Stochastic Volatility

David G. Hobson
School of Mathematical Sciences, University of Bath
9 October 1996

Abstract

The volatility of a financial asset is the variance per unit time of the logarithm of the price of the asset. Volatility has a key role to play in the determination of risk and in the valuation of options and other derivative securities.

The widespread Black-Scholes model for asset prices assumes constant volatility. The purpose of this chapter is to review the evidence for non-constant volatility and to consider the implications for option pricing of alternative random or stochastic volatility models. We concentrate on continuous time diffusion models for the volatility, but we also make comments about certain classes of discrete time models, such as *ARV, ARCH* and *GARCH*.

1 Volatility and the need for Stochastic Volatility models

1.1 Introduction

A common approach in the modelling of financial assets is to assume that the proportional price changes of an asset form a Gaussian process with stationary independent increments. The celebrated (and ubiquitous) Black-Scholes option pricing formula is based on such a premise. The success and longevity of the Gaussian modelling approach depends on two main factors: firstly the mathematical tractability of the model, and secondly the fact that in many circumstances the model provides a reasonable and simple approximation to observed market behaviour.

An immediate corollary of the Gaussian assumption is that the behaviour of the asset price can be summarised by two parameters, namely the mean and the standard deviation of the Gaussian variables. In finance-speak the standard deviation is renamed the volatility. Volatility is a key concept because it is a measure of uncertainty about future price movements, because it is directly related to the risk associated with holding financial securities and hence affects consumption/investment decisions and portfolio choice, and because volatility is the key parameter in the pricing of options and other derivative securities.

This chapter is concerned with the estimation of volatility and the implications of the empirical observation that volatility appears non-constant over time. Some of the evidence for this claim is given in the next subsection. In Section 2 we review the continuous-time models which have been introduced to reflect the non-constant volatility phenomenon, including level-dependent and stochastic diffusion models for the volatility. Since the fundamental problem in mathematical finance is to price derivative securities such as options we focus in Section 3 on the implications of these alternative models for derivative pricing. The types of discrete-time models (such as ARV, $ARCH$ and $GARCH$) favoured by econometricians to model stochastic volatility are the subject of Section 4. Some comments and conclusions on the importance of accurate modelling of volatility are given in a conclusion.

1.2 Simple models for Asset Prices

The canonical continuous Gaussian process is Brownian motion. In the same way that the normal law arises as the limit of a normalised sum of independent random variables, so Brownian motion arises as the limit of a random walk as, simultaneously, step sizes are reduced and step frequency increased. The links between Brownian motion and finance are long and illustrious; and indeed (mathematician's) Brownian motion was devised by Bachelier (1900) as a model for French stock prices.

The increments of a Brownian motion are normal random variables. As a consequence Brownian motion can and does become negative which makes it an unsatisfactory model for limited liability stocks. Instead a more reasonable model was proposed by Osborne (1959) and Samuelson (1965) who took the asset price to be an exponential (or geometric) Brownian motion. Thus the logarithm of the asset price is a Brownian motion.

To be more formal, the standard and classical model for the behaviour of the price of a financial asset, such as a stock, assumes that the price process $(P_t)_{t\geq 0}$ is the solution to a stochastic differential equation (SDE)

$$dP_t = P_t(\mu dt + \sigma dB_t) \tag{1}$$

where t is measured in units of one year, B_t is a Brownian motion and μ, the mean, and σ, the volatility, are constant parameters of the model. The time convention is chosen to ensure that σ can be interpreted as an annualised volatility. This SDE has solution

$$P_t = P_0 \exp\{\sigma B_t + (\mu - \tfrac{1}{2}\sigma^2)t\}. \tag{2}$$

For a highly readable introduction to stochastic differential equations see Øksendal (1985). The discrete time analogue of (1) based on a daily sequence of observations $(P_n)_{n\geq 0}$ is

$$\ln P_n - \ln P_{n-1} \equiv \Delta(\ln P_n) = \nu + \sigma Z_n. \tag{3}$$

where (Z_n) is a sequence of independent normal random variables with zero mean and variance (1/365). Again this choice of normalisation ensures that σ can be interpreted as an annualised volatility. The assumption that the innovations have a normal distribution means that the increments have a natural nesting property. For example the proportional price changes of a weekly time series also have a normal distribution. Thus the dynamics are not dependent on the choice of timescale.

Two contradictory philosophies are available here. It is possible to view the discrete time series as a δ-skeleton of the underlying continuous Markov process given by (1) with the understanding that even tick data provides only an approximation to the inherently unobservable true process. Alternatively the SDE formulation, whose merit is tractability, can be viewed as a limiting approximation to a discrete stochastic difference equation. Discrete models are suited to qualitative and descriptive analyses, whereas continuous time models provide the natural framework for theoretical option pricing. We shall take the view that the fundamental problem in mathematical finance is the calculation of derivative prices and in particular formulæ which relate the price of an option to the price of the underlying asset and other key variables. Hence we shall concentrate on continuous time diffusion models for the price process and volatility.

In principle, if the continuous time model can be observed perfectly (and in continuous time) then it is possible to read off the instantaneous value of the volatility from the asset price. (The square of the volatility is the quadratic variation of the log-price process.) In practice however the volatility must be estimated from the data. Suppose that the data consists of a series of daily observations of the price of an asset $(P_k)_{k \leq N}$. Our first estimate of the volatility, $\hat{\sigma}$, is called the *historic volatility*. At time n, the historic volatility based on the last J days is the maximum likelihood estimator obtained from the model (3) and the data $P_{n-J-1}, \ldots P_n$:

$$\hat{\sigma}_{n,J} \equiv \sqrt{\frac{365}{J-1} \sum_{j=0}^{J-1} \left(\Delta(\ln P_{n-j}) - \frac{1}{J} \sum_{j=0}^{J-1} \Delta(\ln P_{n-j}) \right)^2} \tag{4}$$

The factor of 365 converts daily volatility into an annualised term. Typically J is taken to be 90 or 180 days. These choices are a compromise between the desire for a large number of observations and a realisation that the dynamics of the price process are unlikely to remain constant over several years.

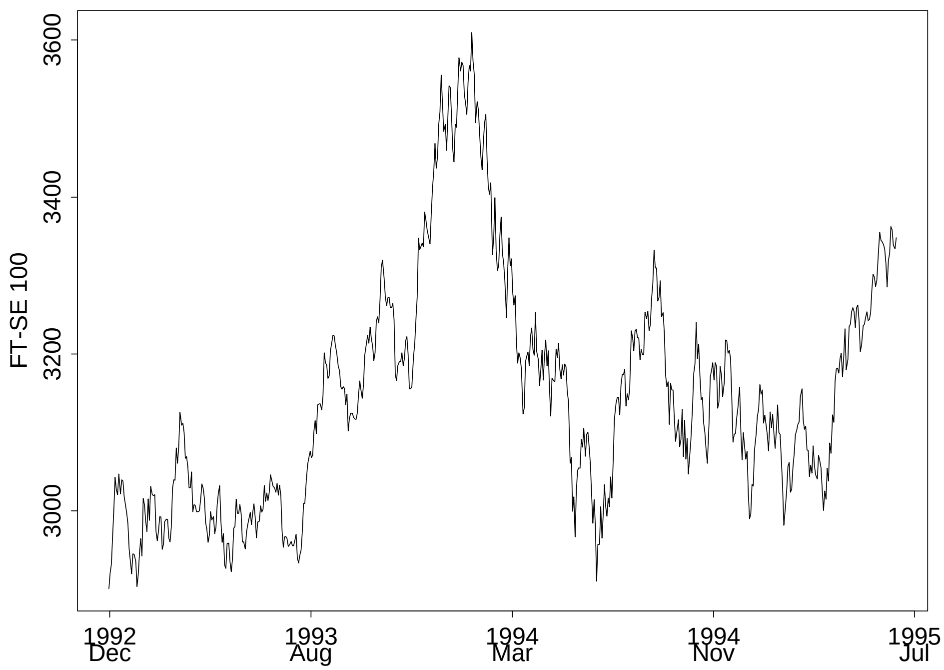

Figure 1: The September 1995 futures price for the *FT-SE* 100 index.

