Lecture 10: Stochastic Volatility Models

10.1 Introduction to Stochastic Volatility

10.1.1 Approaches to Stochastic Volatility Modeling

The local volatility model we covered in some detail was in fact a special case of a stochastic volatility model; the local volatility $\sigma(S, t)$ varies with stock price *S*, and *S* itself is stochastic. In local volatility models, therefore, volatility is 100% correlated with the stochastic stock price. Obviously, of course, volatility is stochastic. In index markets implied volatility tends to be correlated with the stock price or index level, but the correlation coefficient is not 100% except during large market crashes. Therefore our next step is to investigate models where volatility can vary independently of stock price.

There are several approaches to modeling stochastic volatility.

1. The most obvious is to allow the instantaneous stock volatility σ itself to be a stochastic variable, as it clearly is. These models involves two stochastic factors, the stock price and the volatility.

   Within this class there are two subclasses: (i) models in which σ is independent of S and therefore have no skew unless volatility is stochastic, and so obtain the skew via the correlation between stock price and volatility; and (ii) models (like local volatility models) that have $\sigma = \sigma(S)$ and so begin with a skew even when volatility is deterministic, and then add volatility to that skew.
2. A second approach is to allow the Black-Scholes implied volatilities $\Sigma(K, t)$ of options themselves to be stochastic. These are the analog of market models in interest-rate modeling. There are then strong constraints on the evolution of the B-S implied volatilities in order to avoid arbitrage.
3. A third approach is that of stochastic implied tree models that begin with a local volatility model, which already projects the future no-arbitrage implied tree from a snapshot of current market options prices, and then allow these trees themselves to vary stochastically. Here again there are strong no-arbitrage conditions on the evolution.

Modeling stochastic volatility is therefore going to be potentially much more complex than modeling local volatility. In this section I want to develop the models of 1(i) and 1(ii) above, and study the character of the solutions and the smile that results from them. Some good references are

- Wilmott, “Derivatives” (Several chapters on stochastic volatility).
- Chapter 2 of Fouque, Papanicolaou and Sircar, “Derivatives in Financial Markets with Stochastic Volatility,” Cambridge University Press.
- Alan Lewis “Option Valuation Under Stochastic Volatility.” Financial Press.

- Hull and White: Journal of Finance XLII, No 2, June 87, pp 281-300.
- Gatheral: The Implied Volatility Surface

Wilmott is perhaps the easiest place to start. Gatheral's compact book has lots of details on the analytic solutions to these models and their properties.

10.1.2 The Stochastic Differential Equation for Stochastic Volatility Models

Let's start by describing how the volatility of a stock or index itself might move, beginning from GBM for the stock price and assuming that all of the smile is due to the stochastic nature of the volatility in the GBM.

$$\frac{dS_t}{S_t} = \mu_t dt + \sigma_t dZ_t$$
$$d\sigma = \text{several possibilities discussed below}$$

We will often write $V = \sigma^2$ for the variance.

The Hull-White stochastic volatility model is one of the simplest and earliest models. It describes the stochastic evolution of the variance by geometric Brownian motion:

$$\frac{dV_t}{V_t} = \alpha_t dt + \xi dW_t \qquad \text{Hull-White} \qquad \text{Eq.10.1}$$

The parameter ξ is the volatility of volatility; typical fluctuations of volatility can be very large.

Volatility is a parameter of markets, like interest rates or credit spreads. Interest rates determine bond prices through the discounting formula; variances determine options prices through the Black-Scholes model. Both parameters are derived from prices rather than being prices themselves, and, like most market parameters, tends to be range-bound or mean-reverting, rising in periods of excitement which usually correlate with falling stock prices, and falling as the world calms down, then returning to its long-term mean. Realized and implied volatility are likely both mean-reverting variables, correlated with moves the underlying index or stock price.

