1

Introduction

In the Black–Scholes option pricing model the stock price dynamics are assumed to follow an Itô process with constant characteristics. This key hypothesis, dating from a 1965 paper by Paul Samuelson, adapts ideas from a remarkable doctoral thesis by the French mathematician Louis Bachelier in 1900. The model makes various simplifying assumptions about the market, not all of which are borne out by market data. Nonetheless, the Black–Scholes prices of European derivatives provide benchmarks against which prices quoted in the market can be judged.

We turn first to a description of the continuous-time price processes for the assets that comprise the basic single-stock Black–Scholes model.

1.1 Asset dynamics

The market model contains two underlying securities.

- The **risk-free** asset (money-market account), described by a deterministic function

  $$dA(t) = rA(t)dt,$$

  with $A(0) = 1$ (for convenience), where $r > 0$ is the risk-free rate.

This is an ordinary differential equation $A'(t) = rA(t)$ but for consistency with stock prices, which are assumed to be Itô processes, we use differential notation. The equation has a unique solution:

$$A(t) = \mathrm{e}^{rt}.$$

- The **risky** asset, thought of as a stock, is represented by an Itô process of the form

$$dS(t) = \mu S(t)dt + \sigma S(t)dW(t), \tag{1.1}$$

with $S(0)$ given, where we call $\mu \in \mathbb{R}$ the **drift**, and $\sigma > 0$ the **volatility**, of the stock price S.

The sign of σ is actually irrelevant. If σ is negative then we change W to $-W$ and we have an equation with positive σ but with respect to $(-W)$, which is again a Wiener process. The probability space underlying W will be denoted by $(\Omega, \mathcal{F}, P)$, and the associated filtration is given by $\mathcal{F}_t^S = \sigma(S(u) : u \leq t\})$. Writing out (1.1) we see that

$$S(t) = S(0) + \mu \int_0^t S(u)du + \sigma \int_0^t S(u)dW(u).$$

The stochastic differential equation (1.1) has a unique solution since the coefficients are Lipschitz with linear growth:

$$\mu S(t) = a(t, S(t)), \quad a(t,x) = \mu x,$$
$$\sigma S(t) = b(t, S(t)), \quad b(t,x) = \sigma x,$$

so that

$$|a(t,x) - a(t,y)| \leq |\mu||x-y|,$$
$$|b(t,x) - b(t,y)| \leq \sigma|x-y|,$$

linear growth being obvious, and we can apply the existence and uniqueness theorem for stochastic differential equations, proved in [SCF] as Theorem 5.8.

We can determine the solution immediately: it takes the form

$$S(t) = S(0)\exp\{\mu t - \frac{\sigma^2}{2}t + \sigma W(t)\}. \tag{1.2}$$

Exercise 1.1 Show that this process solves (1.1).

As the solution is unique, S given by (1.2) is the unique solution of (1.1). Note that the filtration $\mathcal{F}^S$ governing the random fluctuations in the stock price S coincides with the natural filtration of W, where $\mathcal{F}_t^W = \sigma(W(u) : u \leq t)$ for each $t \in [0,T]$, since (1.2) shows that W is the only source of randomness in S.

Exercise 1.2 Find the probability that $S(2t) > 2S(t)$ for some $t > 0$.

Model parameters

To understand the role of the parameters μ, σ in this model we compute the expectation of $S(t)$. Recall that for a normally distributed random variable X with $\mathbb{E}(X) = 0$ we have

$$\mathbb{E}(\exp\{X\}) = \exp\{\frac{1}{2}\text{Var}(X)\}. \tag{1.3}$$

We apply this with $X = \sigma W(t)$, so that $\text{Var}(X) = \sigma^2 t$ (we write the expectation of X with respect to P simply as $\mathbb{E}(X)$ rather than $\mathbb{E}_P(X)$ when there is no danger of confusion):

$$\begin{aligned}\mathbb{E}(S(t)) &= S(0)\mathbb{E}(\exp\{\mu t - \frac{1}{2}\sigma^2 t + \sigma W(t)\}) \\ &= S(0)\exp\{\mu t - \frac{1}{2}\sigma^2 t\}\mathbb{E}(\exp\{\sigma W(t)\}) \\ &= S(0)\exp\{\mu t\}.\end{aligned}$$

Clearly, if $\mu = 0$ then the expectation of $S(t)$ is constant in time.

