Recent Developments in Mathematical Finance: A Practitioner's Point of View

Marcus Overhaus
Ana Bermudez, Hans Buehler, Andy Ferraris,
Christopher Jordinson, Aziz Lamnouar, Anwar Puthu

Quantitative Products Analytics
Deutsche Bank AG

Global Markets Equity
Deutsche Bank AG London
1 Great Winchester Street
London EC2N 2EQ, United Kingdom

Version December 11, 2005

Abstract

This article is an introduction into the principles of financial mathematics with a strong view on how the theory is actually used in practise. Using a very basic model, we explain the importance of replication of a financial contract. Afterwards, we show that the concept of replication is an integral part not only of the famous Black&Scholes model but also of more general diffusion based models.

The ideas are then applied to the pricing and hedging of options on variance. In this context, we present recent results on the modelling of the term-structure of variance swaps.

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 [B00], but it has gained significant attention only with the publication of the works from Black, Scholes and Merton (BSM), [BS73] and [M73a], in 1973, who were the first to apply systematically 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 may want to write a contract depending on the movement of a stock price index (a "contingent claim" or "option"). From a economic point of view, the main difference between the number of car crashes and the price level of some index is that we can *trade* in the stock. An insurance company has no means of reacting to market information such as an increased number of car crashes once the contract has been issued. However, 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 tries to trade in the index itself such that the value of the trading portfolio is always equal to the value of 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 [M73b], Harrison and Kreps [HK79] and Harrison and Pliska [HP81] for markets which are "free of arbitrage" in the sense that there is no[1] trading strategy whose execution bears no costs, but which has a non-zero probability of a positive payoff (we call such a strategy a "free lunch").

A key concept here is the idea of a *martingale*: this is a stochastic process such that the expected 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.[2] Moreover, if this *martingale measure* is unique, then all contingent claims can be replicated and their unique prices are given as discounted expectations under this measure. In this sense, the initial "historic" or "observable" measure plays no other role than to identify which are the possible and which are the impossible 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.

The success of Black, Scholes and Merton's model cannot be underestimated: ever since the publication of their articles, huge markets of liquid (i.e. 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 supply and demand and cannot be explained anymore by the original BSM model.[3]

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 role 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 be used to price advanced structured products.

2 The Paradigm of Replication

To illustrate the idea of replication, let us first consider a quite simple scenario: hedging of a call option in a one-period model.

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 in the market. However, if S_T is larger than K, then we can buy the stock for K and 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

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

(with $x^+ := \max(0, x)$).

Let us assume that a share is trading today at around $S_0 =$ €100. We also assume that it is well-known in the market that the stock price S_1 in one year will either be €200 with a probability of 95%, or €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.

In the language of probability theory, we use a probability space $(\Omega, \mathcal{A}, \mathbb{Q})$ with two states of the world, $\Omega = \{\omega_u, \omega_d\}$, a set of events $\mathcal{A} = \mathcal{P}(\Omega)$ (the power 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 cash account $B = (B_0, B_1)$ is deterministic. We also assume that we can invest any amount we want in shares or the cash account, including fractional and negative amounts. Moreover, we assume there are no transaction costs, taxes etc.

Under these assumptions, we want to sell a call with strike $K =$ €100. Such a call contract looks like quite a good investment opportunity: it pays out €100 in 95% of all possible future scenarios.

[1] admissible

[2] This is true for discrete time markets; for continuous time markets a more refined result is due to Delbaen/Schachermayer [DS94].

[3] In fact, the underlying mathematical model in BSM's work was originally proposed by Samuelson [S65]. However, it is customary today to refer to it as "Black-Scholes-Merton"- or simply "Black-Scholes"-model.

The question at hand is: what is a "fair" price of the call ?

If we were thinking in insurance terms, we might suggest to use as a price the average return: 95% €100 = €95. That may sound reasonable since, *on average*, neither party will make a sure gain or loss. However, we missed the main point made in the introduction: we can trade in the underlying. Can we implement a trading strategy which allows us to deliver the promised payoff without any risk? In our case, the "strategy" can only be an initial investment of Δ into the stock and an investment of β into the cash account. The "fair" price would then be the cost of the implementation of the strategy – hence, the real question is:

How much money do we 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:

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

Now note that it is far more natural to use discounted values for the value of a contract and the stock price since this allows us to compare cash flows based on a single reference date. We therefore divide the above equation by the cash value B_1 and obtain an equivalent equation in terms of the discounted values $\hat{S}_t := S_t/B_t$ and $\hat{C}_t := C_t/B_t$:

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

Since $S_1(\omega_u) \neq S_1(\omega_d)$, we can invert $\hat{X}_1$ (which does not depend on the payoff C_1) and compute

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

The initial price of this strategy is therefore

$$\hat{C}_0 = (\,\hat{S}_0\,,\,1\,)\hat{X}_1^{-1} \begin{pmatrix} \hat{C}_1(\omega_u) \\ \hat{C}_1(\omega_d) \end{pmatrix} = p_0(\omega_u)\,\hat{C}_1(\omega_u) \;+\; p_0(\omega_d)\,\hat{C}_1(\omega_d) \;.$$

with

$$p_0(\omega_u) = \frac{\hat{S}_0 - \hat{S}_1(\omega_d)}{\hat{S}_1(\omega_u) - \hat{S}_1(\omega_d)} \quad \text{and} \quad p_0(\omega_d) = \frac{\hat{S}_1(\omega_u) - \hat{S}_0}{\hat{S}_1(\omega_u) - \hat{S}_1(\omega_d)} \;.$$

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 $p_0(\omega_d) = 1 - p_0(\omega_u) = 93.46\%$. Consequently, the fair 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, it 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, then we borrow money to buy the shares and reap a sure profit. Reversely, we sell the stock and invest the proceeds in the cash account.[4] Now imagine one of the two situations would occur. In that case, everybody would rush to exploit the arbitrage opportunity. As an effect, it would disappear quickly, and we conclude intuitively that in a frictionless market no arbitrage opportunity should exists. Under this assumption, that "there is no free lunch", we have therefore shown that (6) must hold.

If (6) holds, $p_0(\cdot)$ can be interpreted as *probabilities* of the outcomes ω_u and ω_d under some measure. Equation (5) can then be written very intuitively as an expectation

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

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 process: for each single node, the situation 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 backwards to the evaluation date $t = 0$ and yields today's price of the contract. This price can also be written as an expectation,

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

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

The obvious question is then what happens if the number of time-periods n goes to infinity. Under some regularity conditions,[5] the tree model above will converge against the aforementioned Black/Scholes/Merton model: this is a continuous time model based on Brownian motion.

