MASTERING MATHEMATICAL FINANCE

The Black–Scholes Model

MAREK CAPIŃSKI
EKKEHARD KOPP

CAMBRIDGE

The Black–Scholes Model

MAREK CAPIŃSKI
AGH University of Science and Technology, Kraków, Poland

EKKEHARD KOPP
University of Hull, Hull, UK

CAMBRIDGE
UNIVERSITY PRESS

CAMBRIDGE UNIVERSITY PRESS
Cambridge, New York, Melbourne, Madrid, Cape Town,
Singapore, São Paulo, Delhi, Mexico City

Cambridge University Press
The Edinburgh Building, Cambridge CB2 8RU, UK

Published in the United States of America by Cambridge University Press, New York

www.cambridge.org
Information on this title: www.cambridge.org/9781107001695

First published 2012

Printed and bound in the United Kingdom by the MPG Books Group

A catalogue record for this publication is available from the British Library

Library of Congress Cataloguing in Publication data
Capiński, Marek, 1951–
The Black-Scholes model / Marek Capiński, Ekkehard Kopp.
p. cm. – (Mastering mathematical finance)
Includes index.
ISBN 978-1-107-00169-5 (hardback)
1. Options (Finance) – Prices – Mathematical models. I. Kopp, P. E., 1944– II. Title.
HG6024.A3C364 2013
332.64′53 – dc23 2012023180

ISBN 978-1-107-00169-5 Hardback
ISBN 978-0-521-17300-1 Paperback

Additional resources for this publication at www.cambridge.org/9781107001695

The Black–Scholes Model

The Black–Scholes option pricing model is the first, and by far the best-known, continuous-time mathematical model used in mathematical finance. Here, it provides a sufficiently complex, yet tractable, testbed for exploring the basic methodology of option pricing.

The discussion of extended markets, the careful attention paid to the requirements for admissible trading strategies, the development of pricing formulae for many widely traded instruments and the additional complications offered by multi-stock models will appeal to a wide class of instructors. Students, practitioners and researchers alike will benefit from the book's rigorous, but unfussy, approach to technical issues. It highlights potential pitfalls, gives clear motivation for results and techniques, and includes carefully chosen examples and exercises, all of which makes it suitable for self-study.

MAREK CAPIŃSKI has published over 50 research papers and eleven books. His diverse interests include mathematical finance, corporate finance and stochastic hydrodynamics. For over 35 years he has been teaching these topics, mainly in Poland and in the UK, where he has held visiting fellowships. He is currently Professor of Applied Mathematics at AGH University of Science and Technology in Kraków, Poland, where he established a Master's programme in mathematical finance.

EKKEHARD KOPP is Emeritus Professor of Mathematics at the University of Hull, UK, where he taught courses at all levels in analysis, measure and probability, stochastic processes and mathematical finance between 1970 and 2007. His editorial experience includes service as founding member of the Springer Finance series (1998–2008) and the Cambridge University Press AIMS Library Series. He has taught in the UK, Canada and South Africa, and he has authored more than 50 research publications and five books.

Mastering Mathematical Finance

Mastering Mathematical Finance is a series of short books that cover all core topics and the most common electives offered in Master's programmes in mathematical or quantitative finance. The books are closely coordinated and largely self-contained, and can be used efficiently in combination but also individually.

The MMF books start financially from scratch and mathematically assume only undergraduate calculus, linear algebra and elementary probability theory. The necessary mathematics is developed rigorously, with emphasis on a natural development of mathematical ideas and financial intuition, and the readers quickly see real-life financial applications, both for motivation and as the ultimate end for the theory. All books are written for both teaching and self-study, with worked examples, exercises and solutions.

[DMFM] *Discrete Models of Financial Markets*,
Marek Capiński, Ekkehard Kopp

[PF] *Probability for Finance*,
Ekkehard Kopp, Jan Malczak, Tomasz Zastawniak

[SCF] *Stochastic Calculus for Finance*,
Marek Capiński, Ekkehard Kopp, Janusz Traple

[BSM] *The Black–Scholes Model*,
Marek Capiński, Ekkehard Kopp

[PTRM] *Portfolio Theory and Risk Management*,
Maciej J. Capiński, Ekkehard Kopp

[NMFC] *Numerical Methods in Finance with C++*,
Maciej J. Capiński, Tomasz Zastawniak

[SIR] *Stochastic Interest Rates*,
Daragh McInerney, Tomasz Zastawniak

[CR] *Credit Risk*,
Marek Capiński, Tomasz Zastawniak

[FE] *Financial Econometrics*,
Marek Capiński, Jian Zhang

[SCAF] *Stochastic Control Applied to Finance*,
Szymon Peszat, Tomasz Zastawniak

Series editors Marek Capiński, *AGH University of Science and Technology, Kraków*; Ekkehard Kopp, *University of Hull*; Tomasz Zastawniak, *University of York*

Preface

The development of modern financial markets can be traced back to two events in the USA in 1973, both of which revolutionised market practice, for very different reasons. One of these revolutions was essentially institutional: the opening of the world's first options exchange in Chicago allowed options to be exchanged in much the same way as stocks (that is, through a regulated exchange) rather than having to be traded 'over the counter' as separate contracts between buyer and seller. The second upheaval was purely theoretical: the publication in the *Journal of Political Economy* of the now famous paper by Fischer Black and Myron Scholes (extended by Robert Merton in the same year), which developed arbitrage techniques for pricing and hedging options, and presented the now ubiquitous Black–Scholes formula for the rational pricing of European call options.

