Financial Mathematics:
The Value of a Derivative is the Cost of its Hedge*

Global Quantitative Research Team

Version 15 June 2005

Contents

*plus the money we milk out of the investor x Correlliertes Ereignisse -¿ handeln !

1 Introduction

The rise of the financial markets in the last decades has spurred a rapid development of the relatively young field of mathematical finance. Its origins lie in Bachelier's works around 1900 [**?**], but it has gained significant attention only with the publication of the works from Black, Scholes and Merton (BSM) [**?**] and [**?**] in 1973, who were the first to establish the most fundamental principle of today's applied finance: the idea of *replication*.

While an insurance company may write contracts on events such as people's deaths, fire accidents, car crashes etc., a financial institution wants to write a contract depending on the movement of a major stock price index (a "contingent claim" or "option"). From a mathematical point of view, the main difference between the number of car crashes and the price level of some index is not so much the discontinuous nature of the car crashes.

Instead, the real difference is that we can *trade* in the stock. An insurance company has no means of reacting on market information such as increased number of car crashes once the contract is issued. But the writer (issuer) of a contingent claim is able to trade frequently in the underlying stock price throughout the life of the contract.

This is where the idea of replication starts: the issuer of a contract on the index trades in the index itself such that the value of the trading portfolio is always equal to the obligations arising from the contract. If this is possible, then the "fair" price of the contract must be the cost of the associated replication strategy and since there is no risk involved in executing the strategy, this price is independent of any risk-preference.

This fundamental insight has been formalized by Merton [**?**], Harrison and Kreps [**?**] 1979 and Harrison and Pliska [**?**] 1981.

A key concept here is the idea of a "martingale": this is a stochastic process such that the expectation of the value of the process at a later time coincides with the known value at the observation time. Hence, it resembles the idea of a "fair" game. Mathematically, the "First Fundamental Theorem of Asset Pricing" states that the market is free of arbitrage if and only if there exists a measure, equivalent to the initial measure, under which all discounted price process are martingales.[1] Moreover, if this *martingale measure* is unique, then all contingent claims can be replicated and their unique prices are given as discounted expectation under this measure. In this sense, the initial "historic" or "observable" measure plays no other role as to identify the impossible and the at all possible states of the world: its actual probabilities do not enter into the construction of a replication strategy and therefore the computation of a contingent claim's price.

[1] This is true for discrete time markets; for continuous time markets are more refined result is due to Delbaen and Schachermayer [**?**].

The success of Black, Scholes and Merton's model cannot be underestimated: ever since its publication, huge markets of liquid (exchange traded) options have developed. By now, these markets have moved far beyond BSM and the prices of so-called "vanilla" options are subject to demand and supply and cannot be explained anymore my the original BSM model.

This situation poses a new challenge to financial mathematics: Since liquid options can be traded on the exchange, there is no immediate need anymore to replicate them. This has become the job of the "market". Rather, they become by themselves *assets* and we can use them alongside the underlying share or index to hedge risks appearing in more complex structures. This in turn requires us to research models which are consistent with a range of observed liquid options which then can used to price ever more structured products.

2 The Paradigm of Replication

To illustrate the power of replication, let us first consider quite a simple scenario: hedging of a call in a one-period model. In such a a contract, you and your counter-party agree to exchange in one year's time one share for a today agreed fixed cash price K (the "strike").

For example, assume that a share is trading today at around $S_0 = €100$. Let us assume that it is well-known in the market that the stock price S_1 in one year will either be $S_1^u = €200$ with a probability of 95%, or $S_1^d = €98$ with 5%. If we want to borrow €1 from the bank today, we will have to pay back €1.05 in one year. We therefore set the value of the bank account today to $B_0 = €1$ and in one year to $B_1 = €1.05$.

The stock clearly looks like a good investment opportunity:

Graph with line up and down to illustrate the bias

In the language of probability theory, we used a probability space $(\Omega, \mathcal{A}, \mathbb{Q})$ with $\Omega = \{\omega_u, \omega_d\}$, a set of events $\mathcal{A} = \mathcal{P}(\Omega)$ (the potential set) and a probability measure $\mathbb{Q}[\{\omega_u\}] = 95\%$ and $\mathbb{Q}[\{\omega_d\}] = 5\%$. The "process" $S = (S_0, S_1)$ is defined only for the discrete times $t = 0, 1$ and it has the values $S_1(\omega_u) = €200$ and $S_1(\omega_d) = €98$. The bank account $B = (B_0, B_1)$ is deterministic.

The question at hand is: which is the minimal cash strike K we would be ready to accept in return for delivering one share in one year's time?

Well, if we were thinking in insurance terms, we might suggest to use the average return $K_1 := 95\%\, €200 + 5\%\, €98 = €194.9$. That may sound reasonable, since *on average*, neither party will make a sure gain or loss.

However, the notion of "on average" is slightly misleading: for this particular contract, the underlying stock goes either up or down. No matter how often we sell the contract, the profit/loss arising from our trade will *not* average out in any way. In particular, if the stock drops, our client will pay €194.4 for a share which is worth only €98.

Indeed, for us, a strike of €194.4 is surprisingly profitable: since we are *sure* that we will get this amount of cash in a year's time and because we are also sure that we will have to deliver one share in return, why should we not buy the share today for a meagre €100? Since we did not yet receive any money from the investor, we of course have to borrow the required amount of €100 from the bank.

After one year, the debt in our bank account is €105. Now the forward contract is exercised: we give away our share and receive the promised €194.4 from the investor. Hence, we end up with a very good-looking balance of €194.4 - €105 = €89.9 on our bank account. Even better, all the previous steps where deterministic – so the profit is completely risk-free.

Needless to say that we will not find any investor who will enter this contract: everybody else is able to do the same computations and nobody will be willing to provide us with such a "free lunch".

Indeed, we call K a "fair" strike if we are able to offer both sides of the contract to the investor. In the above example, the fair strike is just the terminal cost which occur from exercising the hedging strategy, so $K =$ €105. By construction, it is free of any risk-preference.

2.1 A Formal Approach to Hedging

The example above is of course quite simple. However, it provides valuable insight on how to hedge more general payoffs depending on S_1. Consider for example a (European) *call*: a call is the *right to buy* the stock S at time T for a fixed price K. Clearly, if the stock price S_T at T is below K, we would not make use of this right since we can buy the stock cheaper on the market. However, if S_T is larger than K, then we buy the stock for K and can sell it instantly for S_T, thereby locking in a profit of $S_T - K$. Hence, we can write the value of the payoff of a call with maturity T and strike K as

$$C_T(\omega) := \Big(S_T(\omega) - K\Big)^+ \ . \tag{1}$$

(with $x^+ := \max(0, x)$). Note that the concept of a "fair strike" does not make sense in case of a call: for any strike, the payoff is strictly positive. We will therefore aim to compute the *fair price* of the contract: a price is *fair* in our model, if we would accept it both as buyer or seller of the call.

The fundamental question in our two-period model above is now:

How much money do I need to invest in both stock and cash account to be able to replicate the terminal payoff (1) with $T = 1$ and $K =$ €100 ?