Figure 1 gives a plot of the September 1995 futures price of the *FT-SE* 100 index (the *Financial Times — Stock Exchange* index of the stock prices of 100 leading UK companies) over the period December 1992 to July 1995 and Figure 2 gives an estimate of the 90-day historic volatility based on this data. The advantage of considering the futures price (the amount which it is agreed now is to be paid in September 1995 for

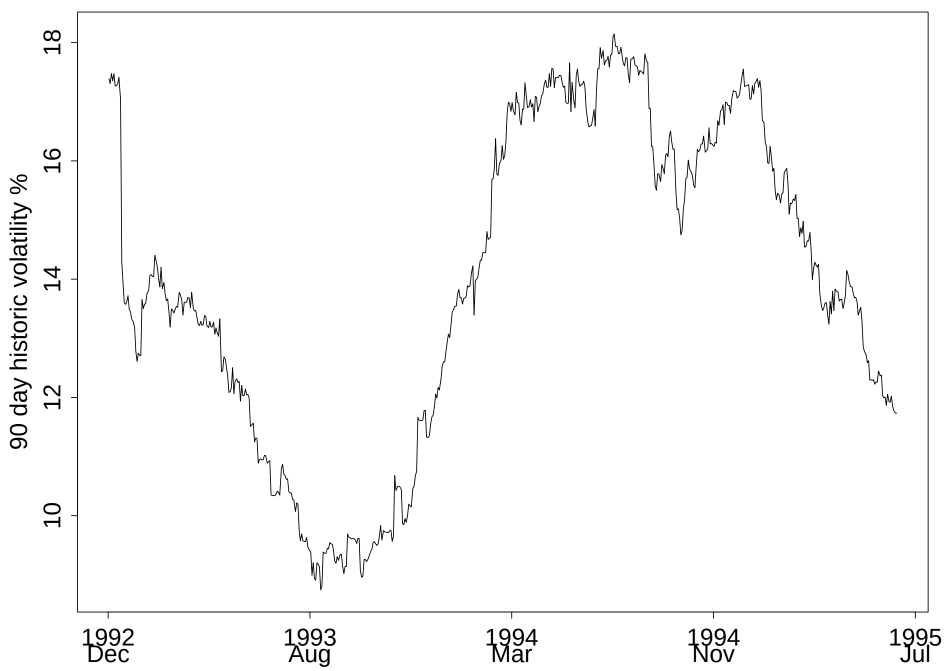

Figure 2: 90-day historic volatility for the *FT-SE* 100 index based on the *FT-SE* price data in Figure 1.

delivery, again in September 1995, of (the cash value of) the basket of stocks in the *FT-SE* index) is that our analysis is not confounded by interest rate and discounting effects.

Figure 2 gives an estimate of the 90-day historic volatility based on the above data. This limited evidence supports the contention that stock volatility is not constant and moreover that volatility shocks persist through time. This conclusion was reached by Mandelbrot (1963), Fama (1965), Blattberg and Gonedes (1974), and Scott (1987) amongst others. Stochastic volatility models are needed to describe and explain volatility patterns.

1.3 The Black-Scholes paradigm and Option pricing

One of the key contributions of mathematics to finance has been the development of formulæ for the pricing of options and other derivative securities. Black and Scholes (1973) showed that, subject to certain modelling assumptions, there is a strategy for risklessly hedging options in the sense that it is possible to perfectly replicate the payoff of the option through dynamic trading. Thus there is a unique preference-independent rational price for an option. This price corresponds to the fortune needed to purchase the initial portfolio which is required to hedge the option. This observation has revolutionised financial markets and contributed greatly to the explosion in the volume of trading in derivative securities.

The purchaser of a European *call option* on an asset with *strike* K and *expiry* T has the right, but not the obligation, to buy one unit of the asset at time T for a price K. (An American call option conveys the right to buy the asset at any time before T; the option is European style if the right to buy is restricted to the time T alone.) This right will only be exercised if the price P_T of the asset at time T is above K; otherwise at expiry the option is worthless. It is often convenient to think of an option as a derivative security which at time T pays the cash amount $(P_T - K)^+ \equiv \max\{P_T - K, 0\}$. The fundamental problem in mathematical finance is to find the fair price of such an option at a time t prior to expiry.

In order to price this option it is necessary, following Black and Scholes (1973) to make a number of regularity assumptions about the financial market in which the underlying asset is traded. In particular the market is assumed to be perfect and frictionless, so that there are no transaction costs, there is no taxation, and the underlying asset is available in arbitrary amounts. There is a constant rate of interest r for both borrowing and lending, there are no dividends and there are no restrictions on short selling of stock provided that the net wealth of the trader remains non-negative. In particular a trader may sell stock or bonds that he does not own provided that by the end of the trading period he has repurchased sufficient quantities to cover his obligations. Many of these assumptions can be weakened. For example it is easy to relax the assumption about constant interest rates to an assumption of deterministic interest rates. Finally Black and Scholes assume that the asset price process is given by the solution to (1) with constant and known parameter values μ and σ.

The Black-Scholes price C of a call option is given by

$$C(P_t, t; K, T; \sigma, r) \equiv C = Ke^{-r(T-t)}(M_t\Phi(d_1) - \Phi(d_2)), \tag{5}$$

where $M_t \equiv (P_t/Ke^{-r(T-t)})$ is the *moneyness* of the option and d_1 and d_2 are given by

$$\begin{aligned} d_1 &= \frac{\ln(M_t) + \frac{1}{2}\sigma^2(T-t)}{\sigma\sqrt{T-t}} \\ d_2 &= d_1 - \sigma\sqrt{T-t} \end{aligned}$$

respectively. (The term moneyness refers to the fact that if $M_t > 1$ then the option is said to be in-the-money and if the futures price remains unchanged then the option will make a payout on expiry. An option for which $M_t < 1$ is said to be out-of-the-money since unless the underlying value of the asset increases the option payout will be zero.) In the expression (5) $\Phi(.)$ denotes the cumulative normal distribution function, and K and T are the strike and expiry as before. In the sequel we will be flexible in deciding which of the quantities P_t, t, K, T, σ are to be considered as variables, and which are fixed parameters.

There are several important remarks which should be made about this formula. Firstly the justification for calling C the fair price of the option is based on the fact that the quantity C can be used to finance a trading strategy which at maturity is guaranteed to match the payoff of the option. Models with this replication property are said to be complete. Secondly the drift parameter μ does not appear. Indeed it is as if the option price was calculated as the discounted expected payoff of an option on an asset whose dynamics are given by the SDE

$$\frac{dP}{P} = \sigma dW + r dt$$

rather than (1). In particular the call price can be expressed as a conditional expectation given the current price:

$$C \equiv C(P_t, t) = e^{-r(T-t)}\tilde{\mathbb{E}}[(P_T - K)^+ | P_t] \tag{6}$$

where $\tilde{\mathbb{E}}$ denotes expectation with respect to the risk-neutral probability measure. This is the measure under which W is a Brownian motion. Thirdly, and as a direct consequence of the second remark, there is a single unknown parameter in the Black-Scholes formula. The strike K and time to expiry $(T-t)$ are part of the specification of the option; the interest rate r is assumed known and the current price P_t is observable. Thus the price of the option depends solely on the value of the volatility. Moreover the option price depends on the volatility only through the quantity $\sigma^2(T-t)$, which is the integrated squared volatility over the remaining lifetime of the option. This illustrates a more general comment that the Black-Scholes model can easily be adapted to allow for time varying parameter values for the volatility parameter, provided that the behaviour is deterministic, and provided that the term $\sigma^2(T-t)$ is replaced by $\int_t^T \sigma_s^2 ds$.

The call option pricing function $C \equiv C(\sigma)$ is an increasing function of the volatility σ. This observation can be verified by differentiation of (5), or is immediate from the representation (6); see also Figure 5. This means that not only can we calculate the price of an option given a value for the volatility parameter, but also that given the price of an option it is possible to deduce the unique value of the volatility which must be substituted into the Black-Scholes formula to obtain the observed option price. We define the *implied volatility* $\tilde{\sigma}$, to be the value of the volatility parameter σ which is consistent with the Black-Scholes formula and the observed call price.

Thus we have a new measure of volatility. Implied volatility is a market assessment of the expected future volatility over the lifetime of the option. Implied volatility is a useful device because it provides a convenient shorthand for expressing the option price, and because it facilitates price comparisons of options with different characteristics.

