Stochastic Volatility Models: Considerations for the Lay Actuary¹

Phil Joubert Coomaren Vencatasawmy

(Presented to the Finance & Investment Conference, 19-21 June 2005)

Abstract

Stochastic models for asset prices processes are now familiar to actuaries. Many of the models used in life office and pension fund valuation and asset-liability modelling studies assume deterministic volatility parameters.

Empirical evidence however, suggests that volatility in asset prices varies with time. Further, volatilities implied by traded option prices show a term structure for implied volatility, as well as an apparent dependence on the "moneyness" of the option. These observations seem to be at odds with a constant volatility assumption.

In this paper we present some empirical observations concerning volatility, and consider the impact of volatility on actuarial work. We then review some of the common models which incorporate stochastic volatility and consider issues related to parameterising such models.

¹ We are indebted to Stuart Jones for his patience and for providing invaluable technical assistance. We would also like to thank Dylan Brooks and Carmela Calvosa for providing data and fruitful discussions.

1. Introduction

Volatility is central to many applied issues in finance and financial engineering, ranging from asset pricing and asset allocation to risk management. Financial economists have always been intrigued by the very high precision with which volatility can be estimated under the diffusion assumption routinely invoked in theoretical work. The basic insight follows from the observation that precise estimation of diffusion volatility does not require a long calendar span of data; rather, volatility can be estimated arbitrarily well from an arbitrarily short span of data, provided that returns are sampled sufficiently frequently. This contrasts sharply with precise estimation of the drift, which generally requires a long calendar span of data, regardless of the frequency with which returns are sampled. There is also the baffling range of volatility terms used: "Historical Volatility", "Implied Volatility", "Forecast Volatility", etc.. In this paper, the first two terms are most important.

Historical volatility is a measure of the previous fluctuations in share price (crudely: an indicator of the share's *up/downess*). There is much discussion over the best method of calculating the historic volatility. The most usual method is the standard deviation of the log of price returns - this procedure is fairly standard and can be found in most textbooks. While the calculation itself is straight-forward, it is accurate only within the parameters of each calculation (e.g. the specific time period: 3 months, 3 years etc.). There is great scope for analysing the share price behaviour over different time periods, and thereby calculating different historic volatilities.

Instead of inputting a volatility parameter into an option model (e.g. Black-Scholes) to determine an option's fair value, the calculation can be turned round, where the actual current option price is *input* and the volatility is *output*. The term *implied volatility* is obviously self-explanatory - that level of volatility that will calculate a fair value actually equal to the current trading option price. This calculation can be very useful when comparing different options. The implied volatility can be regarded as a measure of an option's "expensiveness" in the market, and is used by traders setting up combination strategies, where they have to identify relatively cheap and expensive options (even though these options have different terms). It is perhaps useful to note that implied volatility only has any meaning in the context of a particular option model (it is not intrinsic to the option itself). So, although options have existed for a long time, implied volatility has only had any meaning since the option pricing model of Fisher Black and Myron Scholes (devised in the early 1970's) stated that the value of an option was a function of the volatility of the underlying share price.

To calculate the fair value, an option model requires the input of volatility, or, more precisely, the input of: *forecast volatility of the share price over the period to expiry of the option*. The big question (the *art*) of option theory is how to estimate this forecast volatility. This single stage provides gainful employment for a legion of academics, analysts and traders. An *estimate of future share price fluctuations* - plenty of room for "wooliness" there!

Below we explore the relationship between these three concepts, but first we motivate our paper: why volatility, and why bother with volatility models?

The paper is organised as follows: Section one gives a brief history of the concept of volatility, and introduces some of the products traded in the broader financial markets. Section two discusses the definitions of, and differences between various concepts of volatility, and considers problems related to reliably estimating their values. In section three we review evidence that volatility is random, and in section four consider several models of asset prices which attempt to capture this. Section five comprises of the application of these models to actuarial problems – as alluded to above, the problem of market-consistent valuation of life insurance business, and effective risk control of the same. We summarise our conclusions in section six.

Note that we will often refer to "options" in this paper. Everything applicable to exchange traded options can be read as being applicable to life contracts with guarantees. We take this as dictated by the regulators, and do not enter into the debate as to whether this is the most appropriate methodology.

1.1. Why volatility?

Recently "volatility" has become part of the standard actuarial jargon. Particularly for life company actuaries, discussions about volatility – what it is, the appropriate value it should take, how it behaves, and what it does, are now part of the job.

Modern (Life Company) actuarial work consists broadly of two tasks – the valuation of assets and liabilities, and risk control.

For those involved in valuation, the move to market-consistent valuation has meant applying option pricing techniques to life insurance contracts. At the heart of these valuation techniques is the concept of volatility, so volatility has been placed directly on the balance sheets of life companies – both on the asset and liability side. As a result of this auditors are asking increasingly sophisticated questions about volatility parameters – are they appropriate for the contracts in question, will they reproduce market prices at a fine level of detail, or only in aggregate, and if so, why? Key concepts to answering these questions are: the assets in a given fund, the terms of the liabilities, and the levels of guarantees of those liabilities. Each combination will produce a different "market consistent" value for the liabilities, and requires a different value of implied volatility. In [24] Sheldon and Smith provide methodologies to make the appropriate choice of implied volatility based on the assets in a fund. In this paper we will address the remaining two issues.

Risk control involves calculating the possible movements of assets and liabilities on an insurer's balance sheet. Again, volatility is a key parameter, and in this case it is internal risk managers, and possibly the regulator, who should be asking tough questions about valuation techniques and particularly about volatility parameters.

A longer term concern is the increasing number of investment banks developing structured products aimed for sale to insurers. Many of these products will contain complex derivative instruments. Understanding the two related concepts of volatility – volatility influencing asset price movements, and implied volatility

governing derivative values, will allow us make informed decisions on the value and risks of such products.

The fact that the well-known Black-Scholes model does not fit empirical observations will be familiar to most readers. Below we consider a class of models which attempt to more closely explain the data. In particular: option prices given by the Black-Scholes model do not fit those observed in the market, and asset price movements have fat-tails. Stochastic volatility models may allow the valuation actuary to achieve a closer match to the relevant market-traded assets, and may allow the risk control actuary to capture an area of risk which might otherwise be ignored.

1.2. A brief history of volatility

It has become traditional recently in any paper concerning options to make a historical reference to Louis Bachelier. Since most actuaries will be unfamiliar with Bachelier, we provide a brief summary.

In his 1900 paper "*Théorie de la Spéculation*" [2], Bachelier considered a model of stock markets where prices follow what is now known as a Wiener process – Brownian motion. In this paper he derived, amongst other things, the price of a barrier option – a full 73 years before Fisher Black and Myron Scholes published their famous paper [4]. This result was remarkable, given that it predates the birth of modern statistics, and even Einstein's 1905 paper on Brownian motion. So far ahead of its time, it was lost until fairly recently, and instead the glory went to Black and Scholes.