The expression for $\mathbb{E}(S(t))$ gives μ as the (annualised) logarithmic return of the expected price

$$\mu = \frac{1}{t}\ln\frac{\mathbb{E}(S(t))}{S(0)}, \tag{1.4}$$

which should not be confused with the expected (annualised) logarithmic return

$$\frac{1}{t}\mathbb{E}(\ln\frac{S(t)}{S(0)}) = \frac{1}{t}\mathbb{E}(\mu t - \frac{\sigma^2}{2}t + \sigma W(t)) = \mu - \frac{\sigma^2}{2}.$$

The variance of the return is

$$\begin{aligned}\text{Var}(\mu t - \frac{\sigma^2}{2}t + \sigma W(t)) &= \text{Var}(\sigma W(t)) \quad \text{(adding a constant has no impact)} \\ &= \sigma^2 t \quad (\text{since } \text{Var}(W(t)) = t).\end{aligned}$$

A natural question emerges of how to find these parameters given some past stock prices. The formula (1.4) suggests taking average prices as the proxy for the expected price, but the accuracy of this is poor, according to statistical theory.

Much more effective is the approximation of volatility provided, for instance, by the following scheme. Consider the process

$$\ln S(t) = \ln S(0) + (\mu - \frac{1}{2}\sigma^2)t + \sigma W(t),$$

which is an Itô process with constant characteristics. Its quadratic variation is equal to $\sigma^2 t$ (see [SCF]) and for a partition of $[0, t]$ given by $0 = t_1 < \cdots < t_n = t$, with small mesh $\max |t_{k+1} - t_k|$, we have

$$\sum_k (\ln S(t_{k+1}) - \ln S(t_k))^2 \approx \sigma^2 t.$$

Hence if the times t_k represent past instants at which we know the prices, then we can take

$$\sigma = \sqrt{\frac{1}{t} \sum \ln \frac{S(t_{k+1})}{S(t_k)}}$$

as our estimate of the volatility coefficient, a positive number called the sample volatility.

Exercise 1.3 Find the formula for the variance of the stock price: $\mathrm{Var}(S(t))$.

Exercise 1.4 Consider an alternative model where the stock prices follow an Ornstein–Uhlenbeck process: this is a solution of $dS_1(t) = \mu_1 S_1(t)dt + \sigma_1 dW(t)$ (see [SCF]). Find the probability that at a certain time $t_1 > 0$ we will have negative prices: i.e. compute $P(S_1(t_1) < 0)$. Illustrate the result numerically.

Exercise 1.5 Allowing time-dependent but deterministic σ_1 in the Ornstein–Uhlenbeck model, find its shape so that $\mathrm{Var}(S(t)) = \mathrm{Var}(S_1(t))$.

Exercise 1.6 Let L be a random variable representing the loss on some business activity. Value at Risk at confidence level a is defined as $v = \inf\{x : P(L \leq x) \geq a\}$. Compute v for $a = 95\%$, where L is the loss on the investment in a single share of stock purchased at $S(0) = 100$ and sold at $S(T)$ with $\mu = 10\%$, $\sigma = 40\%$, $T = 1$.

1.2 Methods of option pricing

We consider a possible line of attack for pricing options in the Black–Scholes market. To make progress we impose various assumptions, and in doing so we survey the range of tasks required to solve the option pricing problem.

Recall that a European derivative security is a contract where the seller promises the buyer a random payment H at some prescribed future time T, called the exercise time. In our pricing model H is a random variable defined on a probability space $(\Omega, \mathcal{F}, P)$ supporting the Wiener process W, equipped with its natural filtration $(\mathcal{F}_t^W)_{t\in[0,T]}$, and we may assume in addition that $\mathcal{F}_T^W = \mathcal{F}$. The natural filtration of W coincides with the filtration generated by the Itô process $S = (S(t))_{t\in[0,T]}$ described above. We call S the underlying security – with S as defined above, $(\mathcal{F}_t^S)_{t\in[0,T]}$ is simply the natural filtration of W. This measurability is the only link between H and the underlying. If $H = h(S(T))$ for some Borel function h, the derivative security is path-independent and it of course satisfies the measurability condition, but not the other way round, since the σ-field $\mathcal{F}_T$ is generated by the entire price process, not simply by $S(T)$. A familiar path-independent security is the European call option with strike K, where $h(x) = (x - K)^+ = \max\{0, x - K\}$, so that the option payoff at expiry is $H = (S(T) - K)^+$.