Suppose that the assumptions of the Black-Scholes model are satisfied, so that in particular the price process of the underlying asset is given by the solution to the SDE (1). Provided that the market prices options using the Black-Scholes formula, then the implied volatility $\tilde{\sigma}$ should be identically equal to the true parameter value σ. For all strikes K and maturities T we can define $\tilde{\sigma}(K,T)$ to be the implied volatility of the call option with maturity $T > t$ and strike K. If the Black-Scholes model is correct then a plot of $\tilde{\sigma}(K,T)$ should yield a constant surface.

Analyses of implied volatility patterns have been attempted by Rubinstein (1985), Skeikh (1991), Fung and Hsieh (1991), Heynen, Kemma and Vorst (1994) and Xu and Taylor (1994) amongst others. These authors all find systematic biases in the implied volatility surface. In particular there is strong evidence of an implied volatility smile (so that for the cross section of the implied volatility surface corresponding to a given maturity implied volatility is a convex function of the strike) and some evidence of skews (so that added to convexity of the cross-sectional implied volatility, there is an additional linear relationship). Skews are particularly evident in implied volatilities for indices rather than individual stocks; see Wiggins (1987).

Figure 3 presents implied volatility data for a set of call options traded on the London Financial Futures Exchange on the 18th April 1995. The call options are European style options on the *FT-SE* 100 Index. The plot is presented as a surface parameterised by the date of the expiry of the option and the moneyness of the option. For each of the five expiry dates there are between 13 and 18 options traded with different strikes.

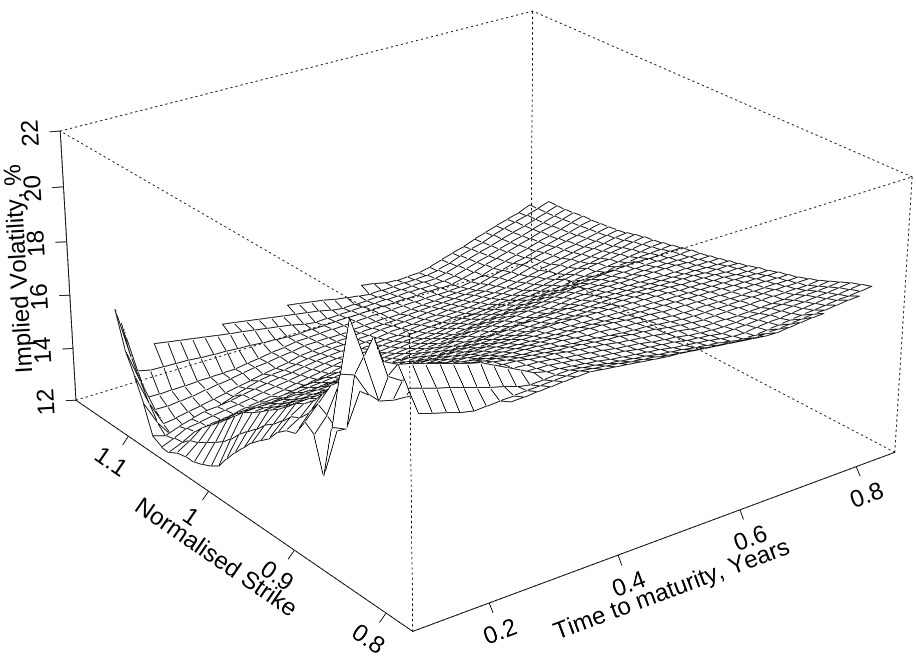

Figure 3: Implied volatilities of European call options on the *FT-SE* 100 Index on the 18th April 1995. The normalised strike is the ratio of the true strike and the relevant futures price of the index. The graph has been interpolated by Splus.

There is clear evidence of a skew and some evidence of a volatility smile. These effects become less pronounced as maturity increases. Finally it appears the implied volatility of an at-the-money option is an increasing function of time to expiry. A modification of the Black-Scholes model is required to account for these effects and stochastic volatility models provide a potential explanation.

2 Non-constant volatility models

In the first section of this chapter we described the standard exponential Brownian motion model for asset prices and noted some of the discrepancies and inconsistencies which arise when this model is compared with market experience. In particular Figure 2 shows time series plots of realised volatility which contradict a stationary normal hypothesis. Moreover plots of implied volatility in Figure 3 are inconsistent with market belief in Black-Scholes with constant volatility. In this section we aim to outline some of the models which attempt to explain, or at least account for, these inconsistencies. Although there are other potential explanations for the observed biases (for example non-zero transaction costs will require modifications of the Black-Scholes formula and liquidity considerations may inflate the prices of options which are away from the money) we will focus on explanations and models in which the volatility becomes non-constant.

In the 1960s empirical studies of asset price behaviour by Mandelbrot (1963) and Fama (1965) found leptokurtosis in the distribution of the daily changes in the log-price. These authors were led to suggest an innovations process consisting of random variables with stable Paretian distributions with characteristic exponents between one and two in an attempt to explain the observed fat-tails of the empirical distribution relative to a normal law. Second and higher moments do not exist for such distributions so that the notion of volatility becomes ill-defined, with serious implications for the pricing of options, at least within the Black-Scholes paradigm.

The stable Paretian hypothesis continues to have its proponents (see for example Peters (1991)). However the tractability that the Black-Scholes model derives from the Gaussian character of its underlying variables allows it to retain it's pre-eminence amongst the class of asset price models as the reference model against which others are compared. Instead of rejecting normality financial economists have searched for alternative explanations for the observed kurtosis and apparent randomness of volatility which rely on modification of the Gaussian framework.

2.1 Subordinators and volume effects

One attractive explanation for the apparent randomness of volatility claims that the asset price process *is* an exponential Brownian motion, but only when the time parameter t is interpreted as an *intrinsic clock* rather than real or calendar time. Relative to real time the daily changes are a mixture of normals. This model was proposed by Clark (1973) who argued that the daily proportional price change is a sum of a random number of within-day price changes and that the number of such changes is related positively to the rate of information flow or the volume of trading. Strong supporting evidence for his general thesis was found by Epps and Epps (1976) and Tauchen and Pitts (1983); in particular a mixture of normals hypothesis was observed to fit the data more accurately than a stable Paretian distribution. Karpoff (1987) documents several studies relating asset price volatility to traded volume.

In general the model is as follows. Let A_t be a subordinator so that A_t is a non-decreasing Markov process with stationary independent increments. The price process $(P_t)_{t\geq 0}$ is a random time-change of an exponential Brownian motion and is given by

$$P_t = P_0 \exp(B_{A(t)} - \mu A_t + \nu t)$$

for a pair of drift parameters μ and ν. For the subordinator $A_t \equiv \sigma^2 t$ we recover the standard Black-Scholes model.

Madan and Senata (1990) proposed a particular choice of subordinator and termed the resulting stock price model the *Variance-Gamma* model. In this model the subordinator is a Gamma process. The Gamma process is a pure jump process and the price process inherits this property also. Madan and Senata note that their model has the following desirable properties: firstly the distribution of proportional price changes is fat-tailed relative to the normal; secondly the distribution has finite moments, at least of lower orders; thirdly the process is consistent with an underlying continuous-time stochastic process; and finally that the model can be extended to a multi-dimensional process. The model also has one serious disadvantage however, the existence of jumps in the asset price makes the pricing and hedging of options very awkward.

2.2 Leverage effects and implied volatility skews

A second observed feature of stock price volatility is a correlation between volatility and price level. This relationship is implicit in implied options prices; the market expects volatility to rise as prices fall. One common explanation for this relates volatility to leverage effects.

Imagine a firm with debt whose share value represents the surplus of the firm's assets over this debt. Suppose that the value of the firm's assets fluctuates like an exponential Brownian motion with constant variance whilst the value of the debt remains fixed. Then the magnitude of the proportional changes in the share value is greater than the magnitude of the proportional change in the asset value, though of course the absolute changes are the same. Hence the volatility of the share price is greater than the volatility of the assets. Moreover if the value of the assets rises then the ratio of the stock value to asset value approaches unity and fluctuations in the value of assets are directly reflected in the stock price. Conversely if the value of the assets falls then the debt factor becomes significant and the effect is to magnify changes in the asset value as represented via the stock price. In this way leverage introduces a negative correlation between volatility and price level.