Most readers will recall that Black and Scholes' surprising (at the time) result was that the value of an option is independent of the expected return of the underlying stock – but explicitly dependent on its (expected) volatility.

Thus the financial-mathematical concept of volatility appears to date back at least 100 years. No doubt the purely financial concept (intuitively the amount of variation manifest in stock prices in a given time period – we consider the concept in some depth below) has been recognised ever since men began trading together in markets.

Although there is some evidence that option-type contracts were used in ancient Greece, Rome and the Arab world, trading in modern Black-Scholes style options (with their explicit dependence on volatility) began on the Chicago Board of Exchange in 1973 with calls written on 16 stocks. Trading in puts started 4 years later [7].

Growth in the options market has exploded since then, with notional principle on outstanding exchange traded derivatives estimated at \$52trn, with OTC (Over The Counter) contracts perhaps accounting for five times that [12].

The growth in trade of options – essentially trading in volatility – has lead to increasingly complicated and sophisticated strategies being undertaken. Market players wishing to gain exposure purely to volatility movements can, for example, adopt positions in options, and then "delta hedge", or adopt a position in the underlying which negates the effect of movements in the price of the underlying –

but leaves exposure to changes in implied volatility. Another possible trade is to take positions in two options with different strikes - known as a "strangle" - a highly risky trade with pure volatility exposure. More recently several standardised products have begun trading which offer cheaper and easier access to pure volatility plays.

In 1998 two new financial products were launched: volatility futures on the Deutsche Terminborse and volatility swaps from Salomon Smith Barney (now part of Citigroup) [15].

The first of these instruments allowed traders to hedge against the movement in the price of options written on the DAX index due to changes in volatility [6]. The second allowed traders to gain exposure to only the volatility of underlying instruments, without labour intensive and expensive delta hedging.

The VIX is an index of implied volatility on the S&P100 index. OTC derivatives are available to trade on this index.

2. Definitions and concepts

We now consider exactly what we mean by volatility. Primarily we must distinguish between two related, but distinct, concepts: the volatility of a financial instrument, and the implied volatility of an option written on such an instrument.

2.1. Underlying volatility

The price of a financial instrument can be thought of as a random variable. In order to describe how much that price might vary over a particular time period we would look for some appropriate statistic.

A natural place to start would be to consider the average of the price movements, measured at some time frequency (every 5 minutes, every day, at year end). Since we are interested in the scale rather than the direction of changes, we would take absolute values, giving a statistic

$$\gamma = \frac{1}{N_d} \sum_{k=1}^{N_d} \left| \Delta S_k \right| \tag{2.1}$$

Where N_d is the number of time periods observed, and ΔS_k is the price change in the kth time period.

Students of finance will, however, be familiar with the fact that it is more often returns that are of interest, rather than absolute price changes. We are then lead to consider log price changes. Then, for reasons related to the differentiability of the modulus function, we could take squares of log price changes, "normalised" by subtracting the mean price change over the observation period. Finally, taking the square root of the final statistic would yield the familiar Root Mean Square (RMS) statistic:

$$\overline{\sigma} = \sqrt{\frac{1}{N} \sum \left(\ln \frac{S_k}{S_{k-1}} \right)^2}$$
(2.2)

The purpose of rehashing what will be long since forgotten material for most is to emphasise that (2.2), despite being the more familiar equation, is in some respects the more artificial. In studying the behaviour of market prices (2.1) can reveal some important aspects – see for example [5], chapter 2.4, where high frequency data is analysed.

Statistics are almost inevitably quoted as RMS or sample variance, which is useful, given that the latter is an unbiased estimator for the "population" variance.

Black and Scholes proposed the following dynamics for asset prices:

$$dS_t = \mu S_t dt + \sigma S_t dW_t \tag{2.3}$$

This equation says that the instantaneous change in the price of an asset is driven by a deterministic average component, and a random component given by a normal random variable. Calculating (2.2) for a large number of observations of price changes would give an accurate estimate for σ directly. The situation becomes more complicated when σ is non-constant however – for example if it varies randomly as in the models considered below.

2.2. Implied volatility

Above in equation (2.3) we recalled the Black-Scholes model for asset prices. The Black-Scholes price of a call option written on an asset is then given by the well known formula

$$C_{t} = S_{t}N(d_{+}) - Ke^{-r(T-t)}N(d_{-})$$
where
$$d_{\pm} = \frac{\ln\frac{S_{t}}{K} + (r \pm \sigma^{2})(T-t)}{\sigma\sqrt{T-t}}$$
(2.4)

It was Black and Scholes original belief that a historically estimated σ would be used to derive a single, objective value for an option. The market quickly proved them wrong.

Some elementary calculus will show that (2.4) is a monotonic function of σ , meaning that we can invert it. Now the price of an option is not dictated by the above equation based on some external value of σ (as Black and Scholes initially thought), but is set, as with all prices, by supply and demand. So if we invert the function, we reveal the value of σ "implied" by market prices. This is referred to as the implied volatility.

As stated above, the prices of all financial instruments are set by supply and demand. There are only a certain number of shares in a company. The price of those shares balances supply and demand. Now consider the following argument: derivatives written on the shares of a company are different – there is no limit to the number of derivative contracts written on the shares, they can be "created" at no cost in infinite amounts (assuming settlement in cash). So excess demand for option contracts should be immediately matched by market makers attempting to make a profit selling such derivatives. Assuming that the market is competitive, margins should be driven to zero, leaving implied volatility as the value at which there is zero net supply of options. This value should be the market expectation of future volatility.

The above argument, whilst attractive, is wrong. Firstly, the argument seeks to relate equation (2.3) to (2.4) through market prices. However, as discussed below, given that (2.3) is not an adequate model of the market, this is not necessarily a robust argument. Further, option markets do not have the postulated perfect elasticity of supply. Writing options requires capital, which is a scarce resource, and will constrict supply, as will many other frictional costs.

There are in fact many reasons to suppose that implied volatility would be at best a biased estimator of future volatility, including the facts that option prices will contain loadings for capital costs and possibly profit. See [24] for a more detailed list. We prefer to think of implied volatilities as "normalised" option prices. Just as we can compare the price of a 3 month 5% coupon bond with that of a five year 10% coupon bond by comparing the yield, so we can compare prices of options with different strikes, maturities, etc, by comparing implied volatilities.

2.3. Volatility as risk

Note that we have not made the common identification of volatility with risk. This identification dates back to the CAPM and beyond, and arises from the assumptions of such early finance models – essentially that investor preferences or asset returns are adequately described by two parameters.

In a world of normal returns all risk measures are equivalent – a portfolio selected using standard deviations of returns (our underlying volatility above) and one selected using Value at Risk (V@R) as optimisation parameters will be identical.

However, as discussed below, we do not live in a normal world. In the real world not all risk measures give identical results, and in particular, the standard deviation of returns is not an adequate measure of risk. V@R (or one of its "coherent" relatives) is preferred.