2.1 Stochastic Calculus

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". The BSM-model (based on Samuelson's work [S65]) is written in differential form as

$$\begin{aligned} dS_t(\omega) &= \mu S_t(\omega)\, dt + \sigma S_t(\omega)\, dW_t^{\mathbb{Q}}(\omega) \qquad && S_0 \in \mathbb{R}^+ \\ dB_t &= rB_t\, dt && B_0 = 1 \ . \end{aligned} \tag{9}$$

The idea is that the stock price S has some deterministic drift μ and that it 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 (see figure 2.1).

[4] One of the assumptions was that we can sell a stock which we do not own.

[5] See, for example Föllmer/Schied [FS04].

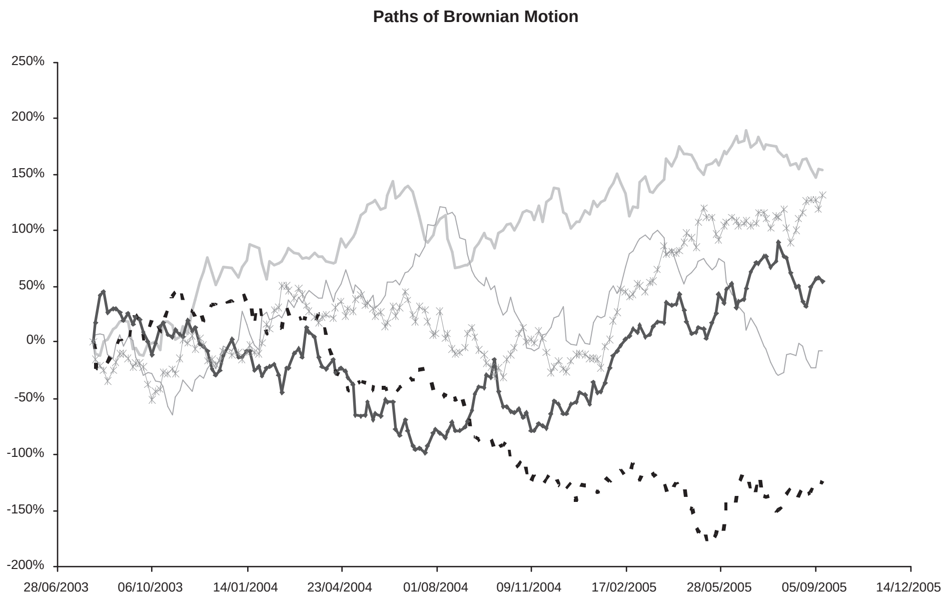

Figure 1: Sample paths of Brownian motion

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 Itô [I44] and, less well-known, Doeblin [D40]. This theory develops the notion of stochastic diffusion processes which are solutions to stochastic differential equations of the form

$$dX_t(\omega) = a\big(t; X_t(\omega)\big)\,dt + \sum_{j=1}^{n} b^j\big(t; X_t(\omega)\big)\,dW_t^j(\omega) \qquad X_0 = x \in \mathbb{R} \tag{10}$$

where $W = (W^1, \ldots, W^n)$ is a vector of independent Brownian motions and where the "drift" vector a and the "volatility" matrix b are measurable functions. If they are sufficiently well-behaved (for example, if they are globally Lipschitz) then (10) has a unique solution.[6]

Even though the variation of such a *diffusion* X is generally not finite, the quadratic variation

$$\begin{aligned} \langle X\rangle_T(\omega) &:= \lim_{m\uparrow\infty} \sum_{i=1}^{m} \Big(X_{i\frac{T}{m}}(\omega) - X_{(i-1)\frac{T}{m}}(\omega)\Big)^2 \\ &= \sum_{j=1}^{n} \int_0^T b^j\big(t; X_t(\omega)\big)^2\,dt \end{aligned} \tag{11}$$

of such a process is a suitable integrator. As a result,[7] the famous *Itô-formula* holds for $f \in C^{1,2}$:

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

[6] E.g. see Protter [P04].
[7] See Föllmer [F81].

$$+\sum_{j=1}^{n}\int_0^T \partial_x f(t,X_t(\omega))\, b^j\big(t;X_t(\omega)\big)\, dW_t^j(\omega)$$
$$+\frac{1}{2}\sum_{j=1}^{n}\int_0^T \partial_{xx}^2 f(t,X_t(\omega))\, b^j\big(t,X_t(\omega)\big)^2\, dt\ .$$

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

$$\begin{aligned} df(t,X_t) &= \Big(\partial_t f(t,X_t) + \partial_x f(t,X_t)\, a(t,X_t) + \tfrac{1}{2}\partial_{xx}^2 f(t,X_t)\, b(t,X_t)^2\Big)\, dt \\ &\quad + \partial_x f(t,X_t)\, a(t,X_t)\, dW_t\ . \end{aligned} \tag{12}$$

Extensions to the case where X takes values in $\mathbb{R}^m$ are straight-forward. From here, the entire theory is built up: an integral

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

with respect to X is again a well-defined diffusion for all "non-anticipating" processes $\varphi = (\varphi_t)_{t\in[0,T^*]}$ which are suitably integrable.[8] "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. The set $\mathbb{F}$ is called a *filtration* and the processes X and φ are required to be *adapted* to $\mathbb{F}$.

Using Ito's formula (12) it is straight forward to confirm that

$$\begin{aligned} S_t(\omega) &= S_0 \exp\Big\{(\mu - \tfrac{1}{2}\sigma)t + \sigma W_t^{\mathbb{Q}}(\omega)\Big\} \\ B_t &= e^{rt} \end{aligned}$$

is the solution to BSM's model (9). Such a process S is called a *geometric Brownian motion with drift* μ *and volatility* σ and it is strictly positive for all finite t. In the special case where the drift vanishes and where σ is one, S is called *stochastic exponential* or *Doléans-Dade exponential* because it solves the fundamental equation

$$dS_t(\omega) = S_t(\omega)\, dW_t^{\mathbb{Q}}(\omega)\ .$$

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

2.2 Hedging in Black, Scholes, Merton

Given BSM's model for the evolution of the stock price, the task at hand is now to develop replication 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 account.