The above arguments provide a direct inspiration for the equity price models of Geske (1979), Rubinstein (1983) and Bensoussan, Crouhy and Galai (1994). In essence these models propose that the stock price is the solution of a SDE

$$\frac{dP_t}{P_t} = \sigma(P)dB_t + \mu dt \tag{7}$$

which is a modified form of (1) in which the volatility component is allowed to depend on the price level. Geske deduces an explicit form for the function $\sigma(P)$ which is related to the debt structure of the firm, but it is also possible to consider models which begin by specifying arbitrary forms for equations (7). One class of such models is the Constant Elasticity of Variance (CEV) class of models proposed by Cox and Ross (1976) for which $\sigma(P) = \sigma P^{\alpha-1}$ for some $\alpha \in (0,1)$. In the CEV model there is a negative correlation between the volatility and price level.

The class of models of the form (7) have several desirable features. Firstly, for suitable choices of $\sigma(P)$ the dependence between volatility and price level can be modelled. Secondly, the model is complete and as in the Black-Scholes model there is a unique preference independent price for an option.

Given a new model it is illuminating to compare the options prices from this model with those from the Black-Scholes formula. Consider the following exercise; calculate the prices of call options under the alternative model for a range of different strikes and exercise dates. Use these prices to derive the Black-Scholes implied volatility of each option. Finally plot the resulting implied volatility surface, and compare with market implied volatility data.

Figure 4 illustrates the results of such an exercise for the CEV model. Motivated by the calculations of Schroder (1989) we set $\alpha = 2/3$ so that there are simple explicit expressions for the prices of call options. There is a negative skew in the implied volatility surface and this factor dominates any convexity or smile effects. There is also a small increase in implied volatility with maturity.

In general leverage effects result in a negative skew in the shape of the implied volatility smile, and they may help explain some of the observed biases in market data. However it is not possible to capture volatility smiles in models which are motivated by leverage considerations alone.

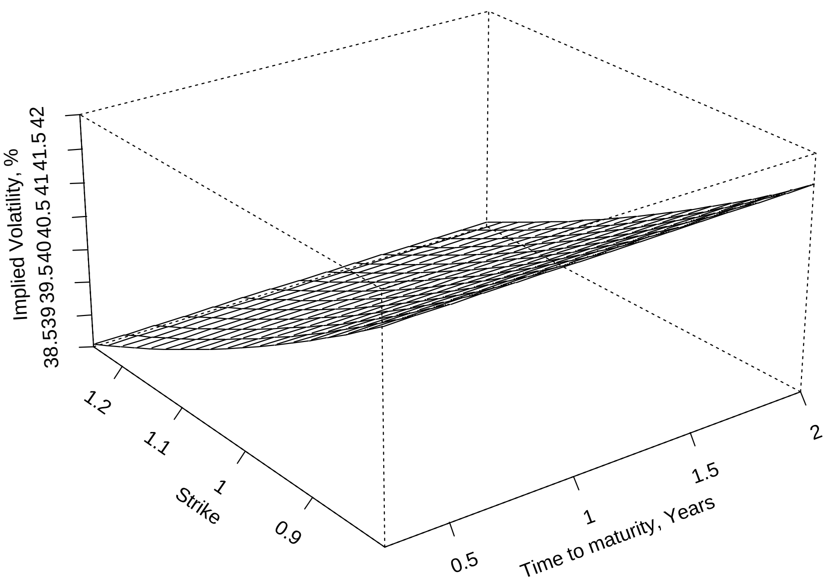

Figure 4: Implied volatilities for the CEV model. Note the strong inverse relationship between implied volatility and strike, which decreases only slightly with time.

2.3 Models of Stochastic Volatility

Volume and leverage effects can partially account for the observed patterns in volatility. However these explanations remain incomplete and more sophisticated models for the volatility are required which allow for further random changes in the level of volatility. In response to this need a series of models for asset price processes were proposed in the late 1980s which took volatility as an exogenous stochastic process.

Scott (1987), Wiggins (1987), Hull and White (1987, 1988), Stein and Stein (1991) and Heston (1993) each proposed models of the form

$$\frac{dP_t}{P_t} = \sigma_t dB_t + \mu dt \tag{8}$$

where σ_t, the stochastic volatility process, is itself the solution of a stochastic differential equation. Several candidate SDEs for the volatility process have been suggested. The candidate models have generally been motivated by intuition, convenience and a desire for tractability, rather than because of an empirical relationship with realised volatility. In particular the following models have all appeared in the literature:

$$d\sigma_t = \sigma_t(\alpha dt + \gamma dW_t) \tag{9}$$

$$d\sigma_t = \sigma_t((\alpha - \beta\sigma_t)dt + \gamma dW_t) \tag{10}$$

$$d\sigma_t = \beta(\alpha - \sigma_t)dt + \gamma dW_t \tag{11}$$

$$d\sigma_t = \left(\frac{\delta}{\sigma_t} - \beta\sigma_t\right) dt + \gamma dW_t \tag{12}$$

In each case W is a Brownian motion, perhaps correlated with the Brownian motion B which forms part of the specification (8). Denote this correlation by ρ so that $(dB_t dW_t) \equiv \rho dt$. We will assume that ρ is a constant with modulus less than one.

The model (9) was introduced by Hull and White (1987) who took $\rho \equiv 0$ and Wiggins (1987) who considered the general case. The volatility is an exponential Brownian motion (or equivalently the logarithm of the volatility is a drifting Brownian motion). Scott (1987) considered the case (10) in which the logarithm of the volatility is an Ornstein-Uhlenbeck (OU) process. The discrete time analogue of an OU process is an $AR(1)$ time series, see Section 4.1 below. The models specified by (9) and (10) have the advantage that the volatility is strictly positive for all time. However even though the model (11) allows the process σ to become negative, this need not be a major handicap since (8) remains well defined for negative values of σ, and it is possible to define volatility as the positive square root of the process σ_t^2. This third model was proposed by Scott (1987) and further investigated by Stein and Stein (1991). In both these articles the authors specialised to the case $\rho = 0$. In this model the volatility process itself is an OU process with mean reversion level α. The final model (12) was proposed by Hull and

White (1988) and Heston (1993). In this model the volatility is related to the square-root process of Cox, Ingersoll and Ross (1985) and σ can be interpreted as the radial distance from the origin of a multidimensional OU process.

Two other models of note were proposed by Johnson and Shanno (1987) who modelled both the price and volatility as CEV processes, and Melino and Turnbull (1990) who took the price to be a CEV process and the logarithm of the volatility to be an OU process.

2.4 Transition densities

Consider the model

$$\frac{dP_t}{P_t} = \sigma_t dB_t + \mu dt \tag{13}$$

$$d\sigma_t = \gamma(\sigma_t) dW_t + \nu(\sigma_t) dt \tag{14}$$

where, for the moment, B and W are *independent* Brownian motions. Then σ and B are independent and, conditional on $(\sigma_s)_{0\le s\le t}$, we have that $\int_0^t \sigma_s dB_s$ is a Gaussian random variable with zero mean and variance $V_t \equiv \int_0^t \sigma_s^2 ds$. In particular, from the analogue of the representation (2), we have that $P_t = P_0 \exp\{Z\}$ where Z is a Gaussian random variable with mean $\mu t - \frac{1}{2}V_t$ and variance V_t. Thus the transition density is a 'mixture of normals', with the mixing distribution depending on the autonomous stochastic process σ. If the value of the volatility is related to the rate of transactions then we recover the volume of transactions model described in Section 2.1, with a volume described by a random process.

Thus, for a stochastic volatility model in which the volatility is independent of the Brownian motion which drives the SDE for the price process, it is sufficient to characterise the law of V_t in order to derive the transition law for the price. Stein and Stein (1991) illustrate this result when the volatility process is an OU process and give an explicit form for the transition density.

When ρ is non-zero the interactions between the volatility and the driving Brownian motion complicate the analysis. However there is strong empirical evidence that ρ is non-zero. A negative value of ρ provides one method of capturing the observed negative correlation between volatility and price. Hence it is worthwhile to pursue the general case and to resort to numerical methods if necessary (see Johnson and Shanno (1987) and Wiggins (1987)). Heston (1993) has devised an efficient method for calculating options using characteristic functions.

It is possible to recover the level dependent volatility models of, for example, Cox and Ross and Geske by taking $|\rho| = 1$ and choosing an appropriate, though potentially unwieldy, specification for the parameters γ and ν in (14). If $\rho = 1$ then the diffusion (P_t, σ_t) is degenerate and there is a deterministic relationship between the processes P

and σ. The model is then similar in spirit to *GARCH* models, see Section 4.2. See also Hobson and Rogers (1996) who define a continuous time model of the form (8) in which B and W are perfectly correlated.