Volatility is important for risk control, as we will see, but as a risk <u>factor</u>, not a measure of risk. We point the interested reader to the excellent reference [5].

3. Empirical evidence of Stochastic Volatility

Having defined what we mean by volatility, we now motivate the remainder of our discussion by discussing evidence that volatility varies stochastically. We (broadly) follow [18] here in presenting some empirical observations of observed price behaviour in both the cash and options markets. We consider some economic explanations, and relate them to the topic at hand:

3.1 Fat tails

It is now generally accepted that the empirical distribution of asset returns is leptokurtic – meaning (roughly) that the fourth moment about the mean is greater than the same statistic for a normal distribution with the same variance. See for example [5] chapter 2.

This means that more extreme returns, and fewer "midrange" returns are observed, than would be expected under a Gaussian distribution.



Figure 3.1 – Empirical daily S&P log return distribution 1 June 1988 – 31 December 2004 vs. Gaussian PDF. The higher peak and fatter tails of the empirical distribution are evident.

3.2 Volatility clustering & persistence

A glance at a financial time series often immediately reveals periods of high volatility and periods of low volatility.



Figure 3.2 – S&P daily log return absolute value, 2004. Period of "high" volatility circled in red, period of low volatility in green.

In fact, fat tails and volatility clustering are two sides of the same coin. It is well known that a "mixture" of distributions, for example price changes distributed according to a normal distribution, but with a random variance, can replicate fat tails.

However, both fat tails and volatility clustering may be equally well explained by directly modelling the underlying price distribution as having fat-tails.

Another empirical "fact" is the persistence of volatility regimes – there are <u>periods</u> of high volatility, and <u>periods</u> of low volatility, not just random incidences. This observation indicates something about any proposed model of volatility. See [16] for an interesting characterisation of this behaviour, or [17] for further development.

3.3 Leverage effects

The empirical observation that volatility and share prices are negatively correlated is well known. The term "leverage effect" was first coined in [3]. The argument is that falling share prices increase the debt-to-equity ratio of firms. This leads to higher uncertainty or risk, which increases the volatility of the share price. Hence price movements and volatility are negatively correlated.

In market lore the same phenomenon is often explained by the "fear and greed" effect. When times are good and prices are rising, traders become lazy, and are happy to sit back as their P&L's increase. Fewer trades means lower volatility. When prices start dropping however, traders (and their clients) start panicking, rush to cover positions, and generally create more market activity (and hence volatility).

3.4 Information arrivals and market activity

Movements in share prices occur due to the arrival of information (this is essentially the efficient markets hypothesis). Clearly news does not arrive as a steady stream; hence the timing of price movements is random. However, the random process governing the arrival of information is such that it seems to be incompatible with simple models such as the Black-Scholes geometric Brownian motion model of share prices.

One can consider a trading day: the markets open at 08:00, and there is a flurry of activity as traders and their clients look to act on information read in the morning newspapers. After an hour or two, the market calms down a little, with traders keeping an eye on the Reuters monitor. During the mid-afternoon traders may go out for a coffee, few big trades will go through and the market will generally be quiet. Then there will be another flurry of activity just before the market closes, as people close out open positions, companies try to announce unflattering information at the last possible minute, etc.

In fact, studies show precisely these sorts of intraday "seasonal" activity effects. Similar effects are observed on longer time intervals. Further studies show that volatility and market activity are correlated. **3.5** Volatility co-movements

We observe that volatility is dependent across markets. The old adage that when New York sneezes, London catches a cold holds in the scale of price movements as well as the direction. When modelling multiple asset classes it may be important to capture this behaviour.

3.6 Implied volatilities

A study of implied volatility, as discusses above, is primarily a study of "normalised" option prices. However, implied volatilities of options on a given underlying, but with differing strikes and maturities, reveals some interesting behaviour, which can be directly linked to the observations above.

Firstly, implied volatilities vary with term to expiry – we call this the "term structure" of implied volatility.

If we accept that implied volatilities are, in part, an estimate of future volatility, then it seems reasonable that a trader might have a different estimate depending on the time horizon. The volatility of the share price of a company before it makes its next earnings report might reasonably be expected to be lower than after the report is made. Hence options maturing before the next report date would have a lower (time averaged) implied volatility than those maturing later.

We could also relate this observation to 3.2: if traders think we are currently in a low volatility environment, they would expect volatilities to increase. If they can make a rough guess (based on past experience) as to the rate of increase, they will adjust their expectations of future volatilities accordingly.

The "volatility smile" is a well known feature of option prices. Essentially the volatility smile shows that options which further into or out of the money are undervalued by the Black-Scholes formula. This shows that the market expects the options to be exercised with greater probability than indicated by the geometric Brownian motion assumption. This is a direct consequence of the fat tails displayed by the price process discussed in 3.1.



Figure 3.3 Implied volatilities of short dated options on the FTSE100 showing the implied volatility smile

An additional feature in some markets is the fact that this smile is asymmetric – the so-called "volatility skew", or sometimes "volatility smirk". This shows that deep out of the money puts are valued by the market more highly than similarly deep out of the money calls. This observation seems to be linked to 3.3.

The skew was apparently only observed after the 1987 stock market crash. Institutions who had written deep out of the money puts were in the worst position at the crash – not only had their own portfolios lost money, but they had to pay out claims to others as well. Many such players did not survive. After the crash anyone writing puts demanded a premium against the possibility of having to pay out at the worst possible time.

Economic pricing theories (see the excellent [8] for example) would place a higher relative price on assets which pay out well when all other assets are making losses. The high price the market puts on these options should tell us something about the frequency of such crashes (or of people's aversion to them)!



Figure 3.4 Implied volatilities of 2 year FTSE100 options showing the skew of implied volatilities

Items 3.1 - 3.5 are relevant to the risk control actuary, who seeks a realistic model of the markets. Item 3.6 is of interest to the valuation actuary. In placing a market consistent value on a book of insurance contracts, he is seeking to replicate the prices of market traded options, and hence must capture the smile effect.

4. Stochastic volatility models

In the previous section we discussed the reasons why we may want to model volatility as a random variable. In this section we consider first consider a simple but effective extension of the Black-Scholes model, before describing true stochastic volatility models, and some of the mathematics involved.

It is not our intention to reproduce a full derivation of all relevant results. Instead we hope to offer a "bluffer's guide" to the subject, with some qualitative discussion of the important results. Readers wanting full proofs are pointed towards the references.

4.1. Local Volatility models

Local volatility models as a concept were first suggested by Dupire in [14]. These models are not stochastic volatility models, in that they do not add any further sources of risk (or randomness). Instead they are an attempt to modify the basic Black-Scholes model to fit observed option prices.

We include them, not only as a historical note on the development of full stochastic volatility models, but also because they are still in use today on some trading floors, and may be of some use to actuaries attempting to value life contracts in a market-consistent way.