A payoff H_T with maturity T is an $\mathcal{F}_T$-measurable non-negative random variable H_T: the measurability condition formalizes the idea that H_T must be determined by the path of S up to T, while non-negativity (which can be replaced by "bounded from below") ensures integrability of the payoff under all probability measures.

A good example is the call (1) with strike K and maturity T,

$$H_T(\omega) := C_T(\omega) = \Big(S_T(\omega) - K\Big)^+\ . \tag{13}$$

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 many shares we will hold at any time t and β tells us how much

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

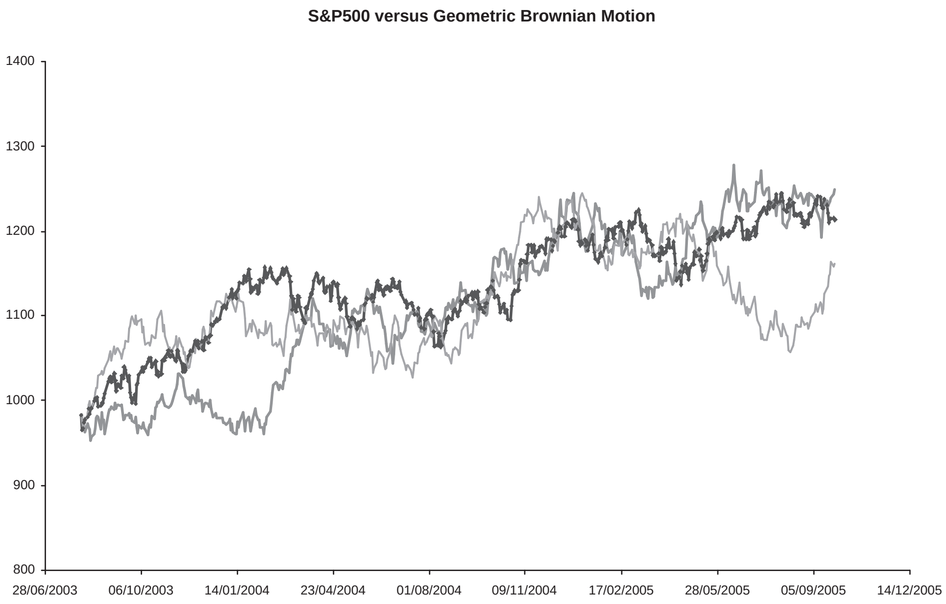

Figure 2: The graph shows a few paths of a geometric Brownian motion (with $\sigma = 15\%$) alongside historical S&P 500 prices. Which one is the index?

money we will invest in the cash account. In other words, the value of our portfolio of shares and cash account at time $t \in [0,T]$ is given as

$$U_t(\omega) := \Delta_t(\omega) S_t(\omega) + \beta_t(\omega) B_t \ .$$

If φ moreover replicates H_T, then this implies that

$$H_T(\omega) = U_T(\omega) \ .$$

However, we do not want to allow all possible strategies (Δ, β). Indeed, all costs which may incur during the hedging should be covered by the initial cost of the hedge, $U_0 = \Delta_0 S_0 + \beta_0$, and no further cash injection should be required after the inception of the contract. Such a strategy is called *self-financing*. Mathematically, this means that the change value of the portfolio over time is the result of the change of the value of stock and cash account:

$$dU_t(\omega) = \Delta_t(\omega)\, dS_t(\omega) + \beta_t(\omega)\, dB_t \ , \tag{14}$$

i.e.

$$U_t(\omega) = U_0 + \int_0^t \Delta_u(\omega)\, dS_u(\omega) + \int_0^t \beta_u(\omega)\, dB_u \ . \tag{15}$$

If then φ replicates H_T, it is appropriate to call

$$H_t(\omega) := U_t(\omega)$$

the *price* of H_T at time $t \in [0,T]$.

We have transformed the task of finding a hedging strategy to the task of determining U and Δ such that $H_T = U_T$: once they are specified, we can always make our portfolio self-financing by borrowing or investing an appropriate amount of money in the cash account. Formally, this means that β is given as the ratio

$$\beta_t(\omega) = \frac{1}{B_t}\Big(U_t(\omega) - \Delta_t(\omega) S_t(\omega) \Big) \ . \tag{16}$$

It is relatively straight-forward to see that this yields via (15) the representation

$$\hat{U}_t(\omega) = \hat{U}_0 + \int_0^t \Delta_t(\omega)\, d\hat{S}_t(\omega) \ . \tag{17}$$

in terms of the *discounted values*

$$\hat{S}_t(\omega) := B_t^{-1} S_t(\omega) \ , \quad \hat{B}_t \equiv 1 \quad \text{and} \quad \hat{U}_t(\omega) := B_t^{-1} U_t(\omega) \ .$$

The discounted stock price process $\hat{S}$ satisfies the SDE

$$d\hat{S}_t(\omega) = (\mu - r)\hat{S}_t(\omega)\, dt + \sigma \hat{S}_t(\omega)\, dW_t^{\mathbb{Q}}(\omega) \quad S_0 \in \mathbb{R}^+ \ . \tag{18}$$

We will also write $\hat{H}_t := \hat{U}_t$ for $t \in [0,T]$ if $\hat{U}_T = \hat{H}_T$.

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

Black, Scholes, Merton

To find now the required Δ-hedge in BSM's model, one can make use of the *representation property* of Brownian motion: this property says that every sufficiently integrable random variable $\hat{H}_T$ can be uniquely written as

$$\hat{H}_T(\omega) = \hat{H}_0^{\mathbb{Q}} + \int_0^T \phi_s^{\mathbb{Q}}(\omega)\, dW_s^{\mathbb{Q}}(\omega) \tag{19}$$

for a integrable process $\phi^{\mathbb{Q}}$ and an initial value $\hat{H}_0^{\mathbb{Q}}$.[9]

This expression does not help us very much yet, since the Brownian motion $W^{\mathbb{Q}}$ itself cannot be traded in the market. However, given (18) 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

$$\hat{H}_T(\omega) = \hat{H}_0^{\mathbb{Q}} + \int_0^T \frac{\phi_t^{\mathbb{Q}}(\omega)}{\sigma \hat{S}_t(\omega)}\, d\hat{S}_t(\omega) \quad - \int_0^T \frac{\phi_t^{\mathbb{Q}}(\omega)(\mu - r)}{\sigma} dt \ . \tag{20}$$

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 (recall that $d\hat{B}_t \equiv 0$). 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 (20) with (17).