3 Option pricing for Stochastic Volatility Models

In this section we consider the option pricing implications of diffusion models for the volatility. In particular it is no longer true that there are unique preference independent options prices. Insted the model is incomplete and economic considerations (such as risk aversion) must be introduced to obtain pricing formulæ.

Suppose that P and σ are defined as in Section 2.4 above, without the assumption that B and W are independent. Indeed write $W_t \equiv \rho B_t + \sqrt{1-\rho^2} Z_t$ for a Brownian motion Z which is independent of B. Suppose that the aim is to price an option, and that the price of that option is given by a (differentiable) function H which depends on the current value of the asset, the current volatility and the time to go. Then we can apply Itô's formula to $H(P_t, \sigma_t, T-t)$ to obtain

$$dH = H_1 dP + H_2 d\sigma + \Lambda dt,$$

where suffices denote partial differentiation with respect to the relevant co-ordinate of H and

$$\begin{aligned}\Lambda dt &= \tfrac{1}{2} H_{11}(dP)^2 + H_{12}(dP)(d\sigma) + \tfrac{1}{2} H_{22}(d\sigma)^2 - H_3 dt \\ &= \{\tfrac{1}{2} H_{11} P^2 \sigma^2 + \rho\gamma H_{12} P\sigma + \tfrac{1}{2}\gamma^2 H_{22} - H_3\} dt.\end{aligned}$$

If volatility were a traded asset then it would be possible to invest in volatility and the stock to form a riskless hedge portfolio for the option. However this is not the case so there is no riskless hedge and the prices of options will depend on the risk preferences of investors. These preferences may be expressed via a utility function (see Hodges and Neuberger (1989) or Karatzas, Lehoczky, Shreve and Xu (1991)), or via a local-risk minimisation criterion (Hofmann, Platen and Schweizer (1992) or Platen and Schweizer (1994)).

Substituting for $d\sigma$ we obtain

$$\begin{aligned}dH &= H_1 P \frac{dP}{P} + H_2 \left(\nu dt + \frac{\gamma\rho}{\sigma} \left[\frac{dP}{P} - \mu dt \right] + \gamma \sqrt{1-\rho^2} dZ \right) + \Lambda dt \\ &= \left(H_1 P + \frac{\gamma\rho H_2}{\sigma} \right) \frac{dP}{P} + H_2 \gamma \sqrt{1-\rho^2} dZ + \left(H_2 \left[\nu - \frac{\gamma\rho\mu}{\sigma} \right] + \Lambda \right) dt.\end{aligned}$$

Now define $\Psi = \Psi(P_t, \sigma_t, T-t)$ via

$$\Psi \equiv - \left(H_1 P + \frac{\gamma\rho H_2}{\sigma} \right)$$

Observe that the martingale term of $(dH + \Psi(dP/P))$ only involves dZ so that a portfolio consisting of a call and an amount Φ invested in the stock is uncorrelated with the stock. Asset pricing models imply that the rate of return on this portfolio must be r with an extra return for risk:

$$\mathbb{E}\left[dH + \Psi\frac{dP}{P}\right] = r(H+\Psi)dt - \lambda^* H_2\gamma\sqrt{1-\rho^2}dt$$

where λ^* is the market price of volatility risk associated with dZ. Typically the value $H + \Psi$ of the portfolio is negative which explains the sign convention for the market price of risk. Equating finite variation terms we obtain

$$H_2\left(\nu - \frac{\gamma\rho\mu}{\sigma}\right) + \Lambda = r(H+\Psi) - \lambda^* H_2\gamma\sqrt{1-\rho^2}$$

Finally some algebraic manipulation of this equation yields the stochastic volatility option pricing partial differential equation for H:

$$\begin{aligned} &\tfrac{1}{2}H_{11}P^2\sigma^2 + \rho\gamma H_{12}P\sigma + \tfrac{1}{2}\gamma^2 H_{22} - H_3 - rH \\ &\quad + rH_1P + H_2\left(\nu - \frac{\gamma\rho(\mu - r)}{\sigma} + \lambda^*\gamma\sqrt{1-\rho^2}\right) = 0 \end{aligned} \tag{15}$$

subject to the boundary condition $H(x, y, 0) = (x - K)^+$. Thus the price of an option has an interpretation as the expected payoff of the option under a model in which the price process and the volatility satisfy the SDEs

$$\begin{aligned} \frac{dP}{P} &= \sigma_t dB + r dt \\ d\sigma_t &= \gamma(\sigma_t)dW + \tilde{\nu}(\sigma_t)dt \end{aligned}$$

where

$$\tilde{\nu}(\sigma) = \nu(\sigma) - \frac{\gamma(\sigma)\rho(\mu - r)}{\sigma} + \lambda^*(\sigma)\gamma(\sigma)\sqrt{1-\rho^2}.$$

The option pricing equation (15) has an analogue in expressions given by Wiggins (1987, Equation (8)), Scott (1987, Equation (4)) and Stein and Stein (1991, Equation (14)). In principle we solve (15) to deduce theoretical options prices. Before we comment on the discussion in the literature on the (numerical) solutions of (15) some general comments are in order.

Firstly suppose that $\gamma \equiv 0 \equiv \nu$ so that the stochastic process (σ_t) is in fact a deterministic constant. Then we can view σ as a constant parameter of the model rather than a stochastic variable, and the option pricing equation (15) for the price $C \equiv C(P, u)$ of a call option as a function of the price P of the underlying asset and the time to go u reduces to

$$\tfrac{1}{2}P^2\sigma^2 C_{PP} - C_u + rPC_P - rC = 0$$

with boundary condition $C(x,0) = (x-K)^+$. This is the Black-Scholes partial differential equation, and C is the Black-Scholes price of an option.

Secondly, if volatility is stochastic but uncorrelated with the asset, (so that $\rho = 0$), then the option price can be expressed as

$$\begin{aligned} H(P_t, \sigma_t, T-t) = \tilde{\mathbb{E}}[(P_T - K)^+] &= \tilde{\mathbb{E}}[\tilde{\mathbb{E}}[(P_T - K)^+ | (\sigma_s)_{t \leq s \leq T}]] \\ &= \tilde{\mathbb{E}}\left[C\left(\left(\frac{1}{T-t}\int_t^T \sigma_s^2 ds\right)^{1/2}\right)\right]. \end{aligned}$$

Thus the option price is an average of Black-Scholes prices. To investigate this relationship further consider the dependence of the Black-Scholes formula on σ. Suppose $t = 0$ and define $\sigma_I = \sqrt{(2/T)|\ln m|}$ where m is the moneyness of the option. Thus σ_I is zero for at the money options. Then C is convex in σ for $\sigma \leq \sigma_I$, and concave for $\sigma \geq \sigma_I$; see Figure 5. Thus for an at the money option with Black-Scholes implied volatility $\tilde{\sigma}$ it follows from Jensen's inequality that

$$C(\tilde{\sigma}) \equiv \tilde{\mathbb{E}}[C((V_T/T)^{1/2})] \leq C(\tilde{\mathbb{E}}((V_T/T)^{1/2})$$

where $V_T = \int_0^T \sigma_s^2 ds$. By monotonicity of the Black-Scholes formula, for an at the money option, the Black-Scholes implied volatility is less than the expected average volatility, under the risk-neutral pricing measure. Conversely, for a far in or out of the money option, then for σ_0 sufficiently small

$$C(\tilde{\sigma}) \equiv \tilde{\mathbb{E}}[C((V_T/T)^{1/2})] \geq C(\tilde{\mathbb{E}}((V_T/T)^{1/2}).$$

Thus we expect that the implied volatility for away from the money options will exceed the expected average volatility, and that there will be an implied volatility smile.