Such a security is sold at time 0 and the first task we tackle is to find its price at that time – this is known as the option premium.

Risk-neutral probability approach

In the finite discrete-time setting discussed in [DMFM] the key assumption was the absence of arbitrage. This economic hypothesis was given mathematical form by the first fundamental theorem of asset pricing, which showed that the No Arbitrage Principle was equivalent to the existence of a measure Q, with the same null sets as P, under which the discounted price process is a martingale. This result, together with the fact that the transform (or 'discrete stochastic integral') of a martingale is again a martingale, allowed us to identify the value process of a path-independent European derivative with that of a 'replicating' trading strategy involving only stocks and the money market account.

In continuous-time models the analogue of the first fundamental theorem is rather sophisticated and we shall not pursue it directly, but will instead reformulate the No Arbitrage Principle in more detail later. For our present purposes we state three assumptions that suffice to explain the approach to pricing that will enable us to derive the Black–Scholes formula and related results. This section is intended simply to give the flavour of the arguments that will be deployed.

Assumption 1.1
There exists a pair (x, y) of processes, adapted to the filtration $(\mathcal{F}_t^S)_{t\in[0,T]}$, producing portfolios consisting of holdings in the stock and the money market account, with values

$$V(t) = x(t)S(t) + y(t)A(t)$$

assumed to match the option payoff at maturity

$$V(T) = H$$

and therefore (x, y) is called a replicating strategy.

The condition we impose on the trading strategies employed is a natural continuous time analogue of the self-financing condition demanded of discrete time models, capturing the idea that changes in the values and holdings of assets are the sole drivers of changes of wealth, allowing no inflows or outflows of funds.

Assumption 1.2
There exists a replicating strategy satisfying the self-financing condition:

$$dV(t) = x(t)dS(t) + y(t)dA(t).$$

In the binomial model the construction of a risk-neutral probability was straightforward, in continuous time it will be quite involved and for the time being we impose it as follows.

Assumption 1.3
There exists a probability Q, with the same null sets as P, such that $\tilde{S}(t) = e^{-rt}S(t)$ and $\tilde{V}(t) = e^{-rt}V(t)$ are martingales with respect to Q and the filtration $(\mathcal{F}_t^S)_{t\in[0,T]}$.

A martingale has constant expectation so in particular $V(0) = \mathbb{E}_Q(\tilde{V}(T))$, hence

$$V(0) = \mathbb{E}_Q(e^{-rT}H),$$

which, as we show in Theorem 2.16, must be $H(0)$, the initial price of the derivative with payoff H, since in the case of inequality an arbitrage opportunity emerges: buy the cheap asset and sell the expensive one, invest the profit risk-free, so a riskless profit is maintained at maturity, due to replication, with some care needed to meet some admissibility conditions (the details can be found in Chapter 2).

The PDE approach

To develop an alternative pricing method, the replication condition is formulated in a stronger version: in addition to matching at maturity we assume that the entire process of option prices $H(t)$ is indistinguishable from the value process of the strategy. (Again, this is easily obtained in the discrete-time setting – see Theorem 4.40 in [DMFM].) We make two further assumptions.

Assumption 1.4
There is a self-financing strategy (x, y) such that the option value process can be written in the form

$$H(t) = x(t)S(t) + y(t)A(t).$$

The spirit of the next condition is that there exists a closed form formula for the option price, though we do not yet know its shape. An additional feature is that the price does not depend on the history of stock prices (at this point the reader should recall the Markov property discussed in [SCF]). This is only applicable to path-independent derivatives.

Assumption 1.5
The process $H(t)$ is of the form

$$H(t) = u(t, S(t)),$$

where the deterministic function $u(t, z)$ has continuous first derivative with respect to $t \in [0, T]$ and continuous first and second derivatives in $z \in \mathbb{R}$.