10.1.3 Mean Reversion and its Qualities

Ornstein-Uhlenbeck models are the classic way to describe mean-reverting stochastic variables. The stochastic differential equation for a mean reverting variable Y is

$$dY = \alpha(m - Y)dt + \beta dW \qquad \text{Ornstein Uhlenbeck} \qquad \text{Eq.10.2}$$

Why is this descriptive of mean reversion for the variable Y?

First let's solve it for $\beta = 0$ with zero volatility and no stochastic variability. Equation 10.2 then becomes

$$dY = \alpha(m - Y)dt$$

which has the deterministic solution $Y(t) = m + (Y_0 - m)e^{-\alpha t}$

As *t* gets very large, the effect of any initial position $Y_0 \equiv Y(0)$ away from the long-term mean *m* becomes irrelevant, and *Y* moves towards the long-term mean with a characteristic time $1/\alpha$. Alan Lewis in his book on stochastic volatility that estimates of half-life of volatility is between a few weeks and more than a year. This equation may not describe volatility completely accurately, since volatility tends to jump up when markets crash, and then stay high for a long time. There is a stickiness or persistence to high and low volatilities which is not quite described by Ornstein Uhlenbeck. For proprietary options traders or hedge funds that trade volatility as an asset, understanding the dynamics of volatility is very valuable.

For non-zero volatility in the Ornstein-Uhlenbeck process one can show that the solution is

$$Y(t) = m + (Y_0 - m)e^{-\alpha t} + \beta\int_0^t e^{-\alpha(t-s)}dW_s \qquad \text{Eq.10.3}$$

That is, the sum of previous random increments to *Y* damp out exponentially with time. The contribution of a random previous move to the long-term value of Y(t) damps out exponentially.

This solution satisfies the stochastic differential equation, as shown below.

$$\begin{aligned} dY(t) &= -\alpha(Y_0 - m)e^{-\alpha t} + \beta dW_t - \beta\alpha\int_0^t e^{-\alpha(t-s)}dW_s \\ &= -\alpha\left[Y(t) - m - \cancel{\beta\int_0^t e^{-\alpha(t-s)}dW_s}\right] + \beta dW_t - \cancel{\beta\alpha\int_0^t e^{-\alpha(t-s)}dW_s} \\ &= \alpha[m - Y(t)] + \beta dW_t \end{aligned}$$

Let's look at the behavior of the mean $\overline{Y(t)}$ of Y(t) at time *t*, averaged over all increments dW_s. From Equation 10.3, since each Brownian increment has a mean of zero, we see that

$$\overline{Y(t)} = m + (Y_0 - m)e^{-\alpha t}$$

so that the average displacement at time *t* is just the deterministic one.

We can also calculate the variance of the displacements at time *t* by making use of the fact that the dW_s are independent increments and so

$$dW_s dW_u = du ds \delta(u - s)$$

Therefore,

$$\begin{aligned}
[Y(t) - \overline{Y(t)}]^2 &= \beta^2 \int_0^t \int_0^t e^{-\alpha(t-s)} e^{-\alpha(t-u)} dW_s dW_u \\
&= \beta^2 \int_0^t \int_0^t e^{-2\alpha t} e^{\alpha(s+u)} du ds \delta(u - s) \\
&= \beta^2 \int_0^t ds e^{-2\alpha t} e^{2\alpha s} \\
&= \frac{\beta^2}{2\alpha}(1 - e^{-2\alpha t})
\end{aligned}$$

For small times *t*, a Taylor series expansion shows that the variance behaves like $\beta^2 t$, which is like standard Brownian motion and grows with time. But variance of this process is asymptotically finite and equal to $\frac{\beta^2}{2\alpha}$ as $t \to \infty$., which is the effect of mean reversion. As the mean reversion coefficient α gets larger, the variance gets smaller. Here is a rough sketch of the distribution of the process over time.