Essentially Dupire's contribution was to demonstrate that, given a set of option prices (subject to certain reasonable constraints), one can find a deterministic function b(S,t) of the underlying and time, such that the price of the underlying can be written as a diffusion-type equation:

$$dS = rSdt + b(S,t)SdW$$

and the option prices implied by this equation fit the observed prices.

To demonstrate how this may be done, we consider a simple example (in fact, this example predates Dupire's work).

Recall in 3.1.6 we discussed the term structure of implied volatilities. We note the value of the at-the-money (say) implied volatility at each time t, and write it $\overline{\sigma}(t)$. We could then construct a (deterministic) process $\sigma(t)$ such that

$$\overline{\sigma}^2(t) = \frac{1}{t} \int_0^t \sigma^2(s) ds$$

i.e. each observed implied volatility is the average of a "volatility function" – analogous to observed spot interest rates being the average of (unseen) forward rates. We then consider the process

$$dS_t = rS_t dt + \sigma(t)S_t dW_t$$

which has solution (by applying Ito's Lemma)

$$S_{t} = S_{0} \exp\left(\left(r - \frac{\overline{\sigma}^{2}(t)}{2}\right)t + \int_{0}^{t} \sigma(s)dW_{s}\right)$$

the price of a European call on S can be seen to be given by the usual Black-Scholes formula with volatility term $\sqrt{\overline{\sigma}^2(t)}$ - and thus we retrieve the implied volatilities we started with.

To see this consider that the Black-Scholes formula gives the price of an option assuming that the terminal price of the underlying is log-normally distributed. Note that

$$\int_{0}^{t} \sigma(s) dW_{s} = \lim \sum \sigma(s_{i}) \Delta W_{s}$$

and recalling that the sum of normally distributed random variables is again normal, we see that S is indeed log-normally distributed.

Dupire showed that this procedure can also be used to find consistent functions to describe σ to fit the smile and skew of observed implied volatilities. In fact, as he noted, there is little trouble obtaining this fit because the number of possible parameters is large. For example, to capture a skew effect, one could write σ as a linear (affine) function of S.

The term structure of implied volatilities for swaptions derived from some interest rate models will be the implicit product of a similar mechanism.

Note that, in general, since local volatility models do not add any randomness, the resulting distribution of asset prices will be Gaussian – so for example there will be no allowance for fat tails in the resulting model.

4.2. The general form of stochastic volatility

Continuous time financial models are written as diffusion processes using stochastic differential equations. The general form of the models we are investigating is

$$dS_{t} = \mu_{t}S_{t}dt + f(\sigma_{t}, S_{t})S_{t}dW_{t}^{S}$$
and
$$d\sigma_{t} = m(\sigma_{t}, t)dt + \eta(\sigma_{t}, t)dW_{t}^{\sigma} \qquad (4.1)$$
with
$$dW_{t}^{S}dW_{t}^{\sigma} = \rho dt$$

These equations mean that the instantaneous return on S is given by some deterministic term μ plus some random noise, the "scale" of which is given by $f(\sigma)$. σ itself follows similar (but more general) random dynamics.

Before continuing, we pause to consider some desirable qualities of a model of volatility. Drawing on section 3 and intuition, we would presume that a model of volatility should:

- Be always positive
- Revert to some mean value
- Display a term structure
- Have some form of negative relationship with price movements

The first three qualities would lead us to consider interest rate models, which share these characteristics, as an appropriate starting point for models of the volatility process.

4.3. The Heston model

The Heston model [21] is the "classic" model of stochastic volatility – the model which has perhaps come closest to matching the success of Black-Scholes. Most stochastic volatility models are benchmarked against Heston, and Bloomberg offers a Heston implementation as standard. We explain the model in some detail below, in order to examine the workings of a "typical" stochastic volatility model.

Heston assumed that the "spot" variance process σ_t^2 obeys the dynamics:

$$d\sigma_t^2 = -\lambda(\sigma_t^2 - \overline{\sigma}^2)dt + \eta\sigma_t dW_t^2$$

This is of course the process proposed by Cox, Ingersoll and Ross in [9] to model the short interest rate. The model is mean reverting.

The parameters may be interpreted as:

- $\overline{\sigma}$ the long-run mean level of volatility of the asset price.
- λ the mean-reversion rate of the process a higher value means that volatility will revert back to its long-term mean faster from a given perturbation.
- η the "vol-of-vol", i.e. the scale of changed in the volatility process itself. This parameter essentially controls the depth of the implied volatility smile, together with λ .
- σ_0 the initial level of the volatility process. This sets the level of the smile.
- ρ in the Heston model shocks to the asset price and volatility process may be correlated, as empirical evidence suggests. A negative correlation will result in an implied volatility skew.

The solution to the system of equations (4.1) will have an asset price volatility term of the form

$$\sqrt{\int_{0}^{t} \sigma_{s}^{2} ds}$$

So the volatility of price changes in separate time periods will be autocorrelated – a desirable feature of a volatility model if we are to capture the market behaviour described in 3.2.

Having chosen a process for volatility, the next task is to use the model to price options. Readers will recall that the Black-Scholes formula for the price of an option is obtained by solving the Black-Scholes partial differential equation (PDE). The PDE is derived by considering a portfolio consisting of a derivative instrument and offsetting positions in the underlying asset and a risk free asset. Some manipulation using Ito's Lemma gives the PDE:

$$\frac{\partial C}{\partial t} + rS \frac{\partial C}{\partial S} + \frac{1}{2} \sigma^2 S^2 \frac{\partial^2 C}{\partial S^2} - rC = 0$$

A similar argument can allow for other sources of risk – such as stochastic volatility – giving the Heston PDE. To avoid awkward notation we write v for σ^2

 σ_t^2 , and the resulting equation reads

$$\frac{\partial C}{\partial t} + rS \frac{\partial C}{\partial s} + \frac{1}{2}vS^2 \frac{\partial^2 C}{\partial s^2} + \left(\Psi_{-\lambda(v-\overline{\sigma}^2)}\right)\frac{\partial C}{\partial v} + \frac{1}{2}\eta^2 \frac{\partial^2 C}{\partial v^2} + \rho vS\eta \frac{\partial^2 C}{\partial v\partial s} - rC = 0$$

Heston proceeded to solve this equation, giving a (semi) closed form solution for the price of a call. Note that the first three and last terms of this equation form the Black-Scholes PDE. The remaining terms are due to the volatility risk, and a cross-risk term.

Each of the partial derivatives in the Black-Scholes equation has a familiar name – delta, gamma, etc, which we associate with a source of risk – movement in the price of the underlying, risk free rate, etc. In a stochastic volatility world there are additional sources of risk and hence additional partials derivatives. These also have names – Volga is the second partial derivative with respect to volatility, Vanna is the second partial derivative with respect to both volatility and the underlying.