Well, since we only have two possible states of the world, we simply obtain two linear equations for the amount Δ to invest in the share and the amount β to invest in the cash account:

$$\begin{pmatrix} C_1(\omega_u) \\ C_1(\omega_d) \end{pmatrix} = \begin{pmatrix} S_1(\omega_u) & B_1 \\ S_1(\omega_d) & B_1 \end{pmatrix} \begin{pmatrix} \Delta \\ \beta \end{pmatrix} \tag{2}$$

Division by the cash bound B_1 yields in terms of the discounted values $\hat{S}_t := S_t/B_t$ and $\hat{C}_t := C_t/B_t$ the equivalent formulation

$$\begin{pmatrix} \hat{C}_1(\omega_u) \\ \hat{C}_1(\omega_d) \end{pmatrix} = \underbrace{\begin{pmatrix} \hat{S}_1(\omega_u) & 1 \\ \hat{S}_1(\omega_d) & 1 \end{pmatrix}}_{=:\hat{X}_1} \begin{pmatrix} \Delta \\ \beta \end{pmatrix} . \tag{3}$$

(It is far more natural to use discounted values for the value of a contract and the stock prices since this allows us to compare cash flows based on a single reference date.)

Since $S_1(\omega_u) \neq S_1(\omega_d)$, we can invert $\hat{X}_1$ and the inverse is given as

$$\hat{X}_1^{-1} = \frac{1}{\hat{S}_1(\omega_u) - \hat{S}_1(\omega_d)} \begin{pmatrix} 1 & -1 \\ -\hat{S}_1(\omega_d) & \hat{S}_1(\omega_u) \end{pmatrix} .$$

This in turn allows us to solve (3), and we obtain that the hedging ratios for the call are given as

$$\begin{pmatrix} \Delta \\ \beta \end{pmatrix} = \hat{X}_1^{-1} \begin{pmatrix} \hat{C}_1(\omega_u) \\ \hat{C}_1(\omega_d) \end{pmatrix} . \tag{4}$$

As before, investing Δ into the stock and β into the cash account will replicate the payoff $\hat{C}_1$. The fair price $\hat{C}_0$ of the call at time zero is by definition the price of exercising the replication strategy,

$$\begin{aligned}
\hat{C}_0 &= \left(\hat{S}_0 \, , \, 1 \right) \hat{X}_2^{-1} \begin{pmatrix} \hat{C}_1(\omega_u) \\ \hat{C}_1(\omega_d) \end{pmatrix} \\
&= \left(\frac{\hat{S}_0 - \hat{S}_1(\omega_d)}{\hat{S}_1(\omega_u) - \hat{S}_1(\omega_d)} \, , \frac{\hat{S}_1(\omega_u) - \hat{S}_0}{\hat{S}_1(\omega_u) - \hat{S}_1(\omega_d)} \right) \begin{pmatrix} \hat{C}_1(\omega_u) \\ \hat{C}_1(\omega_d) \end{pmatrix} \\
&=: \left(p_0(\omega_u) \, , p_0(\omega_d) \right) \begin{pmatrix} \hat{C}_1(\omega_u) \\ \hat{C}_1(\omega_d) \end{pmatrix} \\
&= p_0(\omega_u) \, \hat{C}_1(\omega_u) \; + \; p_0(\omega_d) \, \hat{C}_1(\omega_d) \; .
\end{aligned}$$

Remarkably, neither $\hat{X}_1^{-1}$ nor the quantities $p_0(\cdot)$ depend at all on the particular payoff. That means that the fair price of any payoff H with values $H_1(\omega_u)$ and $H_1(\omega_d)$ is uniquely given as

$$H_0 = p_0(\omega_u) \, \frac{H_1(\omega_u)}{B_1} \; + \; p_0(\omega_d) \, \frac{H_1(\omega_d)}{B_1} \; . \tag{5}$$

In our example above, we obtain

$$p_0(\omega_u) = \frac{1}{1.05} \frac{1.05 \,€100 - €98}{€200 - €98} = 6.54\% \; .$$

and $p_0(\omega_u) = 1 - p_0(\omega_d) = 93.46\%$. Consequently, the price of the call which will pay out €100 in 95% of all cases is surprisingly low:

$$C_0 = 6.54\% \frac{€200 - €100}{1.05} = €6.22 \ .$$

There is another important aspect in the previous computation: in equation (5) above, the quantities $p_0(\omega_u)$ and $p_0(\omega_d)$ are between zero and one as long as

$$\hat{S}_0 > \hat{S}_1(\omega_d) \quad \text{and} \quad \hat{S}_1(\omega_u) > \hat{S}_0 \ . \tag{6}$$

If this is not the case, that means that the stock will either always outperform an investment in the cash account or always underperform it. But this is clearly an arbitrage-situation because if the share always outperforms the cash account, nobody would invest in the cash account. Reversing the argument means that nobody would trade in the stock if it always underperforms. Hence, both situations are impossible we conclude that (6) must hold.

Hence, $p_0(\cdot)$ are actual *probabilities* and equation (5) can be written very intuitively as

$$H_0 = \frac{1}{B_1} \mathbb{E}_{\mathbb{P}} [\, H_1 \,] \tag{7}$$

in terms of the so-called *risk-neutral* measure defined by $\mathbb{P}[\{\omega_u\}] := p_0(\omega_u)$ and $\mathbb{P}[\{\omega_d\}] := p_0(\omega_d)$.

It is a general concept in finance that a price of a payoff H_1 is usually given as the discounted expectation under a so-called *risk-neutral* measure $\mathbb{P}$.

Of course, a two-state model is too limited for real-life applications. A first improvement is to allow more than one time-period: in each period $0 = t_0 < \cdots < t_n = T$, the stock can either go up or down, so we obtain a tree-like picture.

Tree

For each single node, the picture is locally the same as above: since we know the value of the contract H at maturity (it is just the payoff), we can compute the value of the contract on each node just before maturity using the same ideas around (5) above. This procedure is then iteratively applied up to to the evaluation $t = 0$ and yields today's price of the contract. This price is then also given as an expectation (8),

$$H_0 = \frac{1}{B_T} \mathbb{E}_{\mathbb{P}} [\, H_T \,] \tag{8}$$

in terms of the iteratively defined measure $\mathbb{P}$.

The obvious question is then of course what happens if the number of time-periods n goes to infinity. Under some regularity conditions, the tree model above will converge against the aforementioned Black&Scholes model: this is a continuous time model based on Brownian motion.

2.2 In the beginning, there was Brownian Motion

We now work on a probability space $\mathbb{W} = (\Omega, \mathcal{A}, \mathbb{F}, \mathbb{Q})$. A standard Brownian motion $W^{\mathbb{Q}}$ on $\mathbb{W}$ with horizon T^* is a stochastic process $W^{\mathbb{Q}} = (W_t^{\mathbb{Q}})_{t\in[0,T^*]}$ with continuous trajectories $t \mapsto W_t^{\mathbb{Q}}(\omega)$ and independent, normally distributed increments $W_t - W_s$ with mean zero and variance $t-s$. Such a process formalizes the concept of "random shocks". Indeed, BSM's model is written in differential form as

$$\left.\begin{array}{rcl} dS_t(\omega) &=& \mu S_t(\omega)\,dt + \sigma S_t(\omega)\,dW_t^{\mathbb{Q}}(\omega) \\ dB_t &=& rB_t\,dt \quad (B_0 = 1) \end{array}\right\} . \tag{9}$$

The idea is that the stock price S has some deterministic drift μ. In addition to this yield, the stock is subject to sudden random shocks $dW^{\mathbb{Q}}$ with "volatility" σ. The cash account accrues interest with an instantaneous rate of r.

The problem with (9) is that the continuous paths $t \mapsto W_t^{\mathbb{Q}}(\omega)$ of Brownian motion are very rough: they are almost surely not differentiable.

