Chapter 16 Stochastic Volatility

The needs for stochastic volatility models come from the observations that volatilities (historical or implied) of asset prices as time series certainly are stochastic in nature, and constant volatility models such as the Black-Scholes are incapable of addressing market price behaviors, many of them directly attributed to the constant volatility assumption. In this chapter, we first address the motivations for stochastic volatility, then discuss the fundamental pricing methodologies when volatility is no longer a constant in a model, and finally we mention several derivative products that are specifically targeting at the volatility as the underlying.

16.1 Motivations and Needs in Pricing and Hedging

One obvious and major limitation in the classic Black-Scholes-Merton model is its assumption that the stock price follows a geometric Brownian motion with constant volatility. Even though there is no perfect way to determine the volatility of a stock, one thing we know for sure is that it varies in time in some *random* fashion. The implication of the constant volatility assumption leads to a log price that is supposed to have a normal distribution, a claim that is easily invalidated in reality. For actual pricing and trading of stock derivatives, this limitation gives rise to the phenomenon of so-called "volatility smile" or "volatility skew", where the implied volatilities observed on the market vary according to the strike price and the expiration of individual contracts, which means that the most important input (underlying volatility) in pricing an option needs to be adjusted in practice.
The overall purposes of developing stochastic volatility models are twofold: we want to model the actual volatility as realistic as possible, so that the stock price distribution used in the model is close to the observed data (for instance, the tails of the lognormal distribution are too thin for most observed stocks so other distributions are preferred); at the same time we want to develop a tool that fills in the gap in the implied volatility data set so that an appropriate volatility value can be used in derivative pricing and hedging.
First of all, let us revisit the concept of implied volatility σ_{imp}, which is a value that is associated with a call or put price. In the case of a call, this value is obtained by solving the following equation:

$$c_{BS}(t, S(t); K, T, \sigma_{imp}) = c_{market}, \tag{1}$$

where c_{market} is the observed call price, and c_{BS} is the Black-Scholes-Merton formula for European call. A similar relation defines the implied volatility for a put. The need for stochastic volatility becomes obvious when we notice that σ_{imp} would be different for different strike K, everything else being equal. This clearly contradicts the assumptions of the Black-Scholes model, where σ is just the volatility for the underlying stock, which has nothing to do with the strike price, and it introduces an extra source of ambiguity in option pricing. For example, if a 3-month call on stock X with strike \$50 is being traded at implied volatility of 20%, while a similar call on the same stock with strike \$45 is being quoted at 22%, what σ

value would you use in the Black-Scholes-Merton formula when you need to sell another call at strike \$55? As we know that the volatility is the single most crucial parameter in pricing, this lack of preciseness will be pivotal in the eventual profit and loss of the trading.
Another situation where a stochastic volatility model will be clearly appreciated is hedging derivatives with a significant vega exposure, that is, the derivative has a strong dependence on the volatility of the underlying. We can imagine that the derivative value depends on S_t and σ_t. As time elapses, the value of the derivative

$$V(S_t, \sigma_t, t) \longrightarrow V(S_{t+\Delta t}, \sigma_{t+\Delta t}, t + \Delta t)$$

Just like the delta hedge will balance of the change in the underlying, we would like to have a hedge to balance the change in σ, which would require a model for σ_t, as we require a model for S_t to delta hedge.
When we develop stochastic volatility models, we need to keep in mind some key phenomena observed in stock price data: (1) volatility clustering, and (2) the common highly peaked, fat tail stock return distributions. One breakthrough in option pricing in the last 20 years is the realization of the fact that volatility as a process is mean reverting in time: abnormally high or low levels of volatility cannot be sustained for very long time, so mean-reverting processes (such as Ornstein-Uhlenbeck) are natural candidates for modeling stochastic volatility. Once we have a stochastic volatility model, we would hope that the model would allow us to

- obtain certain shape of the volatility skew/smile and its evolution in time;
- possible hedging strategies in regards to the volatility risk.