The "new" term Ψ in the Heston PDE arises as a "price of volatility risk". This term is unknown. The reader will recall that a "complete" market is one in which all derivative claims can be hedged, in an "incomplete" market some risks cannot be hedged. So in a complete market, all risks are priced by the market, and the market price of volatility risk will be known, resulting in a unique solution to the PDE. In an incomplete market there will be many, possibly an infinite number, of solutions. Heston assumed that the price of volatility risk was proportional to volatility. Other researchers have assumed that the volatility risk commands no premium (an unlikely scenario). Recent research uses the price of other traded options, or of volatility swaps, to "complete" the market and derive a unique price. See [17].

Heston solved the above PDE for the price of a call. This solution has the form:

$$C_t = SP_1 - Ke^{-r(T-t)}P_2$$

where P1 and P2 are pseudo- or risk-neutral probabilities.

Readers will note the similarity with the Black-Scholes formula. In fact the price of a call will always have this form, in any model. The full formula is long, and hence we have relegated it to an appendix.

Closed form formulae for the prices of options are considered essential for the success of a model by those working in short term finance. Given that banks must mark their books to market and run risk control overnight, it is easy to see why.

4.4. Other models

Many alternative models have been proposed to Heston. The Hull-White [22] model is similar to the Heston model, with dynamics for the variance process given by a mean-reverting geometric random walk.

A recent and popular model is SABR (Stochastic Alpha Beta Rho) [20], which combines some features of stochastic volatility and local volatility models.

Fouque *et al.* propose a model in [17] where volatility is driven by an Ornstein-Uhlenbeck (mean-reverting arithmetic random walk) process, Y_t , and positivity is achieved by setting $\sigma = \exp(Y_t)$.

4.5. GARCH models

Another approach for modelling the variability of returns over time is to let the conditional variance be a function of the squares of previous observations and past variances. This leads to the autoregressive conditional heteroscedasticity (ARCH) models. ARCH processes have proved to be an extremely popular class of non-linear models for financial time series. The importance of ARCH processes in modelling financial time series is seen most clearly in models of asset pricing which involve agents, maximising expected utility over uncertain events. Analogous to Stochastic variance models being discrete approximations to continuous time option valuation models that use diffusion processes, ARCH models can also approximate a wide range of stochastic differential equations.

The ARCH(1) model can be written as

$$\varepsilon_t | x_{t-1}, x_{t-2}, \dots \sim NID(0, \sigma_t^2)$$
 (4.5*a*)

with

$$\sigma_t^2 = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 \tag{4.5b}$$

where

 $\varepsilon_t = x_t - \mu$

The ARCH(1) model can be seen as an extension to linear time series models by adding a component for the variance that varies with past values of the time series. New information has been added to the model and it is expected that the model fits the data better. The presence of ARCH can lead to serious model misspecification if it ignored: as with all forms of heteroskedasticity, analysis assuming its absence will result in inappropriate parameter standard errors, and these will be typically too small

A practical difficulty with ARCH models is that for large lags, unconstrained estimation will often lead to the violation of the non-negativity constraints that need to ensure that the conditional variance is always positive. To obtain more flexibility the generalised ARCH (GARCH) process was proposed. The GARCH(p,q) process has the conditional variance function (replacing equation 4.5b)

$$\sigma_t^2 = \alpha_0 + \sum_{i=1}^{t-p} \alpha_i \varepsilon_{t-i}^2 + \sum_{i=1}^{t-q} \beta_i \sigma_{t-i}^2$$

For positive variances all the coefficients must be positive. Because ARCH processes are thick tailed the conditions for weak stationarity are often more stringent [23]. Many extensions of the simple GARCH model have been developed in the literature.

4.6. A brief note on parameters and parameter estimation

The devil, as always, is in the details. And the detail of all financial models includes parameter estimation or calibration. Hence it is often far more difficult to obtain information about calibration techniques than to obtain specifications of models. However, discussions with practitioners reveal some information.

To a trading desk quant, the problem is to obtain a set of parameters such that his chosen model most closely replicates the market prices of the calibration instruments. A typical calibration may use three options as calibration instruments – generally the at-the-money forward, and options with strikes at 105% and 95% of this level. The model is then calibrated by minimising the square difference between the model predicted prices and the observed prices, using the Levenberg-Marquardt algorithm, or a more robust algorithm known as Differential Evolution. Minimisation is complicated somewhat by restrictions on some parameter values – in the Heston model above for example, obviously ρ must be between -1 and +1.

This methodology is inappropriate for a risk control environment, as it will give risk adjusted parameters. The real-world parameters may be estimated from historical data, or a combination of historical data and option data, with some assumption about market prices of risk allowing risk-neutral parameters to be "inverted" to give real-world ones.

Historical estimation generally proceeds iteratively – in the Heston model μ and the mean σ may be estimated first assuming no stochastic volatility by standard regression techniques. Then the full model will be assumed and the remaining parameters estimated.

4.7. An important result - The 1st theorem of stochastic volatility

We now present an important result which relates the return on a book of options (or a book of life contacts with guarantees) to the return predicted by Black-Scholes. Recall that in the Black-Scholes world a derivative position can be hedged by taking offsetting positions in the underlying and in a risk-free savings account. However, in the real world the implied volatility of the derivative may change, without any change in the value of the underlying. The "hedged" portfolio would no longer be hedged.

Now assume we have hedged our portfolio by using a Black-Scholes model, with implied volatility σ . Suppose that the actual volatility between time 0 and t instead follows a process β_t , for example it may follow one of the processes given above. We will expect our portfolio to have value zero at all times in the future. In fact, what we find is that the "hedging error" evolves as a random process with dynamics

$$Z_t = \int_{\alpha}^{t} e^{r(t-s)} \frac{1}{2} S_s^2 \Gamma_s^2 (\hat{\sigma}^2 - \beta_s^2) ds$$

where the reader will recall that Γ is the 2nd derivative of the option price with respect to the underlying.

A delta-hedged portfolio is not hedged at all, although it may take a while before you notice! This is an important point to bear in mind if considering delta-hedging a life insurance fund.

From the result we can deduce several things: Firstly, if the implied and actual volatilities are close "most of the time", the hedging error will be small. Secondly if the gamma of the option (or the overall gamma of a book of options) is small, then again, the hedging error will be small.

The result shows how it is possible for a derivatives market to exist, even in the absence of an exact model for price dynamics. It also explicitly demonstrates the effect of volatility risk, a subject we return to later.

Despite its importance, this result is not found in many text books on option pricing. [10] is a good source for proof and further discussion.

5. Stochastic Volatility models in action

In this section we implement one of the models described above, and use it to demonstrate how stochastic volatility models may be of use in actuarial work. Finally we discuss some concerns related to performance of the class of models.

The Heston model is undoubtedly the best known stochastic volatility models, and was one of the earliest. It could be described as the "classic" model, and is the one which has come closest to emulating the success of the Black-Scholes model. Indeed, Bloomberg and other information providers/ brokers offer automated Heston valuation of options to traders.

We have relegated the actual formulas for the Heston model to appendix A, for reasons which will be evident to the reader who ventures that far. Readers interested in the "guts" of the model are directed there, or to the original paper [21], or the more accessible material in [19].