Paths of BM

Indeed, the paths $\omega \mapsto W_\cdot^{\mathbb{Q}}(\omega)$ are not of finite variation, so the Fundamental Theorem of Calculus does not hold. This observation is the starting point for developing the theory of stochastic calculus which has been pioneered by Ito [**?**] and, less well-known, Doeblin [**?**]. This theory develops the notion of a *stochastic integral* with respect to "semi-martingales" $X = (X_t)_{t\in[0,T^*]}$ (semi-martingales are the most general objects for which the stochastic integral can be defined).[2] Even though their variation is generally not finite, the quadratic variation $\langle X\rangle = (\langle X\rangle_t)_{t\in[0,T^*]}$ of such a process is a suitable integrator. As a result, the famous *Ito-formula* holds for $f \in C^{1,2}$:

$$\begin{aligned} f(T, X_T(\omega)) - f(0, X_0) &= \int_0^T \partial_t f(t, X_t(\omega))\,dt \\ &+ \int_0^T \partial_x f(t, X_t(\omega))\,dX_t(\omega) \\ &+ \frac{1}{2}\int_0^T \partial^2_{xx} f(t, X_t(\omega))\,d\langle X\rangle_t(\omega) \ . \end{aligned}$$

It is often written in differential form without notion of the ω argument,

$$df(t, X_t) = \partial_t f(t, X_t)\,dt + \partial_x f(t, X_t)\,dX_t + \frac{1}{2}\partial^2_{xx} f(t, X_t)\,d\langle X\rangle_t \ . \tag{10}$$

From there, the general theory of stochastic integration is built up. An integral

$$\int_0^T \varphi_t(\omega)\,dX_t(\omega)$$

[2] See Protter [**?**] for details.

with respect to a semi-martingale X is well-defined for all "non-anticipating" processes $\varphi = (\varphi_t)_{t\in[0,T^*]}$ which are suitably integrable[3] and is itself a semi-martingale. "Non-anticipating" means that the value of φ_t should not incorporate information beyond time t. This concept is formalized using a set $\mathbb{F} = (\mathcal{F}_t)_{t\in[0,T^*]}$ of σ-algebras $\mathcal{F}_t$, each of which represents the information available at time t.[4] The set $\mathbb{F}$ is called a *filtration* and the processes X and φ are required to be *adapted* to $\mathbb{F}$.

With this framework at hand, we can compute a solution to BSM's model (9). It is given as

$$\left.\begin{array}{lcl} S_t(\omega) & = & S_0 \exp\left\{(\mu - \frac{1}{2}\sigma)t + \sigma W_t^{\mathbb{Q}}(\omega)\right\} \\ B_t & = & e^{rt} \end{array}\right\} .$$

and S_t is strictly positive for all finite t. The process

$$X_t(\omega) := e^{W_t^{\mathbb{Q}}(\omega) - \frac{1}{2}t}$$

is called *geometric Brownian motion* and it is the *stochastic exponential* or *Dolean-Dade-exponential* of $W^{\mathbb{Q}}$: it solves the fundamental equation

$$dX_t(\omega) = X_t(\omega)\, dW_t^{\mathbb{Q}}(\omega) .$$

A picture of the paths of geometric Brownian motion illustrates better than any mathematical or financial argument why the stochastic exponential is widely used in financial modeling:

Graph of path of exponential Brownian motion and DAX

2.3 Hedging in Black, Scholes, Merton

Given BSM's model of the evolution of the stock price, the task at hand is now to develop replication strategies or *hedging strategies* for payoffs based on the underlying stock: if we promised a certain payoff to the buyer of a contract, we aim to replicate its value by continuous trading in stock and cash bond.

A payoff H_T with maturity T is a $\mathcal{F}_T$-measurable and integrable random variable $H_T \in L^1(\mathcal{F}_T)$: the measurability conditions formalizes the idea that H_T must be determined by the path of S up to T. A good example is the call (1) with strike K and maturity T,

$$H_T(\omega) := C_T(\omega) = \Big(S_T(\omega) - K\Big)^+ . \tag{11}$$

As before, we search for a hedging strategy $\varphi = (\Delta, \beta)$ which replicates the payoff H_T: the random process $\Delta = (\Delta_t)_{t\in[0,T]}$ specifies how much money we

[3] It is required that $\int_0^T \varphi_t^2\, d\langle X\rangle_t < \infty$ and that φ is locally bounded. The latter property avoids doubling-strategies.

[4] Formally, $\mathcal{F}_t = \sigma(W^{-1}(B) : B$ is a Borel-set from $C[0,t]\,)$.

will invest at any time t in the stock and β tells us how much money we will invest in the bank account. If φ replicates H_T, that means that

$$H_T(\omega) = H_0 + \int_0^T \Delta_t(\omega)\, dS_t(\omega) + \int_0^T \beta_t(\omega)\, dB_t \ . \tag{12}$$

However, we do not want to allow all possible strategies (Δ, β): To justify calling H_0 the *price* of H_T, there should be no further cash injection required after the inception of the contract. Such a strategy is called *self-financing*. It means that at any point in time, the value of our hedging efforts so far is equal to the current value of our position in S and B. Mathematically, this is expressed as

$$H_t(\omega) = \Delta_t(\omega) S_t(\omega) + \beta_t(\omega) B_t \tag{13}$$

with

$$H_t(\omega) := H_0 + \int_0^t \Delta_u(\omega)\, dS_u(\omega) + \int_0^t \beta_u(\omega)\, dB_u \ .$$

In effect, (13) reduces the freedom of finding a hedging strategy to determining Δ alone: once it is specified, β is given as the ratio

$$\beta_t(\omega) = \frac{1}{B_t}\Big(H_t(\omega) - \Delta_t(\omega) S_t(\omega)\Big) \ . \tag{14}$$

It is relatively straight forward[5] to see that this and (12) yield the representation

$$\hat{H}_t(\omega) = \hat{H}_0 + \int_0^t \Delta_t(\omega)\, d\hat{S}_t(\omega) \ . \tag{15}$$

in terms of the *discounted values* $\hat{S}_t := B_t^{-1} S_t$ and $\hat{H}_t := B_t^{-1} H_t$. The discounted stock price $\hat{S}$ satisfies the SDE

$$d\hat{S}_t(\omega) = (\mu - r)\hat{S}_t(\omega)\, dt + \sigma \hat{S}_t(\omega)\, dW_t^{\mathbb{Q}}(\omega) \ , \tag{16}$$

while the discounted bank account is simply the constant $\hat{B}_t \equiv 1$.

The idea of expressing the financial quantities S and H in relative to the value of the bank account is very natural: in fact, the value of any contract is always relative to the risk-free interest we can earn by putting our money directly in the bank account. Therefore, it makes perfect economic sense to use B itself as a unit for our computations. Mathematically, both approaches are equivalent.

[5]Using (12) we find

$$dH_t = \Delta_t \left((\mu - r) S_t\, dt + \sigma S_t\, dW_t^{\mathbb{Q}}\right) + H_t\, r dt = B_t \Delta_t \left((\mu - r)\frac{S_t}{B_t}\, dt + \sigma \frac{S_t}{B_t}\, dW_t^{\mathbb{Q}}\right) + B_t \frac{H_t}{B_t}\, r dt \ .$$

Result (15) follows from a simple rearrangement to

$$B_t^{-1} dH_t + H_t dB_t^{-1} = \Delta_t d\frac{S_t}{B_t} \ ,$$

and identifying the right hand side of this equation as $dB_t^{-1} H_t$.