The solution to this hurdle comes in the form of a further fundamental theorem of stochastic calculus: the *Cameron-Martin-Girsanov theorem.* 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[10] probability measure

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

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

[10] Two measures $\mathbb{P}$ and $\mathbb{Q}$ are equivalent on some σ-algebra $\mathcal{A}$, if their impossible events are the same, i.e. if $\mathbb{P}[A] = 0$ if and only if $\mathbb{Q}[A] = 0$ for $A \in \mathcal{A}$.

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

$$W_t^{\mathbb{P}}(\omega) := W_t^{\mathbb{Q}}(\omega) - \int_0^t \lambda_s(\omega)\, ds \tag{21}$$

is a Brownian motion.

How does this help in our situation? Well, equations (21) and (18) show that in terms of this Brownian motion, the discounted stock price follows the SDE

$$d\hat{S}_t(\omega) = \hat{S}_t(\omega)(\mu - r + \sigma\lambda_t(\omega))\, dt + \hat{S}_t(\omega)\, \sigma\, dW_t^{\mathbb{P}}(\omega)\ . \tag{22}$$

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

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

Thus we have gained an additional degree of freedom by our choice of λ. Since our aim is to nullify the right hand integral above, we chose the measure $\mathbb{P}$ given by the process

$$\lambda_t := (r - \mu)/\sigma\ , \tag{23}$$

(called the "market price of risk") for which we get the desired replication

$$\hat{H}_T(\omega) = \hat{H}_0^{\mathbb{P}} + \int_0^T \Delta_t(\omega)\, d\hat{S}_t(\omega) \quad \text{with} \quad \Delta_t(\omega) := \frac{\phi_t^{\mathbb{P}}(\omega)}{\sigma \hat{S}_t(\omega)}\ . \tag{24}$$

With β defined in (16), this gives us the desired hedging strategy (Δ, β).

2.3 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 it the *fair price* of H_T. But how can we efficiently calculate $H_0^{\mathbb{P}}$?

To this end, the notion of a *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":

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

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 some regularity conditions, also a martingale.[11] That implies in particular that $\mathbb{E}[N_t] = N_0 = 0$ for all $t \in [0, T]$.

In the BSM-model, we note that (22) becomes

$$d\hat{S}_t(\omega) = \hat{S}_t(\omega)\, \sigma\, dW_t^{\mathbb{P}}(\omega)\ , \tag{25}$$

so "under $\mathbb{P}$", $\hat{S}$ has no drift. In integral form, (25) is written as

$$\hat{S}_t(\omega) = S_0 + \int_0^t \hat{S}_s(\omega)\, \sigma\, dW_s^{\mathbb{P}}(\omega) \tag{26}$$

and for some payoff H_T, we have

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

[11] In general, the integral is just a local martingale; details on the subject can be found in Revuz/Yor [RY99].

That means that both the discounted stock and the discounted price processes $\hat{H}$ of all suitably integrable payoffs are martingales.[12] As a consequence, the price of H_T is given as

$$H_0^{\mathbb{P}} = \mathbb{E}_{\mathbb{P}}\left[\hat{H}_T\right] = \frac{1}{B_T}\mathbb{E}_{\mathbb{P}}\left[H_T\right] , \tag{27}$$

just as in (7) for our discrete model. At some later time t, the martingale property equally yields

$$\hat{H}_t^{\mathbb{P}} = \mathbb{E}_{\mathbb{P}}\left[\hat{H}_T \,\middle|\, \mathcal{F}_t\right] . \tag{28}$$

(Note that our discussion also implies that each martingale M can be written as an integral of W.) These quantities can now be computed using either analytical or numerical methods [BFGLMO99]. For more complicated models, we will usually resort to Monte-Carlo or finite difference schemes to evaluate (27).

In case of BSM's model, we can 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

$$\hat{S}_T = S_0 \exp\left\{\sigma W_T^{\mathbb{P}} - \frac{1}{2}\sigma^2 T\right\} . \tag{29}$$

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

$$H_0 = S_0\mathcal{N}(d^+) - e^{-rT}K\mathcal{N}(d^-) \tag{30}$$

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 cumulative standard normal distribution function).

2.4 The Fundamental Theorem of Asset Pricing

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

According to Harrison and Pliska [HP81], such markets can be characterized in discrete time[13] as follows:

Theorem 2.1 (First Fundamental Theorem of Asset Pricing) *A market is free of arbitrage if and only if there exists a measure $\mathbb{P}$ equivalent to the original measure $\mathbb{Q}$ such that all tradable assets are 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 $\mathbb{P}$). Clearly, $\mathbb{P}$ does not need to be unique. In this case, the market is called *incomplete* and there is no unique price for some payoffs. In our discussion later on, we will nonetheless be able to "complete" the market by considering further tradable hedging instruments. Intuitively, we need one linearly independent hedging instrument for every source of randomness.[14]

[12] The regularity conditions hold for the BSM-model.

[13] In continuous time, the results are more complex and out of the scope of this article. See Delbaen/Schachermayer [DS94] for details.

[14] Not all markets can be completed; genuinely random jump processes produce markets where replication with finitely many instruments is not possible.

2.5 The Hedging Strategy in Markovian models

We have shown with (28), that we can compute a price $H_t^{\mathbb{P}}$ of a contingent claim in a complete market at every time $t < T$ as a conditional expectation. Let us now turn to the actual computation of the hedging ratio $\Delta^{\mathbb{P}}$.

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

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

For example, it follows from (29) that BSM's stock price can be written as

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

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

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

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 on $S_T = \hat{S}_T B_T$ 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 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 (28) and the Markov property of $\hat{S}$ that there exists a *price-function* h such that

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

Usually, h is smooth enough, so we can apply Itô's formula (12) and get

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

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

We also know that (24) 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:

$$\Delta_t^{\mathbb{P}}(\omega) \equiv \partial_S h\big(t, \hat{S}_t(\omega)\big) \tag{31}$$

(the "Delta" of an option is the sensitivity of the price of an option to a change in the stock price). We can also imply that the remaining terms in the above equation must cancel. This yields the following PDE for the price function h:

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

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

The consequences of relation (31) are remarkable: it means that as long as our underlying model is largely in line with the actual market, we can compute a reliable hedging strategy for any payoff by differentiating its price function with respect to 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 was widely 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 day.