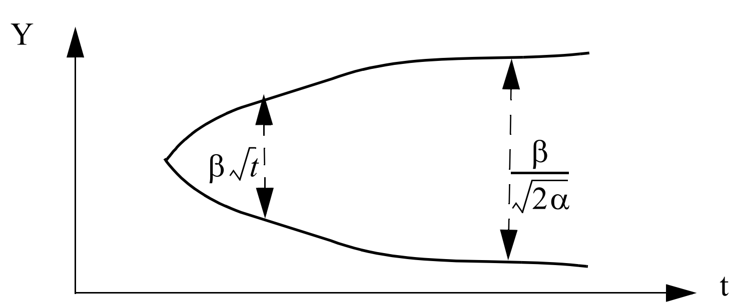

At time $t \approx \frac{1}{2\alpha}$, the characteristic time, the variance grows no larger. In contrast, for regular Brownian motion, the linear dependence of the variance on t continues for all time.

10.1.4 Some stochastic volatility models

Most of these models assume traditional geometric Brownian motion for the stock price, so that in the absence of stochastic volatility, there is no smile:

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

They then supplement this evolution by making the volatility σ (or the variance) stochastic too.

The simplest mean-reverting stochastic volatility one can write is

$$d\sigma = \alpha(m - \sigma)dt + \beta dW$$

The trouble with this is that the volatility σ can become negative, since it's similar to arithmetic Brownian motion. We'd prefer to avoid that, even if the probability of the volatility becoming negative is small.

Models of the type

$$dV = \alpha(m - V)dt + \beta V dW$$

have mean-reverting variance V whose arithmetic volatility βV vanishes as the variance becomes zero. Therefore, the variance can never become negative.

A very popular model, because it's analytically soluble, is that of Heston, which can be written in the form

$$dV = \alpha(m - V)dt + \xi\sqrt{V}dW \qquad \text{Eq.10.4}$$

This square-root factor in the stochastic behavior of the variance originates in an earlier model of interest rates by Cox, Ingersoll and Ross. The stochastic variance in the Heston model is mean-reverting and the volatility of variance vanishes as volatility gets small, so that, again, variance is always non-negative. Analytic solutions and their derivation are available in Heston's original paper, as well as in the books of Lewis and Gatheral.

All of these models have two stochastic variables, S and σ, and allow a correlation between them so that

$$d\dot{Z}dW = \rho dt \quad \text{Eq.10.5}$$

This is quite different from Black-Scholes. In all of these models the option value is not given by the Black-Scholes solution or value. because of the presence of another stochastic variable. Even standard options in these models prices depend on the volatility evolution process, and the models must be calibrated to observed options prices.

Obviously there is something more realistic about a model with stochastic volatility. With only a few parameters -- volatility of volatility and its correlation, you get a rich structure which, while it may not exactly fit the implied volatility surface, gives a generally sensible dynamics.

On the other hand, volatility evolution is even less well understood than stock price movement, and these models are unlikely to describe volatility accurately. Furthermore, one could argue that correlation is at least as stochastic as volatility itself, so that choosing a constant correlation to describe a stochastic volatility is itself inconsistent and inadequate.

10.1.5 Risk-neutral valuation and stochastic volatility models

The principle behind arbitrage-free options valuation is that you must be able to hedge away all the risk of the option instantaneously. You can do that if there are enough securities to span all the possible states of the world dynamically. If so, the riskless hedged portfolio must earn the risk-free rate in order to avoid violations of arbitrage.

Let's look at this for the Black-Scholes model, first. In the binomial model, when on the stock price S is stochastic, you can use the stock and the riskless bond to create all the state-dependent Arrow-Debreu securities, Π_u and Π_d, that you need to span the space of payoff states. Two securities can span two final states. Hence you can value any instrument with arbitrary payoffs one period in the future; in particular you can value an option. As a result, the actual expected return of the stock price is irrelevant, because one can replicate the payoff of the option in all future states.