Black, Scholes, Merton

To find now the required Δ hedge, BSM's model makes use of the *representation property* of Brownian motion: This property says that every random variable $\hat{H}_T \in L^1(\mathcal{F}_T)$ can be uniquely written as

$$\hat{H}_T(\omega) = \hat{H}_0^{\mathbb{Q}} + \int_0^T \phi_s^{\mathbb{Q}}(\omega)\, dW_s^{\mathbb{Q}}(\omega) \tag{17}$$

for a integrable process $\phi^{\mathbb{Q}}$ and an initial value $\hat{H}_0^{\mathbb{Q}}$.[6]

This expression does not yet help us very much yet, since the Brownian motion $W^{\mathbb{Q}}$ itself cannot be traded in the market. However, given (16) we may write

$$dW_t^{\mathbb{Q}}(\omega) = \frac{d\hat{S}_t(\omega) - (\mu - r)\hat{S}_t(\omega)\, dt}{\hat{S}_t(\omega)\sigma} \ ,$$

so we obtain

$$\hat{H}_T(\omega) = \hat{H}_0^{\mathbb{Q}} + \int_0^T \frac{\phi_t^{\mathbb{Q}}(\omega)}{\sigma \hat{S}_t(\omega)}\, d\hat{S}_t(\omega) \quad \underbrace{- \int_0^T \frac{\phi_t^{\mathbb{Q}}(\omega)(\mu - r)}{\sigma} dt}_{(*)} \ . \tag{18}$$

Unfortunately, this approach does not seem to work either: while it is perfectly possible to trade the stock according to the strategy $\phi_t^{\mathbb{Q}}/\sigma\hat{S}_t$, we cannot realize the right hand integral $(*)$ by dynamic trading in some tradable market instrument. Hence, we would have to inject additional cash into the hedge, which implies that the strategy is not self-financing. This can also be seen by comparing (18) with (15).

The solution to this hurdle comes in form of a further fundamental theorem of stochastic calculus: the *Cameron-Martin-Girsanov formula.* It dictates how processes change if we move from one measure $\mathbb{Q}$ to an alternative measure measure $\mathbb{P}$: assume that $\lambda = (\lambda_t)_{t \in [0,T^*]}$ is a bounded integrable process and define

$$D_t(\omega) := \exp\left\{ \int_0^t \lambda_s(\omega)\, dW_s^{\mathbb{Q}}(\omega) - \frac{1}{2}\int_0^t \lambda_s^2(\omega)\, ds \right\} \ .$$

This process has unit expectation under $\mathbb{Q}$ and can therefore be used to define a new equivalent[7] probability measure

$$\mathbb{P}[\ A\ ] := \mathbb{E}_{\mathbb{Q}}[\ 1_A\, D_t\ ] \qquad \text{for } A \in \mathcal{F}_t.$$

The *Cameron-Martin-Girsanov*-theorem states that under this measure, the process

$$W_t^{\mathbb{P}}(\omega) := W_t^{\mathbb{Q}}(\omega) - \int_0^t \lambda_s(\omega)\, ds \tag{19}$$

[6] The notation should highlight that both $\phi^{\mathbb{Q}}$ and $\hat{H}_0^{\mathbb{Q}}$ depend on $W^{\mathbb{Q}}$.

[7] To measures $\mathbb{P}$ and $\mathbb{Q}$ are equivalent their unlikely events are the same, i.e. if $\mathbb{P}[A] = 0$ if and only if $\mathbb{Q}[A] = 0$ for $A \in \mathcal{A}$.

is a Brownian motion.

How does this help in our situation? Well, equations (19) and (16) show that in terms of this Brownian motion, the discounted stock price follows the SDE

$$d\hat{S}_t(\omega) = \hat{S}_t(\omega)(\mu - r + \sigma\lambda_t(\omega))\,dt + \hat{S}_t(\omega)\,\sigma\,dW_t^{\mathbb{P}}(\omega)\ . \tag{20}$$

Moreover, $W^{\mathbb{P}}$ also has the representation property, hence we find a new process $\phi^{\mathbb{P}}$ and some $\hat{H}_0^{\mathbb{P}}$ such that

$$\hat{H}_T(\omega) = \hat{H}_0^{\mathbb{P}} + \int_0^T \frac{\phi_t^{\mathbb{P}}(\omega)}{\sigma\hat{S}_t(\omega)}\,d\hat{S}_t(\omega) \quad \underbrace{-\int_0^T \frac{\phi_t^{\mathbb{P}}(\omega)(\mu - r + \sigma\lambda_t(\omega))}{\sigma}dt}_{=:(**)}\ .$$

Thus we have gained an additional degree of freedom by our choice of λ. Our aim is to nullify $(**)$, we chose the measure $\mathbb{P}$ given by the process

$$\lambda_t := (r - \mu)/\sigma\ , \tag{21}$$

for which we get the desired replication

$$\hat{H}_T(\omega) = \hat{H}_0^{\mathbb{P}} + \int_0^T \Delta_t(\omega)\,d\hat{S}_t(\omega) \quad \text{with} \quad \Delta_t(\omega) := \frac{\phi_t^{\mathbb{P}}(\omega)}{\sigma\hat{S}_t(\omega)}\ . \tag{22}$$

With β defined in (14), this gives us the hedging strategy (Δ, β) we were looking for.

2.4 The Price of a Contract as an Expectation

By construction, the strategy we just computed is self-financing, hence the costs of the hedging strategy must be covered by the initial position $H_0^{\mathbb{P}}$. Therefore we call the *fair price* if H_T. But how can we efficiently calculate $H_0^{\mathbb{P}}$?

To this end, the notion of *martingale* becomes important: a process $M = (M_t)_{t\in[0,T]}$ is called a martingale under $\mathbb{P}$, if M is $\mathbb{F}$-adapted, if $M_T \in L^1$ and if "the best forecast of a future value of M given the past is its current value":

$$\mathbb{E}_{\mathbb{P}}\left[\, M_{t+u} \,\middle|\, \mathcal{F}_t \,\right](\omega) = M_t(\omega) \qquad (u \geq 0)\ .$$

For example, Brownian motion itself is a martingale (this is easy to prove using the independence if its increments). Moreover, if $M = (M_t)_{t\in[0,T]}$ is a martingale, then each Integral $N_t := \int_0^T \varphi_t\,dM_t$ is under regularity conditions also a martingale.[8] That implies in particular that $\mathbb{E}[N_t] = N_0 = 0$ for all $t \in [0, T]$.

Applied to BSM's model, we note that (20) becomes

$$d\hat{S}_t(\omega) = \hat{S}_t(\omega)\,\sigma\,dW_t^{\mathbb{P}}(\omega)\ , \tag{23}$$

[8] In general, the integral is just a local martingale; details on the subject can be found in Revuz/Yor [?].

so "under $\mathbb{P}$", $\hat{S}$ has no drift. In integral form, (23) is written as

$$\hat{S}_t(\omega) = S_0 + \int_0^t \hat{S}_s(\omega)\,\sigma\,dW_s^{\mathbb{P}}(\omega) \tag{24}$$

and form some payoff H_T, we anyway have

$$\hat{H}_t(\omega) = \hat{H}_0^{\mathbb{P}} + \int_0^t \phi_s^{\mathbb{P}}(\omega)\,dW_s^{\mathbb{P}}(\omega)\ .$$

That means that both the discounted stock and the discounted price processes $\hat{H}$ of *all* payoffs are martingales.[9] As a consequence, the price of H_T is given as

$$H_0^{\mathbb{P}} = \mathbb{E}_{\mathbb{P}}\left[\,\hat{H}_T\,\right] = \frac{1}{B_T}\mathbb{E}_{\mathbb{P}}\left[\,H_T\,\right]\ ,$$

just as in (8) for our discrete model. At some later time t, the martingale property equally yields

$$\hat{H}_t^{\mathbb{P}} = \mathbb{E}_{\mathbb{P}}\left[\,\hat{H}_t\,\middle|\,\mathcal{F}_t\,\right]\ . \tag{25}$$

In case of BSM's model, this expression allows us to compute the value of a call $H_T(\omega) := (S_T(\omega) - K)^+ = (\hat{S}_T(\omega)B_T - K)^+$ explicitly: as we saw before, the discounted stock price is given in terms of $W^{\mathbb{P}}$ as

$$\hat{S}_T = S_0 \exp\left\{\sigma W_T^{\mathbb{P}} - \frac{1}{2}\sigma^2 T\right\} \tag{26}$$

Given that $W_T^{\mathbb{P}}$ is normal with mean zero and variance T, it is straight-forward to derive the famous *Black&Scholes*-formula

$$H_0 = S_0\mathcal{N}(d^+) - e^{-rT}K\mathcal{N}(d^-) \tag{27}$$

with

$$d^{\pm} := \frac{\ln S_0/K + \left(r \pm \frac{1}{2}\sigma^2\right)T}{\sigma\sqrt{T}}$$

(we used $\mathcal{N}(x)$ to denote the standard normal cumulative distribution function).