5.1. The Heston model – results

Before applying the Heston model to actuarial problems, we present some graphs which demonstrate the distribution of asset returns following a Heston process.

A Heston process with vanishing volatility of volatility (vol-vol in the jargon) should reduce to a Black-Scholes style model. The parameters of the mean-reversion term in the volatility process may lead to a volatility which changes over time:



Figure 5.1 – The probability distribution function of the returns from Black Scholes model



Figure 5.2 - The probability distribution of the returns from the Heston model with volatility of volatility assumed zero.

The vol-vol term determines the fatness of the tails of the distribution (the kurtosis):



Figure 5.3 - The probability distribution of the returns from the Heston model with no skewness

Finally, the correlation term determines the skewness of the distribution:



Figure 5.4- The probability distribution of the returns from the Heston model

For comparison we show the daily empirical return distribution, derived from the total return index on the S&P500 over the period 01-June 1988 to 31-December 2004:



Figure 5.5. The empirical distribution of the data

All distributions were generated by simulation (the S&P returns were observed), and normalised to have zero mean and unit variance. We have made no attempt to optimise the PDF shown in Figure 5.4 to match that in Figure 5.5. See [13] for a study where precisely this was done – they find that the empirical return on the S&P500 is indeed consistent with a correctly parameterised Heston model.

5.2. Valuation

We consider a simplified model of a life insurance fund. This fund contains a pool of assets. The liabilities of the fund are the assets deemed to be owned by the policy-holders (asset shares or value of unit funds), and some promise to pay a minimum guaranteed level of benefits at certain dates in the future. We assume all policies have a single maturity date, so as to avoid the influence of factors which are not of interest to us in this demonstration.

Modern actuarial orthodoxy (now enshrined in the UK by new FSA regulations) holds that this fund should be viewed as a long pool of funds, short the assets shares, and short a put option on those asset shares. Put-call parity tells us that the liability side is equivalent to a short bond (face value to guarantee) and a call on some proportion of the value of the assets.

Now in most funds there will be a mix of business written at different times in the past, and subjected to different levels of declared bonus. Hence there will be a range of different levels of guarantee in the fund. In the parlance of traded options, we would say that there are a range of strikes.

As discussed above, a portfolio of call options with a range of strikes will show a range of different implied volatilities. A Black-Scholes implementation, either analytic or Monte-Carlo, used to value the liabilities of the fund will not give market consistent answers, in that the values obtained will not be consistent with the market prices of traded options.

However, we could instead use the Heston model. The Heston model has a (semi-) closed formula for the prices of plain-vanilla options, which we use here. The results are shown below.

Note that we can replicate these values by assuming that the assets of the fund follow a Heston process – something we cannot do with a Black-Scholes model where we have changed the volatility for each strike level.

5.2.1. LifeCo ltd

LifeCo is our mock life insurance company. We will consider one of the company's funds, holding five blocks of business, divided by levels of guarantee, in turn based on past declared bonuses. The fund is assumed to be invested 100% in the FTSE100, and currently has £500m in assets. All business under consideration will mature in 2 years time.

We assume that asset shares and guarantees have been aggregated in a meaningful way:

Business	Aggregate Asset	Aggregate	Guarantee Present Value
1			
1	£90,000,000	£86,968,848	£80,057,008
2	£90,000,000	£92,734,738	£85,364,655
3	£90,000,000	£96,578,665	£88,903,086
4	£90,000,000	£102,344,556	£94,210,733
5	£90,000,000	£108,110,446	£99,518,380

Table 5.1 – LifeCo Asset shares and guarantees by business block

Guaranteed amounts have been discounted at a risk free rate of 4.615%. At maturity in 2 years time the fund will pay out the greater of asset shares and the guaranteed level. The liabilities of the fund are therefore the guaranteed amount, plus a call option on the asset shares, struck at the level of the guarantee.

We assume no further bonuses will be paid. Our task now is to value the call options.

5.2.2. Black-Scholes valuation

The assets of the fund are invested in the FTSE100, so we must look to options written on the FTSE100 to determine what constitutes a "market consistent" valuation.

We have used data on Euronext FTSE100 options, dated 28/02/2005. We will assume this is our valuation date.



Below we show the implied volatilities for 2 year options.

Figure 5.6 Implied volatilities of 2 year FTSE100 options

Note that the options show a distinct skew, but no smile, i.e., the surface does not turn up at higher strikes. This is a common feature of long dated options (in banking parlance, 2 years is a long time. Only options with less than 9 months to maturity trade in a truly deep and liquid market).

We now have several options. We can value each of our 5 blocks of business separately, using different implied volatilities for each as appropriate. This will give a market consistent valuation. However, frequently in a real-world insurance implementation, we would be using a Monte Carlo model of the whole fund. In this case this option would not be available to us.

We could instead use a single implied volatility value. The at-the-money-forward implied volatility (the value for options whose strike is equal to the underlying forward price) is given as 11.34%.

Alternatively we could use an average implied volatility, possibly weighted by asset share or level of guarantee. In a Monte Carlo implementation using the Black-Scholes (or similar) model this is certainly what we would have to do. Auditors and other interested parties are likely to question this method. Primarily they will want to know if a valuation method which is market consistent on average, but not at a detailed level, can be said to be market consistent at all.

Business block	Strike volatility	ATM volatility – 11.34%	Asset share weighted volatility – 11.5%
1	£12,339,077	£11,558,625	£11,627,379
2	£8,341,349	£7,926,537	£8,016,234
3	£6,054,191	£5,962,919	£6,059,798
4	£3,325,006	£3,702,605	£3,798,270
5	£1,576,051	£2,170,670	£2,253,473
Total	£31,635,674	£31,321,356	£31,755,154
Error		0.4%	-1.0%

In the table below we show the Black-Scholes value of call options under each of the possible methods.

Table 5.2 Value of call component of liabilities using Black-Scholes and several possible choices for volatility parameters

The strike volatility values in Table 5.2 are the true "market consistent" values in that they are consistent with the values of traded options with the same characteristics. The alternative methods will produce answers which are "wrong" -0.4% too large and 1% too small respectively. The business block level errors are in many cases greater.

5.2.3. Heston valuation

We now consider valuing the same block of business using a Heston model instead. Again, given the complicated nature of real insurance business, in practise this would probably involve a Monte Carlo model of the fund. We will simply use the analytical formula.

We calibrated the Heston model to the prices of options with strikes at (approximately) 95%, 100% and 105% of forward, as we believe is customary on derivatives desks. Our calibration involved minimising the square distance of market from model prices, and used the Levenberg-Marquardt algorithm. The resulting parameters are:

σ_0	0.129296
mean σ	0.123964
Λ	0.756694
Р	-0.63314
Н	0.227985

Table 5.3 Heston parameters

Below we show the values place on the business by the Heston model. We have shown the "true" market consistent values as before.

