Stochastic Volatility

A Gentle Introduction

Fredrik Armerin

Department of Mathematics
Royal Institute of Technology, Stockholm, Sweden

Contents

Chapter 1

Introduction

1.1 Volatility

The purpose with these notes is to give an introduction to the important topic of stochastic volatility. *Volatility* is the standard deviation of the return of an asset. In the model used by Black & Scholes the stock price S_T at time T is assumed to fulfill

$$S_T = S_t e^{(\mu - \sigma^2/2)(T-t) + \sigma(W_T - W_t)},$$

where μ and $\sigma > 0$ are constants and W is a standard Brownian motion (or Wiener process). The return between time t and $t + \Delta$ can thus be written

$$\ln \frac{S_{t+\Delta}}{S_t} = \left(\mu - \frac{\sigma^2}{2}\right)\Delta + \sigma(W_{t+\Delta} - W_t).$$

From this we see two assumptions included in the Black & Scholes model:

- The returns are normally distributed, and
- the standard deviation of the returns (i.e. the volatility) is constant.

If one looks at time series of stock returns, it is quite easy to see that both these assumptions are not consistent with the data.[1]

Example 1.1.1 We have a time series with 267 observations of the daily return of the IBM share. The price and the return series are plotted in Figure 1.1. Looking at them doesn't say much. Let us now look at the 30-day volatility. By this we mean the historical volatility we get if we use the 30 latest observations. It is then scaled to give the yearly volatility. From this picture it looks like the volatility changes over time. We have also plotted a histogram of the return; see Figure 1.2. We have estimated the daily mean $\widehat{\mu}$ and standard deviation $\widehat{\sigma}$ of the

[1] It is possible to use econometric methods to give precise statistical meanings to these facts, see e.g. Campbell et al [3] and Cuthbertson [4].

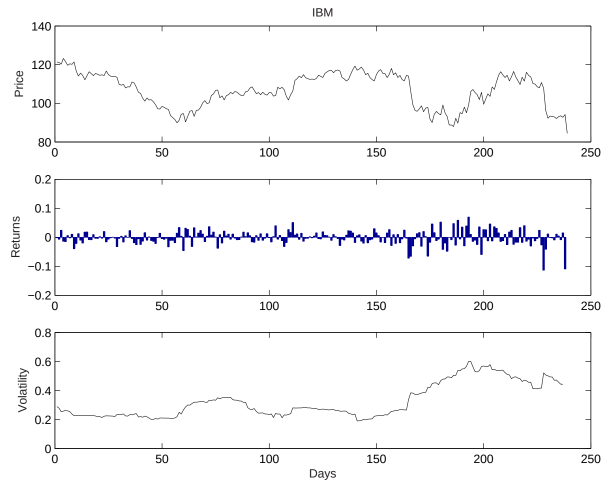

Figure 1.1: The IBM data.

returns using the whole time series and got the values

$$\widehat{\mu} = -8.38 \cdot 10^{-4} \quad \text{and} \quad \widehat{\sigma} = 0.0234.$$

In Figure 1.2 is also drawn the normal density with mean $\widehat{\mu}$ and standard deviation $\widehat{\sigma}$. Looking at the figure it seems unlikely that data is normally distributed. To convince ourself that this is the case, we will now conduct a test. For a random X we let

$$\beta_1 = \frac{E\left[X^3\right]^2}{\text{Var}(X)^3} \quad \text{and} \quad \beta_2 = \frac{E\left[X^4\right]}{\text{Var}(X)^2}.$$

If X is normally distributed, then one can show that

$$W = n\left[\frac{\beta_1}{6} + \frac{(\beta_2 - 3)^2}{24}\right] \overset{\text{As}}{\sim} \chi^2(2)$$

(see Greene [7] p. 309). With our data set consisting of $n = 267$ observations we get $\widehat{\beta_1} = 0.6311$ and $\widehat{\beta_2} = 6.592$. This gives

$$\widehat{W} = n\left[\frac{\widehat{\beta_1}}{6} + \frac{(\widehat{\beta_2} - 3)^2}{24}\right] = 171.6.$$

This value is so high that we can reject the hypothesis of normally distributed returns on any reasonable level. This is due to the fact that the two lowest returns are so extremely unlikely under the normality assumption. If we discard them, arguing that they may have occurred on an extreme day, we have a sample of $n = 265$ observations with new estimated values $\widehat{\beta_1}' = 3.225 \cdot 10^{-4}$ and $\widehat{\beta_2}' = 3.942$. This gives $\widehat{W}' = 9.813$, and corresponds to a P-value of 0.0074. Thus, we can in this case reject the hypothesis of normality on all levels down to 0.0074. $\square$

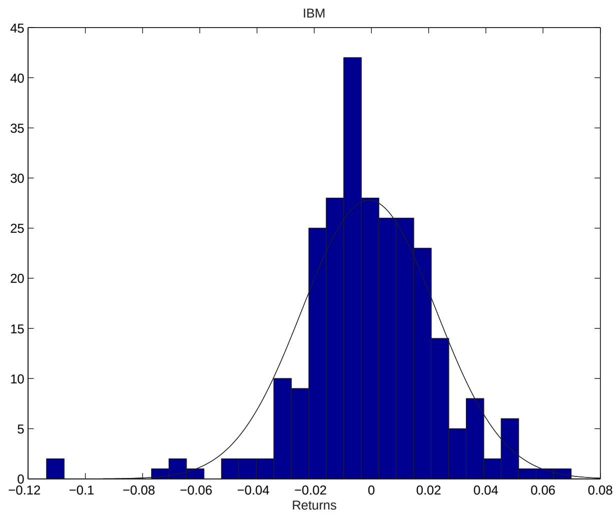

Figure 1.2: The histogram of IBM returns and the fitted normal density.