The observation that $\hat{S}$ and all price processes $\hat{H}$ are (local) martingales under some measure $\mathbb{P}$ which is equivalent to $\mathbb{Q}$, is a very general property if the market is free of arbitrage. Let us define an arbitrage-opportunity as an investment strategy Δ such that the value of the portfolio

$$\hat{H}_T^{\Delta}(\omega) = \int_0^T \Delta_t(\omega)\,d\hat{S}_t(\omega)$$

is non-negative with a non-zero probability of being strictly positive. The key point here is that the cost of running Δ are initially zero, so the existence of

[9] The aforementioned regularity conditions hold for BSM's model and relevant payoffs.

such a Δ would imply that we could produce a "free lunch" in form of a risk-less potential profit. If such opportunities are excluded, we say that the market is free of arbitrage.

According to the Delbaen/Schachermayer's seminal paper [**?**], such markets can be characterized as follows:

Theorem 2.1 (First Fundamental Theorem of Asset Pricing) *A market is arbitrage-free if and only if there exists a measure $\mathbb{P}$ equivalent to the original measure $\mathbb{Q}$ such that all tradable assets are local martingales under $\mathbb{P}$.*

In all commonly used models, the discounted stock price is a true martingale under some measure $\mathbb{P}$. In fact, it is quite common to model $\hat{S}$ directly under the measure $\mathbb{P}$.

2.5 The Hedging Strategy in Markovian models

Now that we have shown with (25), that we can compute a price $H_t^{\mathbb{P}}$ of a contingent claim at every time $t < T$ as a conditional expectation, let us turn to the actual computation of the hedging ratio $\Delta^{\mathbb{P}}$. In the general setup above, this can be very complicated, but there is an important class of models in which it becomes feasible.

To this end, let us once more consider BSM's model (9). The discounted prices process $\hat{S}$ has the so-called *Markov*-property: intuitively, that means that at any time t, the possible future values of $\hat{S}_{t+u}$ only depend on the current state $S_t(\omega)$, but not on any past information. Mathematically, it means that for every bounded function F and all $T \geq t$,

$$\mathbb{E}_{\mathbb{P}}\left[\, F(\hat{S}_T) \,\middle|\, \mathcal{F}_t \,\right](\omega) = \mathbb{E}_{\mathbb{P}}\left[\, F(\hat{S}_T) \,\middle|\, \sigma(\hat{S}_t) \,\right](\omega) \ . \tag{28}$$

The right hand expression is the expectation of $f(\hat{S}_T)$ conditional on the knowledge of $\hat{S}_t(\omega)$ at time t (this is obviously far less information than is contained in $\mathcal{F}_t$). Since it only depends on the value of $\hat{S}$ at time t, it can also be written in terms of a function of $S_t(\omega)$,

$$f\big(t, S_t(\omega)\big) := \mathbb{E}_{\mathbb{P}}\left[\, F(\hat{S}_T) \,\middle|\, \sigma(\hat{S}_t) \,\right](\omega) \ .$$

For example, it follows from (26) that BSM's stock price can be written as

$$\hat{S}_T(\omega) = \hat{S}_t(\omega) \exp\left\{ \sigma \left(W_T^{\mathbb{P}}(\omega) - W_t^{\mathbb{P}}(\omega) \right) - \frac{1}{2}\sigma^2 T \right\} \ .$$

The independence of the increments $W^{\mathbb{P}}$ under $\mathbb{P}$ and the fact that they are normally distributed indeed means that

$$\mathbb{E}_{\mathbb{P}}\left[\, F(S_T) \,\middle|\, \mathcal{F}_t \,\right](\omega) = \int_{\mathbb{R}} F\big(S_t(\omega) e^{\sigma\sqrt{T-t}y - \frac{1}{2}\sigma^2(T-t)}\big)\, d\mathcal{N}(y) =: f\big(t, S_t(\omega)\big)$$

where $\mathcal{N}$ denotes the standard normal distribution function.

In financial terms the consequence of the Markov-property is that if a contingent claim H_T depends only of $S_T = \hat{S}_T B_T$ at the maturity of the deal (for example in the case of a call $H_t(\omega) = (S_T(\omega) - K)^+$), then the value of the claim at time $t < T$ cannot depend on any past data before t, but will only depend on $\hat{S}_t(\omega)$ today. Hence, let us assume that H_T can be written as $\hat{H}_T(\omega) = F(S_T(\omega))$. It follows from (25) and the Markov-property of $\hat{S}$ that there exists a *price-function* h such that

$$h\big(t, \hat{S}_t(\omega)\big) = \hat{H}_t^{\mathbb{P}}(\omega) \ .$$

If h is smooth enough, we can apply Ito's formula (10) and get

$$\begin{aligned}\hat{H}_T(\omega) - \hat{H}_t^{\mathbb{P}}(\omega) &= \int_t^T \partial_S h(u, \hat{S}_u(\omega))\, d\hat{S}_t(\omega) \\ &\quad + \int_t^T \left(\frac{1}{2} \partial^2_{SS} h_u(\hat{S}_u(\omega))\, \hat{S}_u^2(\omega) \sigma_u^2\, du + \partial_t h(u, \hat{S}_u(\omega)) \right) du \ ,\end{aligned}$$

where we made use of $d\langle \hat{S} \rangle_t = \hat{S}_t^2 \sigma_t^2\, dt$.

The point is now that we already know that (22) holds. Hence, we can identify the process Δ of the claim H_T at time t simply as the derivative of its price with respect to the stock:

$$\Delta_t^{\mathbb{P}}(\omega) \equiv \partial_S h\big(t, \hat{S}_t(\omega)\big) \ . \tag{29}$$

This also implies that the remaining du-terms in the above equation must cancel. This yields that the PDE

$$0 = \partial_u h(u, s) + \frac{1}{2} \partial^2_{ss} h(u, s)\, s^2 \sigma_u^2$$

must hold with boundary condition $h(T, s) = F(s)$. This equation is known as the (discounted) *Black&Scholes-PDE*. Any unique solution h to it is a valid price function for the contingent claim $\hat{H}_T(\omega) = F(\hat{S}_T(\omega))$.

The effect of relation (29) is remarkable: it means that as long as our underlying model is roughly in line with the actual market, we can compute a reliable hedging strategy for any payoff by deriving its price function in the stock variable $\hat{S}_t$. This way we are able to provide a real life hedging signal to the trading desk. Moreover, we can aggregate various different "delta positions" across different products. The resulting *delta-hedging* is exactly what is done every day on the trading desks of big financial institutions.

3 Option Markets

Black, Scholes and Merton's model was the first model which has widely been used in the financial industry. Not least due to their work, the markets for contingent claims (or "options") developed rapidly. Today, millions of dollars and euros worth of options are traded each month.

But of course, the real world has left BSM's model behind: so-called "vanilla" options such as the call (1), are today liquidly traded on exchanges. This means that they have moved from being a risky contract to being an *asset* on their own right – obviously, there is no need to replicate a traded contract. Rather, we can make use of them to improve our hedging of more complex, so-called *exotic* payoffs.