Renault and Touzi (1995) show that, again in the case $\rho = 0$, the volatility smile is symmetric. Consider the Black-Scholes call price C as a function of the moneyness M_t and the time to go u, then (5) yields

$$\begin{aligned} C(M^{-1}, u) &= Ke^{-rT}(M^{-1}\Phi(-d_2) - \Phi(-d_1)) \\ &= \frac{Ke^{-rT}}{M}((1-M) + M\Phi(d_1) - \Phi(d_2)) \\ &= M^{-1}C(M,u) + Ke^{-rT}(M^{-1} - 1). \end{aligned}$$

Hence there is a simple expression relating the prices of in and out of the money calls. Moreover, if we think of the stochastic volatility option pricing function H as a function of moneyness, the current value of the volatility σ_t, and the time to expiry u, then an investigation of the solutions to (15) yields that, provided $\rho = 0$,

$$H(M^{-1}, \sigma_t, u) = M^{-1}H(M, \sigma_t, u) + Ke^{-rT}(M^{-1} - 1)$$

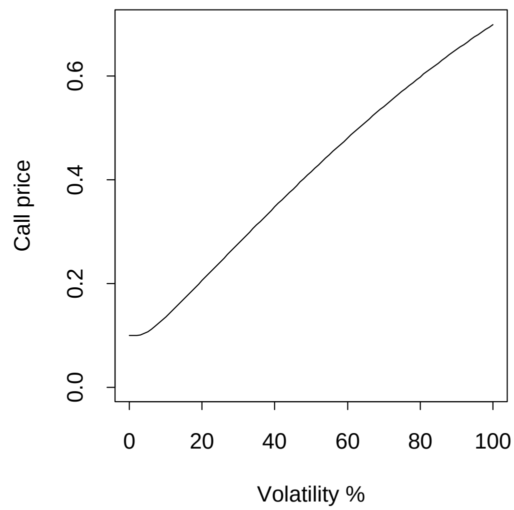

Figure 5: The price of an option as a function of volatility. The plot is based on an option with strike $K = 0.9$, and expiry $T = 3$, for an underlying asset whose price is unity. Thus $\sigma_I = 26.5\%$.

To verify this claim observe that $MH(M^{-1}, \sigma, u) + Ke^{-rT}(1 - M^{-1})$ also solves (15) and the same boundary condition. From this it is a simple exercise to deduce that if the option with moneyness M has an implied volatility $\tilde{\sigma}_M$, so that $H(M, \sigma_t, u) = C(M, u; \tilde{\sigma}_M)$, then also $H(M^{-1}, \sigma_t, u) \equiv C(m^{-1}, u; \tilde{\sigma}_M)$. Now, since $\tilde{\sigma}_{M^{-1}}$ is the value of the implied volatility for which $H(M^{-1}, \sigma_t, u) \equiv C(m^{-1}, u; \tilde{\sigma}_{M^{-1}})$, we must have that $\tilde{\sigma}_M \equiv \tilde{\sigma}_{M^{-1}}$.

In more general situations with non-zero correlation ρ the picture is more complicated. Several authors have attempted to solve (15) in this case. Hull and White (1988) consider solutions which take the form of power series expansions in the volatility of volatility parameter γ:

$$H(M, \sigma_t, u) = C(M, u) + f_0(M, \sigma_t, u) + \gamma f_1(M, \sigma_t, u) + \gamma^2 f_2(M, \sigma_t, u) + \dots$$

Explicit, though complicated, forms can be deduced for the functions $f_0, f_1, f_2 \dots$. See Figure 6 for the predicted implied volatility surface based on an expansion to second order of the Hull and White option pricing series. Alternatively Johnson and Shanno (1987), Wiggins (1987) and Heston (1993) calculate numerical solutions to (15). In each case the authors find that when the correlation is negative out of the money call options are relatively more expensive under a stochastic volatility model when compared with Black-Scholes prices. This is consistent with the biases found in Rubinstein (1985), Wiggins (1987), Heynen, Kemma and Vorst (1994), and Figure 3. Wiggins attempts to derive an estimate for ρ. His estimates support the hypothesis that ρ is negative, but the precise estimates vary widely depending on the particular method he uses. However he does provide evidence that the negative correlation is more pronounced for indices rather than individual stocks.

Continuous-time stochastic volatility provides an attractive and intuitive explanation for observed volatility patterns and for observed biases in implied volatility. In particular smiles, skews and upward and downward implied volatility term structures arise naturally from a stochastic volatility model. However the fact that stochastic volatility models fit empirical patterns is not conclusive evidence that those models are correct and the biases in market prices may be the result of other factors such as liquidity problems.

4 Discrete-time Models

Whilst continuous time models provide the natural framework for an analysis of option pricing, discrete time models are ideal for the statistical and descriptive analysis of the patterns of daily price changes. There are two main classes of discrete-time models for stock prices with volatility. The first class, the *autoregresive random variance* (ARV) or stochastic variance models are a discrete time approximation to the continuous time diffusion models we outlined in Sections 2 and 3. The second class of *autoregressive conditional heteroskedastic* (ARCH) models and its descendents are motivated by an attempt to explain volatility clustering and the habit of large price changes to be followed by further large changes.

4.1 ARV models

Let $Y_n = \ln P_n$ so that Y_n denotes the log price. Then the natural discrete time analogue of (13) and (14) is to take

$$Y_n = Y_{n-1} + \nu + \sigma_{n-1} Z_n \tag{16}$$

where $(Z_n)_{n\geq 0}$ is a sequence of independent standard normal variables and σ_n is the solution of a stochastic difference equation. Many authors including Chesney and Scott (1989) and Duffie and Singleton (1993) consider a model of the form

$$\ln \sigma_n = \alpha - \phi(\ln \sigma_{n-1} - \alpha) + \theta z_n \tag{17}$$

for parameters α, ϕ, θ and z_n a sequence of independent identically distributed random variables such that (Z_n, z_n) forms a bivariate normal sequence with correlation ρ. Equation (17) is a direct analogue of (10). The model specified by (16) and (17) is called an *ARV* model (Taylor (1986)).

The *ARV* model is stationary if $|\phi| < 1$ and then $\ln \sigma_t$ has mean α and variance $\beta = \theta^2/(1-\phi^2)$. Provided that $\rho = 0$ the unconditional distribution of the return $Y_t - Y_{t-1}$

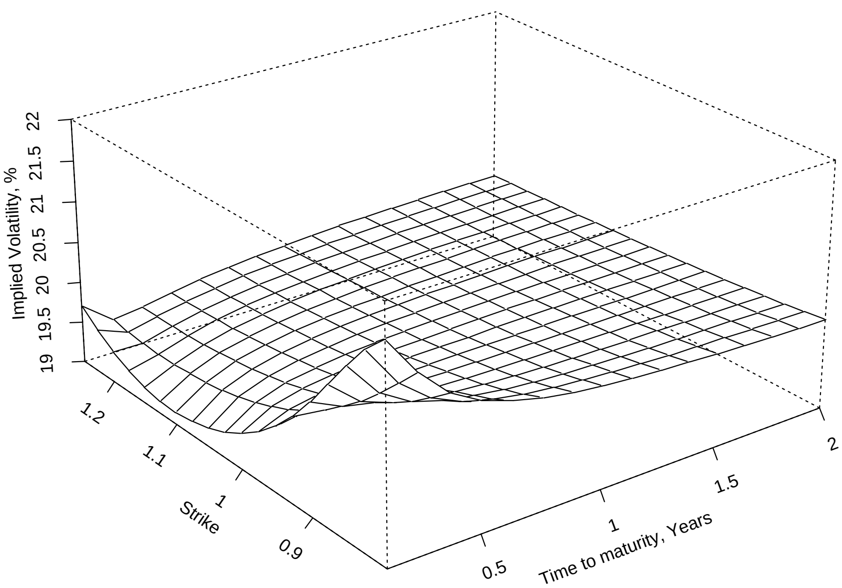

Figure 6: Implied volatilities from the Hull White expansion to second order for an option on an underlying asset whose current price is unity. Note that $\rho = -0.2$ and that for options which are close to maturity there is a pronounced volatility smile, and some evidence of an additional inverse relationship between strike and implied volatility. These effects decrease with maturity.

is a mixture of normal distributions, with analogues in the rate of transaction models of Clark (1973) and Tauchen and Pitts (1983).

In the *ARV* model the volatility process is unobservable which contrasts with the continuous time situation in which the instantaneous value for volatility can be inferred from the quadratic variation of the log-price. As an unfortunate consequence most *ARV* models lack one-step transition densities for the process Y_n. This means that it is frequently not possible to obtain maximum likelihood estimates for parameter values.

Instead parameter values are frequently estimated using methods of moments techniques, see Taylor (1986), Melino and Turnbull (1990) and Duffie and Singleton (1993). Of particular interest is the autoregressive co-efficient ϕ which governs the persistence of volatility shocks. According to Taylor (1994) most estimates of this parameter which are based on daily observations yield values greater than 0.95. Harvey, Ruiz and Shephard (1996) find that a multivariate *ARV* model fits well to to exchange rates data and captures movements in volatility, though for certain currencies they are led to suggest a heavy tailed distribution for the innovations process. See Ghysels, Harvey and Renault (1996) for a thorough discussion of *ARV* models and their statistical properties.