When one finds that the model one is using is not in conformity with the data, the natural thing to do is of course to modify the model. The first idea would be to allow for a deterministic but time dependent volatility. It turns out that although we do not have constant volatility any more, we still will have normally distributed returns (this is shown in Chapter 3 below). To further expand the model we can allow the volatility to depend both on time and on the state; by 'state' we mean the current stock price. State dependent volatility is random in the sense that we do not know at time t what the volatility at same later time t' will be. At time t however, the volatility for the next 'very short' time epoch will be approximately $\sigma(t, S_t)$, which is known at time t. In a stochastic volatility model, however, the volatility is random in the sense that the volatility at time t' is 'totally unknown' at some earlier time t. The above reasoning is heuristic, but the idea is that the difference between the state dependent and the stochastic volatility is that the former is 'locally known', while the latter model does not have this feature. We can summarize the different models as follows, where the complexity increases as we go downwards.

Constant volatility: σ

$\downarrow$

Time dependent volatility: $\sigma(t)$

$\downarrow$

Time-and state dependent volatility: $\sigma(t, S_t)$

$\downarrow$

Stochastic volatility: $\sigma(t, \omega)$

We know that when we price contingent claims, we work under an equivalent martingale measure. It is important to realize that when we change measure from the original one to an equivalent martingale measure, we change the drift but not the volatility of the stock price process. This is why we can use data from the real world to improve our pricing models.

The aim with these lecture notes is to cover one lecture on stochastic volatility to students familiar to the basic Black & Scholes model and the elementary stochastic calculus needed to reach the risk-neutral valuation formula. Due to this fact the list of references mostly consists of textbooks, and we refer to these for research articles on stochastic volatility.

1.2 Assumptions and notation

To make the exposition easy we will only consider models on a finite fixed time interval $[0, T]$, where $T > 0$. We will further assume that we have a given filtered probability space $(\Omega, \mathcal{F}, P, (\mathcal{F}_t)_{0 \leq t \leq T})$, were the filtration is the filtration generated by the stochastic processes of our model. We will also assume that the drift term of the stock price process is a constant times the stock price. There is no real loss in generality in doing this, since we will mostly be concerned with the behavior of the stock price process under equivalent (risk-neutral) martingale measures. Unless otherwise stated, a contingent claim X is an $\mathcal{F}_T$-measurable random variable fulfilling necessary integrability conditions. We also assume that every process and stochastic integral we are considering is well behaved and fulfills measurability and integrability conditions needed. We write $X \sim N(\mu, \sigma^2)$ to mean that the random variable X is normally distributed with mean μ and variance σ^2.

Chapter 2

The Black & Scholes model

To get started we recall some facts about the model used by Black & Scholes when they derived their formula for the price of a European call option.[1]

2.1 Valuation of contingent claims

The market is assumed to consist of risk-free lending and borrowing with constant interest rate r and a stock with price process given by a geometric Brownian motion. Let B_t and S_t denote the price processes of 'money in the bank' and the stock respectively and let W_t be a (standard) Brownian motion. The model can then be written

$$\begin{cases} dB_t &= rB_t dt; & B_0 = 1 \\ dS_t &= \mu S_t dt + \sigma S_t dW_t; & S_0 > 0, \end{cases}$$

where $r > 0$ ($r > -1$ is necessary), $\mu \in \mathbb{R}$ and $\sigma > 0$. The theory of pricing by no-arbitrage gives us the following expression for the price of the European contingent claim giving the stochastic amount X at the expiration time T:

$$\Pi_X(t) = e^{-r(T-t)} E^Q \left[X | \mathcal{F}_t \right], \quad 0 \leq t \leq T,$$

where Q is the equivalent risk-neutral measure under which S_t is a geometric Brownian motion with drift equal to rS_t (Theorem 6.1.4 in Bingham & Kiesel [2]).

To further simplify we will assume that X has the form $X = f(S_T)$ for some 'nice' function f. To be able to write the price of $X = f(S_T)$ in a more explicit way, we start by noting that the solution to the SDE satisfied by S_t under Q is

$$S_t = S_0 \exp\left(\left\{ r - \frac{\sigma^2}{2} \right\} t + \sigma W_t \right), \tag{2.1}$$

[1] Although (as is commented on in Bingham & Kiesel [2] p. 152 ff.) the model was not invented by Black & Scholes, we will refer to it as the Black & Scholes model.

and we get

$$S_T = S_t \exp\left(\left\{r - \frac{\sigma^2}{2}\right\}(T-t) + \sigma(W_T - W_t)\right). \tag{2.2}$$

By using this and the Markov property of Itô diffusions we can write

$$\Pi_X(t) = e^{-r(T-t)} E^Q\left[f\left(S_t \exp\left(\left\{r - \frac{\sigma^2}{2}\right\}(T-t) + \sigma(W_T - W_t)\right)\right)\middle| S_t\right]$$

(recall that we have $X = f(S_T)$). Equation (2.2) now gives that *conditioned on* S_t we have

$$Z = \ln\left(\frac{S_T}{S_t}\right) \sim N\left(\left\{r - \frac{\sigma^2}{2}\right\}(T-t), \sigma^2(T-t)\right). \tag{2.3}$$

Thus we can write

$$\Pi_X(t) = e^{-r(T-t)} E^Q\left[f(S_t e^Z)|S_t\right],$$

where Z is the random variable defined above. For further use we let $\Pi_X^{BS}(t;\sigma)$ denote the price at time t of the claim X, given that the stock price follows a geometric Brownian motion with volatility σ.

2.2 Implied volatility

In the Black & Scholes model the price c of a European call option is given by

$$c = S_t\Phi(d_1) - Ke^{-r(T-t)}\Phi(d_2),$$

where

$$d_1 = \frac{\ln\left(\frac{S_t}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)(T-t)}{\sigma\sqrt{T-t}} \quad \text{and} \quad d_2 = d_1 - \sigma\sqrt{T-t}.$$

We wee that the price depends on the following six quantities:

- today's date t,
- the stock price today S_t,
- the volatility σ,
- the interest rate r,
- the maturity time T, and
- the strike price K.

The interest rate and the volatility are model parameters, the valuation time t is chosen by us, and the stock price S_t is given by the market. The maturity time and strike price, finally, are specific for every option. Among these quantities, the only one that is difficult (indeed very difficult) to estimate is the volatility σ. Now assume that we observe the market price of a European option with maturity time T and strike price K. We denote this observed price by $c_{\text{obs}}(T, K)$.