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 the hedging of more complex, so-called *exotic* payoffs.

3.1 The End of BSM: The Implied Volatility Skew

Why is 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 we know the interest rates r, then we can invert the BSM call price function (30) 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 $\hat{\sigma}(T, K)$. (Note that using the concept of implied volatility does not imply that practitioners actually believes that the market follows BSM's model. Rather, it is a convenient way of parameterizing the liquid option prices. It also indicates the "level of risk" of a liquid instrument.)

If we compute the implied volatility at a fixed maturity for a range of strikes, we will find a picture similar to the one displayed in figure 3.

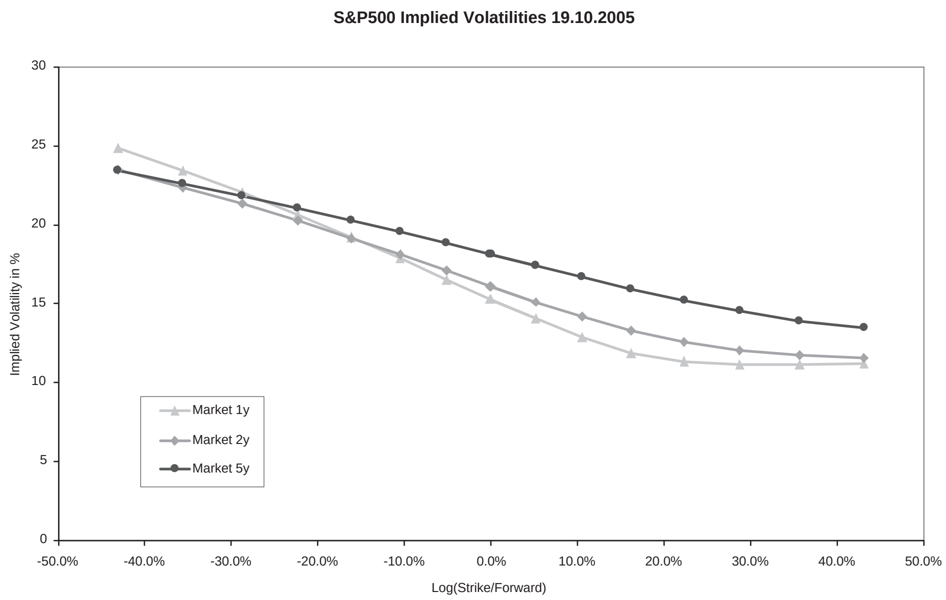

Figure 3: The graph shows S&P 500 implied volatilities for a few maturities. Implied volatilities tend to be much higher on the downside.

If "the market" behaved like a BSM model, this curve would be 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.

The reason is that the stock price usually gets more volatile when it falls: investors are more uneasy about a drop in a share price than about a rise. The effect is that "the volatility" itself is stochastic and instantaneously anti-correlated to the stock price. As we will see, this implies that we cannot replicate a payoff by purely trading in the stock: there are additional risk-factors in the economy which need to be hedged by other means. Indeed, the presence of listed options allows us to do just that. Moreover, under the assumption of an arbitrage-free market, the Fundamental Theorem of Asset Pricing proves that there is some measure $\mathbb{P}$ such that the stock price and all traded options prices are (local) martingales.

Since the options are themselves tradable assets, they become available as hedging instruments. 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.[15]

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". We assume now that the underlying probability space $(\Omega, \mathcal{A}, \mathbb{P})$ with martingale measure $\mathbb{P}$ supports 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 model the stock and other quantities directly under a martingale measure $\mathbb{P}$.[16]

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 non-negative payoff H_T measurable with respect to $\mathcal{F}_T$ can be written in terms of a sum of integrals as

$$\hat{H}_t(\omega) = \hat{H}_0 + \sum_{j=1}^n \int_0^t \varphi_s^j(\omega)\, dW_s^j(\omega) \ . \tag{32}$$

As we mentioned before, we can also write X in such a form, hence there is a random vector ρ with $dX_t(\omega) = \sum_{j=1}^n \rho_t^j(\omega)\, dW_t^j(\omega)$ such that

$$d\hat{S}_t(\omega) = \sqrt{\zeta_t(\omega)}\, \hat{S}_t(\omega) \sum_{j=1}^n \rho_t^j(\omega)\, dW_t^j(\omega) \tag{33}$$

In a general *stochastic volatility model*, we specify the dynamics of $\zeta = (\zeta_t)_{t\in[0,T^*]}$ in a way which we hope yields reasonable dynamics for the stock price and any reference instruments.

The best-known classic stochastic volatility model is probably that of Heston [H93]. In this two-factor model ($n = 2$), the stochastic variance process is the solution to the SDE

$$d\zeta(\omega) = \kappa(\theta - \zeta_t(\omega))\, dt + \nu\sqrt{\zeta_t(\omega)}\, dW_t^1(\omega) \tag{34}$$

for constants $\kappa > 0$, $\theta > 0$, $\nu > 0$ and starting at $\zeta_0 > 0$. 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" $\rho \in (-1, 1)$

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

[16] For notational convenience, we will therefore omit the notion of the underlying measure $\mathbb{P}$ for the remainder of the article.

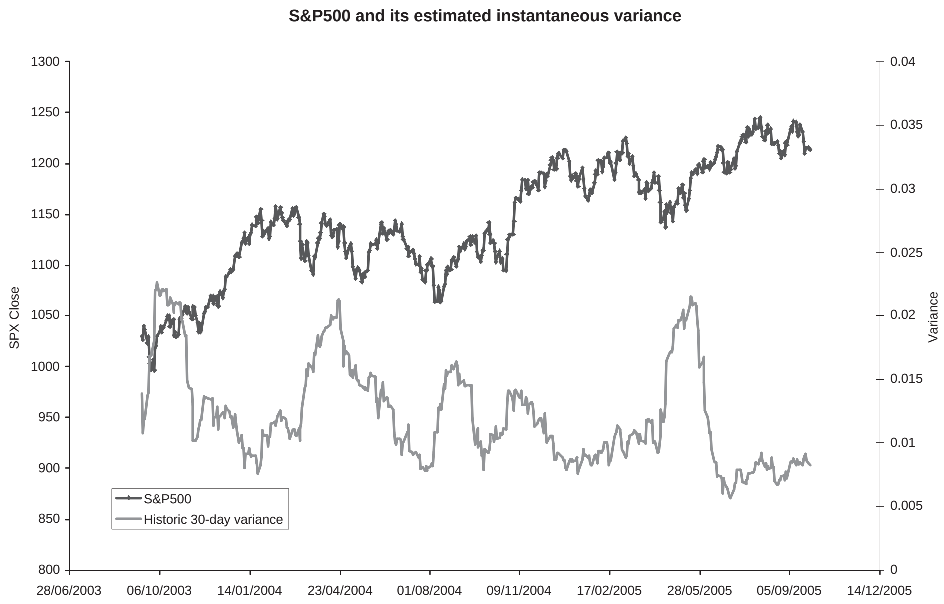

Figure 4: The graph shows S&P 500 and its estimated 30-day variance.

