2, 127–141.

[32] T. Ito and W.L. Lin, *Lunch break and intraday volatility of stock returns*, Economics Letters **39** (1992), no. 1, 85–90.

[33] Prem C. Jain, *Response of hourly stock prices and trading volume to economic news*, Journal of Business **61** (1988), no. 2, 219–231.

[34] Charles M. Jones, Gautam Kaul, and Marc L. Lipson, *Information, trading and volatility*, The Journal of Financial Economics **36** (1994), no. 1, 127–154.

[35] ———, *Transactions, volume, and volatility*, The Review of Financial Studies **7** (1994), no. 4, 631–651.

[36] Jonathan M. Karpoff, *A theory of trading volume*, Journal of Finance **41** (1986), 1069–1087.

[37] ———, *The relationship between price changes and trading volume: A survey*, Journal Of Financial and Quantitative Analysis **22** (1987), 109–126.

[38] Albert S. Kyle, *Continuous auctions and insider trading*, Econometrica **53** (1985), no. 6, 1315–1336.

[39] Christopher G. Lamoureux and William D. Lastrapes, *Heteroskedasticity in stock return data: Volume versus garch effects*, The Journal of Finance **XLV** (1990), no. 1, 221–229.

[40] ———, *Endogeneous trading volume and momentum in stock return volatility*, The Journal of Business and Economic Statistics **12** (1994), 253–260.

[41] Alan L. Lewis, *Option valuation under stochastic volatility: with Mathematica code*, Finance Press, 2000.

[42] Thomas Lux, *Multi-fractal processes as models for financial returns: A first assessment*, Social Science Research Network Working Paper Series, August 1999.

[43] Benoit B. Mandelbrot, *The variation of certain speculative prices*, Journal of Business **36** (1963), no. 4, 394–419.

[44] Rosario N. Mantegna and H. Eugene Stanley, *Stochastic processes with ultraslow convergence to a gaussian: The truncated Lévy flight*, Physical Review Letters **73** (1994), no. 22, 2946–2949.

[45] S. McKee, D.P. Wall, and S.K. Wilson, *An alternating direction implicit scheme for parabolic equations with mixed derivative and convective terms*, Journal of Computational Physics **126** (1996), 64–76.

[46] Daniel B. Nelson, *Conditional heteroskedasticity in asset returns: A new approach*, Econometrica **59** (1991), no. 2, 347–370.

[47] Maureen O'Hara, *Market microstructure theory*, Blackwell, 1995.

[48] Vasiliki Plerou, Parameswaran Gopikrishnan, Luis A. Nunes Amaral, Martin Meyer, and H. Eugene Stanley, *Scaling of the distribution of price fluctuations of individual companies*, Physical Review E **60** (1999), no. 6, 6519–6529.

[49] E. Renault and N. Touzi, *Option hedging and implied volatilities in a stochastic volatility model*, Mathematical Finance **6** (1996), no. 3, 279–302.

[50] Stephan A. Ross, *Information and volatility: The no-arbitrage martingale approach to timing and resolution irrelevancy*, The Journal of Finance **44** (1989), no. 1, 1–17.

[51] Gennady Samorodnitsky, *Stable non-gaussian random processes*, Chapman & Hall, 1994.

[52] George E. Tauchen and Mark Pitts, *The price variability-volume relationships on speculative markets*, Econometrica **51** (1983), no. 2, 485–505.

[53] Paul Wilmott, *Derivatives: The theory and practice of financial engineering*, John Wiley and Sons, 1998.

[54] Paul Wilmott, Jeff Dewynne, and Sam Howison, *Option pricing: Mathematical models and computation*, Oxford Financial Press, 1993.