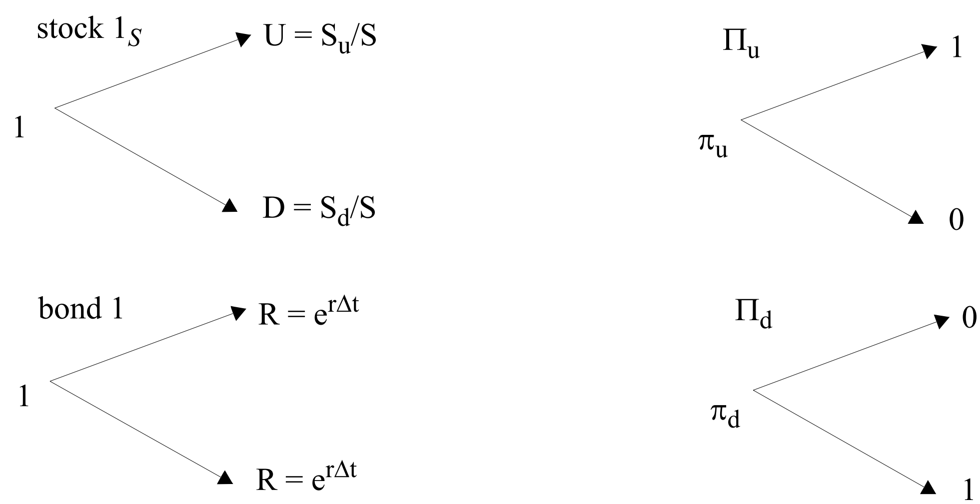

Now let's consider the case where volatility (and hence variance) itself is stochastic: S and σ (or V) can vary. This evolution is shown schematically in the figure below, where the difference between σ_u and σ_d is related to the volatility of volatility, and there is a correlation between S and σ

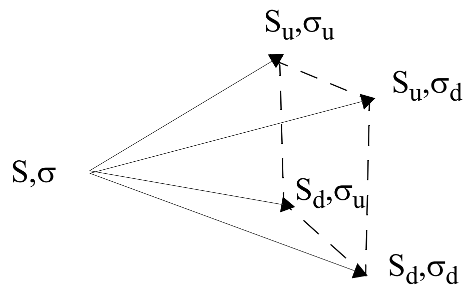

Now there are four possible final states one period later in the state space; to value an option in an arbitrage-free or preference-free way, we need to create four Arrow-Debreu securities, whose prices we know today, that respectively pay $1 in each of the four final states, and zero in the other three. But we still know only the current price of the stock S and the riskless bond B, which allow us determine only two of the final state prices. We would need to know the *price of the volatility* σ today in order to span the other states. We need to know the risk-neutral drift of volatility and the volatility of volatility and its correlation with S. If we knew the price of a $1 payoff in each final state, we could value and hedge options in the presence of stochastic volatility.

But, unlike the stock price S, volatility is not a security or a traded variable,. You cannot simply go out and buy volatility; what you can buy are securities, like standard options, that *depend* on the value of volatility and its future behavior. You have to use options (or perhaps other instruments like variance swaps) to span the space. Or, to put it another way, if you want to create a riskless hedge for an option on a stock whose price and whose volatility is stochastic, you must trades shares of the stock to hedge the stock-price variability *and* you must trade another option whose price is also sensitive to volatility to hedge the volatility variability.

For those of you who have studied the Vasiçek interest-rate model in courses on term structure, this is rather similar. Though you can hedge the stock exposure of an option with stock, you cannot hedge the interest-rate exposure of a bond with "interest rates", because interest rates don't trade. Instead, you must hedge the interest-rate sensitivity of one bond with another bond in order to create a riskless hedge, You cannot actually trade interest rates themselves but only things that depend on them.

If we hedge options only with shares of stock, then we don't have a complete market, and perfect replication of the option payoff is impossible. Since one cannot create a riskless portfolio, the principle of no riskless arbitrage cannot lead to a unique value for the option. Your utility or tolerance for risk will affect the option value, an unpleasant truth.