This process has a number of properties which can also be observed in real markets: the first is "mean-reversion", i.e. in the absence of random shocks it has an exponential decay with reversion speed κ towards the long-variance level θ. The second is a level-dependent "volatility of volatility" $\nu\sqrt{\cdot}$ which implies that ζ is more volatile if the overall volatility level is high. Finally, the correlation parameter governs the instantaneous co-movement of the stock and the variance process. It is typically negative, say, -0.7, which means that if the stock moves down, variance increases. As mentioned, this behavior is also present in real markets and it also implies that the implied volatility of an option priced with this model exhibits the desired skew.

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, 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 minimized. 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, in order to implement the parameter-calibration using a numerical minimization scheme, it is necessary that the computation of the expectation of the option payoffs be reasonably quick.

Note that the approach also shows that in this sense the original "historical" measure $\mathbb{Q}$ does not play a role in the implementation of a model if its parameters are determined via calibration. This implies that the value of a contract for any institution that hedges its option risk will be different from its value for a investor who is optimizing his investment according to his risk preference. Therefore, a price offered by a financial institution can be "fair" but still be a good investment for a particular investor (recall the seemingly cheap price of the call in the initial example on page 3).[17]

[17] Refer to Föllmer/Schied for details on preferences and related concepts.

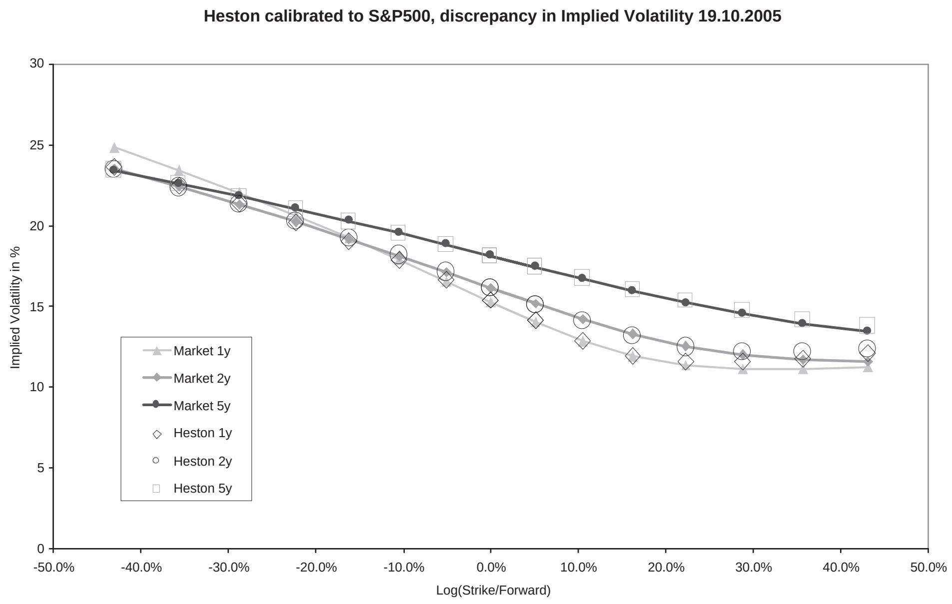

Figure 5: The result of calibrating Heston to S&P 500 option prices. The graph shows the market implied volatilities and those from the calibrated model (the calibrated parameters are $\zeta_0^2 = 0.15$, $\theta^2 = 0.34$, $\kappa = 0.13$, $\rho = -0.76$ and $\nu = 0.23$).

Hedging in General Stochastic Volatility Models

What happens now to the 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 $\zeta_t^1 > 0$). As a first step, we can write

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

which, using equation (32), gives us the desired hedging result

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

just as in (24).

However, it is also intuitively clear that this will not work if $n > 1$: there are more sources of randomness than market instruments in which we can invest. The market with only the stock and the bank account considered as tradable instruments is *incomplete* because we cannot replicate all possible payoffs. This is where the liquid options come in: if in addition to the stock, $(n-1)$ "reference instruments" with discounted price processes $\hat{M}^2, \ldots, \hat{M}^n$ are traded in the market, then we should be able to hedge the remaining uncertainty in the market.

Let $\hat{M}^1 := \hat{S}$. The first step is to use the representation property of W for the vector M. This gives us a random "volatility matrix" $\eta = (\eta^{i,j})_{i,j=1,\ldots,n}$ such that

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

To recover the vector W we need to be able to invert the matrix η. This confirms the intuition that the liquid options in consideration need to be linearly independent. It also means that if we want to recover

W up to some T, none of the options is allowed to expire: 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, (32) shows that each payoff H can now be replicated as

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

by dynamically trading in stock and reference instruments using the strategy $\varphi = (\Delta^1, \dots, \Delta^n\ \beta)$; the investment strategy for the cash account is once more determined by the self-financing requirement. According to (16), this means

$$\beta_t(\omega) = \frac{1}{B_t}\Big(H_t(\omega) - \sum_{j=1}^{n} \Delta_t^j(\omega) M_t^j(\omega) \Big)\ . \tag{35}$$

The 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 prices of the market instruments,

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

we can compute, just as in (31), the hedging ratios as

$$\Delta_t^j(\omega) \equiv \partial_{M^j} h\big(t; M_t^1(\omega), \dots, M_t^n(\omega)\big)\ . \tag{36}$$