16.2 A General Form of Modeling Volatility

A general form of stochastic volatility models can be expressed as a process for the stock price

$$\frac{dS(t)}{S(t)} = \mu dt + \sqrt{V(t)}\, dW_t^{(1)}, \tag{2}$$

where $V(t) = \sigma^2(t)$ is the instantaneous variance of stock return at time t, and V itself can be modeled by a process

$$dV(t) = \mu_V\,(S, V, t)\, dt + \sigma_V(S, V, t) V^{\alpha}\, dW_t^{(2)} \tag{3}$$

with a correlation between two Brownian motions $W^{(1)}$ and $W^{(2)}$ given by

$$< dW^{(1)} \cdot dW^{(2)} >= \rho\, dt, \qquad -1 \le \rho \le 1. \tag{4}$$

Notice that the parameter α is there to allow us the additional properties. For example, if $\alpha = 1/2$, we have a square-root process for V for which with a few technical conditions imposed on the other parameters, we will have V to remain positive almost surely given $V(0) > 0$. This will become important in the analytics later.
Once we have a stochastic volatility model, the following issues will arise:

- What is the probability measure related to $W^{(2)}$?
- How do we hedge the volatility risk?
- What are the special features we should include in a model, such as the mean reversion behavior?
- How do we get an estimate for the correlation? What kind of impact does correlation have on pricing and hedging?

16.3 Pricing in Stochastic Volatility Models

As in any stochastic models for financial assets, the price of a contingent claim (derivative) can be expressed as an expected value of discounted payoff, with respect to an appropriate probability measure. This principle also applies to a stochastic volatility model and we have the following approaches.

16.3.1 Closed-form Expectation

The time zero price as a discounted expectation of a payoff at time T:

$$V_0 = e^{-rT}\tilde{\mathbb{E}}\left[F(S_T)\right]$$

here the expectation is taken with the joint probability measure for S_t and σ_t (S_T depends on $\sigma_t, 0 < t < T$).
The expectation approach is only practical only in some specific situations. Here is a very simple example: the stock price follows a Black-Scholes model with the volatility used in the model

$$\bar{\sigma}^2 = \begin{cases} \sigma_1^2 & \text{with probability } p \\ \sigma_2^2 & \text{with probability } 1-p \end{cases} \tag{5}$$

and

$$\bar{\sigma}^2 = \frac{1}{T}\int_0^T \sigma_t^2\, dt$$

If this is deterministic, $C_t = C_{BS}(S_t, K, \bar{\sigma}, T-t)$. Now that $\bar{\sigma}$ is a Bernoulli rv,

$$\begin{aligned} C &= \mathbb{E}\left[C_{BS}(S, K, \bar{\sigma})\right] \\ &= pC_{BS}(S, K, \sigma_1) + (1-p)C_{BS}(S, K, \sigma_2) \\ &= C_{BS}(S, K, I(p, K)) \end{aligned} \tag{6}$$

The last two lines define an implied volatility for this model $I(p, K)$ and we can show that it does have a smile property:

$$\frac{\partial I}{\partial K} < 0 \ \text{ for } K < K_{\min}, \quad \frac{\partial I}{\partial K} > 0 \ \text{ for } K > K_{\min}.$$

for some strike $K_{\min}$ which is close to the spot price.

We can extend this to more general distributions for $\bar{\sigma}^2 = \bar{V}$:

$$C = \int g(\bar{V}) C_{BS}(\bar{V})\, d\bar{V} \tag{7}$$

under the following conditions:

- It's a European option, only the distribution of S_T is required in the expectation calculation;
- The distribution of S_T depends on $\int_0^T \sigma_t^2\, dt$ only;
- σ_t does not depend on S_t.

In practice, a moment-matching method can be used to approximate S_T by a random variable with a distribution we are familiar with (such as log-normal) so that a closed-form solution can be obtained. The approximation is made based on the matching of the first few moments.