If you can also use other options as well as shares of stock to hedge your stock option, and if you know (i.e. if you *assume* that you know) the stochastic process for option prices as well as stock prices, then you can hedge one option's exposure to volatility with another option, you can derive an arbitrage-free formula for options values. But how well do we actually know the stochastic process for options prices?

Nevertheless, we now proceed to analyze such models. Stochastic volatility models produce characteristic smiles which we'll derive and discuss. The solutions to stochastic volatility models can often be written as averages over a distribution of Black-Scholes prices with a range of volatilities, which makes analysis easier and more intuitive.

10.2 A Preliminary Look at the Properties of Stochastic Volatility Models

In this section we use the Black-Scholes formula to understand the qualitative behavior of stochastic volatility models.

Let's begin by considering a long position in a call option in a Black-Scholes world with constant volatility σ, with value $C_{BS}(S, t, K, T, r, \sigma)$. Then, writing $\tau = T - t$

$$\frac{\partial C}{\partial \sigma} = \frac{Se^{-\frac{1}{2}\left(\frac{\ln S/K}{\sigma\sqrt{\tau}} + \frac{\sigma\sqrt{\tau}}{2}\right)^2}\sqrt{\tau}}{\sqrt{2\pi}}$$

is always positive, and

$$\frac{\partial^2 C}{\partial \sigma^2} = \frac{S\sqrt{\tau}N'(d_1)}{\sigma}\left(\frac{(\ln S/K)^2}{\sigma^2\tau} - \frac{\sigma^2\tau}{4}\right)$$

which is positive except when S/K is close to unity, and is plotted below for some typical parameters.

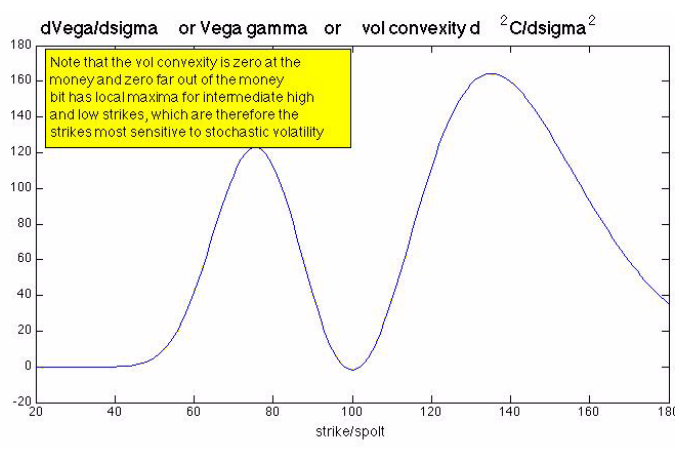

Note the predominantly positive convexity, with peaks on either side. The means that, within the Black-Scholes framework, a long position in a call is also a long position in volatility of volatility, because of the convexity w.r.t. volatility itself, just as we once saw that a hedged option is long volatility. This means that, if volatility is volatile, then the convexity in volatility adds value to the option away from at-the-money, and adds value to out-of-the-money options relative to at-the-money options in a U-shaped smile.

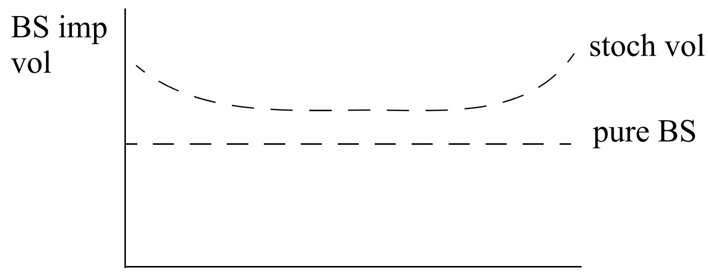

Similarly, one can plot the Black-Scholes $\frac{\partial^2 C}{\partial\sigma\partial S}$, the rate of change of vega with spot S:

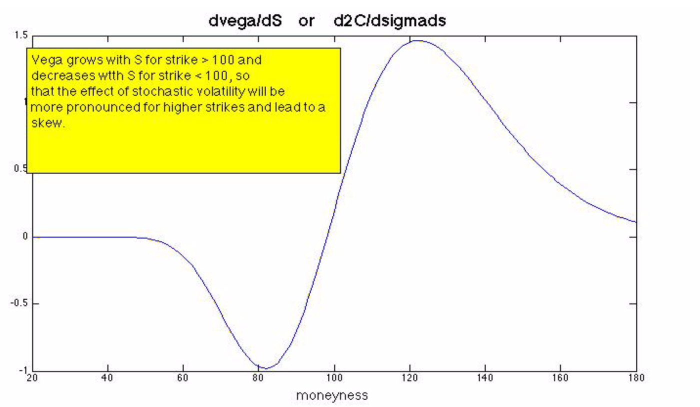

The volatility sensitivity is asymmetric about moneyness, and thus when volatility is volatile and correlated with the stock price, adds asymmetric value to the options and hence produces an asymmetric tilted smile, i.e. a skew.

These are crude but usefully intuitive ways to understand the effect of stochastic volatility on the smile.

10.3 From Local Volatility to Stochastic Volatility

I want to go step by step towards stochastic volatility. Since we already understand local volatility models and the smile they produce, let's see the effect of adding a stochastic element to a local volatility model.

Consider the parametric model

$$\begin{aligned} \frac{dS}{S} &= \alpha S^{\beta-1} dW \\ d\alpha &= \xi\alpha dZ \\ dZdW &= \rho dt \end{aligned} \qquad \text{Eq.10.6}$$

This is the SABR model of Hagan et al, published in Wilmott, 2002.

For $\beta = 1$ this reduces to the usual geometric Brownian motion with no smile; for other value of β, this is the (parametric) local volatility CEV model which produces a skew. The parameter α is the stochastic part of the smile-inducing) local volatility, and ξ is the volatility of volatility.

Let's assume $\rho = 0$ and β close to 1, and work perturbatively to get an idea of the resultant skew in this stochastic volatility model.

We know (roughly) that for $\xi = 0$, close to at-the-money, the implied volatility is the average of the local volatilities, so that

$$\Sigma_{LV}(S, t, K, T, \alpha) = \frac{\alpha}{2}[S^{\beta-1} + K^{\beta-1}] \qquad \text{Eq.10.7}$$

and for β close to (but less than) 1 we can write a Taylor expansion to obtain

$$\Sigma_{LV}(S, t, K, T, \alpha) \approx \frac{\alpha}{S^{1-\beta}}\left[1 + (\beta - 1)\ln\frac{K}{S}\right] \qquad \text{Eq.10.8}$$

This is a linear skew in $\ln\frac{K}{S}$ with negative slope, and with an at-the-money volatility that increases as the stock price drops. It is easy to check that this model satisfies the Rule of 2.

Now let's switch on the stochastic volatility by letting ξ be non-zero. Then there are a range of possible α values and we can think (roughly) of the call price C_{SLV} in this stochasticized local volatility model as being an average of the BS options prices with the implied volatility of Equation 10.8 over the distribution of α:

$$C_{SLV} = \int C_{BS}(\Sigma_{LV}(S, t, K, T, \alpha))\phi(\alpha)d\alpha$$

We can expand this about the mean of the distribution $\bar{\alpha}$

$$\begin{aligned} C_{SLV} &= \int C_{BS}(\Sigma_{LV}(S, t, K, T, \bar{\alpha} + (\alpha - \bar{\alpha})))\phi(\alpha)d\alpha \\ &\approx \int \left\{ C_{BS}(\Sigma_{LV}(S, t, K, T, \bar{\alpha})) + \frac{\partial C_{BS}}{\partial \alpha}(\bar{\alpha})(\alpha - \bar{\alpha}) + \frac{1}{2}\frac{\partial^2}{\partial \bar{\alpha}^2}(C_{BS}(\bar{\alpha})(\alpha - \bar{\alpha})^2) \right\} \phi(\alpha)d\alpha \\ &\approx C_{BS}(\bar{\alpha}) + \frac{1}{2}\frac{\partial^2}{\partial \alpha^2}(C_{BS}(\bar{\alpha}))var(\alpha) \end{aligned} \quad \text{Eq.10.9}$$

where the small volatility of volatility of volatility is assumed to allow us to cut off the Taylor series at the second term.

To see the approximate effect of this on the implied Black-Scholes volatility, we can write

$$\begin{aligned} C_{SLV} \equiv C_{BS}(\Sigma_{SLV}) &\approx C_{BS}(\Sigma_{LV}\overline{(\alpha)} + \{\Sigma_{SLV} - \Sigma_{LV}\}) \\ &\approx C_{BS}(\Sigma_{LV} + \{\Sigma_{SLV} - \Sigma_{LV}\}) \\ &\approx C_{BS}(\bar{\alpha}) + \frac{\partial C_{BS}}{\partial \Sigma_{LV}}(\Sigma_{SLV} - \Sigma_{LV}) \end{aligned} \quad \text{Eq.10.10}$$

Comparing the above two equations, we obtain

$$\Sigma_{SLV} \approx \Sigma_{LV}(\bar{\alpha}) + \frac{\frac{1}{2}\frac{\partial^2}{\partial \bar{\alpha}^2}(C_{BS}(\bar{\alpha}))var(\alpha)}{\frac{\partial C_{BS}}{\partial \Sigma_{LV}}}$$

We can evaluate the BS derivatives above for small times to expiration

to obtain

$$\Sigma_{SLV} \approx \Sigma_{LV}(\bar{\alpha})\left\{1 + \frac{1}{2}\left[\frac{\xi}{\Sigma_{LV}(\bar{\alpha})}\right]^2\left[\ln\frac{S}{K}\right]^2\right\} \quad \text{Eq.10.11}$$

The local volatility smile is altered by the addition of a quadratic term in $\ln\frac{S}{K}$

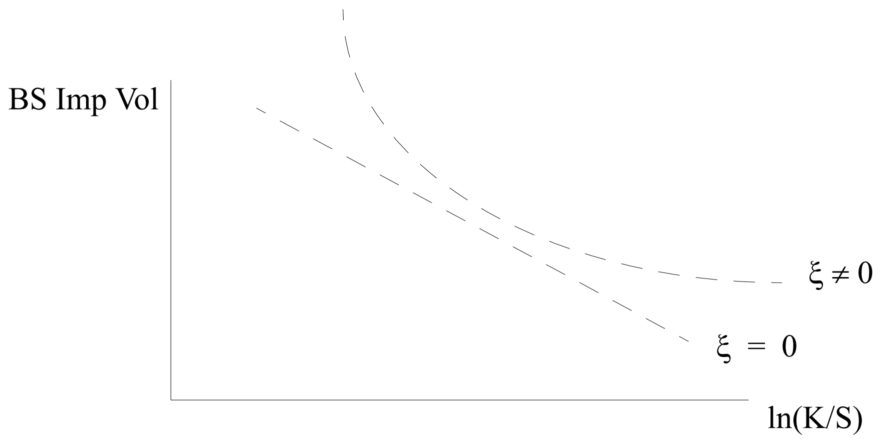

Here, because we began with a local volatility smile, there is no need for correlation between volatility and stock price in order to obtain a smile. But correlation will add to subtract to the smile too.ga

You can find more details on solving the SABR model more exactly via asymptotic perturbation theory with large volatility of volatlity in the original paper of Hagan et al., and in Wilmott's book, Chapter 55.