By the late 1970s the basis of their arguments, and the link with martingale theory in particular, had become well enough understood to allow the rapid development of this theoretical breakthrough, which has, since the 1980s, pre-occupied a host of financial economists and mathematicians (principally probabilists) and has given rise to the new profession of quantitative analyst (or 'quant'), which has attracted into the finance sector a large section of the best graduates with mathematics, physics, statistics or computer science degrees. This, in turn, has spawned a host of postgraduate courses emphasising market practice and taught in business schools, but increasingly also courses attached to mathematical sciences departments, focusing on the underlying mathematics, much of which is of comparatively recent origin.

At the same time, finance practitioners have led the explosive, largely unregulated, growth of new financial instruments, grouped together under the term 'derivative securities', which are constantly being devised to meet (or create) demand for specific tailor-made financial products in the banking, currency, insurance, energy and mortgage markets. Hedge funds, which specialise in trading these highly leveraged products, involving huge sums, have become major players in most developed economies. While the mathematical theory underlying their activities is based firmly on market models that exclude arbitrage opportunities (colloquially, a 'free lunch'), in practice much of the motivation comes from the search for risk-free profits, or, perhaps more accurately, the exploitation of market imperfections which briefly create highly marginal profits available through rapid, large-scale trading. This leads to secondary markets whose size overshadows global primary trade – by 2007 the annual volume of derivative trades had reached one quadrillion (10^{15}) US dollars, ten times the global industrial output over the past century – and where incautious, sometimes politically driven, decisions can leave banking institutions exposed to colossal losses, as was demonstrated painfully by the global banking crises of 2008–9 that continue to haunt the global economy.

All this suggests that a more thorough understanding of the principles underlying market practice is essential both for the improvement of that practice and for its regulation. Like nuclear power or the combustion engine, modern financial markets cannot be un-invented; instead, clear insight into their purpose, workings and potential benefits, which necessarily involves mastery of their mathematical basis, is a pre-requisite for adjusting market practice and preventing its abuse.

We will focus attention on the development of the Black–Scholes pricing model and its ramifications. Unlike its much simpler discrete-time binomial counterpart, the Cox–Ross–Rubinstein model (see [DMFM]), a proper understanding of this model requires substantial mathematical tools, principally from the stochastic calculus, which are developed carefully in [SCF]. The random dynamics of stock prices in the Black–Scholes model are based upon the Wiener process (often called Brownian motion). Despite its greater mathematical complexity, the continuous-time model produces a unique pricing formula for vanilla European options which is simpler than its discrete-time counterpart (the CRR formula described in [DMFM]), and has been universally adopted as a standard tool by finance practitioners.

Chapters 1–3 present the basic single-stock model for a general European derivative, with a focus on the explicit formulae for pricing calls and puts, and give a careful account of restrictions on admissible trading strategies. Since arbitrage opportunities usually involve trading in derivatives, the assets held in such strategies should include holdings in the derivatives being priced, and we show that, in our model, the prices of derivative and the replicating strategy must coincide if arbitrage is to be avoided. Option prices are derived in detail for vanilla European options and the unique admissible replicating strategy is constructed and related to the Black–Scholes PDE and to sensitivity measures for the option price relative to its parameters. The key roles of the risk-neutral probability and the representation of martingales by stochastic integrals are highlighted.

Chapter 4 extends and applies the Black–Scholes model in a variety of settings: options on foreign currencies, on futures and on other options. A structural model of credit risk is shown to fit the option pricing setting, a pricing model with time-dependent parameters is introduced, American call options are considered briefly, and the chapter closes with a description of the growth-optimal portfolio. Chapter 5 extends the discussion to the more exotic barriers, lookbacks and Asian options. A two-asset Black–Scholes model is first considered in Chapter 6 before presenting a general multi-asset pricing model, requiring more general versions of the Lévy and Girsanov theorems.

We restrict ourselves to the Black–Scholes setting and its immediate generalisations throughout this volume, working with the natural filtration of a given Wiener process and keeping our reliance on general martingale theory to a minimum. Notable features include the justification of derivative prices by means of replicating strategies and the care taken at the outset in defining the class of admissible trading strategies. The emphasis is on honest proofs of the results we discuss, with much attention given to specific examples and calculation of pricing formulae for different types of options. As usual, the many exercises, whose solutions are made available on the linked website www.cambridge.org/9781107001695, form an integral part of the development of the theory and applications.

We wish to thank all who have read the drafts and provided us with feedback.

Contents

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

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

We apply this with $X = \sigma W(t)$, so that $\mathrm{Var}(X) = \sigma^2 t$ (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

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

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 $z \in \mathbb{R}$.