3.3 Hedging Options on Realized Variance

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

From a theoretical point of view, this is satisfactory. However, from a practical point of view we now have to find some model which, on one hand, provides a good fit to the market prices of our reference instruments and, on the other hand, 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 particularly hot topic these days, we shall use these 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

$$\sum_{i=1}^{n} \big(\log S_{t_i}(\omega) - \log S_{t_{i-1}}(\omega) \big)^2 \tag{37}$$

where $0 = t_0 < \cdots < t_n = T$ are the business days up to T.[18] As in (11), this quantity will converge for $n \uparrow \infty$ to the quadratic variation of the logarithm of the stock price, and it is common to assume that realized variance is indeed measured using

$$V_T^T(\omega) := \langle \log S \rangle_T(\omega) = \int_0^T \zeta_t(\omega)\, dt\ . \tag{38}$$

The most basic product which is based on realized variance is a so-called *variance swap*. A variance swap pays V_T^T/T in exchange for a fixed cash strike K; since the strike K is fixed, we will ignore it in the

[18] Usually, the sum (37) is scaled by the constant $252\, T/n$ in order to annualize the variance with respect to the number of business days per year, normalized to 252. For details on variance swap contracts, see Demeterfi et al [DDKZ99].

sequel. We denote the price at time t of a variance swap with maturity T by V_t^T. It is as usual given as the expectation of the payoff under the risk neutral measure $\mathbb{P}$:

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

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 their "volatility" i.e. as $\sqrt{V_0^T/T}$. This notation is also used in the graphs below.

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,

$$H_T(\omega) := \left(\frac{1}{T} \int_0^T \zeta_t(\omega)\, dt - K \right)^+ \tag{39}$$

It makes sense to buy such a call if the investor believes that the variance of the stock price will go up (recall that the investor is using the historical measure $\mathbb{Q}$): 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 reasonably assume that the variance will pick up some time soon. If this is not the case, we can sell instead a put,

$$\left(K - \frac{1}{T} \int_0^T \zeta_t(\omega)\, dt \right)^+ \ .$$

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

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

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

$$V_t^T(\omega) = \int_0^t \zeta_s(\omega)\, ds + \mathbb{G}(Z_t(\omega); T-t) \ . \tag{40}$$

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 reasonably. 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 for all T. If we assume that Z is the solution to an SDE of the form

$$\begin{pmatrix} dZ_t^1 \\ \vdots \\ dZ_t^m \end{pmatrix} = \begin{pmatrix} \mu^1(Z_t) \\ \vdots \\ \mu^m(Z_t) \end{pmatrix} dt + \begin{pmatrix} \varsigma^{1,1}(Z_t) & \cdots & \varsigma^{1,n}(Z_t) \\ \vdots & \ddots & \vdots \\ \varsigma^{m,1}(Z_t) & \cdots & \varsigma^{m,n}(Z_t) \end{pmatrix} \begin{pmatrix} dW_t^1 \\ \vdots \\ dW_t^n \end{pmatrix} , \tag{41}$$

then it can be shown using Itô's formula that under some regularity conditions (see [B05]), the function $\mathbb{G}$ and the coefficients μ and ς must satisfy

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

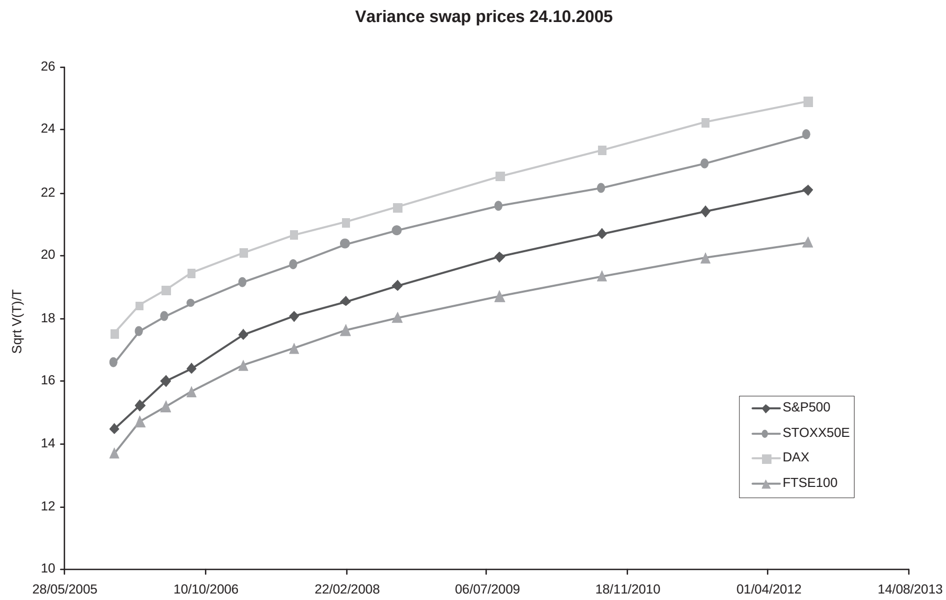

Figure 6: Variance swap market prices on various indices, quoted in "volatility". Note that the general shapes of the curves are quite similar.

With standard 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 price processes of all variance swaps are proper martingales. Moreover, differentiating (40) with respect to the maturity T of the variance swap shows that

$$\zeta_t(\omega) = \partial_\tau \mathbb{G}(Z_t(\omega), 0)\ .$$

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

Double Mean-Reversion

An good example of the above approach is the "Double mean-reversion" 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.[19] The function $\mathbb{G}$ can be reduced to the respective functional $\mathbb{H}$ given by Heston's model (34) by setting $z_2 = z_3$,

$$\mathbb{H}(z;\tau) := z_3\tau + (z_1 - z_3)\frac{1-e^{-\kappa\tau}}{\kappa}\ .$$

During the course of the last two years, $\mathbb{G}$ has shown to fit well to observed market prices. Figure ref-fig:vswapfit shows an example[20]

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

[20] In theory, the constants κ and c have to be calibrated only once per underlying, but in practise they are updated upon strong movements of the market.