Since the *ARV* model is an approximation to diffusion models of stochastic volatility there is a correspondence between options prices in an *ARV* model and numerical solutions of the stochastic volatility option pricing equation (15). Thus options prices in *ARV* models are preference dependent, and an ARV model can account for smiles and skews in implied volatility.

4.2 ARCH and GARCH models

Autoregressive conditional heteroskedastic models were introduced by Engle (1982) in an attempt to model persistence in volatility shocks by assuming an autoregressive structure for the conditional variances. Retaining the convention that $Y_n = \ln P_n$ an *ARCH* model assumes that

$$Y_n = Y_{n-1} + \nu + \eta_n \epsilon_n \qquad \epsilon_n \text{ i.i.d. } D(0,1) \tag{18}$$

where D is a general distribution with zero mean and unit variance, and η_n is a function of the past proportional price changes. The simplest *ARCH* model, an *ARCH*(1) combines (18) with

$$\eta_n^2 = \alpha + \beta(Y_{n-1} - Y_{n-2} - \nu)^2 = \alpha + \beta\eta_{n-1}^2\epsilon_{n-1}^2 \tag{19}$$

ARCH models have the advantage that it is straightforward to write down the log-likelihood and hence to derive maximum likelihood estimators for the parameters.

In empirical applications higher order *ARCH*(q) models with a large number of parameters are often needed to characterise the behaviour of financial time series. To circumvent this problem Bollerslev (1986) devised a class of *generalised ARCH* or *GARCH* models which allow the conditional variance to depend directly on previous values. In a *GARCH*(1,1) model (independently proposed by Taylor (1986)) we have that

$$\eta_n^2 = \alpha + \beta(Y_{n-1} - Y_{n-2} - \nu)^2 + \gamma\eta_{n-1}^2 = \alpha + \beta\eta_{n-1}^2\epsilon_{n-1}^2 + \gamma\eta_{n-1}^2 \tag{20}$$

Other extensions are also possible, see Bollerslev, Engle and Nelson (1994), Harvey, Ruiz and Shephard (1994) or Shephard (1996) for comprehensive surveys.

As defined in (19) and (20) the updates of the conditional variance depend on the squares of the residual process. Hence these simple models cannot capture leverage effects. However the exponential *ARCH* model of Nelson (1991) does not treat positive and negative innovations symmetrically and can allow for a correlation between volatility, as expressed by η, and price level.

The natural candidate distribution for $D(0,1)$ is standard normal. However some empirical studies of stock prices, including for example Bollerslev (1986) and Bollerslev, Engle and Nelson (1994), have found that the standardised innovations process $(Y_n - Y_{n-1} - \nu)\eta_n^{-1}$ displays excess kurtosis. Taylor (1994) suggests use of a scaled t-distribution or a generalised error distribution. With these choices there are two sources of the kurtosis in the unconditional distribution for the log-price, namely the kurtosis from the price innovations and the changes in the underlying volatility level η.

GARCH models have been extremely successful in the modelling of equity markets. Highly significant test statistics have been reported by Engle and Mustafa (1992) in an analysis of stock returns, and Schwert (1990) for futures markets. See Bollerslev, Chou and Kroner (1992) for an extensive survey of articles reaching similar conclusions. The autoregressive structure imposed by *GARCH* model for the conditional variances allows volatility to persist over time and captures the observed clustering of large price movements. Note however that Lamoureux and Lastrapes (1990) find that daily trading volume has a significant explanatory power regarding the variance of daily returns and that furthermore *ARCH* effects tend to disappear when volume is included in the variance equation.

In an *ARV* model the volatility process is an autonomous process. In contrast in *GARCH* models the volatility process is a deterministic function of the innovations ϵ_n. Neverthless Nelson (1990) has shown that with judicious choice of parameter values the continuous time limit of a *GARCH* process is a diffusion model with stochastic volatility of the form (8) and (11). However the rate of convergence to the diffusion limit is much slower than that from an *ARV* model. Bollerslev, Engle and Nelson (1994) show both how a *GARCH* model can be used to approximate a diffusion and how a diffusion process can be used to approximate a *GARCH* model.

Consideration of the SDE high frequency limit of *GARCH* processes raise the problem of temporal aggregation of GARCH processes. The non-linearities of GARCH models mean that if a low-frequency sample is taken from a high frequency *GARCH* model, then the resulting time series is not *GARCH*, at least in the sense defined above. Although Drost and Nijman (1993) have introduced the concept of *weak GARCH* models which are stable under temporal aggregation, in general the frequency of observations has an important bearing on the statistical properties of the model. For example an i.i.d. innovations process in the definition of the price process $(Y_n)_{n\geq 0}$ will generally result in time dependence of the innovations of $(Y_{kn})_{n\geq 0}$ for $k > 1$.

4.3 GARCH option pricing

Even for discrete time models the pricing of options remains an important issue, and there has been much recent interest in *GARCH* option pricing formulæ which is summarised in the paper by Duan (1995).

Duan assumes a model in which the innovations ϵ_n are normal variables and ν takes the form

$$\nu = r + \lambda^* \eta_{n-1} - \tfrac{1}{2}\eta_{n-1}^2$$

where λ^* is a volatility risk premium, and the $-\frac{1}{2}\eta_{n-1}^2$ term ensures that when $\lambda^* = 0$ the discounted price process is a martingale. With this specification the price process evolves as

$$P_n = P_{n-1} \exp\{r + \lambda^* \eta_{n-1} + \eta_{n-1}\epsilon_n - \tfrac{1}{2}\eta_{n-1}^2.\}$$

The concept of a *locally risk neutral valuation relationship* is used to argue that options should be priced as the discounted expected payoff under a model in which the price and volatility update according to the stochastic difference equations

$$\begin{aligned} Y_n &= Y_{n-1} + r - \tfrac{1}{2}\eta_{n-1} + \eta_{n-1}\tilde{\epsilon}_n \\ \eta_n^2 &= \alpha + \beta\eta_{n-1}(\tilde{\epsilon}_{n-1} - \lambda^*)^2 + \gamma\eta_{n-1}^2 \end{aligned}$$

for an i.i.d. sequence $\tilde{\epsilon}_n$ of standard normal variables.

Duan estimates the parameters of the model from market prices and uses Monte Carlo techniques to obtain options prices. He finds implied volatility smiles which become weaker as time to maturity increases. Depending on the initial value of the conditional volatility the term structure of implied volatility of an at the money option can be either downward or upward sloping.

In general prices from the *GARCH* option pricing model are consistent with the biases found by Rubinstein (1985) and Skeikh (1991). However although this evidence supports the *GARCH* modelling hypothesis, it cannot guarantee the veracity of the model. Moreover Monte Carlo techniques are a computationally expensive technique for calculating options prices.

5 Conclusions

The Black-Scholes exponential Brownian motion model provides an approximate description of the behaviour of asset prices and a benchmark against which other models can be compared. The volatility parameter is a crucial component of the model and stochastic volatility models aim to reflect the apparent randomness of the level of volatility, as observed in empirical studies.

To this extent stochastic volatility models are partially successful and moreover they can capture, and potentially explain, some of the observed biases in the Black-Scholes formula for options. Both diffusion models and $GARCH$ models can account for smiles, skews and term structures which have been observed in market prices for options, and stochastic volatility models are widely used in the financial community as a refinement of the Black-Scholes model. Exotic options are frequently even more sensitive to levels of volatility than standard calls, and as trading in such instruments blossoms, those financial institutions which have models with the ability to reasonably and consistently price and hedge derivatives will have a competitive advantage.

References

BACHELIER, L. (1900): Théorie de la speculation. *Annales Scientifiques de l'Ecole Normale Supérieure, Troisiéme serie* **17**, 21-88. Translated in *The random character of stock market prices*, ed: Cootner, P., MIT Press (1965).

BENSOUSSAN, A., CROUHY, M. and GALAI, D. (1994): Stochastic equity volatility and the capital structure of the firm. *Philosophical Transactions of the Royal Society of London*, Series A, **347**, 449-598.

BLACK, F. and SCHOLES, M. (1973): The pricing of options and corporate liabilities. *Journal of Political Economy,* **81**, 637-659.