With a fixed interest rate r and time t, implying that we also have a fixed S_t, we can write the theoretical Black & Scholes price of the call option as a function $c(\sigma, T, K)$. Since we cannot observe the volatility σ, a natural question is: given the observed price c_{obs}, what does this tell us about the volatility σ?

Definition 2.2.1 The *implied volatility* I of a European call option is a strictly positive solution to the equation

$$c_{\text{obs}}(T, K) = c(I, T, K). \tag{2.4}$$

The implied volatility is thus the volatility we have to insert into the Black & Scholes formula to get the observed market price of the option. Note that, with r, t and S_t still being fixed, I is a function of T, K and the observed option price c_{obs}. In the definition of the implied volatility we speak of *a* solution to Equation (2.4), and this raises the question of how many solutions there really are. This issue is resolved in the following proposition.

Proposition 2.2.2 *There can only exist at most one solution (i.e. zero or one) to Equation (2.4), and if*

$$c_{\text{obs}}(T, K) > c(0, T, K),$$

then there exists exactly one strictly positive solution.

Proof. We have (see Bingham & Kiesel [2] p. 196)

$$\frac{\partial c}{\partial \sigma} = S_t \sqrt{T-t}\,\varphi\left(\frac{\log(S_t/K) + (r+\sigma^2/2)(T-t)}{\sigma\sqrt{T-t}}\right) > 0,$$

so c is strictly increasing as a function of σ. Due to this fact, there will always exist exactly one strictly positive solution to Equation (2.4) as long as $c_{\text{obs}}(T, K) > c(0, T, K)$, and none if $c_{\text{obs}}(T, K) \leq c(0, T, K)$. □

Now also fix the time to maturity T. If the Black & Scholes model was correct, the implied volatility would be equal to the constant volatility σ specified in the model. Empirical results indicate that this is not always the case. The implied volatility as a function of K is most often not a flat curve. Instead we can typically get a 'smile' (a U-shaped curve), a 'skew' (a downward sloping curve), a 'smirk' (a downward sloping curve which increase for large K) or a 'frown' (an up-side-down U-shaped curve). The empirical evidence thus shows that the market does not price European call options according to the Black & Scholes model, that is, it does seem plausible for the volatility σ to be constant. This leads us towards the models where volatility is not constant but dependent on time.

Chapter 3

Extending the Black & Scholes model

In this chapter we will discuss two extensions of the original model of Black & Scholes that proceed the models with stochastic volatility. These extensions will be used to bridge the gap between the original Black & Scholes model and the ones with stochastic volatility.

3.1 Time dependent volatility

Let the stock price be modelled (under the risk-neutral measure Q) as

$$dS_t = rS_t dt + \sigma(t) S_t dW_t;\ S_0 = s > 0, \tag{3.1}$$

where $\sigma : [0, T] \to (0, \infty)$ is a deterministic function. The value of the bank account is again assumed to follow

$$dB_t = rB_t dt;\ B_0 = 1.$$

The solution to the SDE governing the dynamics of the stock price is given by

$$S_t = S_0 \exp\left(\int_0^t \left\{ r - \frac{\sigma^2(s)}{2} \right\} ds + \int_0^t \sigma(s) dW_s \right).$$

Defining

$$\overline{\sigma^2}(t, T) = \frac{1}{T - t} \int_t^T \sigma^2(s) ds,$$

we see that we can write the solution as

$$S_t = S_0 \exp\left(\left\{ r - \frac{\overline{\sigma^2}(0, t)}{2} \right\} t + \int_0^t \sigma(s) dW_s \right).$$

Furthermore, we see that we have

$$S_T = S_t \exp\left(\left\{r - \frac{\overline{\sigma^2}(t,T)}{2}\right\}(T-t) + \int_t^T \sigma(s)dW_s\right),$$

and that the distribution of $\ln(S_T/S_t)$ *conditioned on* S_t is given by

$$\ln\left(\frac{S_T}{S_t}\right) \sim N\left(\left\{r - \overline{\sigma^2}(t,T)/2\right\}(T-t), \overline{\sigma^2}(t)(T-t)\right). \tag{3.2}$$

Thus, with $X = f(S_T)$ (again we assume that f is a nice function) we see, comparing this expression with (2.3), that we can use the same pricing formula as in the standard Black & Scholes case. We only have to replace σ^2 with $\overline{\sigma^2}$, and doing so we arrive at, for $t \leq T$,

$$\Pi_X(t) = \Pi_X^{BS}\left(t; \sqrt{\overline{\sigma^2}(t,T)}\right),$$

where as usual $\Pi_X(t)$ denotes the price at time t of the contract X. Again everything is easy to get hold of, except for the volatility. In this case, the volatility is not merely a number, but a whole function. It turns out, however, that given the implied volatilities on the market, it is possible to derive the function $\sigma(t)$.

3.1.1 Getting $\sigma(t)$ from the implied volatility

By fixing a strike price K, we can look at the implied volatility as a function of the time to maturity T only. It will of course also be dependent on the observed options prices, but since these are given by the market, and not possible for us to choose, we regard them as parameters and suppress their dependence on the implied volatility. To conclude, we let $I(T)$ denote the implied volatility given by the observed price of some European option with given strike price K and time to maturity T.

By observing the implied volatility at some fixed time t_0 as it varies over times to maturity T, we can recover the time-dependent volatility $\sigma(t)$ for $t \geq t_0$. We will make the assumption that there exists an option with maturity time T for every $T \geq t_0$. The idea is to equate the theoretical volatility under the model given by Equation (3.1), the LHS in the next equation, with the observed implied volatility:

$$\sqrt{\frac{1}{T-t_0}\int_{t_0}^T \sigma^2(s)ds} = I(T).$$

We can write this as

$$\int_{t_0}^T \sigma^2(s)ds = I^2(T)\cdot(T-t_0).$$

Differentiating both sides with respect to T (recall that we have fixed t_0) we get

$$\sigma^2(T) = 2I(T)I'(T) \cdot (T - t_0) + I^2(T).$$

By changing $T \to t$ and taking the square root we get

$$\sigma(t) = \sqrt{2I(t)I'(t) \cdot (t - t_0) + I^2(t)} \quad \text{for every } t \geq t_0.$$

Thus, what we have achieved is an explicit formula, showing how to extract the volatility function $\sigma(t)$ from the observed implied volatilities. The problem is, from a practical point of view, that the assumption that there exists an option which mature at any given time $T \geq t_0$ is unrealistic. Most often we only have a finite number of maturity times for a European call option with strike price K. By making the assumption that $\sigma(t)$ is piecewise constant or linear we can still be able to extract the information we want from the implied volatility. See Willmott [8] Section 22.3 for more on this.