Business	Strike	Heston
block	volatility	value
1	£12,339,077	£12,300,595
2	£8,341,349	£8,341,322
3	£6,054,191	£6,054,185
4	£3,325,006	£3,324,997
5	£1,576,051	£1,570,263
Total	£31,635,674	£31,591,361
Error		-0.1%

Table 5.4 Heston valuation results

As can be seen, the Heston model provides a good fit to market values across a range of strikes.

Note that for LifeCo we had the freedom to choose our levels of guarantee. In fact, we selected them such that we could mark the business blocks to market using our data. Of course, we then calibrated to that data, so it is unsurprising that we obtain a close fit. Note however, that blocks 1 and 5 were not calibrated to, and yet a close fit is still obtained. When using such a model in the real world it is unlikely that the guarantees in your book of life contracts will align so well with the strikes of market traded options! However, despite this, the ability of the Heston model to fit a range of strikes is an advantage over the alternatives of using different implied volatilities for each level of guarantee, or using an implied volatility which is correct in aggregate, but incorrect at each level of guarantee.

5.3. Risk Control

We continue to consider the same fund in LifeCo, but we now consider risk control. To be precise, we consider how the realistic valuation (i.e. the market-consistent valuation) of the fund might vary over a single year.

We use Value at Risk (V@R) methodology to compare the level of risk calculated assuming the assets in the fund follow:

- a) A Black-Scholes process
- b) A Heston process

The fund is solvent at time zero, with £500million in assets.

5.3.1. Simulations

We generated 10,000 simulations from a Black-Scholes process, assuming a flat risk-free rate of 4.615%, and a 2% equity risk premium. For volatility we used the ATM forward volatility of 11.34%.

For the Heston process we use the same parameters as above. However, we must move from the estimated pricing or risk-neutral parameters to real world parameters. This can be done in an analogous fashion to the way it is done in the Black-Scholes model – by adding back a risk premium.

First we increase the asset price drift parameter, as in the Black-Scholes model. We use a 2% risk premium again to be consistent.

We now need to include a premium for volatility risk. As investors are risk-averse and dislike volatility, when pricing they assume that the long-run average volatility will be higher than it is, and that it will return to this value at a faster rate. As discussed briefly above, we could attempt to derive an exact value for the adjustment by considering the prices of volatility related assets (e.g. delta-hedged options) or by considering the average return on options as compared to their prices. However, we use an ad-hoc adjustment of $\kappa = -0.04$, applied in the following formulae:

$$\lambda = \lambda^{RN} - \kappa$$
$$\overline{\sigma}^2 = \frac{\lambda_{RN} \overline{\sigma}_{RN}^2}{\lambda}$$

Where the subscript "RN" denotes risk-neutral values and variables with no subscript are real-world parameters.

We considered the realistic balance sheet at the end of one year, with liabilities valued using a Black-Scholes volatility of 11.34%, and the same Heston parameters as in section 5.2.3. We calculated the 0.5^{th} percentile of realistic net assets.

5.3.2. Results

The results are as follows:

Model	V@R _{0.005}
Black-Scholes	£76,967,075
Heston	£159,227,058

Table 5.5: A comparison of the outputs from the Black-Scholes model and the Heston Model

Clearly the greater value given by the Heston model indicates that the Black-Scholes model is severely understating the risk of the fund.

A priori one would expect this, the short comings of Black-Scholes as a risk measurement tool are well known. We would expect a greater level of risk from the Heston model to arise from two factors: the fat-tailed and skewed asset returns affecting the asset side of the balance sheet, and the shift in volatility affecting the liability side of the balance sheet. However, we present separate asset and liability results below:

Model	Assets 0.5 th percentile	Option liabilities 99.5 th percentile
Black-Scholes	£392,427,162	£179,777,263
Heston	£310,322,042	£146,895,098

Table 5.6: A comparison of the results from the Black-Scholes model and the Heston Model

As can be seen, the Heston model shows a significant asset-side risk. However, on the liability side, the risk is greater from the Black-Scholes model. This seems to be because the options are most valuable when asset prices have risen at the same time as spot volatilities. However, our calibration of the model returned a large negative correlation in volatilities and asset prices, so the "perfect storm" scenario seems to happen with negligible probability.

This issue illustrates a problem with the Heston model – the parameters which replicate option prices do not seem to generate realistic projections. In particular, some alternative mechanism to generate an appropriate match to the observed skew of implied volatilities is suggested.

An alternative parameterisation may be used which will result in simulations which more closely match historical observation. [13] perform such an exercise, with good results. However, this type of calibration is unlikely to give a good match to market prices at time zero. This is a situation which will be familiar to actuaries who have been involved in similar risk measurement projects.

5.4. Performance

Recall from our earlier introduction that many models for volatility are derived from previously existing models for interest rates. Those who are familiar with such models will be aware of the drawbacks of those models. Stochastic volatility models may fail to accurately describe price movements in other ways as well.

We summarise some of the common failings below:

- Negative values. Processes following an Ornstein-Uhlenbeck dynamic can go negative an undesirable situation for both volatility and interest rates.
- Values near zero. Many processes are constrained by their form to be positive – both the Hull-White and Heston models fit into this category. However, generally this is achieved by reducing the volatility of the process as it reaches zero – with the result that the process eventually becomes stuck near zero.

Models which achieve positivity through other mechanisms generally pay by losing something in tractability.

On the trading floor the advantages of these models are thought to outweigh these disadvantages – especially as most trading floor pricing is done by analytical techniques, not Monte Carlo methods, and the time horizon is short. However, in the actuarial world the opposite is true – time horizons are long and Monte Carlo techniques are prevalent. Having a model which can plausibly project asset prices (and volatilities) over many years is a must. We believe that none of the models discussed in this paper can fulfil these criteria.

Actuarial models built for long term projection of interest rates (see for example, [6]), may be pressed into service to project volatility instead.

• However, all is not well on the short end of option pricing either. The parameters required to reproduce observed implied volatility smiles are clearly unreasonable – correlations of close to -1, and vol-vol of 70% are common for the Heston model for example. Clearly observed volatilities do not follow such and extremely calibrated Heston dynamic. Instead this points to an alternative dynamic for prices – short dated options are mostly priced on jump risk – that is, the possibility of a large jump in asset prices shortly before the option matures. There is large amount of resources currently being put into implementing jump models on trading floors. In this respect actuarial techniques have run slightly ahead of market techniques – we are aware of at least one model used for actuarial work which is based on exactly this type of process.

5.5. The State of the Art – Unified Volatility Models

As described above, stochastic volatility models do not seem to provide a complete solution to option pricing and asset value projection problems. In section 4.4 we briefly mentioned SABR, a recent model which combines local volatility models (as discussed in section 4.1) with stochastic volatility. A lot of current research focuses on combining local volatility, stochastic volatility, and jump models. [1] is a recent paper describing one promising method of doing this.