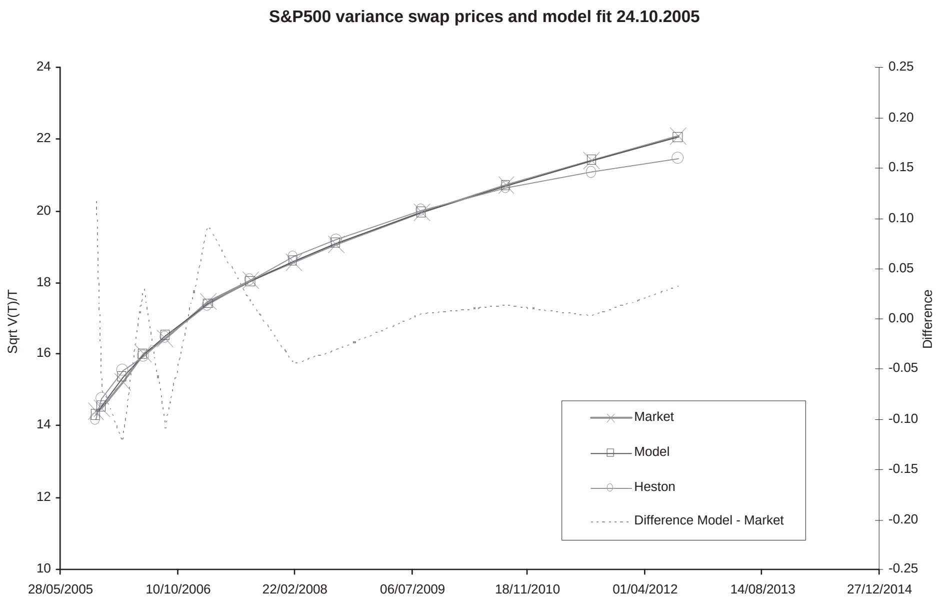

Figure 7: Fit of our model function $\mathbb{G}$ and Heston's function $\mathbb{H}$ to S&P 500 variance swap market prices. The scale of differences between the calibrated function $\mathbb{G}$ and the market is indicated on the right hand axis.

A consistent process Z with SDE (41) has the general form

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

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 idea is that the stochastic variance process ζ is again mean-reverting. In extension to Heston's model (34), the level of mean-reversion is itself stochastic and mean-reverting to just another stochastic long-term level. This parameterizes the variance swap prices in terms of a short, medium and long level of variance. Such an approach is justified given historic data, which often show a mean-reversion-like behaviour around a moving mean (it should be noted that the improved flexibility of the model comes at a cost: its numerical implementation, in particular of the calibration itself, is far more challenging than for Heston's model for which fast and accurate methods are available).

Once this model is implemented, we can finally hedge our options on variance. Consider for example a call on variance (39) and recall (36). 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

$$H_t(\omega) = \tilde{h}\big(t; S_t(\omega), Z_t^1(\omega), \ldots, Z_t^3(\omega)\big) \ . \tag{42}$$

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

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

is invertible with smooth inverse $\tilde{G}^{-1}$. Hence, we can write (42) as

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

and therefore as

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

This gives us the hedging strategy

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

according to (36). The computation of these ratios can be implemented using numerical differentiation (usually, a simple centered difference with a specified width is sufficient).

4 Conclusions

We have introduced basic ideas of mathematical finance and shown that in contrast to an insurance company, a financial institution which trades options on stocks can hedge itself and try to replicate the target payoffs. It is an appealing and intuitive concept to characterize the fair price of a payoff as the cost of the implementation of this replication strategy since it does not involve any risk. It may mean that we have to consider further reference instruments in addition to the stock – in our example, we have used variance swaps alongside the underlying stock to hedge so-called options on realized variance. It is a great challenge today to develop models which capture the relevant behavior of the markets; advanced mathematical tools are applied on a daily basis to develop to ever more realistic and complex models.

Of course, the framework presented here is idealistic: in reality, the dynamics of the hedging instruments are not easy to determine and we have to pay transaction costs and taxes. Extraordinary market events often make it impossible to execute a hedging strategy and so forth. Many of these issues have been addressed elsewhere or are the subject of current research.

However, the success of the financial industry also indicates that the general approach of pricing by replication works and that it is a powerful tool to manage the risk of exotic payoffs. Mathematics is a key contributor to this success.

References

[B00] *L.Bachelier:*
"Theórie de la spéculation. Ann.Sci.École Norm Sup.17, pp.21-86, 1900

[BS73] *F.Black, M.Scholes:*
"The Pricing of Options and Corporate Liabilities", Journal of Political Economy, 81, pp. 637-59, 1973

[BFGLMO99] *O.Brockhaus, A.Ferraris, C.Gallus, D.Long, R.Martin, M. Overhaus:*
Modelling and hedging equity derivatives. Risk, 1999

[B05] *H.Buehler:*
"Consistent Variance Curve Models", working paper, June 2005
http://ssrn.com/abstract=687258

[DS94] *F.Delbaen, W.Schachermayer:*
"A General Version of the Fundamental Theorem of Asset Pricing", Mathematische Annalen 300, p. 463-520 (1994)

[DDKZ99] *Demeterfi, K., Derman, E., Kamal, M., Zou, J.:* "More Than You Ever Wanted To Know About Volatility Swaps", Quantitative Strategies Research Notes, Goldman Sachs, 1999

[D40] *W.Döblin:*
"Sur l'equation de Komogoroff", C.R. Ser. I 331, 1059-1102, 1940

[F81] *H.Föllmer:*
"Calcul d'Itô sans probabilités.", Seminaire de probabilites XV, Univ. Strasbourg 1979/80, Lect. Notes Math. 850 (1981), 143-150.

[FS04] *H.Föllmer, A.Schied:*
Stochastic Finance. Second Edition. de Gruyter, 2004

[HK79] *J.Harrison, D.Kreps:*
"Martingales and Arbitrage in Multiperiod Securities Markets", Journal of Economic Theory, 20, pp. 381-408, 1979

[HP81] *J.Harrison, S.Pliska:*
"Martingales and Stochastic Integrals in the Theory of Continuous Trading", Stochasic Processes and their Applications, 11, pp. 215-60, 1981

[H93] *S.Heston:*
"A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options", 1993

[I44] *K.Itô:*
"Stochastic integral", Proc. Imperial Acad. Tokyo, Vol. 20, pp.519-524, 1944

[M73a] *R.Merton:*
"An intertemporal capital asset pricing model", Econometrica, Vol. 41, pp.867-888, 1973

[M73b] *R.Merton:*
"The theory of rational option pricing", Bell J. Econ. Manag. Sci.4, pp.141-183, 1973

[P04] *P.Protter:*
Stochastic Integration and Differential Equations. Second Edition, Springer 2004

[RY99] *D.Revuz, M.Yor:*
Continuous martingales and Brownian motion. Third edition, Springer 1999

[S65] *P.Samuelson:*
"Rational Theory of Warrant Pricing", Industrial Management Review 6, pg.13-31 (1965)