3.1.2 Conclusions

With the approach of a deterministic but time-dependent volatility we have moved away from the constant volatility model of Black & Scholes. But we see from Equation (3.2) that the returns still will be normally distributed. Since this empirically is not the fact, we must move on, trying to find a model where the returns are not normally distributed.

3.2 Time- and state dependent volatility

It is possible to model the volatility as $\sigma(t, x)$, where we insert S_t in place of x in the SDE for the stock price. The difference between this approach and the stochastic volatility one is that although the volatility is random we do not introduce any more randomness. The volatility σ is a function of S_t, which in turn is driven by the Brownian motion W_t – representing the only source of randomness in our model. We can, as in the case with time-dependent volatility, deduce $\sigma(t, S_t)$ from the implied volatilities $I(K, T)$ (now depending on both strike price and maturity time). The function $\sigma(t, x)$ consistent with observed implied volatilities $I(T, K)$ is called the *local volatility surface*. The calculations in this case is more involved than in the time-dependent case and we do not present them here. The interested reader is referred to Willmott [8] Sections 22.5–22.7.

Chapter 4

Models with stochastic volatility

The main idea with models where we have stochastic volatility is that we introduce more randomness beyond the Brownian motion driving the stock price. In a sense the models where the volatility is state-dependent is also stochastic, but for a model to be called a stochastic volatility model, we have to introduce additional randomness.

4.1 The market model

The stochastic volatility model we will use is not the most general one, but it will be sufficient for our purposes. For a slightly more general model see Section 7.3 in Bingham & Kiesel [2]. To begin with, let (W_t^1, W_t^2) be a 2-dimensional Brownian motion (remember that W^1 and W^2 then are independent) and let

$$\begin{cases} dS_t &= \mu S_t dt + \sigma(Y_t) S_t dW_t^1; \;\; S_0 > 0 \\ dY_t &= m(t, Y_t)dt + v(t, Y_t)\left(\rho dW_t^1 + \sqrt{1-\rho^2} dW_t^2\right); \;\; Y_0 \text{ given.} \end{cases} \tag{4.1}$$

Here all functions are assumed to be well behaved. Especially we will demand that $\sigma(y) > 0$ for every $y \in \mathbb{R}$. The constant parameter ρ, interpreted as the constant instantaneous correlation between dS_t/S_t and dY_t/Y_t, is further assumed to fulfill $\rho \in [-1, 1]$. If we define $Z_t = \rho W_t^1 + \sqrt{1-\rho^2} W_t^2$ then we can write $dY_t = m(t, Y_t)dt + v(t, Y_t)dZ_t$. Since W^1 and W^2 are independent Brownian motions, Z is also a Brownian motion. It further holds that $d\langle Z, W^1 \rangle_t = \rho dt$. The reason for not using (W^1, Z) instead of (W^1, W^2) is that we will make a 2-dimensional Girsanov transform later on, and then it is advantageous to have two independent Brownian motions.

We think of Y_t as some underlying process which determines the volatility. Note that $\sigma(Y_t)$ is the volatility of the stock price. We will use the short hand notation $\sigma_t = \sigma(Y_t)$. A common belief is that volatility is mean-reverting. By simply assuming that the volatility is a mean-reverting Ornstein-Uhlenbeck (OU) process will get us into trouble since we would get negative volatility with positive

$\sigma(y)$	Y_t	
e^y	$dY_t = a(b - Y_t)dt + \beta dZ_t$	(OU)
$\sqrt{y}$	$dY_t = aY_t dt + \beta Y_t dZ_t$	(GBM)
$\sqrt{y}$	$dY_t = a(b - Y_t)dt + \beta\sqrt{Y_t}dZ_t$	(CIR)

Table 4.1: Examples of pairs of a volatility function and a process driving the volatility. Here $dZ_t = \rho dW_t^1 + \sqrt{1-\rho^2}dW_t^2$.

probability. Instead we could assume that Y_t is an OU process and then let $\sigma(y) = e^y$, so that $\sigma_t = e^{Y_t}$, to avoid the problem of getting negative volatility. This idea is carried through in Fouque et al [6]. Other examples of underlying process Y_t include the geometric Brownian motion and the Cox-Ingersoll-Ross process. Finally we assume that the market contains a risk-free asset with dynamics

$$dB_t = rB_t dt;\ B_0 = 1.$$

4.2 Pricing

To price contingent claims in this model we could either use the idea of constructing a locally risk-free portfolio and then equate the return of this portfolio with the risk-free rate r, or we could look for an equivalent martingale measure. We will not proceed according to first approach (the interested reader can find this program carried through in Section 2.4 in Fouque et al [6]). Instead we will use the equivalent martingale measure approach.

Before deriving a pricing formula, we must be aware of the fact that a stochastic volatility model is an incomplete model. Recall that we say that a model is free of arbitrage if there exists at least one equivalent martingale measure. It may happen in a model that is free of arbitrage that there are more than one equivalent martingale measure. If this is the case, we have to choose one of all these measures to price the contingent claims. There is a vast literature on the subject of choosing martingale measure when the underlying model is incomplete. For a general introduction to the theory of pricing and hedging in incomplete markets, see Bingham & Kiesel [2] Section 7.1 and 7.2 respectively.