3.1 The End of BSM

Of course, wherefrom do we know that BSM's model is not sufficient anymore? After all, it is a robust approach which is well-known to many market participants.

The reason lies in what is known as the "volatility skew" or "volatility smile": if we are able to observe a call option price with maturity T and strike K, and the interest rates r, then we can invert the BSM call price function (27) and back out the only unknown parameter, σ. If it is determined in this way, this parameter is called the *implied volatility* of the option and denoted by $\hat{\sigma}(T, K)$. Using the notion of implied volatility does not imply that we actually believe that the market follows BSM's model. Rather, it is a means of assigning the "risk" in the option price a standardized measure: if implied volatility is zero, then there is no uncertainty in the future movements of the stock price.[10]

If we compute the implied volatility at a fixed maturity for a range of strikes, we will find a picture similar to the following one:

DAX implied vol for 1y

If "the market" would behave close to a BSM model, this curve would be close to a flat line. The fact that it has a very pronounced "skew" indicates that whatever cumulative pricing mechanism is at work, it is not close to BSM's model. Rather, we see that options with strikes below 100% are comparatively more expensive than those with strikes above 100% spot. Let us investigate why this could be the case:

To this end, consider a *put* contract with strike K and maturity T. Such a contract gives you the right to *sell* the underlying share S at maturity for the fixed strike price K. This is conceptually very similar to the call contract (1). Let us see how it works: if the share price S_T is below K, the investor will happily buy the share for S_T and sell it for K, thereby locking in a profit of $K - S_T > 0$. On the other hand, if the share trades above S_T, there is no money to be made and the put expires worthless. This is captured in the payoff formula

$$P_T(\omega) := \Big(K - S_T(\omega)\Big)^+$$

[10]Strictly speaking, there is additional uncertainty in the interest rates. This is neglected here.

Given this payoff and the call (1) with the same (T, K), it is easy to see that the so-called *Put-Call-parity* must hold

$$C_T(\omega) - P_T(\omega) = S_T(\omega) - K \ .$$

The right hand side is just of this equation can easily be replicated by shorting K cash in the bank account and buying one share. This hedging "strategy" is risk-less and does not depend on the actual volatility of S. Hence, pricing a put is equivalent to pricing a call. Accordingly the implied volatility $\hat{\sigma}(T, K)$ for a call and a put must be equal.

It is easy to see why put contracts are attractive to investors who want to shield themselves against a drop in the share price: if the investor holds one share and a put with strike K, the overall portfolio at time T is

$$S_T(\omega) + (K - S_T(\omega))^+ = \max \left\{ S_T(\omega), K \right\} \ .$$

Hence, the possible loss of the portfolio is floored at K.

The market's preference for this type of "capital protection" on the "downside" (i.e. strikes below 100% of spot) is the reason why options with a lower strike are relatively more expensive than higher strike options: people are simply prepared to pay more for protection than for bets on the upside in the form of calls.[11]

That may sound as if we can make money if there is a discrepancy between the option prices and the prices implied by the "real" stock price movement. That may well be. But such discrepancies will disappear quickly as "arbitrageurs" will exploit the differences instantly. This is a concern for option traders and "market makers" (who provide liquidity for listed option markets), but if we want to price and hedge complex "exotic" deals, we may well assume that the market is free of arbitrage – in that case, the Fundamental Theorem of Asset Pricing proves that there *is* some measure $\mathbb{P}$ such that the stock price and all options prices are (local) martingales.

Since the options are themselves tradable assets, they become available by themselves for our hedging strategy. Of course, if we chose a set of tradable instruments for hedging purposes, it is necessary that our model not only assigns the correct market prices to them but also produces a realistic dynamical behavior.[12]

3.2 General Stochastic Volatility Models

If we want to improve the model of the stock price to take into account the observed implied volatility skew, we enter the field of "stochastic volatility".

[11] In some equity markets such as the Japanese Nikkei, the skew is more a smile. Also, a smile is common or options on exchange rates.

[12] Note that a good "fit" to the market prices alone is not good enough. This can already be seen for the stock itself: a perfect fit is achieved with a model set is not stochastic at all, i.e. $S_t = S_0 e^{\mu t}$. But this "model" does not at all capture the risk of the movement of the stock price. We can imagine how happily other people would like to buy puts with downside strikes from us.

We assume now that the underlying probability space $(\Omega, \mathcal{A}, \mathbb{P})$ with martingale measure $\mathbb{P}$ supports a n independent standard Brownian motions $W^1, \dots, W^n$ which generate the underlying filtration $\mathbb{F} = (\mathcal{F}_t)_{t\in[0,T^*]}$; we call the vector $W = (W^1, \dots, W^n)$ an n-dimensional Brownian motion. (We will usually model the stock and other quantities directly under a martingale measure $\mathbb{P}$. For notational convenience, we will therefore omit the notion of the underlying measure for the remainder of the article.)

If we now assume that the discounted stock price process $\hat{S} = (\hat{S}_t)_{t\in[0,T^*]}$ is a strictly positive continuous $(\mathbb{F}, \mathbb{P})$-martingale, it can be shown that there exists a *stochastic variance process* ζ and a Brownian motion X such that, in differential form,

$$d\hat{S}_t(\omega) = \sqrt{\zeta_t(\omega)}\ , \hat{S}_t(\omega)\, dX_t(\omega)\ .$$

As before, this equation has the solution

$$\hat{S}_t(\omega) = S_0 \exp\left\{ \int_0^t \sqrt{\zeta_s(\omega)}\, dX_s(\omega) - \frac{1}{2}\int_0^t \zeta_s(\omega)\, ds \right\}\ .$$

The vector W once more has the representation property in the sense that every payoff $H_T \in L^1(\mathcal{F}_T)$ can again be written in terms of a sum of integrals as

$$\hat{H}_t(\omega) = \hat{H}_0 + \sum_{j=1}^{n} \int_0^t \varphi_s^j(\omega)\, dW_s^j(\omega)\ . \tag{30}$$

This is particularly true for X_{T^*}, hence there is a vector ρ such that

$$d\hat{S}_t(\omega) = \sum_{j=1}^{n} \sqrt{\zeta_t(\omega)}\, \hat{S}_t(\omega)\, \rho_t^j(\omega)\, dW_t^j(\omega) \tag{31}$$

In a general *stochastic volatility model*, we will now specify the dynamics of ζ in a way which we hope yields reasonable dynamics for the stock price and any reference instruments. So best-known classic volatility model is doubtless that of Heston [7]. In this two-factor model ($n = 2$), the stochastic variance process is the solution to the SDE

$$d\zeta(\omega) = \kappa(\theta - \zeta_t(\omega))\, dt + \nu\sqrt{\zeta_t(\omega)}\, dW_t^1(\omega) \tag{32}$$

for constants $\kappa > 0$, $\theta > 0$, $\nu > 0$ and starting in $\zeta_0 > 0$. The Brownian motion of the stock is defined as $X_t := \rho W_t^1 + \sqrt{1-\rho^2}\, W_t^2$ for a "correlation coefficient" $\rho \in (-1, 1)$

This process has a "mean-reversion" drift, i.e. in the absence of random shocks it has an exponential decay with reversion speed κ towards the long-variance level θ. The "volatility of volatility" ν controls the strength of the random shocks and the correlation parameter governs the instantaneous co-movement of the stock and the variance process. It is typically negative of, say, -0.7, which means that if the stock moves down, variance increases. This behavior is also present in the real market. It also implies that the implied volatility of an option priced with this model exhibits the desired skew.