The danger with building such more advanced models is that as the models become more complex, the number of parameters increases. One may end up with a situation where the number of degrees of freedom to calibrate a model is so great, that one simply ends up wrapping a model around the data. In this situation one may have a model which closely describes, for example, the implied volatility surface on a particular day, but which is unable to project a situation very far from the one to which it is calibrated – an undesirable situation in a risk-control project.

6. Conclusion

In this paper we considered evidence from asset price processes and option prices which indicates that a constant volatility parameter, as suggested by Black-Scholes and other models, is not an appropriate modelling choice. In particular, the empirical distribution of asset price changes shows fat tails, and the Black-Scholes implied volatility shows a term and strike dependent structure. This is of interest to actuaries involved in the valuation and risk control functions of life companies.

We considered a class of models which may describe these factors – stochastic volatility models. We implemented the Heston model, and put it to use to value life insurance liabilities and to estimate value at risk in a life insurance fund. We find that the model may be useful – both in calculating market consistent liabilities across a range of guarantees, and in revealing an extra element of risk.

However, we find that the model is not without problems; in particular the parameters obtained in a calibration exercise seem unrealistic. Models combining jumps and stochastic volatility may be the way forward.

A. Appendix – Formula for a Call option in the Heston model

The price of a plain vanilla call option on a share following the dynamics of the Heston model is given by the following formula:

$$C_t = SP_0 - Ke^{-r\tau}P_1$$

We have written τ for time to maturity, and below x for the log forward-strike ratio. The pseudo probabilities are:

$$P_{j} = \frac{1}{2} + \frac{1}{\pi} \int_{0}^{\infty} \Re\left(\frac{\exp(C_{j}(k,\tau)\overline{\sigma}^{2} + D_{j}(k,\tau)\sigma_{0}^{2} + ikx)}{ik}\right) dk$$

Re() indicates the real part of the argument. The imaginary functions arise as a result of solving the PDE – done using Fourier transforms. The functions C and D are given by

$$C(k,\tau) = r_{-} \frac{1 - e^{-d\tau}}{1 - g e^{-d\tau}}$$
$$D(k,\tau) = \lambda \left(\tau r_{-} - \frac{2}{\eta^{2}} \ln\left(\frac{1 - g e^{-d\tau}}{1 - g}\right)\right)$$

With the further definitions

$$g = \frac{r_{-}}{r_{+}}$$

$$r_{\pm} = \frac{\beta \pm \sqrt{\beta^{2} - 4\alpha\gamma}}{2\gamma} = \frac{\beta \pm d}{\eta^{2}}$$

$$\alpha = -\frac{k^{2}}{2} - \frac{ik}{2} + ijk$$

$$\beta = \lambda - \rho\eta j - \rho\eta ik$$

$$\gamma = \frac{\eta^{2}}{2}$$

And that's it! The fact that the formula involves an integral means that it is not truly a "closed-form" solution, but numerical integration is much faster than Monte Carlo methods.

About the Authors

Phil Joubert (FFA)

Phil Joubert is a manager at Deloitte Actuarial and Insurance Solutions. He has worked extensively with life insurers on risk based capital, market consistent valuation, and other projects.

He holds a BSc (Hons) in Pure Mathematics from the University of the Witwatersrand, South Africa, and is currently on sabbatical to pursue an MSc in Mathematics and Finance at Imperial College. He qualified as an FFA in 2003.

pj104@imperial.ac.uk

Deloitte A&IS, Stonecutter court, 1 Stonecutter st EC4A 4TR

Coomaren P. Vencatasawmy

Coomaren Vencatasawmy is a Risk Financing Analyst at Jardine Lloyd Thompson. He has worked in research and development and as a statistical consultant to insurance companies. He provides structured finance solutions for insurance companies.

He has a PhD in Statistics from the University of Sheffield and an MBA from Edinburgh Business school.

Coomaren_Vencatasawmy@JLTGROUP.COM

Bibliography

- [1] Albanese, C., Kuznetsov, A. "Unifying the three volatility models", 2003
- [2] Bachelier, L. "Theorie de la Speculation", 1900
- [3] Black, F. "Studies in Stock Price Volatility Changes", Proceedings of the 1976 Business Meeting of the Business and Economic Statistics Section, American Statistical Association. 1976
- [4] Black, F., Scholes, M. "*The pricing of options and corporate liabilities*." Journal of Political Economy, 7: 637-54, 1973.
- [5] Bouchaud, J-P & Potters, M. "Theory of Financial Risks", Cambridge University Press, 2000
- [6] Cairns, A.J.G., and Garcia Rosas, S.A. "A Family Of Term-structure Models with Stochastic Volatility", 2003
- [7] Chicago Board Of Exchange "History of CBOE" http://www.cboe.com/AboutCBOE/History.aspx
- [8] Cochrane, J. "Asset Pricing", Princeton University Press, 2005
- [9] Cox, J., Ingersoll, J., Ross, S., "A theory of the term structure of interest rates.", Econometrica, 53: 385-408, 1985.
- [10] Davis, M.H.A. "Stochastic Volatility: The Hedger's Perspective", 2001
- [11] Davis, M.H.A., "Complete-Market models of Stochastic Volatility", 2005
- [12] Dodd, Randall "Derivatives Markets: Sources of Vulnerability in U.S. Financial Markets", Financial Policy Forum, 2005
 (Special Policy Report 8 of the Financial Policy Forum, forthcoming in a volume of policy studies from the Political Economy Research Institute entitled Financialization and the World Economy and published by Edward Elgar Publishers, also forthcoming in The ICFAI Journal of Derivatives Markets.)
- [13] Dragulescu, A., Yakovenko, D., "Probability distribution of returns in the Heston model with stochastic volatility", Quantitative Finance, Vol 2 (443-453), 2002
- [14] Dupire, B. "Pricing and Hedging with Smiles", 1993
- [15] "Pre packaged Volatility Plays" Derivatives Strategy, April 1998
- [16] Exley, J; Mehta, S; Smith, A, "Mean Reversion" FIB Conference 2004
- [17] Fouque, J.P., Papanicolaou, P., Sircar, K. "Derivatives in Financial Markets with Stochastic Volatility", Cambridge University Press, 2000
- [18] Ghysels, E, Harvey, A & Renault, E "Stochastic Volatility" Serie Scientifique, 1995
- [19] Gatheral, J. "Lecture 1: Stochastic Volatility and Local Volatility", Case Studies in Financial Modelling Course Notes, 2002
- [20] Hagan, P., Kumar, D., Lesniewski, A., Woodward, D. "Managing Smile Risk", Wilmott Magazine, 2002
- [21] Heston, S. "A closed form solution for Options with Stochastic Volatility with applications to Bond and Currency Options", The Review of Financial Studies, 1993
- [22] Hull, J.; White, . "The pricing of options on assets with stochastic

volatilities." J. Finance, 42(2):281{300, 1987.

- [23] Mills, T.C. 1999. "The Econometric modelling of financial time series"
- [24] Sheldon, T.J., Smith, A.D., "Market Consistent Valuation of Life Assurance Business", Institute of Actuaries, 2004