We are now ready to approach the problem of pricing in this incomplete stochastic volatility model. As in the original Black & Scholes model we change measure, moving from our original measure P to an equivalent martingale measure Q. This is performed using a Girsanov transform, and since we have two Brownian motions in our model, we make a 2-dimensional Girsanov transform. Now recall that under an equivalent martingale measure, every discounted price process should be a martingale. Looking at Equation (4.1), we see that in order for the discounted stock price process to be a martingale under any equivalent martingale measure it must have drift rS_t under this measure; precisely as in the original Black & Scholes model. If (W_t^1, W_t^2), for $t \in [0,T]$, is a 2-dimensional Brownian motion, then

$$(\widetilde{W}_t^1, \widetilde{W}_t^2) = \left(W_t^1 + \int_0^t \theta_1(s,\omega)ds, W_t^2 + \int_0^t \theta_2(s,\omega)ds \right), \text{ for } t \in [0,T],$$

is a 2-dimensional Brownian motion under the measure Q^{θ_1,θ_2}, where the Radon-Nikodym derivative $dQ^{\theta_1,\theta_2}/dP$ is given by

$$\frac{dQ^{\theta_1,\theta_2}}{dP} = \exp\left(-\frac{1}{2} \int_0^T \left(\theta_1^2(t) + \theta_2^2(t)\right) dt - \int_0^T \theta_1(t) W_t^1 - \int_0^T \theta_2(t) W_t^2 \right).$$

(This is theorem 5.8.1 in Bingham & Kiesel [2].) The previous discussion regarding the drift of the price of any traded asset under an equivalent martingale measure implies that we have

$$\begin{aligned} \theta_1(t,\omega) &= \frac{\mu - r}{\sigma(Y_t(\omega))} \\ \theta_2(t,\omega) &= \gamma(t,\omega). \end{aligned}$$

Here γ is a stochastic process that we must choose. The theory gives, however, no answer to the question of how we should choose γ. Since the process Y driving the volatility is not the price of a traded asset, we need not impose the martingale condition. Instead, and this is the core of incomplete models in terms of Girsanov transforms, we can let γ be *any* enough regular process. Thus, for every choice of γ we get an equivalent martingale measure which we can use to price contingent claims. Looking at θ_1, we see that this is the market price of risk. Due to this we call γ the *market price of volatility risk*. Since the equivalent martingale measure Q^{θ_1,θ_2} only depends on $\theta_2 = \gamma$, we will denote it by Q^γ. The expectation of a random variable X with respect to the measure Q^γ is denoted $E^\gamma[X]$. We will further let $\Pi_X^\gamma(t)$ denote the price of the claim X at time t under the equivalent martingale measure Q^γ:

$$\Pi_X^\gamma(t) = e^{-r(T-t)} E^\gamma\left[X | \mathcal{F}_t\right].$$

Generally γ can be any (sufficiently nice) adapted process. If we make the additional assumption that γ has the form

$$\gamma_t = \gamma(t, S_t, Y_t),$$

then we can derive a Black & Scholes-like PDE. Again assume that the claim we want to price is given by $X = f(S_T)$ for some function f. If we let

$$F(t,x,y) = e^{-r(T-t)} E^\gamma\left[f(S_T) | S_t = x, Y_t = y\right],$$

then it turns out that F solves the following PDE:

$$
\begin{cases}
\frac{\partial F}{\partial t} + rx\frac{\partial F}{\partial x} + \frac{1}{2}\sigma^2(y)\frac{\partial^2 F}{\partial x^2} - rF & + \quad \rho v(t,y)x\sigma(y)\frac{\partial^2 F}{\partial x \partial y} - v(t,y)\Lambda(t,x,y)\frac{\partial F}{\partial y} \\
 & + \quad m(t,y)\frac{\partial F}{\partial y} + \frac{1}{2}v^2(t,y)\frac{\partial^2 F}{\partial y^2} \\
 & = \quad 0 \\
F(T,x,y) = f(x), &
\end{cases}
$$

where

$$\Lambda(t,x,y) = \rho\frac{\mu - r}{\sigma(y)} + \gamma(t,x,y)\sqrt{1-\rho^2}.$$

We see that Λ is a (non-linear) combination of the market price of risk and the market price of volatility risk. Note that the PDE above is a Feynman-Kac PDE for the 2-dimensional diffusion (S, Y). For a derivation of it using hedging arguments, see Fouque et al [6] Section 2.4. We will now group the different parts of this PDE.[1]

$\boxed{1.}$

$$\frac{\partial F}{\partial t} + rx\frac{\partial F}{\partial x} + \frac{1}{2}\sigma^2(y)\frac{\partial^2 F}{\partial x^2} - rF \equiv \mathfrak{L}_{BS}(\sigma(y))F.$$

If this is set equal to 0, we get the ordinary Black & Scholes equation with volatility $\sigma(y)$. The differential operator $\mathfrak{L}_{BS}$ is often called the Black & Scholes operator.

$\boxed{2.}$

$$\rho v(t,y)x\sigma(y)\frac{\partial^2 F}{\partial x \partial y}.$$

This term comes from the fact that we have correlation between the two driving Brownian motions. If the correlation is zero (i.e. $\rho = 0$), this term disappears.

$\boxed{3.}$

$$v(t,y)\Lambda(t,x,y)\frac{\partial F}{\partial y}$$

This part comes from the risk premium of volatility.

$\boxed{4.}$

$$m(t,y)\frac{\partial F}{\partial y} + \frac{1}{2}v^2(t,y)\frac{\partial^2 F}{\partial y^2}$$

Recall that the (infinitesimal) generator A_X of the diffusion

$$dX_t = \mu(X_t)dt + \sigma(X_t)dW_t$$

[1] The following comments on the PDE follows Fouque et al [6] p. 46.

is given by

$$A_X f(x) = \mu(x) f'(x) + \frac{1}{2}\sigma^2(x) f''(x).$$

From this we see that this last part of the PDE is nothing but the generator of the volatility driving process Y.

Using the notation presented in the above list, we can write the PDE determining the price of a contingent claim more compactly as

$$\left\{ \mathcal{L}_{BS}(\sigma(y)) + \rho v(t,y) x \sigma(y) \frac{\partial^2}{\partial x \partial y} + v(t,y)\Lambda(t,x,y)\frac{\partial}{\partial y} + A_Y \right\} F = 0.$$

4.3 Choosing the martingale measure

Since we must choose a process γ, how do we do it? The answer is that we have to look at market prices, and from these prices try to estimate γ. In Fouque et al [6] Section 2.7 the following scheme is suggested. Choose a model for the volatility and assume that γ is a constant. Calculate the theoretical prices of European call options with different strike prices and maturity times. Then go out to the market and observe the actual prices $c_{\text{obs}}(K,T)$ for these options. Finally use the method of least-squares to estimate the parameters, i.e. solve the problem

$$\min_{\psi} \sum_{(K,T)\in\mathcal{K}} \left(c(K,T;\psi) - c_{\text{obs}}(K,T) \right)^2,$$

where ψ denotes the vector of parameters of our model, $\mathcal{K}$ is the set of strike price-maturity time pairs and $c(K,T;\psi)$ is the theoretical price for an European call option with strike price K and maturity time T under the model with parameter vector ψ. Note that it may be hard to calculate these theoretical prices.