DAX implied vol for 1y and Heston implieds

Heston's model is popular because it provides a parsimonious and intuitive description of the short variance process. Moreover, it is possible to compute the fair price of a call in this model using relatively quick numerical methods. This last point is of big importance: indeed, assume if we observe a set $C = (C^{T,K})_{(T,K)\in\mathcal{K}}$ of market call prices (given by an implied volatility skew as the one above), then we can try to find parameters $(\kappa, \theta, \nu, \rho, \zeta_0)$ such that

$$\sum_{(T,K)\in\mathcal{K}} \left\| \ C^{T,K} - B_T^{-1}\mathbb{E}_{\mathbb{P}}\left[(S_T - K)^+ \right] \ \right\|$$

is minimal. This procedure of "calibrating" a model is widely used in practise and guarantees that a model is well-fitted to relevant observed option prices. However, to implement the parameter-calibration using a numerical minimization scheme, is clearly necessary that the computation of the expectation of the option payoffs is reasonably quick.

Note that the approach also shows that the original "historic" measure $\mathbb{Q}$ does not play any role in the implementation of a model if the parameters of it are determined via calibration.[13]

Hedging in General Stochastic Volatility Models

What happens to our replication argument from the previous section?

Well, if we are in a one-factor case $n = 1$, it actually goes through in the same way it did before (for ease of exposure, let us assume that $\sigma_t^1 > 0$): as a first step, we can write

$$dW_t^1(\omega) = \frac{1}{\sigma_t^1(\omega)\hat{S}_t(\omega)}\ d\hat{S}_t(\omega)$$

which, using equation (30), gives us the desired hedging result

$$\hat{H}_T(\omega) = \hat{H}_0 + \int_0^T \frac{\varphi_s^1(\omega)}{\sigma_t^1(\omega)\hat{S}_t(\omega)}\ d\hat{S}_t(\omega)\ .$$

just as before (22).

However, it is also intuitively clear that this will not work anymore if $n > 1$: there are more sources of randomness than market instruments we can invest in. This is where the liquid options come in: if in addition to the stock, $(n-1)$ "reference instruments" with discounted price processes $\hat{M}^2, \dots, \hat{M}^n$ are traded in the market, then we should be able to hedge the remaining uncertainty in the market.

[13]The measure $\mathbb{Q}$ is used in the design and back-testing of a model: for example, in an investigation of the more general model

$$d\zeta(\omega) = \kappa(\theta - \zeta_t(\omega))\,dt + \nu\zeta_t(\omega)^\alpha\,dW_t^1(\omega)\ ,$$

Aït-Sahalia/Kimmel [?] have estimated that α is between 0.59 and 0.85.

For ease of notation, we now set $\hat{M}^1 := \hat{S}$. The first step is to use the representation property of W for the vector M. This gives us a random "volatility matrix" η such that

$$d\begin{pmatrix} \hat{M}_t^1(\omega) \\ \vdots \\ \hat{M}_t^n(\omega) \end{pmatrix} = \underbrace{\begin{pmatrix} \eta_t^{1,1}(\omega) & \cdots & \eta_t^{1,n}(\omega) \\ \vdots & \ddots & \vdots \\ \eta_t^{n,1}(\omega) & \cdots & \eta_t^{n,n}(\omega) \end{pmatrix}}_{=:\eta_t} d\begin{pmatrix} W_t^1(\omega) \\ \vdots \\ W_t^n(\omega) \end{pmatrix} .$$

To recover the vector W we obviously need to be able to invert the matrix η. This confirms the intuition that the liquid options in consideration need to be sufficiently linearly independent. It also means that if we want to recover W up to some T, none of the option is allowed to be expired: otherwise, the price process would be constant (equal to the terminal value of the payoff) and the respective row in the volatility matrix would be zero.

Under the assumption that η is invertible, using (30) shows that each payoff H can now be replicated

$$H_T(\omega) = H_0 + \sum_{j=1}^n \int_0^T \Delta_t^j(\omega)\, dM_t^j(\omega) + \int_0^T \beta_t(\omega)\, dB_t$$

by dynamically trading in stock and reference instruments using the strategy $\varphi = (\Delta^1, \dots, \Delta^n\ \beta)$; the investment strategy for the cash bond is once more determined by the self-financing requirement. According to (14), this means

$$\beta_t(\omega) = \frac{1}{B_t}\Big(H_t(\omega) - \sum_{j=1}^n \Delta_t^j(\omega) M_t^j(\omega) \Big) . \tag{33}$$

Our strategy is risk-free, which means that H_0 is the fair price of the payoff H_T. Once again, if the price H_t is a sufficiently smooth function of the market instruments,

$$H_t(\omega) = h\big(t; M_t^1(\omega), \dots, M^n(\omega)\big) ,$$

we can compute just as in (29) the hedging ratios as

$$\Delta_t^j(\omega) \equiv \partial_{M^j} h\big(t; M_t^1(\omega), \dots, M^n(\omega)\big) . \tag{34}$$

3.3 Hedging Options on Realized Variance

Well it seems all the work is done: if there are finitely many continuous risk factors W and sufficiently many independent liquid options, then we have seen that we can still hedge all payoffs.

From a theoretical point, this is satisfactory. However, from a practical point of view we now have to find some model which one hand side fits the well to the market prices of our reference instruments and which, on the other hand side, produces reasonable dynamics for those instruments.

Hence, there are choices to be made which depend on the products we want to hedge with our model. Since so-called "options on realized variance" are a particular hot topic these days, which we will use those as an example.

We will assume for the remainder of the article that $r = 0$, i.e. $B_t = 1$ such that discounted values and actual values coincide.

The "realized variance" of a stock process in the period $[0, T]$ is the estimated daily variance of the returns of the stock, commonly measured using

$$\sum_{i=1}^{n} \big(\log S_{t_i}(\omega) - \log S_{t_{i-1}}(\omega)\big)^2$$

where $0 = t_0 < \cdots < t_n = T$ are the business days up to T. This quantity will converge for $n \uparrow \infty$ for each ω to the quadratic variation of the logarithm of the stock price, and it is customary to assume that realized variance is indeed measured using

$$V_T^T(\omega) := \langle \log S \rangle_T(\omega) = \int_0^T \zeta_t(\omega)\, dt \ . \tag{35}$$

The most basic product which is based on realized variance, is a so-called *variance swap*. A variance swap simply pays V in exchange for a fixed cash strike K; since the strike K is fixed, we can assume without loss of generality that it is zero.[14] We denote the price at time t of a variance swap with maturity T by V_t^T. It is as usual given as

$$\begin{aligned} V_t^T(\omega) &:= \mathbb{E}_{\mathbb{P}}\left[\int_0^T \zeta_s\, ds \,\Big|\, \mathcal{F}_t \right](\omega) \\ &= \int_0^t \zeta_s(\omega)\, ds + \mathbb{E}_{\mathbb{P}}\left[\int_t^T \zeta_s\, ds \,\Big|\, \mathcal{F}_t \right](\omega) \ . \end{aligned}$$

The last equation follows because the variance up to t is measurable with respect to the stock price observations. Note that the prices of variance swaps are usually quoted in terms of volatility i.e. as $\sqrt{V_0^T/T}$.

The key about variance swaps is that they are reasonably liquid for most major indices. This means we can use them as reference instruments, which makes particular sense if we want to price payoffs based on the realized variance.