16.3.2 Monte Carlo simulations

The idea is straightforward as we just need to simulate S_t and $V_t = \sigma_t^2$ simultaneously. But the actual implementation can be complicated as we try very hard to use various variance reduction techniques to make simulations efficient. To be more specific, we simulate S and V according to the following:

$$\frac{S_{n+1} - S_n}{S_n} = \mu\, \Delta t + \sqrt{V_n}\, \Delta W^{(1)} \tag{8}$$

$$V_{n+1} - V_n = \mu_V\, \Delta t + \sigma_V V_n^{\alpha} \Delta W^{(2)} \tag{9}$$

with $\left\langle \Delta W^{(1)} \cdot \Delta W^{(2)} \right\rangle = \rho \Delta t$. In the end, we average the discounted payoff

$$C = \frac{e^{-rT}}{M} \sum_{j}^{M} F(S_N^{(j)}), \tag{10}$$

with M the number of paths, and $S_N^{(j)}$ the simulated stock price at time t_N, along the jth path. To achieve sufficient convergence, we can expect a huge increase in cost, as this is trying to approximate a multidimensional integral with dimension N^2, in comparison with the dimension N in the constant volatility model.
What can we do to speed up the convergence? There are a couple of obvious ideas:

- Using larger time step Δt to cut down on the number of steps. However we must be careful if μ_V and σ_V are time dependent, in which case a large step will cause additional errors.

- In the case $\rho = 0$, and the coefficients μ_V and σ_V are independent of S, we can simulate V separately. If the option payoff depends on S_N only, we only need to simulate S_N which may depend only on the total variance. For example, if we have

$$\log S_N \sim N\left(\int_0^T \mu\, dt, \int_0^T \sigma^2\, dt\right)$$

then we will only need a simulated value for

$$\int_0^T \sigma^2\, dt \approx \sum_{n=0}^{N-1} V_n\, \Delta t$$

and use it in a single step simulation for S_N.

16.3.3 PDE approach

To derive the pricing equation for derivatives under the stochastic volatility model, we follow a similar no-arbitrage based approach to form a portfolio that will be hedged for risks in both stock price and volatility. More precisely, we seek a portfolio that is free of both $W^{(1)}$ and $W^{(2)}$ components. To do that, we need the Itô's formula in two dimensions and apply it to the price of a portfolio that consists of two derivatives and the underlying itself,

$$\Pi(t) = U(t) - \Delta(t)S(t) - \Delta_1(t)U_1(t), \tag{11}$$

where $U(t)$ is the price of the asset in question, and U_1 is the price of another asset that is dependent on the same underlying $S(t)$. The reason for two derivative assets in the portfolio is obvious: we have two risk factors, so more than one asset is needed in order to cancel the multi-factors in their prices. A key point in differentiating $\Pi(t)$ is to notice that we assume self-financing, therefore

$$d\Pi = dU - \Delta\, dS - \Delta_1\, dU_1 \tag{12}$$

and the dependences are on S and V. After the application of Itô's formula and regrouping, we have two terms with W_1 and W_2 dependences, and we can set them both to zero by requiring

$$\frac{\partial U}{\partial S} - \Delta_1 \frac{\partial U_1}{\partial S} = \Delta, \tag{13}$$

$$\frac{\partial U}{\partial V} - \Delta_1 \frac{\partial U_1}{\partial V} = 0. \tag{14}$$

With risk factor dependences eliminated, we must have a riskless return for the portfolio:

$$d\Pi = r\Pi\, dt \tag{15}$$

where r is the riskless rate. However, this equation alone does not directly lead to a PDE for U or U_1, as there are two unknown functions U and U_1 involved.