BLATTBERG, R.C. and GONEDES, N.J. (1974): A comparison of the stable and Student distributions as statistical models for stock prices. *Journal of Business*, **47**, 244-280.

BOLLERSLEV, T. (1986): Generalised autoregressive conditional heteroskedasticity. *Journal of Econometrics*, **31**, 307-327.

BOLLERSLEV, T., CHOU, R.Y. and KRONER, K.F. (1992): *ARCH* modelling in finance: a review of the theory and empirical evidence. *Journal of Econometrics*, **52**, 5-59.

BOLLERSLEV, T., ENGLE, R.F. and NELSON (1994): *ARCH* models. In *The Handbook of Econometrics, Vol 4*, ed Engle, R.F. and McFadden, D.L.. Elsevier, Amsterdam.

CHESNEY, M. and SCOTT, L.O. (1989): Pricing European options: a comparison of the modified Black-Scholes model and a random variance model. *Journal of Financial and Quantitative Analysis,* **24**, 267-284.

CLARK, P.K. (1973): A subordinated stochastic process model with finite variance for speculative prices. *Econometrica*, **41**, 135-155.

COX, J.C., INGERSOLL, J.E. and ROSS, S.A. (1985): An intertemporal general equilibrium model of asset prices. *Econometrica*, **53**, 363-384.

COX, J.C. and ROSS, S.A. (1976): The valuation of options for alternative stochastic processes. *Journal of Financial Economics*, **3**, 145-166.

DROST, F.C., and NIJMAN, T.E. (1993): Temporal aggregation of GARCH processes. *Econometrica*, **61**, 909-927.

DUAN, J-C. (1995): The *GARCH* option pricing model. *Mathematical Finance*, **5**, 13-32.

DUFFIE, D. and SINGLETON, K.J. (1993): Simulated moments estimation of Markov models of asset prices. *Econometrica*, **61**, 929-952.

ENGLE, R.F. (1982): Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, **50**, 987-1007.

ENGLE, R.F. and MUSTAFA, C. (1992): Implied *ARCH* models from options prices. *Journal of Econometrics*, **52**, 289-311.

EPPS, T.W. and EPPS, M.L. (1976): The stochastic dependence of security price changes and transaction volumes: implications for the mixture-of-distributions hypothesis. *Econometrica*, **44**, 305-321.

FAMA, E.F. (1965): The behaviour of stock market prices. *Journal of Business*, **38**, 34-105.

FUNG, W.K.H. and HSIEH, D.A. (1991): Empirical analysis of implied volatility: stocks, bonds and currencies. *Proceedings of the fourth FORC conference, University of Warwick.*

GESKE, R. (1979): The valuation of compound options. *Journal of Financial Economics*, **7**, 63-81.

GHYSELS, E., HARVEY, A. and RENAULT, E. (1996): Stochastic Volatility. In *Handbook of Statistics, Vol 14, Statistical Methods in Finance.*

HARVEY, A., RUIZ, E. and SHEPHARD, N. (1994): Multivariate stochastic variance models. *Review of Economic Studies*, **61**, 247-264.

HESTON, S.L. (1993): A closed-form solution for options with stochastic volatiltiy with applications to bond and currency options. *Review of Financial Studies*, **6**, 327-343.

HEYNEN, R., KEMMA, A. and VORST, T. (1994): Analysis of the term structure of implied volatility. *Journal of Financial and Quantitative Analysis*, **29**, 31-56.

HOBSON, D.G. and ROGERS, L.C.G. (1996): Complete models of stochastic volatility. To appear in *Mathematical Finance.*

HODGES, S.D. and NEUBERGER, A. (1989): Optimal replication of contingent claims under transaction costs. *Review of Futures Markets*, **8**, 222-239.

HOFMANN, N., PLATEN, E. and SCHWEIZER, M. (1992): Option pricing under incompleteness and stochastic volatility. *Mathematical Finance*, **2**, 153-187.

HULL, J. and WHITE, A. (1987): The pricing of options on assets with stochastic volatilities. *Journal of Finance*, **42**, 281-300.

HULL, J. and WHITE, A. (1988): An analysis of the bias in option pricing caused by a stochastic volatility. *Advances in Futures and Options Research*, **3**, 29-61.

JOHNSON, H. and SHANNO, D. (1987): Option pricing when the variance is changing. *Journal of Financial and Quantitative Analysis*, **22**, 143-151.

KARATZAS, I., LEHOCZKY, J.P., SHREVE, S.E. and XU, G.-L.. (1991): Martingale and duality methods for utility maximisation in an incomplete market. *SIAM Journal of Control and Optimisation*, **29**, 702-730.

KARPOFF, J.M. (1987): The relation between price changes and trading volume. *Journal of Financial and Quantitative Analysis*, **22**, 109-126.

LAMOUREUX, C.G. and LASTRAPES, W.D. (1990): Heteroskedasticity in stock return data: volume versus *GARCH* effects. *Journal of Finance*, **45**, 221-229.

MADAN, D.B. and SENETA, E. (1990): The variance gamma (V.G.) model for share market returns. *Journal of Business,* **63**, 511-524.

MANDLEBROT, B.B. (1963): The variation of certain speculative prices. *Journal of Business,* **36**, 394-416.

MELINO, A., and TURNBULL, S.M. (1990): Pricing foreign currency options with stochastic volatility. *Journal of Econometrics*, **45**, 239-265.

NELSON, D.B. (1990): *ARCH* models as diffusion approximations. *Journal of Econometrics*, **45**, 7-38.

NELSON, D.B. (1991): Conditional heteroskedasticity inasset returns: a new approach. *Econometrica*, **59**, 347-370.

ØKSENDAL, B. (1985): *Stochastic Differential Equations*, Springer-Verlag, Berlin.

OSBORNE, M.F.M. (1959): Brownian motion in the stock market. *Operations Research*, **7**, 145-173.

PETERS, E.E. (1991): *Chaos and order in the capital markets.* Wiley, New York.

PLATEN, E. and SCHWEIZER, M. (1994): On smile and skewness. *Discussion paper No B-302, University of Bonn .*

RENAULT, E. and TOUZI, N. (1995): Option hedging and implied volatilities in a stochastic volatility model. Preprint.

RUBINSTEIN, M. (1983): Displaced diffusion option pricing. *Journal of Finance*, **38**, 213-217.

RUBINSTEIN, M. (1985): Nonparametric tests of alternative option pricing models using all reported trades and quotes on the 30 most active CBOE option classes from Aug. 23 1976 through Aug. 31 1978. *Journal of Finance*, **40**, 455-480.

SAMUELSON, P.A. (1965): Rational theory of warrant pricing. *Industrial Management Review*, **6**, # 2, 13-39.

SCHRODER, M. (1989): Computing the constant elasticity of variance option pricing formula. *Journal of Finance*, **44**, 211-219.

SCHWERT, G.W. (1990): Stock volatility and the crash of 87. *Review of Financial Studies*, **3**, 77-102.

SCOTT, L.O. (1987): Option pricing when the variance changes randomly: theory, estimation and an application. *Journal of Financial and Quantitative Analysis,* **22**, 419-438.

SHEIKH, A. (1991): Transaction data tests of S&P 100 call option pricing. *Journal of Financial and Quantitative Analysis*, **26**, 459-475.

SHEPHARD, N. (1996) Statistical aspects of *ARCH* and stochastic volatility. In *Time Series Models in Econometrics, Finance and other fields.* Editors: Cox, D.R., Hinkley, D.V. and Barndorff-Nielsen, O.E.. Chapman and Hall Monographs in Statistics and Applied Probability, 65.

STEIN, E.M. and STEIN, J.C. (1991): Stock price distributions with stochastic volatility: an analytic approach. *Review of Financial Studies*, **4**, 727-752.

TAUCHEN, G.E. and PITTS, M. (1983): The price variability-volume relationship on speculative markets. *Econometrica*, **51**, 485-505.

TAYLOR, S.J. (1986): *Modelling financial time series.* Wiley, Chichester, UK.

TAYLOR, S.J. (1994): Modelling stochastic volatility: a review and comparative study. *Mathematical Finance*, **4**, 183-204.

WIGGINS, J.B. (1987): Option values under stochastic volatility. Theory and empirical estimates. *Journal of Financial Economics,* **19**, 351-372.

XU, X and TAYLOR, S.J. (1994): The term structure of volatility implied by foreign exchange options. *Journal of Financial and Quantitative Analysis*, **29**, 57-74.