An example of such an "option on variance" is a call on variance,

$$H_T(\omega) := \left(\int_0^T \zeta_t(\omega)\, dt - K \right)^+ \tag{36}$$

It makes sense to buy such a call if the investor believes that the variance of the stock price will go up: a call is cheaper than the plain variance swap. For example, overall variance levels are relatively low these days (at least if compared with the end of the last decade in the internet-bubble), so an investor might

[14] The value of K paid at T is just $B_T^{-1} K$.

reasonably assume that the variance will pick up some time soon. Anyway, if this is not the case, we can maybe make a point and sell instead a put,

$$\left(K - \int_0^T \zeta_t(\omega)\,dt\right)^+ .$$

Once the option is sold, we will have to hedge ourselves. If we assume that there is more uncertainty in the market than a plain one-factor model could capture (where are risks can be hedged with by trading in the stock alone), then a reasonable approach is to use variance swaps to hedge our exposure to the variance risk.

The first step is therefore to specify a model which fits well the traded variance swap prices.

S&P market prices

Assume we find a $C^{2m,2}$ function $\mathbb{G} : \mathcal{Z} \in \mathbb{R}^m \times [0,T] \to \mathbb{R}^+$ which is seen to interpolate the market prices of variance swaps reasonably well for varying parameter z from the open set $\mathcal{Z}$. An obvious approach would then be to construct a process $Z = (Z_t)_{t\in[0,T]}$ such that

$$V_t^T(\omega) = \int_0^t \zeta_s(\omega)\,ds + \mathbb{G}(Z_t(\omega); T-t) \ . \tag{37}$$

This ensures that the variance swap prices will always have the desired functional shape. By assigning convenient dynamics to Z, we can also ensure that the prices of the variance swaps behave reasonable. Of course, the considerations of the previous chapters imply that the process Z cannot be chosen arbitrarily. In particular, $V^T = (V_t^T)_{t\in[0,T]}$ must be a martingale. Under some regularity conditions, it can be shown [**?**], that if we assume that Z is the solution to an SDE of the form

$$\begin{pmatrix} dZ_t^1 \\ \vdots \\ dZ_t^m \end{pmatrix} = \begin{pmatrix} \mu^1(Z_t) \\ \vdots \\ \mu^m(Z_t) \end{pmatrix} dt + \begin{pmatrix} \varsigma^{1,1}(Z_t) & \cdots & \varsigma^{1,n}(Z_t) \\ \vdots & \ddots & \vdots \\ \varsigma^{m,1}(Z_t) & \cdots & \varsigma^{m,n}(Z_t) \end{pmatrix} \begin{pmatrix} dW_t^1 \\ \vdots \\ dW_t^n \end{pmatrix} , \tag{38}$$

then the function $\mathbb{G}$ and the coefficients μ and ς must satisfy

$$\partial_z \mathbb{G}(z;\tau) = \sum_{i=1}^m \mu^i(z)\partial_{z^i}\mathbb{G}(z;\tau) + \frac{1}{2}\sum_{i=1}^m\sum_{j=1}^n \partial^2_{z^i z^j}\mathbb{G}(z;\tau)\sum_{k=1}^n \varsigma^{i,n}(z)\,\varsigma^{j,n}(z) \ .$$

With appropriate conventions, this can be written more clearly as

$$\partial_z \mathbb{G}(z;\tau) = \mu(z)\partial_z\mathbb{G}(z;\tau) + \frac{1}{2}\varsigma(z)^2\,\partial^2_{zz}\mathbb{G}(z;\tau) \ .$$

If this equation is satisfied for a "consistent" pair $(\mathbb{G}, Z)$, then the variance swaps a proper martingales. Moreover, deriving (37) in the maturity T of the variance swap shows that

$$\zeta_t(\omega) = \partial_\tau G(Z_t(\omega), 0) \ ,$$

hence our approach produces a stochastic volatility as introduced above (31): it merely remains to specify the instantaneous correlation vector $\rho = (\rho^1, \ldots, \rho^n)$ for the Brownian motion X. We will assume that this vector is a constant.

Double Mean-Reversion

An good example of the above model is the function $\mathbb{G} : \mathbb{R}^{+3} \times [0, \infty) \to \mathbb{R}^+$ defined as

$$\mathbb{G}(z;\tau) = z_3\tau + (z_1 - z_3)\frac{1 - e^{-\kappa\tau}}{\kappa} + (z_2 - z_3)\frac{\kappa}{\kappa - c}\left(\frac{1 - e^{-c\tau}}{c} - \frac{1 - e^{-\kappa\tau}}{\kappa}\right)$$

for strictly positive constants κ and c.[15] This function fits well to the observed data:

S&P market prices and fit of function

A consistent process Z with SDE (38) has the general form

$$\begin{aligned} dZ_t^1 &= \kappa(Z_t^2 - Z_t^1)\,dt &+& \ \varsigma^1(Z_t)^T\,dW_t \\ dZ_t^2 &= \kappa(Z_t^3 - Z_t^2)\,dt &+& \ \varsigma^2(Z_t)^T\,dW_t \\ dZ_t^3 &= && \ \varsigma^3(Z_t)^T\,dW_t \end{aligned}$$

where the volatility vectors $\varsigma^1, \ldots, \varsigma^3$ must ensure that $Z^1, \ldots, Z^3$ stay strictly positive. From the definition of $\mathbb{G}$ it follows that $\zeta_t = Z_t^1$.

The intuition behind such a process is indeed that just like Heston (32), the stochastic variance process ζ is again mean-reverting. The difference is now that the level of mean-reversion is itself stochastic and mean-reverting. The third parameter Z^3 allows additional dependency on this "long term" level of Z^2.

Once this model is implemented, we can finally hedge our options on variance. Consider a call on variance (36) and recall (34).

We proceed as follows: by construction, the model vector $(S, Z^1, \ldots, Z^3)$ is Markov. That means that the price H_t of each payoff H_T is given in terms of a function H as

$$H_t(\omega) = \tilde{h}\big(t; S_t(\omega), Z_t^1(\omega), \ldots, Z_t^3(\omega)\big) \tag{39}$$

Unfortunately, Z is not itself tradable. Hence, we chose three reference variance swaps $V^1, \ldots, V^3$ with maturities $T_1, \ldots, T_3$. The function

$$\tilde{\mathbb{G}} : z \in R^{+3} \to \Big(\mathbb{G}(z; T_1), \mathbb{G}(z; T_2), \mathbb{G}(z; T_3)\Big)$$

[15] The limit $\kappa = c$ exists.

is (numerically) invertible with smooth inverse $\tilde{G}^{-1}$. Hence, we can write (39) as

$$H_t(\omega) = \tilde{h}\big(t; S_t(\omega), \tilde{G}^{-1}(V_t^1(\omega) - V_t^t(\omega), \ldots, V_t^3(\omega) - V_t^t(\omega))\big)$$

and therefore as

$$H_t(\omega) = h\Big(t; S_t(\omega), V_t^1(\omega), \ldots, V_t^3(\omega); \int_0^t \zeta_s(\omega)\, ds\Big) .$$

This gives us the hedging strategy

$$\Delta_t^S := \partial_S h(\cdots) \quad \text{and} \quad \Delta_t^j := \partial_{V^j} h(\cdots) \;\; j = 1, \ldots, 3$$

according to (34). The computation of these ratios can therefore implemented using a numerical differentiation (usually, a simple centered difference with a specified width is sufficient).

References

[1] *Global Quantitative Research Team:*
"Control Variates", August 2004

[2] *Global Quantitative Research Team:*
"Stationary Heston", August 2003

[3] *Global Quantitative Research Team:*
"Double Heston", January 2005

[4] *P.Carr, D.Madan:*
"Option Valuation Using the Fast Fourier Transform", March 1999

[5] *D.Lamberton, B.Lapeyre:*
Introduction au Calcul Stochastique applique a la Finance, ellipses, 1997

[6] *Kloeden, E.Platen:*
To be filled in

[7] *S.Heston:*
"A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options", 1993