4.4 Uncorrelated processes

If there is no correlation between S and Y (i.e. $\rho = 0$) we can use iterated expectations to get back to the case with time-dependent volatility discussed in Section 3.1 above. Again let the contingent claim be given by $X = f(S_T)$. In this case with two processes we must condition on the filtration generated by both S and Y; it is not enough only to consider the filtration generated by the stock price process alone. Our filtration is in this case given by

$$\mathcal{F}_t = \sigma\left(S_u, Y_u; 0 \le u \le t\right), \ t \in [0,T].$$

Due to independence we can write[2]

$$\mathcal{F}_t = \sigma(S_u; 0 \le u \le t) \vee \sigma(Y_u; 0 \le u \le t), \ t \in [0,T].$$

[2] If $\mathcal{F}$ and $\mathcal{G}$ and are two σ-algebras, then $\mathcal{F} \vee \mathcal{G}$ denotes the smallest σ-algebra containing all sets of $\mathcal{F}$ and $\mathcal{G}$.

Since we cannot see into the future, the information generated by Y from t to T is not in $\mathcal{F}_t$. Let

$$\sigma\{Y\} = \sigma(Y_t; 0 \leq t \leq T)$$

denote the σ-algebra generated by the whole trajectory of Y from 0 to T. Then

$$\mathcal{F}_t \subset \mathcal{F}_t \vee \sigma(Y_u; t \leq u \leq T) = \sigma\{Y\} \vee \sigma(S_u; 0 \leq u \leq t),\ t \in [0, T],$$

so the following equality follows from iterated expectations (the smallest σ-algebra wins) and the Markov property

$$\begin{aligned}
\Pi_X^\gamma(t) &= e^{-r(T-t)} E^\gamma \left[f(S_T) | \mathcal{F}_t \right] \\
&= e^{-r(T-t)} E^\gamma \left[E^\gamma \left[f(S_T) | \sigma\{Y\} \vee \sigma(S_u; 0 \leq u \leq t) \right] \Big| \mathcal{F}_t \right] \\
&= e^{-r(T-t)} E^\gamma \left[E^\gamma \left[f(S_T) | \sigma\{Y\} \vee \sigma(S_t) \right] \Big| \mathcal{F}_t \right].
\end{aligned}$$

But the inner expectation is nothing but the Black & Scholes price with time-dependent volatility $\sigma(Y_t)$ at time $t \in [0, T]$, that is

$$e^{-r(T-t)} E^\gamma \left[f(S_T) | \sigma(Y) \vee \sigma(S_t) \right] = \Pi_X^{BS} \left(t; \sqrt{\frac{1}{T-t} \int_t^T \sigma(Y_u) du} \right). \qquad (4.2)$$

Combining this we get

$$\Pi_X^\gamma(t) = E^\gamma \left[\Pi_X^{BS} \left(t; \sqrt{\frac{1}{T-t} \int_t^T \sigma(Y_u) du} \right) \Bigg| \mathcal{F}_t \right].$$

4.5 Correlated processes

This case is not so easy as the previous one. We have to solve the PDE, which may not be possible analytically. In Fouque et al [6] an approximate method for solving the pricing PDE is presented. Their book (an excellent starting point for the study of stochastic volatility) also contains a step-by-step guide to how to use their method.

4.6 The leverage effect

A well known fact is that generally $\rho < 0$ for stocks (i.e. is a negative correlation between return and volatility). That negative returns are associated with increasing volatility is known as the leverage effect. The essence of the leverage effect consists of the argument that a drop in stock price increase the volatility. Assume that the value of a firm at one time is V. This value consists of the value

$V(D)$ of the firm's debts and the value $V(E)$ of its equity: $V = V(D) + V(E)$. The equity is what is left of the firm's value after having paid the debts. That is, if we shut the firm down and pay back the debt then the equity is what is left. Thus, if the firm has N number of stocks and the stock price today is S then $V(E) = NS$ and we have

$$V = V(D) + NS.$$

The *leverage* of a firm is defined as the proportion the debt has of the firm value:

$$\text{Leverage} = \frac{V(D)}{V} = \frac{V(D)}{V(D) + NS}.$$

Now assume that we have a drastic drop in the firm's stock price. Then the leverage increase, that is, the proportion of debt of the value increases. The firm is now more sensitive against a negative change in the terms with its bond holders (the ones who have borrowed the firm its debt). Thus, one could argue that the firm should be considered a more risky one now than before the drastic drop. Since we measure risk in terms of volatility, we expect the volatility to increase, thus giving a negative correlation between the stock return and volatility.

Chapter 5

Hedging and stochastic volatility

This chapter is devoted to the question of what will happen if we believe that the volatility is some constant σ, but the true volatility is given by the stochastic process $\beta(t,\omega)$. The view we take is that of a hedger who wants to hedge his position.[1] In this chapter our market model is

$$\begin{cases} dB_t &= rB_t dt; & B_0 = 1 \\ dS_t &= \mu S_t t + \beta_t S_t dW_t; & S_0 > 0. \end{cases}$$

5.1 The cost process

Assume that we have a strategy, specified by the number of bonds and stocks we hold at time t, and denoted h_t^B and h_t^S respectively. Then the value V_t of this portfolio at time t is given by

$$V_t = h_t^B B_t + h_t^S S_t,$$

and the dynamics of the value is given by (notice that we do not impose the self-financing condition on our portfolio)

$$dV_t = h_t^B dB_t + h_t^S dS_t + B_t dh_t^B + S_t dh_t^S.$$

Now define the *cost process* as

$$C_t = \int_t^T \left(B_u dh_u^B + S_u dh_u^S \right).$$

Then $C_T = 0$ and $dC_t = -(B_t dh_t^B + S_t dh_t^S)$. Given a strategy (h^B, h^S) we interpret C_t as the cumulated cost we have at time t in order to maintain our strategy. With this definition we have

$$dV_t = h_t^B dB_t + h_t^S dS_t - dC_t,$$

[1] For a more comprehensive study of this problem see Davis [5].