Another important observation is to be made, which is motivated by a separation of variable technique. Here we need to separate U terms and U_1 terms into two sides of the equation, such as

$$\mathcal{L}U = \mathcal{L}_1 U_1, \tag{16}$$

for some differential operators $\mathcal{L}$ and $\mathcal{L}_1$. Since U and U_1 are arbitrary derivative prices, both sides must be independent of either U or U_1 and we have

$$\mathcal{L}U = \mathcal{L}_1 U_1 = -f(S, V, t). \tag{17}$$

Motivated by the definition of market price of risk, a specific form of f is chosen:

$$f = \alpha - \phi\beta\sqrt{V} \tag{18}$$

Notice that α and β notations are generic and they are not necessarily related to Eq.(3). The rationale behind this form is the following: if we form a delta hedged portfolio

$$\Pi_1 = U - U_S S \tag{19}$$

and the excess return over dt is

$$d\Pi_1 - r\Pi_1 dt = \sqrt{V} U_V \left(\phi\, dt + \sigma_V\, dW^{(2)}\right). \tag{20}$$

Here ϕ represents the excess return over the risk-less interest, scaled by the volatility of volatility η. It is natural to call it the market price of risk in regard to the volatility factor. As in the risk-neutral setting for constant volatility model, if we set $\phi = 0$, the PDE for U can be written as

$$\frac{\partial U}{\partial t} + \frac{1}{2} V S^2 \frac{\partial^2 U}{\partial S^2} + \rho\sigma_V S V \frac{\partial^2 U}{\partial S \partial V} + \frac{1}{2}\sigma_V^2 V \frac{\partial^2 U}{\partial V^2} + rS\frac{\partial U}{\partial S} + \mu_V \frac{\partial U}{\partial V} = rU \tag{21}$$

The most popular model in the above form is probably Heston's model (1993), where a mean-reversion is built into the volatility process:

$$\begin{aligned} \frac{dS(t)}{S(t)} &= \mu dt + \sqrt{V(t)}\, dW^{(1)}(t), && (22) \\ dV(t) &= \kappa(\theta - V)dt + \sigma_V \sqrt{V(t)} dW^{(2)}(t). && (23) \end{aligned}$$

One advantage of this model is that we can easily interpret various parameters: κ represents the speed of mean reversion, and θ is the long time equilibrium of V. It should be pointed out that the process for V is just one example of the CIR process, where a crucial condition for the parameters is known ($\kappa\theta \geq \sigma_V^2/2$) in order to make sure that V stays above zero if it starts out positive.
The resulting PDE for Heston's model is

$$\frac{\partial U}{\partial t} + \frac{1}{2} V S^2 \frac{\partial^2 U}{\partial S^2} + \rho\sigma_V S V \frac{\partial^2 U}{\partial S \partial V} + \frac{1}{2}\sigma_V^2 V \frac{\partial^2 U}{\partial V^2} + rS\frac{\partial U}{\partial S} + \kappa(\theta - V)\frac{\partial U}{\partial V} = rU \tag{24}$$

The success of this model is due to the fact that a closed-form solution to the above equation exists for the European call and put. There are several steps involved in solving this equation for the particular form of payoff functions (call or put).

- Step 1: Introduce changes of variables

$$x = \log(Se^{-r(T-t)}/K), \quad \tau = T - t, \quad c(x, V, \tau) = U(S, V, t),$$

so we can eliminate the variable dependence in the coefficients.

- Step 2: Motivated by the Black-Scholes formula, suggest a solution form

$$c(x, V, \tau) = K \left[e^x P_1(x, V, \tau) - P_0(x, V, \tau) \right],$$

so we have P_1 and P_0 satisfying two equations that are similar.

- Step 3: The terminal conditions for P_1 and P_0 have a nice form

$$\lim_{\tau \to 0} P_j = H(x) = \begin{cases} 1 & x > 0 \\ 0 & x \leq 0 \end{cases} \qquad j = 0, 1. \tag{25}$$

- Step 4: After the treatment regarding the terminal conditions, Fourier transform can be used to convert the PDE into several ODEs.

- Step 5: Once the ODEs are solved, take the inverse Fourier transform to express the solution to the original problem in a principal-value integral. The calculation is rather tedious but it is indeed a closed-form solution.