(We write $u(t,z)$ rather than $u(t,x)$, in order to avoid confusion with the use of x for the process of stock holdings in the trading strategy.)

Applying the Itô formula, we find that the process $H(t)$ is an Itô process and has the following stochastic differential

$$dH = \left(u_t + \mu S u_z + \frac{1}{2}\sigma^2 S^2 u_{zz}\right) dt + \sigma S u_z dW. \tag{1.5}$$

With the given form of $dS(t)$ and $dA(t)$ the self-financing condition reads

$$dH = (x\mu S + ryA)\, dt + x\sigma S\, dW. \tag{1.6}$$

Now use the fact that the representation of an Itô process is unique, so the characteristics on the right-hand sides of (1.5) and (1.6) must agree:

$$\begin{aligned} u_t + \mu S u_z + \frac{1}{2}\sigma^2 S^2 u_{zz} &= x\mu S + ryA, \\ \sigma S u_z &= x\sigma S. \end{aligned}$$

(For better readibility we sometimes omit the arguments of the functions and processes.)

The second line immediately gives

$$x(t) = u_z(t, S(t)),$$

which, inserted into the first line, reads

$$u_t(t, S(t)) + \frac{1}{2}\sigma^2 S^2(t) u_{zz}(t, S(t)) = ry(t)A(t),$$

providing a formula for the second component of the replicating strategy

$$y(t) = \frac{1}{rA(t)}\left(u_t(t, S(t)) + \frac{1}{2}\sigma^2 S^2(t) u_{zz}(t, S(t))\right).$$

The replication condition $H(t) = x(t)S(t) + y(t)A(t)$, with the above expressions for x and y inserted, gives

$$u(t, S(t)) = u_z(t, S(t))S(t) + \frac{1}{r}\left(u_t(t, S(t)) + \frac{1}{2}\sigma^2 S^2(t) u_{zz}(t, S(t))\right),$$

since $H(t) = u(t, S(t)))$ by our assumption. After replacing $S(t)$ by a general variable z this equation can be reorganised into

$$u_t(t,z) = -\frac{1}{2}\sigma^2 x^2 u_{zz}(t,z) - rzu_z(t,z) + ru(t,z) \quad \text{for } 0 < t < T,\ z \in \mathbb{R}.$$

Clearly, at time T the value $H(T)$ must agree with the option payoff so we impose the terminal condition

$$u(T,z) = h(z) \text{ for } z \in \mathbb{R}.$$

The above partial differential equation (PDE) is known as the Black–Scholes PDE and solving it will give us the function u, and so the pricing problem will have a sucessful conclusion. With some knowledge from the theory of PDEs one may show that this problem has a unique solution given by a closed-form expression.

In outlining this approach to the pricing problem we have made some powerful assumptions. It will turn out for the Black–Scholes model we are able to prove all these statements – the assumptions will be converted into theorems. We will find the expression for u as well, and then we will revisit the above PDE to see that this is indeed a solution.

Cambridge University Press
978-1-107-00169-5 - The Black–Scholes Model
Marek Capiński and Ekkehard Kopp
Excerpt
More information

2

Strategies and risk-neutral probability

We begin executing some of the elements of the programme outlined in the previous chapter. The first goal is to discuss the martingale properties of strategies. For this we must first construct a risk-neutral probability. Next we define self-financing strategies by analogy with the discrete case and observe that in the present setting pathologies can emerge which are absent in discrete-time models. This leads to conclusions about the class of trading strategies that are admissible in continuous-time models.

2.1 Finding the risk-neutral probability

We consider the discounted stock price process $\tilde{S}(t) = e^{-rt}S(t)$, which, by the definition of $S(t)$, becomes

$$\begin{aligned}\tilde{S}(t) &= \exp\{-rt\}S(0)\exp\left\{\mu t - \frac{1}{2}\sigma^2 t + \sigma W(t)\right\}\\ &= S(0)\exp\left\{(\mu - r)t - \frac{1}{2}\sigma^2 t + \sigma W(t)\right\}.\end{aligned}$$

To explore situations where this becomes a martingale with respect to the probability P, we compute conditional expectations.