and we see from this that a portfolio is self-financing if and only if $dC_t = 0$, which is equivalent to $C_T - C_t = 0$. But since $C_T = 0$, we see that we have in fact proven the following proposition:

Proposition 5.1.1 *A portfolio strategy (h^B, h^S) is self-financing if and only if the cost process associated with the strategy is identically 0.*

Notice that given *any* process (Δ_t) representing the number of stocks we want to have at time t, and *any* process V_t representing the value we want the portfolio to have at time t, we can always find a portfolio (Δ, h^B) such that

$$V_t = \Delta_t S_t + h_t^B B_t \text{ for every } t \in [0, T]$$

(simply by letting $h_t^B = (V_t - \Delta_t S_t)/B_t$), but that in general the portfolio (Δ, h^B) will not be self-financing.

5.2 Hedging a call option

Now assume that we are faced with the following situation. We have sold a call option with strike price K and maturity time T for an amount c_0 at time 0, and want to hedge this position. We believe that the volatility is some constant σ, but the true volatility is given by the stochastic process $\beta(t, \omega)$. Thus, we believe that the price of the option at time t, which we denote by P_t, is given by $P_t = F(t, S_t)$, where F solves the Black-Scholes equation

$$\begin{cases} \frac{\partial F}{\partial t} + rx\frac{\partial F}{\partial x} + \frac{1}{2}\sigma^2 x^2 \frac{\partial^2 F}{\partial x^2} &= rF \\ F(T, x) &= (x - K)^+. \end{cases}$$

We hedge our position by using the delta hedge given by

$$\Delta_t = \frac{\partial F}{\partial x}(t, S_t) \text{ and } h_t^B = \frac{F(t, S_t) - \frac{\partial F}{\partial x}(t, S_t) S_t}{B_t}.$$

If σ was the true value of the volatility, then this (continuously rebalanced) delta hedge is perfect in the sense that the value of our portfolio perfectly matches the value of the option at any time $t \in [0, T]$. This portfolio will generally not be self-financing. The dynamics of P is given by

$$dP_t = \Delta_t dS_t + h_t^B dB_t + S_t d\Delta_t + B_t dh_t^B = \Delta_t dS_t + h_t^B dB_t - dC_t.$$

Using the Itô formula we get

$$\begin{aligned} dP_t = dF(t, S_t) &= \frac{\partial F}{\partial t} dt + \frac{\partial F}{\partial x} dS_t + \frac{1}{2}\frac{\partial^2 F}{\partial x^2} d\langle S \rangle_t \\ &= \left[\frac{\partial F}{\partial t} + \frac{1}{2}\beta_t^2 S_t^2 \frac{\partial^2 F}{\partial x^2}\right] dt + \frac{\partial F}{\partial x} dS_t. \end{aligned}$$

The cost process associated with this strategy is thus given by, where we use the expressions for Δ_t and h_t^B from above,

$$\begin{aligned}
-dC_t &= dP_t - (\Delta_t dS_t + h_t^B dB_t) \\
&= \left[\frac{\partial F}{\partial t} + \frac{1}{2}\beta_t^2 S_t^2 \frac{\partial^2 F}{\partial x^2}\right] dt + \frac{\partial F}{\partial x} dS_t - \left(\frac{\partial F}{\partial x} dS_t + r\left[F - \frac{\partial F}{\partial x} S_t\right] dt\right) \\
&= \left[\left\{\frac{\partial F}{\partial t} + rS_t \frac{\partial F}{\partial x} - rF\right\} + \frac{1}{2}\beta_t^2 S_t^2 \frac{\partial^2 F}{\partial x^2}\right] dt.
\end{aligned}$$

Now we use the fact that F solves the Black-Scholes equation, which means that we can substitute the expression in the curly parenthesis with $-\frac{1}{2}\sigma^2 S_t^2 \frac{\partial^2 F}{\partial x^2}$, to arrive at

$$-dC_t = \frac{1}{2} S_t^2 \frac{\partial^2 F}{\partial x^2}(t, S_t)\left(\beta_t^2 - \sigma^2\right) dt.$$

Integrating this from 0 to T and using the fact that $C_T = 0$ implies that

$$C_0 = \frac{1}{2}\int_0^T S_t^2 \frac{\partial^2 F}{\partial x^2}(t, S_t)\left(\beta^2(t,\omega) - \sigma^2\right) dt.$$

Since $\partial^2 F/\partial x^2$ (the 'Gamma'; see Bingham & Kiesel p. 196) is strictly positive for a European call option we see that if $\sigma \geq \beta_t$ for every $t \in [0, T]$, then the cost process is non-positive for every t, i.e. we will never have to add any money in addition the amount c_0 we got at time 0; we only have to collect the surplus we gain on the hedge. But what σ should we choose? Obviously, the higher constant σ we choose, the smaller the cost will be. But now recall that we have sold the option at time 0 for the amount c_0. Let σ_{imp} denote the implied volatility representing the price c_0 (we assume that the implied volatility is strictly positive, see Proposition 2.2.2). Then we can think of this σ_{imp} as the volatility we want β_t to be below. To be concrete, let us consider the following example.