It should be noted that in many applications the condition $\kappa\theta \geq \sigma_V^2/2$ is sometimes violated. One treatment to fix this artifact is to introduce a boundary condition (such as the reflection condition) for the V process, and impose a corresponding condition for the PDE at $V = 0$.

16.4 Hedging Issues

Suppose we have a portfolio consisting of only three assets: α shares of stock S, β shares of option C on S, and γ units of the riskless bond B, the total value is

$$P = \alpha S + \beta C + \gamma B$$

We can choose the number of shares so that $P(0) = 0$, $\partial P/\partial S = 0$, but we won't be able to have $\partial P/\partial \sigma = 0$ (vega neutral) unless we do not own any option. In order to have a vega neutral portfolio, it is necessary to have at least two different options (either strike, maturity, or both). Suppose we have two call options C_1 and C_2 with shares β_1 and β_2 respectively. In order to have a vega neutral position in the Black-Scholes model, we must have

$$\beta_1 \frac{\partial C^{(1)}}{\partial \sigma} + \beta_2 \frac{\partial C^{(2)}}{\partial \sigma} = 0,$$

where $\partial C/\partial \sigma = S_0 \sqrt{T} N'(d_1)$ is the vega of the call in Black-Scholes model. With different options, there will be different d_1 values used in the vega calculations. In the real world, the σ values used in d_1 (the implied volatility) should also differ, which makes it even more uncertain to satisfy the hedge condition. The change in portfolio value due to a volatility change is

$$\Delta P = \beta_1 \frac{\partial C_{BS}^{(1)}}{\partial \sigma} \Delta\sigma_1 + \beta_2 \frac{\partial C_{BS}^{(2)}}{\partial \sigma} \Delta\sigma_2$$

where $\Delta\sigma_1$ and $\Delta\sigma_2$ are the changes in implied vol for these two options respectively, which are not available unless there is a reliable model for the smile dynamics. A stochastic volatility model is required to generate such a smile dynamics.
It becomes obvious that the smile/skew dynamics is so important in balancing a portfolio to withstand the vol changes. It is also important to point out some general features of the smile/skew corresponding to specific situations:

- If the correlation between the underlying and its volatility is negative, the slope of the skew tends to be negative, and vise versa;
- For stochastic volatility models, the smile/skew is usually more pronounced for longer expiration than for shorter expiration. This is in a sharp contrast with jump diffusion models, where the presence of jumps impacts short terms more than long terms.

16.5 Volatility Contracts

Volatility behavior is such a central issue in financial markets and there is no doubt that various contracts have been developed over the years to attract investors with a serious exposure to the volatility risk. We just mention two types in this section.

1. Variance Swaps

   These contracts allow the investor to bet on the realized variance over a future period of time. The payoff is proportional to

   $$(\text{realized variance over certain future period of time}) - K$$

   One possible scenario is that the market just experienced some unexpected events and the volatility went surging. A seasoned investor would know that this may be a temporary condition and we should not wait too long before the volatility comes back to its normal level. He/she can certainly take a short position on the volatility by selling options at high implied vols, and hope to close out these positions later when the vol comes to its senses. Doing this will require a delta hedge (buying/selling stocks to eliminate the impact of stock price changes), which can add up the costs quite a bit. The variance swaps are pure bets on the volatility and they are going to be more effective than selling options.

   In Heston's model, we can calculate an expectation of the future variance

   $$\frac{1}{T} E\left[\int_0^T V_t \, dt\right] = m + \frac{1 - e^{-\kappa T}}{\kappa T}$$

Instead of bets on future variance, there are also contracts written on future volatility (square root of the variance), which are similar to variance swaps and the difference can be estimated by some convexity adjustments.

2. VIX contracts

   There is a major index to summarize the market volatility level (S&P 500) and it is called the VIX price, sometimes dubbed as "fear index". You can also look at the options on VIX to extract some information about the volatility of the volatility.