Example 5.2.1 Assume that someone wants to buy from us a European call option with strike price 90 and maturing in 3 months. The price of the stock today is 94, and the risk-free rate is 4.5 %. The volatility today is estimated to be 35 %, and we believe that it will never go beyond 70 % during the 3 months the option is alive. Using the Black & Scholes formula we get the following prices (again $c(\sigma)$ denotes the price of a European call option in the Black & Scholes model if the constant volatility is σ)

$$c(0.35) = 9.19 \quad \text{and} \quad c(0.70) = 15.37.$$

If the volatility was known to remain constant at 35 % the buyer would probably not like to pay much more than 9.19. When the volatility is stochastic, it is likely that he he is prepared to pay more than 9.19 (due to the uncertainty of future volatility). But how much more is he prepared to pay? It is possible that when presented with our suggestion of 15.37, he thinks that this is too high a price, and that he is not willing to pay more than, say, 12.50. A call option price of 12.50 corresponds to an implied volatility of 53.84 %. Now we have to decide whether to accept this offer or not. If we use the constant volatility of 53.84 % when hedging, we may end up loosing money, even though the volatility stayed below 70 %. □

5.3 Hedging general contingent claims

One reason for the fact that we will always be on the safe side (i.e. having a non-positive final cost C_0) when we hedge a European call option using a constant volatility that dominates the stochastic volatility β_t, is that the price function is convex in the stock price. Going through the arguments of the previous section, we see that we will have a non-positive final cost C_0 if

(1) $\sigma \geq \beta(t, \omega)$ for every $t \in [0, T]$, and

(2) the price function

$$F(t, x) = e^{-r(T-t)} E^Q \left[f(S_T) | S_t = x \right]$$

is convex in x.[2]

5.4 The Black-Scholes-Barenblatt equation

In this section we will consider the case when the volatility $\beta(t, \omega)$ is assumed to belong to the band $[\underline{\sigma}, \overline{\sigma}]$:

$$\underline{\sigma} \leq \beta(t, \omega) \leq \overline{\sigma} \quad \text{for every} \quad t \in [0, T],$$

where $0 < \underline{\sigma} < \overline{\sigma} < \infty$. We introduce the function

$$U(x) = \begin{cases} \overline{\sigma} & \text{if } x \geq 0 \\ \underline{\sigma} & \text{if } x < 0, \end{cases}$$

and let $F^+(t, x)$ be the solution to the equation

$$\begin{cases} \frac{\partial F^+}{\partial t} + rx \frac{\partial F^+}{\partial x} + \frac{1}{2} U \left(\frac{\partial^2 F^+}{\partial x^2} \right) x^2 \frac{\partial^2 F^+}{\partial x^2} - rF^+ & = & 0 \\ F^+(T, x) & = & f(x). \end{cases}$$

[2] Recall that a twice continuously differentiable function g is convex in x if $\frac{\partial^2 g}{\partial x^2} > 0$.

Note that this is a non-linear PDE. It is called the *Black-Scholes-Barenblatt equation*. We introduce $U(x)$ for the reason that

$$\beta_t \frac{\partial^2 F^+}{\partial x^2} \leq U\left(\frac{\partial^2 F^+}{\partial x^2}\right)\frac{\partial^2 F^+}{\partial x^2}$$

holds. Now assume that we have sold a claim at time 0 having the payoff $X = f(S_T)$, and consider the strategy consisting of holding

$$\Delta_t = \frac{\partial F^+}{\partial x}(t, S_t)$$

stocks at time $t \in [0, T]$, and letting $h_t^B = (F^+(t, S_t) - \Delta_t S_t)/B_t$. Then

$$\begin{aligned} F^+(T, S_T) &= F^+(0, S_0) + \int_0^T \Delta_t dS_t + \int_0^T h_t^B dB_t + C_0 \\ &= F^+(0, S_0) + \int_0^T \Delta_t dS_t + \int_0^T r\left(F^+(t, S_t) - \Delta_t S_t\right) dt - C_0, \end{aligned}$$

where C_0 is the cost at time 0 of the strategy (Δ, h^B). Itô's formula on $F^+(t, S_t)$ gives

$$\begin{aligned} F^+(T, S_T) &= F^+(0, S_0) + \int_0^T \left(\frac{\partial F^+}{\partial t} + \frac{1}{2}S_t^2\beta_t^2\frac{\partial^2 F^+}{\partial x^2}\right) dt + \int_0^T \frac{\partial F^+}{\partial x} dS_t \\ &\leq F^+(0, S_0) + \int_0^T \left(\frac{\partial F^+}{\partial t} + \frac{1}{2}S_t^2 U\left(\frac{\partial^2 F^+}{\partial x^2}\right)\frac{\partial^2 F^+}{\partial x^2}\right) dt + \int_0^T \frac{\partial F^+}{\partial x} dS_t \end{aligned}$$

Using the facts that F^+ solves the Black-Scholes-Barenblatt equation and that we have $\Delta_t = \partial^2 F^+/\partial x^2(t, S_t)$ we get

$$F^+(T, S_T) \leq F^+(0, S_0) + \int_0^T r(F^+(t, S_t) - S_t\Delta_t)dt + \int_0^T \Delta_t dS_t.$$

But the cost C_0 is nothing but the right-hand side minus the left-hand side of the previous relation. Thus we have shown that with the strategy (Δ, h^B) above we are always guaranteed a non-negative cost if the volatility stays within the band $[\underline{\sigma}, \overline{\sigma}]$. A hedging strategy of this type, where the cost always is non-negative, is known as a *superhedge*. In the previous section we showed how to hedge a claim which is convex is the present stock price. The method described in this section works for any claim as long as the volatility stays within $[\underline{\sigma}, \overline{\sigma}]$. For more on this approach see Avellaneda et al [1].

Bibliography

[1] Avellaneda, M., Levy, A. and Parás, A. (1995), 'Pricing and hedging derivative securities in markets with uncertain volatility', *Applied Mathematical Finance*, **2**, 73-88

[2] Bingham, N. H. & Kiesel R. (2000), 'Risk-Neutral Valuation: Pricing and Hedging of Financial Derivatives', *Springer-Verlag*

[3] Campbell, J. Y., Lo. A. W. & MacKinley, A. C. (1997), 'The Econometrics of Financial Markets', *Princeton University Press*

[4] Cuthbertson, K. (1996), 'Quantitative Financial Economics', *Wiley*

[5] Davis, M. H. A. (2001), 'Stochastic Volatility: the Hedger's Perspective', *Working Paper*

[6] Fouque, J-P., Papanicolau, G. & Sircar, K. R (2000) 'Derivatives in Financial Markets with Stochastic Volatility', *Cambridge University Press*

[7] Greene, W. H. (1997), 'Econometric Analysis', 3rd Ed., *Prentice-Hall*

[8] Willmott, P. (1998), 'Derivatives: The Theory and Practice of Financial Engineering', *John Wiley & Sons*